

Meta-Information to Support Sensemaking by Developers

Amber Horvath

CMU-HCII-24-103

August 2024

Human-Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Committee

Brad A. Myers, Chair
Laura Dabbish
Aniket Kittur
Elena Glassman, Harvard University
Andrew Macvean, Google Inc.

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Human-Computer Interaction.*

Copyright © 2024 Amber Horvath

This work was supported by the National Science Foundation, under NSF grant CCF-2007482, and by Google. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect those of the National Science Foundation or Google.

Keywords: Human-Computer Interaction, Software Engineering, Sensemaking, Code Comprehension, Annotation, Developer Tools, Meta-Information, Developer Experience (DX)

Abstract

Software development requires developers to juggle and balance many information-seeking and understanding tasks. From understanding how a bug was introduced, to choosing what application programming interface (API) method to use to resolve the bug, to determining how to properly integrate this change, even the smallest implementation tasks can lead to many questions. These questions may range from hard-to-answer questions about the rationale behind the original code to common questions such as how to use an API. Software development, in contrast to other sensemaking domains, has the unique property that many information artifacts are created at different points during the development process (e.g., output data). Once this challenging sensemaking is done, this rich thought history is often lost given the high cost of externalizing these details, despite potentially being useful to future developers.

In this thesis, I explore different systems and methods for authoring and using this rich thought history as *meta-information* about code. Specifically, I have developed systems for *annotating* to support developers' natural sensemaking when understanding information-dense sources such as software documentation and source code. I then demonstrated how this meta-information can be captured and harnessed for new tasks, including for assessing the trustworthiness of documentation, for creating long-form documentation as stories about the code, for capturing design rationale and provenance data of code, and for supporting developer implementation tasks such as debugging.

This thesis begins by exploring methods for externalizing developers' thoughts in a form that is lightweight yet contextualized. We explore annotating as a method for simultaneously allowing developers to offload their mental processes, while presenting that information in-context for later developers to utilize. We developed two prototype annotation systems, Adamite and Catseye, which showed the promise of annotating for assisting developers both in overcoming issues with using developer documentation and debugging code. The dynamic nature of code and its connection to annotated materials introduced unique design challenges in which information can quickly become outdated and disconnected, leading both to changes in the Catseye user interface and to the insight that leveraging the (lack of) connectivity between annotation and documentation can support other documentation-related tasks, which inspired Sodalite. The final two systems, Meta-Manager and MMAI, explore capturing other forms of already-authored meta-information, such as edit traces and log data, for question-answering support, with MMAI utilizing large language models to make that querying possible in natural language.

The series of work introduced in this thesis points to the need to treat users' thoughts and intents as first-class entities and that meta-information is a way of preserving that information. I show that developers' thought histories can be represented in the form of code-related meta-information and, through proper tooling, can be used by later developers to accelerate their sensemaking of code.

Acknowledgements

Completing this dissertation was a marathon, not a sprint. And, during this marathon, I have had the tremendous opportunity to grow as a researcher, thanks in part to the endless support from my network of friends, family, academic and industry partners, and collaborators.

Firstly, I would not be here without my amazing advisor, Brad Myers. Brad has been with me through this whole journey, fostering my academic growth through all of the ups and downs that come with completing a Ph.D. There were times when I wanted to quit, but his belief, support, and resilience, even when I could not muster up such feelings, kept me here and allowed me to reach this point – this achievement is as much his as it is mine. I am proud to be a node on his famous [advisee tree](#).

I am also very fortunate to have a large thesis committee of brilliant researchers to further assist in my academic growth: Laura Dabbish, Aniket (Niki) Kittur, Elena Glassman, and Andrew Macvean. Their varied expertise has lead to thoughtful questions and lines of inquiry I would not have otherwise considered. They, along with Brad, have acted as prime models of how to be an excellent researcher and collaborator – lessons I will take with me moving forward.

Beyond Brad and my committee members, I have been privileged over the course of my Ph.D. to have opportunities to receive mentorship from the wider HCI, Software Engineering, and Carnegie Mellon University communities. I would like to thank Margaret Burnett who advised me at Oregon State University and introduced me to this wonderful, chaotic world of academic research that I am now a part of – her mentorship has been invaluable. I would also like to thank Emerson Murphy-Hill, my internship host at Google in 2019, who has continued to provide me support and guidance over the last few years. Special thanks also to Andrew Begel, Patrick Carrington, Jodi Forlizzi, Jess Hammer, Ken Holstein, Geoff Kaufmann, Queenie Kravitz, Bogdan Vasilescu, John Zimmerman, and countless others with whom I had the pleasure of learning from during this journey.

I am eternally grateful to my collaborators without whom this research would not be possible. I would like to give special thanks to my amazing lab mates, Michael Coblenz, Matthew Davis, Mary Beth Kery, Toby Li, Jenny Liang, Michael Xieyang Liu, and Daye Nam for their support and contributions to this work. I also had the fortunate opportunity to work with undergraduate and masters students, all of whom helped tremendously with this work – thank you to Shannon Bonet, Kazi Jawad, Emma Paterson, Imtiaz Rahman, Connor Shannon, Matthew Shu, and Lai Wei.

I would additionally like to thank the sponsors of this research, my paper reviewers, and my study participants – all of whom served as the backbone of this work. Specifically, I thank the National Science Foundation (CCF-2007482) and Google for the provided financial support.

My heartfelt love and appreciation extends to all of the amazing friends I have made over the course of the last six years: Karan Ahuja, Lea Albaugh, Mark Buttwieiler, Alex Cabrera, Julia Cambre, Tianying Chen, Erica Cruz, Wesley Deng, Jonathan Dinu, Will Epperson, Morgan Evans, Matt Ho, Roger Iyengar, Sung Jang, Hyeonsu Kang, Pranav Khadpe, Rushil Khurana, Lynn Kirabo, Andrew Kuznetsov, Tom Magelinski, Courtney Miller, Steven Moore, Huy Nguyen, Wode “Nimo” Ni, Napol Rachata-sumrit, Sam Reig, Melrose Roderick, Vivian Shen, Jaemarie Solyst, Jordan Taylor, Stephanie Valencia Valencia, Danny Weitekamp, Kristin Williams, Franceska Xhakaj, Nur Yildirim, and many more. Thank you all for the laughs and good memories. Special thanks to my dear friends Sarah and Gareth Baldrice-Franklin for our weekly-ish movie nights that helped me stay sane during COVID and beyond. And, of course, Emily Miller, who has been like a sister to me for years.

I would also like to take space to acknowledge and appreciate a valued member of our community who we lost far too soon: Sujeath Paredy. As part of our 2018 cohort of Ph.D. friends, his loss has felt like a giant, gaping hole as our group moves towards graduation – I would like to dedicate this dissertation to him. I remember you, Sujeath, and wish you were still here with us.

Lastly, my deepest gratitude goes to my family – Diane Pearson, Steve Horvath, and Katelyn Horvath. They have stood by me through every up and down, providing immeasurable emotional support. I could not have done it without them. I would also like to thank my extended family and my “bonus” family, Lisa Hendriksen, Casey Keating, Suzanne Robinson, Brittany Thackeray, and Renate Woods, who all stood by and joined in the celebrations. I would also like to thank my pet family, Ringo the Pomeranian, Eevee the cat, and Pichu the cat – I know they cannot read, but their cuddles and love has gotten me through many exhausting paper deadlines. Last, but certainly not least, I am forever indebted to my fiancé, River Hendriksen, who has truly been everything to me over the course of this dissertation and beyond – collaborator, friend, shoulder to cry on, co-author and partner. He helped me carve out the life here in Pittsburgh that I have grown to love so much, including happy hours at our local Primanti Bros., evenings at Pittsburgh Penguins games, and nights at Dependable Drive-In. I love you.

Contents

Abstract	iii
Acknowledgements	v
1 Introduction	1
1.1 Why Meta-Information?	3
1.1.1 Meta-Information Properties	4
1.1.2 Source Properties	6
1.2 Overview	8
2 Background and Related Work	13
2.1 Making Sense of Code	13
2.1.1 Developer Tools	13
2.2 Meta-Information About Code	15
2.2.1 Writing About Code	15
2.2.2 Other Meta-Information	17
3 Adamite: Meta-Information as Annotations on Documentation	19
3.1 Overview	19
3.2 Preliminary Studies and Design Goals	20
3.2.1 Lab Study with Hypothesis	20
3.2.2 Corpus Analysis of Hypothesis Annotations	21
3.2.3 Design Goals	22
3.3 Overview of Adamite	23
3.4 Lab Study	25
3.4.1 Method	25
3.4.2 Results	29
3.5 Limitations	35
3.6 Discussion	36
4 Catseye: Meta-Information for Sensemaking About Code	39
4.1 Overview	39
4.2 Catseye	41
4.2.1 Overview of Catseye	41
4.2.2 Background and Design Goals	43
4.2.3 Implementation Notes	45

4.3	Lab Study	46
4.3.1	Method	46
4.3.2	Participants	48
4.3.3	Analysis	48
4.4	Results	49
4.4.1	What Information Do Developers Keep Track of with Annotations and Artifacts?	49
4.4.2	How Do Developers Use Their Annotations and Artifacts?	52
4.4.3	How Did Participants Identify and Fix Their Bugs?	53
4.4.4	Continued Usage	54
4.5	Discussion	55
4.6	My Usage of Catseye	57
4.7	Conclusion and Future Work	59
5	Curating Ephemeral Meta-Information with Catseye	61
5.1	Overview	61
5.2	Background and Related Work	62
5.3	Preliminary Exploration of Adamite Annotation Curation	64
5.4	Preliminary Study of Catseye Annotations	65
5.5	Design Probe: Catseye Annotation Curation through Re-Anchoring	68
5.5.1	Algorithmic Re-Anchoring	69
5.5.2	User Interface for Re-Anchoring	71
5.6	Design Probe: Catseye Annotation Curation through Batch Processing	73
5.6.1	User Interface for Batch Processing	74
5.6.2	Merging Annotations	75
5.7	Discussion and Future Work	77
6	Sodalite: Meta-Information to Support Documentation Management	81
6.1	Overview	81
6.2	Background and Related Work	82
6.3	Sodalite	83
6.3.1	Templates	83
6.3.2	Code Links and Suggestions	85
6.3.3	Support for Reading	86
6.3.4	Support for Maintenance	87
6.4	Evaluation of Sodalite	89
6.4.1	Study Design	89
6.4.2	Study Results	90
6.5	Discussion and Future Work	92
7	Meta-Manager: Meta-Information for Question-Answering	95
7.1	Overview	95
7.2	Overview of Meta-Manager	96

7.2.1	Developer Information Needs	97
7.2.2	Scenario	98
7.2.3	Detailed Meta-Manager Design	100
7.2.4	Implementation	105
7.3	Lab Study	106
7.3.1	Method	106
7.3.2	Participants	109
7.3.3	Quantitative Results	110
7.3.4	Qualitative Results	111
7.4	Discussion	114
7.5	Limitations and Threats to Validity	115
7.6	Future Work	116
7.7	Conclusion	118
8	MMAI: Accelerating Sensemaking with Logs and LLMs	119
8.1	Overview	119
8.2	Background and Related Work	120
8.2.1	Terminology and Background	120
8.2.2	Print Debugging	122
8.3	Exploratory Interview Study	123
8.3.1	Method	123
8.3.2	Results and Discussion	124
8.4	Overview of MMAI	126
8.4.1	Scenario	127
8.4.2	Detailed Design	130
8.4.3	Implementation	134
8.5	Discussion	137
8.6	Future Work and Conclusion	138
9	Conclusion and Future Work	143
9.1	Summary of Contributions	143
9.2	Discussion and Future Work	145
9.2.1	Designing with Information Ephemerality In Mind	145
9.2.2	Designing for the Software Engineer of Tomorrow	146
9.2.3	Designing for Meta-Information as a First-Class Entity	147
9.3	Concluding Remarks	148
A	Meta-Manager Architecture and History Model	151
B	MMAI GPT-4 Prompts	155
B.1	Pre-Processing History Prompt	155
B.1.1	Prompt	155
B.2	Default User Query Prompt	156

B.2.1	Prompt	156
C	Chapter 3 Study Replication Materials	157
C.1	Preliminary Study Materials	157
C.1.1	Protocol	157
C.1.2	Task Source Code	161
C.2	Lab Study Materials	161
C.2.1	Adamite Authoring Condition Protocol	161
C.2.2	Adamite Using Condition Protocol	168
C.2.3	Control Condition Protocol	175
C.2.4	Post-Task Survey	178
C.2.5	Task Source Code	179
D	Chapter 4 Study Replication Materials	181
D.1	Catseye Condition Protocol	181
D.1.1	Consent	181
D.1.2	Part 1: Tutorial	182
D.1.3	Part 2: Task	189
D.1.4	Part 3: After Task	189
D.2	Control Condition Protocol	190
D.2.1	Consent	190
D.2.2	Part 1: Tutorial	191
D.2.3	Part 2: Task	194
D.2.4	Part 3: After Task	194
D.3	Post-Task Survey	195
D.3.1	Notes	195
D.3.2	Survey	195
E	Chapter 6 Study Replication Materials	199
E.1	Sodalite Study Protocol	199
E.1.1	Introduction	199
E.1.2	Tutorial	200
E.1.3	Study	203
E.2	Post-Task Survey	204
F	Chapter 7 Study Replication Materials	207
F.1	Meta-Manager Study Protocol	207
F.1.1	Tutorial	207
F.1.2	Task	210
F.1.3	End Task	211
F.2	Post-Task Survey	211

G Chapter 8 Interview Protocol	215
G.1 Protocol	215
G.1.1 Questions	215
Bibliography	217

List of Figures

- 1.1 An exemplar instance of meta-information (in this case, an annotation) attached to a source (here, code), as discussed in Chapter 4. Key properties of meta-information, along with properties derived from the source document, are presented as a way of structuring some of the key features we discuss and use when designing our systems described throughout this thesis. **Blue** properties are associated with the information source, while **green** properties are associated with the meta-information artifact itself. 3
- 1.2 The “virtuous cycle” of code meta-information generation during software development (top arc) and how that meta-information may be used by later readers (bottom arc). The various forms of meta-information each of my systems generates and exemplar questions that that information may answer are color-coded: **Adamite** (Chapter 3) is shown in **green**, **Catseye** (Chapter 4) is shown in **red**, **Sodalite** (Chapter 6) is shown in **orange**, **Meta-Manager** (Chapter 7) is shown in **blue**, and **MMAI** (Chapter 8) is shown in **purple**. 8
- 3.1 Adamite’s sidebar (on the right) open on an already-annotated web page in the browser. (1) shows the pop-up for when a user selects some text – at this point they can begin creating a new annotation by selecting an annotation type. (2) shows the menu of question annotation prompts users can choose from. (3) shows a published normal annotation with two anchors. (4) shows how the annotated text appears on the web page. (5) shows Adamite’s search and filter pane. (6) shows the pinned annotation list button. 23
- 3.2 The correct output for the task. Each number refers to the step number. (1) creates and renders the 4 images. (2) puts the images in 2 rows. (3) arranges images by a user-defined property using the `arrangeBy` method. (4) requires the user to set a label on their data, such that elements with the same label will have a matching stripe along the bottom of the picture. 26
- 3.3 The number of annotations removed for each reason, along with the annotations kept, out of the 91 total annotations. 2 highlight annotations were retained as the users edited them to add text, making them semantically identical to normal annotations. 27

3.4	The difference between reading and each of the other two conditions is statistically significant, but the difference between control and authoring is not. The average number of steps completed is in the center of each box.	29
3.5	The proportions for each type of the annotations made in the authoring condition (out of 91), with the exact count for each type above the bar and the proportion in parentheses.	30
3.6	Which participants made what type of annotation. A1 through A10 refer to the 10 authoring condition participants. R1 and R2 are the two reading condition participants who created annotations.	33
4.1	Catseye as it appears in Visual Studio Code. (1) shows how the annotation appears in the editor – the code is highlighted with a light gray box and, when hovered over, the annotation content appears in the pop-up with any other documentation. Clicking on the “Show Annotation” button opens and brings into focus the Catseye pane if it is not already visible, then scrolls to the annotation. (2) shows the annotation location(s) in the scroll bar gutter in a light green. (3) is a search bar for searching across the user’s annotations. (4) is the Catseye pane – the pane is segmented into sections corresponding to the annotations’ locations in the file system, with the “Current File” section currently open. (5) is an annotation – the top of the annotation shows the author and creation time information on the left, and buttons for various actions. (6) shows the two code anchors for the annotation. (7) is the content the user added as an annotation to the code snippets. (8) is a snapshot of the code at a previous version with a comment added by the author about this version of the code. (9) is a reply to the original annotation.	40
4.2	An annotation made by a participant in our study. (1), (2), and (3) show the 3 different code anchors the participant created across multiple files, with the first anchor (“gameloader.js”) as the site of their question, and the remaining two anchors and reply (4) answering their question. The annotation was pinned.	42
4.3	The average number of annotations and artifacts participants created during the study.	50

- 5.1 An example scenario in which an annotation’s anchor is updated in multiple ways after a `git pull`. In the case only the position changes, the annotation content should stay as-is and the anchor will update to represent the new start and end positions. If the anchor is deleted, it is unclear whether to delete the annotation or to re-attach the annotation at a different, semantically similar point. Likewise, if the content within the anchor’s bound changes (`this._copyVscodeMetadata` to `this.getCopyMetaData()`, in the figure), under different circumstances it may make sense to delete the annotation, keep the annotation on the new anchor content, or re-attach the annotation elsewhere. 63
- 5.2 How the Adamite website appears. (1) is the search and filter pane for finding relevant annotations. (2) is the list of groups the user is a part of. (3) is the sort bar – currently the list is sorted by time in descending order. (4) are the batch operations a user can do on their set of selected annotations. (5) is an annotation – it is currently selected, shown at the green checkmark in the top left corner. (6) is another annotation with two anchors (the two green rectangles with text and a URL) and annotated text. 64
- 5.3 How the re-anchoring algorithm expands its search outwards, given no sufficient match at each step. (1) is the original anchor content and location, (2a) and (2b) are the 5 lines above and below, respectively, the original anchor location, (3) is all of the code belonging to the parent node of the original anchor in the AST, and (4) is the whole file. 69
- 5.4 How running the matching algorithm at the token level with `this._copyVscodeMetadata` across multiple tokens in multiple code lines would yield different results, given the string difference and location difference. Note that the most reasonable re-anchoring spot from Figure 5.1 has the highest match score. 70
- 5.5 The user interface when re-anchoring an annotation. (1) shows the last-known anchor, along with the surrounding code in green. (2) is the Show/Hide Suggestions button, allowing the user to leave the annotation unanchored, if the user so desires. (3) is the Manually Reanchor option – the user can select some code in the editor, then click the “Manually Reanchor” button to set the anchor to their selected code. (4) is the first of multiple candidate anchors. (5) are the options for what a user can do with a candidate anchor – they can either remove that particular anchor as an option, “Reanchor” their annotation to that anchor, or click through the carousel (i.e., 4 gray dots) to view their other options. (6) is the annotation. 72

- 5.6 The user interface for batch operating on annotations in Catseye. (1) is the search bar, (2) is the set of sorting and filtering operations including sort by location or time, scope of annotations to include (including annotations made on this file, this project, or across all projects), who authored the annotation, what annotation type it is (similar to the Adamite annotation types), and whether or not the annotation has been marked as “resolved”, is pinned, or is anchored. (3) is the set of batch operations a user can perform on any selected annotations, with options for (from left to right) merging (Figure 5.7), pinning, sharing, resolving, or deleting. (4) are collapsed annotations which are currently checked, meaning they will be included in any batch operations selected at (3). (5) are the buttons for operations a user can perform on a single annotation. 74
- 5.7 The UI for merging annotations. (1) are batch operations for creating the merged annotation through importing all of the content from the original annotations. (2) is a preview of the annotation anchors, with the option to remove them (trash can icon to their right). (3) is the regular annotation authoring UI, such that users can add types and text that may not be in the original annotations. (4) is the first annotation that is being used to create the merged annotation – the checkmark shows the anchor will be included. (5) points at the icons for importing annotation content and replies into the resulting annotation – annotation replies can be added to the merged annotation’s main body text by clicking the double chevron icon, while the single upwards arrow will bring the reply’s content up as a reply. (6) is the reply content for this annotation. (7) is the other annotation to be included in the merged annotation – since it is a highlight annotation, it has no annotation text or replies. 76
- 6.1 The editor for authoring a story using Sodalite – this is a simplified and anonymous recreation of P2’s story for demonstrative purposes. (1a) and (1b) show the relationship between the code editor and the story editor – in this case, the user has clicked the code Link, prompting the suggestions pane to show identifiers related to Link. (2) is the template list with the selected “Overview” template highlighted. The “Code References” at (3) are represented in a list (4a, 5a, and 6a) and their locations are shown in the story text at (4b), (5b), and (6b), respectively. The title of the story is (7). 84

- 6.2 How Sodalite appears after a story has been authored. (1) shows the hover text when a code link is interacted with in the code editor. (2) is a code link that has been marked as “needs review” given that the original code (shown in (3)) and the code in the editor (at (1)) are different. (4) are two other collapsed code stories. 86
- 7.1 Meta-Manager as it appears within Visual Studio Code: the pane appears in the bottom area of the editor, with the left area displaying a visualization of the history of the code file over time, while the right area displays information about a particular code version. 96
- 7.2 How the code box looks when expanded to show a code version – in this case, a “Paste” event version. (1) shows the buttons specific to a “Paste” code version, including the “See Copy” button which will navigate the user to the corresponding copy event on the timeline (if the copy happened in a different file, then the code box will update with a preview of how the code in the other file looked at the time of the copy, which can be clicked on to change to that file); (2) shows the text explaining what happened with this particular paste event — clicking in this area will open the editor tab showing what the code file looks like now; (3) shows the code for this version, along with a light blue highlight on the code that was pasted. 101
- 7.3 A zoomed-in portion of the timeline shown in Figure 7.1. This zoomed-in portion shows around 120 edits between Version 710 and Version 830, with the scrubber set around Version 740, when a user pasted code from Stack Overflow. 105
- 7.4 Each question scored by participants in terms of how often they encounter similar questions in their own programming experiences. . . . 111
- 8.1 MMAI as it appears within Visual Studio Code: the pane appears in the bottom area of the editor, with the left area displaying a visualization of the history of user-selected code nodes that the user wants to inspect, while the right area displays information about each code node, now including additional output information and an interactive chat component. The two code version boxes for `calculateBoxValue` and `App` are collapsed (see expanded version in Figure 8.3). 127
- 8.2 How `mmlog` calls appear in the editor, including star icons in the gutter to either include or exclude the calls and their corresponding nodes from the rendered “Joint History” visualization (Figure 8.1-1) and code version boxes (Figure 8.1-2b and 3b) in the MMAI pane. Filled-in star icons mean the nodes will be included. Note that these two functions, `calculateBoxValue` and `App`, are in different files. 128

- 8.3 How an expanded code version appears in MMAI. (1) is similar to Meta-Manager (see Figure 7.2), in that information about the edit is shown including when it occurred and what the code looked like at that point in time, along with any additional meta-information (with this version containing no additional meta-information). (5) shows the optional, supplementary information when `mmlog` is used, including the log statement (6) (in case multiple `mmlog` statements exist in one version), most recent log statement and values along with aggregate counts and a button to ask about that value (7), and computed queries that can be sent to the LLM, given the `mmlog` statement and value (8), along with the option for the developer to write their own query (9). 129
- 8.4 How the chat interface appears within the MMAI interface, below the log and version data. (1) is the user’s query – in this case, this is a pre-generated query, created by pressing the “Compare to current version button” (Figure 8.3-C). (2) is the LLM’s response. (3) is a link to the code version the system is looking at when coming up with its answer. (4) are the log values the system found in its history search. (5) are the generated follow-up questions given the answer the system generated. (6) are the references the system used, including links to code versions and a summary of the span of history it looked at when coming to its answer. 132
- 8.5 How an `mmlog` version appears in the MMAI database. One code version owns a collection of these `mmlog` versions dependent upon how many `mmlog` statements the MMAI node has at a particular version and how many times each `mmlog` statement is invoked. 136
- A.1 A simplified version of how Meta-Manager’s data model and architecture look on a file called “example.ts”. 152
- A.2 How a version of a node appears in the Meta-Manager database on Firestore [61]. Note that the id of the version includes the node ID and the time at which the version was captured. This version does not include any additional meta-information, such as web activity. . . 153

List of Tables

3.1	Counts of each issue and question annotation that identified or was caused by an issue discussed in [249]. Note that two code examples did not work because they suffered from a fragmentation issue, so they are coded both as a poor code example and a fragmentation issue. Similarly, all fragmentation questions were caused by fragmented code examples, so they are also coded as both a fragmentation question and code example question.	31
4.1	The bugs present in the two games. “Value” refers to a construct in the program, such as an operator, boolean, or variable.	47
4.2	How the study task encapsulates the types of information Catseye supports.	48
4.3	The annotations and artifacts participants created during the study while working on each bug. The experimental condition made 40 annotations while working on bugs, while the control condition created 35 artifacts. The last 2 columns refer to the proportion of annotations made about that bug out of the 40 annotations made while debugging, and the proportion of control condition artifacts made about that bug out of the 35 artifacts made while debugging, respectively.	53
7.1	Each question that was asked during the task, along with what information need from prior literature it corresponds to, the steps that could be taken in Meta-Manager to answer the question, and how participants performed on the question in terms of correctness and time spent (in minutes). Note that some questions represent more than one information need, such as Q5, which both asks what code is related to the commented out loop, but also why the loop was commented out, which is a rationale question.	109

Chapter 1

Introduction

Developing software requires developers to keep track of many types of information while performing various interleaved tasks. For example, to debug some code, a developer must first navigate through the code to determine what part or parts of the code are responsible for the bug [129, 142, 197], run the code and diagnose the output to understand how and when the code fails [131, 142, 197], write code and possibly research solutions online to fix the bug [29, 147, 263], then test the code again to ensure the solution worked [73]. Typically, the developer's mental model of this rapidly-evolving problem space is not externalized [33, 151, 191], which can be problematic as working memory is limited [99, 160, 243]. Further, development processes are often iterative [126], can be interrupted [189, 191], and may just be one task amongst many that the developer is handling [127], leading to further cognitive strain.

Given this significant cognitive cost, many research projects have enumerated the challenges in both writing and understanding code. Just some of the challenges include maintaining task awareness [159, 189, 191], making sense of external libraries and their documentation [4, 104, 174] (with software documentation, itself, having its own set of well-known problems [4, 143, 180, 212, 249]), understanding the rationale behind the current code [127, 138, 159, 226], determining what part or parts of the code are relevant to the change the developer is introducing [126, 132], and designing that change and implementing it [132, 147]. Completing these smaller information-intensive tasks often results in the creation of knowledge that is lost, either because it is not externalized by the progenitor of the information or because it is not logged.

Some qualities of programming make tracking this information especially challenging in comparison to other complex sensemaking domains. For one, a developer's end-state is not often reaching an answer or decision given these implicit questions, but is to utilize that information to complete some programming-related task, such as fixing a bug. Secondly, the information landscape is often changing as the developer introduces new code, which can produce new runtime behaviors and introduce new questions. This means that:

- the developer does not normally have the mental bandwidth to externalize this information beyond, perhaps, a short note [107, 151]

- this information is implicitly associated with the corresponding code that stems from the gained knowledge [107, 147]
- as the code changes, the information the developer learns is highly contextualized to the time in which it was discovered or created [17].

We call this information created as a byproduct of authoring or making sense of developer materials (e.g., code, web resources, documentation, etc.) *meta-information*.

In software development, one of the most ubiquitous and well-known meta-information types are code comments. Often created as a way of keeping track of bugs [240] and open to-do items [237, 238], documenting code [220], or informal versioning of code [119], code comments can serve as rich information about what a developer is doing when managed with tooling [247]. Other forms of meta-information about code include Git commit messages, code documentation, code versions, and so on. The act of developing software produces a vast and dense information landscape comprised of these various forms of meta-information. Meta-information also exists in other information authoring domains such as Microsoft Word or Google Docs comments.

In this thesis, I¹ explore developing software systems to capture and present programming-related meta-information to assist in overcoming known challenges in developing software. This work is predicated on the knowledge that many forms of meta-information are generated at different times during the software development life cycle and I hypothesize that this knowledge, when captured and presented in a comprehensible manner that leverages the code or other source information within the developer's working context, can be useful. By associating this meta-information with code, the connection between what the developers care about (i.e., the code) and all of this other information that normally is not captured but can actually answer developers' implicit questions about code (e.g., visited web pages for *where* some code came from, code edit history for *when* some change occurred, added and removed and commented out code for reasoning about *what* a developer has tried, etc.) is an effective way to keep the developers within their working context, while providing them the information they need to answer their questions.

My thesis aims to investigate the following claim:

*By tightly coupling meta-information to its source origin through **anchoring** to a particular location at a point in time, and by leveraging properties of the source environment, information that would otherwise be unavailable or intractable can be useful for developers in both the short-term and long-term.*

For the initial author in the short-term, I show that externalizing their thoughts and questions as meta-information can help with keeping track of information while completing cognitively-demanding software engineering tasks when anchored to a

¹The projects discussed in this thesis were lead by myself but were developed in collaboration with other researchers – out of respect for their contributions, I will predominately use “we” when discussing the various projects and papers that comprise this thesis.

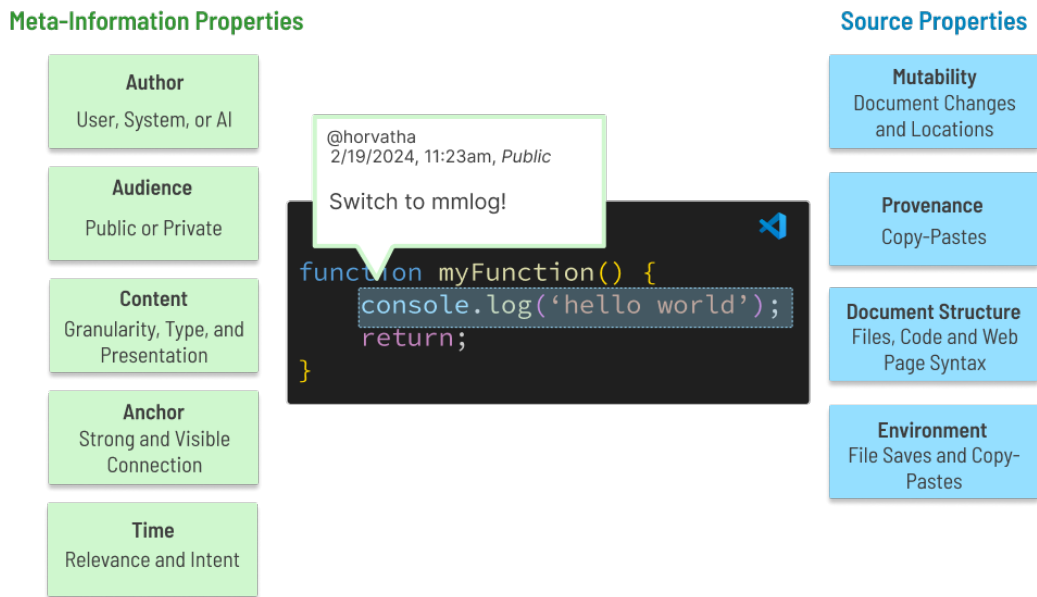


FIGURE 1.1: An exemplar instance of meta-information (in this case, an annotation) attached to a source (here, code), as discussed in Chapter 4. Key properties of meta-information, along with properties derived from the source document, are presented as a way of structuring some of the key features we discuss and use when designing our systems described throughout this thesis. Blue properties are associated with the information source, while green properties are associated with the meta-information artifact itself.

source material, such as code or documentation. Over a longer span of time, these traces of the context in which the initial development occurred can help with answering otherwise unanswerable questions about code and for better understanding unfamiliar code and learning resources. I also show that code meta-information can be used, not only for reasoning about some code, but also for evaluating written materials about code, given the connection between meta-information and its corresponding code. An overarching goal of this thesis is to treat information about code as a first-class entity, not unlike the code itself, given that the rich history of code and its development is often the *answer* to developers' questions, but is not typically captured in a systematic or easily explorable manner. I achieve this through creating a tight coupling between meta-information and its source, bringing the meta-information into the rich source environment and, thus, presenting it as a first-class entity.

1.1 Why Meta-Information?

To better introduce code meta-information and its potential value for code comprehension, along with the source document and environment it can be associated with, I will briefly discuss some of the properties and values we found throughout this

thesis and our exploration of meta-information systems (see Figure 1.1). These properties both represent facets of our meta-information systems and design considerations that we had to intentionally support given each system and its goal. Ultimately, we believe these to be some of the most important properties to consider when designing a meta-information system, especially with respect to supporting developer sensemaking.

1.1.1 Meta-Information Properties

Author

Different types of meta-information can come from different sources – the user, system, or an intelligent agent. Broadly, our systems have found success through supporting all 3 authors, given the information type. We support allowing more personal meta-information to come from the user (Chapters 3 and 4), which can be curated, scaffolded and transformed with tooling (Chapters 5 and 6). For more descriptive and laborious-to-track information, such as visited web pages and output values, we have found success in automating the collection and contextualizing that information with system-authored meta-information (Chapters 7 and 8). AI can then make sense of that information to assist in higher-order reasoning tasks (Chapter 8), in turn creating additional meta-information that may be stored.

Audience

With self-authored meta-information, such as annotations, notes are commonly intended to be primarily for the author’s personal benefit (such as a “to-do” note). Oftentimes, the author must make a conscious decision to transform the note into something shareable, either because the text requires some “clean up” due to terseness and/or lack of formality or because not enough information was included in the original note in order to be comprehensible by someone else [221]. To support cleaning the text, this process can be supported through curation methods (Chapter 5) and through structuring the text from the outset such as with templates (Chapter 6). To convey more context, intentional system design decisions, such as supporting multiple anchors and annotation types, can remove the need for conveying this additional context in text through conveying the context through the note’s properties (Chapter 3). We also found a benefit in designing a system entirely for a more private and personal context with our code annotation system, which was initially designed primarily for an individual’s information tracking when developing software (Chapter 4).

Content

Meta-information, by definition, is *any* information associated with some artifact – thus it can be truly anything. This intentional ambiguity leads to design questions

of what granularity to capture information at, what types of information to support, and how to present this information in a comprehensible manner. We found success in supporting generalized information “workspaces” by using the flexible properties of annotations – we found we could support high-level categorization with types, connecting of resources with multiple anchors, and prioritization with pinning (Chapters 3 and 4). We then explored capturing other domain-specific forms of meta-information such as code versions through the unifying and general purpose nature of annotations (Chapter 4). This large amount of meta-information allows us to transform the meta-information into newer, more comprehensive forms such as curated or merged documents (Chapters 5 and 6).

Choosing what granularity to capture meta-information at and how to present it are highly-dependent upon the sensemaking task we are aiming to support. For making sense of code history, including log data, we sought to capture as much potentially-useful information as possible at a low-level of granularity considering it is unclear what will be helpful at a later point in time (Chapters 7 and 8). However, this leads to issues of scale resulting in a needs for a complementary presentation. Broadly, we found success in coupling together complementary views for sensemaking – a high-level visualization to flatten a potentially large set of data into a singular view, with hints of what information may be most relevant given the tasks we are supporting, and a lower-level detail view for in-depth exploration (Chapter 7). Designing meta-information systems allows us to also capture information that is not normally stored due to issues of scale and contextualization – we explore capturing large amounts of log and history data and leverage the unique properties of anchoring information in order to both scale down the amount of information to store and present it in a meaningful manner later.

Anchor

Arguably the most important property we explored in our work, information “anchors” act as the binding between the meta-information and its source. With annotations, the source anchor(s) are the connection between the note and some text or code – we found success in leveraging that connection for clustering related information (Chapters 3 and 4), conveying context for later users (Chapter 3), and leveraging this tight coupling for later curation-related tasks such as merging information and creating documents (Chapters 5 and 6). This connectivity can also support more domain-specific tasks such as navigating a code base (Chapters 4 and 6). In later systems, we automatically generated connections using the unique source structure of code and connected that code to meta-information. In this way, semantically-meaningful blocks of code became anchors for various forms of meta-information across time. The tight coupling afforded by anchoring meta-information to its progenitor was a powerful enough design that it is a central component in every system discussed in this thesis.

Time

The time at which some meta-information is created is often an important facet of the meta-information itself. Some types of meta-information, such as notes, are often authored with the expectation that they will only be used in the moment and, thus, do not need to persist across time. Other forms are fully bound to the time in which they are made and are immutable, such as code versions where their purpose is to capture some code's state at a particular point in time. This work explores the spectrum of how to think about time and its relative impact on meta-information given the type of meta-information. For our annotation systems, time was a key factor in understanding relevance of annotations across time (Chapters 3, 4 and 5), especially with respect to source code and changing source content over time. Our later history systems treat time as an immutable construct by which meta-information is bound to, given a code anchor (Chapter 7) or an execution of a log statement (Chapter 8).

1.1.2 Source Properties

We consider the "source" to be the space in which the progenitor of some meta-information exists. In this thesis, this is most often the integrated development environment (IDE) and its source code, but may also refer to web-based information, including documentation.

Mutability

Directly related to the meta-information properties of "time" and "anchors", source mutability makes it particularly challenging to keep some piece of meta-information anchored over time. Our systems have explored various ways of handling this design challenge – with our annotation systems we wanted to provide the user direct control over their anchoring points, leading to a large amount of user control and flexibility (Chapters 3 and 4). However, that came at the cost of having less certainty as to how to handle updating the anchor across document changes – see Chapter 5 for a whole discussion on this challenge. Later systems switched to utilizing the inherent structure of the source code to derive anchor points without requiring user intervention. This approach still requires some text-matching and inference, which can be wrong, and loses the ability for a user to directly control their anchor points – yet we had more success in keeping meta-information attached across document changes than in the prior systems. Ideally, a design can be as automatic as possible and minimize user intervention, unless the user wants to control their anchoring locations in which case a system should attempt to assist in the process through utilizing other contextual information.

Provenance

A key form of meta-information about the source is the provenance of the source, or *where* and *when* the source material came to be. By taking advantage of the working context of the source, we can sometimes automatically derive provenance data by recording copy-paste events (Chapter 7). However, capturing a full “trail” of provenance information can be difficult if, e.g., some code in my IDE originally came from Stack Overflow which came from some other developer’s IDE which came from a Git repository, and so on. Composing and capturing even part of this trail requires extensive domain-specific tooling, which we explore in Chapter 7. When looking forward with respect to meta-information systems, I expect information provenance to become one of the more important pieces of meta-information in order to ascertain how trustworthy some information is, given whether or not it came from an AI system or from Stack Overflow, and how much research a developer did prior to coming to an implementation decision [153].

Document Structure

An advantage in providing tailored meta-information support in a specific working context is being able to provide domain-specific features. An example of this is leveraging the source document structure in order to provide additional benefits. In our systems, we largely utilized the source code document structure to assist in (re-)anchoring our meta-information to some code (Chapters 4, 5, 6, 7, and 8). We additionally derived page-level details from the structure of the web page in order to send additional meta-information about, e.g., a ChatGPT thread to complement our other historical data (Chapter 7). For supporting other domains, it is important to consider what meta-information can be extracted given the source material itself and how it may be leveraged to support information comprehension tasks.

Environment

Meta-information is often either authored by a user or derived from interaction traces created by the user as a byproduct of completing their primary task. Similar to the discussion on document structure, having detailed and comprehensive knowledge of a user’s working context and environment allows us to support richer meta-information capture and presentation. We tightly coupled many of our systems functionalities to the practices developers employ in the IDE. For example, annotation anchors are updated and code versions and log data are pushed on file save (Chapters 4, 7, and 8) and we capture and push provenance data, both within IDE and within browser, on copy and paste (Chapter 7), respectively. Detailed knowledge of a user’s working environment and the activities they are completing in that environment can lead to the automatic capturing and anchoring of more forms of meta-information.

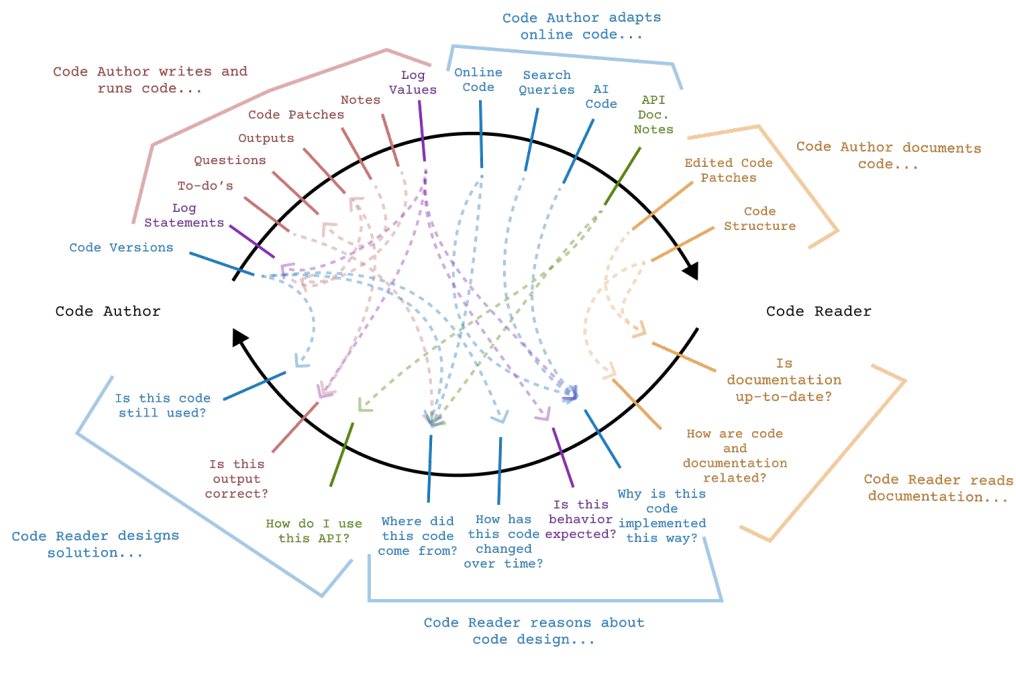


FIGURE 1.2: The “virtuous cycle” of code meta-information generation during software development (top arc) and how that meta-information may be used by later readers (bottom arc). The various forms of meta-information each of my systems generates and exemplar questions that that information may answer are color-coded: *Adamite* (Chapter 3) is shown in green, *Catseye* (Chapter 4) is shown in red, *Sodalite* (Chapter 6) is shown in orange, *Meta-Manager* (Chapter 7) is shown in blue, and *MMAI* (Chapter 8) is shown in purple.

1.2 Overview

To begin exploring to what extent meta-information can be used to help software developers overcome their sensemaking barriers given the aforementioned design properties of meta-information and its relationship to a source document, we developed a series of systems that exemplify the qualities we expected would be helpful for developer sensemaking, including question-asking and answering. Figure 1.2 shows an example “cycle” of virtuous meta-information creation and usage given a developer’s initial work and subsequent comprehension by a later² developer. This figure acts as a complement to Figure 1.1, in that it situates the example meta-information and its properties into the larger context of what types of meta-information the systems support capturing and example usages of that meta-information for comprehension. Thus, for each item on the top arc, we developed support for the various design properties shown and discussed in Figure 1.1 and discussed in Section 1.1, based upon the type of comprehension support the corresponding tool(s) was designed to provide.

²Note that a “later” developer, despite having questions about the initial work, may be the *same* developer as the initial author – research has shown that code comprehension by an initial developer about their own code, after some time has passed, can be as difficult as understanding unfamiliar code written by a different developer [159].

We began by exploring *annotations* as a vehicle for associating a user-authored note with a document as a way of crowd-sourcing documentation notes to lower documentation-usage barriers. To this end, we designed the Adamite system (Chapter 3), a web browser-based annotation tool specifically designed for annotating software documentation with features designed to combat known documentation usage barriers. We chose to focus on documentation, given the large body of software engineering research discussing particular pain points that we expected meta-level notes to address. For example, some prior literature has discussed the challenges in finding pertinent information in documentation and how this information may be fragmented across multiple locations within the documentation [4, 52, 249] – to address this barrier, Adamite allows developers to add multiple text anchor points to their annotation, such that all relevant parts of the documentation may be connected to one note. Adamite demonstrated the efficacy of utilizing developer-authored meta-information presented as annotations to assist later developers in their own programming tasks – developers using Adamite annotations performed significantly better on a programming task.

Adamite was successful in helping later developers utilize meta-information about code to overcome their information barriers – however, we found that the initial authors did not experience a significant effect, despite appreciating the annotation tool. Given the need for more support for the initial author in utilizing their meta-information and the increased need for information tracking support when actively writing code in an IDE, I developed Catseye, a Visual Studio Code [166] extension for annotating code (Chapter 4). Catseye not only supported adding free-form text to code that was abstracted away from the original source code (as opposed to code comments) but also supported other developer information-tracking activities including collecting output data and lightweight versioning of code, with the option to annotate these other forms of meta-information. Developers using Catseye performed better when attempting to keep track of information while debugging some purposefully-confusing code.

In active usage of Catseye, it became clear that the extremely mutable nature of code can lead to information within the annotations quickly becoming out-of-date. Keeping annotations attached to the correct code is surprisingly non-trivial – while some simple text edits can be easily interpreted as small arithmetic changes to the line number and offsets along the bounds of an annotation’s anchor, for large scale or batch changes, such as a Git merge, the question of what, exactly to do with the annotation and where, if at all, should the annotation be re-attached becomes more complex. Further, given that the nature of these annotations is often, but not always, ephemeral, therein lies a need to support developers in actively managing their meta-information. This lead us to wonder how we could support developers in *curating* their annotations (Chapter 5), given developers’ perceived value of the annotations and to explore tooling approaches within Catseye that would allow

developers to perform this curation, including algorithmically re-anchoring annotations after significant file changes and batch-level operations to act on annotations for, e.g., removing large sets of out-dated annotations.

Our exploration into how annotations go out-of-date led to the insight that the relationship between annotations and their anchor points can be a signal of out-of-dateness. Considering out-of-dateness is a large problem in documentation as well, the insight that the connection between code and text is both always changing and the accuracy of the link can be used as a signal for the trustworthiness of the document led to us extending Catseye for long-form documentation with the system Sodalite (Chapter 6). Long-form documentation can be a useful information source, given that it co-evolves with the code, but practice shows that does not always happen [143]. Sodalite uses the connection between code and text and its ability to connect the code to the text as a metric for how “healthy” the document is in the user’s current programming context. This system showcases how meta-information can be used for new activities beyond simply displaying data.

Adamite, Catseye, and Sodalite were all successful in supporting developers in authoring their own meta-information, but, in their designs, were reliant upon the programmer choosing to actually write some information. Further, what developers have to say about code is only one type of meta-information and is incapable of answering some of the types of questions developers may have, such as *when* a particular change was introduced. To expand the types of questions we could support answering with meta-information and to lower the cost of authoring useful information, we developed the Meta-Manager, a Visual Studio Code extension with a supplementary Google Chrome extension, that automatically listens for and captures editing events of interest in both the editor and the browser (Chapter 7). Some editing events of interest occur *within* the editor including copy-pastes between code patches, which may inform where some code came from and when (i.e., provenance), and code commenting, which may show the author trying different solutions or the author adding or changing documentation. Other edits of interest occur in the *web browser* – we are particularly interested in copy-pastes from the browser (e.g., code snippets taken from GitHub, ChatGPT code snippets, or answers from Stack Overflow posts) into the editor, which may inform *why* some code is written the way it is, given relevant Google search queries and/or visited web page(s). AI-generated code from ChatGPT, in particular, contains a rich collection of interesting additional meta-information about the development of the code, such as the initial query, chat title, and possibly a back-and-forth conversation between the user and the agent, which may inform what the developer’s original intent was. Indeed, in our study developers were able to use the Meta-Manager to answer otherwise unanswerable questions about code using this rich version history, with information about AI-generated code presented as code meta-information being particularly valued by participants.

With the rise of large language models in recent years, meta-information is only

going to become more valuable, not only as a means for potentially training these models, but also as a mechanism for potentially accelerating a developer in their individual sensemaking journey. Code meta-information can be used as a dataset to pull from when utilizing LLMs for software development tasks as a form of retrieval-augmented generation (RAG). A context in which this ability to revisit and reason about code changes may be particularly valuable is print debugging, such that developers can use that history to understand how or why some code no longer works and compare that code to the current code. In order to support that use case, I extended the capabilities of the Meta-Manager to allow for direct querying on its historical dataset while expanding the type of information it captures, including outputs produced by `console.log`. I performed a small qualitative needs-finding interview to determine what questions developers would most like to ask their code and output history and designed functionalities within the new Meta-Manager (hereafter referred to as “MMAI”) to support those tasks (Chapter 8). With this new architecture in place, I showcase that utilizing meta-information for RAG can accelerate otherwise onerous tasks and assist in reasoning when debugging.

In conclusion, I reflect on the contributions made and the insights gained through this body of work, and discuss what outstanding questions and subjects exist for future research (Chapter 9). Broadly, this thesis points towards a future in which meta-information, including historical details and user-authored notes, are tightly coupled given anchoring to support the personalization of, shareability of, and comprehension of dense and often confusing materials. I envision a space in which source code can be shared like an old textbook handed down from a mentor, annotated with all of the rich insights, thoughts, provenance data, and questions I would otherwise be unaware of in order to understand the full story. This information can then be utilized by LLMs for summarization, retrieval and querying, and further personalization tailored towards my specific sensemaking needs given my particular tasks. The software engineer of tomorrow is going to be working in a space where, perhaps, writing code is not as paramount as understanding it and meta-information points towards a way of capturing and presenting the information to assist in that understanding.

Chapter 2

Background and Related Work

Developers are tasked with managing many different types of information when both authoring and making sense of code. In the following sections, I discuss prior literature about how developers make sense of code, tooling support for that process, meta-information about code (including documentation), and prior systems to support the creation and management of some of these types of meta-information.

2.1 Making Sense of Code

Researchers in both Human-Computer Interaction (HCI) and Software Engineering have extensively studied how developers make sense of or comprehend code [27, 29, 58, 66, 127, 129, 130, 131, 134, 159, 187, 213, 214, 223]. These studies are typically situated in the context of understanding unfamiliar code [58, 127, 134, 159, 214], which is a common and significant challenge for developers [174]. For example, developers need to understand unfamiliar code when learning a new API [56, 63, 66, 108, 151, 174, 176, 212], joining a new code base [18, 51, 113, 210, 235, 248], adapting code from online sources [12, 29, 151, 177], taking over a code base from a departed coworker [178, 209, 210], and so on. Some of the challenges developers must overcome when understanding unfamiliar code include maintaining their task awareness [157, 170, 189, 191, 238]; developing mental model of the code [27]; questions, hypotheses, and facts learned about the code [78, 138, 226]; the codes' different versions and outputs [119, 142]; and "working set" of code patches [27]. How developers actually complete these difficult information tracking tasks is further confounded by developers' varying experience levels [64, 133, 140, 145, 159], learning styles and preferences [16, 49, 64, 145, 164], and the nature of their tasks [197]. Even when an implementation task is small or the code is more familiar, developers must continually keep track of these varied information types for maintaining software and have varying degrees of success [126, 147].

2.1.1 Developer Tools

Given the diversity of types of information developers must keep track of, a number of research tools have focused on supporting these different activities. One challenge developers experience when making sense of code is navigating through

code and finding and relating code patches of interest (sometimes referred to as the “working set”) [27, 45]. Researchers have investigated various solutions to ease navigation costs including clustering code patches as “bubbles” [27] or in a dedicated workspace [1, 45], changing navigation techniques away from clicking and into zooming and scrolling [59], and augmenting code comments to serve as navigational way points [89, 237, 238, 247]. Another way of easing navigation that is now widely adopted is a “thumbnail” image attached near the scroll bar that shows the whole file in a minimized format [57]. Other projects have attempted to ease navigation through suggesting code patches to support various tasks including completing a pull request [79, 118], fault localization [73, 196, 197] and for helping newcomers to code projects find where to begin [17, 51, 100]. In all of these cases, the code patches have a higher-level semantic meaning as to why they are relevant to the developer – a relationship which is only sometimes made clear given the tool. All of my tools attempt to make the relationship between this meta-information and its corresponding code traceable and bi-directional to assist in comprehension and navigation.

Relevant to my work are projects that have focused on leveraging the large amount of already-written code and developer resources that exist online to assist in overcoming cognitive barriers incurred when programming. These projects commonly suggest code snippets taken from documentation [28, 162, 182], open source repositories [245], or question-answer forums [110, 273], given a developer’s query. The underlying rationale for this design, namely that developers care about code first and foremost, also drives my design approaches, but my systems capture other forms of information beyond just code since code, alone, cannot answer every type of question. Other systems lower the barrier of foraging for programming-related information on the web through integrating web-based functionalities (e.g., web search) into the IDE [47, 83, 92, 188]. My systems, especially the Meta-Manager, attempt to also better connect the activities that are naturally happening both in the editor and the web browser, but do not require the user to adopt a new browsing or development environment. Meta-Manager also has an advantage over these prior tools through not only connecting web activity and the IDE but contextualizing these activities to particular code versions and supporting search to support foraging [233] through queries [217], such that the developer can understand not only what happened with some code, but when and why.

Over the last few years, a new class of developer tools have risen to prominence for making sense of code: LLM-powered coding assistants. Some of these popular tools include ChatGPT [183] and GitHub Copilot [82]. The allure of these tools and, subsequently, their rapid adoption is perhaps unsurprising – unlike many prior developer tools, they are easily accessible, requiring minimal to no set-up, and can immediately provide a benefit to developers through supporting both code generation and code understanding tasks. Some preliminary reports even suggest that the nature of software development itself may be shifting as these tools become more popular [44], with programmers reporting around 31 percent of their new code comes

from AI tools [148], they use these tools for high-level guidance on how to approach a task [122], and they spend more time reviewing AI-generated code [24, 169]. Since the public release of these tools in late 2022, there have been many studies looking at the quality of the generated code [55, 71, 242, 267, 272], novice versus expert usage of these tools and the potential harm they may introduce for novices and their learning [179, 202], along with how these tools impact perceived productivity [230, 257]. These tools can also be used to “short-cut” code comprehension through directly answering questions about unfamiliar code [36, 115, 227], including APIs [175], and assisting in cognitively-demanding tasks such as debugging [43]. My work extends this line of inquiry through utilizing developer meta-information as a repository unto which an LLM can query upon to support code sensemaking when debugging.

2.2 Meta-Information About Code

2.2.1 Writing About Code

Of particular relevance to my work, especially the earlier work around annotations is how developers write about their code. One strategy developers employ when keeping track of information is writing down, either through code comments or external notes, what they want to keep track of. In the case of code comments, research has explored what types of information are in code comments [220, 237, 240, 258], whether the code comments are actually useful [26, 201, 258], and how these comments are later used [74, 204] and cleaned up [240, 241]. An analysis of 2,000 GitHub projects found 12 distinct categories of comments developers write and found that information designed for code authors and users appeared less often than more formal types of documentation, suggesting that developers are not frequently commenting for their own benefit or these comments are removed prior to submitting the code to publicly-viewable repositories [220]. Other meta-information about code that is sometimes included in code comments are links to where some code originated from, in the case that the code was copy-pasted from online [11, 93].

Developers also employ note-taking for tracking information about code [46, 151, 157, 159, 189]. These notes are commonly used as memory-aids [46, 151, 189] both for themselves and for communicating insights to other teammates [39, 159], for keeping track of useful resources [151], and for keeping track of their progress on a development task [157, 189]. Notably, these notes often lose their initial utility given their lack of context, while others are authored with the full expectation that they will be thrown away [151] as a form of “information scrap” [21]. Of particular relevance, particularly to Catseye, are tooling solutions for developers that directly support note-taking to assist in code comprehension, such as Codepad [190] or Connotator [114]. Connotator explores supporting different levels of code annotations to support program comprehension, but does not allow annotating partial selections

such as expressions or strings nor does it allow the developer to write free-form content, instead only supporting pre-determined tags [114].

Code development and coordination with a software team commonly results in the creation of other types of written works. This includes version control system (VCS) commit and pull request messages [149, 154, 246], computational notebooks [94, 120, 198, 254, 255], bug reports [23, 205], and code review messages [9, 87, 135, 172, 173, 216]. Developers may also write about their code with the intent of helping others, whether that be a blog post [184, 185, 192], question-and-answer forum post [8, 262, 276], or tutorial [94, 96, 192, 244, 253]. Most of these written artifacts suffer from the same challenges that other written works about code face in that they are typically abstracted away from the original code context, thus making them difficult to discover or glean answers from.

Documentation

Perhaps the most ubiquitous and studied type of writing about code is software documentation. Much prior literature has investigated what information is written in documentation [97, 158], what the problems are with that information [3, 4, 35, 68, 143, 163, 180, 212, 214, 249], how this documentation is authored [53, 222, 236], and automating documentation processes to offset authoring costs and standardize the information present in documentation [2, 91]. Notably, the majority of this documentation work is in the context of software documentation for end users of a software library, e.g. API documentation [174]. Slightly less work has focused on documentation created for other developers working on the same code base, such as internal documentation about the code base [210, 222] or open source on-boarding files [244], where the primary goal is to help other developers understand and contribute to the code base.

Another form of documentation relevant to our previous discussion of navigation are diagrams that attempt to visualize the digital repository and space of code into a more physical and spatial form. Examples include code maps [60], which describe each code component as a connected node along a hierarchical graph, and code cities [259], which visualize object-oriented programs as cities that can be traversed and interacted with. Both of these approaches attempt to more directly visually situate a developer in the documentation with respect to the part of the code they are working on. Sometimes these diagrams are made for social purposes on the fly for communicating with teammates, given the teammates' respective information needs and knowledge [39].

Once some software documentation is made, researchers have studied how developers maintain those documents and use that information. In studies of usage, researchers have identified many problems of documentation that lead to the documentation being less trustworthy [153], including questions about how up-to-date the information is [4, 143, 249] and how complete the information is [212, 214].

Maintaining documentation has also been found to be challenging, with developers reporting on the costly time spent on updating documentation [249] and a lack of clarity on where and how to update the documentation given code changes [76].

2.2.2 Other Meta-Information

Code History

Developers keep track of many other types of information that are not or cannot be written down. One ubiquitous form of meta-information about code is code history, i.e., code versions across time. Most software engineering utilizes a VCS (e.g., Git) for keeping track of code releases and to make collaboration across potentially many different versions of code possible. However, in between formal check-ins of code, developers have reported a need for keeping track of their code versions. Whether that be in the context of data science programming where intermittent versions may represent small-scale experiments [95, 119, 120, 224] or when maintaining a code project and trying different solutions for fixing a bug [269] or adapting a code example [29], developers employ a variety of methods for keeping track of these intermittent versions such as commenting out the code [119]. Some systems (e.g., [186, 255]) have leveraged discussion threads about code and mapped the threads to the relevant code – an approach not dissimilar from Catseye, Meta-Manager, and Sodalite. Both these less formal intermittent code versions and committed code versions can serve as meta-information about the “current” version of the code in service of, e.g., understanding how some code evolved over time.

To support sensemaking of code history, prior research has explored creating visualizations of code history. Most code history visualizations reserve the x-axis for representing time, but the y-axis and presentation of the code and edits varies. Some tools adopt a stream visualization in which each “stream” represents a block of code [261]. The stream expands, retracts, and moves up and down along the y-axis as lines of code are added, removed, and the location of the code moves, respectively. Notably, this visualization approach was inspired by systems visualizing other historical, text-based data including Wikipedia page edit histories [252] and Google Doc histories [256]. Other visualizations reserve the y-axis and data points along the visualization for edits that happened to some code at a particular time [269, 270, 271]. Some visualization approaches remove the timeline-style presentation in order to support other activities, such as Quickpose [206], which uses a node-based graph to support actively modifying and annotating the history. Seesoft [65] and Augur [80] choose to visualize the line of code itself and imbue additional meta-information such as who most edited the code and what part of the code structure is this history a part of. Meta-Manager and MMAI differ through combining the approaches of both aggregating information through a stream visualization and showing lower-level details in an additional widget.

Outputs and Logs

Another form of code meta-information closely related to code versions are outputs generated by code. Outputs may be as simple as a `console.log` message (i.e., log statement and value) or as complex as a full graphical user interface. In either case, developers must often keep track of how these outputs change with respect to their code edits [119]. This may happen when the developer must revert to a prior, functional version after a bug is introduced or to compare differences in outputs with the intention of choosing an optimal version [269]. Some tools have been developed to keep track of outputs with respect to the code versions that generated them [120], while others have focused on making the output itself more useful through additional affordances and improved presentation [112].

Log statements and values, in particular, have been extensively studied [81]. Logs are particularly challenging for developers to manage given the scale of logs often accrued in an enterprise system [219], their presentation in plain text making them hard to work with [112], and their decoupled nature from the code and runtime context in which they were generated [62]. Further, log statements, not unlike code comments, often suffer from the “information scrap” [22] challenge, in that they can be written for a single, throw-away line of inquiry as part of print debugging [19]; otherwise, they may be meant to exist indefinitely to capture code performance metrics [81], but their intent from the outset is not always apparent. Given these challenges, some prior tools have attempted to mitigate these problems through reducing scale [219] and changing log presentation [112]. Despite the short-comings of logging, they are nonetheless ubiquitous, especially in the context of “print statement debugging” [144, 260]. Given this ubiquity as a form of code meta-information and their applied nature when debugging, we developed MMAI as a way to better manage and make sense of their contents using AI while leveraging the architecture of the Meta-Manager to address challenges of scale and lack of context.

In our work, we design tooling approaches to support the authoring and management of programming-related meta-information for sensemaking. Whether this meta-information is in service of externalizing and tracking some thought about code, an intermittent code version, a piece of useful documentation, or the history of a function, this work explores how to properly capture, present, and anchor that information to make it more useful for an individual user and across time.

Chapter 3

Adamite: Meta-Information as Annotations on Documentation

This chapter is adapted from my paper:

[104] Amber Horvath, Michael Xieyang Liu, River Hendriksen, Connor Shannon, Emma Paterson, Kazi Jawad, Andrew Macvean, and Brad A. Myers. “Understanding How Programmers Can Use Annotations on Documentation”. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 16 pages.

3.1 Overview

Application programming interfaces (APIs), including libraries, frameworks, toolkits, and software development kits (SDKs), are used by virtually all code [174]. Programmers at all levels must continually learn and use new APIs in order to complete any project of significant size or complexity [63]. In learning APIs, developers depend upon the documentation, including tutorials, reference documentation, and code examples, along with question-and-answer sites like Stack Overflow [143]. However, developer documentation is known to be problematic with common and significant issues including incompleteness of information, out-of-date information, and fragmented information [3, 4, 212, 249].

In this chapter, we investigate how developer-authored meta-information about documentation, presented as annotations, may help developers overcome some of these barriers. We hypothesized that the highly-contextualized nature of developers notes when anchored to text within the documentation, can help the initial annotator better utilize their notes about documentation. Further, when given proper tooling support, this information can be leveraged by later users of the documentation to overcome their documentation barriers.

To explore the concept of annotations as a way of supporting short notes on documentation that are useful both for the author and for later readers, we started with a preliminary lab study that explored the concept of annotations on documentation using an off-the-shelf Chrome extension, Hypothesis [111], and then we performed a corpus analysis of annotations on documentation created using Hypothesis. Given what we learned from these preliminary analyses, we developed our own documentation-specific annotation tool, Adamite¹. Next, we ran a two-pass user study where we explored the kinds of annotations developers *authored* when learning a new API and then had another set of developers *read* those annotations while attempting to complete the same API learning task using Adamite. We compared these participants to a *control* condition which had no annotations. From these studies, we provide evidence that annotations are useful in helping developers overcome documentation-related issues.

3.2 Preliminary Studies and Design Goals

3.2.1 Lab Study with Hypothesis

To explore the efficacy of annotations as a useful learning device for API learning tasks using documentation, we ran a preliminary study where people learned an unfamiliar API while using Hypothesis [111]. In summary, we found that developers are able to use annotations in the ways we envisioned, but that annotation authoring and reading could be improved for developers by adding additional tooling features.

Study Design

The preliminary study had two distinct phases: the first phase was focused on understanding how developers *author* annotations during an API learning task, while the second phase focused on how developers *read* annotations that are already attached to documentation. In each phase, the 4 participants completed an API learning task adapted from my previous study [103] that required them to forage through API documentation to complete a programming task. The full set of study materials can be found in Appendix C, Section C.1.

In the *authoring* condition, participants were given the documentation with no annotations, and instructed to add annotations when they learned anything useful, had questions about the content in the documentation, or had any other thoughts about the documentation. In the *reading* condition, participants were given the same documentation, but with annotations added. The annotations included annotations that I authored that were designed to be helpful for this task given what developers were confused about in a previous study [103], and other annotations that were

¹Adamite stands for Annotated Documentation Allows for More Information Transfer across Engineers and is a green mineral.

designed to be “distracting” to simulate the more realistic case where not all annotations would be relevant. In total, we had 23 “helpful” annotations and 44 “distractor” annotations, totalling 67 unique annotations.

Results

In the authoring condition, the 4 participants together authored a total of 19 unique annotations. Each participant, on average, authored 4.75 annotations, with annotations averaging 4.41 words. Hypothesis also allows users to simply highlight a piece of text on a web page without adding any text content to the anchor (hereafter referred to as “highlight” annotations) — out of the 19 unique annotations authored, 5 were these simple highlight annotations.

Many of the annotations that participants created showed a part of the documentation that illustrated how to achieve some part of their current task. Other annotations served as task reminders or open questions the author had about the documentation content.

In our analyses, we also found that participants had many questions about the documentation (on average, 10.8 questions per participant) which were not annotated. While trying to answer these questions, participants routinely encountered more confusing information, resulting in them losing track of their original questions. Given this confusion, participants struggled to answer their questions, with only 29% of questions definitely answered. Notably, Hypothesis does not have any way of marking or following up on a question.

From this initial preliminary study, we found evidence that annotations may enhance the original text. Additionally, we found support for different types of information needs that are not directly supported by Hypothesis, such as keeping track of open questions. In the reading condition, we also learned through our “distractor” annotations that annotations need to be easy to skim and relatively short in length and we need more support for discovering annotations, either through more anchor text points or filter and search.

3.2.2 Corpus Analysis of Hypothesis Annotations

In order to supplement our preliminary study, we queried Hypothesis’s API to get a list of public annotations which developers have already made on official API documentation including public APIs from Google, Microsoft, Oracle, and Mozilla, along with other developer learning resources including Stack Overflow, W3Schools, and GitHub.

Across these sites, we found 1,995 public annotations. Of the 1,995 annotations, 196 were questions about the content of the documentation, and 995 were highlight type annotations. Of the 1,000 annotations with content, 16 of the annotations were to-do items the author wanted to follow up on, 43 pointed out problematic aspects

of and potential improvements to the documentation, and 79 were created to specifically call out important or useful parts of the documentation². These annotations were authored by 298 unique users (average = 6.694 annotations per user, minimum = 1, maximum = 677) across 1,143 unique web pages. The authored annotations were, on average, 8.79 words long and were anchored to text that averaged 12.35 words.

We believe some of the 1,000 annotations that contained content could benefit other developers with additional tooling support. For example, a Hypothesis user annotated the text “you can pass the path to the serve account key in code” and asked how they can do that. This user then later annotated a code example showing how to achieve this behavior at a different point in the documentation and said “finally found it”. While these two annotations depend upon one another in order to make sense and point to different parts of the documentation, Hypothesis does not allow for these annotations to reference one another, suggesting a need for better tooling support for multiple anchors for annotations and keeping track of open questions.

These annotations provide support for our claim that some developers are willing to write annotations and attach them to documentation, as they are already doing this. Moreover, the annotations we found follow some of the patterns we identified in our preliminary study, such as open questions and issues. However, Hypothesis’s general-purpose annotation system does not have enough support to effectively utilize these annotations.

3.2.3 Design Goals

Given prior literature and what we discovered in our preliminary studies, we developed the following design goals for our system:

- **Support developer note taking through annotations.** Developers sometimes take notes on what they have learned [151, 157, 159, 189], even when using documentation (Section 3.2.2, [151]).
- **Support developers’ question-asking and answering.** Developers have many questions about unfamiliar APIs and their documentation ([63, 226], Sections 3.2.1 and 3.2.2).
- **Support diagnosing documentation issues.** Developers commonly identify documentation issues including obsolete information [4, 249], incorrectness ([4, 249], Section 3.2.2), incompleteness [4, 37, 52, 211, 212, 249, 278] and ambiguities [4, 211, 212, 249] but have no way of sharing that information.

²These counts were generated by counting instances of phrases like “incorrect” and “todo” in the annotation content, then manually reviewing all of the annotations that contained those phrases to determine if they were actually e.g., an issue or todo item.

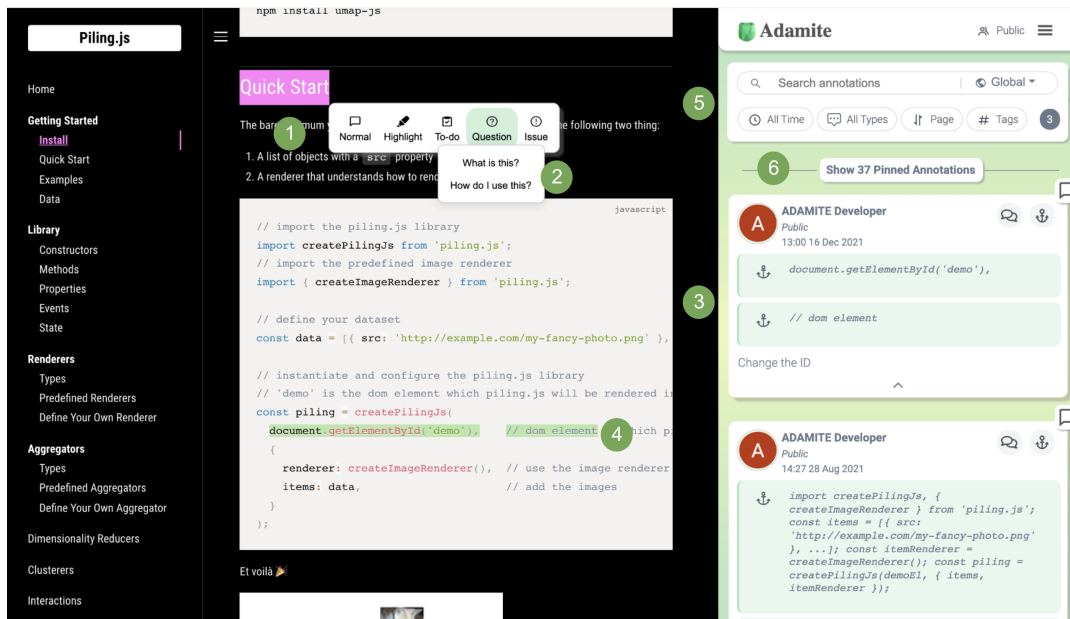


FIGURE 3.1: Adamite’s sidebar (on the right) open on an already-annotated web page in the browser. (1) shows the pop-up for when a user selects some text – at this point they can begin creating a new annotation by selecting an annotation type. (2) shows the menu of question annotation prompts that users can choose from. (3) shows a published normal annotation with two anchors. (4) shows how the annotated text appears on the web page. (5) shows Adamite’s search and filter pane. (6) shows the pinned annotation list button.

- **Support developer task-tracking.** Developers take notes on open tasks that they must work on, especially when interrupted [189], and occasionally take notes on tasks they must complete that are related to parts of the documentation they are reading (Section 3.2.2).
- **Support connecting related parts of the documentation.** Developers need to build up a mental representation of an API [108, 124, 151] and connect related resources [4, 52, 249].
- **Support better discovery of important parts of the documentation.** Developers, especially selective [30] and opportunistic [29] learners, want to quickly find information that is relevant to them ([163], Section 3.2.1).

3.3 Overview of Adamite

We designed Adamite, a Google Chrome browser extension, specifically to help developers keep track of important information, organize their learning, and share their insights with one another. To create an annotation, a developer, who we call the annotation “author” simply needs to open the Adamite sidebar, highlight some text on the web page (called the “anchor”), select the type of annotation, optionally add text to a rich text editor that appears in the sidebar, and click on the publish button. Once published, the text that the user annotated will be highlighted on the web page and the annotation will appear in the Adamite sidebar – see Figure 3.1.

Users may also add tags and additional anchors to the annotation. Annotations can be published publicly, privately, or to a group of Adamite users. Once an annotation has been published, it may be replied to by others, and edited or deleted by the original author. Clicking on the anchor icon on the annotation will scroll to the part of the web page the annotation is anchored to (or will open a new tab if the anchor is on a different page) – conversely, clicking on the highlighted text on the web page will scroll to the corresponding annotation in the sidebar.

One goal of Adamite is helping developers structure and share what they learn in the documentation in a way that is useful both to themselves and for later developers. To achieve this, we developed annotation types. In addition to the typical “normal” (with a user-written comment) and “highlight” (just the anchor and no comment) annotations, Adamite supports question, issue, and to-do annotations. We chose these three annotation types to assist developers in keeping track of their questions, to point out and possibly attempt to rectify issues found in the documentation, and to help them keep track of their tasks. Issue annotations have a button intended to alert key stakeholders, such as the documentation writers, of the described problem with the documentation. Question annotations are stateful, meaning unanswered annotations will stay available until the developer either marks the question as “answered” (at which point the answer will be appended to the original question), or marks the question as “no longer relevant”. To-do annotations are also always available until they are marked as complete.

Question and to-do annotations are always available using Adamite’s “pinning” mechanism. Web annotation systems typically only show annotations that are on the user’s current web page. However, considering that documentation may be spread across many pages and developers may visit many web pages when attempting to complete a programming task, we added in the ability to *pin* an annotation, such that it is always available in a list at the top of the sidebar. To-do and question annotations are pinned by default, since the developer is unlikely to find their answer or finish their task while they are on the same web page.

Developers can further structure and connect their information using *multiple anchors*. A common documentation problem is fragmented information, so developers can compose together related pieces of information using the annotation as a joining method. This feature can be used to connect parts of the documentation that the user feels should be presented together, or to better contextualize their annotation. Developers may also use anchors as a way of collecting multiple parts of the documentation that they feel are related to one another given the developer’s task and their evolving understanding of the API.

For later users of the annotated documentation (who we call annotation “readers,” but can be the same person as the annotation authors), it is likely that not all of the annotations are relevant to what the developer is trying to do. To help readers find the most relevant annotations, we support search (using Elasticsearch [67]) and filters (see Figure 3.1-5). Readers can search across a web page, website, or across all

of Adamite’s annotations and filter on the annotation type, when the annotation was created, and what tags the author has tagged the annotation with. Readers may also sort the annotations by their location on the page or by the time at which the annotation was authored. Adamite also follows the design of other web-based annotation systems by only showing annotations on the current web page, by default.

3.4 Lab Study

In order to understand the role that annotations play in developers’ documentation usage while learning a new API and using Adamite, we ran a lab study with three conditions to understand how developers create and use annotations. Participants in one condition *authored* annotations while completing an API learning task, and participants in the second condition *read* these participant-authored annotations. The third condition was a *control* condition where participants completed the same API learning task using just the documentation. The lab study consisted of a training task, a programming task, and a survey to assess the participant’s background. The full protocol for the lab study across each condition can be found in Appendix C.2.

3.4.1 Method

Training

Each condition included a training exercise using Tippy, a React library for making tooltips, and its documentation to either familiarize the participants with Adamite and its functionality (Adamite conditions) or to familiarize them with thinking aloud while reading through documentation (control). Participants in the Adamite conditions learned how to create an annotation of each type, reply to an annotation, add an additional anchor to an existing annotation, search, filter, edit and delete an annotation and practiced thinking aloud while performing these tasks. The control condition practiced thinking aloud when they had a question, found an answer to their question, and identified an issue in the documentation.

Task

For the task, participants were asked to complete an image aggregation and organization task using Piling.js (hereafter referred to as “Piling”), a JavaScript library for handling visualizations [195]. Piling was chosen as it is a relatively small library, meaning the participants would have adequate time to gain a high-level understanding of the library during a lab study.

The task was to use Piling to take a set of four provided images and render and sort the images (see Figure 3.2 for the output and detailed steps and Appendix C.2.5 for a link to the code). The task was chosen as, despite its apparent simplicity, it requires the participant to learn some of Piling’s core concepts. Participants were objectively graded upon how many of the 4 steps they were able to complete correctly.

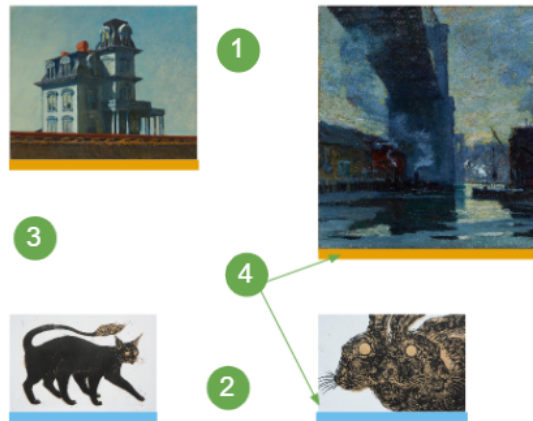


FIGURE 3.2: The correct output for the task. Each number refers to the step number. (1) creates and renders the 4 images. (2) puts the images in 2 rows. (3) arranges images by a user-defined property using the `orderBy` method. (4) requires the user to set a label on their data, such that elements with the same label will have a matching stripe along the bottom of the picture.

To start, participants were given a JavaScript file containing comments stating the goal of each step.

In addition to completing the programming task, participants were asked to think aloud and to pretend as though they were in a small team learning Piling. Dependent upon the condition, further instructions differed slightly. *Control* condition participants were told that they needed to relay what they had learned to their teammates in whatever way they would normally do so, such as note taking. *Adamite authoring* participants were instructed to create annotations with any questions or thoughts they had about the documentation, issues they found in the documentation, and thoughts they wanted to follow up on and that these annotations would be shared with their future teammates. Each authoring participant started with an un-annotated version of the documentation. In the *Adamite reading* condition, participants were given annotated documentation and were told to pretend that the annotations were created by a teammate who had already learned Piling and were instructed to speak aloud when an annotation was helpful or unhelpful. Participants in the reading condition were not required (but were allowed) to create annotations or interact with the annotations present in the documentation. The full set of study materials can be found in Appendix C.

Annotation Selection

In choosing annotations to include in the reading condition, two researchers separately coded each annotation created during the authoring condition for whether or not to include it. Inclusion criteria included identifying matching annotations across participants to select which one was the clearest, most appropriately anchored, and concise – qualities informed by our preliminary study and others [5]. The predominant cause for the majority of annotations to be removed was redundancy – participants commonly annotated information related to the first two steps of the task




FIGURE 3.3: The number of annotations removed for each reason, along with the annotations kept, out of the 91 total annotations. 2 highlight annotations were retained as the users edited them to add text, making them semantically identical to normal annotations.

(see Figure 3.3)³. We also excluded annotations that the participant later stated were incorrect or that the participant deleted, and annotations that lacked sufficient context (including all highlight annotations). Finally, we omitted to-do annotations, as they are designed only for the original author’s usage. The researchers had a 71% agreement – in the cases where the researchers disagreed, they had a discussion until agreement was reached. Through this process, we were left with 31 annotations. One additional annotation was added by myself to assist with the final step of the task, as no participant in either the authoring condition or control condition got to that point of the task. After this process, we had 32 annotations, with 18 normal type annotations, 10 issue type annotations, and 4 question annotations, 3 of which were answered⁴. We did *not* omit any annotations due to relevance or correctness, since we wanted to leave in anything that at least one participant wanted to comment on to be more realistic.

Participants

We recruited 31 participants using departmental mailing lists at our university and social media. One participant could not finish the study due to technical difficulties, so we only report on the 30 who completed the whole study. Each condition included 10 participants and were randomly assigned between the authoring and control condition – the reading condition occurred after the other two conditions so all remaining participants who signed up were assigned to that condition.

All of the participants were required to have some amount of experience using JavaScript, not to have used Piling before, and to have been programming for at least

³Note that this is due to being a lab study – in a realistic situation, people would likely read an existing annotation and not create a redundant one.

⁴We included one unanswered question to account for the realistic situation that not all questions would be answered and because the question asked was a common question among participants – notably, answering this question was not necessary for succeeding in the task.

1 year (actual minimum: 1 year, maximum: 20 years, average: 7.98 years). The participants' professions included graduate students in computer science-related fields, user experience researchers with a computer science background, and professional programmers. The gender composition of our study was 19 men, 10 women, and 1 non-binary person. Participants across each condition had a similar amount of JavaScript experience and years of programming experience.

All study sessions were completed remotely using video conferencing software. Participants were audio and video recorded, and each participant's session took approximately 90 minutes, with 45 minutes of that time allotted for the programming task. Each participant was given access to the Piling documentation and a CodeSandbox.io [31] project which had JavaScript, HTML, and CSS files with Piling installed and a photo of the output, along with written-out steps for the task. Participants were compensated \$25 for their time, save for 2 participants who elected not to be compensated.

Analysis Methods

Across all of the conditions, we objectively graded participants on whether or not they succeeded in completing each of the 4 steps outlined in the task instructions. In the Adamite conditions, we analyzed the video recordings and log data to count how many annotations participants authored, and how often they interacted with Adamite and its annotations.

We qualitatively coded the annotations developers made in order to characterize developers' annotating strategies. Using an open coding method, two authors coded the normal type annotations by independently coding each annotation and refining categories based upon their individual codes. For issue and question type annotations, we coded the annotations dependent upon what issue in a list of commonly defined documentation issues [4] was identified in the annotation (issue type) or what issue caused the participant's confusion with issue types including: incompleteness, fragmentation, incorrectness, poor code example, and ambiguity. Two of the authors independently coded the annotations and reached 75% agreement when coding the issue annotations and 73% for the question annotations – remaining annotations were discussed until agreement was achieved.

In the annotation reading condition, we analyzed how often participants said that an annotation was helpful or unhelpful in order to better understand what annotations succeeded in helping participants. We calculated average helpfulness by how many participants said an annotation was helpful and dividing by how many participants encountered the annotation. We ranked annotations from most helpful to least helpful by how many participants said the annotation was helpful subtracted by how many said it was unhelpful⁵

⁵We chose this method of ranking in order to account for the "impact" of an annotation – for example, if only one participant encountered a specific annotation and found it helpful (100% helpful),




FIGURE 3.4: The difference between reading and each of the other two conditions is statistically significant, but the difference between control and authoring is not. The average number of steps completed is in the center of each box.

In the control condition, we kept track of whether and how the participant chose to relay their information to their teammates. We also referenced the auto-generated transcripts to find and count whenever a participant stated a question.

3.4.2 Results

On average, participants in the control completed 1.5 of the 4 steps, authoring participants completed 1.4 steps, and the reading condition completed 2.5 steps (see Figure 3.4). Participants in the reading condition performed significantly better than participants in the control and authoring conditions (paired T-test versus control, $p < .01$, paired T-test versus authoring, $p < .01$). This provides evidence annotated documentation helps developers in using API documentation.

In the control condition, 1 participant chose to take notes in a Google Doc and 1 participant made comments in their code as notes for their future teammates. 2 participants spoke aloud when they had a thought that they would want to share as a note to their teammates, but did not actually write any notes down. The remaining 6 control participants did not take notes or verbally indicate the intent to take notes at any point during the study. This suggests that without a mechanism for externalizing their thoughts, these 6 participants may not have been able to actually share what they learned and only 2 participants had any artifact to share with their future teammates.

How Developers Annotate Documentation

The 10 participants in the annotation authoring condition created 91 annotations across all five of the annotation types (see Figures 3.5 and 3.6). The annotations that participants authored were, on average, relatively short in length at 9.31 words (minimum = 0, maximum = 34, median = 8). On average, each participant authored 9.1 annotations (median = 8, standard deviation = 4.094, minimum = 5, maximum = 18),

we did not want that to be seen as “more helpful” than another annotation that helped 5 out of the 6 people that encountered it (83.3% helpful).




FIGURE 3.5: The proportions for each type of the annotations made in the authoring condition (out of 91), with the exact count for each type above the bar and the proportion in parentheses.

with the most used annotation type being the normal-type at 32 authored annotations (35.1%). 2 participants in the annotation *reading* condition created 6 annotations (3 normal, 2 highlights, and 1 question), resulting in 97 annotations across all conditions. The rest of the analyses just look at the 91 annotations from the authoring condition.

Considering the large amount of normal type annotations and how normal annotations can contain nearly any type of information, we sought to characterize the content of these annotations. Through open coding, two coders defined 5 categories – “note to self” in which the participant made a note about the documentation’s content that was primarily for themselves, “explanation of code” in which the participant tried to better explain what a particular code example was doing, “hypothesis” in which the participant guessed about how some part of Piling works, “important to task” in which the participant highlighted a particular part of the documentation as critical for one of the steps of the task, and “other” for any annotations that did not fit into the previous categories. With this categorization, we had 12 “note to self” annotations, 10 “explanations of code”, 7 “hypotheses”, 2 “important to task” annotations, and 1 “other” annotation. The 1 “other” annotation was an annotation with no content that was created purely as a navigational aid. Participants commonly used “note to self” to keep track of their information and contextualize it to their open task, and used “hypotheses” as ephemeral thoughts about the documentation – these activities are more directly supported in our tool, Catseye, within the context of the code.

The “note to self” annotations typically served as reminders to the author to externalize an important detail about the API or a code example. For example, one participant annotated a call to `document.getElementById('demo')` with “remember to change the ID” as a reminder to themselves, as they were in the process of adapting the example. This note could also benefit future users of the documentation as a note that the code example will not work without some modifications.

Unexplained or poorly explained code examples are a frequent problem in documentation [4, 249] and Piling was no different, so our participants attempted to explain some of the code examples and, sometimes, contextualize them to the goals

Documentation Issue	Issue Annotations	Question Annotations About Issue	Percent Questions Answered
Incompleteness	3	6	50%
Fragmentation	4	3	33%
Incorrectness	4	2	0%
Poor Code Example	10	7	43%
Ambiguity	2	12	33%

TABLE 3.1: Counts of each issue and question annotation that identified or was caused by an issue discussed in [249]. Note that two code examples did not work because they suffered from a fragmentation issue, so they are coded both as a poor code example and a fragmentation issue. Similarly, all fragmentation questions were caused by fragmented code examples, so they are also coded as both a fragmentation question and code example question.

of the task. The most helpful and second most helpful annotations are both explanations of code with the most helpful explaining how to use Piling’s row property to create columns, and the second most helpful annotation explaining how the code example for `piling.arrangeBy` works and how to adapt the code example to work using a callback function – both necessary steps for completing the task. Further, in our qualitative coding of issue and question annotations, poorly explained code examples were the most common type of identified issue, suggesting that participants valued the ability to express lightweight thoughts about code in context.

Participants also often annotated code examples which did not work as an issue type annotation, with 10 issue annotations being labeled in the closed-coding as a “poor code example”. Poor code examples were the most frequently identified issues (Table 3.1). These code examples typically did not work either because a variable in the example was undefined and thus the code could not be just copy-pasted (7/10) or because the documentation did not show an output of what the code example actually did (3/10). Some of the undefined variable issues occurred because the variable was defined in a different part of the documentation (2/7) – a documentation fragmentation issue as well as a code example issue. One participant was able to use multiple anchors to suggest where the definition for the variable should be moved to in order to make the code example work.

Developers also had many questions that relate to documentation issues reported by prior studies [4, 249]. Ambiguity and poor code examples were the source of the majority of developers’ questions, which matches the findings reported in [249], with ambiguity, in particular, standing out as a common and severe blocker for developers. Ambiguity and fragmentation issues also resulted in questions that were difficult for annotation authors to answer, with only 33% of questions caused by ambiguity and 33% of questions caused by fragmentation being answered. Considering some participants were able to solve fragmentation issues using multiple anchors with Adamite, this suggests these developers’ questions may have been answered if they had been presented with these annotations. In fact, in the Adamite reading condition, 2 participants had their issue of `aggregateColorMap` not compiling solved by an annotation that used multiple anchors to link to the part of the documentation that defines `aggregateColorMap`.

Issue-type and question-type annotations accounted for roughly half of all the authored annotations. Nearly all issue annotations succeeded in identifying at least one of the issues identified in [249] (see Table 3.1), save for one issue annotation that stated that a particular part of Piling is “super high maintenance for a simple use case”, which is not an issue with the documentation, but with the library itself. Notably, this issue annotation was the third most helpful annotation in the reading condition with participants appreciating that it warned them about a part of the library they were thinking of using, suggesting that issue annotations are useful beyond identifying documentation problems.

Participants also hypothesized about parts of Piling, including how the library worked and what various constructs were relevant to the task. One participant annotated a code example that used an undefined parameter and said “I think that `k` is equal to the number of photos in the data set. I could be wrong though - TBD”. Another participant was trying to determine what function to use to sort their images and annotated `piling.groupBy` with the text “This might be helpful”, but, upon finding `piling.arrangeBy`, annotated that method with “Actually, maybe this”. These hypotheses along with “note to self” annotations demonstrate how annotating can be a lightweight technique for jotting down thoughts as a developer is gaining familiarity with an unfamiliar library.

Considering roughly half of the authored normal annotations are primarily beneficial to the original author (i.e., notes to self and hypotheses) and every participant made a personal annotation (i.e., notes to self, hypotheses, to-do’s, and highlights), we find evidence that annotating is an effective mechanism for externalizing information and helpful for the author. Further, all participants revisited at least one of their annotations at least once (min = 1, max = 36, average = 10.225 revisits per participant), suggesting participants were able to get some utility out of their annotations. We also included 6 notes to self and 3 hypotheses in the reading condition to see whether these thoughts could benefit other developers. Notes to self, in particular, were the most revisited type of annotation by their authors, and participants, on average, revisited these notes 1.8 more times – more than any other annotation type or coded normal annotation types.

Some of the annotations were used in conjunction with Adamite’s other novel features, resulting in the annotations being more useful. The most commonly revisited annotation, a note to self, had 5 anchors with each anchor describing a necessary step in order to properly instantiate the `piling` object. The participant pinned this annotation such that they could reference the anchor steps in their CodeSandbox project (which was open in a separate Chrome tab) – this annotation was also useful in the annotation reading condition with one participant replying to thank the author. This shows that Adamite’s features not only support the creation of lightweight notes but also allow developers to utilize their and other people’s notes in context.

The annotations that participants authored were, on average, relatively short in length at 9.31 words (minimum = 0, maximum = 34, median = 8). The short length of




FIGURE 3.6: Which participants made what type of annotation. A1 through A10 refer to the 10 authoring condition participants. R1 and R2 are the two reading condition participants who created annotations.

annotations makes them relatively easy to author. The annotations included in the reading study averaged 11.94 words and the 10 most helpful annotations were, on average, 13 words long. These results suggest that short notes are able to help future users of documentation, while not requiring a large amount of effort on the author’s part to create. These annotation lengths are also consistent with the annotations authored using Hypothesis in our corpus analysis, suggesting these annotations are similar to the types of annotations authored in the wild.

In terms of task completion, participants in the annotation authoring condition on average completed 1.4 steps out of the 4 steps required to complete the task (standard deviation = 0.69, minimum = 1, maximum = 3, median = 1) (see Figure 3.4). There is no statistically significant difference between the control condition completion rate of 1.5 and the 1.4 in the authoring condition (two-tailed T-test, $p = 0.78$), suggesting that annotating the documentation, while not increasing their performance, also did not require so much overhead that participants were unable to complete the task in the same amount of time as if they had not been annotating. Moreover, considering how often developers’ revisited their notes, specifically their “notes to self”, this suggests authors were able to successfully use annotations as an externalization of their thoughts. Chapter 4 further explores how to make annotations more beneficial for the initial author.

How Developers Use Annotated Documentation

Participants in the reading condition read, on average, 23.7 annotations (including revisiting annotations they had already read, with 72% of annotations read more than once), and read, on average, 15.6 unique annotations. Participants found 45%

of the annotations that they encountered helpful, and only 8% not helpful. The top-performing 6 participants in the reading condition also reported the highest proportion of helpful annotations, suggesting that their success may be attributed to their use of the annotations.

Of the 32 annotations included in the reading condition, the most helpful type of annotation for readers of the documentation was answered question annotations, with, on average, 54% of participants who encountered them stating they were helpful. Normal-type annotations were the second-most helpful type of annotation (avg. 47% helpful) and issue type annotations helped on average 35% of the time. Some issue type annotations were more helpful than other issue annotations including annotations that identified poor code examples, which were helpful, on average, 45% of the time. Even though these issue annotations did not necessarily suggest a solution, they did work in confirming the participant's suspicions that the documentation itself was incorrect and not the participant's implementation. Sometimes, participants found useful annotations through search – participants searched a total of 78 times and 11 of these searches returned an annotation that the user immediately found useful.

The most helpful annotation was anchored to the text “columns” and simply states “Use this to create rows” — a short, 5 word annotation that explains how this property can achieve an effect required by the second step of the task that is not immediately clear when reading the documentation. The second most helpful annotation also succeeded in elucidating how to use part of the API that is relevant to the task through using multiple anchors and clarifying a code example for `orderBy` — a method necessary for the third step. In the reading condition, participants were more successful in completing these two challenging steps, with 9 participants able to complete step 2 and 4 participants able to complete step 3. This increase in performance suggests that participants were able to utilize what the annotation authoring condition learned in order to more effectively complete their task.

Annotations that were not as immediately relevant to the task could also be helpful. Two issue type annotations were the third and fourth most helpful annotations, each warning participants about unhelpful and incorrect parts of the documentation. For example, the fourth most helpful annotation, which helped 4 participants, stated that a code example in the documentation throws an error that a variable is not declared — participants found this annotation useful as it deterred them from using that code example or, if they did use it, reassured them that they were not doing something wrong, since another user had the same problem.

Conversely, the *least* useful annotations were the ones that lacked enough context to be reusable. For example, an unhelpful annotation was an annotation anchored to the text “columns 10”, stating that the default value of `column` is 10, which is redundant with the text of the anchor. The original author annotated this as the reason their 4 images showed up in a single row since the `column` parameter needs to change to make 2 rows, however, the annotation is missing this full context. Future

annotation systems designed to help programmers should explore automatically inferring additional context to make the annotation more comprehensible to later users — if Adamite were to be integrated with the developer’s integrated development environment (IDE), we may be able to capture the code and its output before and after the user created the annotation to better explain why they made the annotation and what they were trying to achieve. Adamite and Catseye (Chapter 4) share the same database, so supporting synchronicity between the two applications is technically possible – ultimately, this path was not fully explored due to a lack of clear design vision in terms of how a between-application, contextually-aware annotation platform could support developer sensemaking. Meta-Manager (Chapter 7) also explores bridging together browser and IDE activities in order to support developer question-asking and answering.

Participants especially appreciated the normal type annotations that were explanations of code, with participants finding them helpful 63% of the time. Code explanations typically elucidated what a code example was illustrating or explained how to adapt a particular code example for the purposes of the task. “Notes to self” were also surprisingly useful, with participants finding them helpful 53% of the time – given that the notes to self typically represented a thought or reminder the developer had about the documentation while completing the task, these results suggest that the participants in the reading study had similar thoughts about the documentation. Conversely, hypotheses were not very helpful, with only 16% of participants finding them helpful – given the uncertainty of these annotations, participants may have found them less trustworthy. These results suggest that explanations of code and developers’ personal notes can be useful if they are framed in a knowledgeable fashion, while hypotheses are more useful for the original author.

3.5 Limitations

Given Piling’s complex documentation, Adamite may not be as helpful when the documentation is simpler or clearer, so the study cannot necessarily be said to apply to those situations. Our lab study was also constrained to a single forty-five minute session, so it is unclear how developers’ API learning and annotation authoring and reading behavior would change over a longer period of time. Piling also has a small user-base, so we have less evidence that Adamite would be useful for APIs with better documentation or APIs with a large user-base that can provide useful crowd-sourced information on Stack Overflow or mailing lists. Future work should see how developers use Adamite in the wild with more popular APIs over a longer period of time.

Considering we selected the annotations to be included in the reading condition, there is an additional limitation that this curating process would not happen in the wild, and, since the last annotation added was created by a researcher, we cannot say that every annotation was participant-authored. The most common reason for

removing an annotation was due to the annotation's content being redundant with another annotation which would be less likely to occur in the real world where users can see other users annotations and will most likely be performing different tasks. While annotations unrelated to the user's task may be distracting, some annotations, such as issue annotations, may be useful to any developer using the construct(s) the annotation references. Further, Adamite supports tagging and filtering which could be used to filter out annotations that are unrelated to what the user is working on. Annotations could also be curated in order to ensure higher quality annotations are more commonly seen using crowd-sourcing methods (e.g., supporting voting and editing other users' contents [6]) which could be added to Adamite.

Adamite as a tool is also limited by its inability to work on dynamic web pages such as Google Docs since dynamic web pages do not have stable anchor points for our highlighting algorithm and the content on these web pages often changes, causing the annotation to lose its original context. Considering developer documentation is relatively static, Adamite works well in this situation, but we do not claim that Adamite will work on more volatile pages. Adamite also does not work on PDFs, despite some documentation existing in the form of a PDF. API documentation is a good use case for Adamite, though, since there are many well-known issues with documentation that annotations can address and API documentation is commonly presented on a website.

3.6 Discussion

Our results suggest that annotations are useful for documentation readers in overcoming some of the known barriers of documentation and that the act of annotating when learning a new API can help developers keep track of their thoughts and open questions. Creating annotations was also useful to the author as a form of self-explanation, which has been shown to be useful for learning in prior studies [32, 40], and these self-explanations, or "notes to self", were useful to others. The novel features of Adamite, especially types, multiple anchors and pinning, helped annotation authors better structure and contextualize their information and helped annotation readers find relevant information.

Participants particularly enjoyed that the annotations had types, and also envisioned future enhancements. 7 participants in the authoring condition noted that they enjoyed the question-type annotations, with 2 specifically mentioning the two built-in question menu items, suggesting that assisting in annotation authoring may be a fruitful avenue for future annotation systems. One participant made an issue type annotation, but wanted the issue to only be shared with documentation writers, while 2 other annotators wanted the "issue" type annotation to be less "confrontational" and instead frame the annotation as a "suggestion" to the documentation maintainers.

Having types for annotations also resulted in two completely separate classes of annotations that users made. As demonstrated in our qualitative coding, the kinds of information developers noted in their normal annotations (i.e., notes to self, important to task, hypotheses, and explanations of code) are very different from the information that developers noted in their issue and question annotations, which were primarily documentation-focused. This suggests that annotation typing is an effective way to elicit information through annotations that may not otherwise be noted.

Reflecting upon Adamite, some of the most significant lessons-learned and open challenges from this work that influenced future work are as follows:

- **Context is key.** One may ask “why annotations and why developer documentation?” – the two research areas feel disparate, yet annotations provide a key function that makes them particularly well-suited to addressing documentation barriers, and that is the context that is inherent to the annotation anchoring point. Consider some of the previously-discussed helpful annotations – they were only helpful because their terse message was bound to a rich information context, that being the documentation’s text. Documentation issues often stem from a lack of context (e.g., “poor code examples” are often missing context about when to use the code, what the code means, etc.) so providing a vehicle by which additional context can be built upon by users of the documentation is a successful approach. Further, annotations often lost their utility when context was not conveyed well, usually due to poor anchor location (e.g., an earlier example about the annotation’s text being redundant with the annotated text). In the later works discussed in this dissertation, we further explore the power of providing tight couplings between user-generated or user-derived data and challenging source materials such as code and documentation.

The tight coupling and its ability to convey context can largely be attributed to the powerful properties of *anchoring*. We additionally gave design consideration for the annotation *content* and *audience*, given the more social, team-based context of the tool and study. This led us to design features such as multi-anchoring for conveying context across time and audience through coupling annotation content with different parts of the source document, and structuring content with different types to implicitly communicate with later developers about the original intent behind an annotation.

- **What is the incentive to annotate?** When beginning this work, we were motivated by the knowledge that developers sometimes take notes when working on code [157, 159], including when understanding documentation [151]. However, the behavior of the control participants, I believe, presents, perhaps, a more realistic look at how this behavior manifests in the “real world”. The majority of control participants did not take notes or even verbally indicate

at what point they would want to keep track of some information. Developing good information tracking practices is difficult for even the most organized and disciplined engineer, so the benefits of using a tool such as Adamite must be high enough that it is worth pausing ones work to annotate. With a system like Adamite that has an implicit crowd-sourcing model, getting enough upfront users to develop a community that can be self-sustaining is a known challenge [6]. Our later work explored this question from multiple angles, including moving working contexts and designing additional tooling to more directly benefit the initial annotation author (Chapter 4), and eschewing direct developer note authoring all-together in favor of capturing already-authored information that may be useful, including web browser and code editing events (Chapter 7) and log statements and values (Chapter 8).

In the creation and evaluation of Adamite, we sought to explore to what extent annotations may help annotation authors and readers in overcoming previously-reported shortcomings of documentation. Through this exploration, we have evidence that developers are able to identify documentation issues using annotations and are able to answer some of their documentation questions. Specifically, our participants were able to identify “poor code examples”, “incompleteness” and “ambiguity” issues, some of the largest blockers when using documentation [249]. Other developers can make use of these answered questions and issue annotations, with answered questions as the most helpful annotation type and explained code examples also helping annotation readers. However, annotations cannot solve every documentation issue. If the API and its documentation are updated, the annotations may go out of date, at which point they may be more harmful than helpful. While our algorithm attempts to reattach the annotation to its anchor point, the annotation content will not change to reflect that reattachment, at which point the content may be incorrect. This tension when updating content between developer-authored meta-information and the original source it is connected to is further explored with our tools Catseye and Sodalite, along with our design exploration of curating annotations through re-anchoring in Chapter 5.

Chapter 4

Catseye: Meta-Information for Sensemaking About Code

This chapter is adapted from my paper:

[107] Amber Horvath, Brad A. Myers, Andrew Macvean, and Imtiaz Rahman. 2022. “Using Annotations for Sensemaking About Code”. In *The 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*, October 29–November 2, 2022, Bend, OR, USA. ACM, New York, NY, USA, 16 pages.

4.1 Overview

With our previous tool, Adamite, we found evidence that developer meta-information, in the form of annotations, was useful in overcoming known barriers in documentation usage. In particular, *later* users of the documentation experienced a significant increase in task completion. However, the *initial* authors did not experience a significant effect. Despite this, we had evidence that developers did appreciate the ability to annotate programming-related documents with their thoughts, questions, and hypotheses that they wanted to keep track of. With this in mind, we wanted to explore how to make annotating more beneficial for the initial user.

One context in which keeping track of information is possibly even more important than in the context of using developer documentation is when actively programming. Prior work has discussed some of the challenges developers must face when making sense of code. A developer must maintain their own task context [170], while also keeping track of the various questions [226] and hypotheses [161] they have about the code, the answers they find to these questions [138, 226], their code locations of interest (commonly referred to as the “working set” [27, 45, 126, 274]), the different versions of the code they try [119, 269], and how those various versions produce differing outputs [119]. All of these information tasks can become even more difficult to manage when a developer is interrupted [156, 189, 191]. Strategies for keeping track of these various forms of information are rarely successful and are

often siloed into different tools or approaches, given the different types of information to track.

In this chapter, we explore the concept of generalizing annotating for capturing different forms of meta-information using a singular, unified mechanism in the IDE with our tool Catseye. Catseye adopts and extends many of the features introduced in Adamite, such as multiple anchors and pinning, while adding in support for programming-specific information tracking, including lightweight versioning and output capturing. Catseye has advantages over other forms of information tracking commonly employed by developers, such as code comments for tracking open tasks [237], by both allowing for a higher level of specificity through anchoring the meta-information to the relevant code yet abstracting that information away from the source code, such that it does not clutter the source with potentially erroneous information. Our lab study suggests that Catseye did assist in information tracking, with participants using Catseye performing better on a debugging task when compared to participants using their own information tracking strategies.

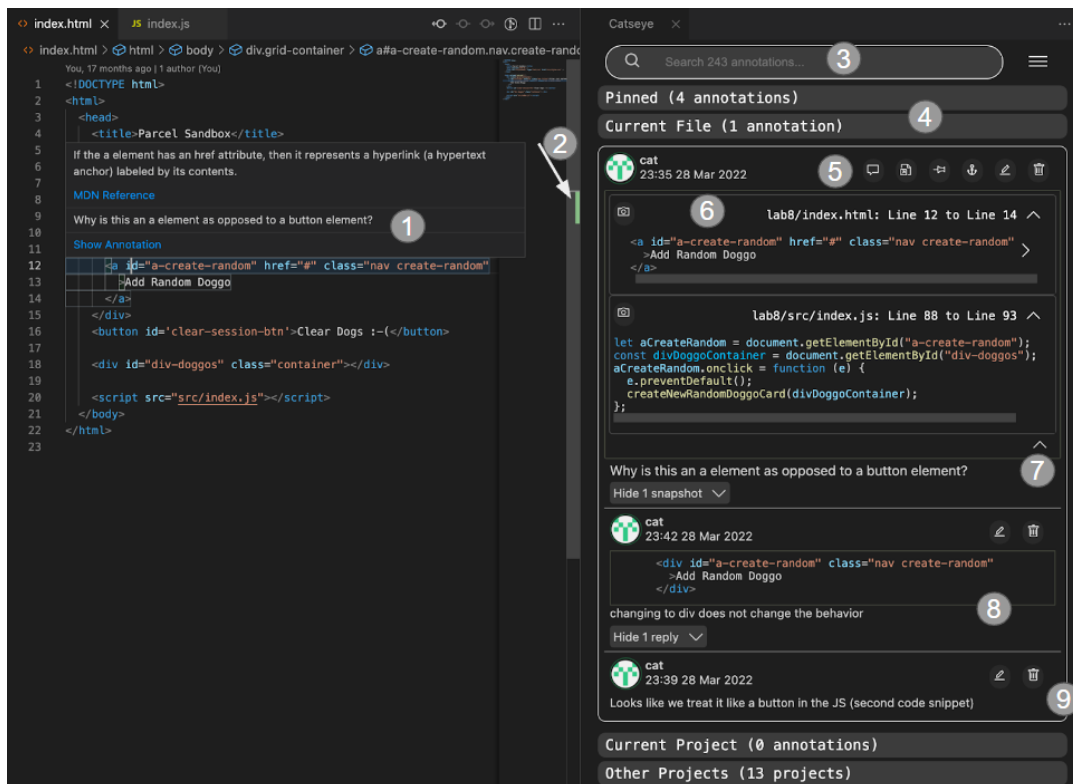


FIGURE 4.1: Catseye as it appears in Visual Studio Code. (1) shows how the annotation appears in the editor – the code is highlighted with a light gray box and, when hovered over, the annotation content appears in the pop-up with any other documentation. Clicking on the “Show Annotation” button opens and brings into focus the Catseye pane if it is not already visible, then scrolls to the annotation. (2) shows the annotation location(s) in the scroll bar gutter in a light green. (3) is a search bar for searching across the user’s annotations. (4) is the Catseye pane – the pane is segmented into sections corresponding to the annotations’ locations in the file system, with the “Current File” section currently open. (5) is an annotation – the top of the annotation shows the author and creation time information on the left, and buttons for various actions. (6) shows the two code anchors for the annotation. (7) is the content the user added as an annotation to the code snippets. (8) is a snapshot of the code at a previous version with a comment added by the author about this version of the code. (9) is a reply to the original annotation.

4.2 Catseye

4.2.1 Overview of Catseye

We developed Catseye— an extension for the Visual Studio Code editor [166] — that allows developers to keep track of their tasks, open questions and hypotheses, answers to these questions, and more in the form of annotations attached to one or more snippets of code (see Figure 4.1)¹. We chose to adopt some of the features of the Adamite system [104] for Catseye, given the similar goals of the systems with Adamite based around supporting developers’ information tracking on the web. Adamite showed the benefits of multiple anchors and pinning, and we expected that these features would help with code comprehension issues, such as managing a working set. We also introduced code-specific information tracking capabilities, including lightweight versioning and output capture.

To create an annotation, a developer selects a snippet of code in the editor and, using a keyboard shortcut, the context menu, or Visual Studio Code’s Command Palette, indicates that they want to create an annotation. The Catseye pane will update with a preview of the annotation, where the developer can add text and choose whether or not to pin the annotation. Once the annotation has been created, it will appear in the Catseye pane and the editor will update with a light gray box around the annotated code at the anchor point (see Figure 4.1-1). With an annotation, a developer can click on it to jump to the anchor point in the code (and vice versa), build upon it through adding additional code snippets as “anchors”, capture versions of the code and the code output, “pin” the annotation for easier navigation, “reply” to the annotation with more information, search for the annotation, export the annotation as a code comment, and edit and/or delete the annotation.

Given our high-level goal of creating an annotation system where annotations serve as ephemeral notes when making sense of code, we explicitly designed annotations to function similarly to Google Doc comments or Microsoft Word comments. Annotation anchors in Catseye move around with the code as the code and its location in the editor change over time, and the annotations appear in their own designated area that is detached from the developer’s editor. We chose this design metaphor to emphasize the point that annotations are not code comments — they are separate, meta-level notes that are attached to but abstracted from the developer’s working context. Annotation anchors update whenever the developer edits their code, and the copy of the code at the anchor that is shown in the annotation (Figure 4.1-6) is updated whenever the developer saves their code.

Annotation code anchors can also be used as navigational aids. The developer can click on the code or the file path in the annotation (Figure 4.1-6 and Figure 4.2-1) which will open that file in a new tab if it is not already open, bring that file’s tab to

¹Note that this figure and remaining figures in this chapter show the version of Catseye that existed at the time of publication and can be considered “version 1” — “version 2”, which is the current version, differs slightly in terms of its user interface and can be seen in Chapter 5.

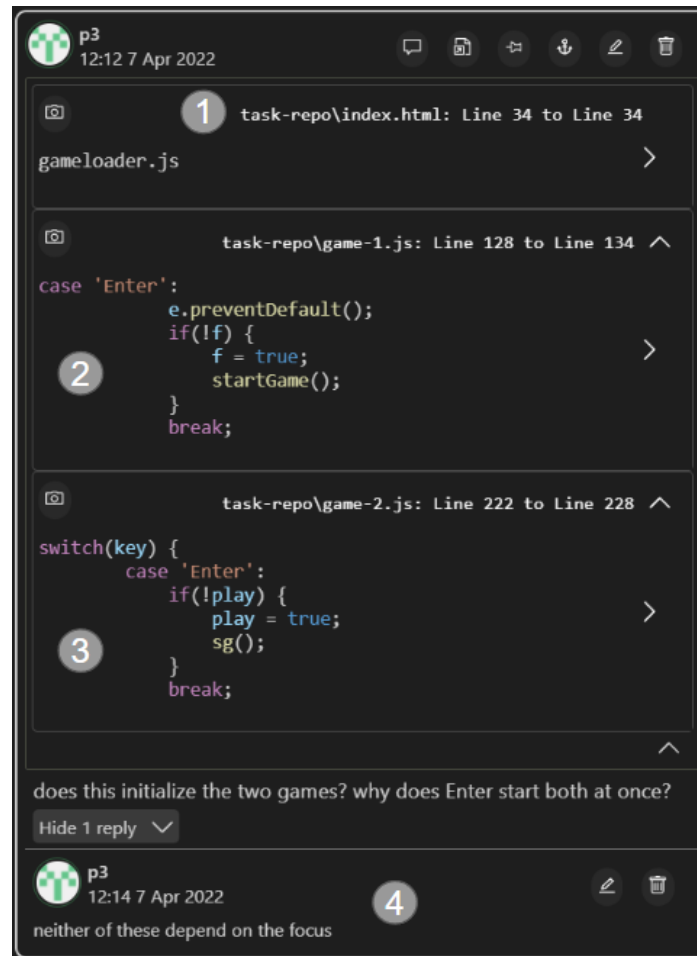


FIGURE 4.2: An annotation made by a participant in our study. (1), (2), and (3) show the 3 different code anchors the participant created across multiple files, with the first anchor (“gameloader.js”) as the site of their question, and the remaining two anchors and reply (4) answering their question. The annotation was pinned.

the front if it was not already, and scroll to the code’s location in the file. Additional navigational affordances are provided for pinned annotations – a developer can use a keyboard shortcut to cycle through each pinned annotation’s location to help with navigating through important code patches. Figure 4.2 shows a pinned annotation one participant in the user study created to help with managing their working set.

Given the mutable nature of code, keeping annotations attached introduced some design challenges unique to Catseye in comparison to other annotation systems designed for more static information. Since code is expected to change, retaining the original anchor point becomes more important as the annotation’s content is more likely to become out-of-date – we choose to store a copy of the user’s original anchor point for reference as the code changes. The developer can also deliberately save a version of the code by clicking the snapshot button on the code anchor box (see Figure 4.1-6). Once the snapshot has been created, the developer can edit the snapshot to add their own meta-information, such as what output that version of the code produced (Figure 4.1-8). Notably, these design choices were reasonable for a first-pass prototype system, but the tension between annotation and anchor contents is

focused on more in Chapter 5.

Another design challenge is how to handle the case where the user deletes the code that an annotation is attached to – should the annotation be removed as well or should it persist? Notably, this was not a challenge we had to consider with Adamite, given that an annotation author in most cases cannot directly edit the web page and, thereby, delete their annotation’s anchor. Different annotation systems do different things in the case of anchor deletion – Overleaf comments will persist when the anchor is deleted, with a line showing where the text used to be, while Google Doc comments will be removed if their anchor is removed. We chose to follow Google Doc’s design choice, with the rationale that, if the user wants the annotation content to persist, they can attach an additional anchor to the annotation or export the annotation as a code comment. Again, this design challenge is more directly discussed and approached in Chapter 5.

A related issue is what to do with annotations on code which is copy-and-pasted. Like with deleting, different annotation systems do different things. We decided to *not* copy the annotation with the code with the justification that we want to reduce as much potential annotation clutter as possible (thus we choose to never create an annotation without the user’s explicit request). These design choices also turned out to be the most useful way for these features to work in my personal usage of the tool (see Section 4.6). Notably, handling copy-pasted code is addressed more directly with our later tool, Meta-Manager, where the relationship between the copied code and pasted code is retained and tracked but in service of supporting slightly different design goals.

Related to questions of what to do with copy-paste, we also had to consider what to do when the developer uses Visual Studio Code’s undo or redo functions. Currently, annotation creation is not put on the undo stack, thus a user cannot undo creating an annotation. Annotations may also be anchored to code that is subsequently undone, thus removed from the code base – this works the same as deleting the code, in that the annotation is removed. However, Catseye currently does not handle redo, thus it will *not* detect if some removed code has returned due to a redo and re-attach the annotation.

Oftentimes, when a developer wants to keep track of some information, they will want to resolve or build upon the information they initially felt was worth jotting down. Catseye follows a similar model to other annotation systems where “following up” on the content of a note is kept very general, so the developer can either edit their original note, or reply to it with their additional thoughts.

4.2.2 Background and Design Goals

In creating Catseye for Visual Studio Code, we were particularly interested in helping developers capture and keep-track of their ephemeral thoughts, questions, concerns, and open action items related to their code, since this is the least well addressed aspect of previous tools. We envision that annotating will help with the

following use cases:

- **Keeping track of developers' questions and hypotheses about code.** Sillito et al. [226] found that, during software maintenance tasks, developers reported over 40 different kinds of questions they had about the code and that there is little tooling support for finding answers to those questions. They also found that there is limited tooling support for helping developers keep track of these questions as they come to an answer. Given that Catseye's annotations retain the original context of the code, and allow for composition of multiple code snippets through multiple-anchoring, we hypothesize that annotations may help with keeping track of these complex questions and the eventual answer a developer finds.
- **Keeping track of facts developers learn about code.** During any task that requires comprehending code, developers will naturally collect a body of knowledge about the code base [138]. Only some of these facts are appropriate as documentation, either because the behavior of the code is expected to change (such as during debugging) or because there is uncertainty about the veracity of the fact. We hypothesize that Catseye and its annotations will help with externalizing these thoughts while not requiring laborious clean up, since the source code is unaffected.
- **Keeping track of developers' open to-do items.** In trying to complete a complex programming task, a developer needs to keep track of a multitude of both high level goals and lower-level implementation steps in order to achieve that goal which the developer may forget, especially when interrupted [189]. Developers can use annotations to mark the code to change with the details of their "todo" item, compose snippets that are related to the change using multiple anchors, and can pin and un-pin the annotation as a way of marking whether or not the task still needs to be addressed.
- **Helping developers navigate their code.** An oft reported difficulty in programming is navigating the code base, especially when it is large. Developers typically discover a "working set" of task-relevant code fragments [27, 45, 126], then spend time navigating among these fragments as they implement their change. This navigation takes up a large amount of time, especially since these fragments can be difficult to return to [126]. We expect that clustering annotation anchors using multiple anchors, pinning these annotations for easier tracking, and using the code anchors as quick links will make this navigation easier.
- **Keeping track of localized changes.** When a developer is implementing a change, they often try multiple versions of the code in order to investigate the differences in output and ensure that the change works. These changes can be relatively small (i.e., less than 5 lines of code), may not be tracked in version

control [119], and switching between these versions can be difficult if the prior versions are not retained, especially since they may be inaccessible through undo commands [269]. Catseye allows developers to snapshot their code for versioning, such that developers can keep track of the different changes they try and can optionally associate these versions with the output they produced.

- **Keeping track of changing system output.** While testing changes, developers have reported a need for keeping track of what version of their code produced what output [119] and have used strategies such as copy-pasting the output into text files. We provide annotations as a place to store these outputs – developers can either reply to their annotation with the output values or edit their snapshots with the output which comes from that version of the code, thereby leveraging the context of the code.

Notably, the use cases described above are designed for the benefit of the *original author of the annotation*. We chose to focus on supporting the initial annotation author given the goal of supporting a developer’s tracking of information, which is largely localized to a single author and their implementation session(s). Whether or not code annotations can be helpful across time or between developers is discussed and explored more in Chapter 5.

4.2.3 Implementation Notes

Catseye is implemented as a Visual Studio Code extension and is written in TypeScript [165] with the implementation using React for the user interface [69] and Google Firestore [61] for storing annotations on the database, along with authenticating the user.

One particularly important aspect of Catseye is managing the code anchors as the user actively edits the code. Annotation anchors are kept up-to-date using Visual Studio Code’s document change event handler. Whenever the user modifies a file that contains an annotation, the Visual Studio Code API generates a change event object that we interpret. For simple cases, such as adding a new line at the top of the file, updating the anchors is trivial, but, in the case of more complex changes, such as pulling in a new version of code from GitHub or formatting a code file using a package such as Prettier.js [203], updating the annotation anchor becomes complex. The Visual Studio Code API treats these batch changes as many small edits applied in rapid succession. These rapid, successive edits leads to compounding and cascading errors in terms of the computations used to evaluate an annotation’s anchor point. If this occurs, we detect the anchor as invalid and delete the annotation². Given that tracking source locations has been a long-standing challenge in software engineering [208], we explore mechanisms for better handling re-attaching annotations in Chapter 5. Additionally, Sodalite expands upon the code anchoring

²Note that this was never intended to be a permanent design decision, as we knew we were going to look into more intelligently re-anchoring annotations, as discussed in Chapter 5.

functionality thru introducing new types of code anchors that are designed to be more flexible, such that un-anchoring is less common (see Chapter 6).

4.3 Lab Study

In order to understand how developers keep track of information while making sense of code when using their own strategies and when using annotations, we ran a small lab study. Participants in the *experimental* condition authored annotations while using Catseye to help them keep track of information, while participants in the *control* condition used whatever strategies they normally would employ. The lab study consisted of a training task, then a debugging task, and ended with a survey to assess the participants backgrounds, their experience with Catseye if in the experimental condition, and their experience completing the task. We chose to use a between-subjects design as opposed to a within-subjects design due to the nature of the task. The study took around 90 minutes, so adding another training session and 45 minute task would make the study too long. Further, as discussed in [128], since these are problem-solving tasks that you can only do once, creating 2 tasks which are independent but of equal difficulty is challenging. The full set of replication and study materials can be found in Appendix D.

4.3.1 Method

Training Task

Both conditions included a training task using a repository of website templates³ to either familiarize the participants with Catseye (experimental condition) or to show-case how the participants currently keep track of information when programming (control condition). Participants in the Catseye condition learned how to create and edit an annotation, pin and reply to an annotation, navigate and version their code using annotations, and collect system output, with all functionalities contextualized to how they may be useful for keeping track of different kinds of information. Participants in the control condition were asked to describe how they currently keep track of the different types of information we expect Catseye to support. In this way, we tried to make sure that both groups were primed about the kinds of activities that Catseye is designed to support.

Main Task

For the main task, participants were instructed to understand and attempt to debug a website. Participants were told to imagine that they were a new developer on a team and that they were tasked with understanding and debugging some code. For the first 15 minutes, the participants were not allowed to edit the pre-existing code

³<https://github.com/ShauryaBhandari/Website-Templates>

Game	Bug	Minimal Solution
Both	Unable to Play Games Independently	Change one of the event listeners to a different key (1 value change)
Snake	Screen Does Not Refresh	Adapt Tetris's screen clearing function to Snake (10 line change)
Snake	Snake is Too Fast	Adapt Tetris's timing function to Snake (10 line change)
Snake	Snake is Drawn Incorrectly	Change the constant value for the snake segment length (1 value change)
Snake	Food Collision Check is Incorrect	Change the ORs in the boolean to ANDs (2 value change)
Tetris	Blocks Falls in Last Key Press Direction	Set current direction of block fall to "down" on each game loop (3 value change)
Tetris	Rotating Square Causes Square to Move Upwards	Add conditional to prevent square from being rotated (3 line change)
Tetris	Game Does Not End	Change conditional to whether the stack of blocks is at the top of the screen (1 value change)
Tetris	Game Calculates Score Incorrectly	Increment user's number of cleared rows instead of setting to last clear row value (1 value change)

TABLE 4.1: The bugs present in the two games. "Value" refers to a construct in the program, such as an operator, boolean, or variable.

(but they could add comments and print statements) as part of the scenario in which they are new to a team and should spend time familiarizing themselves with the code prior to contributing changes. This also allowed us to investigate differences in the kinds of annotations made while understanding versus debugging and editing. After the 15 minutes of understanding and testing the code, participants had 30 minutes to attempt to use what they learned to resolve issues they had discovered.

The website included buggy implementations of Snake and Tetris. Each game had 4 bugs, with an additional bug that affected both games, totalling 9 bugs (see Table 4.1). We chose these two games since they are both relatively well-known, are event-based which makes understanding their structure less straightforward, and have clear requirements such that testing the games takes less time in comparison to actually debugging their logic. Similar tasks have been used in related studies [189, 191].

The code was specifically designed to be confusing in order to make keeping track of information particularly important (see Table 4.2). Given prior literature around what makes code confusing [226], we purposefully included bad code smells such as poorly-named variables, global variables, lack of organization amongst methods, and no documentation. Since participants only had 45 minutes for the task, we wanted to necessitate keeping track of information while also keeping the task semi-realistic through using known issues when comprehending unfamiliar code. To further validate the realism of the code, we included two questions in our post-task survey that asked participants how similar the code they saw in the study is to code they have encountered during their time as developers and how frequently they have encountered such code. Participants reported the code is similar to code they have encountered before⁴ but that they (fortunately) do not encounter code like this very frequently⁵.

⁴average = 3.4 out of 5, using a 1-to-5-point Likert scale from very dissimilar to very similar

⁵average = 2.6 out of 5, using a 1-to-5-point Likert scale from never to very frequently

What Information to Track	What Aspect of the Task	Explanation
Questions, Hypotheses, and Answers	Debugging task	Debugging naturally leads to many questions and hypotheses about the program behavior but subsequent answers may be lost or forgotten [226]
Facts	Poorly-written and documented code	Developers are tasked with learning what the code constructs are and how they are used
Open Tasks	15-30 time split	Allow developers to discover many bugs, then have them decide which ones to focus on and how to fix them
Navigation	Poorly-organized code	Constructs, including methods and classes, are spread across multiple files including some files the participant cannot edit
Localized Changes and Output	Arcade Games	By having two arcade games, participants are tasked with making changes and seeing how that affects each game and how that affects each game's output

TABLE 4.2: How the study task encapsulates the types of information Catseye supports.

4.3.2 Participants

We recruited 13 participants (5 women and 8 men) using study recruitment channels at our institution, and advertisements on social media. Participants were randomly assigned between the control and experimental conditions, with 7 participants in the experimental condition and 6 in the control condition – participants in the experimental condition are referred to as “P1” through “P7” and control participants “C1” through “C6”.

All of the participants were required to have some amount of experience using JavaScript, to have a GitHub account, and to regularly use Visual Studio Code. The participants’ professions included graduate students in computer science-related fields, undergraduate students in computer science, and professional programmers. On average, participants had 10.2 years of programming experience, 5.2 years of professional programming experience, and rated their familiarity with JavaScript at 4.5 out of 7. Participants in the control condition had more experience and more professional experience, on average, than experimental participants, but not significantly more.

4.3.3 Analysis

Across both conditions, we objectively coded what bugs the participant succeeded in fixing (see Table 4.1). In the experimental condition, we analyzed the video recordings and log data to count how many annotations each participant authored and how often they interacted with their annotations to assess the utility of the annotations for keeping track of information. We additionally logged whether or not any annotations were made in the first 15 minutes and, for annotations created during the debugging part of the task, what bug the participant was attempting to solve at the time of creation. We analyzed the videos in the control condition to log the same types of interactions including the artifacts developers made in that condition, such as code comments and external notes.

We additionally labeled the annotations and control condition notes with the type of information it was being used to help keep track of. We objectively coded this conservatively based off the content. If an annotation’s or artifact’s content was

phrased as a question or had a question mark, it was coded as a question; if the content had words such as “might” or “seems like”, it was coded as a hypothesis; if the content was phrased as an objective such as “change this”, it was coded as a task; and if the content was stated as a fact (e.g., “game-1.js is snake”) it was coded as a fact (even if the fact was incorrect). The same process was used for annotation replies.

We counted items as used for “versioning” when they either contained a snapshot (annotation) or were used to mark a change they made to the code base (annotation or control artifact). For navigation, we counted an annotation that is pinned and/or had multiple anchors as used for navigation.⁶ We counted an annotation or control artifact as being used for output if the participant used it to store or comment upon the game output. If the content of an annotation or artifact did not fit into any of these categories, it was marked as “Other”. For replies, we also labeled whether or not a reply served as an answer to their question annotation – a reply was considered an “answer” if its content was a direct response to the question’s content that supported or refuted it.

4.4 Results

Participants in the experimental condition fixed, on average, 1.85 bugs (min = 0, max = 4), while participants in the control condition fixed 0.67 bugs (min = 0, max = 1), a significant difference (*two-tailed T-test*, $p = .04$). To further explore these results, we investigate what types of information participants chose to keep track of through annotations versus what information control participants used their artifacts to keep track of, how participants used their information when completing the debugging tasks, and how participants performed on the debugging task.

4.4.1 What Information Do Developers Keep Track of with Annotations and Artifacts?

Experimental condition participants created 84 annotations, with each of these participants creating, on average, 12 annotations (min = 6, max = 21, median = 10, std. dev. = 5.446). 44 of the annotations were made in the first 15 minutes (a rate of 2.93 annotations per minute) and 40 were made in the last 30 minutes (a decreased rate, at about 1.33 annotations per minute). This slight drop-off in annotating is not particularly surprising, given that the types of activities developers were performing changed as they moved from understanding code to more actively writing and fixing code. The size of the annotations averaged 12.2 words (min = 1, max = 45, median = 12.5) and they were attached to code averaging 29.3 characters. Each anchor was, on average, 1.59 lines long, with the majority of annotations attached to one line or less of code (71/84).

⁶Since multiple anchors and pinning are unrelated to the text content of an annotation, this means an annotation could be marked as both “navigation” and, for example, “fact”.




FIGURE 4.3: The average number of annotations and artifacts participants created during the study.

Across those 84 annotations, developers had a variety of types of information they chose to keep track of through annotations (see Figure 4.3). The most common usage for an annotation was to keep track of open questions developers had with 28 out of the 84 annotations being questions (33.3%). 16 of these questions were made during the first 15 minutes, while the remaining 12 were created in the last 30 minutes. 6 out of those 28 questions were definitively answered, while 2 had follow-up hypotheses given program output behavior, and 1 had a follow-up question associated with the original question, resulting in 9 out of the 28 questions being followed-up on in some way. Given the complexity of the task, the amount of answered questions is not particularly surprising, but the fact that participants followed-up on their questions at all provides support for the argument that annotations can serve as dedicated spaces for these questions. In contrast, as discussed below, only 1 of the control condition’s 17 questions were followed-up on or answered.

The second most common type of information experimental condition participants kept track of in their annotations were facts they discovered about the code, with facts comprising 27 out of the 84 annotations (32.3%). 23 of these 27 facts clarified information that was explicitly designed to be confusing. For example, P1 annotated `const c = document.getElementById('t')` with “this is the canvas of the tetris game”. 18 of these 27 annotations were made in the first 15 minutes, when participants were reading through and understanding the code.

Experimental condition participants also utilized annotations to keep track of their open tasks (21.4% of their annotations) and to navigate the code (17.9% of their annotations). These annotations typically served as reminders to the participant about places in the code base they suspected were related to the bugs they identified

in the code. The majority of task annotations were made during the 30 minute debugging phase (12/18) suggesting that there was more of a need for keeping track of their areas of interest in the code once they were developing, as opposed to when they were trying to understand the logic.

Experimental condition participants did not use their annotations for keeping track of their code versions, with no participants using the snapshot feature. Participants did create annotations to comment on parts of the code they added or modified, with 6 annotations made on the participant's own code that they added. Considering that participants, in general, did not edit the code very much, since the bugs did not require large modifications to fix, there may have been less of a need to keep track of small localized changes. Further, the code that participants chose to annotate was usually code that the participants did not edit, with only 12 of the annotations' corresponding code being edited. 3 annotations were used to keep track of output – the small number of output annotations may also be due to the minimal amounts of changes participants made to the code. Further, since the program was a computer game, much of the output changes were graphical which is not output that Catseye can capture currently.

In the control condition, the participants created a total of 100 different artifacts, averaging 16.67 artifacts per participant (min = 2, max = 27, median = 15, std. dev. = 9.771) – which is slightly more artifacts per participant than in the experimental condition, but the difference is not statistically significant ($p = .76$, T-test). This may partially be due to the fact that control participants were primed to think about and show how they keep track of information. Further, all of the control participants had some note-taking strategy that they described using in their daily work.

Their artifacts included 78 code comments, 14 external notes (with 5 of the notes being created on a tablet computer, 1 being created in a Notepad document, and the remaining 8 created using pen and paper), and 8 Git commit message⁷. The artifacts averaged 5.95 words (min = 1, max = 22) – notably shorter than the annotations, which averaged 12.2 words per annotation.

The information that control participants chose to keep track of through artifacts differs from the information that was annotated (see Figure 4.3). While, in both conditions, facts, questions, and open tasks were the three most commonly kept track of information types, participants in the control condition kept track of facts the most, while participants in the experimental condition kept track of questions most often. Participants in the control condition also favored “task” artifacts more so than experimental condition participants, with most in the form of “TODO” code comments (66.6%) consistent with prior work [237]. Only one control participant actively attempted to keep track of different versions of their code, and none of the control participants kept track of the output of their code or used any specific mechanisms to help navigate their code, aside from traditional code search (which participants in the experimental condition also used). These results suggest that annotations can

⁷All of the Git commit messages were created by one participant.

promote more entering of questions and answers in comparison to traditional notes and annotations can keep track of other types of information that may otherwise not be captured.

4.4.2 How Do Developers Use Their Annotations and Artifacts?

We quantify usage of annotations or artifacts by counting whenever a user interacted with their annotation or artifact in some way. On average, experimental participants revisited 5.71 unique annotations, and revisited their annotations 11.14 times over the course of the study. In contrast, the control condition, on average, revisited 3.5 artifacts a total of 4.67 times suggesting that the annotations were more successful in encouraging participants to follow-up on their information. One of the two top participants, both of whom fixed 4 bugs, also created the most annotations (21) and revisited them the most, revisiting 14 annotations 32 times. This suggests that annotation usage may have contributed to his success.

Experimental participants often revisited their annotations to add replies to their annotations, with participants creating 16 replies and 5 out of 7 participants creating at least 1 reply. Replies typically served as an extension of the annotation's original content, with some annotations serving as answers to the original question (6), hypotheses about the behavior of the code (3), and follow-up tasks that they wanted to complete related to the original annotation (3). For example, P5 made an annotation about Snake where the initial annotation just said "Game 1: Snake" and created two replies, with the first reply explaining what the two Snake-related files did, and the second reply listing all of the bugs she had encountered with Snake – she then revisited this annotation 3 times over the course of the study to keep track of her bugs. Replies also sometimes functioned as places to discuss the behavior of the code *after* the participant attempted to fix a bug associated with the annotated code, with 2 replies commenting on whether their implementation worked or not.

Participants also used pinning, multiple anchors, and anchor clicking as a way of supporting their navigation while working on the task. 5 annotations had multiple anchors, 3 annotations were pinned, and the participants used the anchors to navigate the code base 19 times. In contrast, the control condition did not use any of their artifacts to help them navigate the code base.

Three experimental participants also chose to delete their annotations once they were "done" with them, with these participants deleting a total of 14 annotations. The most successful participant deleted 12 of his 21 annotations over the course of the study – whenever he fixed a bug he would find each annotation that related to that bug and delete it, while keeping open the annotations that were still unresolved. His usage of Catseye suggests that annotations can function similarly to comments in systems like Google Docs where, even if the content of the comment is not necessarily a "to-do item", the comments can still be resolved in a similar manner. Control participants only deleted their artifacts, on average, 0.8 times while experimental participants averaged 2.16 deletions, further suggesting that a Google

Bug	# of Experimental Participants Who Fixed This Bug	# of Control Participants Who Fixed This Bug	% of Debugging Annotations Made About Bug	% of Debugging Control Artifacts Made About Bug
Unable to Play Games Independently	2	2	12.5%	0%
Snake Screen Does Not Refresh	3	0	12.5%	0%
Snake is Too Fast	2	2	15%	88.6%
Snake is Drawn Incorrectly	1	0	2.5%	0%
Snake Food Collision Check is Incorrect	1	0	7.5%	5.7%
Tetris Blocks Falls in Last Key Press Direction	1	1	35%	5.7%
Tetris Rotating Square Causes Square to Move Upwards	0	0	0%	0%
Tetris Game Does Not End	1	0	12.5%	0%
Tetris Game Calculates Score Incorrectly	0	0	2.5%	0%

TABLE 4.3: The annotations and artifacts participants created during the study while working on each bug. The experimental condition made 40 annotations while working on bugs, while the control condition created 35 artifacts. The last 2 columns refer to the proportion of annotations made about that bug out of the 40 annotations made while debugging, and the proportion of control condition artifacts made about that bug out of the 35 artifacts made while debugging, respectively.

Docs-style design encourages more clean-up than regular code comments or external notes.

4.4.3 How Did Participants Identify and Fix Their Bugs?

All bugs were identified by at least 1 experimental participant and had at least 1 annotation created about it, save for the Tetris square rotation bug. No participants in the control condition identified the Tetris rotating square bug or the Tetris score calculation bug, so no control participants made artifacts about those bugs.

When struggling with difficult bugs, participants seemed to create more annotations and artifacts. For example, the “Snake is Too Fast” bug, which required 10 lines of code to change, was only successfully completed by 3 out of the 7 participants who attempted it, and resulted in the majority of control condition artifacts to be about this bug, along with some annotations (see Table 4.3). Conversely, some of the simpler bugs to fix, such as “Snake is Drawn Incorrectly” had fewer annotations made about them as there was less need for participants to externalize their thought processes.

Some particularly complex bugs led participants in the control condition to utilize their notes in different ways than their experimental counterparts did. 3 control participants made a total of 3 external notes that were either visual diagrams of how they thought the games should function or were algorithmic step-by-step instructions for how to design their bug fix. Since the experimental condition created no similar notes, this suggests that future versions of Catseye may better support users by including a way to attach and create visual diagrams, screenshots, or drawings to annotations and support richer interactions for checking off completed steps in an algorithm.

Participants in the experimental condition occasionally made annotations that documented how they fixed a bug, with 4 annotations created for this purpose. For

example, P6 wrote some code to try and fix the Snake Screen Does Not Refresh bug and annotated their code with the text “Attempt at clearing the score” and edited their implementation 3 times to try and achieve the correct behavior. P6 then made 2 more annotations on other code snippets that they were referencing when trying to fix their implementation, hypothesizing about how they could adapt the functionality of that code to solve their bug. Their usage suggests annotations can help with marking and documenting code while debugging, including code the user has added that attempts to fix the bug.

4.4.4 Continued Usage

Four participants asked to continue using Catseye after the study, with 2 participants creating more annotations in their own code after their session. One of these 2 participants reported back on her usage of the tool in her daily work. She found the tool useful for externalizing her “design-oriented notes-to-self” such as “maybe I should do X instead, if I do decide to do that, this is the code that needs to be edited to make it happen”. Notably, this is the type of information she says she would normally write down on a piece of paper and *not* use code comments for since she does not want to create clutter that will only confuse her or her collaborators later. She found Catseye valuable for acting as a space for capturing this “thought history” that leverages the context of the code. Her experience lends further support to our claim that creating a space for keeping track of developers’ thoughts in a programming tool is useful.

Further support for our claim that annotating can be useful is the adoption of Catseye by a small amount of developers following the conference publication of the work. 7 developers created 35 annotations across 9 unique projects. 8 of these annotations were deleted, 5 had multiple anchors, and 10 of these annotations had a total of 19 replies by the author (1 of which was deleted). The annotations were, on average, attached to code anchors that spanned 3.5 lines (notably longer than the anchors found in the study, which averaged around 1 line of code), with the longest anchor spanning 19 lines of code. The annotation’s content averaged 68 characters (minimum: 0, maximum: 315).

One of these later users authored the majority of the annotations, creating 23 annotations – their usage across this larger amount of annotations is consistent with the sorts of annotation types and behaviors we found in the study and in my usage of the tool (see Section 4.6). In reviewing their annotations (all of which were made across a week-long span), it became clear that they had multiple threads of inquiry, mostly relating to trying to track down a memory-consumption bug. In anecdotally applying the codes used in the study to their inquiry (ignoring the 2 they deleted), they created 7 “fact” type annotations, 7 “question” annotations (4 which were answered), 2 “hypotheses” (both of which were followed-up on with replies), 1 “task” (which was also followed-up with a reply stating the result of the task), and 3 other

annotations. The 3 “other” annotations did not fit into the previously-defined categories: 1 just said “one more candidate” about some code (perhaps in reference to their task of identifying where a bug was originating from), 1 was a comment stating “seems kind of sketch, but this is research!” on some code (which is more in line with the emergent “reaction”-type annotations discussed in Section 4.6), and the last just says “True” on some variable, which is too terse to be comprehensible for me, an outsider, to code. 3 of the fact annotations were made with the intention of recording design rationale. Both hypotheses were followed up on, with one being replied to 3 days after the initial annotation was made. In fact, this user replied to nearly half of their annotation (9 out of the 21 not-deleted annotations) with 18 replies. Their replies often followed-up on hypotheses, tasks, and questions (7 replies), proposing new tasks (1), hypotheses (4) or questions (2) with some replies serving as answers to those follow-up lines of inquiry (4). This participant also used multiple anchors and referenced them in their replies – 3 of their annotations had 2 anchors, 2 of which were explicitly referenced as evidence when coming to an answer. Sometimes these lines of inquiry would span across multiple days, with the longest span of time being 3 days between an annotation’s initial creation, and the eventual reply that closed the line of inquiry. While we do not have full insight into their usage of the tool or motivation for using it (nor why they chose to stop using it), their self-motivated usage does provide additional evidence for some of the behaviors we believe annotating to be helpful for (i.e., *replies* for following-up on content and *multiple anchors* for clustering and referencing related code).

4.5 Discussion

Our experiment lends support to the concept that annotations may be used as a lightweight way of capturing and following-up on information that may not otherwise be kept track of when programming. Participants succeeded in creating questions, following up on those insights, and revisiting these notes in order to fix their bugs and had more success, on average, than the control condition.

Participants found the annotating metaphor familiar and understandable, despite the amount of complex activities participants could use the annotations to support, with participants in the post-task survey saying the system was very easy to learn how to use.⁸ Prior work has noted that annotations’ flexible nature and structure allows them to be used in a variety of ways [5, 21] – we build upon this by showing that annotations can be used in new ways, including to store output, to store and capture versions of code, and as navigational aids. Typically, attempts to support these different activities are siloed into different research tools; annotation systems show a promising alternative where, by utilizing annotations’ flexible

⁸“I consider it easy for me to learn how to use Catseye”, average score 6.83 out of 7, with 7 being “Strongly Agree”

nature, they can act as a more general “workspace” for storing and thinking about contextualized information a developer cares about.

Another way that participants used annotations as a general “workspace” for thinking was through using their annotation as a “placeholder” when navigating their code. For example, P3 created a question annotation asking why a certain method gets called twice. They followed-up with a hypothesis stating “seems like it’s because it calls the draw function, which has some special logic that only occurs if play is true”, but, while writing this reply, they paused to continue exploring the file while reflecting on their hypothesis, before returning to the annotation and finishing their initial thought with a guess, “but you still want to call `window.requestAnimationFrame` i guess?”, given what they had learned while exploring. Three other participants paused while creating their annotations to explore the files and think critically about what they were choosing to annotate, which suggests that the choice to have the annotations in their own dedicated pane separate from the context of the code may better support this kind of self-reflection. Notably, these self-reflections have been shown to improve learning outcomes [41].

Some participants in the Catseye condition thought that the output and version capturing features would be particularly useful, despite not using them in the study. Two experimental condition participants reported in the post-task survey that they would use the snapshot feature in their own programming, since they found it difficult to go back to GitHub to see other versions of their code and they sometimes created many small changes that were not tracked in version control. A third experimental participant noted that they wanted to use this feature to capture output when they are performing maintenance tasks like refactoring and need to keep track of many “moving parts” and how their changes affect the behavior of their code. Since running the study, we enhanced Catseye to automatically capture intermittent output through connecting into the Visual Studio Code debugging API to capture and store run time data as replies when an annotated line of code is run.

Two participants in the control condition and 2 participants in the experimental condition created artifacts that, while phrased as a fact, were incorrect. The control participants added comments above functions incorrectly stating what the functions’ purposes were. The 2 experimental participants incorrectly assumed what a function and variable were used for, respectively. These annotations, while incorrect, are only visible to the original annotator, while the code comments could, in theory, be viewed by any collaborator, which could potentially misinform them. Even if the annotations were viewable to collaborators, they would not be in the code acting as documentation, lessening their potential to be harmful. An incorrect annotation could be a learning opportunity with the reply feature, where a collaborator could clarify or correct a misinterpretation of the code.

Two control participants had problems managing their code comments. C1 created a comment noting that a particular part of the Snake code looked like it was used for initialization, then discarded the commit that contained that comment. 10

minutes later, they searched for that comment, forgetting that they had discarded the comment. C2 marked a part of the code with the comment “REVISIT” but then undid a series of changes in order to revert to an older version of the code, removing that comment in the process, and then never revisited that part of the code. Participants in the Catseye condition did not create any code comments (2 participants started to make code comments before removing them and manually converting them into annotations) and did not lose any of their annotations during the study – annotations’ meta-nature may serve as a safeguard from erroneously removing them. However, Catseye users could, in theory, lose their annotations by erroneously deleting the code with which the annotation is associated. Currently, Catseye does *not* put annotation creation, editing, or deletion into the Visual Studio Code undo stack, another area where other annotation systems differ – Overleaf similarly does not, whereas Google Doc puts comment creation in the undo stack (but *not* comment editing or resolving!). In my own usage of Catseye, putting annotation creation on the stack would be beneficial, given that I normally annotated code I had just worked on – this means that performing an undo operation is likely to modify the code anchor the annotation corresponds to, which may change the annotation’s meaning or cause it to become detached if the code did not exist in the prior edit. Annotation deletion would also make sense to go on the undo stack, especially in the case that the annotation is deleted because its code anchor was deleted, given that an undo operation would bring back that code anchor, thus should bring back the annotation. This is how similar systems, such as Google Docs, works.

4.6 My Usage of Catseye

I used Catseye while developing Catseye as a form of “dogfooding.” Anecdotally, we report on the annotations created by myself using the same methodology of labelling each annotation by the primary type of information it was meant to keep track of. We omit annotations made without any text content, as they do not have enough context for labeling, and annotations made purely for testing the application, considering they do not represent “real” usage of the tool. All of these annotations were created by myself to help myself, either in the short term (so they were ephemeral) or for when I returned to the code later.

Over 12 months, I created 182 “real” annotations in the Catseye repository across 25 source files, with each annotation averaging 28.95 words (min. = 2, max. = 143, std. dev. = 27.99). 42 annotations had a total of 58 replies, 8 annotations had 18 snapshots, 3 annotations had multiple anchors,⁹ and 9 annotations were pinned.¹⁰ 110 of these 182 annotations were deleted as I finished open tasks and iterated over

⁹Multiple anchors were added late in development, so the lack of multiple anchor usage is primarily due to the short amount of time to use the feature.

¹⁰This is a conservative count, considering we used our log data and the log does not count whenever an annotation is pinned or un-pinned, just whether the annotation was pinned the last time it was updated in the database.

the code, including removing the code the annotation is anchored to which deletes the annotation.

The content in my annotations differs slightly from the annotations created in the lab study, perhaps due to the different nature of development. The most common annotation type was “task” type annotations, with 48 of the 182 annotations reporting some open action item I needed to act on – in contrast to the task annotations made in the study, these task annotations served as reminders for places to change when performing maintenance tasks like refactoring, as opposed to parts of the code that may have a bug. I also created many question annotations (47 out of 182) – 12 were replied to, with 4 of these replies answering the original question, and 17 questions were deleted. These questions typically pondered the system behavior, previous design choices, or details of the Visual Studio Code API. The third most common code was “Other”, with 16 annotations. All of these “other” annotations served as a *reaction* to the code, with these reactions sometimes pondering the former design rationale and sometimes expressing frustration with implementation challenges. For example, one annotation, anchored to a particularly confusing function said “This is a pain”. No annotations were made like this in the lab study, suggesting that actively implementing and writing code may elicit different types of information than code understanding and debugging tasks.

Actively using Catseye also led to changes in the tool’s design. Initially, an annotation would be copied if the developer copied the annotation’s code anchor point. However, given that code was often copied to be used as a template, in active development, copy-pasting the annotations resulted in duplicated annotations with irrelevant content as the code changed, which I found to be more distracting than helpful. This experience informed our design decision to not copy the annotation content when the user copy-pastes the anchor code. However, there may be situations in which it would be better to copy the annotation with the code, such as when the annotation serves to document behavior about the code. Future versions of Catseye might benefit from allowing the user to choose whether or not to have the annotation copied with the code, or alternatively, to add a new anchor to the original annotation.

I found the tool useful for externalizing and thinking through problems and stopped using code comments in favor of annotations. The ephemeral nature of the annotations was particularly useful in the development of Catseye since the code was nearly constantly changing which lead to a lot of uncertainty about the design and implementation details – information that I did not want to commit to the code base as comments in case the information ended up becoming obsolete (a common problem with code comments [74, 201, 240, 241]).

4.7 Conclusion and Future Work

In our development of Catseye, we found evidence that note taking and collecting of code-related meta-information was helpful for developers when tracking information. Further, we found evidence that these activities could be supported using a singular, lightweight interaction design – annotation. These findings extended our previous findings from Adamite by generalizing our approach to a different developer sensemaking domain while tracking other types of information. Below, I briefly reflect on some take-away lessons and challenges from this work that moved the research forward:

- **Working in a dynamic environment with static data introduces complexity.** Before, in Chapter 3, I claimed “context was key” and, as we showed in this work, it still is. Annotations and their anchor points were commonly revisited and helped both myself and our study participants think through and reason about problems that were subsequently solved. However, with a mutable source document, “context” is now multi-dimensional – the source location and the surrounding information is constantly changing across time while the annotation, if a developer chooses not to edit it, is stuck back in time. We largely, albeit unintentionally, avoided this challenge in this study, given that participants did not edit the source code very much and the study was a short, finite window of time in which the annotation could not “age” out of its context. However, in my own usage of the tool, I experienced firsthand how annotations, while useful in the moment, would eventually lose their context and become clutter, either due to my memory slowly eroding over time or the source code changing so much such that the annotation content itself no longer made sense. Small-scale design decisions, such as choosing to delete an annotation when its anchor is removed and not copy-pasting an annotation if its anchor is copied, helped a bit but are not enough to fully address the problem. How do we build systems that are designed with information mutability in mind? Chapters 5 and 6 probe on this challenge. This challenge speaks particularly to the earlier discussion of environment and source *mutability* and *provenance* introduced in Section 1.1 and how that impacts reliability of *anchoring* across *time*.
- **Lightweight unification of tooling that operates at a meta-level is a rich and relatively untapped market.** Annotating, as a practice, has largely been localized to and studied in the context of reading materials. However, choosing to apply annotating to a dynamic space opens up a myriad of opportunities for more domain-specific tooling. While this chapter only explores source code and some of the types of meta-information you may want when working with source code, such as code versions or outputs, one can imagine generalizing the design approach put forward by this work to other domains. For example,

Figma has a commenting system, but it is relatively limited – what if we “supercharged” these comments to additionally version themselves and the designs they are attached to to help with reasoning about design rationale over time and across collaborators? What if these comments could be used to extract and aggregate different parts of a design together to create a new prototype? I believe part of the reason commenting and annotating are so ubiquitous across information authoring and management platforms is because there is a lot of power in being able to directly point at something and say “this is what I am talking about”. Information closeness is a strong undercurrent of this work and Catseye and Adamite show the power of this in both communicating across time (Adamite) and in binding otherwise-disparate pieces of information together (Catseye). In this way, we unify *content* through the common reference point of the *anchor*. For those interested in building their own annotation systems, I strongly suggest understanding the types of information management tasks your clientele are actually trying to complete when annotating and more directly supporting those tasks through increasing both the types of information annotations can hold and giving more functionality in terms of what an annotation can do. Such annotating mechanisms can leverage the user’s working *environment* and, potentially, the *document structure* in order to lower information authoring costs and support other tasks, such as tracking data *provenance*.

Much of the success of the annotation approach can be attributed to the localized nature of the annotations and their connection to the source material through their anchor point. This connection allows for the note to serve as a navigational way point, allows for commenting on different types of code with the same interaction method (e.g., commenting on a variable within a line of code versus commenting on a whole conditional that spans multiple lines), and supports code-specific activities including versioning and output capture. Despite these powerful attributes, much of that power stems from the tight coupling between the annotation and its anchor point – thus, that power is lost when the annotation becomes unanchored. We further explore questions of anchoring, along with longer-term benefits and management of annotations in the next chapter on annotation curation. We then extend our code anchoring support and the types of information that may be attached to the code anchor points through extending Catseye into a long-form documentation tool, Sodalite.

Chapter 5

Curating Ephemeral Meta-Information with Catseye

5.1 Overview

In our creation and development of Catseye and Adamite, a constant design question was how to manage the *scale* of annotations, how to deal with *changing source content*, and *what to do* with annotations across time to increase their utility, given their ephemeral nature. These challenges are not dissimilar to challenges encountered with other “information scraps” (a class annotations fall into) [21, 22] – namely, the ad hoc nature of creating and using information scraps is often at odds with the design of traditional personal information management (PIM) tools, which typically prescribe an organizational and rigid structure for managing these scraps [21, 251]. “Information scraps” are characterized both by their form (e.g., lightweight, often captured on various platforms across different modalities) and by their intent (e.g., sticky note to *remind* you of an upcoming appointment, Google Doc *reference list* of URLs linking to recipes¹, etc.). These scraps are most often used in the moment as a form of cognitive offloading and externalization, but occasionally are meant to persist indefinitely (such as in the case of the reference list – a living document) or until a set time (such as with a reminder of an appointment).

Annotations and our annotation platforms, Catseye and Adamite, both require and allow for some amount of structure (i.e., requiring an anchor point for contextualization and, in Adamite, supporting organization through types and optional tags). However, the content of these annotations and what to do with the annotations falls into the information scrap challenge in that their utility, especially over the long term, is often not obvious (e.g., is a developer using Adamite to annotate this API method as a *reminder* to use it later? Or to mark it as *obsolete*?). Annotations have the additional challenge of being dependent upon a stable anchoring point, which must remain constant or, at least, semantically equivalent to the original anchor point’s content, which introduces additional design challenges in terms of determining what to do with these information scraps after some time has passed

¹This example is taken from my lived experience – my fiancé and I have one of these documents and it is chaotic, to say the least!

and the associated document has changed. Lastly, another design challenge that differentiates our annotation systems from prior work around information scraps and PIM tools is that our tools are meant to be used *collaboratively*. This leads to the challenge of scaling these systems to a point where potentially hundreds of annotations are shared to accelerate other users sensemaking of documentation and code. Given enterprise level systems, with potentially hundreds of engineers using these systems on a shared code repository with Catseye, how does one ensure all annotations are versioned properly, are useful, and have consistent and stable anchoring points?

Given this tension between retaining the pros of annotating as a lightweight way of expressing an idea when sensemaking while ensuring that our design is scalable and robust to changing documents, we began a design probe on how to support annotation curation. We began an early exploration of this concept with Adamite through the creation of an Adamite website² – through that exploration, we found a need for batch processing of annotations. We then focused our attention on Catseye, given the more ephemeral and dynamic nature of annotations in that context – we hypothesized curation would be even more important in that space. We modified Catseye to support large-scale curation³. Our curation work with Catseye is two-pronged – one prong explores how to deal with annotations *en masse* as a way of cutting down on potentially hundreds of irrelevant annotations, e.g., to throw away after their moment of utility is over. Another prong explores annotation’s connectivity to a source document and how to algorithmically re-attach the annotation across numerous scenarios, as an additional form of curating one’s set of annotations (see Figure 5.1). We then implemented these designs into Catseye– all screenshots in this chapter represent the current version of the Catseye system (in contrast to the screenshots in the previous chapter). Through this exploration, we created a set of prototypes and design insights that informed our later work, especially with respect to Sodalite.

5.2 Background and Related Work

Management of lightweight information, such as notes and annotations, has been extensively studied in both the fields of psychology and Human-Computer Interaction [14, 21, 22, 25, 54, 136, 215]. Researchers in psychology have largely studied this phenomenon with respect to both the cognitive underpinnings of how humans make sense of and organize this information [14, 85, 137, 215] and in terms of the motivation behind creating these information scraps [85, 116, 136, 137]. HCI research has explored this phenomenon from a more applied standpoint, including how users collect and create this information [54, 84, 116], annotate it [84], and use it for later tasks [25, 86, 124]. A predominant theme across most of this work is that the cost

²This work is currently unpublished. It was done in collaboration with Emma Paterson, River Hendriksen, Kazi Jawad, Connor Shannon, Andrew Macvean, and Brad A. Myers.

³This work is currently unpublished. The Catseye curation work was done in collaboration with Shannon Bonet, Matthew Shu, Andrew Macvean, and Brad A. Myers.



FIGURE 5.1: An example scenario in which an annotation's anchor is updated in multiple ways after a `git pull`. In the case only the position changes, the annotation content should stay as-is and the anchor will update to represent the new start and end positions. If the anchor is deleted, it is unclear whether to delete the annotation or to re-attach the annotation at a different, semantically similar point. Likewise, if the content within the anchor's bound changes (`this._copyVscodeMetadata` to `this._copyVscodeMetadata`, in the figure), under different circumstances it may make sense to delete the annotation, keep the annotation on the new anchor content, or re-attach the annotation elsewhere.

of externalizing this information must be kept low [54, 116] (otherwise, users have been known to favor pen and paper [54]) and mechanisms for organizing and curating the information must be flexible [54, 117, 251]. Curation typically occurs when a user wants to transform their information scraps into something that can be used by another person, such as a blog post [157] or decision table [151]. A challenge in curation is that it is typically cognitively expensive and time-consuming [152] and a user's mental model of a problem space is often evolving, so the curated representation can become out-dated [72, 98, 124]. However, failing to curate these sets of information scraps can result in the information becoming intractable [54].

Less work has explored annotations, specifically, and their relationship with source documents (i.e., anchors), especially with respect to curation. A prior literature review [5] found that the flexible nature of annotations allows them to serve a variety of purposes, including supporting in-context commenting and creating connections among parts of text. Other work noted that annotations may be seen as a conversational tool among the document users, as well as with the document creators [75, 88]. The nature of annotating and its tight coupling of user commentary and source document has made it useful in a variety of crowd-sourcing contexts, such as for question-answering on web pages [42, 102, 104] and in classrooms [70, 77, 279]. However, most of these works have been in the context of static materials, such as books, and *not* dynamic materials, such as source code, which inspired us to further explore this often-overlooked challenge in utilizing annotations for long-term information collection and sharing. Our work expands on this annotation and information management work through exploring mechanisms for both automatically re-attaching annotations given document changes and curating annotations with user interaction techniques.

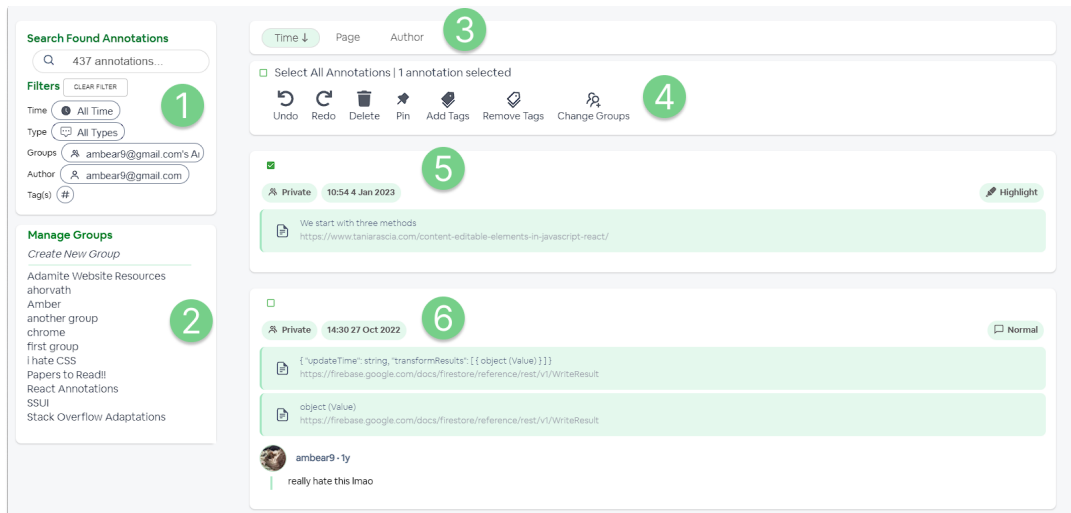


FIGURE 5.2: How the Adamite website appears. (1) is the search and filter pane for finding relevant annotations. (2) is the list of groups the user is a part of. (3) is the sort bar – currently the list is sorted by time in descending order. (4) are the batch operations a user can do on their set of selected annotations. (5) is an annotation – it is currently selected, shown at the green checkmark in the top left corner. (6) is another annotation with two anchors (the two green rectangles with text and a URL) and annotated text.

5.3 Preliminary Exploration of Adamite Annotation Curation

Prior to investigating annotation curation with Catseye, we sought to explore to what extent annotation management could be supported on the web with Adamite. Our primary motivations for supporting annotation management in the context of developer documentation annotations created with Adamite were as follows:

- Annotations can become obsolete as a developer learns more about a particular API and completes various implementation tasks, so developers should have an easy and effective mechanism to remove annotations.
- Since a user's annotations are most-likely spread across multiple web pages, developers should have a centralized hub where they can view annotations in aggregate.
- Annotations have the potential to be synthesized together into groups or documents which can only feasibly be done in a dedicated workspace outside of the context of the Adamite sidebar.
- Annotating is a social activity – thus having a space in which annotations can be shared and managed prior to sharing with others is paramount.

These design goals lead us to the development of the Adamite website, which can be found at <https://adamite.netlify.app>.

Supporting web-based annotation curation and management on a website introduced some design challenges. A key attribute of annotations is their presentation and existence in context – with a web-based browser extension, that context is the

web page and text the annotation is attached to on the web page. However, when the annotation is presented out-of-context on a stand-alone website, it is difficult to convey that context outside of merely displaying the anchor text and annotation (see Figure 5.2-5 and 6). In designing annotations for a web page, there is a question of how much context from the original web page to display – part of an annotation’s power lies in its brevity and ephemeral nature. Adding in too much information to the website version of an annotation, such as more surrounding text, may make the annotation difficult to parse. However, choosing to not provide additional context can make the annotation difficult to understand. Ultimately, we chose to preserve the annotation’s brevity and mirror its appearance to how it would appear in the Adamite sidebar with the hope that the familiar presentation would make the annotation more comprehensible. An alternative design could be to extract the surrounding HTML and styling, similar to how Unakite functions [151].

Considering one of the key goals of this work was to support curating through batch removal and synthesis, the Adamite website is where we first explored batch processing of annotations (see Figure 5.2-4). The batch processing on the Adamite website functions similarly to how batch processing on other information management systems, such as E-mail inboxes, work – a chosen operation will act on any annotations that are selected (in that their checkboxes are in the checked state). For the website, we chose to support mass deletion, pinning, tag addition and removal, and group management, where groups are a way of sharing a specific set of annotations with other Adamite users. In this way, we wanted to support moving annotations depending upon where they exist with respect to an individual developer’s sensemaking – if they are irrelevant, delete, if they are important, pin, and if they need to be re-categorized in some way, move groups or change the annotation tags. We also support undo and redo, such that developers do not, e.g., accidentally remove an important annotation and have no way of getting it back. We further explore batch processing of annotations in Section 5.6 in the different context of developer code annotations where the curation is done in-editor and the annotations have slightly different properties than the Adamite web annotations. Some of the batch processes returned, such as deletion, while others did not – namely group management, since Catseye does not have a concept of “groups”.

Through this exploration, we developed techniques for acting on annotations in aggregate that would inform our later design of batch processes for Catseye annotations. However, some activities we wanted to support, such as transforming a set of annotations into a new form or document, were not supported with the Adamite website, which we support in the new version of Catseye (see Section 5.6.2).

5.4 Preliminary Study of Catseye Annotations

As an initial exploration of what annotations to curate and how, we explored annotations already created naturally by myself as part of dog-fooding Catseye (see

Section 4.6). A subset of these annotations (122 out of the 181) were created from the fall of 2021 to the spring of 2022 and were developed while implementing and debugging various parts of the Catseye code base. In anticipation of our summer undergraduate researchers joining the research team, I hand-coded each authored annotation by whether or not I believed it would be useful to someone unfamiliar to the code base. This was done in service of finding patterns in terms of the annotation’s structure and content and whether or not I believed it would be beneficial. I chose to do this coding independently for multiple reasons – for one, this internal decision-making of whether or not to “share” a potentially “helpful” annotation is normally how such curation would happen “in the wild” [77], as people are not typically working together to decide whether or not to make an annotation public. Further, determining whether or not an annotation may be helpful requires knowledge both of the annotation’s content (which can be difficult to understand given their terseness) and the code base (which, at that point, had only been developed by myself), thus finding other researchers with the requisite knowledge to assist in the coding was not possible. Only annotations which were not already-deleted were included in the analysis since deletion suggests the annotation was no longer relevant to me, let alone relevant to someone else. Signals of usefulness included whether or not it documented some part of the code, how comprehensible the annotation was (with many of my annotations being informal in nature), and whether the code and annotation content were still relevant or not. Considering I could not definitively say whether or not each annotation would be helpful for the summer students, the codes were indefinite and included the codes “Probably”, “Maybe”, and “Probably Not”. For each determination, I included an explanation as to why I believed the annotation was probably, maybe, or probably not helpful, and, for situations in which I thought the annotation could be helpful (i.e., “Maybe” and “Probably” codes), I included an additional code for in what development context I believed the annotation could be helpful. Development working contexts in which I anticipated annotations could be helpful for included software maintenance, extension testing, API usage, understanding design rationale, or feature requests.

Through this preliminary analysis of the 91 non-deleted annotations that existed in the code base at the time, the majority of annotations were “probably not” helpful (49), while 33 were “maybe” helpful, and 9 were “probably” helpful. The most common reason for annotations not being helpful was due to the content no longer being relevant, either because the code had changed enough such that annotation’s content was out-dated (18) or because the task or bug the annotation was referring to was complete (16), suggesting that a mechanism for easily removing completed task annotations, like Adamite, would be useful in Catseye. However, detecting code-versus-annotation content inconsistencies is less trivial and part of the inspiration of our re-anchoring work. Additionally, some annotations were marked “not helpful” because the way in which they were written was too informal to be easily shareable by myself and comprehensible by another developer (11). Prior annotation research

in the classroom context has discussed that sometimes, when moving an annotation from private to public, annotation authors had to go through a pre-processing step prior to making their annotations public, in which they typically make the annotation longer and more formal by expanding on ideas that were previously just terse reminders to themselves [77]. Such a step may also be required when curating annotations prior to sharing them for code comprehension.

The 9 annotations I had the most faith in potentially being helpful were annotations related to specifications on how to implement or maintain some part of the code base (5), how to use some part of the VS Code API (2), or documented some unintuitive part of the code base (2). Similar to what we saw in the Adamite study, the majority of the potentially helpful annotations were answers to questions I had at some point (4/9) and, like higher-quality Stack Overflow answers [177], 2 of the 4 referenced external resources I used that informed my answer. 2 of the 9 annotations used the annotated code itself as an example of how to complete a task (e.g., “it’s really this simple to authenticate with github”) while 2 other annotations bound a task to the code as a reminder (e.g., “Need to update this [code] too whenever we update the annotation model”) – all 4 of these examples leverage the tight coupling of information and code to transform the code into a representation of an implementation constraint or code example. These qualities make these annotations seemingly more appropriate to be shared than others, thus ones to curate for other users.

The remaining 33 annotations were less obvious as to whether or not they should be shared. A common theme across these annotations was a lack of certainty in the subject matter, with 8 being questions and 5 hypotheses. Notably, this is similar to a finding from the Adamite study, in that unanswered questions were one of the least helpful types of annotations. Other annotations would only be potentially helpful if the student was working on that particular part of the code base and would otherwise be too confusing to be of any use (10). For example, 6 of the 10 annotations were related to how the anchoring algorithms did or did not work. The mechanism for keeping anchor points up-to-date while a user is actively editing code is easily the most complex part of Catseye and also the most prone to bugs, leading to many annotations about those parts of the code base – if the student was working on anchoring, then these annotations could be helpful, otherwise they may draw attention to this part of the code base that is difficult to understand. In this way, curating annotations for later developers may need to change dependent upon what task the later developer is performing and whether or not the annotations are relevant to them, with an annotation’s location a particularly powerful signal of relevance. Other reasons for annotations being “maybe” helpful included acknowledgements of potentially unintuitive design and implementation decisions (9) and potential feature ideas (4).

Through this exercise, we came up with the following design considerations when supporting annotation curation:

- Annotations, in order to be shareable, must be *relevant* to both the current code in the IDE and to the future reader.
 - This suggests a need for both author-side and reader-side curation functionalities.
 - This suggests a need for the relationship between the annotation content and code content to be algorithmically or deterministically evaluated, prior to showing the annotation to readers.
 - Relevance is often lost over time as code is changed, tasks are completed, and so on.
- Some amount of pre-processing must be done by the annotation author as part of the curation, prior to sharing with later readers.
- Signals such as answered questions and annotations that contain reference URLs suggest higher-quality annotations that may be candidates for sharing. Signals of lower-quality annotations include lack of certainty (e.g., “I’m not sure about this...”) and a lack of formality in tone.

5.5 Design Probe: Catseye Annotation Curation through Re-Anchoring

Initially, we did not consider the act of (re-)attaching or removing an annotation given the source code changing (hereafter referred to as “re-anchoring”) a form of curation. Instead, we considered it a technical problem – how do we find the most appropriate point to re-attach an annotation given a change in which the original code anchor was removed, moved to a completely new location, or changed in such a way the code is semantically similar but textually different (e.g., `myList.forEach(t => t + 1)` to `for(const x of myList) x + 1`)? While this technical component of the project is, indeed, a challenge, upon reviewing the Catseye annotations (including how many were no longer relevant due to code changes) and considering various scenarios in which re-anchoring is necessary (Figure 5.1), it became clear that the mere act of choosing to either remove an annotation or re-anchor it given a code change is a form of curation. With this in mind, we sought to both design an algorithm that could find potential re-attachment points and a user-interface for re-attaching or removing the annotation, given uncertainties around whether the annotation should persist or not.

5.5.1 Algorithmic Re-Anchoring

We designed our algorithm with the goal of finding the most-likely re-attachment points, given an unanchored annotation. We considered an annotation to be unanchored if its last logged anchor point that we stored in the database was either invalid⁴ when attempting to attach on system launch or the text that the anchor contains is too dissimilar to the last-logged anchor content.

```

55 export class Container {
56
57 }
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

```

FIGURE 5.3: How the re-anchoring algorithm expands its search outwards, given no sufficient match at each step. (1) is the original anchor content and location, (2a) and (2b) are the 5 lines above and below, respectively, the original anchor location, (3) is all of the code belonging to the parent node of the original anchor in the AST, and (4) is the whole file.

If an appropriate match is not found at the original anchor location (Figure 5.3-1), the search expands outwards to the 5 lines above and below the last-known location (Figure 5.3-2a and 2b), given research that source locations, given a change, are often within a small line range [208]⁵. At this point, we begin utilizing our text and distance-matching algorithm to determine whether the searched range of text contains a reasonable, potential match.

The algorithm begins by tokenizing the last known anchor string by stripping white space and syntactical symbols. This same tokenization is performed on each code line within the search space. We first see whether there is an exact match in terms of token content and order within the search space – if so, we return the source code line’s line numbers and offsets. If no exact match is found, we attempt to find the closest match.

For the closest match, we combine a variety of heuristics in order to calculate closeness. This includes the edit distance between each token in the anchor text and the source document text, the

⁴An anchor can be considered invalid when one or both of the bounds of the anchor (i.e., its starting line, starting offset, ending line, and ending offset) are not contained within the document’s starting and ending lines and offsets.

⁵Note that this approach is contingent upon Catseye having been run semi-recently, thus having a recent “memory” of where the annotation should be anchored. In situations in which the code has changed many times but Catseye has not been used, we additionally keep track of the last known Git commit the annotation was on, making it possible to move forwards in time along the commit tree and run the re-anchoring algorithm at each commit.

distance from the last known location of the anchor and the source location, and the similarity of the surrounding text in the source document to the surrounding text that is saved in our database. The surrounding text is used in combination with the saved anchor text in order to prevent the algorithm from favoring common keywords, in the case the user annotated code composed primarily of said keywords (e.g., `for`, `while`, `function`, etc.). Given these heuristics, a weighted average is created for each token, given its comparison point, and these weights are averaged by line. The algorithm then chooses the sequence of contiguous tokens that yields the highest average value – if the anchor is above a certain threshold, the anchor is added to a list of candidate anchors for the user to review. This process is repeated for each line within the search space. If no anchors meet the quality threshold, the search is expanded to the code within the parent node of the AST (Figure 5.3-3) and, if that search also does not yield sufficient anchors, the search is expanded to the whole file.

```

this._onPaste.fire({ 0.375
  text: pastedText, 0.166
  location, 0.125
  time: Date.now(), 0.10
  vscodeMetadata: this.getCopyMetadata(),
}); 0.645

```

Token: `this._copyVscodeMetadata`

FIGURE 5.4: How running the matching algorithm at the token level with `this._copyVscodeMetadata` across multiple tokens in multiple code lines would yield different results, given the string difference and location difference. Note that the most reasonable re-anchoring spot from Figure 5.1 has the highest match score.

There are situations in which our algorithm may not work. One feature of Catseye is that the user can choose to annotate any code of any size, whether or not their anchor consists of syntactically correct expressions – this means that the user can choose to annotate, for example, just the letter “u” in the keyword “function”. While it is unlikely a developer would choose to do that, the terseness of the anchor would make finding candidate anchors more difficult. We attempt to mitigate this issue through utilizing the surrounding code and taking the anchor’s original location, including line and offsets, into account.

Conversely, a user may choose to annotate a large amount of code, such as a whole file or function. Fuzzy anchoring [50] the text, given that amount of content and taking into account the additional code-structure information we have (such as what the parent AST node is and what the original name of the method was) should yield reasonable results, but it is unlikely the anchor content will ever be exactly the same in these cases. The algorithm may not find reasonable candidate locations if the file has changed too significantly since the last time the annotation was anchored. This is why we chose to show the user candidate locations and allow for other re-anchoring options, including manually choosing a new anchor or leaving

the annotation unanchored, given that it is difficult to determine if the annotation should be deleted or not in that case. Additionally, we have reason to believe that both of these types of anchors (too small or too big) are not particularly common, given that most anchors in our Catseye evaluation (Section 4.4) consisted of one line of code.

Our algorithm is also run under the assumption that there exists at least one reasonable candidate anchor location within the source code and the algorithm's "job" is to find that location. However, this assumption may not be true in cases in which the anchor point was completely removed from the code. In this case, it is unlikely that the algorithm would find candidate anchors that would match its knowledge of the original anchor code and its surrounding content and yield suggestions. Failing to find new candidate anchors may be okay in the case that the annotation should not exist anymore given the removal of the code it is attached to. However, there may be situations in which case re-attaching would make sense – for example, if a user annotated a function call, reasonable re-attachment points may include other code locations where that function was called or the definition for the function. Given the ambiguity of these situations, we ultimately leave the final decision of what to do with their annotation up to the user.

5.5.2 User Interface for Re-Anchoring

Upon finding a set of reasonable new anchor locations for an unanchored annotation, the user interface for Catseye updates such that the user can curate these annotations (see Figure 5.5). When re-anchoring an annotation, a user can view what anchor candidate(s) Catseye automatically found on launch (Figure 5.5-1 and 5) or choose to leave the annotation unanchored by dismissing the re-anchoring interface (Figure 5.5-2). In this way, we wanted to support the lightweight nature of Catseye by not forcing a user to act on the annotation if they do not want to. Catseye also sorts the candidate anchors by how confident it is about each anchor given the weighted averages, with the highest-weighted anchor shown first. In the case of many candidate anchors, the system only shows the top 10, such that users are not overwhelmed with many potentially incorrect choices. The system additionally states its confidence in the top-left corner of the suggested anchor (in this case, it determined this anchor to be a "very similar match"). Additionally, if the code has either changed so significantly that none of the system-suggested anchors are correct or the user wants to re-contextualize their annotation to this new working context, they can choose to manually re-anchor the annotation at any arbitrary point within the code base by selecting some code and clicking the "Manually Reanchor" button (Figure 5.5-3). As discussed in our earlier scenarios (Figure 5.1), there may be situations in which there is no longer an appropriate anchor point, in which case the annotation can be deleted or left unanchored.

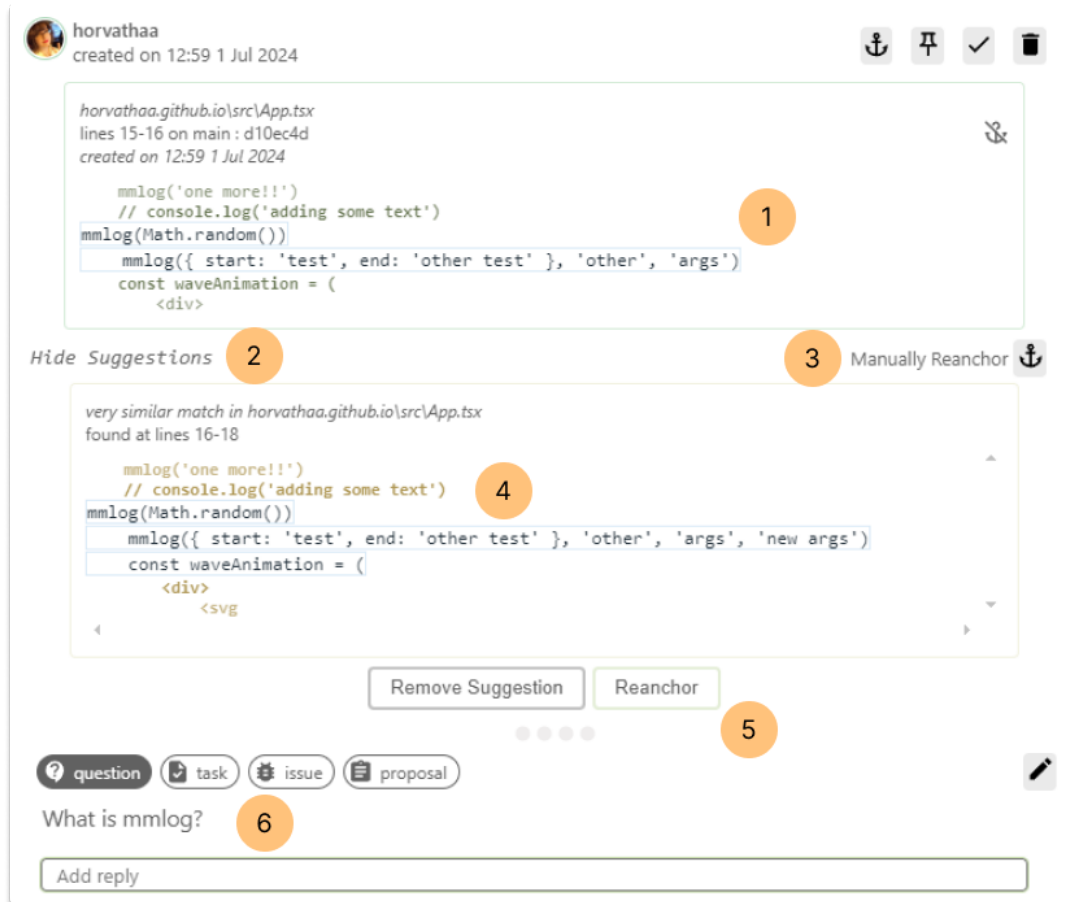


FIGURE 5.5: The user interface when re-anchoring an annotation. (1) shows the last-known anchor, along with the surrounding code in green. (2) is the Show/Hide Suggestions button, allowing the user to leave the annotation unanchored, if the user so desires. (3) is the Manually Reanchor option – the user can select some code in the editor, then click the “Manually Reanchor” button to set the anchor to their selected code. (4) is the first of multiple candidate anchors. (5) are the options for what a user can do with a candidate anchor – they can either remove that particular anchor as an option, “Reanchor” their annotation to that anchor, or click through the carousel (i.e., 4 gray dots) to view their other options. (6) is the annotation.

To further support these anchoring activities, the design of Catseye anchors was updated to include additional meta-information about the anchors to assist in curation. Anchors now display their state via the anchor icon in the top right corner to further clarify to users whether their anchor exists and is being tracked currently (i.e., anchored) or not. Catseye will also show the last-known location of the anchor, along with what Git branch and commit this anchor was last seen on and the time at which it was created (Figure 5.5-1, top left corner). While the text surrounding the anchor is primarily used as supplementary data for the re-anchoring algorithm, we also show it to the user to help them in placing their anchor in context when reviewing their annotations – candidate anchors also include this information to help users in visually assessing whether or not the anchor seems correct.

In designing re-anchoring capabilities for Catseye, a number of design challenges arose. One interesting consideration was whether or not to even have the user review candidate anchors when their annotation becomes unanchored. Indeed, given

high enough confidence in a candidate anchor, it may be an overall better user experience to simply have the system place the annotation at a new anchor without requiring the user to confirm the new location or informing the user it performed this re-anchoring. Ultimately, we decided to loop the user into this re-anchoring process whenever the system does not find an anchor that exactly matches the last-logged anchor. We did this not only to confirm that the resulting anchor is correct, but also as a forcing function for users to perform curation on their annotations. We hypothesized that the point at which an annotation becomes unanchored, such as at a `git merge` or `git pull`, is also a likely time at which the annotation is no longer relevant. We believed this due to the fact that the most common reason an annotation was marked as “probably not” helpful during review was due to the annotation being no longer relevant and this lack of relevance most often stemmed from the code having changed too much. By allowing the user to actively participate in the curation process when this change occurs, we expect that problems noted in earlier PIM and “information scrap” research where users get “lost” amongst a sea of uncurated and irrelevant notes can be avoided [54].

5.6 Design Probe: Catseye Annotation Curation through Batch Processing

As discussed in the introduction to this chapter, a key feature of information scraps is their ephemeral nature. However, when programming, a developer is already under a large amount of mental stress and, thus, doesn’t necessarily have the mental bandwidth to not only manage their code but also their annotations to throw them away. Thus, we sought to improve Catseye’s design to make this activity even more lightweight.

With this in mind, we provide functions for either increasing the utility of an annotation for a later reader or easily removing it from consideration. In this way, we see annotations as a mid-point in terms of a developer’s individual sensemaking journey in which there is a chance the information they gain will no longer be relevant, thus should be thrown away, or is a piece of evidence in a longer stream of information (or information “quest” [90]) that can be later transformed and structured into something useful. To lessen the cost of performing either of these activities, we introduce batch processing of annotations into Catseye. This batch processing works both in tandem with structuring activities, such as adding annotation types (taken from Adamite), and with culling operations such as deletion.

We chose to focus on batch processing over other types of potential information management processes for multiple reasons. The predominant reason is batch processing lends itself nicely to the information management practices we saw in our studies of Catseye and Adamite along with prior research (e.g., [124], [151]) – there is a common pattern in which people performing sensemaking forage for and generate information for a while, then, upon reaching a natural stopping point, pause

and revisit the set of information amassed during their sensemaking episode. For example, in our Catseye study, one participant, upon fixing a bug, would go back to each annotation they created related to that bug, and either delete the annotation or respond to it to “conclude” that thread of inquiry. Batch processing works well in this paradigm in that each “set” of annotations created during some sensemaking can be treated as a batch and operated on to perform the types of activities we hypothesize would be helpful. Other reasons for choosing to explore batch processing includes the ability to adapt already-known successful information management UI techniques such as filtering and because Catseye’s list view of annotation lends itself nicely to acting on items in aggregate.

5.6.1 User Interface for Batch Processing

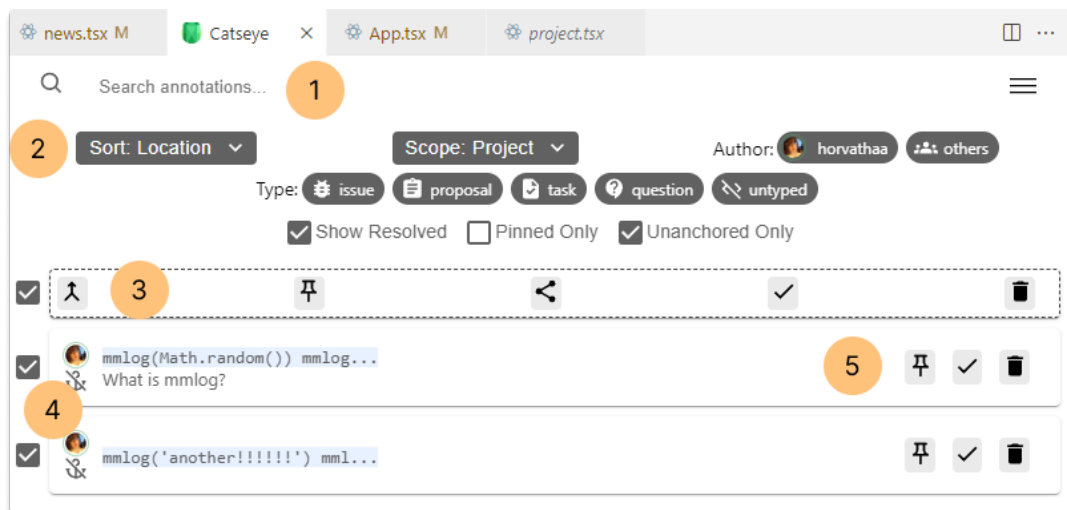


FIGURE 5.6: The user interface for batch operating on annotations in Catseye. (1) is the search bar, (2) is the set of sorting and filtering operations including sort by location or time, scope of annotations to include (including annotations made on this file, this project, or across all projects), who authored the annotation, what annotation type it is (similar to the Adamite annotation types), and whether or not the annotation has been marked as “resolved”, is pinned, or is anchored. (3) is the set of batch operations a user can perform on any selected annotations, with options for (from left to right) merging (Figure 5.7), pinning, sharing, resolving, or deleting. (4) are collapsed annotations which are currently checked, meaning they will be included in any batch operations selected at (3). (5) are the buttons for operations a user can perform on a single annotation.

The user interface for managing annotations *en masse* draws from other systems that utilize large-scale filtering and sorting of information (e.g., shopping websites). The Catseye pane now includes a top bar comprised of operations a user can perform to modify their set of visible annotations through searching (Figure 5.6-1), sorting, and filtering (Figure 5.6-2). Some specific filters, such as “Unanchored Only” and “Pinned Only”, were designed such that users can find annotations that they may want to follow up on, such as for re-anchoring an annotation. Others were designed to remove annotations that may not currently be relevant, such as “Show Resolved”, in that the default is to not show resolved annotations, which are expected to represent finished tasks or answered questions. When finding annotations to assist a

later developer, however, finding such resolved annotations may be useful to isolate answered questions that could be useful, when used in conjunction with the *type* filter to select only question-type annotations. In this way, we sought to improve upon Adamite’s typing, filtering, and sorting functions through combining multiple filters and sorts at once to isolate potentially useful annotations.

Once a user has created a filter or search to show annotations they want to do something with, the user can perform *batch operations*. The currently-supported batch operations (shown at Figure 5.6-3) include merging, pinning, sharing, resolving, and deleting annotations. These batch processes will operate on any currently-selected annotations (Figure 5.6-4). Like with the Adamite website, many of the functions are taken from similar information management systems, with the rationale that such functions would similarly make sense in moving an annotation to the next stage of processing, whether that be when the annotation is no longer useful (batch resolution, un-pinning, and deletion), still useful (batch pinning), or ready to be used by other teammates (batch sharing). However, we found a need to also *transform* annotations prior to their usage by other teammates, leading to the development of our novel merging interface.

5.6.2 Merging Annotations

Merging annotations is the act of taking two or more annotations and combining them to form a single new annotation, while deleting the original annotations (see Figure 5.7). Considering a large part of curating content to be usable by later users is synthesizing information into a more connected and comprehensible form [151], we sought to provide an interface for that process that not only supports this synthesis but also removes the “leftover” content after that work is done, thus supporting two aspects of the curation process in one interface.

The user interface for merging annotations allows for a high level of specificity in terms of allowing the user to decide what parts of each annotation they want to keep and what they want to remove. Each core text-based component of an annotation (i.e., anchor, annotation content, replies) can be imported into the merged annotation. When text from an annotation body or reply is added to the annotation content (Figure 5.7-5), the content is automatically marked with additional meta-information including who authored the information and the time at which it was originally authored such that the original context is available. We chose to support adding reply content to the main annotation body considering many of our most useful annotations, both in this study and in the Adamite study, are question-answer pairs in which the answer was typically in a reply that may be missed, thus adding that relevant and important information to the main annotation content could help a later user’s sensemaking. The user can also choose to expand upon, pare down, or otherwise transform the text, e.g., to be more formal since the information is imported directly into an editable text field.

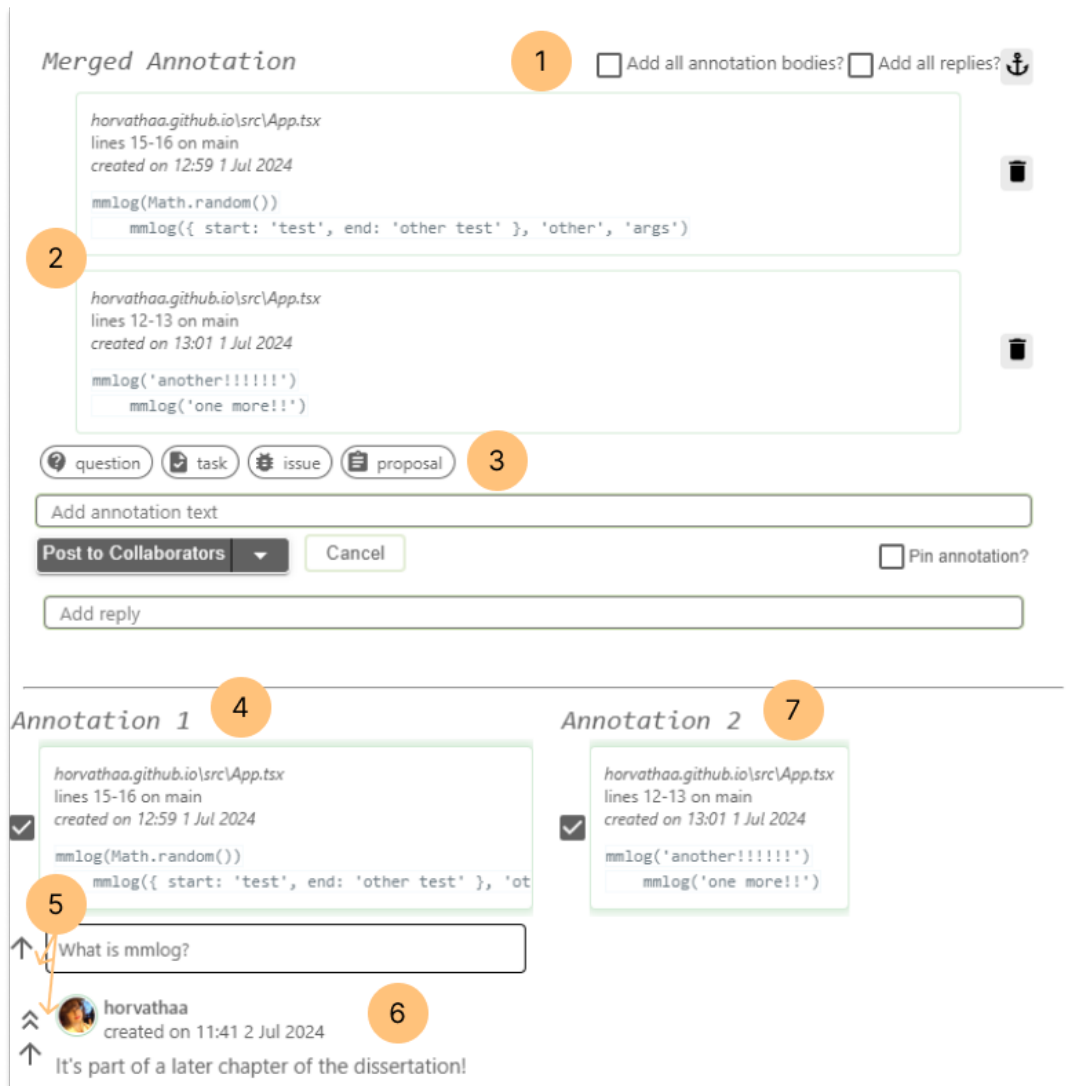


FIGURE 5.7: The UI for merging annotations. (1) are batch operations for creating the merged annotation through importing all of the content from the original annotations. (2) is a preview of the annotation anchors, with the option to remove them (trash can icon to their right). (3) is the regular annotation authoring UI, such that users can add types and text that may not be in the original annotations. (4) is the first annotation that is being used to create the merged annotation – the checkmark shows the anchor will be included. (5) points at the icons for importing annotation content and replies into the resulting annotation – annotation replies can be added to the merged annotation’s main body text by clicking the double chevron icon, while the single upwards arrow will bring the reply’s content up as a reply. (6) is the reply content for this annotation. (7) is the other annotation to be included in the merged annotation – since it is a highlight annotation, it has no annotation text or replies.

We also support batch operations within the merge interface itself. Users can choose to add all annotations or add all replies (Figure 5.7-1) to their resulting annotation. Given the level of complexity of this UI, one can imagine a developer being reticent to use it since the cost of authoring new annotations is relatively low, thus creating a completely new annotation that includes content from other annotations may be easier than using the merge functionality. However, we expect that in the case of more complex merge cases, such as merging many annotations or merging annotations with multiple anchors across different files, these batch operations in

conjunction with the high-level of specificity afforded by the interface would be a better user experience than attempting to make the annotation from scratch.

5.7 Discussion and Future Work

In designing for annotation management, the user interface had to balance potentially-competing goals – support complex activities while still feeling lightweight. In my personal usage of Catseye and in our user studies, developers wanted to be able to jot down an otherwise throw-away thought about their implementation, bug, or something else they discovered in order to offload some of the intense cognitive load they were under. However, that in-the-moment need subsequently leads to more work later through either culling or transforming this information. We do not want to deter developers from generating information due to anxiety about later work, given that this information can be useful. This tension between supporting in-situ lightweight information generation and the subsequent work that lightweight information may incur inspired our exploration of this design space.

A particularly interesting aspect of this design space is the explicit connection between information scrap and document (i.e., annotation and anchor). Prior work in lightweight note-taking and PIM highlighted the importance of information location [21] and how that spatial component of, e.g., placing a sticky note with a grocery list on the fridge can be particularly powerful in reminding and recalling information. This connection between note and artifact has been less-explored in traditional digital spaces due to the 2-dimensional nature of a desktop environment leading to less “space” to work with and a lack of support for general-purpose user-facing customization of digital spaces. With our re-anchoring work, we further explored how this connection between information and source can be complicated in a dynamic space, such as a text editor, where information is constantly changing. This leads to questions of how to best retain the benefits of spatially organizing information when the space is evolving over time – returning to our scenario shown in Figure 5.1, in the case that the code is changed, our re-anchoring algorithm would correctly find the most reasonable new anchor given content and location (Figure 5.4). However, it is still unknown whether the annotation should persist at all – in the example scenario, the user asks “where does this value come from?” about the class property `this._copyVscodeMetadata` and, after the `git pull`, the class property is replaced with a method call (`this.getCopyMetaData()`). Perhaps the addition of a function call answered the user’s question as the call implies the “value” the user was referring to is returned from the function. However, an alternative scenario is that the initial question is expressing confusion around the whole `vscodeMetadata` property, in which case the question may stand. Given annotations’ typically terse nature and the level of cognitive overload a user is typically experiencing when choosing to write such a note, it is unreasonable to expect more exposition that would make classifying these annotations possible. Some very recent work has explored using

LLMs to automatically and intelligently re-anchor annotations with some success, but the approach is expensive and does not fully address the disconnect between ambiguous annotation content and code [168]. In this prototype, we chose to keep the human in the loop given this issue of text-versus-code ambiguity, but our next work, Sodalite, explores additional mechanisms for highlighting text and code inconsistencies with the goal of leveraging the anchoring relationship for document maintenance.

Sodalite additionally builds upon the idea proposed with the batch processing work, namely that annotations can serve as “building blocks” for larger, more curated documents. Merging annotations is one way of taking intermediary information fragments and transforming them into a more shareable and comprehensive form. Sodalite, similarly, allows the user to take locations within the source code and write about them using its specialized mark-up system, but no longer “imports” annotations as a starting point. Sodalite was primarily focused on long-form document maintenance, so leveraging and curating annotations was less relevant. Nonetheless, given how LLMs have evolved since this work, generating documentation through using annotations and their associated code may be a viable alternative to writing documentation from scratch (Sodalite) or merging and editing annotation content (Section 5.6.2).

A challenge in evaluating this type of work is curation can only really happen after a user has amassed a collection of information that necessitates curating. This is the primary reason this work did not result in a published, stand-alone manuscript. An ideal experiment set-up would allow us to evaluate the goals of our curation work: taking a set of information amassed across time and making it useful and comprehensible to both the original author and potential collaborators. A study design that could investigate those requirements could include recruiting multiple participants to create many annotations across a long period of time (i.e., a month or more). After this period of time in the example study, participants would then be instructed to go back and curate that set of information using our tooling. Lastly, we could then evaluate whether that curated set of data helps a new developer better comprehend some code as opposed to other baseline comparisons, , for example, a “raw” set of annotations (what one would get in the old version of Catseye) versus code comments or some other form of comprehension aid (e.g., documentation) versus no annotations nor other comprehension aid. However, this design introduces multiple challenges. One challenge is that such a set up requires recruiting people to create a realistic set of annotations across time and, if the participants choose not to use the tool, the second part of the study would not work as there would be no annotations to curate. Another challenge is establishing a baseline comparison point – to make the proposed example study design work, a participant who created and curated annotations would need to both share their code for the purposes of the comprehension evaluation and would need to create multiple versions of their code,

including a version with code comments, no comments, and just their original annotations prior to the curation. In contrast to related work [124] where participants perform online shopping (an activity anyone can perform on nearly any device, platform, etc.), the problem domain we are working in is highly contextualized to an individual developer's work, making it more difficult to artificially create a sense-making hand-off that is comparable to other sensemaking hand-offs. An alternative study design could be more qualitative in nature and explore the practices developers naturally develop, if any, in curating their annotations, but such a study could not make any claims about whether the curation process was actually helpful.

In pulling out lessons learned from this work, a key takeaway is:

- **Time, anchoring, mutability, and provenance all inform relevance, which is key.** Previous chapters of this dissertation were largely focused on a singular session in which meta-information is relevant. While this allowed us to explore the utility of the features Adamite and Catseye have and how they may assist in developer sensemaking, they did not engage in the complexities that day-to-day usage of an annotation system entails. This work sought to probe on these complexities through looking at, broadly, how to manage annotations across a long period of time, with respect to the anchor, how it changes given document mutability, and how the relevance of the annotation is related to these different factors, including its provenance. In taking into account all of these different factors, a developer performing curation also has to consider the *audience* – are they curating for their benefit in order to reduce clutter? Or, curating to make something comprehensible for a later developer? Most annotations that were deemed irrelevant were irrelevant given time passing, thus potentially having no anchor due to the source changing, and the provenance of the code, including the original intent, no longer being relevant. Distilling all of these factors into a function for assessing relevance may be a path forward for someone wanting to automatically curate a set of annotations.

The connection between document and information is particularly powerful in immediately placing information in context and, through that context, natively supporting the evolution of information over time. By abstracting this information onto its own plane, separate yet connected to the source code, annotations can exist without interfering with the “business logic” of the source code or development environment. This is in contrast to other systems which change the overall structure of the IDE [29] or interfere with the source code [247] in order to support spatial orientation and connectivity between meta-information and code. The invisible binding between source code and information that the re-anchoring work and, more broadly, this thesis explores and highlights how developer's rich and textured inner world of tacit knowledge about code can be translated into meta-information and bound to the code itself, given proper tooling.

Chapter 6

Sodalite: Meta-Information to Support Documentation Management

This chapter is adapted from my paper:

[106] Amber Horvath, Andrew Macvean, and Brad A. Myers. 2023. “Support for Long-Form Documentation Authoring and Maintenance”. In *2023 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. Washington, D.C., October 2-6, 2023. IEEE.

6.1 Overview

The connection between code and text written *about* code has long been studied and explored. From literate programming [48, 125, 225], which has had a recent resurgence with the rise of Jupyter notebooks [121], to various forms of documentation [1, 45], there exists a tension between making code comprehensible through expressing the design intent in a more naturalistic and human-readable form, while allowing for the control that traditional programming languages provide. Our approach with Sodalite¹, sits at the intersection between these two approaches by allowing the source code to be *anchored* to the corresponding written text that seeks to e.g., describe how it functions, or how it came to be.

Once the text is anchored, this relationship can be utilized in new ways. While Catseye and Adamite automatically attempted to update and repair any links between annotation and source text or code that appeared broken, this process sometimes failed and the annotation would be thrown away. This often occurred when the original source document changed in such a way that there was no longer an appropriate anchor location for the annotation, such as during a Git pull or if the

¹Sodalite is a royal blue mineral and stands for Stories for On-boarding as Documentation Authoring, Leveraging IDEs for Text Enhancements.

API documentation changed between releases. While the fact that the annotation is no longer anchored, thus, may no longer be appropriate for the author or a reader to reference, the fact that it became un-anchored itself may be useful. For example, if I have annotated an API method's parameters with the specifics of my particular implementation, it would be useful to know that, in the new release of the API, these parameters no longer exist and my implementation will need to be changed. The challenge of keeping code and documentation in sync with one another has been a long-standing research problem [4, 212, 249].

This chapter explores utilizing the connection between code and text to overcome longstanding barriers in documentation authoring and maintenance. By leveraging code meta-information including the structure, existence, and location of code, Sodalite is able to automatically determine the likelihood of the documentation being out-of-date given how well the system performs at re-attaching each code-text pair to the code in the IDE. Sodalite also uses this meta-information to help the initial documentation author in writing their "code story" through suggesting code patches to document, given the author's already-documented code, and providing templates for bootstrapping the authoring process. We expect this documentation authoring and maintaining process to be particularly useful for developer teams or open source projects, where the documentation and corresponding source code can "live together", with team-oriented documentation suffering from many of the same barriers that end-user facing documentation does [113, 210], including out-of-date information [210], while typically lacking a dedicated team or process for updating the documentation [3, 53].

6.2 Background and Related Work

Prior work about writing software documentation have ranged from understanding what information is in documentation [97, 158], what the problems are with that information [3, 4, 35, 68, 143, 163, 180, 212, 214, 249], how software documentation is authored [53, 222, 236], and automating documentation processes to offset authoring costs and standardize the information present in documentation [2, 91]. Notably, the majority of this documentation work is in the context of software documentation for end users of a software library, e.g. API documentation [174]. Most relevant to our work are studies about documentation created for other developers working on the *same code base*, such as internal documentation about the code base [210, 222] or a code tour for an open source repository [244], where the primary goal is to help other developers understand and contribute to the code base. Our work expands upon this work by contributing a system designed specifically to help create this type of documentation, with a focus on combating some of the known issues in authoring and maintaining this documentation including out-of-date information.

Once some software documentation is made, researchers have studied how developers maintain those documents and use that information. In studies of usage,

researchers have identified many problems of documentation that lead to the documentation being less trustworthy [153], including questions about how up-to-date the information is [4, 143, 249] and completeness [212, 214]. A survey of developers at one company reported that they rarely updated documentation and that they correctly assumed that documentation content is out-of-date [143]. One reason for this lack of maintenance is that finding the appropriate places to update given a change can be challenging, with developers reporting that they would value a tool that helps identify those locations [76]. Our system attempts to reduce some of these costs by having maintenance be a core design consideration by automatically locating and highlighting the out-of-date portions of the document given which code was changed.

6.3 Sodalite

In order to use Sodalite, a user begins by opening the Sodalite webview [167] which appears and functions like any other file in the editor (i.e., it can be dragged, resized, closed, etc.). The user can then view the stories that have been associated with the user’s currently-open GitHub repository or choose to author a new story (see Figure 6.1). When authoring a new story, the user will be presented with a rich text editor, hereafter referred to as the “story editor” (in contrast to the “code editor”, which refers to the Visual Studio Code IDE), and users can utilize code story *templates* and *code links*. For already-authored code stories, the system performs its maintenance algorithm to determine what code links are potentially invalid.

6.3.1 Templates

We included templates for stories to help guide authors about what should be documented (a common problem with documentation authoring is not knowing what is important to document [53]) and to provide built-in mechanisms for allowing developers to find the right code and code-related information to include in the documentation.

We began by identifying various types of developer long-form writings that are not adequately supported by current tooling and would benefit from leveraging the context of the source code. We compiled this list through a combination of reviewing literature and informally consulting with software engineers. For each of these types, we developed a “template” which contains a list of headings, subheadings, and guided instructions with prioritized code links (see Section 6.3.2), that would most likely be included. This list includes:

- **Overview documents**, which serve to introduce a newcomer to a code base [210, 244]. This template prioritizes higher-level information, including references to functions and classes, with the assumption that developers most likely do not want to cover lower-level specifics in an overview. “Overview” also

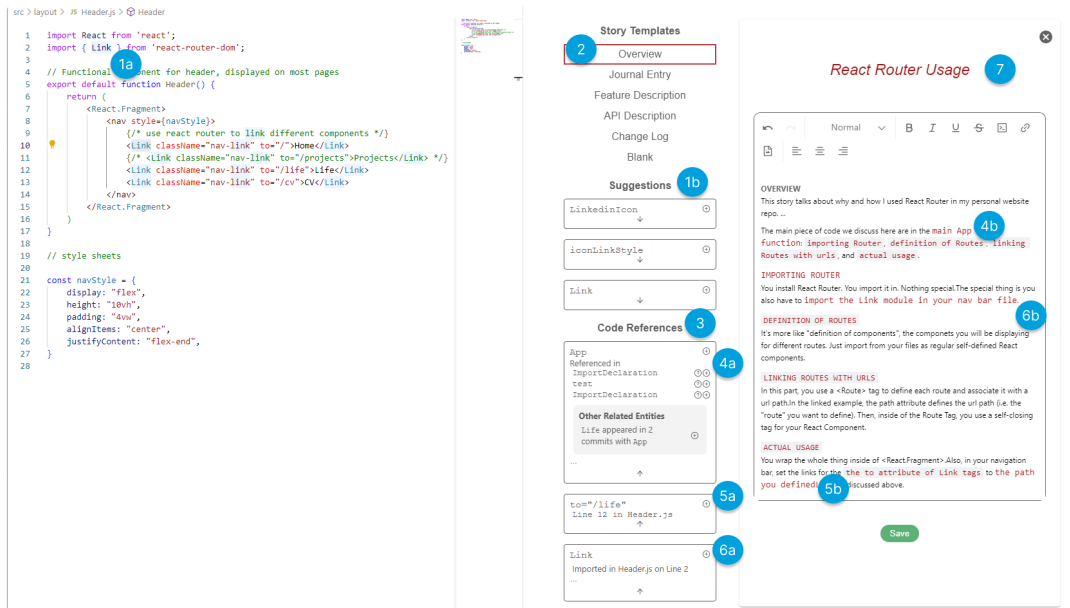


FIGURE 6.1: The editor for authoring a story using Sodalite – this is a simplified and anonymous recreation of P2’s story for demonstrative purposes. (1a) and (1b) show the relationship between the code editor and the story editor – in this case, the user has clicked the code `Link`, prompting the suggestions pane to show identifiers related to `Link`. (2) is the template list with the selected “Overview” template highlighted. The “Code References” at (3) are represented in a list (4a, 5a, and 6a) and their locations are shown in the story text at (4b), (5b), and (6b), respectively. The title of the story is (7).

prioritizes identifiers with certain keywords in their names that are commonly associated with control flow in a web development project, such as “listen” or “handle”, with the rationale that the logical flow of a project may be discussed in an overview.

- **Journal entries**, which serve as logs about what a software engineer completed in a workday and are sometimes required by management [157].
- **Feature descriptions**, in which a developer discusses what a particular part of a code base does and is has been claimed to be one of the most important types of documentation for developers [143].
- **API descriptions**, where a developer catalogs what they have learned about an API to benefit later developers [104].
- **Change logs**, in which a developer enumerates parts of the code base that have changed between releases of a software project [157], but, due to authoring costs, can be incomplete [38].

Note that we are not claiming that this is an exhaustive list and our system allows developers to define their own template if the built-in ones are not adequate by using a predefined JavaScript object notation (JSON) structure that Sodalite reads. The system also defaults to a “Blank” template, such that the user can define their own structure without extending Sodalite.

6.3.2 Code Links and Suggestions

A key feature of Sodalite are *code links*, where the story references back to code. Links can be made manually or using the “Suggestions” pane (see Figure 6.1-1b), which lists code links that the author may include in their code story. Code links come in three varieties:

- **Identifier definitions**, which link to where a specific code entity like a method or variable is first defined. A user can create an identifier definition link by selecting the “plus” button in the suggestions pane at the top right corner of the code link box. Identifier definitions also include additional information about the identifier, including places in which it is referenced, and other classes or functions it references (see Figure 6.1-4a).
- **Identifier references**, which link to a specific instance in which a particular identifier is referenced or used. An author can make an identifier reference by selecting a reference that is listed in an identifier definition’s list of referenced locations. Figure 6.1-5a shows a reference.
- **Code ranges** (see Figure 6.1-6a), which can be any arbitrary range of code that the user has selected in the Visual Studio Code editor. Selected code in the code editor will always appear at the top of the user’s “Suggestions” list, given that selecting code is a strong signal that the user wants to include that code in their story. Code ranges are fundamentally the same as the “code anchors” that Catseye supports.

Code links can be added to a code story in two different ways. If the user has no text selected in the story editor, the code link will either insert the name of the identifier (identifier definitions and references) or the selected code (code range) at the location of the user’s cursor. If the user has selected some text in their code story, that text will be linked to their code (see Figure 6.1-4b, 5b and 6b). Either way, clicking on a code link within the story editor will navigate to wherever that particular code link is located in the code editor. Once a code link has been added to a story, it will appear in a “Code References” list (Figure 6.1-3), such that the link may be used elsewhere in the story.

Code links can also contain additional meta-data Sodalite was able to determine about that part of the code. This includes, for the identifiers, where they are defined, and referenced in different parts of the code. Once a code link has been included in the code story the “Suggestions” pane will include other identifiers that were commonly edited at the same time as that particular identifier. We identify these “co-edits” by parsing the Git commit history for that particular identifier and count when other identifiers appear in the same commit. In this way, we attempt to identify parts of code that are related but not in a way AST parsing would find.

Given these different types of code links, the “Suggestions” pane uses different sources of information to populate its list. Sodalite examines both what the user

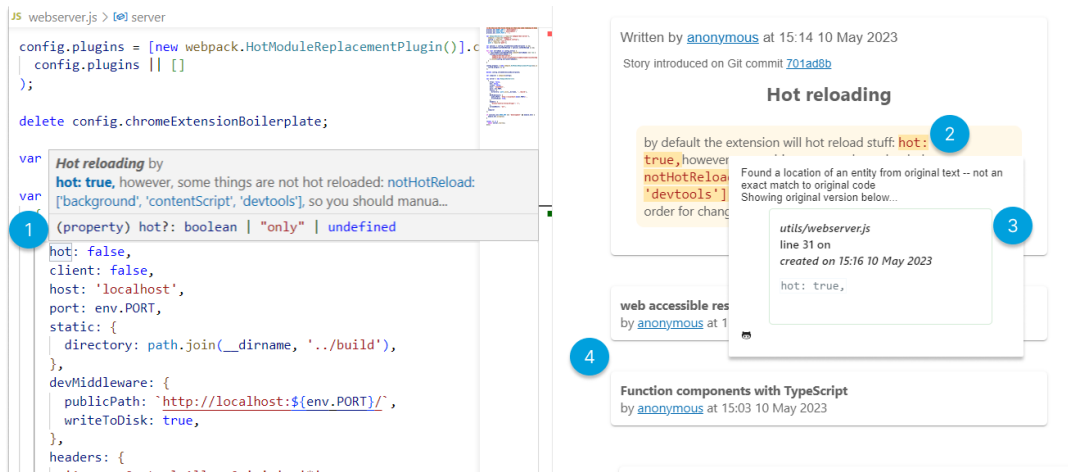


FIGURE 6.2: How Sodalite appears after a story has been authored. (1) shows the hover text when a code link is interacted with in the code editor. (2) is a code link that has been marked as “needs review” given that the original code (shown in (3)) and the code in the editor (at (1)) are different. (4) are two other collapsed code stories.

is doing in the story editor and what they are doing in the code editor. The information the algorithm leverages from the story includes what references are already included in the story and what the user has most recently typed. Identifiers related to references already within the code story will be prioritized, as will identifiers that match some part of the most recently-typed text. The system then complements that information with what it knows about the user’s current location within the code, including if the user is currently selecting an identifier, and, if so, what identifiers are related to the selected identifier, along with other identifiers referenced in the file and, if applicable, in the user’s currently selected scope. Some templates like the “Change Log” prioritize certain types of code links (e.g., code links that have been heavily edited), in which case the system favors those identifiers in the Suggestions list. Identifiers that appear in multiple information sources are pushed to the top of the list. To prevent the user from being overwhelmed with suggestions, we limit the amount of suggestions to the top 10.

When a story is saved, Sodalite generates a JSON file in a system-generated folder, which can be committed to the user’s Git repository and used by other programmers when they read or edit the source code. This JSON file is also used by the system when determining what code links are out-of-date.

6.3.3 Support for Reading

Sodalite has some features designed specifically to help readers of code stories better understand and utilize the documentation. A core feature of Sodalite for readers of code stories is the fact that it is situated within the context of the code. Developers have previously stated they value source code and code comments more than other types of documentation [232] – we hypothesize that bringing the type of information typically in external documentation into the source code will allow for “the

best of both worlds” by staying within the developer’s working context while also supporting longer-form text.

Another feature of Sodalite that we expect to help users of code stories are the bi-directional nature of the code links. When some code has been linked in a story, the code will be highlighted within the IDE, such that users of the code can discover pertinent documentation that is relevant to that code (see Figure 6.2-1). The highlighted code brings up a hover text that shows a preview of the code story, including the surrounding text from the story, the part of the story that is linked to the code in bold, and the name and author of the story. If the surrounding story text contains code links to other parts of the code, those links will navigate the user to the code link’s location. If the user clicks on the link in the hover text, the Sodalite pane will open and scroll to the correct part of the story. This does not interfere with the regular links in Visual Studio Code’s hover text, which still work like normal.

Users of code stories are also presented with additional context about the originally authored story, to further help them validate what information is still relevant. The story shows who authored the story, the date and what Git version the project was on when the story was last authored or modified, and each code link in the story displays additional metadata about what version of the code it was on when authored, along with its original code content and context. This metadata complements the information the out-of-date system provides when the story is validated.

6.3.4 Support for Maintenance

Sodalite leverages being in the code editor and having access to Git versioning information to be able to mark parts of stories as “valid” (in which all code links are valid), potentially “in need of review” (in which the system found a potential match for the code link, but it is not positive – see Figure 6.2) or definitely “invalid” based on how well the system is able to match the code references in the story to the code currently in the user’s project in Visual Studio Code. There are situations in which the documentation may go out-of-date that our system would not capture, but, in a study of documentation problems, [4] found that most cases in which the documentation went out-of-date was due to the code changing, so we focus on that case with Sodalite.

On launch, Sodalite parses every code story file in the user’s current project and builds an internal AST representation of the code in the project to compare the code links against. We use the different types of code links to inform how to re-attach a particular link and whether that attachment is valid or not.

For code anchors, given their flexible code structures (e.g., anything from a string to a full multi-line expression), we begin by evaluating the most optimistic condition, in which nothing has changed, through checking whether the position we have saved contains the same code as the code link. If not, we then look for whether the code from the code link exists anywhere within the document. If that does not

work, we use purely text-based matching mechanisms to discern candidate matching points, since that was the most successful method used in [208]. We use AST information to restrict our search spaces, using a recursive function from most within the last-known scope location.

In the case of multiple matches or no matches, we run our re-anchoring algorithm, partially adapted from [208] and modified slightly from the algorithm discussed in Chapter 5 to account for the new types of code anchors. This algorithm uses the last-known code content, the original location within the file, and the original surrounding code (all of which are saved with each code link) to find the most likely candidate code reference. We then weigh the probability that the location is correct through a combination of calculating the edit-distance between the two versions, the edit-distance between the surrounding lines of code, and the difference between the candidate line(s) of code and the original location, with a closer location being weighted higher.

If there is a low score or no matches, we mark the anchor as invalid and the corresponding text in the story as in need of review. To assist the author in finding either a new location to link the story text to in the code or to discard the link, we present additional metadata about what the code looked like and where it was last located (see Figure 6.2-3, where it says “Showing original version below...”).

For definition references, the system searches the internal representation to see whether the identifier is defined at its last known location. If the definition is not there, the search will expand outwards to see if the identifier is defined in a different file and, if so, attaches to that location, but marks the link as potentially invalid, considering it may be a different identifier that just happens to have the same name. If it is a local definition (e.g., a property named `bar` that is part of class `Foo`) the system will check that both the class `Foo` and the property `bar` exist. If the system cannot find a definition for the particular referenced code entity, the system uses the text-and-location based re-anchoring algorithm. If the algorithm does not return a result of a sufficiently highly-weighted likelihood, we mark the reference as invalid, and the surrounding text within the code story as in need of review by the author.

A similar approach is used for checking identifier references in the code. We begin by seeing whether there are one or more references in the code in the last-known scope in which the reference was used. If there are multiple candidate matches, we find the reference with the most similar AST path to the saved code anchor information. If there are no candidate matches, we once again fall back on the text-matching algorithm and, if the result is inadequate, mark the text within the story as potentially invalid and in need of review.

By attempting to automate the process of re-matching the story content to the corresponding code and suggesting potential edits to make, we attempt to lower the workload for the author in maintaining the documentation. “Valid” code links have no highlights. “In need-of-review” code links (shown in Figure 6.2-2) are highlighted in yellow as a warning to readers. “Invalid” code links are highlighted in red to draw

attention to this part of the story. Like in our discussion of re-anchoring in Chapter 5 and letting an annotation author decide how or if at all to re-anchor their annotation, we ultimately leave fixing code links through applying edits to the code story in the hands of the developer. We chose to do so with a similar rationale to last chapter – the point at which some documentation becomes out-of-date due to broken code links is probably a good time to review the documentation and its content, given that significant changes most likely occurred. Leaving the code links in their invalid or in need of review states, we hypothesize, would motivate that reviewing process by a developer.

6.4 Evaluation of Sodalite

6.4.1 Study Design

In order to assess the usability and usefulness of Sodalite, we ran a small user study to test both the authoring and maintaining aspects of Sodalite. Participants were instructed to use Sodalite to document some code of their choosing. We then observed them as they chose to document whatever information they saw fit.

To evaluate the maintenance features, participants took their documented code and reverted to an *older* version of their code. We considered requiring participants to document an old version of their code, but decided against this since we did not want the study to become a test of how well the participants remembered their old code and developers are not typically documenting their old code. Therefore, they documented the current version, and we pretended that an old version was a newer version.

Participants were instructed to look through their project’s GitHub history to find a version that was many commits behind the version they documented and involved edits to any files that they ended up documenting during the authoring portion of the study. Upon finding an appropriate commit, participants checked out that version of the code and re-ran Sodalite, which triggers the maintenance system to try and find appropriate re-anchoring points. The first author and the participant stepped through each code link within the stories to determine whether the system chose the “correct” action, meaning the participant agreed with the system’s determination of whether the link was “valid”, in “need of review” or “invalid”. The full set of study materials can be found in Appendix E,

We recruited 4 participants (3 men and 1 woman), hereafter referred to as P1 through P4, using study recruitment channels at our institution and advertisements on Twitter. All participants were required to have some amount of experience using JavaScript or TypeScript, to have a project written in one of those programming languages that they were willing to document, and to regularly use Visual Studio Code. All sessions were conducted over video conferencing software and the sessions were

recorded. Participants had, on average, 12.25 years of professional programming experience and considered themselves proficient in JavaScript and TypeScript.

For analyzing the authoring experience, we captured how many stories the participants authored, how long each story was, what type of templates the participants used, how many code links they chose to include across how many files, and what types of code links they were. We additionally reviewed the recorded videos to count whenever the participant navigated through the document using the code links, and counted whenever a participant copied and pasted or deleted a code link. We additionally noted any usability issues participants discovered with the tool.

To assess how well Sodalite did at *maintaining* the code stories, we computed the false positives, false negatives, true positives and true negatives for each code link given the participant's evaluation of whether or not Sodalite correctly re-attached the code link. We additionally logged what the offset was from the code location during the authoring session to the maintenance session, along with the total diff between versions, and what type of re-attachment strategy the system used (e.g., text-based, AST based, etc.).

6.4.2 Study Results

All participants were able to successfully use Sodalite to author at least one code story that utilized Sodalite's features, and agreed with Sodalite's re-attachment decisions, on average, 86.5% of the time. We further discuss participants' experience using Sodalite's authoring and maintenance features.

Authoring

The 4 participants authored, in total, 15 stories, with 2 participants each writing one story, and the other 2 participants authoring 6 and 7 stories, respectively. On average, participant's wrote 238 words (min = 136, max = 312, std. dev. = 75.4) during the 45 minute authoring time. The participants in our study chose code repositories of varying purposes, sizes and complexity to document. Participants, on average, chose to document 3.25 files. Our participants chose to document a Google Chrome extension, a Visual Studio Code extension, a browser game, and a personal website.

Across all stories, participants created 52 code links, with each participant, on average, creating 13 (min = 11, max = 18, std. dev. = 3.36). 18 of the code links were pointers to definitions of identifiers within their code², 7 were identifier references, and the remaining 27 were "code ranges". In their stories, participants, on average created 12.75 code links.

The code links were represented in different ways within the code story. 29 of the 51 code links were attached to text in the code story that was just the name of the identifier or the expression the user had selected. Notably, 18 of those 29 code links were definitions of particular identifiers, where the story sought to explain some

²P1 made exclusively this type of code link, comprising 11 of the 18 total definitions.

detail about that identifier. The remaining 22 code links were attached to text that was part of a sentence within their code story. This split in usages suggests that supporting both of these ways of referring to code helped support different ways in which participants wanted to talk about their code.

Participants sometimes used their code links to help them navigate through their code, using, on average, 41% of their code links at least once to navigate through their code (min = 0, max = .722). Participants would sometimes check, after including a code link, where that link went as a way of assessing the reader experience. P4, in particular, used the navigation features often, revisiting the majority of their code links to review the other places in which the code was referenced and used their links to navigate 13 different times across the session. This suggests that the code links may be useful for navigating through a code base, even when authoring documentation.

Maintaining

Sodalite made the correct decision of either re-attaching or not re-attaching 44 out of the 52 code links, a success rate of 86.5%. Of the incorrect links, 7 were because the system could not find an attachment point, despite there being one within the code, and 1 was because the system erroneously found a location that it thought was correct when it was not. The majority of cases were that Sodalite successfully re-attached a code link, at 36 instances. Eight code links had changed in such a way the system correctly marked them as needing review. The reversions averaged 100 lines of code lost (min = 53, max = 137, std. dev. = 38.92) and 49 lines of code added (min = 3, max = 95, std. dev = 51.97) per file.

Sixteen of the correctly re-attached links were definitions, in which the re-attachment strategy was to find a point within the file in which the named entity was defined. Given the relative broadness of this strategy and the specificity of the named functions or classes participants wanted to document, all but 2 definitions were able to find an appropriate point somewhere in the document. This strategy may not work as well if a method or class is renamed, but that did not happen to occur in any of the participants' versions. The reference re-attachment strategy also worked well, with 8 of the 10 references being successfully re-attached in an appropriate location.

Perhaps more interestingly, the 8 cases that the system correctly marked for review occurred in different situations, including when some code had been commented out, the underlying semantics of the code changed significantly enough such that the text in the story was no longer valid, and when the code links' corresponding definitions did not exist. One such instance is shown in Figure 6.2 in which the original code link and surrounding text states that a value, `hot`, should be set to `true`. However, the code in the version the user has opened has `hot` set to `false`. Two of the cases occurred because the definitions for the linked methods did not exist in that version of the code.

The single false negative case, where Sodalite missed a change, also occurred using the text-based matching algorithm. The anchor was originally placed on an `else` keyword, but, when the prior version was pulled in, the system placed the code anchor on a completely different instance of `else`. This other location was closer to where the original link was and, since `else` is such a common keyword, the system found multiple reasonable locations but selected the incorrect one. For identifier references, Sodalite flags situations like this, but the system does not currently support that for keywords. Future versions of Sodalite would benefit from introducing more heuristics about the context in which the code appears, including any logical dependencies, such as an `else` statement's corresponding `if` statement.

6.5 Discussion and Future Work

Our evaluation lends evidence to the claim that in-editor authoring and maintaining of documentation can be supported with Sodalite. Participants succeeded in creating code stories and, when faced with changed code, the system was very successful in identifying problematic references, and relatively successful in choosing the appropriate action to take.

A current limitation of Sodalite is that all of the mechanisms for determining whether or not the story is out-of-date are contingent upon the story including code links. The expectation is that, given Sodalite being located within the IDE, developers will naturally utilize that context and reference their code within their code stories. Indeed, all participants stated in the post-task survey that they valued the ability to link their text to their code. Nonetheless, future versions of Sodalite may benefit from additional mechanisms for assessing the validity of the text in comparison to the code, for example, through leveraging crowd-sourcing mechanisms for quality control [6, 53].

Another benefit of having even better mechanisms for identifying out-of-date code stories is to notify documentation writers and/or developers when their story has gone out-of-date. It is a known problem that updating documentation is a task that developers typically put off – while Sodalite attempts to make updating easier through marking places for review, the system does not require that the author take any action. One can imagine an additional feature to Sodalite that allows stories to be marked as “very important”, in which case developers could immediately be notified if they make a breaking change, while lower priority stories can be addressed on the developer's “free” time.

In placing Sodalite in the broader context of the work discussed in this dissertation, Sodalite serves as an interesting mid-point between the earlier and later works discussed in this thesis. The system represented a change in mentality from providing new ways of authoring programming-related meta-information and into

defining and understanding new ways of leveraging the connection between meta-information and its source to support richer interactions and functionalities. Cats-eye began by probing on that thread with its unification of functionalities given a shared source location but was still largely contingent upon a developer beginning by authoring some note and only after that point building upon it with more information. Similarly, the curation work discussed in Chapter 5 was also explored with the assumption that a developer has some authored information. Sodalite still requires the developer to author a code document, but leverages the *anchor* connection for new interactions and functionalities – in this way, we adopted a paradigm and mindset shift: we previously fought *against* source mutability and subsequently considered changing and broken anchors as a problem. Meanwhile, in this work, we see a broken relationship as an opportunity to introduce error identification to lower documentation maintenance costs. This approach largely relied upon the *document structure* to make identifying those broken locations possible. The works after this additionally derive value from the document structure such that code anchors can store even more contextual meta-information automatically, as opposed to only after some information has been authored.

Chapter 7

Meta-Manager: Meta-Information for Question-Answering

This chapter is adapted from my paper:

[105] Amber Horvath, Andrew Macvean, and Brad A. Myers. “Meta-Manager: A Tool for Collecting and Exploring Meta Information about Code”. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. Association for Computing Machinery, 2024.

7.1 Overview

When developing and understanding code, as previously discussed, developers are managing many types of information that come in different forms. The previous chapters explore how this information can be externalized and made more useful through tooling approaches that leverage the context of the source code. However, this externalization requires additional cognitive and physical effort by the initial author [15, 34, 99, 124, 218], which can be a deterrent from using these types of tools. Further, these tools also only capture what is expressly written by the developer, which, while useful for capturing information that may exist only within the head of the developer, can miss other interesting or useful contextual meta-information that was created during the sensemaking [151, 152]. Nonetheless, information related to the history and implementation of code, while often not expressed by the developer [147], is often useful in answering historically “hard-to-answer” questions about code, specifically those related to rationale [132, 138].

In this chapter, we explore capturing code editing and history information traces at scale through an automatic, event-driven tooling approach. We isolated editing events that we expected would help developers in answering some of these questions, such as copy-paste events to help with understanding *where* some code originated from and web search events and additional meta-information from web pages

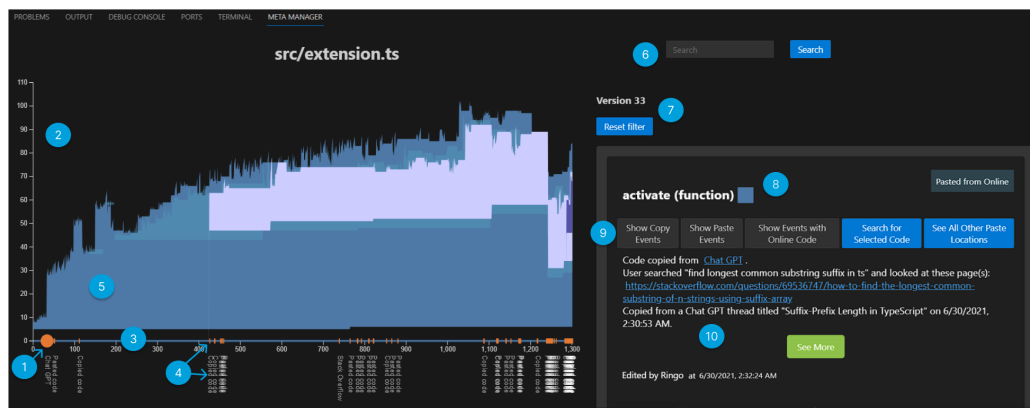


FIGURE 7.1: Meta-Manager as it appears within Visual Studio Code: the pane appears in the bottom area of the editor, with the left area displaying a visualization of the history of the code file over time, while the right area displays information about a particular code version.

including Stack Overflow and ChatGPT for answering *why* some code is written a certain way.

We instantiated this event-driven code history and provenance model in our tool, Meta-Manager. Meta-Manager translates this information into key events along a code history timeline visualization that allows developers to explore how, when, and where some code came to be. We focus on supporting *later developers in answering questions* related to code design rationale, along with code history, provenance, and relationships, given that these questions are significant blockers for developers maintaining code [132, 138, 159]. We specifically designed Meta-Manager with the goal of not only supporting developers in answering these questions but also with respect to dealing with challenges of *scale* (given the sheer amount of edits made to code that may not be of interest but are nonetheless logged [7, 228]), *navigation* (to support developers in finding the edits that are of interest), and *unwritten design rationale* (given developers' reticence to stop their work to document particular decisions [151, 159], despite their generated information trail's potential to be useful [124, 153]). In our evaluation of Meta-Manager, we found participants were able to successfully use it to answer otherwise unanswerable questions about an unfamiliar code base's history.

7.2 Overview of Meta-Manager

In order to capture code provenance, history, and rationale information at scale in a manner that is navigable and informative, we developed Meta-Manager. We begin our discussion by discussing what questions we believe Meta-Manager is well-poised to answer, then show a user scenario on how that history may be generated and used by a later developer to answer a question with Meta-Manager. We then discuss how each feature in Meta-Manager instantiates our design goals and addresses

significant questions developers have about code.

7.2.1 Developer Information Needs

In designing and creating Meta-Manager, we began by reviewing related literature on information needs of developers when working with unfamiliar code [51, 63, 78, 113, 127, 138, 159, 178, 209, 210, 226]. As discussed in Chapter 2, this can happen in many contexts, such as when adopting a code base after a coworker leaves [178, 209, 210] or when maintaining a large code base [78, 127, 138, 159, 226]. In reviewing the literature, we were particularly interested in questions about the *rationale* behind code’s design, given the lack of support for that information need [159], despite its ubiquity and importance [138, 159]. Through our literature review, we identified the following questions as related to code rationale and provenance, and as potentially answerable through directly supporting developer’s sensemaking of code history with tooling:

- **History: How has this code changed over time?** [63, 138, 159, 226] Developers often try and understand the evolution of some code in service of answering a question that is pertinent to their current task. For example, this may help while investigating when a bug was introduced [229], finding when some code was last used in service of understanding how a feature changed over time [266], getting “up-to-speed” on a new code base [113], or finding a snippet of code that was edited repeatedly to understand where the original developer had issues [138, 147, 229]. Isolating when these particular changes happened can be impossible in the case that the intermittent version is not logged in a version control system (which is often the case in situations where a developer is trying out multiple solutions), or very difficult to find even if there [119].
- **Rationale: Why was this code written this way?** [113, 127, 138, 159, 210, 226] A commonly-reported activity among developers when understanding unfamiliar code is reasoning about *why* it is written the way that it is. This information is typically only known by the original author during the time at which the code was written and, if not written down (which is the majority of cases [151, 159]), is lost. On the off chance it is recorded, it is most likely preserved in the form of a random Git commit message or code review comment [193], which are often too difficult to forage through [233]. Developers have stated that attempting to answer these questions are “exhausting” given the lack of tooling support and reticence to ask co-workers [159], yet they must be answered in order to understand design constraints and requirements which will inform later implementation decisions.
- **Relationships: What code is related to this code?** [51, 78, 138, 226] Oftentimes, when contributing a change to a code base, developers must reason about how their new code is related to many other parts of the code beyond simply what

could be found in a call graph. Other relationships that developers reason about are what parts of the code are commonly edited together (the “working set” [27, 45]), and, if introducing a change or refactoring some code, what other parts of the code must be updated. Developers also sometimes wonder what solutions a previous developer already tried when introducing a change, another otherwise untraceable relationship given that such prior solutions are usually commented-out or deleted [138].

- **Provenance: Where did this code come from?** [78, 226] In 2021, Stack Overflow reported that one out of every four users who visit a Stack Overflow question copy some code within five minutes of hitting the page, which totals over 40 million copies across over 7 million posts in the span of only two weeks [200]. Given this ubiquity of online code and developers reliance upon it, researchers have investigated the trustworthiness of code that is sourced from online resources [12], ability to be adapted to a developer’s own working context [277], and correctness of the code in terms of API usage, syntax, and so on [239]. With the rise in LLMs for code generation, research is beginning to focus on the quality of AI-generated code as well [55, 71, 123, 148, 150, 242, 267, 272]. Typically, it is not easy to see what code came from AI or from an online source, versus what was written by developers themselves. While developers occasionally add code comments that cite where some pasted code came from, this does not happen very often [11] and, when it does happen, the links have a tendency to break over time and recreating the context in which that code was initially added and determining whether it is still valid is laborious [93].

It is also worth noting that all of these questions are phrased about “this code”, which is the language used in the original research. This phrasing suggests that developers are typically discussing these questions at the snippet or block level [101], as opposed to the file level, which is what typical version control systems operate at. Our system differs by tracking code at the block level across files and supports drilling down to a specific line or lines of code interactively.

7.2.2 Scenario

Ringo, a software engineer, is working on implementing a calendar widget into his team’s scheduling software. Ringo is using an off-the-shelf React component that provides most of the calendar widget’s functionality and visuals – yet, as he is implementing some of the date verification, he notices that the returned time is incorrect. He begins by *searching Google for how the date verification API works, visits the documentation but does not find any useful code examples, then asks ChatGPT what is wrong with his usage of the API and how to get the API to verify the date correctly*. ChatGPT provides him with a code example, which Ringo *copy-pastes into the code base*. Upon re-running the code, he sees the snippet works and thinks nothing more of it. He, then, *pastes this code into the other parts of the project requiring date verification*.

Many months later, Jeremiah, a software engineer who has recently joined this project team, is working on one of his first pull requests. In doing so, he spends time familiarizing himself with the code base by reading through the code. While reading, Jeremiah notices an odd implementation choice – a particular function uses an earlier version of an API’s method for checking the time of a calendar widget, despite the current version of the calendar API being used elsewhere. Jeremiah is not initially certain whether this confusing implementation decision is intentional or not, as there is no documentation on this line of code, and, given this uncertainty, he is reticent to change the code out of fear of some undocumented design criteria. Jeremiah wonders “*why is this code written this way?*” and launches Meta-Manager to investigate.

Jeremiah notices in the Meta-Manager pane that this particular file has many hundreds of edits and, through the visualization, notices that the particular block with the confusing code was introduced many edits ago. This suggests that Jeremiah’s current teammates would most likely not know why this particular API method is used. Thus, Jeremiah begins using the Meta-Manager by selecting the line of code in question and *searches backwards in time* to see when this line of code was introduced. When the Meta-Manager timeline updates with places in which the line was edited, Jeremiah notices that the line was added with minimal subsequent edits and its first appearance corresponds to a paste from ChatGPT. This tells Jeremiah that *the code has not evolved much over time* suggesting that it was a solution that did not require much tweaking by the author. Jeremiah *inspects the code version by clicking on the “ChatGPT” paste event*. The ChatGPT code version has *additional meta-information including the original developer’s Google search and visited web pages* which shows that they were looking at the API documentation. The thread shows that they asked ChatGPT for a code example that uses the API to verify a date and ChatGPT provided the code using the earlier API method. With this additional context provided by Meta-Manager’s meta-information, Jeremiah now knows *where this odd code came from*, as the older API usage was provided by ChatGPT, and *why the code was written the way it is* – namely, to meet a specification that the newer version of the API does not provide. With this information, Jeremiah no longer needs to ask his teammates about the usage of the old API and feels comfortable leaving it as is – he adds a code comment to the line stating that this line should be updated if the calendar API updates with new date-checking functions.

Jeremiah, lastly, wants to see if there are any other parts of the code using the older version of the API, such that he can similarly mark those parts of the code for updating. In order to *find any code related to his current code*, he looks to see whether this code has been copied and pasted anywhere, and finds that the code was copied and pasted 4 times across history. When looking at those copy events, he navigates to the corresponding pastes and sees that 2 of the 4 pastes no longer exist. For the remaining pastes, he adds a code comment stating the lines should be updated.

7.2.3 Detailed Meta-Manager Design

We now discuss features (labelled with “F” below) of Meta-Manager in terms of its design goals (“D”), and how these features support answering the history and rationale questions about code we have identified in our prior literature review (Section 7.2.1).

[D1] Automatic Code History and Provenance Data

We developed Meta-Manager to support better navigation and sensemaking of code history through a scalable and visualized history view (see Figure 7.1). Meta-Manager supports automated history and provenance data through its organization of data and its history model ([F1]), along with extending its historical data capturing outside of the IDE ([F2]).

[F1] Data Organization. On system launch, Meta-Manager creates an index of the entire code project by traversing through each file and creating an abstract syntax tree (AST) representation of each TypeScript or JavaScript file and, if Meta-Manager has been used with the code project previously, searching the Meta-Manager database to find what code blocks in the current project correspond to the code blocks saved in the database history. In the case that the block does not exist in the database, Meta-Manager will begin tracking its history¹.

We chose to track code history at the block-level, as opposed to the file level, in order to better align with developers’ mental models of code [100] and given that our supported questions are often asked at the block or snippet level. This approach also complements our design goals of combating scale, considering each code block is in charge of its own history, meaning code versions are only captured when a block has changed. By deconstructing the versioning space to each code block and allowing each code block to manage its own history, we can support more fine-grained answering of questions related to *history* and *provenance*.

In order for Meta-Manager to begin logging code versions, the user does not need to take any actions beyond installing the extensions. On each file save, Meta-Manager will log a new version and perform an audit of the file to see if there are any new blocks of code to track. To investigate the code history, the developer can navigate to the “Meta Manager” tab in the bottom area of the editor — doing so will render the edit history of the user’s currently-open file. Whenever the user opens a file, the Meta-Manager will render that particular file’s history. Each code block’s history appears both within the visualization as a colored stream (Figure 7.1-5) and, given the location of the scrubber (Figure 7.1-1) along the timeline, a code box version (Figure 7.1-8) is shown that represents that particular code block at that point in time.

¹For a more comprehensive discussion on how the Meta-Manager node tracking works, see Appendix A.

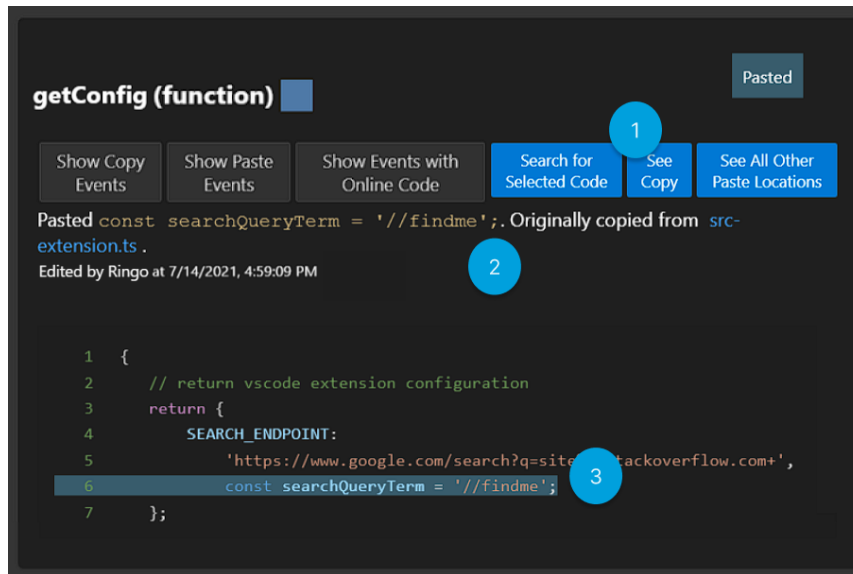


FIGURE 7.2: How the code box looks when expanded to show a code version – in this case, a “Paste” event version. (1) shows the buttons specific to a “Paste” code version, including the “See Copy” button which will navigate the user to the corresponding copy event on the timeline (if the copy happened in a different file, then the code box will update with a preview of how the code in the other file looked at the time of the copy, which can be clicked on to change to that file); (2) shows the text explaining what happened with this particular paste event — clicking in this area will open the editor tab showing what the code file looks like now; (3) shows the code for this version, along with a light blue highlight on the code that was pasted.

[F2] Development Traces Online and in-IDE. Meta-Manager tracks code-related development events within the IDE and online. For certain events, additional meta-information will be shown on those particular code versions with additional affordances. For example, in Figure 7.2, this particular version of the method `getConfig` had a paste event, where the user pasted in the code on line 6. The version adds additional information such as where that copy came from (in this case, the file “src-extension.ts”) and buttons for relevant actions, such as seeing the original copied code (Figure 7.2-1).

In cases where code was pasted from an online source, Meta-Manager will provide additional meta-information about the web page that the code was pasted from, and, if available, what the original user was attempting to do. Meta-Manager’s supplementary browser extension is designed to work with some popular programming learning resources, including Stack Overflow, GitHub, and ChatGPT². If the browser extension detects that the user is on one of these web pages, it will extract website-specific information (e.g., the name of a ChatGPT thread) and listen for copy events. If the Visual Studio Code extension detects a paste which matches the content of the browser extension’s copy, this additional information will be transmitted to the Visual Studio Code extension to be associated with that paste. The hypothesis is that the query text can be a good signal of the developer’s original intent for the code, which has been supported by prior work [124, 153] and our observations.

²We envision this list being substantially expanded to include other commonly-used resources where code is copied from, such as the official documentation for languages and APIs.

Similarly, if the user makes a programming-related Google search³ prior to visiting these websites, their initial query and visited web pages will be included with the meta-information about the pasted code (Figure 7.1-10). Clicking the “See More” button will pull up a preview of the web page in the Meta-Manager pane of the editor, and highlight the part of the code on the web page from where it was copied.

Through automatically capturing this development context that would be too laborious to capture manually, we hypothesize that these pieces of information, when combined and contextualized to when the edit happened, can help developers reason about the *rationale* behind a change and the relevant *provenance*. These features work in conjunction with our data model, which allows each block to track this information. A problem with other methods for keeping track of provenance information, such as code comments that contain links to where some code came from, is that the information can go out-of-date, either in the case the link breaks or the code changes enough such that the code comment is no longer accurate [93]. By having this information versioned, we give the developer the tools to reason about this rationale across time.

[D2] Scalability.

Given the sheer amount of information we are tracking with Meta-Manager, Meta-Manager is designed to support managing large amounts of information. We do this in multiple ways – both *collapsing* information into a visual representation ([F3], [F5]) and *prioritizing* different types of information ([F4]).

[F3] Visualization. The chosen visualization, linked to our data model ([F1]), allows each block to *manage* and *display* its history effectively. The x-axis represents edits, while the y-axis corresponds to file line count, collapsing all edits to illustrate block changes over time. For example, in Figure 7.1, the dark blue stream represents the `activate` function. In the case of nested blocks (e.g., a method within a class), the colors in the visualization will overlap, such as the violet area on the chart covering the dark blue. At the scrubber’s version, the `activate` function grows by approximately 20 lines, reflecting a paste event from “ChatGPT”, suggesting to a user that ChatGPT provided a significant contribution at this time. In this way, the visualization itself can serve to answer some questions about the code’s *history* on its own. The visualization also contextualizes the annotated timeline of events (see Section 7.2.3-[F7]).

[F4] Significant Edit Events. Meta-Manager manages scale by prioritizing certain versions over others. Each code block listens for specific edit events that occur during its history, such that these events may be annotated along the timeline (Section 7.2.3-[F7]). Edit events of interest include copy-paste events, both from online

³We consider a programming-related Google search to be one in which popular programming-related websites appear in the search result list. We acknowledge the potential privacy problems with this feature, and consider the current prototype to mainly be an evaluation of the advantages of doing this, and expect that a more complete tool can provide more control over what is saved from the browser, as in [153].

and from within the IDE, block commenting code, and, given a specific code snippet, when that snippet was edited, added or deleted. When these particular edit types happen, additional meta-information will be captured and shown on the code version, as is the case for the version in Figure 7.1-10 which shows where the code came from, what the user was doing online, who performed the edit, and when it occurred. This meta-information will change given the type of edit (see Figure 7.2 for an example of an in-IDE paste event).

We hypothesize that these edit events will be useful to later developers due to the diverse meta-information they generate, aligning with our earlier discussions on developer information needs. As discussed in Section 7.2.3-[F2], web activities of developers can elucidate code design *rationale* when viewed alongside code versioning. Within the IDE, copy-pasting aids in understanding hidden code *relationships* between the original and pasted sections, assisting in tracking code *provenance*. Block commenting reflects developers exploring different solutions or altering implementation, a code *relationship* typically challenging to trace.

[F5] Zoom and Filter. Another feature Meta-Manager provides to manage scale is through directly interacting with the visualization to reduce the history-space through zooming. Since the number of code versions will increase over time, Meta-Manager allows developers to zoom in to parts of the visualization that they find particularly interesting. The visualization will update to show a slice of the editing history (Figure 7.3), which can be dismissed with a “Reset” button. Users can also *filter* the timeline representation to only show specific edit events in order to further reduce the search space.

[D3] Support Navigation.

In order for the code history to actually be useful for question-answering, developers must be able to *find* the relevant information pieces in service of their questions. Meta-Manager presents this information as code versions that contain meta-information and supports finding these versions through multiple ways.

[F6] Search. Meta-Manager supports searching by both content and by code versions across time. Users can search across time using either code that they have selected in their current code version (Figure 7.2-1, “Search for Selected Code”) or directly through the code editor by selecting some code in their file, then using the context menu to select “Meta Manager: Search for Code Across Time”. These two searches differ slightly from one another, in that the search using the code box will search *forwards* in time from the specific code version, while the search from the code editor will search *backwards* in time (since the editor always shows the current version of the code). Both searches utilize the edit history by modifying the query given how the code changes across each version. This means that the search will attempt to expand if the selection grows, shrink if the selection shrinks, and update the code query content to match on given variable names and other constructs changing over time.

When a search is performed, the timeline will update with events marked “Search Result” for events affecting the specified code, where the code differs in some way from the previous version. This is to prevent the search results from being flooded with events where the code is exactly the same, but has moved as a result of other code above it being edited. When looking at a search result code version, the part of the code that matched the user’s query will be highlighted in orange. The search will also detect significant edits made to the code. This includes when the searched-upon code is initially added, removed, commented out, or commented back in. These events are specifically marked on the timeline with a label corresponding to the type of edit. Searching by content works similarly in that the user can type a query into the search box (Figure 7.2-6) and each code version which includes the searched-upon string will be annotated on the timeline. Searching is fundamental for finding a version that may answer questions of *rationale*, *provenance*, or *history*.

[F7] Annotated Timeline. Meta-Manager leverages the listened-for significant editing events of interest (Section 7.2.3-[F4]) by annotating these events along the visualization’s timeline (see Figure 7.1-4). Clicking on these annotations will navigate the user to that particular code version, further reducing the amount of code versions a user needs to look at in order to find potentially useful information, given our hypothesized information needs that will be met with the meta-information captured during these editing events. The timeline will also be annotated with versions to look at when a user performs a search (Section 7.2.3-[F6]). A large barrier to making sense of code history is the challenge of searching through large histories [233]. Meta-Manager attempts to mitigate this barrier through pulling out the most interesting versions using both its data and history model and through leveraging the user’s interest given a search query.

[F8] Scrubbing. Within Meta-Manager, users can scrub through code versions (Figure 7.1-1). The scrubbing functionality serves multiple purposes: enabling movement between un-annotated versions along the timeline and providing a quick overview of code changes over time. When the code box is expanded and the user is scrubbing, the code will update for each version. This view complements the visualization’s high-level representation of history with its lower-level code history representation and supports varied speeds of historical sensemaking, akin to how a user can scrub through, e.g., a YouTube video and speed up or slow down for targeted viewing. Users can comprehend the code *history* at different levels, aligning with where they are in their sensemaking journey.

We hypothesize that supporting search both by content and across time will help with further bridging the connection between the user’s current working context and the *history* of the code. By supporting this more micro-level investigation, in conjunction with the more macro-level scrubbing and visualization mechanisms for understanding history, users of Meta-Manager can answer their questions at varying levels of granularity.




FIGURE 7.3: A zoomed-in portion of the timeline shown in Figure 7.1. This zoomed-in portion shows around 120 edits between Version 710 and Version 830, with the scrubber set around Version 740, when a user pasted code from Stack Overflow.

7.2.4 Implementation

The Meta-Manager, both the Visual Studio Code editor extension and supplementary browser extension, utilize TypeScript for the logic and React [69] (with D3.js [181] for the chart in the Visual Studio Code extension) for the front end. Firestore [61] is used for authenticating the user, establishing a shared connection between the browser extension and Visual Studio Code extension, and logging the code revisions and metadata in the Meta-Manager database.

The code logging in the editor works by utilizing the TypeScript abstract syntax tree (AST) on system launch to parse each file in the user’s currently-open repository. The system then attempts to match each block within the parsed-AST to known code entities stored in the database. This matching is similar to the re-anchoring algorithm discussed in Chapter 5 and Chapter 6, while leveraging new information afforded to us by the logging of the code to the database. The algorithm uses a variety of heuristics including text-matching using the “bag of words” approach discussed in [261], the last-logged structure of the AST (such that known relationships between blocks are prioritized), the user’s current Git commit and the logged versions’ commits, and the line difference between blocks. In the case that a block in the history is not found in the current version, we consider it deleted, and, to cut down on superfluous versions, do not pull in its history⁴. If an unknown node is found in the current code AST, we assume it is newly created and begin tracking its history. We additionally listen for any document change events to capture whether or not a new block is added or removed to begin tracking or stop tracking the node’s history. Each node within the system subsequently manages its own version history and the events that happen to it, such that users can drill down even further to see, e.g.,

⁴Note that, given the nested structure of the AST, if a missing node’s parent is present in the user’s code, the user can still see the missing code in the parent’s version history and query upon it, if they want.

when a particular code block was introduced or removed. While the user is editing their code, each node keeps what we call a “change buffer” of edits that happen to the node, with each edit parsed to determine whether it is an edit of interest, such as a paste event or a block code comment. In the case of a copy, the system identifies which node experienced the copy such that a connection between the respective versions can be made between that node and the node that receives the paste event, even though the copy event itself doesn’t actually change the code.

For a more comprehensive discussion of the Meta-Manager architecture and data model, please see Appendix A.

7.3 Lab Study

In order to assess how well Meta-Manager performs in helping developers answer historically “hard-to-answer” questions about code history, we ran a small user study. Participants were tasked with using Meta-Manager to explore an unfamiliar code base while using the system to answer questions we asked them about the history of the code, without modifying or running the code. We chose to have a single condition (as opposed to a between or within subjects study design) in which participants used the tool since the questions we asked participants would, without the tool, be unanswerable, meaning there is no real control condition we could grade the experimental condition against. This was done deliberately considering we specifically designed our tool to support answering these types of questions. Thus, ensuring that the tool succeeded in that regard was our primary goal of the study, along with assessing the usability and utility of the tool.

The lab study consisted of a tutorial with Meta-Manager in which the experimenter and participant walked through each feature. Then, the participant and experimenter walked through different parts of the code base and the participant would use Meta-Manager to try and answer each of 8 questions (Table 7.1). Once the participant answered each question, the study ended with a survey to capture participant demographic information, along with their experience using Meta-Manager, and their own history in attempting to answer the types of questions Meta-Manager is designed to help with answering. For the full set of study materials, see Appendix F.

7.3.1 Method

Code History Creation

Given that Meta-Manager has not existed long enough to naturally accrue a history log that would be in line with real, prolonged use of the tool, a code history was artificially created. We did this because we did not want to bias the study in favor of the tool purely because there are a small amount of code versions, thus finding an answer to a question is trivial. The artificial code base is based upon a real code base

[234] for a Visual Studio Code extension created by an external group unaffiliated with Meta-Manager, which functions similarly to CoPilot. This repository was chosen due to the fact that much of the code centers around the Visual Studio Code API, which few developers are familiar with, thus lowering the likelihood of a participant performing well purely due to having more background knowledge in the domain.

Our methodology draws from prior approaches that similarly explored developer sensemaking of code history by using a variety of online sources along with the experimenter's rewriting of the code to create the synthetic code base [120]. Code sourced from different online sources ensures that the code base is unbiased since it does not just use code sourced from one individual who has one implementation style. To create the artificial edits, the first author independently rewrote the code base, following along with the Git commit history in order to capture "real" versions of the code. While writing the code present in each commit, the tool was logging these real versions, but was also recording individual edits (e.g., add 1 line that says `const searchResults = match(searchResults);` in file `search.ts` on commit 4acb) that were then artificially inserted at realistic intervals across each code's history, given the correct file, time period, and node. The first author intentionally did not write "perfect" code that matched what was in each commit, to account for the more realistic intermittent versions the tool would capture in real usage. The author also intentionally added events that we are particularly interested in investigating, such as copy-pastes, across each file's history, along with simulated copy-paste events that match the frequency reported in prior literature on how often developers copy-paste during a normal programming session [109]. We also added realistic copy and pastes from Stack Overflow and a few from ChatGPT (even though the code was actually written before ChatGPT was available) since these will be increasingly important, with these events occurring less frequently than within-editor copy-pastes. To further validate the realism of the code, we followed the same approach as [107] and asked participants how similar the code was to code they had seen in their own work, with participants reporting the code is, on average, similar to code they have encountered before⁵. Using this technique, we generated a code base consisting of 5,661 edits in 1,328 lines of code across 10 files and 28 different code blocks.

Tutorial

The study session began with the experimenter showing the participant how to use each feature in Meta-Manager. This included an explanation of the visualization (including how to zoom in to the visualization), how to use the scrubber to move through the code versions, how to search from both within the code editor and within a code box, how to filter to view only copy events, paste events, or paste events from online, and how to view each corresponding copy and paste between code versions. This tutorial was done in one of the files within the created code base,

⁵average = 3.8 out of 5, using a 1-to-5 point Likert scale from very dissimilar to very similar

such that the participant could see and understand the context of the code base, but none of the code history task questions related to anything in that particular file.

Task

Our main task draws from similar related work [58, 120] in that each participant was required to use Meta-Manager to answer 8 questions. Each question was designed such that it would represent at least one information need we are interested in (see Section 7.2.1) and would require the participant to use some feature of Meta-Manager to answer. Questions also required the participant to perform multiple steps using the tool, such that they would be non-trivial to answer and would represent the more realistic case of using a tool like Meta-Manager, where the full “answer” is multi-faceted and comprised of multiple information pieces. For example, question Q1 asks both what string a regex is matching on and why – “what” refers to the implementation of the regex and is requisite knowledge in order to make a change to the code, while “why” represents the rationale behind the current design and is information that can be used to reason about how a new version should be designed in order to adhere to the original design constraints, goals, and specifications. Table 7.1 lists each question, along with the steps a participant could do in Meta-Manager to answer the question. The solution in the table represents the most efficient way to answer a question, but each question can be answered using other methods. For example, for question Q6, a participant can choose to search forwards in time on the AI generated code, which will pull out instances where that code changed, but they can also scrub through the code versions using the scrubber up to the present to visually see how the code is edited over time. Participants had 10 minutes per question and were not allowed to edit or run the code, or search for information online. When a participant felt they had come to an answer, they were instructed to state their answer and they would move on to the next question.

Questions 1 and 2 were in a file with 90 versions, Question 3 and 4 were in a file with 619 versions, Question 5 was in a file with 727 versions, and questions 6 through 8 were in a file with 1,302 versions.

Analysis

For each participant, we recorded whether or not they got the correct answer for each question and how long it took them to come to the answer. “Correctness” was determined objectively by whether or not they found the correct code or code version that contained the answer and whether the participant’s summation of what they learned was accurate. If a participant got only part of a question right, such as understanding in Q1 *what* the regex is matching on but not understanding *why*, the question was still marked as incorrect. If the participant did not finish within 10 minutes, the question was marked as incorrect. We additionally reviewed the video

Question in Task	Info. Need (Sec. 7.2.1)	Solution	Outcome	Avg. Time Spent If Correct
Q1. In <code>config.ts</code> , there is a regex for search pattern matching. Can you tell me what it is matching on and why?	Rationale	Find paste from ChatGPT, read user's ChatGPT query	6 correct, 1 incorrect	3:28 (min: 0:58, max: 6:32, std. dev.: 1:59)
Q2. There is a bug in the commented out <code>Promise</code> . Can you find where the bug was and what happened?	History	Find where <code>Promise</code> was initially commented out, where <code>Promise</code> came from, and look at versions before that event	4 correct, 3 out-of-time	5:04 (min: 1:58, max: 7:02, std. dev.: 2:24)
Q3. Prior to using <code>parseHTML</code> , the author was using a different API - what was it and why did they stop using it?	Rationale	Search to when <code>parseHTML</code> no longer exists, see what code was there before, and see Stack Overflow post	7 correct	4:35 (min: 2:23, max: 7:38, std. dev.: 1:46)
Q4. Recently, some code was added to <code>search</code> that came from a different file - can you find that code and explain what changed?	Provenance	Filter to see pasted code, find paste event with code copied from a different file, then search for that code in the file	7 correct	4:47 (min: 2:00, max: 8:08, std. dev.: 2:29)
Q5. Look at lines 68 to 70 - there is a commented out <code>forEach</code> loop. Can you find the last time it was used and explain why it was removed and what it was replaced with?	Relationships	Search on commented out code, click on "Commented Out" event, find Stack Overflow post near event with replacement code	7 correct	4:24 (min: 1:57, max: 7:12, std. dev.: 1:44)
Q6. What code was generated by an AI system and what ended up happening to it?	History	Filter to see ChatGPT code, then search forwards in time on that code	5 correct, 2 out-of-time	6:36 (min: 2:02, max: 9:15, std. dev.: 1:53)
Q7. What were all the different things that the programmer tried when setting the <code>match</code> variable?	History	Search backwards in time on <code>match</code> , look at events	5 correct, 2 incorrect	5:18 (min: 2:15, max: 8:17, std. dev.: 2:15)
Q8. Some code from <code>activate</code> was moved into a different file. When did this happen and what was the code that was moved?	Provenance	Filter to code copied in <code>activate</code> , then see corresponding paste locations	7 correct	3:42 (min: 1:14, max: 8:35, std. dev.: 2:17)

TABLE 7.1: Each question that was asked during the task, along with what information need from prior literature it corresponds to, the steps that could be taken in Meta-Manager to answer the question, and how participants performed on the question in terms of correctness and time spent (in minutes). Note that some questions represent more than one information need, such as Q5, which both asks what code is related to the commented out loop, but also why the loop was commented out, which is a rationale question.

recordings to see what features of the tool and strategies participants used when coming to an answer.

7.3.2 Participants

We recruited 7 participants (6 men, 1 woman) using study recruitment channels at our institution, along with advertisements on our social networks. All of the participants were required to have some amount of experience using TypeScript and be familiar with Visual Studio Code. Participant occupations included 4 professional software engineers, 2 researchers, and a financial operations engineer with a computer science background. The average amount of years of professional software engineering was 3.16, self-reported competency with JavaScript was 4.5 (out of 7, where 7 is expert), and an average self-reported competency score of 3 for TypeScript. All study sessions were completed and recorded using Zoom and participants used Zoom to take remote control of the experimenter's computer in order to use the tool. Participants were compensated \$25 for completing the study and the study was approved by our institution's Institutional Review Board. Participants 1 through 7 are hereafter referred to as P1 through P7.

7.3.3 Quantitative Results

Participants, on average, were able to correctly answer their questions 85.7% of the time (48 out of 56), and averaged 4 minutes and 52 seconds per question. No participant got every answer correct, and all participants got at least 6 answers correct. Of the 8 failed questions, 5 occurred because the participant ran out of time, and 3 occurred because the participant came to the wrong answer.

Table 7.1 shows question outcome and how long, on average, getting the correct answer to the question took. Questions 3, 4, 5, and 8 were answered correctly by all participants and did not take relatively long to solve. Participants also solved these questions in the most consistent manner, with all participants starting with the same first step that was outlined in Table 7.1 as the intended solution path. Notably, these questions correspond to 3 of the 4 types of hard-to-answer questions discussed in Section 7.2.1, suggesting that the tool was successful in supporting rationale, provenance, and relationship needs. Participant's success with answering provenance questions supports our hypothesis that copy-paste data can help with reasoning about where and how some code came to be. Additionally, in our post-task survey, participants rated Q5 as the most similar to frequently asked questions they have, suggesting that our tool's ability to support finding relationship and rationale information is particularly valuable.

In the post-task survey, participants reacted favorably to Meta-Manager. Participants agreed that they would find Meta-Manager useful for their daily work (avg. 6.14 out of 7, with 7 being "strongly agree") and enjoyed the features provided by Meta-Manager (avg. 6.57 out of 7). Participants particularly liked the ability to see where code from online came from within the context of the IDE as a way to see what the original developer was doing, with one participant stating that they imagined that this will be how they spend "most of their development time in the future, with more code coming from AI" (P7). This, along with participants overall success on Q1, supports our hypothesis that reasoning about rationale can be supported using information traces from AI code-generation tools and related web activity. Participants desired improvements to the user interface to make some of their interactions a bit more clear (avg., 4.3 out of 7), especially with respect to the interaction between the filter, search, and zoom operations. Participants sometimes lost track of what filters they had in place when searching, and vice versa, which caused confusion. Adopting the UI for sorting and filtering from popular shopping websites might help alleviate these problems.

We additionally asked participants to rate each question asked in the study by how often they have encountered similar questions in their own programming experiences on a 5-point scale from "never ask" to "always ask" (Figure 7.4). Participants reported asking questions similar to Q5, which asked about why some code was introduced to replace some other code, most often, with 4 participants stating they "always ask" questions like this. Notably, that is also one of the questions all participants were successfully able to answer, which suggests the tool is useful in




FIGURE 7.4: Each question scored by participants in terms of how often they encounter similar questions in their own programming experiences.

answering this type of common, hard-to-answer question. Only two questions had some participants state they never asked that question, which were the questions corresponding to reasoning about where some code originated from (an AI system, in this case) and what the previous developer had tried when implementing some change. All questions had at least one participant say that they sometimes ask that question, which is both inline with prior research and contributes additional evidence that tooling to answer these questions is valuable.

7.3.4 Qualitative Results

We now explore participants' qualitative experiences using Meta-Manager in terms of how they used its features to answer each question with respect to Meta-Manager's design goals.

[D1] Automatic Code History and Provenance Data

Participants, overall, enjoyed having access to the code-related history and provenance data, especially in the case of code sourced from online. 6 out of 7 participants explicitly stated in the post-task survey that they valued Meta-Manager's ability to capture what code was sourced from online sources, especially ChatGPT, and that these events were explicitly called out on the timeline and filterable. This preference also manifested in their question-answering strategies with participants commonly defaulting to clicking on any event annotation that came from online, especially if they were stumped on what to do to answer a question. P4 clearly articulated this strategy by saying, after using ChatGPT to solve Q1 and why a regex was written this way, "I'm looking at ChatGPT because that worked well last time." Other participants did not immediately understand that the web-based pastes contained additional meta-information that could help with reasoning about "why" some code is the way it is – P3, in attempting to answer Q3, did not look at the Stack Overflow code version which has a Google query explaining why the user switched API methods, and, instead, brute-force searched through the surrounding code versions

and correctly reasoned that the API methods were swapped due to an asynchronous issue given some type changes made between versions. While this strategy was successful in this case, their usage suggests that some users may not see the connection between web-activity and rationale for changes, suggesting that further highlighting the most pertinent “information cues” from these versions (e.g., Stack Overflow question titles) in the user interface, either through the timeline annotation text or within the version itself, may better serve to highlight the significance of the web activity.

[D2] Make Information Scalable

In terms of managing the sheer scale of the version space participants were operating in, the combination of the visualization, zooming, and filtering worked together well to isolate “sub-histories” of the history to explore. A common strategy in answering history and provenance-based questions, used by 4 participants 12 times, was to use the annotated timeline labels as a boundary for a search space, then “zoom” into this space to look at the intermediate versions. For example, P3, in order to answer Q2, used the visualization to identify that there was a large growth in the code base at the end of the history and there was pasted code added at that time – he then zoomed into the end part of the history at the first instance of pasted code when the lines of code grew in order to reason about how the code changed after the addition of it and prior to it being commented out. In this way, participants were able to leverage the significant editing events, not only for meta-information, but also for their ability to segment the information space. This behavior of orienteering [10] to gain an understanding of part of the information landscape is consistent with behaviors exhibited in other information foraging studies [233], suggesting Meta-Manager’s feature set supports these processes when navigating a large information-space.

[D3] Support Navigation

In our design of Meta-Manager, we were particularly concerned with making the code history space navigable, given this significant challenge in prior work [120]. To this end, we adapted different techniques for moving through the history including annotated timeline labels, scrubbing, and search. However, one interesting aspect of navigation that we did not explore as much, nor has been explored in related literature to the best of our knowledge, is how navigation worked with respect to moving between the “live” version of the code within the IDE and the historical versions housed with Meta-Manager. Through supporting this relationship, we found multiple design challenges and opportunities.

Navigating Through Time. All participants began each question that had an optimal first step of searching by “searching” – the ubiquity of search made it a common strategy. However, one challenge participants faced when searching through

history was going *too far back* in the history and missing the connection of what they were seeing in the prior version versus what was in the IDE. This happened with 2 participants across 3 questions – the participants would search on the current version of the code and then began clicking through the search results starting from the earliest version. Since our algorithm works across time, it begins at the current code and works backwards by adapting its query given identified changes between versions – since participants could not readily see how the query evolved, jumping to the beginning of the search results in the history (which is the *last* match the algorithm found) was sometimes confusing. Evolving the search query is necessary in order to ensure trivial changes are not disregarded as search results (e.g., switching `const match = 'foo'` to `const match = 'Foo'` where the “f” is now capitalized), but Meta-Manager may be improved by supporting more sophisticated ways of summarizing the search over time or refining which matches should be included. This optimization would also help with another issue participants encountered, where the search would perform differently depending upon what code was selected in the IDE – given a question such as Q5 where participants would begin by searching on a commented-out `forEach` loop, some participants would select the whole loop while others would just select the first line, which would result in the search performing differently given these different yet semantically-similar initial strings.

Communicating “Now” versus “Then”. Another challenge of supporting navigation between the IDE and the code history was how to represent the time between the last-logged version and the current version within the IDE. 3 participants did not immediately understand that the end of the code history timeline represented the last-logged version of the code, and not a live view of what was in the IDE. This may be a product of the lab experiment, which did not have participants editing the code which would have let them see how the timeline updates given their edits. Nonetheless, we found that the collapsing of edits by disregarding the amount of time between edits, consistent with prior work [100], did not confuse participants – however, not showing the lapsed time or un-logged edits between the last version and the “live” version within the IDE was confusing, suggesting future work in better showing how the user’s current in-IDE version of the code correlates to the larger history.

Navigating Between and Across Files, Spatially and Temporally. Questions that required participants to reason not only about the history of their current file, but how that history relates to the history of *other* files caused confusion. Q2 required participants to reason about how some code in the current file changed, given its relationship to its original copy source in another file. Understanding the original code’s intent was necessary in order to better reason about why the code from the question was commented out – participants, with the “See Corresponding Copy” button, can see a preview of what the copied code looked like at the time of the paste. 2 successful participants and all unsuccessful participants struggled to reason

about the connection between the “Corresponding Copy” version (which is on a different version within in a different file), the version of the code that received the paste, and how both of these information pieces related to the code in their current IDE. Future systems may investigate how to better support this reasoning across both time and space through supporting more interactive mechanisms for managing versions, which has shown success in other contexts [206].

7.4 Discussion

We now discuss how Meta-Manager is situated in the larger context of making sense of code and its history, and the role meta-information can play in that process. Prior work has investigated how developers make sense of many variants of the same code and its output [233] and the challenges in doing so – the authors note that this foraging process involves managing similar yet disconnected information patches. We showcase Meta-Manager as an improvement upon that model through extracting and utilizing *meta-information* to serve as strong informational cues to both reduce the number of candidate patches to traverse through and to connect the history into a larger, contextualized narrative. P2 noted that they were essentially “recreating the story” of the code when clicking from event label to event label to answer a question about design rationale – notably, [233] also discusses this phenomenon of information foraging being construed by users as assembling a “story”, suggesting that our event labels may serve as one way of structuring these “stories”.

When considering code history as a story, this is not dissimilar to the concept of “literate programming” and its philosophies, originally proposed by Donald Knuth [125]. Knuth believed that code should be more naturalistic, written as an expression of an author’s reasoning behind solving a computational problem. Programming, in its current state, typically relies on documentation as a way of translating between the lower-level code and its higher-level semantic meaning, with this documentation often spread across various platforms and represented using different modalities such as inline code comments, Git commit messages, GitHub pull requests and issues, and formal design documentation. With the rise of LLMs for code generation, there is a new platform and modality for these natural language descriptions of code, which our work has shown are worth capturing as they can be used for reasoning about design rationale. The not-so-distant future of software engineering may consist primarily of this prompting for generating and modifying code – a future in which whole programs may be constructed predominantly through prompts that can be translated into code narratives not unlike the literate programs Knuth described. In this way, the code serves to describe the lower-level implementation but the higher-level goals and reasoning are communicated through the prompts. Meta-Manager begins to probe at how these forms of code-related meta-information may be captured and presented to help construct these narratives.

To the best of our knowledge, our work with Meta-Manager and its study are the first pieces of research to investigate the provenance of AI-generated code. Prior research has cited the importance of this research thread in order to answer questions such as “does AI-generated code leads to fewer (or more?) build breaks”, “what prompts were used to create this code”, and if AI-generated code should be under more or less scrutiny during code reviews [24, 231]. These questions, in theory, could be investigated using Meta-Manager through following the development of AI-generated code throughout its life-cycle. Meta-Manager also demonstrates that, through capturing AI code-generation provenance information, other questions and activities can be supported, such as reasoning about code design rationale. Considering more software will likely have source code written by LLMs, tooling support for maintaining and comprehending this code, along with reviewing it for its applicability and correctness, are likely to become more important.

In summarizing our findings and their implications, we find support for the claims that:

- code history data, when properly versioned, contextualized with meta-information, scaled, visualized, and prioritized to support easier navigation, can be used by developers to reason not only about what, how, and when some change happened, but also *why*;
- capturing information traces during the AI code-generation process can be used to support this reasoning; and
- more generally, some information produced as a by-product of authoring code can be mapped to later developers’ information needs – thus supporting the research framework (Figure 1.2) proposed in Chapter 1.

Previous systems have captured some of this meta-information, such as Mylyn [78], and typically used this information to support code *authoring* tasks, such as localizing code patches to change, but were less concerned with questions of code comprehension by later developers. Other systems, such as the wide array of code visualization systems (e.g., [261, 270]), usually did not imbue additional information about the context in which the code was created. We show, with Meta-Manager, how these two approaches can be complementary to one another in supporting developer sensemaking tasks.

7.5 Limitations and Threats to Validity

Our study is limited by the fact that we did not have a control condition to compare our results to. While the questions that we asked may be impossible to answer without a tool like Meta-Manager, without a control condition, it is difficult to make any definitive claims about whether or not a participant’s ability to answer these questions would result in some measurable difference in terms of code comprehension.

Given the amount of prior work claiming that these are important questions and no question received a “never asked” in the post-task survey, we have evidence to suggest that supporting these questions is useful.

Our study is also limited by the fact that we used an artificial code base, as opposed to a code base generated with prolonged usage of the tool. We feel that choosing to have an artificial code base such that we can simulate the real experience of having many code versions to navigate through allows us to better ensure that Meta-Manager is scalable to support development on a real code project. We attempted to mitigate the potential biases introduced through artificially creating the code base by ensuring that our code changes were produced in a way consistent with real world code editing practices, diversifying where our code came from, and asking our participants whether the code from the study was consistent with code they have seen in their own work. Future work would benefit from assessing how the code versioning works for information seeking by many developers working on the same code base over time.

An additional limitation is that the questions asked during the study were made up by the first author. While each question was derived from previously-reported information needs developers have about unfamiliar code and participants rated each question as a question they at least sometimes ask, additional studies may investigate how developers can use Meta-Manager to answer their own questions about their code in order to better understand how the system supports answering real developers’ questions that were not asked in this study. Further, the study in its current form cannot answer how often Meta-Manager would be useful, as we did not capture the full breadth of questions it can be used to answer and the frequency developers ask those types of questions. Previous work, such as [159], and our own reports from our participants suggest that these questions occur semi-frequently and are challenging to answer – nonetheless, future work would improve upon our work by investigating to what extent the breadth of developers’ information-seeking behaviors are supported with Meta-Manager, and what extensions to Meta-Manager might help it answer more questions. Our study, instead, focuses on to what extent Meta-Manager and its features do work for answering some of the questions we know from prior literature are difficult to answer.

7.6 Future Work

Our lab study provided some evidence that Meta-Manager helps developers answer hard-to-answer questions about code. However, this was in the context of a developer joining an unfamiliar code base with no real contextual knowledge of the code or its history — while this allowed us to best assess how well the system works in assisting developers in answering these questions in, arguably, the most difficult situation, future work would benefit from seeing how Meta-Manager helps developers when they are working on their *own* code. Open questions remain in this situation —

given developers' own mental models of their code base and, most likely, its history, one can imagine that usage of Meta-Manager may change, as developers' questions about the code base may become more specific since they have more information to work off of. Improvements to Meta-Manager to support more personalized information may be a richer querying system that supports project-specific terms or time constraints, such as "show me all edits to this function between bug fix #124 and now", or allowing users to define their own "events" outside of copy-paste, code commenting, etc., that the system will automatically log as an event of interest. Prior work has supported similar team and project-specific tagging of information in software projects to help with source code navigation [238, 247] – extracting project-specific tagged events as timeline events may also help developers with navigating between code versions.

Additionally, there is existing meta-information about code project history that Meta-Manager does not currently show, such as the boundaries between Git commits along the timeline, or event labels for typical historical events like merged-in pull requests. One can imagine changing the visualization to show visual boundaries between these events, which may also help to better visually convey to a user why some block of code has appeared or disappeared as code commits are merged together, pulled in, or checked out. This may also help make this information more easily navigable, considering foraging through many Git events is a known challenge [119]. We additionally want to support more event capturing, such as when documentation-type code comments are edited,⁶ and other interesting code sources, such as CoPilot. CoPilot events, in particular, may be particularly interesting given that user's edits that prompt CoPilot (especially in the case of natural language prompts to CoPilot to generate some code) may serve as strong signals for what the code author was trying to do. Conveying more of this meta-information in conjunction with more structured time information may give developers more useful navigational way points for searching for code across time. Expanding the types of meta-information Meta-Manager supports may help with answering other types of programming-related questions.

However, while expanding the types of information captured may help later developers, an important concern held by developers and researchers we showed this work to is anxiety around sharing information that may be seen as "embarrassing". Multiple researchers asked questions around whether or not there is a way to, perhaps, hide that one copied and pasted from ChatGPT or Stack Overflow since they would not want their coworkers to know they asked a question on one of these sites that they felt they should have known the answer to already. The current version of Meta-Manager was seen as a probe on whether exposing that information to support later developers' sensemaking would be useful and showed the value in

⁶In the current version of Meta-Manager, one can search on a code comment backwards through time in the same way that regular code can be searched on, but better categorizing the type of code comments, such as "TODO's" or "bugs", may make this information more comprehensible and searchable.

doing so, but one can imagine improvements to Meta-Manager that allow for more fine-grained sharing and presentation of information. Ideas include black-listing or white-listing certain websites or subject matters to share information from, choosing to only be listed as “anonymous” on other developers’ Meta-Manager timelines, and/or curating your set of queries and copy-pasted website sources prior to pushing them to the Meta-Manager database.

Currently, Meta-Manager does not support saving or sharing specific code versions, queries, or filter settings. There may be situations in which it would be useful to keep track of that information, such as for communicating with collaborators about how and when a bug was introduced [142] or for saving a code version that a developer may be considering reverting back to [119, 120, 269]. This meta-meta information (meta information about the use of the meta information about the code) could be useful to help others perform similar sensemaking to the current user, based on research [72] that multiple people through time often need to repeat previous people’s work.

7.7 Conclusion

Meta-Manager represents a promising path forwards with respect to collecting and making sense of large amounts of contextual meta-information that is otherwise too difficult to reasonably keep track of. Specifically, the system utilizes an *anchor* derived from the *document structure* which may include *provenance* data and versions based upon the *mutability* of the source across *time*, where all of the information is *authored* by the system. In this system, we additionally wanted to support navigation across many versions, so we leveraged our insights into difficult code rationale and history reasoning in order to prioritize different types of information given its *content*. We further leveraged the developer’s working *environment* through listening for known developer behaviors, including copy-pasting code. Thus, Meta-Manager encapsulates and accounts for 7 of the 8 properties we found to be particularly important in meta-information systems (see Section 1.1) – interestingly, the one property we did not specifically design for (*audience* and how meta-information should be authored and presented for a public versus private audience) was the property we received the most feedback and concern about from the wider research public, with some engineers and academics expressing concern around sharing their meta-information, including Google queries or copy-pasted code from online, if it could be considered as embarrassing. Nonetheless, Meta-Manager overall showed the promise for designing a comprehensive and scalable meta-information system, leading us to leverage its architecture and design principles for our final system, MMAI.

Chapter 8

MMAI: Accelerating Sensemaking with Logs and LLMs

8.1 Overview

In this work, we have investigated methods for improving comprehensibility of code and code-related meta-information through strengthening the relationship between these two rich information sources. We have explored the benefits of bridging code and its meta-information across different sensemaking domains including understanding API documentation and comprehending unfamiliar code, and the results have been promising with developers successfully creating and using the contextualized information to perform better on development tasks and answer relevant questions about code.

Despite this success, I recognize that we are currently in a transitory state of software development with the rise of AI tools for programming assistance. These tools are most often used for code generation, with a recent report from developers claiming that around 31% or more of their new code is generated from AI tools [148]; it is expected that this number will continue to grow as AI tools continually improve and the barriers to using them are lowered. While work with LLMs for programming-assistance mostly revolves around improving code generation and understanding capabilities, I see alternative opportunities for supporting programmers.

One cognitively intensive task that is commonly done by programmers, yet has received surprisingly little attention nor support in Software Engineering research is *print debugging*. A recent report claims that, even though sophisticated IDE-provided debuggers have gotten only more powerful, many programmers still prefer to debug using print statements [19]. Yet, their take-away is not to directly support print debugging but, instead, to make IDE-provided debuggers even better, with the implication that print debugging is sub-optimal and can be eradicated given even more tooling support for sophisticated debuggers. Instead, I argue we should adapt to the tooling developers are already choosing to use in order to make print debugging more powerful. Further, I believe that the meta-information generated during print debugging can be utilized by LLMs to accelerate developers' sensemaking.

As a case study of these hypotheses, I expanded on the Meta-Manager architecture to create MMAI. MMAI builds upon Meta-Manager in order to directly support print debugging through both handling previously-onerous information collection and tracking tasks by supporting the collection of log values through our own logging package, `mmlog`, and supporting reasoning and synthesizing of log values through connecting into the OpenAI API [199] to query on this even richer set of historical data. We chose to support the collection of log data in conjunction with Meta-Manager’s code history and provenance support in order to support more fine-grained question-asking and answering when debugging. For example, consider the query “Why is `foo` equal to `undefined` when it used to be equal to an array of integers?” – with Meta-Manager’s node-based data model and historical knowledge, the system has knowledge of both how `foo` has been used and how `foo` has previously been invoked in the past. This data is then complemented with the `console.log` values that are captured with our package, `mmlog`, such that the LLM has all of the contextual knowledge necessary in order to answer this otherwise vague question. Through our usage of GPT-4, we make asking this question directly actually possible using retrieval augmented generation (RAG), in which GPT-4 takes that query in a pre-formatted prompt and applies it to the set of historical data that we ingest into the query.

To further identify a set of questions MMAI would be well-poised to answer, we completed an exploratory interview study. We asked 10 developers about their practices and questions they have when performing print debugging and the sorts of questions they would ask if given access to an oracle that was aware of their debugging practices (a similar research approach to [121]). Through an open-coding thematic analysis, we developed 7 queries we believe MMAI can answer through direct, AI-assisted history and log support. Through this ideation and exploration, we contribute MMAI as a prototype system for answering debugging and log-related questions with AI, and a set of developer queries asked in service of supporting print statement debugging that have not been previously supported directly via tooling.

8.2 Background and Related Work

In order to disambiguate our discussion of print debugging, which utilizes log statements and log values, from other logging contexts, we will briefly discuss the terminology and definitions derived from prior work that we will be using hereafter. We also briefly discuss some of the background and related work on print debugging in service of further motivating and situating our work.

8.2.1 Terminology and Background

- **Print debugging** [171], also called `printf debugging` [19, 20, 146] or *trace debugging* [146, 207], is the act of placing lines of code that print information to

a console or terminal at run time in service of determining when, where, and how some problem in the code manifests.

- A **log statement** [81] is the line(s) of code within a developer’s source code that invokes a call to print some information to a developer’s console, terminal, or other output location. For example, a log statement in JavaScript could look like this: `console.log('myDogRingo ', myDogRingo())`, while the same example in Python would look like this: `print('myDogRingo ', myDogRingo())`.
- A **log value** is the execution result generated at run time from a *log statement* that appears in the developer’s console or other output location. Continuing from the last example in JavaScript, a log value could look like this in a web console (e.g., Google Chrome Developer Tools): `myDogRingo is very cute`, in which the string “myDogRingo ” is appended to the function call `myDogRingo()`’s return value of “is very cute”. Note that the form in which a log value appears is dependent upon the development environment and the code within the log statement itself, given that log values are an amalgamation of the contents of the log statement. This amalgamation often includes both static values such as strings and dynamic values such as variables. We further clarify that a *log value* is the direct result of a *log statement* – in contrast, we use the term *output* to refer to the actual result of either the entirety or some part of a software system, which may include a full software application, a graphical output, an algorithmic computation, or some other artifact created by the developer’s source code.

Note that the log statements and log values discussed in our work are not the same as the logs discussed in other research about *logging*. Logging is the act of generating long-term output records that characterize the behavior of an often-large enterprise system for later analysis, which contrasts from print debugging, in which log values are not meant to persist indefinitely and are, instead, an act of inquiry by a developer *in situ*. There exists a wide body of Software Engineering research around logging (e.g., a comprehensive literature review that synthesizes over 200 papers [81]), including systems for cleaning messy logs [219] or analyzing and generating workflows of logs [155], and empirical studies of considerations developers have when designing their logs [144]. While some of the problems discussed in the works around dealing with messy logs [219] and reasoning about logs [155] generalize to our work, our design choices for MMAI had to account for in-real-time source code modification, including updating log statements and values, as opposed to running purely post hoc. Further, the nature of inquiry differs between post hoc, large-scale log analysis and ad hoc print debugging – a developer performing post hoc reasoning is often wondering *when* or *how* some observable difference occurred, while a developer debugging wants to know *why* some code *currently* does not work, which may result in reasoning about code changes, but not always. MMAI supports this inquiry through automatically capturing log values that are associated with code

changes, such that the system can draw inferences when comparing the developer's current code to code that previously worked.

8.2.2 Print Debugging

The act of debugging, by definition, is the act of fixing some defect in software. This process can be broken down into steps, called "TRAFFIC" by Zeller [275]: track the problem, reproduce the failure, automate and simplify the test case, find possible infection origins, focus on the most likely origin, isolate the infection, and correct the defect. Most relevant to print debugging are the mid-to-end points of the process (i.e., find, focus, isolate, and correct), with *find* and *focus* being particularly important. Print debugging is often used for fault localization [260], in which developers place log statements at different places within the source code in order to observe whether or not some code is reached and what the log values are at different points during the execution. This foraging process, when debugging, can be time consuming [197]. Once a developer has formed a hypothesis about where a bug is originating from, log statements can similarly be used to observe the behavior of the problematic software and test potential solutions. Given this lightweight and flexible nature in terms of their implementation and the intent behind their usage, log statements remain the tool of choice when debugging [19, 194].

Despite the benefits of log statements for print debugging, there are some drawbacks to their usage. Challenges can be broken down into two predominant themes – log statement authoring and the implications therein, and log value presentation. *Authoring* log statements inherently introduces clutter to the code base with potentially erroneous information (not dissimilar to the code comment discussion from Chapter 4) – developers must be judicious in cleaning up the log statements when they are no longer needed [19, 81]. Interestingly, failing to clean up these log statements has organically led to an emergent class of software repository mining research in which these log statements are mined to understand where bugs are more or less prevalent within code bases [13, 264, 265]. Further, some of these logs come from large enterprise systems such as Google's monolithic code repository [265]. This provides evidence that the analysis of log statements with historical data is useful for reasoning about code faults, even developers in large scale and professional settings are using print debugging despite the existence of more sophisticated techniques, and the act of removing these log statements is not foolproof. Log value presentation suffers from multiple issues, including how log values are (not) ordered [141], the structure (or lack thereof) of log values [112], and the scale of log values to parse through, especially in enterprise contexts [19]. We attempt to mitigate some of these challenges through smart log value summarization, clustering, and filtering, while supporting direct inquiry and log statement and log value grouping with our tool, MMAI.

8.3 Exploratory Interview Study

In order to more comprehensively understand the types of log statement and value reasoning and management activities that MMAI might be helpful with, we completed a formative interview study. The interview results were qualitatively analyzed in order to come up with a set of queries for which we could provide direct tooling support. In summary, we came up with 7 queries that we aimed to support in MMAI.

8.3.1 Method

Interview Design

The interview was semi-structured and included questions and follow-up questions related to debugging with log statements. Participants began by describing their programming background, and then were asked to describe a recent debugging session, in order to help mentally situate them for the following debugging-related questions. For each high-level debugging-related question (e.g., “(In your debugging session), when you were running your code, how did you find the outputs you were interested in in the console?”), there was a follow-up question that asked participants to come up with queries they had when performing that task and how they would express that query to a log statement history and value oracle. In this way, we sought to tightly couple participants’ queries to the tasks that they would support when performing print debugging. Once each high-and-lower level debugging question was asked, participants were asked to share any other debugging or log-statement and value related thought they had that did not organically come up during the interview. The interviews took approximately half an hour, were approved by our institution’s Institutional Review Board, and were recorded with video conferencing software. During each interview, I took notes to supplement the recordings and transcripts. The protocol can be found in Appendix G.

Participants and Recruitment

We recruited 10 participants through our social networks and through our university’s various computer science-related Slack work spaces. Participants’ professions ranged from professional software engineer, to masters students with previous professional programming experience, to fellow PhD students who actively programmed as part of their research. On average, participants had 5 years of professional development experience, and 7.5 years of total experience. Participants’ programming languages of choice ranged from JavaScript and TypeScript to Python, C++, Scala, and Go, with the majority (7 participants) primarily using web-based languages and frameworks.

Analysis

To analyze the data, I used an open-coding, thematic analysis while comparing answers by question and follow-up question. More specifically, I investigated what themes were consistent across each question and what queries participants wanted to ask, given the higher-level task or goal that was presented in the original question. Other emergent themes, such as in what context each task was more or less relevant or what other alternative tooling choices participants used when debugging were also captured.

8.3.2 Results and Discussion

Interview participants will be referred to as P1 through P10. Throughout the results I will discuss high-level, emergent themes, along with the queries that would help in supporting the behaviors or tasks described that are associated with each theme. Each query will be labeled Q1 through Q7.

Participants discussed their log statements as *tools of inquiry*. P7 articulated this sentiment well, stating:

I'm working with a somewhat semi-transparent or opaque system, even if it's my own code, and I'm sort of like poking it, or like trying to interpret it. So I'm asking it to tell me things, and it tells me things that I look at in my console... it's kind of like a theory building exercise. - P7

Thus, the way in which participants structured and placed their log statements within their code to support inquiry varied. Common log statement practices included placing them across branching code to understand control flow and code reachability (4 participants), tracing variable values through execution (4), and checking variable values or data shape (4). While participants' log statements are often semantically related, such as a set of statements checking whether the if or else block was reached, conventional editors and consoles do not support grouping related lines of log statements or values together. Queries that could support this group-level reasoning include:

- **Q1: Group log statement A, B, and C's values**
- **Q2: What are all the log statements and values related to variable A?**

Participants expressed that *finding* places to insert these print statements was often not a challenge, given their familiarity with the code (e.g., "I usually have an intuition as to where the problem is, I'll figure out what the problem is from there." - P4). Instead, a common challenge came from *locating* the relevant log value(s) in the console or log data. Noisy logs, in particular, commonly made locating log values of interest difficult, with participants bemoaning large log value sets due to code being continually re-run in an event loop (P4 and P10) or due to multi-threaded execution

(P6 and P8), large enterprise code bases with many log statements (P2), and asynchronous code creating confusing log sequences (P4 and P7). Participants expressed a desire to filter log values to only log statements that were relevant to their current task, with some participants already using ad hoc strategies to make that filtering possible. Strategies included using a “key-value” system for “labeling” their logs (P1, P2, P3, P6, and P10) in which the “key” will include the variable name and possibly the file name, and the “value” will be the variable itself, such that they can search for that “key” in the console; another more onerous strategy was to remove irrelevant log statements through manually editing the code (P1 and P5), but this strategy broke down if the logs were in files the participant could not edit. Notably, most “relevant” log statements were only the log statements authored within the current debugging session (e.g., “If it’s been more than a little bit of time, [the] logs won’t be helpful anymore because I don’t remember them.” - P1). Thus, we found a need to support the following queries:

- **Q3: Remove all log values that are not related to variable A.**
- **Q4: Remove all log values not created during this debugging session.**

Once the relevant log values (or groups of values) have been found, participants expressed the challenges they had both in interpreting the log values and in using that information to design a solution. Three participants (P5, P7, and P8) expressed essentially the same sentiment when performing this interpretation while debugging – either the code is wrong or their understanding of the code is wrong (“Code doesn’t lie” - P8). Returning to the concept of log statements as tools of inquiry, the lens they provide into the underlying logic and form of data is relatively narrow and fragmented, leading to a large amount of cognitive stress on the participants (“There’s, I guess, a big cognitive load that I’m putting on myself to remember what order everything’s supposed to be firing in” - P4). This synthesis process of reasoning about log statements both in relation to one another and their log values across time was a challenge. Three participants chose to sometimes save log values in other places, such as note pads (P1) or in files with semantically-meaningful names (“[When doing something tricky] I actually save logs into different text files that have descriptive names, [with] one file called ‘multi-thread-on’ and one called ‘multi-thread-off’.” - P6), in order to more directly track values over time or given code versions. Participants wanted support from an oracle that could track and infer correlations between code changes and output changes (P1, P3, P4, P6, P7, and P10), both in support of finding when a bug was introduced (P1 and P3) and for reasoning about why some code is problematic to help with deepening understanding (P3, P4, P6, P7, and P9). We found support for the following queries:

- **Q5: When did this log statement last produce the correct value?**
- **Q6: Why did this log statement produce the last incorrect value?**

- **Q7: Summarize this log statements' values and find when something interesting happened.**

Other themes that emerged were not directly related to queries about log statements or log values, but included sentiments related to classic debuggers and the ephemeral nature of logging. Three participants expressed their distaste for debugging tools (*"I find debug tools very unnatural like 'you should use breakpoints' and 'step in' are not how I think and are confusing."* - P4), while one participant used these tools dependent upon their type of development (P1), and another used them almost exclusively (P8). Multiple participants, when discussing why they do not use these tools, explained that they lacked formal debugging training (P5) or found the tools too heavy-weight, either due to resource intensiveness or due to having to set up specific development configurations (P6). Interestingly, there was a pervasive sentiment of shame amongst some participants when discussing their preference for print debugging over using more sophisticated debuggers, with some claiming perhaps the preference stems from an inflated ego (*"I think to myself, 'I can get away with debugging this without these robust tools'."* - P4) or that they are just not efficient at debugging (P5 and P7). This sense of shame has not only been echoed but reinforced in other research, with a survey respondent in a related study about developer debugging practices [19] stating that people who use print debugging "never learned how to use a debugger". I see developer's reliance on print statements not as a developer-problem but as a tool problem – clearly, print statements are providing some affordances traditional debugging tools are not, so we should adapt the tools to match the developers' habits, not the other way around, which partially inspires this work.

With respect to the ephemeral nature of print debugging, the act of authoring log statements shares many similarities to the act of creating and making sense of the other forms of meta-information that have been discussed in this thesis. Participants discussed how their logs were often only useful in the moment they were authored (not dissimilar from annotations and other information scraps), and how they sometimes must be synthesized together in order to be useful (similar to a driving motivation for the curation work – see Section 5.2). More broadly, these usages and ad hoc practices imply that traditional tooling is not adequate for performing the complex sensemaking activities the participants want to perform with their print statements. Through this meta-information lens, I adopt a similar approach to my previous systems to assist in sensemaking of programming information, with the meta-information of choice this time being log statements and values.

8.4 Overview of MMAI

In order to support more sophisticated reasoning and management of log statements and values as a form of meta-information, we developed MMAI (Figure 8.1), an expansion of the Meta-Manager model. We begin our overview by discussing a

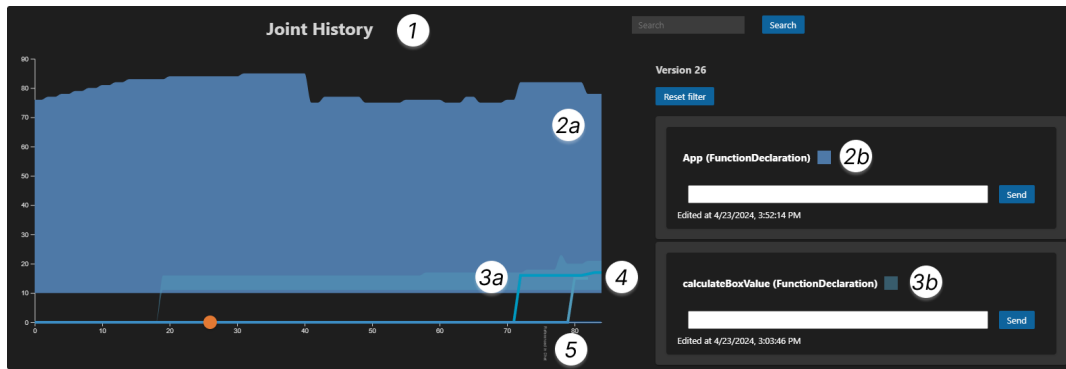


FIGURE 8.1: MMAI as it appears within Visual Studio Code: the pane appears in the bottom area of the editor, with the left area displaying a visualization of the history of user-selected code nodes that the user wants to inspect, while the right area displays information about each code node, now including additional output information and an interactive chat component. The two code version boxes for `calculateBoxValue` and `App` are collapsed (see expanded version in Figure 8.3).

debugging scenario inspired by one of our interviewee’s own experiences to show how MMAI may have been used to overcome some of the barriers they discussed. We then discuss the MMAI design in-depth, including how it leverages and builds upon the Meta-Manager design, and how MMAI directly supports the queries found in our formative study through both specific design choices in the user interface and through leveraging AI.

8.4.1 Scenario

Linda¹ is creating a web application that takes live-stream data from the popular video game streaming platform, Twitch, and renders the streaming video with additional information, including a drawn box around the main player character. She notices that the coordinates of the box are incorrect, thus the box is not rendering in the correct location. She uses MMAI to assist in this debugging episode. Linda has previously inserted MMAI’s `mmlog` statements in this render loop as this part of the code has been problematic in the past. Further, she can choose to filter out the log values they produce as part of MMAI’s log management capabilities, thus choosing to leave the `mmlog` lines in her code introduces no additional cost to her when they are not relevant.

Linda begins by finding and placing relevant `mmlog` calls given her line of inquiry. This includes `mmlog` calls in both the back-end where the computation actually occurs and in the front-end where the box is rendered, given the computed coordinates from the back-end. Normally, the structure of the code and the ordering of the log values in the output would make reasoning about these semantically-related yet disparate parts of the code difficult, but MMAI allows for users to *compose* related code histories, log statements, and log values together into a singular visualization (Figure 8.1-1) through selecting calls to `mmlog` in the editor using icons in the gutter

¹This scenario is inspired by a recent debugging episode shared by P4 – “Linda” is used as a pseudonym.

```

12 function calculateBoxValue(character: Character): number {
13   const box = TwitchOffset(character.getPosition() + 1) // .padStart(
14   const offsetBound = TwitchOffset(character.getPositionBound())
15
16   mmlog(`box: ${box}, offsetBound: ${offsetBound}`, character)
17
18
19   return box + offsetBound
20

```

```

11 function App() {
12   // eslint-disable-next-line no-lone-blocks
13   { ...
20 }
21 mmlog('calculateBoxValue in App', calculateBoxValue(character))
22 const waveAnimation = (...
64 )
65 return (
66   <div className="relative bg-slate-100">
67     <div className="mx-auto max-w-7xl">
68       <div className="relative z-10 pb-8 sm:pb-16 md:pb-20 lg
69         <main className="mx-auto mt-10 max-w-7xl px-4 sm:mt

```

FIGURE 8.2: How `mmlog` calls appear in the editor, including star icons in the gutter to either include or exclude the calls and their corresponding nodes from the rendered “Joint History” visualization (Figure 8.1-1) and code version boxes (Figure 8.1-2b and 3b) in the MMAI pane. Filled-in star icons mean the nodes will be included. Note that these two functions, `calculateBoxValue` and `App`, are in different files.

(Figure 8.2-3). This allows Linda to not only see the histories and log statements and values in a visually grouped manner, but also allows her to query on the composed history of these otherwise-separate parts of the code.

With the previously-captured log values in MMAI’s memory and her set of `mmlog` statements and values of interest, Linda uses MMAI’s querying system to compare her current, incorrect log value to the last-known correct value and the corresponding code. In the current code version, she explores the current code version and its run data in her function of interest. In Linda’s case, *the amount of log values* is a considerable challenge, given that both the live stream data and the rendering loop run continuously. MMAI provides a high-level summary of the current run’s `mmlog` data (Figure 8.3-7) and how it compares to previous code versions’ values of the same data. The LLM performs this comparison – in looking at Linda’s data, the LLM finds that the data in the current run includes negative values for the x-coordinates and suggests comparing code versions (Figure 8.3-8, “Compare to current version” button).

In isolating her problem, she wants to know whether the Twitch API is returning incorrect data, thereby interfering with her calculations, or whether the calculations are implemented incorrectly. Through comparing the earlier working implementation shown to her by the LLM on the timeline to the current implementation, the calculation implementation appears the same. She wants to be thorough in her investigation, so she asks MMAI to find other logs related to her calculations (Figure



FIGURE 8.3: How an expanded code version appears in MMAI. (1) is similar to Meta-Manager (see Figure 7.2), in that information about the edit is shown including when it occurred and what the code looked like at that point in time, along with any additional meta-information (with this version containing no additional meta-information). (5) shows the optional, supplementary information when `mmlog` is used, including the log statement (6) (in case multiple `mmlog` statements exist in one version), most recent log statement and values along with aggregate counts and a button to ask about that value (7), and computed queries that can be sent to the LLM, given the `mmlog` statement and value (8), along with the option for the developer to write their own query (9).

8.3-8, “Find other logs related to box” button), which updates the MMAI pane with other code versions that include `mmlog` calls that include the variables in her current version. Upon inspecting these versions by similarly inspecting the MMAI code version boxes, she finds the same erroneous values. Using the joint history user interface, she similarly ensures that the front-end rendering values are the same as the back-end computations, so she turns her attention towards the Twitch API.

To further dig down into her problem, she uses the MMAI chat system to directly ask the MMAI AI model about her bug. She begins with the pre-formed query of

“Why am I seeing this output” (Figure 8.3-7, question mark button) about her erroneous calculation result. The AI model uses the output values on the current code version as reference and the surrounding context including code, the history of the App function given the joint history, and imported APIs in order to inform its answer (see Section 8.4.3). Using the contextualized information, the model responds that the Twitch API sometimes is laggy in its response time, which could lead to the box coordinate calculation using old data, instead of the current data. MMAI additionally generates follow-up questions given the model’s response, with one of the generated responses asking how to ensure latency with the Twitch API. These follow-up questions are requested as part of the prompt we give MMAI— we instruct MMAI to, given the just-generated answer, along with the additional context in its context window, to come up with the most relevant follow-up questions a developer may want to ask². Linda chooses to ask that question and receives a response about how JavaScript’s `async/await` constructs can resolve this bug. Linda resolves her bug using the model’s suggested code.

8.4.2 Detailed Design

In designing MMAI, we sought to provide functionality that both supports the needs we found in our formative study (Section 8.3.2) and allows for varying levels of inquiry through utilizing the historical information Meta-Manager already captures, while improving upon that architecture using state-of-the-art AI. Below, we discuss how the various parts of MMAI support different stages of the print debugging process with respect to the queries asked at each point.

Locating Output Values of Interest

A recurring theme that arose in our formative study was that methods for isolating log statement(s) and values(s) of interest were ad hoc and often broke down when presented with challenges such as large-amounts of log data. To counteract this challenge, MMAI supports mechanisms for both *finding* and *grouping* log values of interest, *filtering* out irrelevant values, and *summarizing* large sets of output data.

Log values are presented in-context within each code version, such that developers can visually compare different values produced by the same log statement and values produced by multiple log statements within the same function. In this way, MMAI utilizes the Meta-Manager’s block-based organization of code information, which has been successful in combating issues of scale.

The adopted Meta-Manager data organization naturally lends itself well to filtering out values that are not of interest. The user interface already organizes log statements and values by node, meaning a user can directly navigate and inspect log values in the MMAI pane by node and log statement. With this organization, the

²For all prompts used with MMAI, see Appendix B.

user does not need to actively filter out values, and, instead, they can focus on the values they care about.

Developers can further locate related log statements through using MMAI’s query to find logs related to the currently in-focus log and its subject matter (see Figure 8.3-8, “Find other logs...” buttons). MMAI utilizes both the capabilities of its LLM for understanding code and control flow in conjunction with MMAI’s knowledge of the AST (another feature inherited from the Meta-Manager architecture) in order to isolate candidate `mmlog` calls that may be related to the identifiers included in the original `mmlog` argument list (e.g., `box` or `offsetBound` in Figure 8.3). In this way, the user can isolate and synthesize a group of semantically-related but potentially disparate logs and present their values together, addressing Q1 and Q2 from our formative study.

MMAI additionally utilizes the Meta-Manager architecture to improve upon its visualization and presentation of code nodes for synthesizing historical and log data by allowing users to create their own groups of nodes. When the user runs their code, MMAI will detect each executed `mmlog` statement and place a star icon in the gutter next to the statement. A user can then click the star icon(s) (Figure 8.2-3), which will add the associated `mmlog` call and its surrounding node and history to a user-defined set, which will then be rendered in the MMAI pane. Given findings from our formative interview (i.e., Q1: “Show me just log statements A, B, and C”), we wanted to directly support this line of inquiry without requiring the user to have to format their log statement in a particular way to allow for text-based search, like the brittle “key-value” pair strategy we saw in our study. Further, since the MMAI visualization is already organized by time and code version, the traditional time-based ordering of information in a console lends itself naturally to this visualization approach.

Another way in which MMAI leverages the Meta-Manager architecture is through its versioning of log statements and values as part of each code node’s history. When a `mmlog` statement is invoked when the user runs their code, MMAI matches that `mmlog` statement to the known `mmlog` statements in each code node and, if a match is found, appends the generated value to the history for that particular statement. If a match is not found, MMAI will start a new history for that log statement. These values are then stored in the MMAI database on the next file save. Through this data organization, not unlike the code node discussion in Section 7.2.3-[F1], we attempt to cut down on scale through keeping individual database entries as small as possible by having each call to `mmlog` be its own entry in the database (as opposed to aggregating or summarizing all calls that occurred across, e.g., a run or particular code version). This also helps the user find individual instances of calls of interest, since each log statement and value is treated independently, thus a path forward for answering queries like Q5 which asks “When did this log statement last produce the correct value?”. This ability to drill down on a particular `mmlog` statement also makes answering queries such as Q3 (“Remove all log values that are not related to variable A”)

possible by removing all unrelated log values. Nonetheless, if the user is logging a very large value (e.g., a recursive JSON structure), we still run the risk of the entry being too large for the database and/or the value being intractable and thus too difficult for a user to make sense of, especially once some time has passed. We further discuss these limitations and potential solutions in Section 8.6.

A feature unique to MMAI is its historical knowledge of log statements and values that are contextualized to a particular code version and optional supplementary meta-information. With this knowledge, MMAI can both summarize the history through retrieval augmented generation and compare the most recent log value(s) to its understanding of the code's intended result and previously-produced values to assist in isolating anomalies. Through these multiple approaches to finding log values of interest, we expect MMAI will be an improvement over the ad hoc mechanisms developers typically employ.

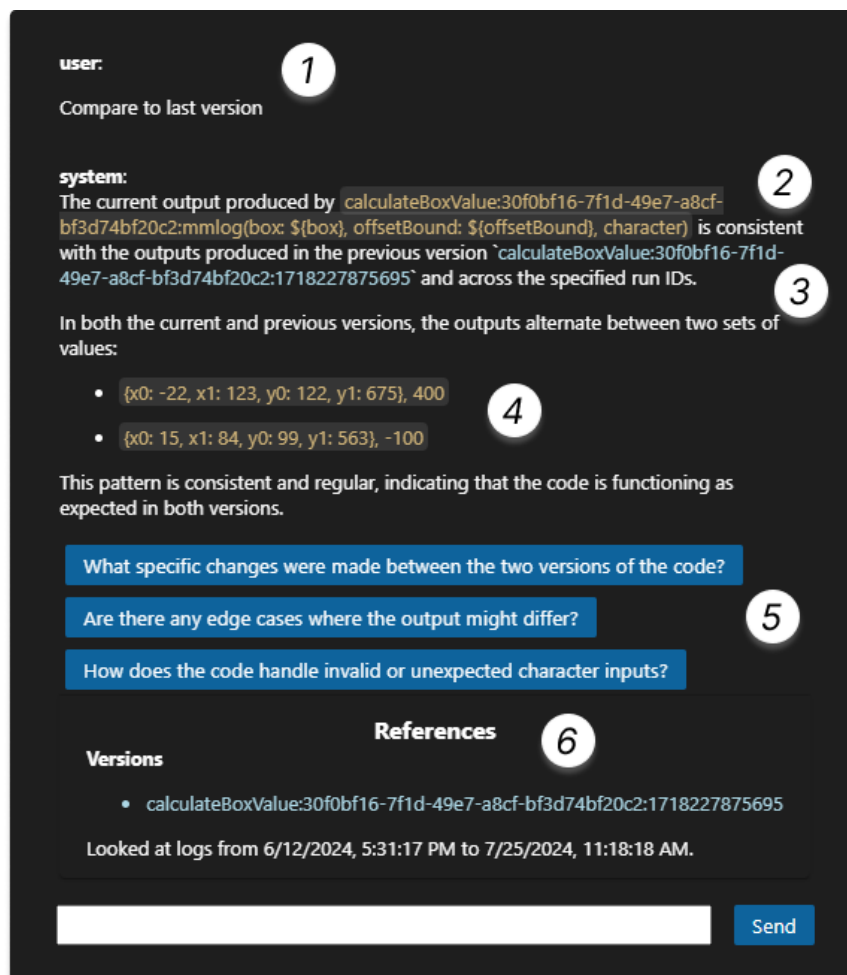


FIGURE 8.4: How the chat interface appears within the MMAI interface, below the log and version data. (1) is the user's query – in this case, this is a pre-generated query, created by pressing the "Compare to current version button" (Figure 8.3-C). (2) is the LLM's response. (3) is a link to the code version the system is looking at when coming up with its answer. (4) are the log values the system found in its history search. (5) are the generated follow-up questions given the answer the system generated. (6) are the references the system used, including links to code versions and a summary of the span of history it looked at when coming to its answer.

Synthesizing and Understanding Log Values

MMAI allows for developers to make sense of their log values in multiple ways. An advantage of existing within the editor is MMAI can utilize more context when allowing users to prompt the LLM model. To assist in this prompt-creation, MMAI has pre-formatted prompts that the user can just press a button to ask, while also supporting free-form prompt generation. The pre-formed prompt associated with the question mark button (Figure 8.3-7) supplies additional meta-information in the prompt including the current code version, a summarized history of the changes before, and information ingested about any external libraries the user may be using, including version numbers taken from the user's `package.json` file. In this way, MMAI attempts to provide as much contextual information as it can to support the LLM in coming up with a reasonable explanation. We additionally require the system to highlight what information (including prior code versions or log values – see Figure 8.4-3 and 6, Figure 8.1-5) it used when coming up with its answer to improve the trustworthiness of the answer and to reduce the likelihood of hallucinations.

Despite our best efforts to improve the quality of generated LLM prompts, it is unlikely that a single LLM answer will completely address whatever issue a developer is encountering. We provide follow-up questions to each LLM-generated answer in order to further utilize the ingested context for reasoning through what the problem may be and further reducing the amount of mental processing the user needs to do (see Figure 8.4-5). These generated questions are created by having the LLM not only come up with an answer to the user's query, but also come up with follow-up questions to the answer given what it found in its historical synthesizing. In the formative study, P7 discussed how he wanted debugging support from a historical oracle to not only function as a repository of prior knowledge but to serve as proactive aide, not unlike a pair programming partner, by coming up with alternative lines of inquiry and suggestions. We use these follow-up questions as a first pass at attempting to provide this pair-programming type assistance when reasoning about what log value(s) may mean and how to address them.

Through directly supporting composition of logs, MMAI can provide more tailored support for reasoning about log statements and values through the creation of semantically-meaningful relationships. When a user creates a group of `mmlog` statements that are not in the same node³, MMAI generates a new, synthesized history that organizes and merges the constituent nodes' versions, log statements and values, and any additional meta-information into a singular history that a new LLM agent can reference when formulating an answer. In this way, the user can ask more complicated questions regarding the relationship of multiple functions, log values, and versions across time that would otherwise be very difficult to ask in a de-contextualized format. For example, this synthesizing of information allows for

³If the `mmlog` statements *are* in the same node, there is no need to generate a synthesized history, given that the node has access to each `mmlog` statements' historical data, including prior versions and log values.

a user to ask “Why does box go from an object to undefined when this function’s code changed?” – with the synthesized history knowledge, our system and the LLM can isolate when the log value was last *not* undefined, what was changed since that point in time between both implementations, and (hopefully) come to an answer.

8.4.3 Implementation

MMAI is a Visual Studio Code extension for the JavaScript and TypeScript programming languages. Considering it leverages the Meta-Manager architecture, some features are identical to that system, including the node matching algorithm utilizing our Firestore database. Logs are now, additionally, stored on the database and associated with a parent code version.

Log Data Design and Processing

In designing MMAI, we extended the Meta-Manager architecture to allow for handling of log statements and values with our own Node.js package for logging. As a brief reminder, Meta-Manager pushes new versions of each node to the database on file save, given whether or not the node has experienced a change that should be saved to the database (see Appendix A for a full discussion on Meta-Manager’s event model). Capturing log values at the code version level introduced some complexities to this model that required design and implementation consideration.

One challenge was how to capture whether or not the code was run and what to count as a run. Given that MMAI is supporting web development, specifically, the event-based nature of such development means that code is nearly constantly running and many paths of execution are only triggered when a certain event, such as a click, occurs. Thus, one version of the code may have hundreds of runs or none at all, depending upon what part of the code the user is testing. For simplicity’s sake, we utilize the VS Code API’s native debugging API event, `onDidStartDebuggingSession`, as our ground truth for whether or not the user is running their code. Each debugging session has a unique id that we associate with any emitted log values from `mmlog` as a way of grouping runs as an additional mechanism for organizing log values. Notably, there are other ways of running code, aside from VS Code’s debugging functionality, such as using the command line, that MMAI does not currently listen for, which future improvements to the system should account for.

When the user is debugging and has `mmlog` calls in their code, `mmlog` transmits information to MMAI which handles the delegation of the log statement and values to the corresponding node. The `mmlog` package acts as a wrapper around `console.log`⁴

⁴Unfortunately, for a user to receive the benefits MMAI provides, despite `mmlog` acting as a wrapper around `console.log`, a user *must* use `mmlog`, given the connectivity between the `mmlog` package and the MMAI VS Code extension and given the matching logic within the MMAI extension for finding invocations of `mmlog` to track.

and transmits any arguments to the `mmlog` call on a socket connection shared between MMAI and `mmlog`. `mmlog` values appear in the user's console as normal, despite the fact that their output is also handled by the MMAI extension and saved to our database. In order to match `mmlog` calls between the user's source code and the actual log values received at run time, `mmlog` additionally transmits the stack trace, which includes the file name and line number. MMAI takes that stack trace and searches for a matching file name in its set of managed user source code files. If a file is found, MMAI transmits the `mmlog` statement and values, along with the line numbers to MMAI's document manager which finds the node that contains that line number and `mmlog` statement, such that the node can add the log data to its history (see Figure 8.5 for an example of how log data looks in history). In situations where a file is *not* found, which may happen if the file is not managed by MMAI (such as if the file is in an external package or the file is not a JavaScript or TypeScript file) or the file has been renamed since MMAI was instantiated (which could be remedied by listening for VS Code's file renaming event), the emitted log value will be ignored. To reduce the likelihood that a particular value is sent to the database too many times and is too large, we run a simple heuristic of whether or not a just-received value is the exact same time as the last seen log value received from the same `mmlog` statement and was just appended to the node's history (i.e., was seen within the last second). Other techniques for guarding against storing too much data will need to be added in the future. Such techniques may include both technical improvements, such as additional heuristics for limiting what is sent to the database, and user interface improvements, such as warning messages stating a particular log value will not be tracked when a user logs a too-large value (such as the whole DOM) to cue the user in that they will need to be more selective if they want a particular value to be tracked.

One advantage of handling log statements and values in a VS Code extension is that we can specifically track the log statement lines in the user's source code. Similar to how nodes are handled, each `mmlog` statement is given a unique identifier derived from the parent node's ID and the contents of the `mmlog` statement. If the user edits their `mmlog` statement (e.g., `mmlog('ringo is', myDogRingo())` to `mmlog('ringo is', myDogRingo(), myCatEevee())`), MMAI detects that edit and updates that `mmlog` statement's ID, both within the extension and in the database, such that future `mmlog` values are associated with the same `mmlog` history.

GPT-4 Support

For the GPT-4 support, we chose to have each node of the AST not only manage its code versions and log statements and values, but also manage an intelligent agent in the form of a unique LLM instance that is given access to that node's history and log data. To reduce the potential monetary and computational cost of having many running LLM agents, only some of which are necessary, we only "activate" (i.e., create the connection to the OpenAI API for that node and give access to history) a node's

```

calculateBoxValue:30f0...mmlog(`box: ${box}, offsetBound: ${offsetBound}`):1718229117103
{
  key: "calculateBoxValue:30f0...:mmlog(`box: ${box}, offsetBound: ${offsetBound}`)",
  mmlog: {
    location: {
      content: "mmlog(`box: ${box}, offsetBound: ${offsetBound}`)",
      fsPath: "c:/Users/linda/twitch/src/utils.ts",
      range: {
        start: {
          line: 16,
          character: 4,
        },
        end: {
          line: 16,
          character: 49,
        }
      }
    }
  },
  value: [
    "box: { x0: -22, x1: 123, y0: 122, y1: 675 }, offsetBound: 400"
  ],
  runId: "2f5b239-69cc-4d72-576d04b990c2",
  time: 1718229117103,
  versionId: "calculateBoxValue:30f0bf16-7f1d-49e7-a8cf-bf3d74bf20c2:1718229117103"
}

```

FIGURE 8.5: How an `mmlog` version appears in the MMAI database. One code version owns a collection of these `mmlog` versions dependent upon how many `mmlog` statements the MMAI node has at a particular version and how many times each `mmlog` statement is invoked.

instance when a user first clicks the star icon for a `mmlog` statement. Each node's code versions (see Figure A.2 for an example version) and `mmlog` version (Figure 8.5) are transformed into a string representation that are joined together to make a long historical record, sorted by time, and fed into that node's GPT-4 instance (see Appendix B for prompts). We chose to do this such that the LLM could ingest more specific data as opposed to all of the history across, e.g., the whole project, and suffer from too much context. Further, ingesting that much data would likely exceed the token limit for the model (as of writing, GPT-4 supports a context window of up to 128,000 tokens). Nonetheless, if a node has a particularly long history or contains many log statements and values, there does exist a possibility that the history may exceed the token limit, in which case we batch calls to the OpenAI API when querying on its history. While batching calls may work, there also exists the possibility that a particular query may not require access to the whole history, such as in the case of finding other related `mmlog` statements – future versions of MMAI may benefit from having more strategic history ingestion and prompting given a user's line of inquiry.

8.5 Discussion

We position MMAI as an exploration of to what extent more robust and domain-specific meta-information collection can assist with developer sensemaking when assisted with LLMs. Without a proper user study, it is hard to know for certain whether developers appreciate the class of features supported by MMAI and whether it does, indeed, cut down on the amount of lower-level information scrounging and synthesizing work typically required during print debugging. Nonetheless, we position this work as an introductory exploration into how we can re-frame AI assistance in programming as an opportunity for cutting down on the tedious work often required when developing code. These small productivity gains can add up over time and relieve cognitive strain that can be better applied to more creatively-fulfilling tasks, such as designing a coding solution to bugs.

MMAI and its design, like the prior systems, were initially motivated by thinking through challenges I have experienced in my own development and how human behavior is often at odds with the design of more formal software development tools. As discussed previously in Section 8.3.2, more sophisticated debugging support does exist and the productivity gains one can experience through becoming well-versed in using debuggers can be significant. Nonetheless, despite having used debuggers previously, I find myself still almost always choosing to use print debugging rather than a debugger, given the much lower cost and the nature of my path of inquiry. The insight that print debugging is popular is not a novel contribution of this work, yet little tooling has attempted to actively support this practice, with the most relevant and recent work focusing primarily on re-shaping log value data but not supporting other aspects of print debugging [112]. This work seeks to celebrate print debugging as a useful practice that should not be shamed or written-off as “inefficient” when compared to more powerful mechanisms.

When placing MMAI in the context of the other work presented in this dissertation, many of the same challenges and themes arise. In returning to the concept of “information scraps” and meta-information, more broadly, log values when print debugging share the ephemeral nature of this class of information, in that it typically has been only useful in the moment it is generated, despite having the potential to answer later questions. However, these log statements made in the moment to support print debugging are often not written with the intention of having long-term utility, leading to terseness that makes their interpretation difficult. Indeed, multiple interview participants expressed that they systematically “clean up” these logs once they are done debugging, such that they do not interfere with the business logic of their code or take up valuable space in the console. This tension between keeping this information to support later sensemaking versus removing the information so it does not interfere with other development-related tasks is partially rectified in MMAI’s adopted architecture from Meta-Manager which airs on the side of “save everything” considering database storage is relatively cheap in comparison to the

human cognitive power of recall. Nonetheless, supporting other information scrap management tasks, such as curation, may make sense when considering extending this work.

Another advantage of the “save everything” design principle in this work is exploring how that data can be utilized by LLMs. A task LLMs are particularly good at is summarizing and retrieving data [268] – we take that idea and utilize it for retrieving log statements and values that can act in service of assisting with a developer’s reasoning about code behavior when print debugging. Considering we are using MMAI largely as a retriever and summarizer of a large repository of information as part of the retrieval augmented generation class of LLM usage, we additionally require MMAI to include references to what information it used when coming up with its answer. In this way, we expect that developers will have an easier time ascertaining whether or not an answer is trustworthy, given that they can reference the citations the system is required to produce upon generation. While this is not necessarily a “fool-proof” solution, it is a first step towards assisting in bug isolation and reasoning in a more precise manner than context-free usage of LLMs, given the large amount of contextual and historical data we are feeding the model.

8.6 Future Work and Conclusion

Some technical issues with MMAI may arise that require more significant consideration. One such issue is what to do when an output value is too large to store, such as if the user chooses to log the entire DOM on a web page. We attempt to mitigate potentially spamming the database by only pushing when the current value is different from the last-saved value and/or some time has passed since the last push to the database, but this approach can still result in large objects being pushed intermittently. Other, more robust potential solutions include fragmenting the object across multiple database entries or otherwise minifying or truncating the representation, such that it can be stored. An alternative solution is to grant the user more control in terms of what is logged – indeed, when discussing logging objects with P7, they mentioned that they are often only interested in one or two properties of a large object, yet specifying those specific properties can be onerous. JavaScript’s `console.log` method logs the actual object in a “collapsed” or summarized form then allows the user to “expand” the object properties they are interested in and lazily loads the referenced values – one can imagine adopting a similar design decision for `mmlog` in which values are only stored indefinitely if the user interacts with them in the user interface, or are, at the very least, prioritized for later retrieval and are marked for the LLM as particularly important. Having further insight into the intent behind *why* a user is choosing to log a particular value could lead to smarter decisions when choosing what information to push to the database and how to use that data.

In the current implementation of MMAI, we only capture values specifically logged using `mmlog`. This was an intentional design decision since we wanted to mirror the development practices developers are doing with, e.g., `console.log` when print debugging. Nonetheless, this choice results in us not capturing other information that may be helpful for a developer and an LLM agent when interpreting `mmlog` values post hoc and further contextualizing buggy code. Additional data future versions of MMAI may consider supporting include other information logged to the console (e.g., error messages), and output data, like the HTML of a webpage. Output data, especially graphical output, could be helpful in further understanding how a bug is manifesting. When considering the scenario described earlier (Section 8.4.1), Linda first realized she had a problem in her code given the graphical output, with a box being rendered incorrectly. If she or another collaborator wants to later understand the bug fix she applied, a capture of the graphical output may be more useful in remembering what the bug was, as opposed to the logged negative numbers. Print debugging, and the log statements and values produced during that process, represent a line of inquiry that is often motivated by both information that is gained through the log values but also information that is generated elsewhere – capturing more of that line of inquiry may result in a more thorough understanding of the developers print debugging practices and may make reasoning about and questioning of history at a later point easier. One can imagine an extension to MMAI that allows for automatic screenshot capturing when the code is run (i.e., when the user runs a new version of their file and the new output is displayed in the web browser) with a web browser extension (not dissimilar to Meta-Manager and its shared connection between browser and VS Code extensions). However, this introduces a new design challenge in terms of understanding *how* that output should be associated with the code and log history – considering the output is an amalgamation of some or all of the various code node’s implementations, how should an output be associated with each of the nodes? One design idea could be adopting the technique used in the WhyLine, in which a developer can “point” to the parts of the output that are relevant [130], but this comes at the cost of introducing additional “work” for a developer, which this work largely tries to avoid, so future considerations are required.

Another effect of only capturing log values generated with `mmlog` is that this requires developers to use an external package and editor extension in order to receive any of the benefits of this work, as opposed to this support being natively provided by current editors and logging functions, like `console.log`. Similar to Meta-Manager and how a developer needs to both have a VS Code extension and Google Chrome browser extension in order to experience the full benefit of that system, we position this work as an exploration of the potential benefits of having this more tightly-integrated working system, in which components of the system are aware of other, relevant contextual meta-information generated in other usually disparate systems. Like Meta-Manager, this approach introduces some risk to developers in

terms of information privacy and having their activities potentially monitored. Future work may explore how to allow for better information protection and filtering on the developer's end.

Some future work from Meta-Manager is still relevant here given the relationship between the two projects, such as supporting annotating or sharing specific code versions with collaborators. In the context of print debugging, one can imagine "saving" a user's path of inquiry that led to their identification and fixing of a bug such that the information can be documented in a bug report or pull request description. This "saving" may also be useful for communicating with collaborators, such as when working in quality assurance testing, where a tester may find a bug and needs to share how a bug is encountered and behaves with a different developer. This practice of transforming a user's sensemaking process into a shareable format is not dissimilar to the work discussed in the curation chapter (Chapter 5) and lessons learned from that work may apply here.

Given MMAI's implementation as a VS Code extension and the system's knowledge of a user's source code, some additional functionalities to support print debugging are possible. Functionalities future versions of MMAI may consider supporting include removing no longer relevant `mmlog` calls if the user wants to clean up their code, directly importing code from LLM responses into the editor, or automatically inserting `mmlog` statements into the code if the user wants to log every instance of a variable or function call. In the current version of MMAI, we chose not to modify the user's code because we did not want to potentially introduce issues given uncertainties around potentially improperly applying an edit and given our lack of certainty in the correctness of our LLM's generated code. Nonetheless, considering a large issue in print debugging is the resulting clean up of log statements and their interference with the console content and business logic of code, automatically removing log statements is an obvious and powerful solution to this problem. We chose to partially address this challenge through supporting filtering of `mmlog` values and allowing users to directly manipulate the MMAI interface to make it clear what `mmlog` statements they wanted to focus on during their current print debugging session.

The space of development tasks that may be offloaded or, at the least, accelerated with AI is vast. Yet, tools focus primarily around code generation and, to a slightly lesser extent, code understanding. In this work, we take AI capabilities and place them in a different context, namely supporting print debugging. An obvious and planned next step for this work is to perform a user study to see to what extent MMAI actually succeeds in helping users make sense of and resolve bugs when performing print debugging.

When looking at the future of software engineering, it is important to keep the "engineer" in engineering. The anxiety that LLMs are introducing into many different practices, including software engineering, about job security and what the future of these industries will look like given the disruptive nature of these technologies is making people apprehensive about even joining the field. As an academic, I believe

it is paramount we try and use these technologies to augment and accelerate the skills developers already have as opposed to introducing technologies that replace the engineer. These technologies are not going away, so we have to learn to use them for good.

Chapter 9

Conclusion and Future Work

9.1 Summary of Contributions

The series of work discussed in this thesis explores the utility of supporting developers in both authoring and utilizing meta-information to better make sense of code and documentation. Specifically, the systems provide platforms for authoring and sharing useful and contextualized information about code (Adamite, Catseye, Sodalite), unify various forms of meta-information into one platform with a single interaction method (Catseye, Meta-Manager), utilize artificial intelligence for accelerating sensemaking with meta-information (MMAI), and demonstrate how meta-information may be used to overcome known barriers in documentation usage and answer long-standing, hard-to-answer questions about code (Adamite, Sodalite, Meta-Manager).

My work has made the following contributions:

- A set of key design properties that should be considered and accounted for when building meta-information systems for programming (chapter 1).
- A review of literature around developer information needs, systems developed to support some of these needs, and how meta-information about code has or can be used to combat these challenges (chapter 2).
- Identifying that short, easy-to-author notes, when attached to programming-related documents including code and documentation with anchors, can address known challenges in information tracking and usage [104, 107] (chapters 3 and 4).
- Adamite, an annotation system designed with specific features to use meta-information as annotations on documentation to overcome known barriers in using software documentation [104] (chapter 3).
- Evidence that developers' notes can be useful for other developers when properly presented given the implicit context communicated through annotated text [104] (chapter 3).

- Identifying that annotations and their anchor points can act as a unified interaction design for capturing and associating different types of code-related meta-information [107] (chapter 4).
- Catseye, a code annotation system, with features to help a developer keep track of code-related meta-information using a unified annotating interaction [107] (chapter 4).
- A literature review of information needs developers have when making sense of and curating ephemeral, short-form information to assist in sensemaking tasks (chapter 5).
- An exploration of design considerations and probes on how to manage large sets of ephemeral data (i.e., annotations) in a complex and oft-changing environment (i.e., source code) where re-attachment and curation are non-trivial (chapter 5).
- Identifying that documentation text, when anchored to code, and the stability of that code anchor can be used to assess how out-of-date some documentation is (chapter 6).
- A documentation authoring and maintenance system, Sodalite, that utilizes code meta-information to support developers in assessing the validity of documentation and finding appropriate parts of the code to document [106] (chapter 6).
- A mapping between historically challenging-to-answer questions about code and the types of meta-information that may be used to answer those questions (chapter 7).
- A system, Meta-Manager, that can collect, index and store a larger collection of code editing and provenance meta-information that has never previously been collected at scale for answering otherwise unanswerable questions about code history [105] (chapter 7).
- Identifying a set of queries that developers have when performing print debugging that are not currently supported in modern tooling (chapter 8.3.2).
- An LLM-powered system utilizing retrieval-augmented generation, MMAI, that captures code meta-information at scale, including output data, to support print debugging (chapter 8).
- A series of lab studies showcasing the usability and utility of these tools in supporting developers' varied information needs in different contexts.

9.2 Discussion and Future Work

In developing the aforementioned systems, including determining design and functionalities, we gained insights into ways in which developer meta-information may be collected, transformed, and presented to assist in developer sensemaking. Below, I briefly reflect on some of the implications and lessons learned from that exploration, along with avenues for further inquiry based upon these implications.

9.2.1 Designing with Information Ephemerality In Mind

When beginning this thesis, I was initially drawn to the idea of supporting lightweight developer sensemaking through notes and annotations due to the seemingly-untapped market of developer note-taking tools. It was only upon beginning to do more literature review that it became clear that developers were, indeed, taking notes, but there was no single tool or practice that was ubiquitous. Instead, developers naturally did what other information workers often do and utilized implicit cues such as open files or tabs for assisting in recall and their natural workspace (i.e., source code and code comments) for jotting down small notes such as open to-do items. Code comments, in particular, seemed, as a “user interface” for note-taking, to be at odds with the paradigm of note taking – code comments are represented within the source code, meaning they will persist indefinitely, whether that be in the current source code or in a commit contained within a VCS. Further, they do not provide support for richer expression of thought or structuring of information, such as creating follow-up comments in the form of a reply.

Catseye, as opposed to code comments, explores more directly supporting information ephemerality with respect to developer note taking. The conventional design approach when creating an information system is to allow information to persist indefinitely, be accessed indefinitely, and be editable indefinitely, such as in the case of Google Drive, where any document you have access to will exist in your drive unless the permissions are changed or it is deleted (and, even if a Drive document is deleted, it is still accessible in a trash bin for a finite period of time). When starting the curation work (Chapter 5), we were initially motivated and bound to a similar design goal – make annotations *persist* as long as possible (i.e., keep them attached to source code) such that the *user* can derive value from them as long as possible and do something with the annotations at a *later point in time*. However, the more I reflect on this work, the more I wonder whether perhaps that is the wrong direction and the correct direction is “how do we get rid of erroneous or no longer useful information as quickly as possible”. Indeed, prior work (e.g., [21, 151]) and my own experiences using Catseye (Section 4.6 and Section 5.4) suggest that much of the information generated is irrelevant past a certain, finite window of usefulness. Not all information is created equally, yet systems often treat the information as such.

Future work may benefit from intelligent mechanisms for automatically curating content. In the context of annotation management, one potentially fruitful avenue

may be intelligently inferring relevance to a particular user and the relative importance of an annotation. Thus, curation becomes a function of two different needs that can be calculated dynamically. Meta-Manager showed the promise of keeping track of information while not requiring anything of the user and MMAI showed that this information can be used by an intelligent agent to assist in querying. Relevance to a user could similarly be calculated through a query or through knowledge about the user's current task, which may be inferred from edit patterns [118] and code locations of interest [45]. How to intelligently determine the importance of an annotation is less obvious – ideas we have had include calculating how often the annotation is revisited, how thorough it is (i.e., does it have multiple anchors? Replies? Long annotation content?), and how recently it was created. A truly intelligent user interface for these annotations could have annotations hidden from the user once they have lost relevance to them, with the expectation they likely will never become relevant again, save for if we can ascertain that the annotation will accelerate a later user's sensemaking.

9.2.2 Designing for the Software Engineer of Tomorrow

Much of my work has been predicated on both understanding the information needs of software developers given the large body of Software Engineering and HCI research that discuss these needs, and considering what developer meta-information may be utilized to address those needs. For example, Catseye, Meta-Manager, and MMAI all explore mechanisms for using meta-information to assist in code authoring, comprehension, and debugging, respectively. However, this meta-information capture and transformation may look vastly different in the coming years if developers are not actively writing code and, instead, guiding AI tools to write code. Developers will most likely still need to understand the existing code to design their changes – what meta-information will be most useful for this next developer to formulate, e.g., the correct prompts for a LLM? How will harnessing this meta-information scale as more of a code base becomes sourced primarily from AI?

How I see software development moving is that understanding code at the syntax level will become less important as understanding the intent behind its design and how it was created will become more important. With code generated by AI systems becoming more common, tooling solutions for tracking the creation, evolution, and intent behind this code will become more relevant. Software development has largely followed a pattern where programming is abstracted further and further away from the actual machinery with tooling to support these higher levels of abstraction. For example, in the move from writing assembly code to writing C code, tools were needed to better help with diagnosing memory allocation bugs – in the move from writing code to guiding AI systems (e.g., ChatGPT, GitHub CoPilot) to write code, I see an expansive and rich design space for improving the usability of these tools and the comprehension of the code created by them through leveraging meta-information generated in this authoring process. My recent work with MMAI

begins to probe on this thread by utilizing LLMs to support software development tasks not directly involving code generation or understanding, which is primarily what AI code-development tools have been used for. Future work, in moving towards an “IDE of the Future”, may similarly explore what emergent tasks come out of software development using LLMs (e.g., prompt (re-)authoring is now a common activity amongst software developers [148]) and what tooling support is necessary to support these tasks versus what classic software development tasks are no longer relevant, if any.

A related research thread, when thinking about how classic software development challenges may change when using LLMs for development, is considering what new “hard-to-answer” questions [139], if any, emerge during this paradigm shift. The Meta-Manager, specifically, was designed to collect contextual, authoring-time meta-information at scale that could be used to answer some of the hard-to-answer questions derived from LaToza and Myers’ seminal paper [139]. Some of the questions described in the paper, such as questions about how to best implement some code or how to refactor some code, can potentially be answered by an LLM agent (whether or not its answer is correct is a different question). Other questions, such as questions of design rationale, which were the largest blockers, rely upon extensive knowledge of the code’s history and sometimes information that is invisible to an LLM (e.g., web activity, discussions amongst team members, GitHub PR descriptions), meaning such questions cannot currently be answered by an LLM. Perhaps it is too early yet to do a follow-up study to LaToza’s work that seeks to explore whether the hard-to-answer questions have shifted in terms of difficulty or if new questions have emerged, given that LLM-authored code has not been around long enough for developers to know how it is reshaping their question-asking and answering. Still, future work would benefit from understanding the information support needs of developers in a future filled with LLM-generated code.

9.2.3 Designing for Meta-Information as a First-Class Entity

The rise of LLMs holds implications for meta-information. Information provenance, especially, is only going to become more important as users in a variety of contexts need to ascertain *where* some information came from. For example, in the education sector, there is ongoing discussions on how to determine whether assignments and tests are actually being completed by a student versus by an LLM – *where* is the content in the assignment originating from. Given concerns around generative AI in a variety of disciplines, tracing information to its origin is becoming more necessary universally. In the context of software development, I showed that this can be possible given extensive, domain-specific tooling with the Meta-Manager and its connectivity between IDE and browser. While this approach was successful, researchers reasonably expressed some hesitance around sharing where and how some information (code, in this case) came to be. Further, Meta-Manager only operates one layer of provenance deep – in other words, if a user copy pastes code from

Stack Overflow, but the code on Stack Overflow was generated by an LLM, Meta-Manager only knows that the code came from Stack Overflow. Designing solutions that can both capture meta-information, including historical and provenance data across platforms and “layers” of sharing, while also respecting a user’s right to privacy will become even more necessary as more generative AI content proliferates our day-to-day lives. More extensive information provenance work can also help with answering other research questions – for example, in the context of software engineering, researchers have already begun wondering whether code produced by LLM’s introduces more build breaks [24], a question which cannot be answered without code provenance meta-information since AI-generated code is typically indiscernible from developer-authored code.

More broadly, my research has shown that collecting large amounts of highly-contextualized meta-information can assist in sensemaking, both for an initial author and later user. Given how storage sizes have ballooned over time and database services have only gotten more powerful, *storing* this meta-information is not necessarily the challenge. The challenge comes in *making this information useful*. One design technique nearly all my systems utilize that has been effective is creating strong, traceable, and visual connections between meta-information and the original information source it is derived from. My annotation work, in particular, succeeded primarily due to the in-context nature and presentation of the annotations. My later works, especially Meta-Manager and MMAI, explored how to introduce new dimensions to that information connectivity through exploring relationships such as time, and new presentations, including visualizations. Sodalite and MMAI additionally showed how utilizing that information connectivity can support new tasks, such as determining documentation accuracy, and that meta-information connectivity and the context such a relationship creates can be utilized by an LLM for completing higher-order intelligence tasks, such as reasoning. Future work would benefit from exploring other ways of making meta-information more valuable through further exploring the design space of how to present meta-information, especially when dealing with competing information dimensions, and how to transform and use meta-information and information connections for serving domain-specific tasks. In Chapter 1, I presented some key design properties that may similarly be valuable to consider when capturing other types of meta-information which were not explored in this thesis, given a particular domain. Collecting more meta-information is paramount in order to serve these design and research goals, so I once again implore researchers to treat meta-information as a first class entity.

9.3 Concluding Remarks

Developing code is hard. Anyone who has written code can attest to this – given the sheer amount of information management work involved in balancing tasks, designing solutions, navigating through an IDE, and so on, developers are under intense

cognitive strain, while often amassing a large set of highly-contextualized information that is subsequently thrown away or remains tacit within the developer's mind.

My research has explored mechanisms for both extracting developer's tacit knowledge in the form of meta-information and presenting it to accelerate a later developer's sensemaking through leveraging design properties of meta-information. Through this research process, I have blended insights from the fields of Human-Computer Interaction and Software Engineering in order to serve developers' specific information needs. HCI research in other domains, such as web browsing and online shopping, has shown that capturing aspects of an initial user's sensemaking journey can accelerate a future user's sensemaking when performing a similar task. In the context of Software Engineering, other research projects primarily focused on assisting developers in completing sensemaking tasks, such as understanding unfamiliar code using in-situ code explanations, while not leveraging this virtuous cycle of sensemaking to assist in other cognitively-demanding aspects of software development. My projects, such as Adamite and Meta-Manager, show that such an approach can be adapted to an information and context-rich discipline such as software development given proper tooling and design.

As the nature of software development evolves with the proliferation of AI tools, there is a need to continue understanding the impact these tools have on developers. Much of my work has been inspired by understanding that human sensemaking, especially when completing a cognitively intense task such as programming, is messy and tooling solutions are often inadequate in terms of supporting the messiness of such information management. There are large unknowns in terms of how AI tools will change software development – my hope is that the creative parts remain in the hands of the developers and AI tools can act as supplementary aids that adhere and adapt to a developer's sensemaking when performing creative work. Such AI support can come in the form of adopting AI tools to the virtuous sensemaking cycle, where meta-information serves as training data to inform the tool about whatever contextual knowledge is necessary to support my development work. By keeping the developer in the loop, I believe we can create a better developer experience and more reliable code for the eventual end user.

Appendix A

Meta-Manager Architecture and History Model

The basic structure of Meta-Manager’s architecture and historical model is as follows:

- On VS Code launch, Meta-Manager attempts to match the project to a known entry in its database through using the GitHub ID (which follows the structure “PROJECT_OWNER/PROJECT_NAME”)¹.
 - If a match is found, Meta-Manager ingests all of the information it has for that project, which follows the structure FILE has many NODES which have many VERSIONS.
 - If no match is found, continue to next step.
- Meta-Manager traverses through the AST for each file that is a legal JavaScript or TypeScript file.
 - If Meta-Manager has found a match for the project, Meta-Manager attempts to match the nodes² it is finding within the AST to known nodes from the database.
 - * Matching nodes utilizes multiple heuristics including text matching between the last known version of the node and the current node, if the node is named (e.g., a non-anonymous function) and whether the names match, and how close the locations between node and database node match (i.e., does the node we are currently comparing have the same parent node(s) as the node from the database?).
 - For each node, either take the closest matching node’s already-unique ID or generate a new unique ID – ID’s include the name of the node, a colon, then a unique string created by the uuid package [250] (e.g., `recursiveSearch:ccaeaf0c-53bc-45dc-a76e-44a08a5a3a2f`).

¹We get the GitHub information using the GitHub API.

²Note that the nodes we discuss are nodes that are direct parents of block nodes (i.e., a left curly brace with code contained therein followed by a right curly brace) – technically, there are many other nodes, such as expression nodes and identifier nodes, that Meta-Manager does not perform this comparison on.

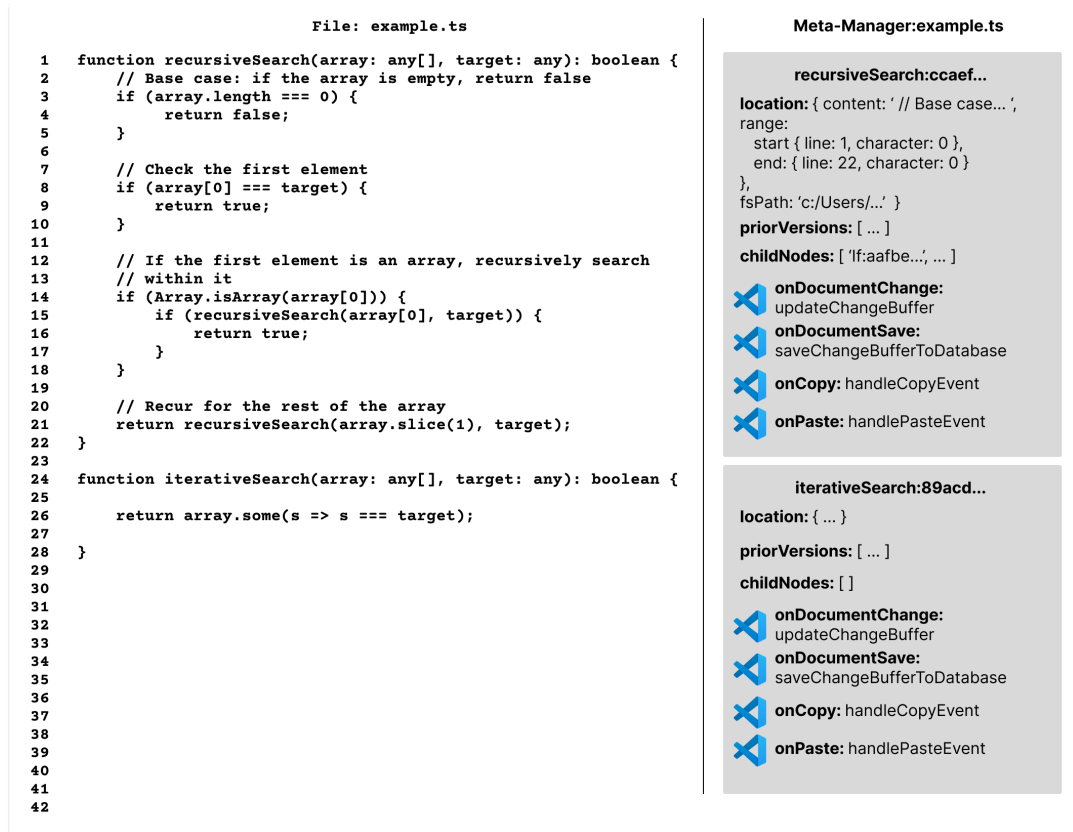


FIGURE A.1: A simplified version of how Meta-Manager’s data model and architecture look on a file called “example.ts”.

- For each node, instantiate the meta-information event listeners and functionality.
 - Some events are derived from native events within VS Code that each node can subscribe to. This includes:
 - * `onDocumentChange`: if some content within the bounds of the node’s location is edited, capture whatever change occurred and add it to the node’s change buffer as a new version.
 - * `onDocumentSave`: if the node has items in its change buffer, push those to the database.
 - * `onCopy`³: if some content within the bounds of the node’s location is copied, capture the code that has been copied and wait for paste event.
 - * `onPaste`: if a paste occurs within the bounds of the node’s location, capture the code that has been pasted and save any additional meta-information, including web-based meta-information or the copied node’s information.

³The VS Code API does not surface the copy or paste events so we use a workaround to emit an event when a copy or paste occurs – see this Stack Overflow post for more information: <https://stackoverflow.com/questions/44598894/in-a-vs-code-extension-how-can-i-be-notified-when-the-user-cuts-copies-or-paste>.



FIGURE A.2: How a version of a node appears in the Meta-Manager database on Firestore [61]. Note that the id of the version includes the node ID and the time at which the version was captured. This version does not include any additional meta-information, such as web activity.

- Other events are specific to Meta-Manager, such as transmitting information to the Meta-Manager pane when the node is selected. For full details on all Meta-Manager events, see the full implementation on GitHub: <https://www.github.com/horvathaa/meta-manager>.

Once the matching algorithm is completed, the Meta-Manager’s internal representation of the data is similar to Figure A.1.

A version of a node saved on the database would look like what is shown in Figure A.2.

Appendix B

MMAI GPT-4 Prompts

B.1 Pre-Processing History Prompt

The following prompt is generated for each `mmlog` version (see Figure 8.5 for an example of how such a `mmlog` version is formatted). The version is joined with the corresponding MMAI code version.

- `$codeLine` is the `mmlog` statement, e.g., Figure 8.5, `mmlog.location.content`.
- `$value` is the `mmlog` value, e.g., Figure 8.5, `mmlog.value`.
- `$codeVersionId` is the unique ID for a particular code version, e.g., Figure A.2, `id` or Figure 8.5, `versionId`.
- `$codeVersionCode` is the code content at a particular code version, e.g., Figure A.2, `location.content`.
- `$runId` is the unique run ID generated by the VS Code Debugging API (see Section 8.4.3), e.g., Figure 8.5, `runId`.
- `$logId` is the unique log ID for the unique `mmlog` statement on that code version, at that time (see Section 8.4.3), e.g., Figure 8.5, `key`.
- `$time` is when the log statement was invoked.

The prompt returns a summarized and indexed version of each `mmlog` statement execution, `$summary`, that is kept in that node's instance of GPT-4's memory.

B.1.1 Prompt

Summarize the output for the following code: `$codeLine`. Across time, this line of code has produced the following values below.

```
### $codeLine HISTORY ###
```

```
for each $value of $codeLine
```

```
The output came from code version $codeVersionId and the following code: $codeVersionCode.
```

```
The output was created on run $runId and generated the following output value:
```

```
$value.
```

```
The output came from code version $codeVersionId and the following code: $codeVersionCode.
```



```

The output was generated on run $runId.
The log ID was $logId.
The time was $time.
end for

```

B.2 Default User Query Prompt

The user query prompt utilizes the output of the Pre-Processing History Prompt, `$summary`, and formats and appends the user's query to it, `$query`, along with additional instructions on how the LLM's output should be formatted for post-processing.

B.2.1 Prompt

You are a code history and output oracle. A user is asking you this: `$query`. Use the provided history data below, under the header `### HISTORY DATA ###` to answer the user's question. Be succinct in your answer. Use code examples when necessary. If any code examples include backticks, replace the backticks with regular quote characters (i.e., `"`). If you provide a block code example, ensure that the beginning of the block code example begins with three back ticks and the end of the block code example ends with three back ticks. Do not discuss the semantics of the log function, `"mmlog"`.

```
### HISTORY DATA ###
```

```
$summary
```

```
### END HISTORY DATA ###
```

Include version IDs, run ids, and log ids under a header `### References ###` that you use when coming to an answer. Below `### References ###`, each referenced version ID should be under a header `### Version ID ###` and each ID should be on its own new line (i.e., no preceding dashes or bullets). Valid version IDs will follow the format `"someFunctionName:someUUID:someTimestamp"`.

Below `### References ###` and `### Version ID ###`, each referenced log should be under a header `### Log ID ###` and each log ID should be on its own new line (i.e., no preceding dashes or bullets). Valid log IDs will follow the format `"someFunctionName:someUUID:someCallToTheMMLOGFunction"`. When printing a log, print the log ID, followed by the run ID referenced for that log (which will be in the form `"someUUID"`), and the time at which the run happened. Example: `"someFunctionName:someUUID:someCallToTheMMLOGFunction on Run #someUUID - 7/24/2024, 2:41:41 PM"`.

If you use a version ID, run ID, or log ID in your answer, you must include it in the `### References ###` section. Ensure that the version ID, run ID, and log ID are accurate and correctly formatted. Do NOT wrap the IDs in backticks or quotes.

Provide 3 follow-up questions a developer may want to ask, given your answer, under the header `### Questions ###`. Each question should be on its own new line (i.e., no preceding dashes or bullets).

Appendix C

Chapter 3 Study Replication Materials

C.1 Preliminary Study Materials

These are the materials used for the preliminary API learning task discussed in Chapter 3, Section 3.2.

C.1.1 Protocol

Notes

This is the script that was used for each study session. Parts **highlighted in yellow** were instructions for myself to complete during the session. The preliminary study consisted of two “phases” – in Phase 1, participants were instructed to *author* annotations, while, in Phase 2, participants were instructed to *use* annotations.

Phase 1 Script

Hello, my name is Amber Horvath. We are interested in understanding the role of note taking during API learning tasks and how developers’ notes may benefit other developers. This study is composed of three parts - the first part will be a short interview about your background and current note-taking habits, if you have any habits at all, the second part will be an API learning task which will have you taking notes about what you learn about the API, and the third part will be another API learning task where you will have access to other developers’ notes. The second and third tasks will be done in the Python programming language using Apache Beam, a relatively new distributed data processing API.

Before we begin with the preliminary interview, I’d like to have you sign this consent form. Let me know if you have any questions. You may also notice there’s a question about recording your audio and your screen - please let me know if you are okay with that or not.

Hand out consent form with particular attention as to whether they’re ok being screen recorded or not

Okay, so we will now begin the interview.

Turn on audio recorder as part of Quicktime**Interview**

- What is your programming background?
- Have you ever been in any group software design sessions, such as during a class project or during an internship? **IF YES**
 - What sort of notes do you take during group meetings?
- What kinds of notes did you take and what did you take those notes with?
- While doing programming tasks, what notes do you take, if any?
- What kinds of notes would you consider these?
- Do you ever revisit your notes? If so, when/why?
- Do you ever share your notes? If so, when/why?
- Do you prepare the notes for sharing in any way?
- **Assuming they have talked about taking notes while programming** When taking notes in the past, did you ever encounter any issues with your workflow?
- Do you have any examples of notes you've taken?

Ok, thank you.

For the task, you will be using an off-the-shelf note-taking application called Hypothesis which allows you to highlight and annotate text and images on any website. **switch to chrome window with hypothesis** The task will involve you attempting to complete a data processing pipeline using Apache Beam. You will have access to some starter code, Apache Beam's documentation, and any other online web resource.

During this time, I would like you to try and write down as many notes as possible, even if this is not something you would normally do during a development task. Try and take notes that you imagine would be useful to yourself in the future or to another developer who is learning Beam. An example of a note may be an inconsistency you found in the documentation or a note highlighting a particularly useful Stack Overflow question or answer. Since we are using an off-the-shelf tool, I would also love for any of your thoughts on how to make an annotation tool like this more useful for development tasks.

Word Count Description

For this task, you will be writing a program using Apache Beam that will, given a text-based input, count how many times each word appears in the inputted document, remove any word that appears more than 5 times, and output a random selection of 10 words and their respective counts. The handling of the input, output, and running the pipeline has already been written - your portion of the task

will be coming up with the PTransforms to use and their respective order within the pipeline. We have also supplied some functions and classes in the script that may be helpful. The code you need to fill in is at line 102 in the provided script.

You will have 45 minutes for this task. Do you have any questions? **Answer questions** Okay, you may begin. **Switch to PyCharm, have Chrome open with Beam documentation**

Post Phase 1 Debrief

- What role did note-taking play during this task for you?
- Do you imagine these notes being useful to you in the future?
- Do you imagine these notes being useful to other developers learning Beam?
- What can you imagine being done to your notes in order to make them more helpful to yourself or other developers?
- Can you envision any changes to Hypothesis that would make note-taking easier and more beneficial for yourself during a development task?

Phase 2 Script

Hello, my name is Amber Horvath. We are interested in understanding the role of note taking during API learning tasks and how developers' notes may benefit other developers. This study is composed of three parts - the first part will be a short interview about your background and current note-taking habits, if you have any habits at all, the second part will be an API learning task which will have you taking notes about what you learn about the API, and the third part will be another API learning task where you will have access to other developers' notes. The second and third tasks will be done in the Python programming language using Apache Beam, a relatively new distributed data processing API.

For this task, I will have you use Hypothesis, an off-the-shelf annotation tool, as you complete an API learning task. You will be familiarizing yourself with Apache Beam, a distributed data processing API, to complete a data pipeline. You will have access to some starter code, Apache Beam's documentation, and any other online web resource.

You will also have access to annotations created using Hypothesis that another developer left that she found particularly helpful while learning Apache Beam. The annotations the user left may be viewed in a panel on the right side of the webpage. Each annotation that has been left has a tag - you may search for a particular tag by typing "tag:" and then the tag name in the search bar on the top of the screen. When reviewing these annotations, you can think of yourself as the second developer in this group and so you should also try and annotate the documentation with things that you have questions about, answers for, and other things you find particularly helpful. These annotations can be set to be viewable to a whole group of users, but

they may also be set to be viewed only by yourself, so feel free to annotate things that you feel would only be beneficial for yourself while also annotating things that may help a multitude of developers. While completing this task, also think about any potential modifications to a tool like Hypothesis that may make it more beneficial for software learning and development tasks.

In addition, we would like you to think aloud as you are working on this task. In particular, please try and ask aloud all of the questions you have about the API, its learning materials, and anything else. If you find an answer, please try and speak that aloud as well.

Word Count Description

For this task, you will be writing a program using Apache Beam that will, given a text-based input, count how many times each word appears in the inputted document, remove any word that appears more than 5 times, and output a random selection of 10 words and their respective counts. The handling of the input, output, and running the pipeline has already been written - your portion of the task will be coming up with the PTransforms to use and their respective order within the pipeline. We have also supplied some functions and classes in the script that may be helpful. The code you need to fill in is at line 102 in the provided script. You will have 45 minutes for this task. Do you have any questions? [Answer questions](#)

Okay, you may begin. [Switch to PyCharm, have Chrome open with Beam documentation](#)

Post Phase 2 Debrief

- What did you think of the annotations that were present in the documentation?
- Were any particularly helpful? Unhelpful? Distracting? Insightful? Inaccurate?
- How did you navigate the documentation with these annotations present?
- Do you think the presence of annotations made the documentation more or less useful to you?
- Did you add any annotations yourself? **IF YES**
 - What role did note-taking play during this task for you?
 - Do you imagine these notes being useful to you in the future?
 - Do you imagine these notes being useful to other developers learning Beam?
 - What can you imagine being done to your notes in order to make them more helpful to yourself or other developers?
 - Can you envision any changes to Hypothesis that would make note-taking easier and more beneficial for yourself during a development task?

C.1.2 Task Source Code

The source code that the participants had to modify as part of the task can be found at: <https://codesandbox.io/p/devbox/musing-lake-sphcmm>. The Code Sandbox must be configured with Apache Beam installed (run `pip install apache-beam`) and running on Python 2.7 to work.

C.2 Lab Study Materials

C.2.1 Adamite Authoring Condition Protocol

Hello, my name is Amber Horvath. We are interested in understanding how developers may learn a new API better while using documentation.

Research Study Overview

We are investigating the act of creating and reading annotations when using documentation. You can think of annotations like Google Doc comments or Microsoft Word Comments, in that they are generally a meta-level note attached to a particular place in a document. Like comments in Google Doc, sometimes the annotations would be just for you, and other times they are to be shared with others.

This study is composed of 2 parts - the first part will be a tutorial to get you familiarized with Adamite, our annotation tool. The second part will be a documentation comprehension task using our experimental tool. For the documentation comprehension task, you will use the JavaScript programming language with Piling.js, a library for aggregating and organizing visualizations. Once you've finished the second task, I will have you complete a short survey about your background and your experience using the tool.

Consent

Before we begin, I'd like to review this consent form. [Show consent form](#) Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form. ****if recording** I will begin recording this session.

Part 1: Tutorial

For the task, you will be using Adamite, our experimental annotation tool, which allows you to highlight and annotate text on any website. Here's a link to download

Adamite: [\[link to Adamite\]](#). Let me know once you've downloaded and unzipped the file named study-build.zip. ****wait** Would you mind sharing your screen?

Installation

****wait for participant to finish unzipping Adamite** To install Adamite, go to "<chrome://extensions/>" or click "Manage Extensions" on the puzzle icon. ****if developer mode isn't enabled** Toggle the "Developer Mode" button at the top right of the screen - when developer mode is enabled, the toggle should be blue and shifted to the right and you should see some options at the top of the screen. There, click the "load unpacked" button and select the unzipped folder, titled "build". You should now see Adamite in your list of extensions. ****check if pinned - if not pinned by default** Click the puzzle icon next to your URL bar (with the left mouse button) and pin Adamite by clicking the pin icon - it may already be automatically pinned, if so, the pin icon will be filled in. Sign Up Now, navigate to a new tab and open the link I just sent in the Zoom Chat (<https://atomiks.github.io/tippyjs/v6/all-props/>) You can press the Adamite badge by the Chrome URL bar to open the Adamite sidebar. You can resize the sidebar by placing your cursor between the end of the sidebar and the beginning of the webpage. Now, please create an account by using the email "a[#]@test.com" and the password "12345678" and hitting the "sign up" button. ****wait for participant to create account**

Adamite Extension Overview

Adamite's sidebar is the home for all of the annotations related to this page. The top part of the sidebar displays your username in the top right corner. Below that, there is a search bar for searching for annotations, a groups menu for showing your groups and the annotations shared between you and your group members, and a filter menu for filtering annotations, which we will discuss later.

Anchor Concept Overview - External vs On-Page

The main portion of the sidebar is the area below that, which is where any annotations on the page or that you have pinned will be displayed. Currently, there is one public annotation visible with an anchor, the selected text which has been annotated, that leads to an external website, the popper.js documentation, shown by the tilted anchor icon and a URL below the anchor text, and an anchor on this page with the text that says "All Props", shown with the upright anchor icon. Clicking on the external anchor icon that leads to popper's documentation will open that page in another tab. Clicking on the anchor icon that is on this page next to the text "All Props" will scroll to that part of the page.

Motivations for Using Annotations

Tippy and the main task's documentation, piling, are both single-page documentation sites but Adamite can also be used to connect multiple documentation pages with external anchors, as shown in this example. Annotations can serve as useful additional information about parts of the documentation, questions you or other developers had about the documentation, connections between related web pages as shown in the popper example, personal to-do items just for you, and more.

Introduction to Tutorial

Now, we will begin the tutorial to learn more about Adamite. In this tutorial, we will go through an exercise where you, a web developer, are using Tippy, a tooltip-creation package for React, to create tooltips in your React project while using Adamite. Note that you do not need to know React or Tippy for this tutorial.

Normal Annotation Task - Creating

As a developer learning Tippy, you decide to annotate the parts of the documentation that seem relevant to your task of adding a tooltip to a button that is on your webpage. You begin by annotating the code in the black box that is the call to `tippy()` as this code example shows how to create a Tippy object. To annotate, highlight the text on the webpage you want to annotate, in this case all the code in the black box. ****wait for participant to highlight** A pop-up should appear showing the different types of annotations you can create. These different types of annotations have different features for better managing your information needs and can also serve as a way of organizing the annotations you create. Whenever you have the sidebar open and select text, this pop-up will appear. Since we are making a general purpose annotation, click the “normal” type button. ****wait for participant to click**

Normal Annotation Task - Editing - Tags, Public/Private

Now, in the sidebar, you should see the start of your new annotation. The top part shows the text you selected while the rich text editor shows where you can add some additional information. In the editor, add the text “This is how I create my Tippy object”. ****wait** You can also add tags to your annotation such that you can organize annotations that are related to one another and filter to only show annotations that have the same tag. Add the tag “code” in the “Add a tag” field and press enter. ****wait** Normally, you could choose whether to publish this annotation privately, to a group of other Adamite users, similar to Google Drive, or publicly, such that anyone with Adamite can see the annotation, but for the purposes of this study, all annotations will be published privately. Now, click the “post as private” button. ****wait** The new annotation will appear in the sidebar. The text that was annotated will also be highlighted on the webpage.

Overview of Viewing Annotations on Sidebar

When you confirm the creation of the annotation, it will be collapsed, meaning it only shows the anchor, which is the text you annotated, and the text you added as an annotation. In this collapsed view, you get a quick summary of the annotation’s content. For an expanded view of the annotation that includes more information such as the tags on the annotation, you can expand the annotation by clicking the down arrow at the bottom of the annotation. Click the expand button. ****wait for them to click**

Pinning Annotations

Let’s say you, a developer learning Tippy, decide that you want to be able to access this annotation no matter what web page you’re on, as you are creating your tippy object and learning what you can do with it. You can pin an annotation such

that it's always accessible in a "pinned annotation" list, which can help with keeping track of important annotations. Try pinning your annotation by clicking the pin icon. ****wait** You should now see a button above the annotations that says "Show 1 Pinned Annotation". All of your pinned annotations, no matter what web page they are on or what filters you have in place, will be accessible in that list. Now, unpin this annotation by clicking the pin icon again.

Editing Annotations

Once you've created an annotation, you can also easily change it. You can edit an annotation's content and tags by clicking the pencil icon to start editing when the annotation is expanded. Suppose you want to add more text to this annotation, to remind you that you need to add props to your constructor - edit your annotation by adding the text "I need to pass in props" and click "Post as private" to confirm the edit. ****wait**

Creating TODO Annotation

Suppose you want your tippy object to have an animation when it appears. There's a prop called "inertia" - search for it on the page - if you find it in the bulleted list, click on it to get to the inertia section showing the code example with the black example code. ****wait** You see that, in order to add this animation, you need to not only pass in the prop "inertia" with a value set to true, but you also need to specify a CSS animation to play using the tippy-box class and the data-inertia attribute, as shown in the second black code example box. Highlight from where code example text says "inertia: true" through the CSS example underneath. ***wait*

Since you want to replicate this behavior on your website, create a to-do annotation to keep track of this code that says "Add this to my code" and click "post as private". ****wait** To-dos are automatically pinned so that they're always accessible and you do not lose track of what you are trying to do. Once you complete your task, you can unpin the annotation or mark the annotation as done with the "Done?" button to archive the annotation.

Creating Issue Annotation

In the CSS code example, you may notice that the animation is a cubic-bezier function. Move your mouse to hover over the box that says "animation: scale" below the code example. Notice that the CSS function is a scaling effect, not a bouncing effect. This makes it difficult to actually adapt this code example as they do not provide an example of how to use cubic-bezier but they also don't show the exact code for how they made the interactive example using the prop "animation: scale". This is a good chance to use the "issue" annotation to mark this example as confusing and possibly incomplete to help future users of the documentation know that this information is confusing. Highlight the text that says "cubic-bezier" and create an issue annotation. Annotate this with "make this match the interactive example below". ****wait**

This annotation has a "Flag for Expert Review" button which, in future versions of Adamite, will notify documentation writers of problems in the documentation.

Creating Question Annotation

Now you want to define the appearing and disappearing behavior of the tooltip. Search “hideOnClick” - there’s a prop for this. The prop’s description references a trigger event, but it’s not clear whether “trigger” is another prop, or something to reference in the hideOnClick event handler. This is a good time to use a question annotation to keep track of this confusing part of the documentation. Add a question annotation asking “how do I use this” on the word “trigger” in the description above the code example. Note that clicking “how do I use this” in the question drop down menu will auto-populate the annotation content with the text “how do I use this”. Create the annotation. ****wait** Notice that this new question annotation is pinned. Question annotations are pinned such that you can keep track of this question and, if you find an answer, you can link that page and its answer back to this page.

Answering Question Annotation

To answer our question, we search “trigger” on this page and see it shows up many times - click until you get to the trigger section showing the trigger example code. We can now answer our question by replying to the annotation with the answer and an anchor to the point that answered our question. Note that you can reply to any annotation, not just question annotations. Expand the question annotation. Click the arrow button left of the pin icon or the button that says “Unanswered Question” to open the reply editor. You can add an anchor to this reply by selecting the text that answered your question - in this case, the header for the “trigger” section - and clicking the “anchor plus” button at the bottom of the reply editor. ****wait** You can write in “trigger is a prop we pass in” as the reply and click “post answer”. ****wait** The annotation will now update with the answer. You can now scroll between the part of the page that had your question and the content which has your answer by clicking on the anchor icons left of the annotated text. ****wait for them to click**

Multiple Anchoring

We can build upon our question through adding more anchors. You want to create an event handler that is triggered when the tooltip appears. Search for “onTrigger” and go to the onTrigger section of the documentation. ****wait** onTrigger has a code example showing how to create an event handler that will fire when the tippy tooltip is triggered. You can add this as an additional anchor to your annotation. Select the whole onTrigger code example and click the “anchor plus” button next to the pin icon. ****wait** This functionality is useful for connecting multiple parts of the documentation for better navigability and other web pages that are relevant to what you are working on or confused about while programming.

Highlight Annotations

When there’s information you think may be useful or interesting but you don’t want to add any text to, you can use the “highlight” annotation. For tippy, it may be useful to remember that the API has a way of dealing with mobile devices using the prop “touch”. The prop “touch” is right above the “trigger” section - scroll to the “trigger” section by clicking your question annotation’s answer that is anchored

to the text that says “trigger”. Add a highlight annotation on the header that says “touch”. ****wait**

Filtering Annotations

At some point, you may feel there are a lot of annotations with only a subset being relevant to what you’re doing. One way of managing many annotations is through filtering. You can filter your set of viewable annotations by clicking on the various dropdowns by the text that says “Filters”. There’s a filter to only look at recently made annotations, a filter for the various annotation types, the sorting function we used earlier, and filtering by tags, which you can do by clicking on a tag on an annotation or through clicking on the button that says “Select Tag”. Try setting a filter such that only to-do annotations are viewable. ***wait for them to filter* Okay, now change the filter back such that all annotation types are viewable.

Part 2: Task

For the task, you will be familiarizing yourself with Piling, an image aggregation and organization library for JavaScript, and the Piling documentation. Imagine you are a developer in a small group that is learning about Piling. You are the first developer to begin using Piling, so you need to learn how to complete your programming task, but you also want to save everyone time by adding annotations that answer the questions you had and use annotations to mark and connect the important parts of the documentation. Assume that anything you are confused or wondering about, your teammates might be as well. Adamite does have a feature for small groups to share annotations motivated by this scenario, but, for the purposes of this study, you will be creating only private annotations.

During the study, you can start with Piling’s reference documentation which is at the link in the Zoom chat: <https://piling.js.org/docs/>. This documentation is far from perfect, so, during the study, I’d like you to annotate all of the questions you have and any answers you find to your questions. I’d also like you to annotate and connect parts of the documentation you feel are related, issues you find in the documentation, and parts of the documentation you want to follow up on. Feel free to annotate both things that you feel would only be beneficial for yourself and things that may help other developers who are trying to learn Piling.

For the programming task, you will be writing a program using Piling that will, given a set of inputted images, load them onto the screen, have images appear on the same row given their subject matter, and have images of the same subject matter have matching borders along the bottom of the image. All animal images should be on one row and all art images are on another row and have different colored bottom borders. Note that this effect should not be achieved through simply changing the order of elements in piling’s data array.

The instructions along with a photo showing the expected output are also provided in the Code Sandbox project, which I am sending in the Zoom chat now: ****wait for them to open** Have you used Code Sandbox before? ****if no** In Code Sandbox, the

left pane shows your code and the right pane shows the output web page. Every time you save your code, your web page preview will update. You can also open that web page in a new tab by clicking the rightmost button after the URL bar. There is also a console that can be opened by clicking the button that says “console” below the web page preview.

The Code Sandbox comes with piling already installed. We have provided a starting script that lists the steps you must complete. Every time you save the script, the outputted web page will refresh. There is also an HTML file and CSS file you may reference, along with a picture that shows the expected outcome. Please have CodeSandbox.io open in one tab and Adamite and the documentation open in another tab and continually share this window during the study. Remember that, while you should try your best to complete the programming task, we are primarily interested in the annotations you create while using Adamite.

In addition to annotating your thoughts about the documentation, we would like you to think aloud as you are working on this task, so I can understand what you are trying to do. In particular, please try and ask aloud all of the questions you have about the Piling library and its documentation, any issues you find in the documentation, parts of the documentation you find useful and want to follow up on, along with your thought process as you try and complete the task. If you begin to forget to annotate or stop thinking aloud, I will continually remind you to annotate and think aloud.

You will have 45 minutes for the task. Do you have any questions? ****wait** Okay, you may begin. ****begin screen recording**

Once task is over Times up! Please save your code and send me the URL of your CodeSandbox.io project in the Zoom chat. ****wait** I will now have you fill out the survey I sent in the Zoom chat (see Section C.2.4) - your Adamite username is a **[tell participant number here]**. Please let me know once you’ve finished the survey.

Once they say they’ve completed the survey The study is now over. If you would like to uninstall the tool, go back to the page <chrome://extensions> and find the Adamite extension. From there, you may click the button that says “remove”.

If you would like to keep Adamite, please be aware that we will still be collecting your user data. If, after 2 months from now, you still have the extension installed, we will notify you that your data is still being collected and will provide instructions on how to remove the application. This will continue until we begin to end this project. If you are still using the application as we begin ending this project, we will notify you 6 weeks in advance of the project ending such that you may prepare to remove the application. We plan to keep Adamite running for at least the next year. When we do plan to end the project, we will send your annotations in the form of a Google Sheet.

We will send your compensation to your email in the form of an Amazon gift card. Thank you for your time!

C.2.2 Adamite Using Condition Protocol

Hello, my name is Amber Horvath. We are interested in understanding how developers may learn a new API better while using documentation.

Research Study Overview

We are investigating the act of creating and reading annotations when using documentation. You can think of annotations like Google Doc comments or Microsoft Word comments, in that they are generally a meta-level note attached to a particular place in a document. Like comments in Google Doc, sometimes the annotations would be just for you, and other times they are to be shared with others.

This study is composed of 2 parts - the first part will be a tutorial to get you familiarized with Adamite, our annotation tool. The second part will be a documentation comprehension task using our experimental tool. For the documentation comprehension task, you will use the JavaScript programming language with Piling.js, a library for aggregating and organizing visualizations. Once you've finished the second task, I will have you complete a short survey about your background and your experience using the tool.

Consent

Before we begin, I'd like to review this consent form. Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form. ****if recording** I will begin recording this session.

Part 1: Tutorial

For the task, you will be using Adamite, our experimental annotation tool, which allows you to highlight and annotate text on any website. Here's a link to download Adamite: [\[link to Adamite\]](#). Let me know once you've downloaded and unzipped the file named study-build.zip. ****wait** Would you mind sharing your screen?

Installation

****wait for participant to finish unzipping Adamite** To install Adamite, go to "<chrome://extensions>". (or they can click "manage extensions" on the puzzle icon) Toggle the "Developer Mode" button at the top right of the screen - when developer mode is enabled, the toggle should be blue and shifted to the right and you should see some options at the top of the screen. There, click the "load unpacked" button and select the unzipped folder, titled "build". You should now see Adamite in your

list of extensions. ****check if pinned - if not pinned by default** Click the puzzle icon next to your URL bar with the left mouse button and pin Adamite by clicking the pin icon - it may already be automatically pinned, if so, the pin icon will be filled in.

Sign In

Now, navigate to a new tab and open the link I just sent in the Zoom Chat (<https://atomiks.github.io/tippyjs/v6/all-props/>). You can press the Adamite badge by the Chrome URL bar to open the Adamite sidebar.

You can resize the sidebar by placing your cursor between the end of the sidebar and the beginning of the webpage. Now, please create an account using the email “r[#]@test.com” and the password “12345678” and hit the “sign up” button. ***wait for participant to sign up**

Adamite Extension Overview

Adamite’s sidebar is the home for all of the annotations related to this page. The top part of the sidebar displays your username in the top right corner. Below that, there is a search bar for searching for annotations, a groups menu for showing your groups and the annotations shared between you and your group members, and a filter menu for filtering annotations.

Anchor Concept Overview - External vs On-Page

The main portion of the sidebar is the area below that, which is where any annotations on the page or that you have pinned will be displayed.

Motivations for Using Annotations

Currently, there are 5 public annotations visible in the sidebar. The first annotation that says “Tippy uses popper” has an anchor, the selected text which has been annotated, that leads to an external website, the popper.js documentation, shown by the tilted anchor icon and a URL below the anchor text. Its second anchor is on this page with the text “All Props”, shown by the upright anchor icon. Clicking on the external anchor icon leads to popper’s documentation and will open that page in another tab. Clicking on the anchor icon that is on this page next to the text “All Props” will scroll to that part of the page. If you hover over the text “All props”, this will highlight that annotation’s anchor in light blue. In the case where you have many annotations and the annotation that corresponds to the anchor is not visible, the annotation will be scrolled to in the sidebar.

Introduction to Tutorial

Now, we will begin the tutorial to learn more about Adamite. In this tutorial, we will go through an exercise where you, a web developer, are using Tippy, a tooltip-creation package for React, to create tooltips in your React project while using Adamite. Tippy and the main task’s documentation, piling, are both single-page documentation sites but Adamite can also be used to connect multiple documentation pages with external anchors, as shown in this example.

Annotations can serve as useful additional information about parts of the documentation, questions you or other developers had about the documentation, connections between related web pages, personal to-do items just for you, and more.

One of your teammates has left you some notes they took when they were learning Tippy so we want to leverage their knowledge so that we can learn Tippy quicker. Note that you do not need to know React or Tippy for this tutorial.

Understanding Annotations

We can begin the tutorial by looking at the annotations our teammate left us. Each annotation is collapsed by default, meaning it only shows the anchor and the text that was added as an annotation. In this collapsed view, you get a quick summary of the annotation's content. For an expanded view of the annotation that includes more information such as when the annotation was created, you can expand the annotation by clicking the down arrow at the bottom of the annotation. Click the expand button on the "Tippy uses Popper!" annotation. ****wait for them to click**

Understanding Annotations - Replying

Now that the annotation is expanded, we can see the name of the author of the annotation, in this case "tippy-dev", the group the annotation was shared with, in this case the "Public" group which includes everyone that uses Adamite, and the time at which the annotation was posted. To the right of the user name, we can see three buttons for interacting with the annotation. The leftmost arrow icon is the reply button which allows us to reply to an annotation. This may be useful for adding a follow-up question to an annotation, thanking the author for pointing out some helpful information, or for answering the author's question. Click the arrow icon to open the reply editor. ****wait for them to click**

In the reply editor, we can add some text in the rich text editor for our reply content. We can optionally add tags such that you can organize annotations that are related to one another and filter to only show annotations that have the same tag. You can also add an anchor to your reply to connect different parts of the documentation together and for adding additional context to your reply by selecting some text on the page and clicking the anchor plus button that is to the left of the cancel button. Once you're satisfied with your reply, you can click the Post Reply button or, if you decided the reply was not necessary, you can click the cancel button to exit out of the reply editor. Click the cancel button. ****wait for them to click**

Understanding Annotations - Pinning

The icon to the right of the reply icon is the pin icon. Pinning annotations causes the annotation to appear in a pinned annotation list, which is accessible on any webpage, and the annotation will ignore any filters in place. This is useful when you have an important piece of information you want to keep track of and potentially follow up on. Question type and to-do type annotations will be pinned by default when you create them. Let's say you, a developer, think the popper documentation is interesting and want to keep track of your teammate's annotation that references it. Try pinning the "Tippy uses Popper" annotation. ****wait** You should now see a button above the annotations that says "Show 1 Pinned Annotation". Any annotations you pin will be in that list. Now, unpin this annotation by clicking the pin icon again. ****wait**

Understanding Annotations - Multiple Anchoring

We can build upon annotations by adding more anchors. Look at the second annotation in the list that has two anchors. Click on the first anchor to scroll to the part of the page the annotation is on. `**wait` In this example, the two anchors serve to explain the user's question by showing that the default value for the delay prop is 0, but the 4th code example usage shows that the hide delay value is null instead of 0. This functionality is useful for connecting multiple parts of the documentation for better navigability and for adding additional context to your annotation. If you find an annotation that you think would benefit from an additional anchor, you can highlight part of the webpage and click the anchor plus button at the top of the expanded annotation, next to the pin icon.

Understanding Annotations - Different Types of Annotations

Now that we know what we can do with annotations, let's look at the annotations our teammate left us. In the sidebar, there are 5 annotations viewable. The first annotation in the list is a normal type annotation, which can be used for sharing additional information about the anchored text, and is denoted by the badge in the top right corner of the annotation with the 4 squares.

The second annotation in the list is an unanswered question annotation, shown by the bright blue question mark badge in the top right corner of the annotation. Question annotations are useful for keeping track of confusing parts of the documentation such that, if we find an answer to our question, we could help future users of the documentation and potentially connect the text or webpage that had our answer to our original question. If we knew the answer to our teammate's question, we could reply to this annotation with an answer.

The next annotation in the list shows an answered question annotation - our teammate didn't know how to use the "trigger" prop in Tippy, but, upon searching, found some example code showing how to use trigger and linked that answer to their original question through adding an anchor to the code example. This answer is then automatically appended by Adamite to the original question annotation's content. Answered question annotations are shown with a dark gray question mark badge in the top right corner.

Below that annotation, there is an issue annotation, which is useful for marking problematic parts of the documentation to help future users of the documentation know that this information is confusing and notify the documentation writers that there may be a problem in the text. In this case, the code example is incomplete and doesn't match the interactive example below the anchored text. This annotation has a "Flag for Expert Review" button which, in future versions of Adamite, will notify documentation writers of problems in the documentation.

Lastly, there's a highlight annotation that is used to highlight a part of the webpage without adding any additional text - these annotations are useful for marking important or potentially useful parts of the documentation for better navigability of

the document. These different types of annotations have different features for better managing your information needs and can also serve as a way of organizing the annotations you create or read.

Normal Annotation Task - Creating

Let's say you want to add more annotations to this webpage. As a developer learning Tippy, you want to annotate the parts of the documentation that seem relevant to your task of adding a tooltip to a button that is on your webpage. You begin by annotating the code in the black box that is the call to `tippy()` as this code example shows how to create a Tippy object. To annotate, highlight the text on the webpage you want to annotate, in this case all the code in the black box. ****wait for participant to highlight**

A pop-up should appear showing the different types of annotations you can create. Whenever you have the sidebar open and select text, this pop-up will appear. Since we are making a general purpose annotation, click the "normal" type button. ****wait for participant to click**

Normal Annotation Task - Editing - Tags, Public/Private

Now, in the sidebar, you should see the start of your new annotation. The top part shows the text you selected while the rich text editor shows where you can add some additional information. In the editor, add the text "This is how I create my Tippy object". ****wait** You can also add tags to your annotation such that you can organize annotations that are related to one another and filter to only show annotations that have the same tag. Add the tag "code" in the "Add a tag" field and press enter. ****wait** Normally, you could choose whether to publish this annotation privately, to a group of other Adamite users, similar to Google Drive, or publicly, such that anyone with Adamite can see the annotation, but for the purposes of this study, all annotations will be published privately. Now, click the "post as private" button. ****wait** The new annotation will appear in the sidebar. The text that was annotated will also be highlighted on the web page.

Filtering Annotations

At some point, you may feel your teammate left a lot of annotations with only a subset being relevant to what you're doing. One way of managing many annotations is through filtering. You can filter your set of viewable annotations by clicking on the various drop downs by the text that says "Filters" at the top of the sidebar. There's a filter to only look at recently made annotations, a filter for the various annotation types, a sorting function to sort annotation either by their location on the page or by when they were created, and filtering by tags, which you can do by clicking on a tag on an annotation or through clicking on the button that says "Select Tag". Try clicking on the "code question" tag to filter by that tag. ****wait for them to filter** Only the two question annotations that were marked with that tag will be viewable. Now change the filter back such that all annotations are viewable by clicking on the x button to the right of the tag name "code question" in the filter menu. ****wait for them to reset filter**

Searching for Annotations

Another way of managing many annotations is through searching to find the annotation you think will have the information you need. For searching, there are 3 modes - globally, meaning across all annotations made using Adamite, on this specific webpage, or on this website, meaning, in this case, any annotations made on atomiks.github.io, the Tippy documentation website. Let's try searching for our teammate's popper annotation - type the text "Popper" in the search bar. ****wait** You will see a preview of the annotation including the type of annotation, the anchor text, and the annotation text and it highlights what part of the annotation matched the search. Click on that annotation in the list. ****wait** Now that annotation will be the only annotation in the sidebar. You can also replace the sidebar's contents with all of the annotations that match your query by clicking the "enter" key after making a search. Clear your search result by clicking the x button in the search bar. ****wait**

Great, that's it for the introduction tutorial. We will now move on to the main task. If at any point during the task you forget how to use any of Adamite's features, feel free to ask me during the task.

Part 2: Task

For the task, you will be familiarizing yourself with Piling, an image aggregation and organization library for JavaScript, and the Piling documentation. I want you to imagine that you are a developer in a small group that is learning about Piling. One of your teammates has been learning Piling and has left some annotations that helped them while working on the task, questions they had while learning Piling, issues they found in the documentation, and more. You want to leverage the notes that they took while trying to complete your programming task and learning Piling yourself, while optionally adding annotations. When you encounter an annotation that is useful to you, I'd like you to speak aloud that the annotation is useful and why the annotation was helpful. Conversely, if an annotation is not helpful to you, I'd like you to speak aloud that the annotation was not helpful and explain why it was not helpful.

During the study, you can start with Piling's reference documentation which is at the link in the Zoom chat <https://piling.js.org/docs/>. This documentation is far from perfect, so, during the study, I'd like you to use the annotations that are present in the documentation. In addition to using the annotations in the documentation, I would like you to think aloud as you are working on this task, so I can understand what you are trying to do. In particular, please try and ask aloud all of the questions you have about the Piling library and its documentation, any issues you find in the documentation, parts of the documentation and annotations you find useful, along with your thought process as you try and complete the task. If you begin to forget to stop thinking aloud, I will continually remind you to think aloud.

For the programming task, you will be writing a program using Piling that will, given a set of inputted images, load them onto the screen, have images appear on

the same row given their subject matter, and have images of the same subject matter have matching borders along the bottom of the image. All animal images should be on one row and all art images are on another row and have different colored bottom borders. Note that this effect should not be achieved through simply changing the order of elements in piling's data array. The instructions along with a photo showing the expected output are also provided in the Code Sandbox project, which I am sending in the Zoom chat now.

****wait for them to open** Have you used Code Sandbox before? ****if no** In Code Sandbox, the left pane shows your code and the right pane shows the output webpage. Every time you save your code, your webpage preview will update. You can also open that webpage in a new tab by clicking the rightmost button after the URL bar. There is also a console that can be opened by clicking the button that says "console" below the webpage preview.

The Code Sandbox comes with piling already installed. We have provided a starting script that lists the steps you must complete. There is also an HTML file and CSS file you may reference, along with a picture that shows the expected outcome. Please have CodeSandbox.io open in one tab and Adamite and the documentation open in another tab and continually share this window during the study. Remember that, while you should try your best to complete the programming task, we are primarily interested in how you use the annotations in the documentation.

You will have 45 minutes for the task. Do you have any questions? ****wait** Okay, you may begin. ****begin screen recording**

Once task is over Times up! Please save your code and send me the URL of your CodeSandbox.io project in the Zoom chat. ****wait** I will now have you fill out the survey I sent in the Zoom chat (see Section C.2.4) - your Adamite username is r[tell participant number here]. Please let me know once you've finished the survey.

Once they say they've completed the survey The study is now over. If you would like to uninstall the tool, go back to the page <chrome://extensions> and find the Adamite extension. From there, you may click the button that says "remove".

If you would like to keep Adamite, please be aware that we will still be collecting your user data. If, after 2 months from now, you still have the extension installed, we will notify you that your data is still being collected and will provide instructions on how to remove the application. This will continue until we begin to end this project. If you are still using the application as we begin ending this project, we will notify you 6 weeks in advance of the project ending such that you may prepare to remove the application. We plan to keep Adamite running for at least the next year. When we do plan to end the project, we will send your annotations in the form of a Google Sheet.

We will send your compensation to your email in the form of an Amazon gift card. Thank you for your time!

C.2.3 Control Condition Protocol

Hello, my name is Amber Horvath. We are interested in understanding how developers may learn a new API better while using documentation.

Research Study Overview

This study is composed of 2 parts - the first part will be a tutorial to get you familiarized with speaking aloud when using documentation. The second part will be a documentation comprehension task. For the documentation comprehension task, you will use the JavaScript programming language with Piling.js, a library for aggregating and organizing visualizations. Once you've finished the second task, I will have you complete a short survey about your background and your experience doing the task.

Consent

Before we begin, I'd like to review this consent form. Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form. ****if recording** I will begin recording this session.

Part 1: Tutorial

Framing

In this tutorial, we will be practicing thinking aloud as a way of training for the main task. Since we are primarily interested in your thought process while learning an unfamiliar API and how you use its documentation while trying to complete a task, I want you to practice speaking aloud all the questions you have about the documentation, the connections you see between different parts of the documentation that are relevant to your task, problems with the documentation, and other important information you learn about the API and its documentation. In this tutorial, we will go through an exercise where you, a web developer, are using Tippy, a tooltip-creation package for React, to create tooltips in your React project. Open the link I sent you in the Zoom chat in a new tab and share your screen: <https://atomiks.github.io/tippyjs/v6/all-props/>.

You, as a new developer learning Tippy, are reading through this documentation to gain an understanding of Tippy. Note that you do not need to know React or Tippy for this tutorial.

Recognizing Useful Parts of the Documentation

As a developer learning Tippy, you begin by learning how to create a Tippy object. At the top of the screen there is a code example in the black box that is a call to `tippy()` showing how to create a tippy object - the first step to adding a tooltip to your webpage. Since this is a necessary piece of information to complete your goal of adding tooltips to your React project, I'd like you to practice speaking aloud that this part of the documentation is useful. ****wait for participant to speak aloud** Great.

Asking Questions

Now that you know how to create a tippy object, you want to define the appearing and disappearing behavior of your Tippy tooltip. Try speaking this goal aloud. ****wait** Now, search "hideOnClick" - there's a prop for this. The prop's description references a trigger event, but it's not clear whether "trigger" is another prop, or something to reference in the hideOnClick event handler. Try speaking aloud this question you have about how to use "trigger". ****wait**

Answering Questions

Since we want to understand how to use trigger, search for "trigger" on this page. ****wait** "Trigger" shows up many times - click until you get to the trigger section showing the trigger example code. We can now answer our question by seeing that the trigger referenced in the hideOnClick prop is another prop we need to pass in to our call to tippy. Try speaking this answer aloud. ****wait**

Recognizing Connected Parts of the Documentation and Speaking Aloud What to Do

Our original goal was to define the appearing and disappearing behavior of our Tippy tooltip. Since we answered our question of how trigger works, we are closer to that goal, but we still need to add an event handler that gets called when the tippy tooltip is triggered. Search for "onTrigger" and go to the onTrigger section of the documentation. ****wait** onTrigger has a code example showing how to create an event handler that will fire when the tippy is triggered. Since this event handler is relevant to your goal, I'd like you to practice speaking aloud that this part of the documentation is related to your goal and to the other components we've looked at. ****wait**

Recognizing Issues in the Documentation

Now you want your tippy tooltip to bounce when it is triggered. There's a prop called "inertia" that makes the tooltip bounce - search for "inertia" on the page. ****wait** In the CSS code example, you may notice that the animation is a cubic-bezier function. Move your mouse to hover over the box that says "animation: "scale". Notice that the CSS function is a scaling effect, not a bouncing effect. This makes it difficult to actually adapt this code example as they do not provide an example of how to use cubic-bezier but they also don't show the exact code for how they made the interactive example using the prop "animation: scale". This is an issue in the documentation - I'd like you to practice speaking aloud that you identified an issue in the documentation. ****wait**

Part 2: Task

For the task, you will be familiarizing yourself with Piling, an image aggregation and organization library for JavaScript, and the Piling documentation. Imagine you are a developer in a small group that is learning about Piling. You are the first developer to begin using Piling, so you need to learn how to complete your programming task, but you also want to save everyone time by sharing what you learned to your teammates. Assume that anything you are confused or wondering about, your teammates might be as well. During the task, I'd like you to do whatever you would naturally do to make sure you are able to share what you learned with your teammates, whether that be taking notes, leaving comments in the code, or some other mechanism.

During the study, you can start with Piling's reference documentation which is at the link in the Zoom chat <https://piling.js.org/docs/>. This documentation is far from perfect, so, during the study, I'd like you to speak aloud all of the questions you have about Piling and its documentation and any answers you find to your questions. I'd also like you to speak aloud the parts of the documentation you feel are related in completing your task. Feel free to speak aloud both things that you feel would only be beneficial for yourself and things that may help other developers who are trying to learn Piling. If you begin to stop thinking aloud, I will continually remind you to think aloud.

For the programming task, you will be writing a program using Piling that will, given a set of inputted images, load them onto the screen, have images appear on the same row given their subject matter, and have images of the same subject matter have matching borders along the bottom of the image. All animal images are on one row and all art images are on another row and have different colored bottom borders. Note that this effect should not be achieved through simply changing the order of elements in piling's data array.

The instructions along with a photo showing the expected output are provided in the Code Sandbox project, which I am sending in the Zoom chat now ****wait for them to open** Have you used Code Sandbox before? ****if no** In Code Sandbox, the left pane shows your code and the right pane shows the output webpage. Everytime you save your code, your webpage preview will update. You can also open that webpage in a new tab by clicking the rightmost button after the URL bar. There is also a console that can be opened by clicking the button that says "console" below the webpage preview.

The Code Sandbox comes with piling already installed. We have provided a starting script that lists the steps you must complete. There is also an HTML file and CSS file you may reference, along with a picture that shows the expected outcome. Please have CodeSandbox.io open in one tab, and the documentation open in another tab. Remember that, while you should try your best to complete the programming task, we are primarily interested in your thought process as you complete the task.

You will have 45 minutes for the task. Do you have any questions? ****wait** Okay, you may begin. ****begin screen recording**

Once task is over

Times up! Please save your code and send me the URL of your CodeSandbox.io project in the Zoom chat. ****wait** I will now have you fill out the survey I sent in the Zoom chat (see Section F.2) - your participant number is c[**tell participant number here**]. Please let me know once you've finished the survey. **Once they say they've completed the survey** The study is now over. We will send your compensation to your email in the form of an Amazon gift card. Thank you for your time!

C.2.4 Post-Task Survey

Please fill out the following scale questions about your thoughts on the usability and usefulness of Adamite. The following page will have some additional questions about your background. The last page will ask about your experience using Adamite.

1. What is your Adamite username? **[required, free response]**
2. For the following 7 statements, rate how much you agree or disagree with the provided statement. **[required, each statement is scored on a 7-point Likert from "Strongly Disagree" to "Strongly Agree".]**
 - I consider Adamite easy to use and easy to achieve what I want
 - I consider my interactions with Adamite to be understandable and clear
 - I consider it easy for me to learn how to use Adamite
 - I would consider Adamite useful for my daily work
 - I enjoyed the features provided by Adamite
 - If possible, I would like to have the tool available for my professional or personal use
 - If possible, I would recommend Adamite to my friends and colleagues doing programming work

Background Questions

Please answer these questions about your programming background.¹

1. Are you familiar using comments in tools such as Google Docs, Microsoft Word, or Overleaf? **[required, yes/no/other]**
2. What notes do you take, if any, when programming? **[required, free response]**

¹For the control condition, their survey solely comprised this section, along with a question asking what their participant number was.

3. What notes do you take, if any, when doing tasks that are not programming? [required, free response]
4. What is your profession? (e.g., student, software engineer, manager, etc.) [required, free response]
5. How many years have you been programming? [required, free response]
6. How many years of professional programming experience do you have? [required, free response]
7. What is your level of expertise in using JavaScript? [required, 7-point scale from “No experience” to “Expert”]

Short Answer Questions

Please write a free response to each question about your experience with Adamite, the annotation tool you used, and your note taking habits.

1. What did you think about the tool, Adamite, in general? [free response]
2. In what situations do you think you would most want an annotation tool like Adamite? [free response]
3. Are there any features of Adamite that you found particularly good? [free response]
4. Are there any features of Adamite that you found particularly confusing? [free response]
5. What new features would you want in an annotation tool? [free response]
6. How do you normally keep track of important information? [free response]
7. Have there been any situations where taking notes or keeping track of information using other strategies have not worked well? If so, please describe. [free response]

C.2.5 Task Source Code

A completed version of the source code that the participants had to modify can be found at: <https://codesandbox.io/p/sandbox/piling-project-script-9rwug>. The version participants would see at the beginning of their session is the exact same, except `index.js` would only have the code comments and not the actual filled-in JavaScript code.

Appendix D

Chapter 4 Study Replication Materials

The following contents are the lab study materials for Chapter 4's lab study. Parts of the protocol which are instructions, notes or are otherwise not said aloud are **highlighted in yellow**.

D.1 Catseye Condition Protocol

Hello, my name is Amber Horvath and I am a fourth-year Ph.D. student in the Human-Computer Interaction Institute at Carnegie Mellon University, working under Dr. Brad Myers. Research on developers has shown that they often need to keep track of lots of information. We are interested in understanding more about this need while completing a programming task.

To support developers' information needs, we are investigating the act of creating and reading annotations when programming. You can think of annotations like Google Doc comments or Microsoft Word Comments, in that they are generally a meta-level note attached to a particular place in a document. Like comments in Google Doc, sometimes the annotations are just for you, and other times they are to be shared with others.

This study is composed of 2 parts - the first part will be a tutorial to get you familiarized with Catseye, our annotation tool. The second part will be a programming task using our experimental tool. For the programming task, you will use the JavaScript programming language and the Visual Studio Code editor augmented with the Catseye plugin. Once you've finished the second task, I will have you complete a short survey about your background and your experience using the tool and completing the study.

D.1.1 Consent

Before we begin, I'd like to review this consent form. Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me

know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form.

****if recording** I will begin recording this session.

D.1.2 Part 1: Tutorial

For the task, you will be using Catseye, our experimental annotation tool, which allows you to annotate your code in the Visual Studio Code editor. Before we can use the Visual Studio Code extension, we need to create an Catseye account. Would you mind sharing your screen? ****wait for screen share**

Create a Catseye Account

Do you already have a Catseye account? ****if no** To create a Catseye account, go to <https://adamite.netlify.app/login>. Then, choose "Login with Google". Once you have logged in, you should be redirected to the Catseye website and logged in to your new Catseye account. ****wait for them to login** Now, click your profile icon and click the "Link GitHub" option from the dropdown. This should open a pop-up where you can sign in and authorize your GitHub account to be used with Catseye.¹

****if yes** We will need to connect your Catseye account to your GitHub account. In order to do that, log in to your Catseye account at adamite.netlify.app/login. Then, once you are logged in, click your profile icon in the top right corner of the screen, and click "Link to GitHub". This should open a pop-up that will prompt you to log into your GitHub account, if it is not already logged in.

Now, please download the Catseye Visual Studio Code extension. Here's a link to download Catseye: [\[link\]](#). Let me know once you've downloaded the file named Catseye.vsix. ****wait**

Installation

****wait for participant to finish downloading Catseye** To install Catseye, first open Visual Studio Code.

Open the extensions pane - you can do this by either using the keyboard shortcuts `ctrl/cmd + shift + x` or by clicking the Extensions button on the left side panel.

Click the three-dot menu in the top right corner of the Extensions side panel. Then, click the "Install from VSIX" menu option - this should open a file explorer window. From there, find and select the Catseye.vsix file we just downloaded and click "Install". This installation may take a while. . . ****wait** While we are waiting, we can clone the repository we will use for the tutorial task. Do you know how to clone

¹Note that this whole log-in flow and installation was greatly simplified after this study.

a Git project? ****if yes** Okay, great, please create a clone of this project and open it in Visual Studio Code <https://github.com/ShauryaBhandari/Website-Templates>

****if not** Okay, copy this GitHub link: <https://github.com/ShauryaBhandari/Website-Templates>. Then, in Visual Studio Code on the Get Started page, click the “Clone Git Repository” link. This should open a text input window at the top of the screen. Paste the link that you just copied into that window and click “Enter”. Then select the location you want this folder to be in - for the purposes of this study, it does not matter where you clone the repository and you can delete this folder once we are done with the study. Once you clone the repository, VS Code should have a pop-up that says “Would you like to open the cloned repository?” in the bottom right corner of the screen - click “Open”.

****once Catseye has finished installing - should prompt for GitHub auth** Okay, it looks like it’s done. Let’s first check that Catseye is in our extensions list - in the Extensions side panel, it should appear in the “Installed” list. ****if there** Okay, now the extension should prompt you to authenticate the app with your GitHub account. ****if not there** Okay, let’s quit and restart Visual Studio Code and then the extension should ask to authenticate with GitHub.

Once GitHub opens in a browser tab, sign in to your account or, if already signed in, follow the instructions for signing in to GitHub with VS Code. These steps should end with you re-launching VS Code. Note that this is a one-time requirement - next time you launch VS Code, you will be logged in and will stay logged in. ****wait** Congratulations, you got Catseye installed!

Open Catseye

Let’s open the project we cloned - open the Website-Templates folder in VS Code. ****wait** Now that we have a project open, we can use Catseye. In this tutorial, we will go through an exercise where you, a web developer, are using this repository of website templates as a starting point for writing some code for your own website.

To open the Catseye pane, you can right-click inside the editor and click the “Catseye: Launch Catseye” item - you can also open Catseye by using the keyboard shortcut `ctrl/cmd + shift + a`. Catseye should open in a separate pane alongside your code. If you want to move Catseye, you can drag the Catseye tab to the right or left until you see a light gray shadow showing where the pane will go. ****wait**

Making our First Annotation

With Catseye open, we can create our first annotation. Let’s open the `index.html` file in the `bizland` folder ****wait** To make an annotation, select the code you want to annotate - in this case, let’s annotate line 20 where the HTML says “title BizLand Theme close title”. ****wait for participant to highlight this line** Next, right click the selected code - in the context menu, you should see 3 Catseye annotating options: “Catseye: create annotation”, “Catseye: create highlight”, and “Catseye: created

pinned annotation”, along with the “Catseye: Launch Catseye” option. Note that there are also keyboard shortcuts to execute these commands - (if Windows) ctrl + alt + a to create an annotation (if Mac) option + cmd + a, ctrl/cmd + alt + h to make a highlight annotation, and ctrl/cmd + alt + p to create a pinned highlight annotation. You can always find these keyboard shortcuts in the context menu when you have some code selected or in the list of keyboard shortcuts which can be found in the File menu, in case you forget. Let’s make an annotation - either click “Catseye: create annotation” from the context menu, or use the ctrl + alt + a keyboard shortcut to make an annotation. ****wait**

Now, in the Catseye pane, you should see the start of your new annotation. The top part shows the code you selected, while the text editor shows where you can add your own thoughts, questions, commentary or other additional information to this code.

Something that may make sense to keep track of through annotations is open tasks you want to work on – in this case, we want to adapt this code to make our own website, so let’s annotate that. In the text box, add the text “change this to a better title”. When creating our annotation, we also have the choice to post the annotation to our collaborators or just for ourselves - for the purposes of this study, we will be posting every annotation privately. We can also choose whether or not to pin an annotation when creating it - for now, we will not pin this annotation but will discuss pinning in more depth later. Now, click “Post as Private” ****wait**. The new annotation will appear in the Catseye pane and the annotated code will have a light gray box around it in the editor. Annotation locations in the code can also be seen in light green in the scroll bar area.

Annotations in the Editor

Catseye annotations can be used to keep track of old versions of the code and to see your thoughts about the code in-context. Let’s hover over the light gray box on line 20. ****wait** You should see the content of your annotation, in this case “Change this to a better title”, along with a link to scroll to your annotation. Clicking on the link will scroll to the annotation in the Catseye pane and highlight it - try clicking on the link now. ****wait** Suppose you actually do change the title of the page - try changing the code from `<title>Bizland Theme</title>` to `<title>My New Website</title>` then save your code. ****wait and see if Catseye pane updates - if it doesn’t prompt them to click the pane to bring in to focus.** The Catseye pane should update with the new version of your code whenever you save your code, so make sure to save often! If you want to see the original code you annotated, click the right arrow to the right of the code - the original version of your code will always be saved. ****wait** You can then go back to the most recent version of the code by clicking the left arrow. ****wait** Catseye’s annotations can be useful for keeping track of your thoughts in-context about the code, and will retain that original context if you change your code.

Keeping Track of Open Questions and Hypotheses

Another use for annotations is to help you keep track of open questions and hypotheses about the code. Now, in the bizland index.html file, let's scroll down to line 316. There, you should see a comment that says "Import JQuery before materialize.js" – the original author of this code does not explain why this is necessary or what problems may arise if you do not import jquery before materialize.js. Since we are hoping to adapt this code to our own needs and may want to import other libraries and need to change the order of the imports, let's mark this part of the code as confusing with an annotation. Select the comment on line 316 `**wait` and create an annotation asking "why do the imports need to be in this order?" `**wait` Annotations are a lightweight way of keeping track of these open questions you may have about unfamiliar code without making comments in the source code that will need to be removed later.

One potential reason for this ordering is that materialize.js may depend on jquery in order to work and since we aren't using npm or another package management system to control our dependencies, we need to do this checking manually. Add a reply to your annotation saying "materialize may depend on jquery" - you can make a reply by clicking on the speech bubble button. `**wait` Replies are a quick and easy way to follow up on the content of your annotation. You can see your reply by clicking on the "Show 1 reply" button. `**wait`

Now, let's see what happens if we change the order of the imports - move line 321, the line with src set to 'js/materialize.min.js', up to line 317, before the jquery import. `**wait` Now save, and open index.html in your browser. `**wait` Once you have index.html open, look at the console message by right-clicking the browser window and selecting "Inspect" or by clicking F12. `**wait` Notice that, in the console, there's an error message saying that \$ is not defined and this error is being thrown by materialize.min.js. This confirms our hypothesis that materialize depends upon JQuery - let's edit our annotation to reflect this fact - press the pencil icon in the top right corner of the annotation, and add the text "confirmed - materialize depends on jquery". `**wait` Now, let's move the code back into the right order - put the materialize.min.js script load below line 327. `**wait` Annotations can be an effective way to keep track of your open questions and your eventual answers to those questions.

Keeping Track of What You Have Learned

Annotations can also be an effective way to keep track of many different kinds of things you learn about the code, but might not want to or be allowed to add as regular comments. Let's scroll to the bottom of the file - lines 352 and 353. `**wait` There, we can see some latitude and longitude coordinates, but it is not immediately clear where those coordinates are pointing to. Let's try putting those coordinates into Google - copy the coordinates on line 353 and paste them into Google. `**wait` Now we can see that these coordinates are in the Boston city area. Let's annotate those

coordinates and say “these coordinates point to the Boston area - we can change them if we want our marker to be elsewhere”. ****wait** Good. If you are learning lots of things in an unfamiliar code base, it may be helpful to create annotations to keep track of these thoughts.

Lightweight Versioning and Capturing System Output

Catseye’s annotations can also be used as a space to keep track of changing system output. Let’s take a look at another website project - open the file called main.js in the Edgedger folder, under js. ****wait** We can see the author is using JQuery to do some things with scroll, but what if we wanted to change the behavior of scrolling? Let’s add an annotation to line 30 on the comment “//Smooth Scrolling” that says “How does this work?” ****wait** Good.

First, let’s see what “hash” is - add a console log below line 35 that prints the value of hash. This will also give us insight into when this function is called. Open index.html in your browser and open your console. ****wait** If we click any of the buttons on the nav bar at the top of the screen, we should see that the console prints out the ID name we are scrolling to. Let’s copy and paste this output ****wait** then go back to our annotation and paste the output as a reply to our annotation. ****wait** Good. Annotation replies are a nice place to add additional information to our annotation, including output values.

If we want to associate specific output with a specific version of our code, it may make sense to save a version of our code then edit the snapshot with the output that version of the code generated. If we look at line 39, we can see that an “animate” function is being called and the value “scrollTop” is being set to some number related to the top of the hash element we investigated before. If we go back to the website and click the nav bar buttons, we can see that the element scrolls but it isn’t very nicely centered on the screen and is too high on the screen – let’s try fixing this. Create an annotation where scrollTop is set that says “improve this”. ****wait** Good.

To fix this, let’s try some different values in place of the “100” they are using. Let’s start with replacing 100 with 500. Let’s also add a console log with the value that we are computing – copy the code that scrollTop is being set to into a console log statement above the animate call. ****wait** Now, save and reload and try clicking one of the buttons. 500 seems to make the element too low on the screen - copy the number we computed and go back to VS Code. ****wait** Now, snapshot your code and edit the snapshot with your copied number – add the text “too low on the screen”. ****wait** Good. Let’s try changing to something in the middle - change 500 to 250. ****wait** Now save and go back to the website – this looks better. Copy and paste the output into a new snapshot of the final version of the code. ****wait** Great - snapshots can be a good place to keep track of which version of your code produces what output.

Further Building Upon and Contextualizing an Annotation

Once we have an annotation, there are multiple ways to build upon it to help you better keep track of your thoughts about the code. One way is to add an additional anchor to an annotation - this may be useful if you find yourself continually going back and forth between locations in the code base, whether that be across files or within one file, or if you want to better contextualize your annotation comment, in the case that it is relevant in multiple places in the code base.

An annotation we made earlier about coordinate locations for Google Maps is relevant in this file too – if we scroll to line 4, you can see that they set some latitude and longitude coordinates. Let’s add this as an anchor to our previous annotation. We can find our other annotations in the “Current Project” section of the Catseye pane – find your latitude and longitude annotation in that section. ****wait** Now select lines 4 through 7, then click the “anchor” button in the top right corner of the annotation. ****wait** You should see your annotation update with another anchor below the original latitude and longitude coordinates. Multiple anchors is one way of creating links between related code in different parts of the code base.

Since this anchor is rather long, it can take up a lot of the Catseye pane - let’s make it smaller by clicking the up arrow directly to the right of the filepath ****wait** - this will truncate the code to only show the first line – even in the collapsed state, you can still navigate to the code by clicking the code or file path in the annotation. You can expand the code by clicking the down arrow to the right of the file path. If the annotation is still taking up too much space in the Catseye pane, you can collapse all of the anchors, such that you can only see the first one – try clicking the up arrow that is below both of the anchors. ****wait** Now you should only see the first code snippet – like when shrinking a single anchor, you can expand the list again by clicking the down arrow.

Navigating Using Annotations

Annotations not only serve as a way of externalizing your thoughts without modifying the source code, but also can serve as prominent navigational aids and quick links to parts of the code you care about.

Annotations in the Catseye pane are sorted by their location in the file system - note that your annotation is in a section called “Current File”. Once you create more annotations across more files, they will either be in the current file, current project, or pinned. Once you use Catseye in multiple projects, you will see your annotations in those projects.

You can navigate to an annotation’s location by clicking the code in the annotation or the filepath. Try clicking your annotation’s code that says “scrollTop”. ****wait** Since you already have this file open, Catseye scrolls the open-editor to the location of the code. Now, try clicking the code latitude and longitude coordinates from the bizland index.html file. ****wait** This opens the bizland index.html file. Note that your

scrollTop annotation is now in the Current Project list, as opposed to the Current File list. If you begin having many annotations in your project, you can hide the annotations in the Current Project list, or any other Catseye pane list, by clicking on the text that says “Current Project” - try closing the list. ****wait** Now open it again by clicking the Current Project text again. ****wait**

If you find yourself continually going to the same part of the code and want an even quicker way to navigate there, it may make sense to create a pinned annotation at that location. You can pin an annotation you’ve already created by clicking the “pin” icon at the top right of the annotation - let’s pin our latitude and longitude annotation. ****wait** Good. You can also create a pinned annotation using the keyboard shortcut ctrl + alt + p or by clicking the “pin annotation” checkbox when creating a new annotation.

Your pinned annotations are in the pinned annotation list and are always accessible there, no matter what project you are working on or what file you have open. With a pinned annotation, you can quickly jump to its location by using the keyboard shortcut ctrl + alt + j. ****wait** Since this annotation has multiple anchors, using the navigation shortcut will jump between the two anchors in this annotation - try clicking ctrl + alt + right square bracket again. ****wait** If you have a set of related code snippets you want to easily navigate between, creating a pinned annotation with multiple anchors is an effective way of doing this.

When we have multiple pinned annotations, we can navigate between the anchors in each pinned annotation. Let’s try pinning the first annotation we made - the My New Website annotation. ****wait** Now, try navigating between the two annotations using the ctrl + alt + right square bracket keyboard shortcut. ****wait** Good.

“Finishing” an Annotation

Once you’re done with an annotation - whether that means you no longer need its content, or you are no longer working on that task anymore, you can delete an annotation by clicking the trash can icon. Try deleting your first annotation about changing the website title. ****wait** Note that the highlight is also removed from the editor. Deleting annotations can help with managing an overabundance of annotations.

In the case where you do have many annotations, you can search to find your annotations – search looks at the annotation content, the code, and the filename, so you can, for example, find all of your annotations made on CSS files or find all the annotations you annotated with the word “todo”. Try searching for “scrollTop”. ****wait** Good.

That is it for the tutorial. Do you have any questions about Catseye? ****wait** Okay, let’s move on to the main task.

D.1.3 Part 2: Task

Now, for the main task, you are going to be using Catseye to help you keep track of your thoughts, questions, comments, and concerns about some unfamiliar code, along with any issues or areas for improvement you find in the code.

I would like you to imagine that you are a new team member on a programming project and you have been tasked with understanding the code, along with finding and fixing bugs in the code. You do not want to edit your teammates' code until you have tested it – thus, for the first 15 minutes of the study, you are not allowed to edit the actual logic of the code, but may annotate the code with what you learn about the code through testing, what bugs you find, areas that can be improved, and any other pieces of information you want to keep track of. You can also run the code, use Google, add print statements and comments, or do anything else you would normally do when testing code, aside from actually editing the code that is already in the repository.

After the 15 minutes, you will have 30 minutes to use your annotations to help you improve upon and debug the code. While you are working on this task, I would also like you to think aloud about what you are doing, how you plan on changing the code, and please speak aloud whenever you discover a bug and whenever you believe you've addressed the bug. I will notify you once your first 15 minutes of reviewing and testing the code are complete, and when your last 30 minutes of debugging are complete.

Do you have any questions? ****wait**

Now, I'd like you to create your own local fork of this repository. <https://github.com/horvathaa/task-repo>. To do that, click the "fork" button at the top right corner of the screen. Once you've created the fork, create a local clone of the fork. Note that you must start with the `mock-cr` branch - you can change to that branch with the git command: `git checkout mock-cr`. You can also find the instructions I explained in the file "mock-cr-task.txt". Now open the website in your browser by opening `index.html`. ****wait** This is a website for playing classic arcade games, including Snake and Tetris. In case you are unfamiliar with these games, the repository has a file called "game-rules.txt" that explains how each game should function.

Are you ready to begin? ****wait**

Okay, I'm starting the timer.

D.1.4 Part 3: After Task

Time's up! Please save all of your files. At the top of the Catseye pane, to the right of the search bar, there's a sandwich menu - please click the sandwich menu, then click "Save Annotations to JSON". Once "output.json" has been created – you should see it at the top level of the project – push your changes to GitHub. ****wait**

Now, please complete the Google Form that I'm sending in the Zoom chat (see [Section D.3 for survey](#)). You can stop sharing your screen at this point. Please let

me know once you've completed the form. ****wait** Alright, with that, the study is officially over.

You can delete the repositories I had you clone - WebsiteTemplates and task-repo. You may uninstall Catseye by going to the Extensions pane in VS Code, finding Catseye in your list of extensions, clicking the gear icon and then clicking "Uninstall" from the dropdown menu.

If you would like to continue using Catseye, please be aware that we will still be collecting your user data. If, after 2 months from now, you still have the extension installed, we will notify you that your data is still being collected and will provide instructions on how to remove the application. This will continue until we begin to end this project. If you are still using the application as we begin ending this project, we will notify you 6 weeks in advance of the project ending such that you may prepare to remove the application. We plan to keep Catseye running for at least the next year. When we do plan to end the project, we will send your annotations to you in the form of a Google Sheet.

I will stop the recording now. Let me know what email you want the Amazon gift card sent to. Do you have any final questions for me? ****wait** Thank you for your time!

D.2 Control Condition Protocol

Hello, my name is Amber Horvath and I am a fourth-year Ph.D. student in the Human-Computer Interaction Institute at Carnegie Mellon University, working under Dr. Brad Myers. Research on developers has shown that they often need to keep track of lots of information. We are interested in understanding more about this need while completing a programming task.

This study is composed of 2 parts - the first part will be a tutorial to practice keeping track of information while programming. The second part will be a programming task. For the programming task, you will use the JavaScript programming language and the Visual Studio Code editor. Once you've finished the second task, I will have you complete a short survey about your background and your experience doing the task.

D.2.1 Consent

Before we begin, I'd like to review this consent form. Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form.

****if recording** I will begin recording this session.

D.2.2 Part 1: Tutorial

In this tutorial, we are going to practice thinking aloud about and externalizing information you are keeping track of while programming. Prior research has identified various situations in which people sometimes need to keep track of information, so in this tutorial, we will go through a scenario where you can explain to me how you would normally keep track of that kind of information. During this tutorial, I would like you to keep track of information in whatever way you normally would when working on a programming task. This could include creating comments in the code, writing information down in a separate document, or keeping all of this information in your head. I will also have you think-aloud.

In order for you to show me how you normally keep track of information while programming, I am going to have you clone a GitHub project that we will walk through together. Do you know how to clone a Git project? ****if yes** Okay, great, please create a clone of this project and open it in Visual Studio Code <https://github.com/ShauryaBhandari/Website-templates> - once you have opened the folder, would you mind sharing your screen?

****if not** Would you mind sharing your screen? Okay, copy this GitHub link:

In this tutorial, we will go through an exercise where you, a web developer, are using this repository of website templates as a starting point for writing some code for your own website. Let's open the index.html file in the bizland folder ****wait**

Keeping Track of Open Tasks

One thing that developers must keep track of when programming is the open tasks they are working on. These tasks may change over time as you, the developer, encounter bugs, or are assigned new tasks. Let's say you want to adapt this code to make your own website - I'd like to do whatever you would normally do to help you keep track of an open task - what do you normally do for keeping track of tasks? ****wait** Okay - throughout the main task, I would like you to do that each time you are trying to keep track of an open task.

Throughout the study, I'd like you to think aloud whenever you are starting a new task and think aloud about how you plan to actualize that task.

Keeping Track of Open Questions and Hypotheses

Another thing developers must keep track of, especially when dealing with unfamiliar code, is their open questions and hypotheses. Confusion surrounding code can lead to introducing buggy code and poor implementation planning, which can block developers from progressing on their work.

In the bizland index.html file, let's scroll down to line 316. There, you should see a comment that says "Import JQuery before materialize.js" - the original author of this code does not explain why this is necessary or what problems may arise if you

do not import jquery before materialize.js. Since we are hoping to adapt this code to our own needs and may want to import other libraries and need to change the order of the imports, I would like you to do whatever you would normally do to help you keep track of an open question or hypothesis - what do you normally do when you have a question about some code? ****wait** Okay - throughout the main task, I would like you to do that each time you have a question about the code.

One potential reason for this ordering is that materialize.js may depend on jquery in order to work and since we aren't using npm or another package management system to control our dependencies, we need to do this checking manually. Let's test this hypothesis by seeing what happens if we change the order of the imports - move line 321, the line with src set to 'js/materialize.min.js', up to line 317, before the jquery import. ****wait** Now save, and open index.html in your browser. ****wait** Once you have index.html open, look at the console message by right-clicking the browser window and selecting "Inspect" or by clicking F12. ****wait** Notice that, in the console, there's an error message saying that \$ is not defined and this error is being thrown by materialize.min.js. This confirms our hypothesis that materialize depends upon JQuery. I'd like you to practice thinking aloud that you found an answer to your initial question. ****wait** Great. Throughout the study, I'd like you to think aloud and do whatever you normally would when you have a question or hypothesis about the behavior of the code, or find an answer to any of your questions.

Keeping Track of What You Have Learned

Programming requires a lot of experimentation and naturally results in you discovering and learning things about the code. To explore this, let's scroll to the bottom of the file - lines 352 and 353. ****wait** There, we can see some latitude and longitude coordinates, but it is not immediately clear where those coordinates are pointing to. Let's try putting those coordinates into Google - copy the coordinates on line 359 and paste this into Google. ****wait** Now we can see that these coordinates are in the Boston city area. I'd like you to practice thinking aloud about how you learned that the coordinates point to the Boston city area. ****wait** Good - please think aloud when you learn things about the code or behavior of the program throughout the study. I would also like you to do whatever you would normally do to help you keep track of something you learned about the code - what do you normally do when you learn something about some code? ****wait** Okay - throughout the main task, I would like you to do that each time you learn something about the code.

Capturing System Output

If we scroll down a little further to line 373, there is a call to the Google Maps API to create a new Google Maps object using the coordinates we just explored. ****wait** Since we are new to the Google Maps API, we are curious what this object holds.

Now, to investigate this question, add a console log statement below line 376 that prints the map object, such as `console.log(map)`. ****wait** Now, let's save and reload `index.html` so we can see what is in the map object. If your console is not already open, open the console. Now, find the console log you added that printed the map object. ****wait** Let's expand the map object and look at its properties. I'd like you to practice thinking aloud about what you see ****wait**. Good - throughout the study, I would like you to think aloud and note what you discover when looking at the program output in whatever way you naturally would. If you notice things changing in the output, what do you normally do when trying to keep track of changing system output? ****wait** Okay, - throughout the main task, I would like you to do that each time you are interested in keeping track of changing system output.

Lightweight Versioning

Sometimes developers make a series of small changes in order to test a value or to try and debug some code – if these small edits do not work, the developer may need to keep track of what they changed or revert to a prior version of the code. If we look at line 374, we can see there's a `document.getElementById` call. If we want to change this code for our own purposes, we may want to change both this ID and the corresponding ID in the HTML. Change the ID to "mywebsiteid" in the JavaScript and the HTML ****wait**. Now, change the coordinates on line 353 to the latitude and longitude coordinates of Pittsburgh - you can find those coordinates by going to Google and searching "Pittsburgh coordinates". ****wait** Now, I would like you to go back to the prior version of the `document.getElementById` call and change the ID in the HTML, as well. ****wait** Good. If you have to revert to a prior version of the code but do not want to revert a whole file, what do you normally do? ****wait** Okay, - throughout the main task, I would like you to do that each time you are interested in reverting to a prior version for some part of the code.

Keeping Track of Patches of Code

When developers are working on fixing a bug or adding a feature, they often need to update and change multiple parts of the code which may be spread across multiple files that they must keep track of. Open `signup.html` in the Cloudlia folder and scroll to line 261 ****wait** - there is a JQuery event listener for adding carousel functionality to a particular element on the page. Let's say you wanted to add this carousel functionality to your `bizland index.html` file in the Testimonials section - how would you normally keep track of these two separate patches of code you are interested in? ****wait** Okay, throughout the main task, I would like you to do that each time you are interested in keeping track of and navigating between different parts of the codebase.

D.2.3 Part 2: Task

For the main task, I would like you to imagine that you are a new team member on a programming project and you have been tasked with understanding the code, along with finding and fixing bugs in the code. You do not want to edit your teammates' code until you have tested it – thus, for the first 15 minutes of the study, you are not allowed to edit the code except to add comments and print statements in the files. In other words, you are only allowed to make edits that do not affect the logic of the code. However, you may not edit the files that have the word “LOCKED” in the filename. You may also take notes on or think aloud about what you learn about the code through testing, what bugs you find, areas that can be improved, and any other pieces of information you want to keep track of in the same way you did during the tutorial. After the 15 minutes, you will have 30 minutes to try to improve upon and debug the code, given what you discovered in the first 15 minutes.

While you are working on this task, I would also like you to think aloud about what you are doing, how you plan on changing the code, and please speak aloud whenever you discover a bug and whenever you believe you've addressed a bug. I will notify you once your first 15 minutes of reviewing and testing the code are complete, and when your last 30 minutes of debugging are complete.

Now, I'd like you to create your own local fork of this repository. <https://github.com/horvathaa/ta-repo>. To do that, click the “fork” button at the top right corner of the screen. Once you've created the fork, create a local clone of the fork. Note that you must start with the mock-cr branch - you can change to that branch with this git command: `git checkout mock-cr`. Now open the website in your browser by opening `index.html`. ****wait** This is a website for playing classic arcade games, including Snake and Tetris. In case you are unfamiliar with these games, the repository has a file called “game-rules.txt” that explains how each game should function.

Do you have any questions? ****wait** Are you ready to begin? ****wait**

Okay, I'm starting the timer.

D.2.4 Part 3: After Task

Time's up! Please save all of your files. At the top of the Catseye pane, to the right of the search bar, there's a sandwich menu - please click the sandwich menu, then click “Save Annotations to JSON”. Once “output.json” has been created – you should see it at the top level of the project – push your changes to GitHub. ****wait**

Now, please complete the Google Form that I'm sending in the Zoom chat (see Section D.3 for survey). You can stop sharing your screen at this point. Please let me know once you've completed the form. ****wait** Alright, with that, the study is officially over.

You can delete the repositories I had you clone - WebsiteTemplates and task-repo. I will stop the recording now. Let me know what email you want the Amazon gift

card sent to. Do you have any final questions for me? ****wait** Thank you for your time!

D.3 Post-Task Survey

D.3.1 Notes

The following survey was used in the Catseye and control conditions – generally, in the Catseye condition, the words “Catseye” and “annotations” were used, while, in the control condition, the questions about specifically Catseye and its usability were omitted and questions about annotations were rephrased to be questions about “notes/comments”. For brevity, I will include just the Catseye survey.

D.3.2 Survey

Please fill out the following scale questions about your thoughts on the usability and usefulness of Catseye. The following pages will have some additional questions about your background. The last page will ask about your experience using Catseye.

1. What is your GitHub username? **[required, free response]**
2. For the following 7 statements, rate how much you agree or disagree with the provided statement. **[required, each statement is scored on a 7-point Likert from “Strongly Disagree” to “Strongly Agree”.]**
 - I consider Catseye easy to use and easy to achieve what I want
 - I consider my interactions with Catseye to be understandable and clear
 - I consider it easy for me to learn how to use Catseye
 - I would consider Catseye useful for my daily work
 - I enjoyed the features provided by Catseye
 - If possible, I would like to have the tool available for my professional or personal use
 - If possible, I would recommend Catseye to my friends and colleagues doing programming work

Background Questions

Please answer these questions about your programming background.

1. Are you familiar using comments in tools such as Google Docs, Microsoft Word, or Overleaf? **[required, yes/no/other]**
2. What notes do you take, if any, when programming? **[required, free response]**

3. What notes do you take, if any, when doing tasks that are not programming? [required, free response]
4. What is your profession? (e.g., student, software engineer, manager, etc.) [required, free response]
5. How many years have you been programming? [required, free response]
6. How many years of professional programming experience do you have? [required, free response]
7. What is your level of expertise in using JavaScript? [required, 7-point scale from "No experience" to "Expert"]

Short Answer Questions

Please write a free response to each question about your experience with Catseye, the annotation tool you used, and your note taking habits. Note that some of these questions are required.

1. What did you think about the tool, Catseye, in general? [free response]
2. In what situations do you think you would most want an annotation tool like Catseye? [free response]
3. Are there any features of Catseye that you found particularly good? [free response]
4. Are there any features of Catseye that you found particularly confusing? [free response]
5. What new features would you want in an annotation tool? [free response]
6. Are the annotations you created in today's study similar to notes/comments you create when programming for your own work? [required, yes/no/"some are similar, some are dissimilar"]
7. If you answered "yes" for the last question, in what ways are the annotations you created in today's study similar to your normal comments/notes? If you answered "no", in what ways are the annotations different? If you answered "some are similar, some are dissimilar", what annotations from the study are similar to your regular notes/comments and what ones are dissimilar? [free response]
8. How similar was the code you encountered in today's study to code you have encountered during your time as a programmer? We define "similar" as code that suffers from similar problems, such as poorly named variables and functions, lack of documentation and lack of organization. [required, 1 to 5 scale from "Not at all similar" to "Very similar"]

9. If you answered that you have encountered code that is similar to the code in today's study, please elaborate on what qualities of the study code are similar to what you have encountered in the past. [free response]
10. How frequently have you encountered code that is similar to the code in today's study? [required, 1 to 5 scale from "Never" to "Very frequently"]
11. Would you have wanted to delete the annotations you made during today's study? [required, yes/no/"I would want to delete some and keep some"/other]
12. How do you normally keep track of important information? [free response]
13. Have there been any situations where taking notes or keeping track of information using other strategies has not worked well? If so, please describe. [free response]

Appendix E

Chapter 6 Study Replication Materials

The following sections are the study protocol for the Sodalite evaluation study and the post-task survey. Parts of the protocol that were not spoken aloud are highlighted in yellow.

E.1 Sodalite Study Protocol

E.1.1 Introduction

Hi, my name is Amber Horvath and I am a 5th year PhD student at Carnegie Mellon University in the Human-Computer Interaction Institute, advised under Dr. Brad Myers. We are researching how developer documentation is authored and maintained and how this process may change when using an in-editor tool.

Research Study Overview

To investigate the documentation authoring and maintenance process, we developed a tool that allows you to create and maintain documentation for your code while staying within the context of your source code.

The study will be composed of two parts. The first part will be a tutorial to familiarize you with our Visual Studio Code extension for authoring and maintaining developer documentation called Sodalite. The second part will be an open-ended session in which you will use our tool to create some documentation for your own JavaScript or TypeScript code. You are free to document as much or as little of your code as you would like using the tool, and can document whatever type of information you feel is important. Upon completing your documentation, there is a 3rd part of the study, in which we will assess the maintenance aspect of the tool by reverting your code to an older version and seeing how well the documentation updates. Lastly, I will have you complete a short survey about your background and your experience using Sodalite.

Consent

Consent Before we begin, I'd like to review this consent form. **send consent form** Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you've submitted the form. ****wait** You will also receive an email with your response to the form.

****if recording** I will begin recording this session.

E.1.2 Tutorial

For the task, you will be using Sodalite, our experimental documentation tool, which allows you to create and maintain your documentation in the Visual Studio Code editor.

We will begin by downloading the VSIX file from this link – https://drive.google.com/file/d/1RxuNM3uoGHAL6k/view?usp=share_link – please right click on the file named “sodalite” and click “download”. It may take a moment. ****wait**

***wait for participant to finish downloading Sodalite** To install Sodalite, first open Visual Studio Code.

Open the extensions pane - you can do this by either using the keyboard shortcuts `ctrl/cmd + shift + x` or by clicking the Extensions button on the left side panel.

Click the three-dot menu in the top right corner of the Extensions side panel. Then, click the “Install from VSIX” menu option - this should open a file explorer window. From there, find and select the `sodalite.vsix` file we just downloaded and click “Install”.

****once Sodalite has finished installing - should prompt for GitHub auth** Okay, it looks like it's done. Let's first check that Sodalite is in our extensions list - in the Extensions side panel, it should appear in the “Installed” list. ****if there** Okay, now the extension should prompt you to authenticate the app with your GitHub account. ****if not there** Okay, let's quit and restart Visual Studio Code and then the extension should ask to authenticate with GitHub.

Start a Sodalite Story

I would like you to open the code base that you are planning on documenting. Note that this codebase must be implemented in the JavaScript or TypeScript programming languages. This is because the system is optimized for those languages. What project are you planning on using for this study? ****wait** Okay. Please open the folder in VS Code. Is this a public repository? ****wait** If so, may I have a link to the project?

I am only interested in this for assessing the size of the repository for analysis purposes. ****ifYes** Okay, great, thank you – please send the link in the chat.

****ifNo** That’s okay! Can you give me a rough estimate as to how large the project is, such as how many files there are or lines of code? ****wait** Thanks.

Note that, depending on the size of your repository, Sodalite may take a little longer to load as it is calculating relationships between different parts of the codebase. Once it is done, you can open the Sodalite panel by using the keyboard short CTRL/CMD + SHIFT + A. Try and open the panel. ****wait** Good. Now that it’s open, you should see a button that says “Create New Story” – I’d like you to click that button. ****wait** You should now see the story editor. Note that, during this session, if you could refrain from editing your code, that will help ensure that the code locations in the story are kept accurate.

Authoring a Sodalite Story

To write a Sodalite story, you can start by selecting a template. Sodalite has certain pre-defined templates to help you start off and structure your particular type of document. Certain templates also have different qualities, meaning different templates will favor different types of code references. For example, the Change Log template will favor functions that have been edited during your most recent coding session. You can also create your own code templates, but we will not be doing that in today’s study. For this brief tutorial, try selecting the Overview template, which can be used to create a high level walk through of your code. ****wait** Great.

Now that you’ve selected a template, you can begin writing your story. One of the first things you can do is give your story a Title – try titling your story “Tutorial”. ****wait** Good.

A core feature of Sodalite is being able to connect your code to your story, such that your documentation can be more clear and complete. To create a connection, try selecting some of the templated text in the story such as the word “describe” in the Summary section. You can select whatever text you want. ****wait** Good. Sodalite tries to leverage both the content in your story and the information in VS Code to suggest code connections. Try now selecting anywhere in VS Code. ****wait to see if code recommendations show up – if not prompt to click elsewhere.**

Notice that Sodalite now has a list of Suggestions on the left. These are different entities within your code. They are collapsed by default – you can expand them by clicking the down arrow. Try clicking the down arrow for the first item in the list ****wait** good. Some of them, such as methods or classes, may also have additional metadata, such as other locations in which they are referenced or places where they are defined. If you click the name of the method, Sodalite will jump you to the location in which that entity is defined. You can add a link to this definition by clicking the “plus” button in the top right corner. Try clicking the plus button. ****wait** Notice how the text that you had selected in the story changed to the color red. If you click that red text, it will take you to the location of the corresponding code.

****once code reference updated** Also notice how the Suggestions pane now has a Code Reference list below it – these are the parts of your code that are currently referenced within the code story.

****if there is Referenced in** In the “Referenced at” section, you can click on the name of different methods or classes in which this variable is referenced. Try clicking on one of the names ****wait** notice that this takes you to the place in which the method is referenced. If you’d like more information about the method that is referencing the entity you are looking at, you can click the question-mark button by the plus button. If you would like to add a code link to that particular place in which that reference is made, such as in a case where you want to discuss how a particular function is called, you can click the plus button by the question mark button.

Sometimes Sodalite is able to compute extra information about your code – you can see that information by clicking on the three-dot button at the bottom of the code reference box. Try clicking on the button.

****if nothing happens** Sodalite was unable to find any additional information about this code – during the course of the study, you can try clicking that button to see additional information.

****if additional information** Aside from just the call graph information, Sodalite computes how often certain parts of code appear in Git commits together, as an additional way of showing connections between the parts of your code. These references are sorted by most common to least common. In your case it appears that ****method name** was in ****number commits** with ****original method name**. You can investigate that method by clicking its name – Sodalite will take you to its definition. You can also add a link to that method by clicking the plus button to the right of the commit information.

In the case that you want to document a specific line or lines of code, you can select a range of code and add that as a code link. Try selecting some code ****wait** notice how the first item in the Suggestion list is now the code you selected. You can then add that code to your story the same way you have before – by selecting some text in the story. If you do not select text in the story and click the plus button, Sodalite will just put the code content directly into your rich text editor. This works alright with things like method names, but in the case of a whole line of code, it may be too much.

Sodalite also functions as a rich text editor. This means it supports adding in links, images, and formatting the text. Feel free to try adding in images or links to make your documentation even better. Note, however, that some features are less supported than others for the code links – copy-pasting a code link or undoing-redoing a change which involves a code link may cause undefined behavior, so I would not recommend doing those actions during the study. Deleting code links is supported, but may also not work 100% perfectly, so I would also not advise doing that unless totally necessary.

To exit the story creation editor, you can either click the “X” button in the top right corner to discard the current story or click the green “Save” button below the editor to save the story. Doing that will create a folder in your repository directory, along with a JSON file in that folder that has information about your story. Make sure to do this at the end of the study! However, in this case, let’s discard this tutorial story by clicking the “X” button.

One final note before starting – given that Sodalite is experimental research software, there is a chance that it may experience some bugs, particularly with the code links. In the case that you experience problems with Sodalite, I would recommend taking whatever you have written, copying it into a plain text editor, clicking the “cancel” button in Sodalite, then re-open a new story in Sodalite and paste in your content. Sorry in advance if this happens, but I would hate for you to lose all of your work! I may also interject if I notice a problem in the way Sodalite is functioning.

Another important note is to not switch templates once you start writing as that may result in you losing your text.

That is it for the tutorial, do you have any questions? ****wait** If, during the course of the session, you forget how to do anything with Sodalite, please feel free to ask me.

E.1.3 Study

We will now begin the main portion of the study. You will be using Sodalite to write some documentation about your code. What you choose to document is completely up to you. Subsequently, the session can go on for as long or short as you would like. You will be compensated at \$25 an hour, for up to 4 hours. Again, if you have any questions about how Sodalite works, please feel free to ask.

Are you ready? ****wait** When you are complete, feel free to tell me. You may now begin.

Anchor Check

You’re done? Alright, please upload the generated JSON file to this link. Let me know once that is done.

****wait** Alright, great. Can you now check out an earlier version of your code? We are also assessing how well our algorithm performs at finding appropriate re-attachment points, given the code links you included in your story. Please go to your GitHub and find a version of your code that is older than 3 months and check out that version of your code. ****Wait** Once you have checked out this code, please create a “.sodalite” folder at the top of your project directory and put the copied JSON file in there. ****wait** Once you have done that, please close then reopen Visual Studio Code and then open Sodalite. ****wait** We will go through each link one at a time, and I would like you to tell me what you think of any messages that Sodalite

provides and whether or not you think the place at which the code link is anchored is appropriate or not. Please open Sodalite. ****wait**

****For each link – if marked as warning or invalid:** Click the link. What do you think of the message? ****wait** Now, do you think this re-attachment spot is valid?

End

The study is now over. I have a couple of questions to ask to get your feedback, and then I will pay you. Here is the link to the Google Form (see Section E.2) – tell me when you are done. You can stop sharing your screen now. ****once completed** Please send me the JSON story that you created – you can do this by uploading it to this Google Drive. Please let me know once it is uploaded. ****wait**

You are also free to uninstall Sodalite – you can do so by going to the Extensions pane, finding Sodalite in the list, then clicking the three dot menu and clicking the “Uninstall” option. You are also free to delete the .sodalite folder in your repository.

Okay, the study is now over. Send me the email you use for PayPal/Amazon, and you will receive your compensation of ****dollars** in the next few days. Thank you so much!

E.2 Post-Task Survey

Please fill out the following scale questions about your thoughts on the usability and usefulness of Sodalite. The following page will have some additional questions about your background. The last page will ask about your experience using Sodalite.

1. What is your GitHub username? **[required, free response]**
2. For the following 7 statements, rate how much you agree or disagree with the provided statement. **[required, each statement is scored on a 7-point Likert from “Strongly Disagree” to “Strongly Agree”.]**
 - I consider Sodalite easy to use and easy to achieve what I want
 - I consider my interactions with Sodalite to be understandable and clear
 - I consider it easy for me to learn how to use Sodalite
 - I would consider Sodalite useful for my daily work
 - I enjoyed the features provided by Sodalite
 - If possible, I would like to have the tool available for my professional or personal use
 - If possible, I would recommend Sodalite to my friends and colleagues doing programming work

Background Questions

Please answer these questions about your programming background.

1. What is your profession? (e.g., student, software engineer, manager, etc.) [required, free response]
2. How many years have you been programming? [required, free response]
3. How many years of professional programming experience do you have? [required, free response]
4. What is your level of expertise in using JavaScript? [required, 7-point scale from “No experience” to “Expert”]

Short Answer Questions

Please write a free response to each question about your experience with Sodalite.

1. What did you think about the tool, Sodalite, in general? [free response]
2. In what situations do you think you would most want an documentation tool like Sodalite? [free response]
3. Are there any features of Sodalite that you found particularly good? [free response]
4. Are there any features of Sodalite that you found particularly confusing? [free response]
5. What new features would you want in an in-editor documentation tool? [free response]

Appendix F

Chapter 7 Study Replication Materials

The following sections are the study protocol for the Meta-Manager evaluation study and the post-task survey. Parts of the protocol that were not spoken aloud are highlighted in yellow.

F.1 Meta-Manager Study Protocol

Hi, my name is Amber Horvath and I am a 6th year PhD student in the Human-Computer Interaction Institute working under Dr. Brad Myers. We are researching how developers manage meta-information related to code provenance and investigating tools to help organize this information.

For the study, you will be using our tool, the “Meta-Manager”, a Visual Studio Code extension, to investigate a code base and to answer some questions about the history of the code. The study will begin with a tutorial to familiarize you with our extension. In the second part, we will walk through a code base and you will use the tool to answer some questions about the code and its history. Note that you are not required to program during this task. Do you have any questions? ****wait** Ok great.

Before we begin, I’d like to review this consent form. **send consent form** Please open the link I just sent in the Zoom chat. Please enter your email at the top of the form. Let me know if you have any questions. When you are ready, there are two optional permissions you must either agree or disagree to at the bottom of the document. ****wait** Please let me know if you agree to me recording this session so I may start the recording or not. You must also consent to the study by typing in your name and date. Please let me know once you’ve submitted the form. ****wait** You will also receive an email with your response to the form.

***if recording** I will begin recording this session.

F.1.1 Tutorial

We will now go through Meta-Manager. I will now give you remote control of my computer — are you ready? ****IF yes, relinquish control — SCREEN SHARE → REMOTE CONTROL → GIVE CONTROL **open ExtractorGithubGist.ts**

Meta-Manager is a code history and information explorer. It is shown in the bottom part of the screen. On the left, there is a timeline chart of the code in this file, in this case “ExtractorGithubGist.ts”. The y-axis represents the line count of the file, and the x-axis represents the various edits that the Meta Manager recorded over time.

On the right side of the screen, you can see information about the various parts of the code as it appeared at that point in time. Currently, we are at Version 0, as shown above the list of code parts, so this is how the code looked when the project started. Below “Version 0”, we can see a box that says “Arrow Function”. Note how “Arrow Function” is highlighted in dark blue — the blue in the chart corresponds to the “Arrow Function”.

To see the code at this version in time, you can click the box, which will expand. Try clicking the “Arrow Function” box. ****wait** Good. Notice how you can see the code.

To see other versions of the code, we can use the scrubber on the bottom axis of the chart. ****circle scrubber** You can click and drag the scrubber ball to flip through different versions of the file across time. Try dragging the scrubber ****wait**. Good. Note how the right side code box has updated with new code, a new time stamp, and a new version number.

Sometimes, the x-axis will have additional points of information that we can investigate. Look at the x-axis at point 12 and point 25 — note how there are labels that says “Copied code”. We can click on these labels to jump to those code versions — try clicking on the label that says “Copied code” at point 12. ****wait** Good. Notice how the code box updated with some additional information and buttons.

These gray buttons act as filters, while the blue buttons are an action. “Show Copied Code” will filter the scrubber timeline to only show “Copied code” events related to this particular function. “Show Pasted Code” will filter to only show code that was pasted in this function. “Show code pasted from online” will filter to only show code that was taken from online sources. Try clicking the different filter buttons and watch how the timeline updates ****wait** good. Now select all the filters again so we can see all of the events along the timeline.

Meanwhile, the blue buttons are different actions that we can take with the code in the box. “See Code Now” acts as a search function — if you have some code selected in the code version, pressing the “Search for Selected Code” button will search forwards in time from the code version to now to see where the code is and where it was. The timeline will update with a list of places in which that code existed and how it changed over time. Try selecting the code `// abstract class function`, then clicking “Search for Selected Code”. ****wait** Notice how the timeline updated with new labels that says “Search Result”. Click the “Search Result” label. ****wait** Now we can see the first instance after our point in time where the code appeared is highlighted in orange. There are also arrow buttons that will move you through each search result — try clicking the right button. ****wait** Good. Once you are done

searching, you can click the “Clear search?” button above the code box. Try clearing the search — good. Let’s look at “Copied code” at version 25.

In the case of “Copied code”, the Meta-Manager keeps track of what code has been copied and where it eventually gets pasted. Notice how, instead of just saying the version number, the text says “Copied sourceUrl: options.url from Arrow Function”. If you scroll down, you will also see that the code that has been copied is highlighted in the code version. With a copied code block, we can get some additional information. The “See All Paste Locations” button at the top of the box will filter the timeline to show instances where that code has been pasted. Try clicking “See All Paste Locations”. ****wait** Good. In instances where the code was pasted *outside* of the currently-open file, the code box will say where the code was pasted, at what version, and what the code looked like at the time of the paste. Let’s look at this paste event. Try clicking the “Pasted code” label. Good.

Meta-Manager similarly keeps track of code that has been pasted from various locations. Like with copy, the text in the box will say what code has been pasted and the code will be highlighted in the code version. Additionally, we can see the corresponding copy event by clicking “See Corresponding Copy”, and we can see what other code was pasted with the same value by clicking “See All Other Paste Locations”. Try clicking “See All Other Paste Locations” ****wait** Good. “See Corresponding Copy” and “See All Other Paste Locations” will also show previews of the copied or pasted code in the case the copy or paste happened in a different file.

Meta-Manager keeps track of additional information for code that was pasted from online. Let’s look at some code that was pasted from Stack Overflow — we can first reset our filter by clicking the “Reset filter?” button so we can see all events again — try clicking that button. ****wait** Good. Now, look at index 27 on the timeline, where the label says “Stack Overflow”. Try clicking that label. ****wait** Good.

Like the other paste events, we can see what code was copied from Stack Overflow at the highlight in the code box, and the same filtering and searching activities are supported. We can also see some information about the Stack Overflow question, including the title, views, votes and when it was posted. If we want even more information, we can click the “See More” button, which will pull up the original question in Visual Studio Code. Try clicking “see More”. ****wait** Good - note how it appeared above as a window — any tabs opened like this will have the title “Meta Manager”, in case you lose it among your other windows. Meta-Manager also supports other websites such as ChatGPT which will have information about the ChatGPT thread some code was copied from. If the user made any programming-related Google searches prior to copy-pasting from these sources, that query will also be included as part of the summary details.

Aside from the code boxes, you can do some other activities with Meta-Manager. Similar to how we can search for code using the “See Code Now” button in the code box, you can also search for code using the search box in the top right corner. Try searching for the word “target”. ****Wait** Note how the timeline updates with

instances where the code “target” appeared in the version.

Another way of searching for code is backwards through time — we can do this by selecting some code in the editor and using the right-click menu to choose “Meta Manager: Search Code Across Time”. Try selecting line 9 and the code `const target equals parse HTML.` ****wait** Then, right click the code and click “Meta Manager: Search for Code Across Time”. ****wait** Notice how the timeline has now updated with labels that says “Search Result” — this search will look for instances in which the code changed over time and when it first appears. If you click at index 4, you should see the first instance of this code line. You can then scrub thru to see what happened to that code — in this case, it was removed up until version 23, which is also marked as a “Search Result” along the timeline. Note that the search is not perfect — your best bet is to start at the rightmost point, thus, most recent, and go through the search results from right to left. Once we are done searching, we can click the “Clear search?” button - try clicking that. ****wait** Good.

You may see that there are many timeline events in close succession — in cases like that, it may be easier to look at each individual version by zooming in on the chart. To zoom, you can set the scrubber at where you want the zoom to begin, then hold down the shift key and click where you want the zoom to end. Try setting the scrubber at index 20 ****wait** — now shift and click at index 32 ****wait**. Now you can see the chart at a closer level and see the marked indices along the timeline more closely. To leave this view, you can click the “Reset version range” button. ****wait** Good.

That does it for the tutorial — do you have any questions? During the study, you are free to ask me any questions about how to use the tool. ****wait** Okay, great.

F.1.2 Task

For the study, we are going to explore this repository. For context, the repository is for a small Visual Studio Code extension that searches the internet to supply in-line suggestions, similar to how the popular CoPilot extension works. A “suggestion” in Visual Studio Code is essentially an auto-complete option.

Your task will be to use the Meta-Manager to explore the code in order to answer some questions about the code history that I will be asking. All of the questions should be answerable with only the tool — thus, there should be no need to add or remove code, or search information online. If you do have questions about some detail of the Visual Studio Code API or TypeScript, you can ask me, but I may not answer if it is information that you can get using the tool. You are also free to ask me to repeat any of the questions. If you feel you are unable to answer a question, let me know and we will move on. I will be allowing 10 minutes for each question.

Additionally, I would like you to think aloud as you use the tool to try and answer the question. Are you ready? ****wait**

Okay, to begin, open “config.ts”. ****wait** This file is used for configuring how the extension should be triggered to produce a suggestion.

First question: In `config.ts`, there is a regex for search pattern matching. Can you tell me what it is trying to match on and why? Let me know when you think you have answered the question and what your answer is.

****Wait**

Second question: There is a bug in the commented out code, can you, without running the code, find where the bug was and what happened.

****Wait**

Great — now let's move to `ExtractorStackOverflow.ts`. Note that, because of Zoom, the extension may be slightly slower to load on these files with many many versions — sorry in advance!

Third question: Prior to using `parseHTML` on line 18, the author was using a different API — what was it? Why did they stop using the previous API? Like before, think aloud and let me know when you believe you've found an answer.

****wait**

Fourth question: Recently, some code was added that came from a different file. Can you find that code and explain what changed?

****wait**

Now, we will go to `search.ts`. ****open search.ts and wait**

Fifth question: Look at lines 65 to 67 — there is a commented out `forEach` loop. Can you find the last time it was used? Please explain why it was removed and what it was replaced with. Let me know once you've come to an answer.

****wait**

Lastly, we will look at `extension.ts`. ****open extension.ts and wait**

Sixth question: What code was generated by an AI system? Additionally, what ended up happening to it? Acceptable answers are "it was deleted", "it was edited", "it was commented out", etc. ****wait**

Seventh question: What were all the different things that the programmer tried for setting `match`? Please find the versions in which `match` was changed. ****wait**

Eighth question: Some code from `activate` was moved into a different file. When did this happen and what was the code that was moved? ****wait**

F.1.3 End Task

That's it! I will now have you complete a survey (see Section F.2) about your background and experience using the tool. You can take control of your own screen again.

F.2 Post-Task Survey

Please fill out the following scale questions about your thoughts on the usability and usefulness of Meta-Manager. The following page will have some additional questions about your background. The last page will ask about your experience using Meta-Manager.

1. For the following 7 statements, rate how much you agree or disagree with the provided statement. [required, each statement is scored on a 7-point Likert from “Strongly Disagree” to “Strongly Agree”].
 - I consider Meta-Manager easy to use and easy to achieve what I want
 - I consider my interactions with Meta-Manager to be understandable and clear
 - I consider it easy for me to learn how to use Meta-Manager
 - I would consider Meta-Manager useful for my daily work
 - I enjoyed the features provided by Meta-Manager
 - If possible, I would like to have the tool available for my professional or personal use
 - If possible, I would recommend Meta-Manager to my friends and colleagues doing programming work

Background Questions

Please answer these questions about your programming background.

1. What is your profession? (e.g., student, software engineer, manager, etc.) [required, free response]
2. How many years have you been programming? [required, free response]
3. How many years of professional programming experience do you have? [required, free response]
4. What is your level of expertise in using JavaScript? [required, 7-point scale from “No experience” to “Expert”]
5. What is your level of expertise in using TypeScript? [required, 7-point scale from “No experience” to “Expert”]

Short Answer Questions

Please write a free response to each question about your experience with Meta-Manager, the tool you used, and the questions you answered during the study. Note that some of these questions are required.

1. What did you think about the tool, Meta-Manager, in general? [free response]
2. In what situations do you think you would most want a code history tool like Meta-Manager? [free response]
3. Are there any features of Meta-Manager that you found particularly good? [free response]

4. Are there any features of Meta-Manager that you found particularly confusing? [free response]
5. What new features would you want in a code history tool? [free response]
6. For the following questions that you answered during the study, please rate how often you have asked similar questions during your own programming. [required, each statement is scored on a 5-point Likert from "Never Asked" to "Always Asked".]
 - What is the design rationale behind some code?
 - Where did a code bug come from and how was it fixed?
 - Why was a particular API used?
 - When did we stop using this code?
 - Where did this code come from?
 - Why was this code added?
 - What website did this code come from?
 - How was this code refactored?
 - What did the original author of this code try when implementing some code?

Appendix G

Chapter 8 Interview Protocol

The following section is the interview protocol used in the exploratory interview discussed in Chapter 8 (see Section 8.3). The interview was semi-structured, so follow-up questions were asked if they organically arose given a participant’s answer to a question. Parts of the protocol that were not spoken aloud are **highlighted in yellow**.

G.1 Protocol

Hi, my name is Amber Horvath and I am a researcher in the HCII at Carnegie Mellon University. We are studying how developers utilize log statements and the subsequent output when understanding code. We are specifically interested in the sorts of *questions* you have when interpreting log statement output and how those outputs relate to code changes over time.

I will ask you questions about how you create log statements, how you manage the statements, and how you interpret their output, with a particular emphasis on the types of questions you are asking at each stage of this process. At any point, if you feel comfortable doing so, you can show your code, log statements, or outputs, if the outputs are saved. Throughout this interview, I may use the terms “log statements”, “print statements”, and “console logs” interchangeably — all of those terms refer to the *source code you wrote to produce some observable output in a console, terminal, or log file*. Similarly, I may use the terms “output” and “log” interchangeably — when I use those terms, I am referring to the *observable content produced by the log statement in the source code you wrote*. Do you have any questions? ****wait**

Here is the consent form *****pause***... Do you consent to being recorded? *****if yes, record***

G.1.1 Questions

1. Can you describe a bit about you, what kind of programming you do, your position, etc.?
 - Can you recall to me your most recent debugging session? Please discuss how, if at all, you used log statements, such as `console.log`, during this session.

2. So now, think about your debugging session. When you were running your code, how did you find the outputs you are interested in?
 - What questions or queries would you like to be able to ask to find log statements and outputs of interest? For example, a query such as “find every output produced by this statement where the function returned undefined”.
3. What questions did you have when you are interpreting logs produced by multiple related log statements? For example, you have a log statement in function A and function B with both logging the same variable C at different points during the code’s execution.
 - How did you find and organize outputs produced by different log statements?
 - What questions or queries would you like to be able to ask to help in reasoning about multiple related log statements?
4. Across different runs of your code, how did you interpret outputs produced by the same log statement?
 - What questions did you have when doing this interpretation?
 - What strategies did you use when doing this interpretation?
 - What questions or queries would you like to be able to ask to help with interpreting log values produced by the same statement across time?
5. How did you think about your log statements in terms of their relationships to one another?
 - Did you think about them in terms of their execution order? In terms of program structure? Some other way?
 - What queries or questions would you like to be able to ask a system in terms of visualizing and organizing log statements given their relationship?
6. How did you keep track of your code edits and how the log statements’ outputs change with respect to those edits?
 - What sorts of questions or queries would you want to ask your code history or output in order to help with this sense making process?
7. Do you have any other log-related thoughts that have not come up organically during this interview? Any other log-related queries you would like to be able to ask?

Okay, with that we are done with the interview. Thank you so much for your time and please let me know if you have any other log-related thoughts.

Bibliography

- [1] Marjan Adeli, Nicholas Nelson, Souti Chattopadhyay, Hayden Coffey, Austin Henley, and Anita Sarma. “Supporting code comprehension via annotations: Right information at the right time and place”. In: *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE. 2020, pp. 1–10.
- [2] Emad Aghajani, Gabriele Bavota, Mario Linares-Vásquez, and Michele Lanza. “Automated documentation of android apps”. In: *IEEE Transactions on Software Engineering* 47.1 (2019), pp. 204–220.
- [3] Emad Aghajani, Csaba Nagy, Mario Linares-Vásquez, Laura Moreno, Gabriele Bavota, Michele Lanza, and David C. Shepherd. “Software Documentation: The Practitioners’ Perspective”. In: *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. ICSE ’20. Seoul, South Korea: Association for Computing Machinery, 2020, 590–601. ISBN: 9781450371216. DOI: [10.1145/3377811.3380405](https://doi.org/10.1145/3377811.3380405). URL: <https://doi.org/10.1145/3377811.3380405>.
- [4] Emad Aghajani, Csaba Nagy, Olga Lucero Vega-Márquez, Mario Linares-Vásquez, Laura Moreno, Gabriele Bavota, and Michele Lanza. “Software Documentation Issues Unveiled”. In: *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. Montreal, QC, Canada: IEEE, 2019, pp. 1199–1210. DOI: [10.1109/ICSE.2019.00122](https://doi.org/10.1109/ICSE.2019.00122).
- [5] Maristella Agosti, Giorgetta Bonfiglio-Dosio, and Nicola Ferro. “A historical and contemporary study on annotations to derive key features for systems design”. In: *International Journal on Digital Libraries* 8.1 (2007), pp. 1–19.
- [6] Mohammad Allahbakhsh, Boualem Benatallah, Aleksandar Ignjatovic, Hamid Reza Motahari-Nezhad, Elisa Bertino, and Schahram Dustdar. “Quality Control in Crowdsourcing Systems: Issues and Directions”. In: *IEEE Internet Computing* 17 (2 2013), pp. 76–81. DOI: [10.1109/MIC.2013.20](https://doi.org/10.1109/MIC.2013.20). URL: <https://doi.org/10.1109/MIC.2013.20>.
- [7] Sven Amann, Sebastian Proksch, Sarah Nadi, and Mira Mezini. “A Study of Visual Studio Usage in Practice”. In: *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. Vol. 1. 2016, pp. 124–134. DOI: [10.1109/SANER.2016.39](https://doi.org/10.1109/SANER.2016.39).

- [8] Muhammad Asaduzzaman, Ahmed Shah Mashiyat, Chanchal K. Roy, and Kevin A. Schneider. "Answering Questions about Unanswered Questions of Stack Overflow". In: *MSR 2013*. San Francisco, CA, USA: IEEE, 2013, pp. 97–100. DOI: [10.1109/MSR.2013.6624015](https://doi.org/10.1109/MSR.2013.6624015). URL: <https://doi.org/10.1109/MSR.2013.6624015>.
- [9] Alberto Bacchelli and Christian Bird. "Expectations, outcomes, and challenges of modern code review". In: *2013 35th International Conference on Software Engineering (ICSE)*. 2013, pp. 712–721. DOI: [10.1109/ICSE.2013.6606617](https://doi.org/10.1109/ICSE.2013.6606617).
- [10] Jeff Baker, Donald Jones, and Jim Burkman. "Using visual representations of data to enhance sensemaking in data exploration tasks". In: *Journal of the Association for Information Systems* 10.7 (2009), p. 2.
- [11] Sebastian Baltes, Richard Kiefer, and Stephan Diehl. "Attribution Required: Stack Overflow Code Snippets in GitHub Projects". In: *2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C)*. 2017, pp. 161–163. DOI: [10.1109/ICSE-C.2017.99](https://doi.org/10.1109/ICSE-C.2017.99).
- [12] Sebastian Baltes, Christoph Treude, and Stephan Diehl. "SOTorrent: Studying the Origin, Evolution, and Usage of Stack Overflow Code Snippets". In: *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. 2019, pp. 191–194. DOI: [10.1109/MSR.2019.00038](https://doi.org/10.1109/MSR.2019.00038).
- [13] Liang Bao, Qian Li, Peiyao Lu, Jie Lu, Tongxiao Ruan, and Ke Zhang. "Execution anomaly detection in large-scale systems through console log analysis". In: *Journal of Systems and Software* 143 (2018), pp. 172–186. ISSN: 0164-1212. DOI: <https://doi.org/10.1016/j.jss.2018.05.016>. URL: <https://www.sciencedirect.com/science/article/pii/S0164121218301031>.
- [14] Deborah K Barreau. "Context as a factor in personal information management systems". In: *Journal of the American society for information science* 46.5 (1995), pp. 327–339.
- [15] David Bawden, Clive Holtham, and Nigel Courtney. "Perspectives on information overload". In: *Aslib proceedings*. Vol. 51. 8. MCB UP Ltd. 1999, pp. 249–255.
- [16] Laura Beckwith, Cory Kissinger, Margaret Burnett, Susan Wiedenbeck, Joseph Lawrance, Alan Blackwell, and Curtis Cook. "Tinkering and gender in end-user programmers' debugging". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '06. Montréal, Québec, Canada: Association for Computing Machinery, 2006, 231–240. ISBN: 1595933727. DOI: [10.1145/1124772.1124808](https://doi.org/10.1145/1124772.1124808). URL: <https://doi.org/10.1145/1124772.1124808>.

- [17] Andrew Begel, Yit Phang Khoo, and Thomas Zimmermann. “Codebook: Discovering and Exploiting Relationships in Software Repositories”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1. ICSE '10*. Cape Town, South Africa: Association for Computing Machinery, 2010, 125–134. ISBN: 9781605587196. DOI: [10.1145/1806799.1806821](https://doi.org/10.1145/1806799.1806821). URL: <https://doi.org/10.1145/1806799.1806821>.
- [18] Andrew Begel and Beth Simon. “Novice Software Developers, All over Again”. In: *Proceedings of the Fourth International Workshop on Computing Education Research*. ICER '08. Sydney, Australia: Association for Computing Machinery, 2008, 3–14. ISBN: 9781605582160. DOI: [10.1145/1404520.1404522](https://doi.org/10.1145/1404520.1404522). URL: <https://doi.org/10.1145/1404520.1404522>.
- [19] Moritz Beller, Niels Spruit, Diomidis Spinellis, and Andy Zaidman. “On the dichotomy of debugging behavior among programmers”. In: *Proceedings of the 40th International Conference on Software Engineering. ICSE '18*. Gothenburg, Sweden: Association for Computing Machinery, 2018, 572–583. ISBN: 9781450356381. DOI: [10.1145/3180155.3180175](https://doi.org/10.1145/3180155.3180175). URL: <https://doi.org/10.1145/3180155.3180175>.
- [20] Moritz Beller, Niels Spruit, and Andy Zaidman. “How developers debug”. In: *PeerJ Preprints* 5 (2017), e2743v1.
- [21] Michael Bernstein, Max Van Kleek, David Karger, and MC Schraefel. “Information scraps: How and why information eludes our personal information management tools”. In: *ACM Transactions on Information Systems (TOIS)* 26.4 (2008), pp. 1–46.
- [22] Michael Scott Bernstein. “Information Scraps: Understanding and Design”. PhD thesis. Citeseer, 2008.
- [23] Nicolas Bettenburg, Sascha Just, Adrian Schröter, Cathrin Weiss, Rahul Premraj, and Thomas Zimmermann. “What makes a good bug report?”. In: *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering*. 2008, pp. 308–318.
- [24] Christian Bird, Denae Ford, Thomas Zimmermann, Nicole Forsgren, Eirini Kalliamvakou, Travis Lowdermilk, and Idan Gazit. “Taking Flight with Copilot: Early insights and opportunities of AI-powered pair-programming tools”. In: *Queue* 20.6 (2023), 35–57. ISSN: 1542-7730. DOI: [10.1145/3582083](https://doi.org/10.1145/3582083). URL: <https://doi.org/10.1145/3582083>.
- [25] Tristan Blanc-Brude and Dominique L. Scapin. “What do people recall about their documents? implications for desktop search tools”. In: *Proceedings of the 12th International Conference on Intelligent User Interfaces. IUI '07*. Honolulu, Hawaii, USA: Association for Computing Machinery, 2007, 102–111. ISBN: 1595934812. DOI: [10.1145/1216295.1216319](https://doi.org/10.1145/1216295.1216319). URL: <https://doi.org/10.1145/1216295.1216319>.

- [26] Jürgen Börstler and Barbara Paech. “The role of method chains and comments in software readability and comprehension—An experiment”. In: *IEEE Transactions on Software Engineering* 42.9 (2016), pp. 886–898.
- [27] Andrew Bragdon, Robert Zeleznik, Steven P. Reiss, Suman Karumuri, William Cheung, Joshua Kaplan, Christopher Coleman, Ferdi Adeputra, and Joseph J. LaViola. “Code Bubbles: A Working Set-Based Interface for Code Understanding and Maintenance”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: Association for Computing Machinery, 2010, 2503–2512. ISBN: 9781605589299. DOI: [10.1145/1753326.1753706](https://doi.org/10.1145/1753326.1753706). URL: <https://doi.org/10.1145/1753326.1753706>.
- [28] Joel Brandt, Mira Dontcheva, Marcos Weskamp, and Scott R. Klemmer. “Example-Centric Programming: Integrating Web Search into the Development Environment”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: Association for Computing Machinery, 2010, 513–522. ISBN: 9781605589299. DOI: [10.1145/1753326.1753402](https://doi.org/10.1145/1753326.1753402). URL: <https://doi.org/10.1145/1753326.1753402>.
- [29] Joel Brandt, Philip J. Guo, Joel Lewenstein, Mira Dontcheva, and Scott R. Klemmer. “Two Studies of Opportunistic Programming: Interleaving Web Foraging, Learning, and Writing Code”. In: *CHI '09*. CHI '09. Boston, MA, USA: Association for Computing Machinery, 2009, 1589–1598. ISBN: 9781605582467. DOI: [10.1145/1518701.1518944](https://doi.org/10.1145/1518701.1518944). URL: <https://doi.org/10.1145/1518701.1518944>.
- [30] Margaret Burnett, Simone Stumpf, Jamie Macbeth, Stephann Makri, Laura Beckwith, Irwin Kwan, Anicia Peters, and William Jernigan. “GenderMag: A Method for Evaluating Software’s Gender Inclusiveness”. In: *Interacting with Computers* 28.6 (Nov. 2016). DOI: [10.1145/3134737](https://doi.org/10.1145/3134737). URL: <https://doi.org/10.1145/3134737>.
- [31] CodeSandbox BV. *CodeSandbox: Online Code Editor and IDE for Rapid Web Development*. CodeSandbox BV. 2021. URL: <https://codesandbox.io/>.
- [32] Paul Chandler and John Sweller. “Cognitive load theory and the format of instruction”. In: *Cognition and instruction* 8.4 (1991), pp. 293–332.
- [33] Joseph Chee Chang, Nathan Hahn, Yongsung Kim, Julina Coupland, Bradley Breneisen, Hannah S Kim, John Hwong, and Aniket Kittur. “When the Tab Comes Due: Challenges in the Cost Structure of Browser Tab Usage”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: [10.1145/3411764.3445585](https://doi.org/10.1145/3411764.3445585). URL: <https://doi.org/10.1145/3411764.3445585>.

- [34] Joseph Chee Chang, Nathan Hahn, and Aniket Kittur. "Supporting Mobile Sensemaking Through Intentionally Uncertain Highlighting". In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. UIST '16. Tokyo, Japan: Association for Computing Machinery, 2016, 61–68. ISBN: 9781450341899. DOI: [10.1145/2984511.2984538](https://doi.org/10.1145/2984511.2984538). URL: <https://doi.org/10.1145/2984511.2984538>.
- [35] Preetha Chatterjee, Manziba Akanda Nishi, Kostadin Damevski, Vinay Augustine, Lori Pollock, and Nicholas A. Kraft. "What information about code snippets is available in different software-related documents? An exploratory study". In: *SANER 2017*. New York City, NY, USA: IEEE, 2017, pp. 382–386.
- [36] Eason Chen, Ray Huang, Han-Shin Chen, Yuen-Hsien Tseng, and Liang-Yi Li. "GPTutor: A ChatGPT-Powered Programming Tool for Code Explanation". In: *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*. Ed. by Ning Wang, Genaro Rebolledo-Mendez, Vania Dimitrova, Noboru Matsuda, and Olga C. Santos. Cham: Springer Nature Switzerland, 2023, pp. 321–327. ISBN: 978-3-031-36336-8.
- [37] J. C. Chen and S. J. Huang. "An empirical analysis of the impact of software development problem factors on software maintainability". In: *Journal of Systems and Software* 82.6 (2009).
- [38] Kai Chen, Stephen R. Schach, Ligu Yu, Jeff Offutt, and Gillian Z. Heller. "Open-Source Change Logs". In: *Empirical Software Engineering* 9.3 (Sept. 2004), pp. 197–210. ISSN: 1573-7616. DOI: [10.1023/B:EMSE.0000027779.70556.d0](https://doi.org/10.1023/B:EMSE.0000027779.70556.d0). URL: <https://doi.org/10.1023/B:EMSE.0000027779.70556.d0>.
- [39] Mauro Cherubini, Gina Venolia, Rob DeLine, and Amy J. Ko. "Let's Go to the Whiteboard: How and Why Software Developers Use Drawings". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '07. San Jose, California, USA: Association for Computing Machinery, 2007, 557–566. ISBN: 9781595935939. DOI: [10.1145/1240624.1240714](https://doi.org/10.1145/1240624.1240714). URL: <https://doi.org/10.1145/1240624.1240714>.
- [40] Michelene TH Chi, Miriam Bassok, Matthew W Lewis, Peter Reimann, and Robert Glaser. "Self-explanations: How students study and use examples in learning to solve problems". In: *Cognitive science* 13.2 (1989), pp. 145–182.
- [41] Michelene TH Chi, Miriam Bassok, Matthew W Lewis, Peter Reimann, and Robert Glaser. "Self-explanations: How students study and use examples in learning to solve problems". In: *Cognitive science* 13.2 (1989), pp. 145–182.
- [42] Parmit K. Chilana, Amy Ko, and James O. Wobbrock. "LemonAid: selection-based crowdsourced contextual help for web applications". In: *CHI 2012*. New York City, NY, USA: ACM, 2012, pp. 1549–1558. DOI: [10.1145/2207676.2208620](https://doi.org/10.1145/2207676.2208620). URL: <https://doi.org/10.1145/2207676.2208620>.

- [43] Bhavya Chopra, Yasharth Bajpai, Param Biyani, Gustavo Soares, Arjun Radhakrishna, Chris Parnin, and Sumit Gulwani. *Exploring Interaction Patterns for Debugging: Enhancing Conversational Capabilities of AI-assistants*. 2024. arXiv: 2402.06229 [cs.HC]. URL: <https://arxiv.org/abs/2402.06229>.
- [44] Matteo Ciniselli, Niccolò Puccinelli, Ketai Qiu, and Luca Di Grazia. *From Today's Code to Tomorrow's Symphony: The AI Transformation of Developer's Routine by 2030*. 2024. arXiv: 2405.12731 [cs.SE]. URL: <https://arxiv.org/abs/2405.12731>.
- [45] Michael J. Coblenz, Amy J. Ko, and Brad A. Myers. "JASPER: An Eclipse Plug-in to Facilitate Software Maintenance Tasks". In: *Proceedings of the 2006 OOPSLA Workshop on Eclipse Technology EXchange*. eclipse '06. Portland, Oregon, USA: Association for Computing Machinery, 2006, 65–69. ISBN: 1595936211. DOI: 10.1145/1188835.1188849. URL: <https://doi.org/10.1145/1188835.1188849>.
- [46] C.R. Cook, J.C. Scholtz, and J.C. Spohrer. *Empirical Studies of Programmers: Fifth Workshop : Papers Presented at the Fifth Workshop on Empirical Studies of Programmers, December 3-5, 1993, Palo Alto, CA*. Human/computer interaction. Norwood, NJ, USA: Ablex Publishing Corporation, 1993. ISBN: 9781567500899. URL: <https://books.google.com/books?id=rMmxq8q0CGYC>.
- [47] Joel Cordeiro, Bruno Antunes, and Paulo Gomes. "Context-based search to overcome learning barriers in software development". In: *2012 First International Workshop on Realizing AI Synergies in Software Engineering (RAISE)*. 2012, pp. 47–51. DOI: 10.1109/RAISE.2012.6227970.
- [48] D. Cordes and M. Brown. "The literate-programming paradigm". In: *Computer* 24.6 (1991), pp. 52–61. DOI: 10.1109/2.86838.
- [49] Carlos J Costa, Manuela Aparicio, and Robert Pierce. "Evaluating information sources for computer programming learning and problem solving". In: *Proceedings of the 9th WSEAS International Conference on APPLIED COMPUTER SCIENCE*. 2009, pp. 218–223.
- [50] csillag. *Fuzzy Anchoring*. Hypothes.is. 2013. URL: <https://web.hypothes.is/blog/fuzzy-anchoring/>.
- [51] D. Cubranic, G.C. Murphy, J. Singer, and K.S. Booth. "Hipikat: a project memory for software development". In: *IEEE Transactions on Software Engineering* 31.6 (2005), pp. 446–465. DOI: 10.1109/TSE.2005.71.
- [52] Alex Cummaudo, Rajesh Vasa, John Grundy, and Mohamed Abdelrazek. "Requirements of API Documentation: A Case Study Into Computer Vision Services". In: *IEEE Transactions on Software Engineering* (2020), pp. 1–1. DOI: 10.1109/TSE.2020.3047088. URL: <https://doi.org/10.1109/TSE.2020.3047088>.

- [53] Barthélemy Dagenais and Martin P. Robillard. “Creating and Evolving Developer Documentation: Understanding the Decisions of Open Source Contributors”. In: *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering*. FSE '10. Santa Fe, New Mexico, USA: Association for Computing Machinery, 2010, 127–136. ISBN: 9781605587912. DOI: [10.1145/1882291.1882312](https://doi.org/10.1145/1882291.1882312). URL: <https://doi.org/10.1145/1882291.1882312>.
- [54] Liwei Dai, Wayne G. Lutters, and Carlie Bower. “Why use memo for all? restructuring mobile applications to support informal note taking”. In: *CHI '05 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '05. Portland, OR, USA: Association for Computing Machinery, 2005, 1320–1323. ISBN: 1595930027. DOI: [10.1145/1056808.1056906](https://doi.org/10.1145/1056808.1056906). URL: <https://doi.org/10.1145/1056808.1056906>.
- [55] Arghavan Moradi Dakhel, Amin Nikanjam, Foutse Khomh, Michel C. Desmarais, and Hironori Washizaki. “Generative AI for Software Development: A Family of Studies on Code Generation”. In: *Generative AI for Effective Software Development*. Ed. by Anh Nguyen-Duc, Pekka Abrahamsson, and Foutse Khomh. Cham: Springer Nature Switzerland, 2024, pp. 151–172. ISBN: 978-3-031-55642-5. DOI: [10.1007/978-3-031-55642-5_7](https://doi.org/10.1007/978-3-031-55642-5_7). URL: https://doi.org/10.1007/978-3-031-55642-5_7.
- [56] Uri Dekel and James D. Herbsleb. “Reading the documentation of invoked API functions in program comprehension”. In: *2009 IEEE 17th International Conference on Program Comprehension*. New York City, NY, USA: IEEE, 2009, pp. 168–177. DOI: [10.1109/ICPC.2009.5090040](https://doi.org/10.1109/ICPC.2009.5090040). URL: <https://doi.org/10.1109/10.1109/ICPC.2009.5090040>.
- [57] R. DeLine, M. Czerwinski, B. Meyers, G. Venolia, S. Drucker, and G. Robertson. “Code Thumbnails: Using Spatial Memory to Navigate Source Code”. In: *Visual Languages and Human-Centric Computing (VL/HCC'06)*. 2006, pp. 11–18. DOI: [10.1109/VLHCC.2006.14](https://doi.org/10.1109/VLHCC.2006.14).
- [58] Robert Deline, Mary Czerwinski, and George Robertson. “Easing program comprehension by sharing navigation data”. In: *VLHCC 2005*. New York City, NY, USA: IEEE, 2005, pp. 241–248. DOI: [10.1109/VLHCC.2005.32](https://doi.org/10.1109/VLHCC.2005.32). URL: <https://doi.org/10.1109/VLHCC.2005.32>.
- [59] Robert DeLine and Kael Rowan. “Code canvas: zooming towards better development environments”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 2*. ICSE '10. Cape Town, South Africa: Association for Computing Machinery, 2010, 207–210. ISBN: 9781605587196. DOI: [10.1145/1810295.1810331](https://doi.org/10.1145/1810295.1810331). URL: <https://doi.org/10.1145/1810295.1810331>.

- [60] Robert DeLine, Gina Venolia, and Kael Rowan. "Software development with code maps". In: *Commun. ACM* 53.8 (2010), 48–54. ISSN: 0001-0782. DOI: [10.1145/1787234.1787250](https://doi.org/10.1145/1787234.1787250). URL: <https://doi.org/10.1145/1787234.1787250>.
- [61] Google Developers. *Cloud Firestore: Store and sync app data at global scale*. Google LLC. 2022. URL: <https://firebase.google.com/products/firestore>.
- [62] Rui Ding, Hucheng Zhou, Jian-Guang Lou, Hongyu Zhang, Qingwei Lin, Qiang Fu, Dongmei Zhang, and Tao Xie. "Log2: A Cost-Aware Logging Mechanism for Performance Diagnosis". In: *2015 USENIX Annual Technical Conference (USENIX ATC 15)*. Santa Clara, CA: USENIX Association, July 2015, pp. 139–150. ISBN: 978-1-931971-225. URL: <https://www.usenix.org/conference/atc15/technical-session/presentation/ding>.
- [63] Ekwa Duala-Ekoko and Martin P. Robillard. "Asking and answering questions about unfamiliar APIs: An exploratory study". In: *ICSE 2012*. New York City, NY, USA: IEEE, 2012, pp. 266–276.
- [64] Ralph H. Earle, Mark A. Rosso, and Kathryn E. Alexander. "User preferences of software documentation genres". In: *Proceedings of the 33rd Annual International Conference on the Design of Communication*. SIGDOC '15. Limerick, Ireland: Association for Computing Machinery, 2015. ISBN: 9781450336482. DOI: [10.1145/2775441.2775457](https://doi.org/10.1145/2775441.2775457). URL: <https://doi.org/10.1145/2775441.2775457>.
- [65] S.C. Eick, J.L. Steffen, and E.E. Sumner. "Seesoft—a tool for visualizing line oriented software statistics". In: *IEEE Transactions on Software Engineering* 18.11 (1992), pp. 957–968. DOI: [10.1109/32.177365](https://doi.org/10.1109/32.177365).
- [66] Daniel S. Eisenberg, Jeffrey Stylos, and Brad A. Myers. "Apatite: A New Interface for Exploring APIs". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: Association for Computing Machinery, 2010, 1331–1334. ISBN: 9781605589299. DOI: [10.1145/1753326.1753525](https://doi.org/10.1145/1753326.1753525). URL: <https://doi.org/10.1145/1753326.1753525>.
- [67] Elastic. *Free and Open Search: Elasticsearch*. Elastic. 2021. URL: <https://www.elastic.co/>.
- [68] Stefan Endrikat, Stefan Hanenberg, Romain Robbes, and Andreas Stefik. "How do API documentation and static typing affect API usability?" In: *ICSE 2014*. New York City, NY, USA: ACM, 2014, pp. 632–642. DOI: [10.1145/2568225.2568299](https://doi.org/10.1145/2568225.2568299). URL: <https://doi.org/10.1145/2568225.2568299>.
- [69] Facebook. *React - A JavaScript library for building user interfaces*. 2023. URL: <https://reactjs.org/>.
- [70] Jingchao Fang, Yanhao Wang, Chi-Lan Yang, and Hao-Chuan Wang. "NoteCoStruct: Powering Online Learners with Socially Scaffolded Note Taking and Sharing". In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing

- Machinery, 2021. Chap. Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–5. ISBN: 9781450380959. URL: <https://doi.org/10.1145/3411763.3451694>.
- [71] Kasra Ferdowsi, Ruanqianqian (Lisa) Huang, Michael B. James, Nadia Polikarpova, and Sorin Lerner. “Validating AI-Generated Code with Live Programming”. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. CHI '24. Honolulu, HI, USA: Association for Computing Machinery, 2024. ISBN: 9798400703300. DOI: [10.1145/3613904.3642495](https://doi.org/10.1145/3613904.3642495). URL: <https://doi.org/10.1145/3613904.3642495>.
- [72] Kristie Fisher, Scott Counts, and Aniket Kittur. “Distributed sensemaking: improving sensemaking by leveraging the efforts of previous users”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, 247–256. ISBN: 9781450310154. DOI: [10.1145/2207676.2207711](https://doi.org/10.1145/2207676.2207711). URL: <https://doi.org/10.1145/2207676.2207711>.
- [73] Scott D. Fleming, Chris Scaffidi, David Piorkowski, Margaret Burnett, Rachel Bellamy, Joseph Lawrance, and Irwin Kwan. “An Information Foraging Theory Perspective on Tools for Debugging, Refactoring, and Reuse Tasks”. In: *ACM Trans. Softw. Eng. Methodol.* 22.2 (2013). ISSN: 1049-331X. DOI: [10.1145/2430545.2430551](https://doi.org/10.1145/2430545.2430551). URL: <https://doi.org/10.1145/2430545.2430551>.
- [74] Beat Fluri, Michael Wursch, and Harald C Gall. “Do code and comments co-evolve? on the relation between source code and comment changes”. In: *14th Working Conference on Reverse Engineering (WCRE 2007)*. IEEE, 2007, pp. 70–79.
- [75] Daniela Fogli, Giuseppe Fresta, and Piero Mussio. “On electronic annotation and its implementation”. In: *AVI 2004*. New York, NY, USA: ACM, 2004, pp. 98–102. DOI: [10.1145/989863.989877](https://doi.org/10.1145/989863.989877). URL: <https://doi.org/10.1145/989863.989877>.
- [76] Andrew Forward and Timothy C. Lethbridge. “The Relevance of Software Documentation, Tools and Technologies: A Survey”. In: *Proceedings of the 2002 ACM Symposium on Document Engineering*. DocEng '02. McLean, Virginia, USA: Association for Computing Machinery, 2002, 26–33. ISBN: 1581135947. DOI: [10.1145/585058.585065](https://doi.org/10.1145/585058.585065). URL: <https://doi.org/10.1145/585058.585065>.
- [77] Adam Fourney and Meredith Ringel Morris. “Enhancing Technical Q&A Forums with CiteHistory”. In: *ICWSM 2013*. Palo Alto, CA, USA: AAAI, 2013, pp. 1–10.
- [78] Thomas Fritz and Gail C Murphy. “Using information fragments to answer the questions developers ask”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering—Volume 1*. 2010, pp. 175–184.

- [79] Thomas Fritz, Jingwen Ou, Gail C. Murphy, and Emerson Murphy-Hill. "A Degree-of-Knowledge Model to Capture Source Code Familiarity". In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1. ICSE '10*. Cape Town, South Africa: Association for Computing Machinery, 2010, 385–394. ISBN: 9781605587196. DOI: [10.1145/1806799.1806856](https://doi.org/10.1145/1806799.1806856). URL: <https://doi.org/10.1145/1806799.1806856>.
- [80] J. Froehlich and P. Dourish. "Unifying artifacts and activities in a visual tool for distributed software development teams". In: *Proceedings. 26th International Conference on Software Engineering*. 2004, pp. 387–396. DOI: [10.1109/ICSE.2004.1317461](https://doi.org/10.1109/ICSE.2004.1317461).
- [81] Sina Gholamian and Paul A. S. Ward. *A Comprehensive Survey of Logging in Software: From Logging Statements Automation to Log Mining and Analysis*. 2022. arXiv: [2110.12489](https://arxiv.org/abs/2110.12489) [cs.SE]. URL: <https://arxiv.org/abs/2110.12489>.
- [82] GitHub. *GitHub Copilot*. Microsoft. 2024. URL: <https://github.com/features/copilot>.
- [83] Max Goldman and Robert C. Miller. "Codetrail: Connecting source code and web resources". In: *Journal of Visual Languages and Computing* 20.4 (2009). Special Issue on Best Papers from VL/HCC2008, pp. 223–235. ISSN: 1045-926X. DOI: <https://doi.org/10.1016/j.jvlc.2009.04.003>. URL: <https://www.sciencedirect.com/science/article/pii/S1045926X09000263>.
- [84] Daniel Gonçalves and Joaquim A. Jorge. "Describing documents: what can users tell us?" In: *Proceedings of the 9th International Conference on Intelligent User Interfaces*. IUI '04. Funchal, Madeira, Portugal: Association for Computing Machinery, 2004, 247–249. ISBN: 1581138156. DOI: [10.1145/964442.964494](https://doi.org/10.1145/964442.964494). URL: <https://doi.org/10.1145/964442.964494>.
- [85] Wayne D Gray and Deborah A Boehm-Davis. "Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior." In: *Journal of experimental psychology: applied* 6.4 (2000), p. 322.
- [86] Wayne D. Gray and Wai-Tat Fu. "Ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '01. Seattle, Washington, USA: Association for Computing Machinery, 2001, 112–119. ISBN: 1581133278. DOI: [10.1145/365024.365061](https://doi.org/10.1145/365024.365061). URL: <https://doi.org/10.1145/365024.365061>.
- [87] Sanuri Dananja Gunawardena, Peter Devine, Isabelle Beaumont, Lola Piper Garden, Emerson Murphy-Hill, and Kelly Blincoe. "Destructive Criticism in Software Code Review Impacts Inclusion". In: *Proc. ACM Hum.-Comput. Interact.* 6.CSCW2 (2022). DOI: [10.1145/3555183](https://doi.org/10.1145/3555183). URL: <https://doi.org/10.1145/3555183>.

- [88] Mark Guzdial and Jennifer Turns. “Effective Discussion Through a Computer Mediated Anchored Forum”. In: *The Journal of the Learning Sciences* 9 (4 2000), pp. 437–469. DOI: [10.1207/S15327809JLS0904_3](https://doi.org/10.1207/S15327809JLS0904_3). URL: https://doi.org/10.1207/S15327809JLS0904_3.
- [89] Anja Guzzi, Lile Hattori, Michele Lanza, Martin Pinzger, and Arie van Deursen. “Collective Code Bookmarks for Program Comprehension”. In: *2011 IEEE 19th International Conference on Program Comprehension*. 2011, pp. 101–110. DOI: [10.1109/ICPC.2011.19](https://doi.org/10.1109/ICPC.2011.19).
- [90] Anja Guzzi, Martin Pinzger, and Arie van Deursen. “Combining micro-blogging and IDE interactions to support developers in their quests”. In: *2010 IEEE International Conference on Software Maintenance*. 2010, pp. 1–5. DOI: [10.1109/ICSM.2010.5609683](https://doi.org/10.1109/ICSM.2010.5609683).
- [91] Sonia Haiduc, Jairo Aponte, Laura Moreno, and Andrian Marcus. “On the use of automated text summarization techniques for summarizing source code”. In: *2010 17th Working Conference on Reverse Engineering*. IEEE. 2010, pp. 35–44.
- [92] Björn Hartmann, Mark Dhillon, and Matthew K. Chan. “HyperSource: Bridging the Gap between Source and Code-Related Web Sites”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’11. Vancouver, BC, Canada: Association for Computing Machinery, 2011, 2207–2210. ISBN: 9781450302289. DOI: [10.1145/1978942.1979263](https://doi.org/10.1145/1978942.1979263). URL: <https://doi.org/10.1145/1978942.1979263>.
- [93] Hideaki Hata, Christoph Treude, Raula Gaikovina Kula, and Takashi Ishio. “9.6 Million Links in Source Code Comments: Purpose, Evolution, and Decay”. In: *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 2019, pp. 1211–1221. DOI: [10.1109/ICSE.2019.00123](https://doi.org/10.1109/ICSE.2019.00123).
- [94] Andrew Head, Elena L. Glassman, Björn Hartmann, and Marti A. Hearst. “Interactive Extraction of Examples from Existing Code”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. Montreal QC, Canada: Association for Computing Machinery, 2018, 1–12. ISBN: 9781450356206. DOI: [10.1145/3173574.3173659](https://doi.org/10.1145/3173574.3173659). URL: <https://doi.org/10.1145/3173574.3173659>.
- [95] Andrew Head, Fred Hohman, Titus Barik, Steven M. Drucker, and Robert DeLine. “Managing Messes in Computational Notebooks”. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. Glasgow, Scotland UK: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: [10.1145/3290605.3300500](https://doi.org/10.1145/3290605.3300500). URL: <https://doi.org/10.1145/3290605.3300500>.

- [96] Andrew Head, Jason Jiang, James Smith, Marti A. Hearst, and Björn Hartmann. “Composing Flexibly-Organized Step-by-Step Tutorials from Linked Source Code, Snippets, and Outputs”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, 1–12. ISBN: 9781450367080. DOI: [10.1145/3313831.3376798](https://doi.org/10.1145/3313831.3376798). URL: <https://doi.org/10.1145/3313831.3376798>.
- [97] Andrew Head, Caitlin Sadowski, Emerson Murphy-Hill, and Andrea Knight. “When Not to Comment: Questions and Tradeoffs with API Documentation for C++ Projects”. In: *Proceedings of the 40th International Conference on Software Engineering*. ICSE '18. Gothenburg, Sweden: Association for Computing Machinery, 2018, 643–653. ISBN: 9781450356381. DOI: [10.1145/3180155.3180176](https://doi.org/10.1145/3180155.3180176). URL: <https://doi.org/10.1145/3180155.3180176>.
- [98] Marti A. Hearst. “What’s Missing from Collaborative Search?” In: *Computer* 47.3 (2014), pp. 58–61. DOI: [10.1109/MC.2014.77](https://doi.org/10.1109/MC.2014.77).
- [99] Ken Hinckley, Xiaojun Bi, Michel Pahud, and Bill Buxton. “Informal Information Gathering Techniques for Active Reading”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, 1893–1896. ISBN: 9781450310154. DOI: [10.1145/2207676.2208327](https://doi.org/10.1145/2207676.2208327). URL: <https://doi.org/10.1145/2207676.2208327>.
- [100] Reid Holmes and Andrew Begel. “Deep intellisense: a tool for rehydrating evaporated information”. In: *Proceedings of the 2008 international working conference on Mining software repositories*. 2008, pp. 23–26.
- [101] Reid Holmes and Andrew Begel. “Deep Intellisense: A Tool for Rehydrating Evaporated Information”. In: *Proceedings of the 2008 International Working Conference on Mining Software Repositories*. MSR '08. Leipzig, Germany: Association for Computing Machinery, 2008, 23–26. ISBN: 9781605580241. DOI: [10.1145/1370750.1370755](https://doi.org/10.1145/1370750.1370755). URL: <https://doi.org/10.1145/1370750.1370755>.
- [102] Lichan Hong and Ed H Chi. “Annotate once, appear anywhere: collective foraging for snippets of interest using paragraph fingerprinting”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2009, pp. 1791–1794.
- [103] Amber Horvath, Sachin Grover, Sihan Dong, Emily Zhou, Finn Voichick, Mary Beth Kery, Shwetha Shinju, Daye Nam, Mariann Nagy, and Brad Myers. “The Long Tail: Understanding the Discoverability of API Functionality”. In: *VLHCC 2019*. New York, NY, USA: IEEE, 2019, pp. 157–161. DOI: [10.1109/VLHCC.2019.8818681](https://doi.org/10.1109/VLHCC.2019.8818681). URL: <https://doi.org/10.1109/VLHCC.2019.8818681>.

- [104] Amber Horvath, Michael Xieyang Liu, River Hendriksen, Connor Shannon, Emma Paterson, Kazi Jawad, Andrew Macvean, and Brad A. Myers. "Understanding How Programmers Can Use Annotations on Documentation". In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. DOI: [10.1145/3491102.3502095](https://doi.org/10.1145/3491102.3502095). URL: <https://doi.org/10.1145/3491102.3502095>.
- [105] Amber Horvath, Andrew Macvean, and Brad A. Myers. "Meta-Manager: A Tool for Collecting and Exploring Meta Information about Code". In: *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24), May 11–16, 2024, Honolulu, HI, USA*. CHI '24. Honolulu, HI, USA: Association for Computing Machinery, 2024. DOI: [10.1145/3613904.3642676](https://doi.org/10.1145/3613904.3642676). URL: <https://doi.org/10.1145/3613904.3642676>.
- [106] Amber Horvath, Andrew Macvean, and Brad A. Myers. "Support for Long-Form Documentation Authoring and Maintenance". In: *VL/HCC 2023*. 2023.
- [107] Amber Horvath, Brad Myers, Andrew Macvean, and Imtiaz Rahman. "Using Annotations for Sensemaking About Code". In: *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST '22. Bend, OR, USA: Association for Computing Machinery, 2022. ISBN: 9781450393201. DOI: [10.1145/3526113.3545667](https://doi.org/10.1145/3526113.3545667). URL: <https://doi.org/10.1145/3526113.3545667>.
- [108] Amber Horvath, Mariann Nagy, Finn Voichick, Mary Beth Kery, and Brad A Myers. "Methods for investigating mental models for learners of APIs". In: *CHI LBW '19*. New York, NY, USA: ACM, 2019, pp. 1–6.
- [109] Daqing Hou, Patricia Jablonski, and Ferosh Jacob. "CnP: Towards an environment for the proactive management of copy-and-paste programming". In: *2009 IEEE 17th International Conference on Program Comprehension*. 2009, pp. 238–242. DOI: [10.1109/ICPC.2009.5090049](https://doi.org/10.1109/ICPC.2009.5090049).
- [110] Qiao Huang, Xin Xia, Zhenchang Xing, David Lo, and Xinyu Wang. "API Method Recommendation without Worrying about the Task-API Knowledge Gap". In: *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. ASE '18. Montpellier, France: Association for Computing Machinery, 2018, 293–304. ISBN: 9781450359375. DOI: [10.1145/3238147.3238191](https://doi.org/10.1145/3238147.3238191). URL: <https://doi.org/10.1145/3238147.3238191>.
- [111] Hypothes.is. *Hypothes.is: Annotate the web, with anyone, anywhere*. Hypothes.is. 2012. URL: <https://web.hypothes.is/>.
- [112] Peiling Jiang, Fuling Sun, and Haijun Xia. "Log-It: Supporting Programming with Interactive, Contextual, Structured, and Visual Logs". In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI '23. Hamburg, Germany: Association for Computing Machinery, 2023. ISBN: 9781450394215.

- DOI: [10.1145/3544548.3581403](https://doi.org/10.1145/3544548.3581403). URL: <https://doi.org/10.1145/3544548.3581403>.
- [113] An Ju, Hitesh Sajnani, Scot Kelly, and Kim Herzig. “A case study of onboarding in software teams: Tasks and strategies”. In: *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 2021, pp. 613–623.
- [114] Ján Juhár. “Supporting Source Code Annotations with Metadata-Aware Development Environment”. In: *2019 Federated Conference on Computer Science and Information Systems (FedCSIS)*. 2019, pp. 411–420. DOI: [10.15439/2019F161](https://doi.org/10.15439/2019F161).
- [115] Md Mahir Asef Kabir, Sk Adnan Hassan, Xiaoyin Wang, Ying Wang, Hai Yu, and Na Meng. *An empirical study of ChatGPT-3.5 on question answering and code maintenance*. 2023. arXiv: [2310.02104](https://arxiv.org/abs/2310.02104) [cs.SE]. URL: <https://arxiv.org/abs/2310.02104>.
- [116] Vaiva Kalnikaitė and Steve Whittaker. “Software or wetware? discovering when and why people use digital prosthetic memory”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’07. San Jose, California, USA: Association for Computing Machinery, 2007, 71–80. ISBN: 9781595935939. DOI: [10.1145/1240624.1240635](https://doi.org/10.1145/1240624.1240635). URL: <https://doi.org/10.1145/1240624.1240635>.
- [117] David R. Karger and Dennis Quan. “Haystack: a user interface for creating, browsing, and organizing arbitrary semistructured information”. In: *CHI ’04 Extended Abstracts on Human Factors in Computing Systems*. CHI EA ’04. Vienna, Austria: Association for Computing Machinery, 2004, 777–778. ISBN: 1581137036. DOI: [10.1145/985921.985931](https://doi.org/10.1145/985921.985931). URL: <https://doi.org/10.1145/985921.985931>.
- [118] Mik Kersten and Gail C. Murphy. “Mylar: A Degree-of-Interest Model for IDEs”. In: *Proceedings of the 4th International Conference on Aspect-Oriented Software Development*. AOSD ’05. Chicago, Illinois: Association for Computing Machinery, 2005, 159–168. ISBN: 1595930426. DOI: [10.1145/1052898.1052912](https://doi.org/10.1145/1052898.1052912). URL: <https://doi.org/10.1145/1052898.1052912>.
- [119] Mary Beth Kery, Amber Horvath, and Brad Myers. “Variolite: Supporting Exploratory Programming by Data Scientists”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. Denver, Colorado, USA: Association for Computing Machinery, 2017, 1265–1276. ISBN: 9781450346559. DOI: [10.1145/3025453.3025626](https://doi.org/10.1145/3025453.3025626). URL: <https://doi.org/10.1145/3025453.3025626>.
- [120] Mary Beth Kery, Bonnie E. John, Patrick O’Flaherty, Amber Horvath, and Brad A. Myers. “Towards Effective Foraging by Data Scientists to Find Past Analysis Choices”. In: *Proceedings of the 2019 CHI Conference on Human Factors*

- in Computing Systems*. CHI '19. Glasgow, Scotland UK: Association for Computing Machinery, 2019, 1–13. ISBN: 9781450359702. DOI: [10.1145/3290605.3300322](https://doi.org/10.1145/3290605.3300322). URL: <https://doi.org/10.1145/3290605.3300322>.
- [121] Mary Beth Kery, Marissa Radensky, Mahima Arya, Bonnie E. John, and Brad A. Myers. “The Story in the Notebook: Exploratory Data Science Using a Literate Programming Tool”. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, 1–11. ISBN: 9781450356206. DOI: [10.1145/3173574.3173748](https://doi.org/10.1145/3173574.3173748). URL: <https://doi.org/10.1145/3173574.3173748>.
- [122] Ranim Khojah, Mazen Mohamad, Philipp Leitner, and Francisco Gomes de Oliveira Neto. *Beyond Code Generation: An Observational Study of ChatGPT Usage in Software Engineering Practice*. 2024. arXiv: [2404.14901](https://arxiv.org/abs/2404.14901) [cs.SE]. URL: <https://arxiv.org/abs/2404.14901>.
- [123] Raphaël Khoury, Anderson R Avila, Jacob Brunelle, and Baba Mamadou Camara. “How Secure is Code Generated by ChatGPT?” In: *arXiv preprint arXiv:2304.09655* (2023).
- [124] Aniket Kittur, Andrew M. Peters, Abdigani Diriyee, Trupti Telang, and Michael R. Bove. “Costs and benefits of structured information foraging”. In: *CHI 2013*. New York, NY, USA: ACM, 2013, pp. 2989–2998.
- [125] D. E. Knuth. “Literate Programming”. In: *The Computer Journal* 27.2 (Jan. 1984), pp. 97–111. ISSN: 0010-4620. DOI: [10.1093/comjnl/27.2.97](https://doi.org/10.1093/comjnl/27.2.97). eprint: <https://academic.oup.com/comjnl/article-pdf/27/2/97/981657/270097.pdf>. URL: <https://doi.org/10.1093/comjnl/27.2.97>.
- [126] Amy J. Ko, Htet Aung, and Brad A. Myers. “Eliciting Design Requirements for Maintenance-Oriented IDEs: A Detailed Study of Corrective and Perfective Maintenance Tasks”. In: *Proceedings of the 27th International Conference on Software Engineering*. ICSE '05. St. Louis, MO, USA: Association for Computing Machinery, 2005, 126–135. ISBN: 1581139632. DOI: [10.1145/1062455.1062492](https://doi.org/10.1145/1062455.1062492). URL: <https://doi.org/10.1145/1062455.1062492>.
- [127] Amy J Ko, Robert DeLine, and Gina Venolia. “Information needs in collocated software development teams”. In: *29th International Conference on Software Engineering (ICSE'07)*. IEEE. 2007, pp. 344–353.
- [128] Amy J Ko, Thomas D LaToza, and Margaret M Burnett. “A practical guide to controlled experiments of software engineering tools with human participants”. In: *Empirical Software Engineering* 20.1 (2015), pp. 110–141.
- [129] Amy J. Ko and Brad A. Myers. “Debugging Reinvented: Asking and Answering Why and Why Not Questions about Program Behavior”. In: *Proceedings of the 30th International Conference on Software Engineering*. ICSE '08. Leipzig, Germany: Association for Computing Machinery, 2008, 301–310. ISBN: 9781605580791.

- DOI: [10.1145/1368088.1368130](https://doi.org/10.1145/1368088.1368130). URL: <https://doi.org/10.1145/1368088.1368130>.
- [130] Amy J. Ko and Brad A. Myers. “Designing the Whyline: A Debugging Interface for Asking Questions about Program Behavior”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: Association for Computing Machinery, 2004, 151–158. ISBN: 1581137028. DOI: [10.1145/985692.985712](https://doi.org/10.1145/985692.985712). URL: <https://doi.org/10.1145/985692.985712>.
- [131] Amy J. Ko and Brad A. Myers. “Finding Causes of Program Output with the Java Whyline”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '09. Boston, MA, USA: Association for Computing Machinery, 2009, 1569–1578. ISBN: 9781605582467. DOI: [10.1145/1518701.1518942](https://doi.org/10.1145/1518701.1518942). URL: <https://doi.org/10.1145/1518701.1518942>.
- [132] Amy J. Ko, Brad A. Myers, Michael J. Coblenz, and Htet Htet Aung. “An Exploratory Study of How Developers Seek, Relate, and Collect Relevant Information during Software Maintenance Tasks”. In: *IEEE Transactions on Software Engineering* 32.12 (2006), pp. 971–987. DOI: [10.1109/TSE.2006.116](https://doi.org/10.1109/TSE.2006.116).
- [133] Amy J. Ko and Yann Riche. “The role of conceptual knowledge in API usability”. In: *2011 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 2011, pp. 173–176. DOI: [10.1109/VLHCC.2011.6070395](https://doi.org/10.1109/VLHCC.2011.6070395).
- [134] Amy J. Ko and Bob Uttl. “Individual differences in program comprehension strategies in unfamiliar programming systems”. In: *11th Annual Workshop on Program Comprehension*. New York, NY, USA: IEEE, 2003, pp. 175–184. DOI: [10.1109/WPC.2003.1199201](https://doi.org/10.1109/WPC.2003.1199201). URL: <https://doi.org/10.1109/WPC.2003.1199201>.
- [135] Oleksii Kononenko, Olga Baysal, and Michael W. Godfrey. “Code Review Quality: How Developers See It”. In: *Proceedings of the 38th International Conference on Software Engineering*. ICSE '16. Austin, Texas: Association for Computing Machinery, 2016, 1028–1038. ISBN: 9781450339001. DOI: [10.1145/2884781.2884840](https://doi.org/10.1145/2884781.2884840). URL: <https://doi.org/10.1145/2884781.2884840>.
- [136] Mik Lamming, Peter Brown, Kathleen Carter, Margery Eldridge, Mike Flynn, Gifford Louie, Peter Robinson, and Abigail Sellen. “The design of a human memory prosthesis”. In: *The Computer Journal* 37.3 (1994), pp. 153–163.
- [137] Mark W Lansdale. “The psychology of personal information management”. In: *Applied ergonomics* 19.1 (1988), pp. 55–66.
- [138] Thomas D. LaToza, David Garlan, James D. Herbsleb, and Brad A. Myers. “Program comprehension as fact finding”. In: *ESEC-FSE 2007*. New York, NY, USA: ACM, 2007, pp. 361–270.

- [139] Thomas D. LaToza and Brad A. Myers. "Hard-to-Answer Questions about Code". In: *Evaluation and Usability of Programming Languages and Tools*. PLATEAU '10. Reno, Nevada: Association for Computing Machinery, 2010. ISBN: 9781450305471. DOI: [10.1145/1937117.1937125](https://doi.org/10.1145/1937117.1937125). URL: <https://doi.org/10.1145/1937117.1937125>.
- [140] Thomas D. LaToza, Gina Venolia, and Robert DeLine. "Maintaining mental models: a study of developer work habits". In: *Proceedings of the 28th International Conference on Software Engineering*. ICSE '06. Shanghai, China: Association for Computing Machinery, 2006, 492–501. ISBN: 1595933751. DOI: [10.1145/1134285.1134355](https://doi.org/10.1145/1134285.1134355). URL: <https://doi.org/10.1145/1134285.1134355>.
- [141] Tom Lauwaerts, Carlos Rojas Castillo, Robbert Gurdeep Singh, Matteo Marra, Christophe Scholliers, and Elisa Gonzalez Boix. "Event-Based Out-of-Place Debugging". In: *Proceedings of the 19th International Conference on Managed Programming Languages and Runtimes*. MPLR '22. Brussels, Belgium: Association for Computing Machinery, 2022, 85–97. ISBN: 9781450396967. DOI: [10.1145/3546918.3546920](https://doi.org/10.1145/3546918.3546920). URL: <https://doi.org/10.1145/3546918.3546920>.
- [142] Joseph Lawrance, Christopher Bogart, Margaret Burnett, Rachel Bellamy, Kyle Rector, and Scott D. Fleming. "How Programmers Debug, Revisited: An Information Foraging Theory Perspective". In: *IEEE Transactions on Software Engineering* 39.2 (2013), pp. 197–215. DOI: [10.1109/TSE.2010.111](https://doi.org/10.1109/TSE.2010.111).
- [143] Timothy C. Lethbirdge, Janice Singer, and Andrew Forward. "How software engineers use documentation: the state of the practice". In: *IEEE Software* 20 (6 Nov. 2003), pp. 35–39. DOI: [10.1109/MS.2003.1241364](https://doi.org/10.1109/MS.2003.1241364). URL: <https://doi.org/10.1109/MS.2003.1241364>.
- [144] Heng Li, Weiyi Shang, Bram Adams, Mohammed Sayagh, and Ahmed E. Hassan. "A Qualitative Study of the Benefits and Costs of Logging From Developers' Perspectives". In: *IEEE Transactions on Software Engineering* 47.12 (2021), pp. 2858–2873. DOI: [10.1109/TSE.2020.2970422](https://doi.org/10.1109/TSE.2020.2970422).
- [145] Hongwei Li, Zhenchang Xing, Xin Peng, and Wenyun Zhao. "What help do developers seek, when and how?" In: *2013 20th Working Conference on Reverse Engineering (WCRE)*. 2013, pp. 142–151. DOI: [10.1109/WCRE.2013.6671289](https://doi.org/10.1109/WCRE.2013.6671289).
- [146] Xiangqi Li and Matthew Flatt. "Medic: metaprogramming and trace-oriented debugging". In: *Proceedings of the Workshop on Future Programming*. FPW 2015. Pittsburgh, PA, USA: Association for Computing Machinery, 2015, 7–14. ISBN: 9781450339056. DOI: [10.1145/2846656.2846658](https://doi.org/10.1145/2846656.2846658). URL: <https://doi.org/10.1145/2846656.2846658>.
- [147] Jenny T. Liang, Maryam Arab, Minhyuk Ko, Amy J. Ko, and Thomas D. LaToza. "A Qualitative Study on the Implementation Design Decisions of Developers". In: *Proceedings of the 45th International Conference on Software Engineering*. ICSE '23. Melbourne, Victoria, Australia: IEEE Press, 2023, 435–447.

- ISBN: 9781665457019. DOI: [10.1109/ICSE48619.2023.00047](https://doi.org/10.1109/ICSE48619.2023.00047). URL: <https://doi.org/10.1109/ICSE48619.2023.00047>.
- [148] Jenny T. Liang, Chenyang Yang, and Brad A. Myers. “A Large-Scale Survey on the Usability of AI Programming Assistants: Successes and Challenges”. In: *Proceedings of the 46th International Conference on Software Engineering*. ICSE '24. To appear. Lisbon, Portugal, 2024. URL: [arXiv:2303.17125](https://arxiv.org/abs/2303.17125).
- [149] Mario Linares-Vásquez, Luis Fernando Cortés-Coy, Jairo Aponte, and Denys Poshyvanyk. “ChangeScribe: A Tool for Automatically Generating Commit Messages”. In: *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*. Vol. 2. 2015, pp. 709–712. DOI: [10.1109/ICSE.2015.229](https://doi.org/10.1109/ICSE.2015.229).
- [150] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. “Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation”. In: *arXiv preprint arXiv:2305.01210* (2023).
- [151] Michael Xieyang Liu, Jane Hsieh, Nathan Hahn, Angelina Zhou, Emily Deng, Shaun Burley, Cynthia Taylor, Aniket Kittur, and Brad A. Myers. “Unakite: Scaffolding Developers’ Decision-Making Using the Web”. In: *UIST 2019*. New York, NY, USA: ACM, 2019, pp. 67–80.
- [152] Michael Xieyang Liu, Aniket Kittur, and Brad A. Myers. “Crystalline: Lowering the Cost for Developers to Collect and Organize Information for Decision Making”. In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. DOI: [10.1145/3491102.3501968](https://doi.org/10.1145/3491102.3501968). URL: <https://doi.org/10.1145/3491102.3501968>.
- [153] Michael Xieyang Liu, Aniket Kittur, and Brad A. Myers. “To Reuse or Not To Reuse? A Framework and System for Evaluating Summarized Knowledge”. In: *Proc. ACM Hum.-Comput. Interact.* 5.CSCW1 (2021). DOI: [10.1145/3449240](https://doi.org/10.1145/3449240). URL: <https://doi.org/10.1145/3449240>.
- [154] Zhongxin Liu, Xin Xia, Christoph Treude, David Lo, and Shanping Li. “Automatic generation of pull request descriptions”. In: *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2019, pp. 176–188.
- [155] Steven Locke, Heng Li, Tse-Hsun Peter Chen, Weiyi Shang, and Wei Liu. “LogAssist: Assisting Log Analysis Through Log Summarization”. In: *IEEE Transactions on Software Engineering* 48.9 (2022), pp. 3227–3241. DOI: [10.1109/TSE.2021.3083715](https://doi.org/10.1109/TSE.2021.3083715).
- [156] Yimeng Ma, Yu Huang, and Kevin Leach. “Breaking the Flow: A Study of Interruptions During Software Engineering Activities”. In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. ICSE '24. Lisbon, Portugal: Association for Computing Machinery, 2024. ISBN: 9798400702174.

- DOI: [10.1145/3597503.3639079](https://doi.org/10.1145/3597503.3639079). URL: <https://doi.org/10.1145/3597503.3639079>.
- [157] Walid Maalej and Hans-Jorg Happel. “From work to word: How do software developers describe their work?” In: *2009 6th IEEE International Working Conference on Mining Software Repositories*. 2009, pp. 121–130. DOI: [10.1109/MSR.2009.5069490](https://doi.org/10.1109/MSR.2009.5069490).
- [158] Walid Maalej and Martin P. Robillard. “Patterns of Knowledge in API Reference Documentation”. In: *IEEE Transactions on Software Engineering* 39.9 (2013), pp. 1264–1282. DOI: [10.1109/TSE.2013.12](https://doi.org/10.1109/TSE.2013.12).
- [159] Walid Maalej, Rebecca Tiarks, Tobias Roehm, and Rainer Koschke. “On the Comprehension of Program Comprehension”. In: *Transactions on Software Engineering* 23 (4 2014), pp. 1–37. DOI: [10.1145/2622669](https://doi.org/10.1145/2622669). URL: <https://doi.org/10.1145/2622669>.
- [160] Catherine C. Marshall and Sara Bly. “Saving and Using Encountered Information: Implications for Electronic Periodicals”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’05. Portland, Oregon, USA: Association for Computing Machinery, 2005, 111–120. ISBN: 1581139985. DOI: [10.1145/1054972.1054989](https://doi.org/10.1145/1054972.1054989). URL: <https://doi.org/10.1145/1054972.1054989>.
- [161] Anneliese von Mayrhauser and A Marie Vans. “Hypothesis-driven understanding processes during corrective maintenance of large scale software”. In: *1997 Proceedings International Conference on Software Maintenance*. IEEE. 1997, pp. 12–20.
- [162] Collin McMillan, Denys Poshyvanyk, and Mark Grechanik. “Recommending Source Code Examples via API Call Usages and Documentation”. In: *Proceedings of the 2nd International Workshop on Recommendation Systems for Software Engineering*. RSSE ’10. Cape Town, South Africa: Association for Computing Machinery, 2010, 21–25. ISBN: 9781605589749. DOI: [10.1145/1808920.1808925](https://doi.org/10.1145/1808920.1808925). URL: <https://doi.org/10.1145/1808920.1808925>.
- [163] Michael Meng, Stephanie M Steinhard, and Andreas Schubert. “How developers use API documentation: an observation study”. In: *Communication Design Quarterly* 7 (2 2019), pp. 40–49. DOI: [10.1145/3358931.3358937](https://doi.org/10.1145/3358931.3358937). URL: <https://doi.org/10.1145/3358931.3358937>.
- [164] Michael Meng, Stephanie Steinhardt, and Andreas Schubert. “How developers use API documentation: an observation study”. In: *Commun. Des. Q. Rev* 7.2 (2019), 40–49. DOI: [10.1145/3358931.3358937](https://doi.org/10.1145/3358931.3358937). URL: <https://doi.org/10.1145/3358931.3358937>.
- [165] Microsoft. *TypeScript: JavaScript with Syntax for Types*. Microsoft. 2023. URL: <https://www.typescriptlang.org/>.

- [166] Microsoft. *Visual Studio Code*. Microsoft. 2023. URL: <https://code.visualstudio.com/>.
- [167] Microsoft. *Webview API | Visual Studio Code Extension API*. Microsoft. 2023. URL: <https://code.visualstudio.com/api/extension-guides/webview>.
- [168] Edward Misback, Zachary Tatlock, and Steven L Tanimoto. “Magic Markup: Maintaining Document-External Markup with an LLM”. In: *arXiv preprint arXiv:2403.03481* (2024).
- [169] Hussein Mozannar, Gagan Bansal, Adam Fourney, and Eric Horvitz. “Reading Between the Lines: Modeling User Behavior and Costs in AI-Assisted Programming”. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. CHI ’24. Honolulu, HI, USA: Association for Computing Machinery, 2024. ISBN: 9798400703300. DOI: [10.1145/3613904.3641936](https://doi.org/10.1145/3613904.3641936). URL: <https://doi.org/10.1145/3613904.3641936>.
- [170] Gail C Murphy, Mik Kersten, Martin P Robillard, and Davor Čubranić. “The emergent structure of development tasks”. In: *European Conference on Object-Oriented Programming*. Springer. 2005, pp. 33–48.
- [171] Laurie Murphy, Gary Lewandowski, Renée McCauley, Beth Simon, Lynda Thomas, and Carol Zander. “Debugging: the good, the bad, and the quirky – a qualitative analysis of novices’ strategies”. In: *Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education*. SIGCSE ’08. Portland, OR, USA: Association for Computing Machinery, 2008, 163–167. ISBN: 9781595937995. DOI: [10.1145/1352135.1352191](https://doi.org/10.1145/1352135.1352191). URL: <https://doi.org/10.1145/1352135.1352191>.
- [172] Emerson Murphy-Hill, Jillian Dicker, Margaret Morrow Hodges, Carolyn D. Egelman, Ciera Jaspán, Lan Cheng, Elizabeth Kammer, Ben Holtz, Matthew A. Jorde, Andrea Knight Dolan, and Collin Green. “Engineering Impacts of Anonymous Author Code Review: A Field Experiment”. In: *IEEE Transactions on Software Engineering* 48.7 (2022), pp. 2495–2509. DOI: [10.1109/TSE.2021.3061527](https://doi.org/10.1109/TSE.2021.3061527).
- [173] Emerson Murphy-Hill, Ciera Jaspán, Carolyn Egelman, and Lan Cheng. “The Pushback Effects of Race, Ethnicity, Gender, and Age in Code Review”. In: *Commun. ACM* 65.3 (2022), 52–57. ISSN: 0001-0782. DOI: [10.1145/3474097](https://doi.org/10.1145/3474097). URL: <https://doi.org/10.1145/3474097>.
- [174] Brad A. Myers and Jeffrey Stylos. “Improving API Usability”. In: *Communications of the ACM* 59.6 (2016), pp. 62–69. DOI: [10.1145/2896587](https://doi.org/10.1145/2896587). URL: <https://doi.org/10.1145/2896587>.
- [175] Daye Nam, Andrew Macvean, Vincent Hellendoorn, Bogdan Vasilescu, and Brad Myers. “Using an LLM to Help With Code Understanding”. In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. ICSE ’24. Lisbon, Portugal: Association for Computing Machinery, 2024. ISBN:

9798400702174. DOI: [10.1145/3597503.3639187](https://doi.org/10.1145/3597503.3639187). URL: <https://doi.org/10.1145/3597503.3639187>.
- [176] Daye Nam, Brad Myers, Bogdan Vasilescu, and Vincent Hellendoorn. “Improving API Knowledge Discovery with ML: A Case Study of Comparable API Methods”. In: *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. 2023, pp. 1890–1906. DOI: [10.1109/ICSE48619.2023.00161](https://doi.org/10.1109/ICSE48619.2023.00161).
- [177] Seyed Mehdi Nasehi, Jonathan Sillito, Frank Maurer, and Chris Burns. “What makes a good code example?: A study of programming Q&A in StackOverflow”. In: *ICSM 2012*. New York, NY, USA: IEEE, 2012, pp. 25–34.
- [178] Mathieu Nassif and Martin P. Robillard. “Revisiting Turnover-Induced Knowledge Loss in Software Projects”. In: *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 2017, pp. 261–272. DOI: [10.1109/ICSME.2017.64](https://doi.org/10.1109/ICSME.2017.64).
- [179] Sydney Nguyen, Hannah McLean Babe, Yangtian Zi, Arjun Guha, Carolyn Jane Anderson, and Molly Q Feldman. “How Beginning Programmers and Code LLMs (Mis)read Each Other”. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. CHI ’24. Honolulu, HI, USA: Association for Computing Machinery, 2024. ISBN: 9798400703300. DOI: [10.1145/3613904.3642706](https://doi.org/10.1145/3613904.3642706). URL: <https://doi.org/10.1145/3613904.3642706>.
- [180] Janet Nykaza, Rhonda Messinger, Fran Boehme, Cherie L. Norman, Matthew Mace, and Manuel Gordon. “What programmers really want: results of a needs assessment for SDK documentation”. In: *SIGDOC 2002*. New York, NY, USA: ACM, 2002, pp. 133–141. DOI: [10.1145/584955.584976](https://doi.org/10.1145/584955.584976). URL: <https://doi.org/10.1145/584955.584976>.
- [181] Observable. *D3 by Observable | The JavaScript library for bespoke data visualization*. 2023. URL: d3js.org.
- [182] Stephen Oney and Joel Brandt. “Codelets: Linking Interactive Documentation and Example Code in the Editor”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’12. Austin, Texas, USA: Association for Computing Machinery, 2012, 2697–2706. ISBN: 9781450310154. DOI: [10.1145/2207676.2208664](https://doi.org/10.1145/2207676.2208664). URL: <https://doi.org/10.1145/2207676.2208664>.
- [183] OpenAI. *ChatGPT*. OpenAI. 2024. URL: <https://chat.openai.com/>.
- [184] Dennis Pagano and Walid Maalej. “How Do Developers Blog? An Exploratory Study”. In: *Proceedings of the 8th Working Conference on Mining Software Repositories*. MSR ’11. Waikiki, Honolulu, HI, USA: Association for Computing Machinery, 2011, 123–132. ISBN: 9781450305747. DOI: [10.1145/1985441.1985461](https://doi.org/10.1145/1985441.1985461). URL: <https://doi.org/10.1145/1985441.1985461>.

- [185] Dennis Pagano and Walid Maalej. “How do open source communities blog?” In: *Empirical Software Engineering* 18.6 (Dec. 2013), pp. 1090–1124. ISSN: 1573-7616. DOI: [10.1007/s10664-012-9211-2](https://doi.org/10.1007/s10664-012-9211-2). URL: <https://doi.org/10.1007/s10664-012-9211-2>.
- [186] Soya Park, Amy X. Zhang, and David R. Karger. “Post-literate Programming: Linking Discussion and Code in Software Development Teams”. In: *Adjunct Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST ’18 Adjunct. Berlin, Germany: Association for Computing Machinery, 2018, 51–53. ISBN: 9781450359498. DOI: [10.1145/3266037.3266098](https://doi.org/10.1145/3266037.3266098). URL: <https://doi.org/10.1145/3266037.3266098>.
- [187] C. Parnin and C. Gorg. “Building Usage Contexts During Program Comprehension”. In: *14th IEEE International Conference on Program Comprehension (ICPC’06)*. 2006, pp. 13–22. DOI: [10.1109/ICPC.2006.14](https://doi.org/10.1109/ICPC.2006.14).
- [188] Chris Parnin. *Programmer Interrupted*. ninlabs research. 2013. URL: <https://blog.ninlabs.com/2013/01/programmer-interrupted/>.
- [189] Chris Parnin and Robert DeLine. “Evaluating Cues for Resuming Interrupted Programming Tasks”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2010, 93–102. ISBN: 9781605589299. URL: <https://doi.org/10.1145/1753326.1753342>.
- [190] Chris Parnin, Carsten Görg, and Spencer Rugaber. “CodePad: interactive spaces for maintaining concentration in programming environments”. In: *Proceedings of the 5th International Symposium on Software Visualization*. SOFTVIS ’10. Salt Lake City, Utah, USA: Association for Computing Machinery, 2010, 15–24. ISBN: 9781450300285. DOI: [10.1145/1879211.1879217](https://doi.org/10.1145/1879211.1879217). URL: <https://doi.org/10.1145/1879211.1879217>.
- [191] Chris Parnin and Spencer Rugaber. “Resumption strategies for interrupted programming tasks”. In: *Software Quality Journal* 19.1 (2011), pp. 5–34.
- [192] Chris Parnin, Christoph Treude, and Margaret-Anne Storey. “Blogging developer knowledge: Motivations, challenges, and future directions”. In: *2013 21st International Conference on Program Comprehension (ICPC)*. 2013, pp. 211–214. DOI: [10.1109/ICPC.2013.6613850](https://doi.org/10.1109/ICPC.2013.6613850).
- [193] Luca Pascarella, Davide Spadini, Fabio Palomba, Magiel Bruntink, and Alberto Bacchelli. “Information Needs in Contemporary Code Review”. In: *Proc. ACM Hum.-Comput. Interact.* 2.CSCW (2018). DOI: [10.1145/3274404](https://doi.org/10.1145/3274404). URL: <https://doi.org/10.1145/3274404>.
- [194] Michael Perscheid, Benjamin Siegmund, Marcel Taeumel, and Robert Hirschfeld. “Studying the advancement in debugging practice of professional software developers”. In: *Software Quality Journal* 25.1 (Mar. 2017), pp. 83–110. ISSN:

- 1573-1367. DOI: [10.1007/s11219-015-9294-2](https://doi.org/10.1007/s11219-015-9294-2). URL: <https://doi.org/10.1007/s11219-015-9294-2>.
- [195] Piling.js. *The Piling.js Docs*. Piling.js. 2021. URL: <https://piling.js.org/docs/>.
- [196] David Piorkowski, Scott Fleming, Christopher Scaffidi, Christopher Bogart, Margaret Burnett, Bonnie John, Rachel Bellamy, and Calvin Swart. "Reactive Information Foraging: An Empirical Investigation of Theory-Based Recommender Systems for Programmers". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, 1471–1480. ISBN: 9781450310154. DOI: [10.1145/2207676.2208608](https://doi.org/10.1145/2207676.2208608). URL: <https://doi.org/10.1145/2207676.2208608>.
- [197] David Piorkowski, Scott D. Fleming, Christopher Scaffidi, Margaret Burnett, Irwin Kwan, Austin Z. Henley, Jamie Macbeth, Charles Hill, and Amber Horvath. "To fix or to learn? How production bias affects developers' information foraging during debugging". In: *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 2015, pp. 11–20. DOI: [10.1109/ICSM.2015.7332447](https://doi.org/10.1109/ICSM.2015.7332447).
- [198] David Piorkowski, Soya Park, April Yi Wang, Dakuo Wang, Michael Muller, and Felix Portnoy. "How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study". In: *Proc. ACM Hum.-Comput. Interact.* 5.CSCW1 (2021). DOI: [10.1145/3449205](https://doi.org/10.1145/3449205). URL: <https://doi.org/10.1145/3449205>.
- [199] OpenAI Platform. *Overview - OpenAI API*. URL: <https://platform.openai.com/>.
- [200] Ben Popper and David Gibson. *How often do people actually copy and paste from Stack Overflow? Now we know*. URL: <https://stackoverflow.blog/2021/12/30/how-often-do-people-actually-copy-and-paste-from-stack-overflow-now-we-know/>.
- [201] Aniket Potdar and Emad Shihab. "An Exploratory Study on Self-Admitted Technical Debt". In: *2014 IEEE International Conference on Software Maintenance and Evolution*. 2014, pp. 91–100. DOI: [10.1109/ICSM.2014.31](https://doi.org/10.1109/ICSM.2014.31).
- [202] James Prather, Brent N. Reeves, Paul Denny, Brett A. Becker, Juho Leinonen, Andrew Luxton-Reilly, Garrett Powell, James Finnie-Ansley, and Eddie Antonio Santos. "'It's Weird That it Knows What I Want': Usability and Interactions with Copilot for Novice Programmers". In: *ACM Trans. Comput.-Hum. Interact.* 31.1 (2023). ISSN: 1073-0516. DOI: [10.1145/3617367](https://doi.org/10.1145/3617367). URL: <https://doi.org/10.1145/3617367>.
- [203] prettier.io. *prettier - npm*. NPM. 2024. URL: <https://www.npmjs.com/package/prettier>.

- [204] Pooja Rani, Mathias Birrer, Sebastiano Panichella, Mohammad Ghafari, and Oscar Nierstrasz. "What do developers discuss about code comments?" In: *2021 IEEE 21st International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE. 2021, pp. 153–164.
- [205] Sarah Rastkar, Gail C Murphy, and Gabriel Murray. "Summarizing software artifacts: a case study of bug reports". In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 1*. 2010, pp. 505–514.
- [206] Eric Rawn, Jingyi Li, Eric Paulos, and Sarah E. Chasins. "Understanding Version Control as Material Interaction with Quickpose". In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI '23. Hamburg, Germany: Association for Computing Machinery, 2023. ISBN: 9781450394215. DOI: [10.1145/3544548.3581394](https://doi.org/10.1145/3544548.3581394). URL: <https://doi.org/10.1145/3544548.3581394>.
- [207] Steven P Reiss. "Trace-based debugging". In: *International Workshop on Automated and Algorithmic Debugging*. Springer. 1993, pp. 305–314.
- [208] Steven P. Reiss. "Tracking Source Locations". In: *Proceedings of the 30th International Conference on Software Engineering*. ICSE '08. Leipzig, Germany: Association for Computing Machinery, 2008, 11–20. ISBN: 9781605580791. DOI: [10.1145/1368088.1368091](https://doi.org/10.1145/1368088.1368091). URL: <https://doi.org/10.1145/1368088.1368091>.
- [209] Peter C. Rigby, Yue Cai Zhu, Samuel M. Donadelli, and Audris Mockus. "Quantifying and mitigating turnover-induced knowledge loss: case studies of chrome and a project at avaya". In: *Proceedings of the 38th International Conference on Software Engineering*. ICSE '16. Austin, Texas: Association for Computing Machinery, 2016, 1006–1016. ISBN: 9781450339001. DOI: [10.1145/2884781.2884851](https://doi.org/10.1145/2884781.2884851). URL: <https://doi.org/10.1145/2884781.2884851>.
- [210] Martin P. Robillard. "Turnover-Induced Knowledge Loss in Practice". In: *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ESEC/FSE 2021. Athens, Greece: Association for Computing Machinery, 2021, 1292–1302. ISBN: 9781450385626. DOI: [10.1145/3468264.3473923](https://doi.org/10.1145/3468264.3473923). URL: <https://doi.org/10.1145/3468264.3473923>.
- [211] Martin P. Robillard. "What Makes APIs Hard to Learn? Answers from Developers". In: *IEEE Software* 26 (6 Oct. 2009), pp. 27–34. DOI: [10.1109/MS.2009.193](https://doi.org/10.1109/MS.2009.193). URL: <https://doi.org/10.1109/MS.2009.193>.
- [212] Martin P. Robillard and Robert DeLine. "A field study of API learning obstacles". In: *Empirical Software Engineering* 16 (6 2011), pp. 703–732.
- [213] M.P. Robillard, W. Coelho, and G.C. Murphy. "How effective developers investigate source code: an exploratory study". In: *IEEE Transactions on Software Engineering* 30.12 (2004), pp. 889–903. DOI: [10.1109/TSE.2004.101](https://doi.org/10.1109/TSE.2004.101).

- [214] Tobias Roehm, Rebecca Tiarks, Rainer Koschke, and Walid Maalej. "How do professional developers comprehend software?" In: *ICSE 2012*. New York, NY, USA: ACM, 2012, pp. 632–542. DOI: [10.1109/ICSE.2012.6227188](https://doi.org/10.1109/ICSE.2012.6227188). URL: <https://doi.org/10.1109/ICSE.2012.6227188>.
- [215] Daniel M Russell, Mark J Stefik, Peter Pirolli, and Stuart K Card. "The cost structure of sensemaking". In: *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. 1993, pp. 269–276.
- [216] Caitlin Sadowski, Emma Söderberg, Luke Church, Michal Sipko, and Alberto Bacchelli. "Modern Code Review: A Case Study at Google". In: *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice*. ICSE-SEIP '18. Gothenburg, Sweden: Association for Computing Machinery, 2018, 181–190. ISBN: 9781450356596. DOI: [10.1145/3183519.3183525](https://doi.org/10.1145/3183519.3183525). URL: <https://doi.org/10.1145/3183519.3183525>.
- [217] Caitlin Sadowski, Kathryn T Stolee, and Sebastian Elbaum. "How developers search for code: a case study". In: *Proceedings of the 2015 10th joint meeting on foundations of software engineering*. 2015, pp. 191–201.
- [218] Bill N Schilit, Gene Golovchinsky, and Morgan N Price. "Beyond paper: supporting active reading with free form digital ink annotations". In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1998, pp. 249–256.
- [219] Donghwan Shin, Domenico Bianculli, and Lionel Briand. "Effective Removal of Operational Log Messages: an Application to Model Inference". In: *arXiv preprint arXiv:2004.07194* (2020).
- [220] Yusuke Shinyama, Yoshitaka Arahori, and Katsuhiko Gondow. "Analyzing Code Comments to Boost Program Comprehension". In: *2018 25th Asia-Pacific Software Engineering Conference (APSEC)*. 2018, pp. 325–334. DOI: [10.1109/APSEC.2018.00047](https://doi.org/10.1109/APSEC.2018.00047).
- [221] Frank Shipman, Morgan Price, Catherine C. Marshall, and Gene Golovchinsky. "Identifying Useful Passages in Documents Based on Annotation Patterns". In: *International Conference on Theory and Practice of Digital Libraries* 2769 (2003), pp. 101–112. DOI: [10.1007/978-3-540-45175-4_11](https://doi.org/10.1007/978-3-540-45175-4_11). URL: https://doi.org/10.1007/978-3-540-45175-4_11.
- [222] Yulia Shmerlin, Irit Hadar, Doron Kliger, and Hayim Makabee. "To document or not to document? An exploratory study on developers' motivation to document code". In: *Advanced Information Systems Engineering Workshops: CAiSE 2015 International Workshops, Stockholm, Sweden, June 8-9, 2015, Proceedings 27*. Springer. 2015, pp. 100–106.
- [223] Ben Shneiderman. *Software psychology: Human factors in computer and information systems (Winthrop computer systems series)*. Winthrop Publishers, 1980.

- [224] Nischal Shrestha, Titus Barik, and Chris Parnin. “Unravel: A Fluent Code Explorer for Data Wrangling”. In: *The 34th Annual ACM Symposium on User Interface Software and Technology*. 2021, pp. 198–207.
- [225] Stephen Shum and Curtis Cook. “Using Literate Programming to Teach Good Programming Practices”. In: *Proceedings of the Twenty-Fifth SIGCSE Symposium on Computer Science Education*. SIGCSE '94. Phoenix, Arizona, USA: Association for Computing Machinery, 1994, 66–70. ISBN: 0897916468. DOI: [10.1145/191029.191059](https://doi.org/10.1145/191029.191059). URL: <https://doi.org/10.1145/191029.191059>.
- [226] Jonathan Sillito, Gail C. Murphy, and Kris De Volder. “Asking and Answering Questions during a Programming Change Task”. In: *IEEE Transactions on Software Engineering* 34.4 (2008), pp. 434–451. DOI: [10.1109/TSE.2008.26](https://doi.org/10.1109/TSE.2008.26).
- [227] Ananya Singha, Bhavya Chopra, Anirudh Khattry, Sumit Gulwani, Austin Z. Henley, Vu Le, Chris Parnin, Mukul Singh, and Gust Verbruggen. *Semantically Aligned Question and Code Generation for Automated Insight Generation*. 2024. arXiv: [2405.01556](https://arxiv.org/abs/2405.01556) [cs.SE]. URL: <https://arxiv.org/abs/2405.01556>.
- [228] Zéphyrin Soh, Foutse Khomh, Yann-Gaël Guéhéneuc, and Giuliano Antoniol. “Noise in Mylyn interaction traces and its impact on developers and recommendation systems”. In: *Empirical Software Engineering* 23.2 (Apr. 2018), pp. 645–692. ISSN: 1573-7616. DOI: [10.1007/s10664-017-9529-x](https://doi.org/10.1007/s10664-017-9529-x). URL: <https://doi.org/10.1007/s10664-017-9529-x>.
- [229] Jeongju Sohn and Shin Yoo. “FLUCCS: Using Code and Change Metrics to Improve Fault Localization”. In: *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis*. ISSA 2017. Santa Barbara, CA, USA: Association for Computing Machinery, 2017, 273–283. ISBN: 9781450350761. DOI: [10.1145/3092703.3092717](https://doi.org/10.1145/3092703.3092717). URL: <https://doi.org/10.1145/3092703.3092717>.
- [230] Illia Solohubov, Artur Moroz, Mariia Yu Tiahunova, Halyna H Kyrychek, and Stepan Skrupsky. “Accelerating software development with AI: exploring the impact of ChatGPT and GitHub Copilot.” In: *CTE*. 2023, pp. 76–86.
- [231] Cleidson R. B. de Souza, Gema Rodríguez-Pérez, Manaal Basha, Dongwook Yoon, and Ivan Beschastnikh. “The Fine Balance Between Helping With Your Job And Taking It: AI Code Assistants Come To The Fore”. In: *IEEE Software* (2024), pp. 1–6. DOI: [10.1109/MS.2024.3357787](https://doi.org/10.1109/MS.2024.3357787).
- [232] Sergio Cozzetti B. de Souza, Nicolas Anquetil, and Káthia M. de Oliveira. “A Study of the Documentation Essential to Software Maintenance”. In: *Proceedings of the 23rd Annual International Conference on Design of Communication: Documenting and Designing for Pervasive Information*. SIGDOC '05. Coventry, United Kingdom: Association for Computing Machinery, 2005, 68–75. ISBN: 1595931759. DOI: [10.1145/1085313.1085331](https://doi.org/10.1145/1085313.1085331). URL: <https://doi.org/10.1145/1085313.1085331>.

- [233] Sruti Srinivasa Ragavan, Sandeep Kaur Kuttal, Charles Hill, Anita Sarma, David Piorkowski, and Margaret Burnett. "Foraging among an overabundance of similar variants". In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016, pp. 3509–3521.
- [234] Captain Stack. *Captain Stack - Code suggestion for VSCode*. URL: <https://github.com/hieunc229/copilot-clone>.
- [235] Igor Steinmacher, Marco Aurelio Graciotto Silva, Marco Aurelio Gerosa, and David F Redmiles. "A systematic literature review on the barriers faced by newcomers to open source software projects". In: *Information and Software Technology* 59 (2015), pp. 67–85.
- [236] Christoph Johann Stettina and Werner Heijstek. "Necessary and Neglected? An Empirical Study of Internal Documentation in Agile Software Development Teams". In: *Proceedings of the 29th ACM International Conference on Design of Communication*. SIGDOC '11. Pisa, Italy: Association for Computing Machinery, 2011, 159–166. ISBN: 9781450309363. DOI: [10.1145/2038476.2038509](https://doi.org/10.1145/2038476.2038509). URL: <https://doi.org/10.1145/2038476.2038509>.
- [237] Margaret-Anne Storey, Jody Ryall, R. Ian Bull, Del Myers, and Janice Singer. "TODO or to bug". In: *2008 ACM/IEEE 30th International Conference on Software Engineering*. 2008, pp. 251–260. DOI: [10.1145/1368088.1368123](https://doi.org/10.1145/1368088.1368123).
- [238] Margaret-Anne Storey, Jody Ryall, Janice Singer, Del Myers, Li-Te Cheng, and Michael Muller. "How Software Developers Use Tagging to Support Reminding and Refinding". In: *IEEE Transactions on Software Engineering* 35.4 (2009), pp. 470–483. DOI: [10.1109/TSE.2009.15](https://doi.org/10.1109/TSE.2009.15).
- [239] Siddharth Subramanian and Reid Holmes. "Making sense of online code snippets". In: *2013 10th Working Conference on Mining Software Repositories (MSR)*. 2013, pp. 85–88. DOI: [10.1109/MSR.2013.6624012](https://doi.org/10.1109/MSR.2013.6624012).
- [240] Lin Tan, Ding Yuan, Gopal Krishna, and Yuanyuan Zhou. "/* iComment: Bugs or bad comments?*". In: *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*. 2007, pp. 145–158.
- [241] Shin Hwei Tan, Darko Marinov, Lin Tan, and Gary T Leavens. "@ tcomment: Testing javadoc comments to detect comment-code inconsistencies". In: *2012 IEEE Fifth International Conference on Software Testing, Verification and Validation*. IEEE. 2012, pp. 260–269.
- [242] Minaoar Hossain Tanzil, Junaed Younus Khan, and Gias Uddin. "ChatGPT Incorrectness Detection in Software Reviews". In: *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. ICSE '24. Lisbon, Portugal: Association for Computing Machinery, 2024. ISBN: 9798400702174. DOI: [10.1145/3597503.3639194](https://doi.org/10.1145/3597503.3639194). URL: <https://doi.org/10.1145/3597503.3639194>.

- [243] Craig S. Tashman and W. Keith Edwards. "Active Reading and Its Discontents: The Situations, Problems and Ideas of Readers". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: Association for Computing Machinery, 2011, 2927–2936. ISBN: 9781450302289. DOI: [10.1145/1978942.1979376](https://doi.org/10.1145/1978942.1979376). URL: <https://doi.org/10.1145/1978942.1979376>.
- [244] Grace Taylor and Steven Clarke. "A Tour Through Code: Helping Developers Become Familiar with Unfamiliar Code". In: *Psychology of Programming Interest Group 33rd Annual Workshop*. PPIG 2022, pp. 114–126.
- [245] Suresh Thummalapenta and Tao Xie. "Parseweb: A Programmer Assistant for Reusing Open Source Code on the Web". In: *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering*. ASE '07. Atlanta, Georgia, USA: Association for Computing Machinery, 2007, 204–213. ISBN: 9781595938824. DOI: [10.1145/1321631.1321663](https://doi.org/10.1145/1321631.1321663). URL: <https://doi.org/10.1145/1321631.1321663>.
- [246] Yingchen Tian, Yuxia Zhang, Klaas-Jan Stol, Lin Jiang, and Hui Liu. "What Makes a Good Commit Message?" In: *Proceedings of the 44th International Conference on Software Engineering*. ICSE '22. Pittsburgh, Pennsylvania: Association for Computing Machinery, 2022, 2389–2401. ISBN: 9781450392211. DOI: [10.1145/3510003.3510205](https://doi.org/10.1145/3510003.3510205). URL: <https://doi.org/10.1145/3510003.3510205>.
- [247] Christoph Treude and Margaret-Anne Storey. "Work Item Tagging: Communicating Concerns in Collaborative Software Development". In: *IEEE Transactions on Software Engineering* 38.1 (2012), pp. 19–34. DOI: [10.1109/TSE.2010.91](https://doi.org/10.1109/TSE.2010.91).
- [248] Jason Tsay, Laura Dabbish, and James Herbsleb. "Influence of Social and Technical Factors for Evaluating Contribution in GitHub". In: *Proceedings of the 36th International Conference on Software Engineering*. ICSE 2014. Hyderabad, India: Association for Computing Machinery, 2014, 356–366. ISBN: 9781450327565. DOI: [10.1145/2568225.2568315](https://doi.org/10.1145/2568225.2568315). URL: <https://doi.org/10.1145/2568225.2568315>.
- [249] Gias Uddin and Martin P. Robillard. "How API Documentation Fails". In: *IEEE Software* 32 (4 Aug. 2015), pp. 68–75. DOI: [10.1109/MS.2014.80](https://doi.org/10.1109/MS.2014.80). URL: <https://doi.org/10.1109/MS.2014.80>.
- [250] uuid. *uuid - npm*. NPM. 2024. URL: <https://www.npmjs.com/package/uuid>.
- [251] Max G. Van Kleek, Wolfe Styke, m.c. schraefel, and David Karger. "Finders/keepers: a longitudinal study of people managing information scraps in a micro-note tool". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: Association for

- Computing Machinery, 2011, 2907–2916. ISBN: 9781450302289. DOI: [10.1145/1978942.1979374](https://doi.org/10.1145/1978942.1979374). URL: <https://doi.org/10.1145/1978942.1979374>.
- [252] Fernanda B. Viégas, Martin Wattenberg, and Kushal Dave. “Studying Cooperation and Conflict between Authors with History Flow Visualizations”. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: Association for Computing Machinery, 2004, 575–582. ISBN: 1581137028. DOI: [10.1145/985692.985765](https://doi.org/10.1145/985692.985765). URL: <https://doi.org/10.1145/985692.985765>.
- [253] April Yi Wang, Andrew Head, Ashley Zhang, Steve Oney, and Christopher Brooks. “Colaroid: A Literate Programming Approach for Authoring Explorable Multi-Stage Tutorials”. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI '23. DOI: [10.1145/3544548.3581525](https://doi.org/10.1145/3544548.3581525). URL: <https://doi.org/10.1145/3544548.3581525>.
- [254] April Yi Wang, Anant Mittal, Christopher Brooks, and Steve Oney. “How Data Scientists Use Computational Notebooks for Real-Time Collaboration”. In: *Proc. ACM Hum.-Comput. Interact.* 3.CSCW (2019). DOI: [10.1145/3359141](https://doi.org/10.1145/3359141). URL: <https://doi.org/10.1145/3359141>.
- [255] April Yi Wang, Zihan Wu, Christopher Brooks, and Steve Oney. “Callisto: Capturing the “Why” by Connecting Conversations with Computational Narratives”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, 1–13. ISBN: 9781450367080. DOI: [10.1145/3313831.3376740](https://doi.org/10.1145/3313831.3376740). URL: <https://doi.org/10.1145/3313831.3376740>.
- [256] Dakuo Wang, Judith S. Olson, Jingwen Zhang, Trung Nguyen, and Gary M. Olson. “DocuViz: Visualizing Collaborative Writing”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, 1865–1874. ISBN: 9781450331456. DOI: [10.1145/2702123.2702517](https://doi.org/10.1145/2702123.2702517). URL: <https://doi.org/10.1145/2702123.2702517>.
- [257] Thomas Weber, Maximilian Brandmaier, Albrecht Schmidt, and Sven Mayer. “Significant Productivity Gains through Programming with Large Language Models”. In: *Proc. ACM Hum.-Comput. Interact.* 8.EICS (2024). DOI: [10.1145/3661145](https://doi.org/10.1145/3661145). URL: <https://doi.org/10.1145/3661145>.
- [258] Fengcai Wen, Csaba Nagy, Gabriele Bavota, and Michele Lanza. “A large-scale empirical study on code-comment inconsistencies”. In: *2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC)*. IEEE, 2019, pp. 53–64.
- [259] Richard Wettel and Michele Lanza. “Visualizing Software Systems as Cities”. In: *2007 4th IEEE International Workshop on Visualizing Software for Understanding and Analysis*. 2007, pp. 92–99. DOI: [10.1109/VISSOF.2007.4290706](https://doi.org/10.1109/VISSOF.2007.4290706).

- [260] Jacqueline Whalley, Amber Settle, and Andrew Luxton-Reilly. “Novice Reflections on Debugging”. In: *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. SIGCSE '21. Virtual Event, USA: Association for Computing Machinery, 2021, 73–79. ISBN: 9781450380621. DOI: [10.1145/3408877.3432374](https://doi.org/10.1145/3408877.3432374). URL: <https://doi.org/10.1145/3408877.3432374>.
- [261] Moritz Wittenhagen, Christian Cherek, and Jan Borchers. “Chronicler: Interactive Exploration of Source Code History”. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, 3522–3532. ISBN: 9781450333627. DOI: [10.1145/2858036.2858442](https://doi.org/10.1145/2858036.2858442). URL: <https://doi.org/10.1145/2858036.2858442>.
- [262] Yuhao Wu, Shaowei Wang, Cor-Paul Bezemer, and Katsuro Inoue. “How do developers utilize source code from stack overflow?” In: *Empirical Software Engineering* 24 (2019), pp. 637–673.
- [263] Xin Xia, Lingfeng Bao, David Lo, Pavneet Singh Kochhar, Ahmed E. Hassan, and Zhenchang Xing. “What do developers search for on the web?” In: *Empirical Software Engineering* 22.6 (Dec. 2017), pp. 3149–3185. ISSN: 1573-7616. DOI: [10.1007/s10664-017-9514-4](https://doi.org/10.1007/s10664-017-9514-4). URL: <https://doi.org/10.1007/s10664-017-9514-4>.
- [264] Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael I. Jordan. “Detecting large-scale system problems by mining console logs”. In: *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*. SOSP '09. Big Sky, Montana, USA: Association for Computing Machinery, 2009, 117–132. ISBN: 9781605587523. DOI: [10.1145/1629575.1629587](https://doi.org/10.1145/1629575.1629587). URL: <https://doi.org/10.1145/1629575.1629587>.
- [265] Wei Xu, Ling Huang, and Michael Jordan. “Experience mining Google’s production console logs”. In: *Workshop on Managing Systems via Log Analysis and Machine Learning Techniques (SLAML 10)*. 2010.
- [266] Yinxing Xue, Zhenchang Xing, and Stan Jarzabek. “Understanding Feature Evolution in a Family of Product Variants”. In: *2010 17th Working Conference on Reverse Engineering*. 2010, pp. 109–118. DOI: [10.1109/WCRE.2010.20](https://doi.org/10.1109/WCRE.2010.20).
- [267] Burak Yetiştirgen, Işık Özsoy, Miray Ayerdem, and Eray Tüzün. *Evaluating the Code Quality of AI-Assisted Code Generation Tools: An Empirical Study on GitHub Copilot, Amazon CodeWhisperer, and ChatGPT*. 2023. arXiv: [2304.10778](https://arxiv.org/abs/2304.10778) [cs.SE]. URL: <https://arxiv.org/abs/2304.10778>.
- [268] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supritha Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, and John Zimmerman. “Creating Design Resources to Scaffold the Ideation of AI Concepts”. In: *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. DIS '23. Pittsburgh, PA, USA: Association for Computing

- Machinery, 2023, 2326–2346. ISBN: 9781450398930. DOI: [10.1145/3563657.3596058](https://doi.org/10.1145/3563657.3596058). URL: <https://doi.org/10.1145/3563657.3596058>.
- [269] Young Seok Yoon and Brad A. Myers. “A longitudinal study of programmers’ backtracking”. In: *2014 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 2014, pp. 101–108. DOI: [10.1109/VLHCC.2014.6883030](https://doi.org/10.1109/VLHCC.2014.6883030).
- [270] YoungSeok Yoon and Brad A. Myers. “Semantic zooming of code change history”. In: *2015 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 2015, pp. 95–99. DOI: [10.1109/VLHCC.2015.7357203](https://doi.org/10.1109/VLHCC.2015.7357203).
- [271] YoungSeok Yoon and Brad A. Myers. “Supporting Selective Undo in a Code Editor”. In: *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*. Vol. 1. 2015, pp. 223–233. DOI: [10.1109/ICSE.2015.43](https://doi.org/10.1109/ICSE.2015.43).
- [272] Xiao Yu, Lei Liu, Xing Hu, Jacky Wai Keung, Jin Liu, and Xin Xia. *Fight Fire with Fire: How Much Can We Trust ChatGPT on Source Code-Related Tasks?* 2024. arXiv: [2405.12641](https://arxiv.org/abs/2405.12641) [cs.SE]. URL: <https://arxiv.org/abs/2405.12641>.
- [273] Alexey Zagalsky, Ohad Barzilay, and Amiram Yehudai. “Example Overflow: Using social media for code recommendation”. In: *2012 Third International Workshop on Recommendation Systems for Software Engineering (RSSE)*. 2012, pp. 38–42. DOI: [10.1109/RSSE.2012.6233407](https://doi.org/10.1109/RSSE.2012.6233407).
- [274] Iyad Zayour and Timothy C Lethbridge. “A cognitive and user centric based approach for reverse engineering tool design”. In: *Proceedings of the 2000 conference of the Centre for Advanced Studies on Collaborative research*. 2000, p. 16.
- [275] Andreas Zeller. *Why programs fail: a guide to systematic debugging*. Morgan Kaufmann, 2009.
- [276] Haoxiang Zhang, Shaowei Wang, Tse-Hsun Chen, Ying Zou, and Ahmed E. Hassan. “An Empirical Study of Obsolete Answers on Stack Overflow”. In: *IEEE Transactions on Software Engineering* 47.4 (2021), pp. 850–862. DOI: [10.1109/TSE.2019.2906315](https://doi.org/10.1109/TSE.2019.2906315).
- [277] Tianyi Zhang, Di Yang, Crista Lopes, and Miryung Kim. “Analyzing and Supporting Adaptation of Online Code Examples”. In: *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 2019, pp. 316–327. DOI: [10.1109/ICSE.2019.00046](https://doi.org/10.1109/ICSE.2019.00046).
- [278] J. Zhi, V. Garousi-Yusifoglu, B. Sun, G. Garousi, S. Shahnewaz, and G. Ruhe. “Cost benefits and quality of software development documentation: a systematic mapping”. In: *Journal of Systems and Software* 99 (2015).
- [279] Sacha Zyto, David Karger, Mark Ackerman, and Sanjoy Mahajan. “Successful classroom deployment of a social document annotation system”. In: *CHI 2012*. New York, NY, USA: ACM, 2012, pp. 1883–1892. DOI: [10.1145/2207676.2208326](https://doi.org/10.1145/2207676.2208326). URL: <https://doi.org/10.1145/2207676.2208326>.