Study of the Utility of Text Classification Based Software Architecture Recovery Method RELAX for Maintenance

Daniel Link
University of Southern California
Los Angeles, California, USA
dlink@usc.edu

Kamonphop Srisopha
University of Southern California
Los Angeles, California, USA
srisopha@usc.edu

Barry Boehm
University of Southern California
Los Angeles, California, USA
boehm@usc.edu

ABSTRACT

Background. The software architecture recovery method RELAX produces a concern-based architectural view of a software system graphically and textually from that system's source code. The method has been implemented in software which can recover the architecture of systems whose source code is written in Java.

Aims. Our aim was to find out whether the availability of architectural views produced by RELAX can help maintainers who are new to a project in becoming productive with development tasks sooner, and how they felt about working in such an environment.

Method. We conducted a user study with nine participants. They were subjected to a controlled experiment in which maintenance success and speed with and without access to RELAX recovery results were compared to each other.

Results. We have observed that employing architecture views produced by RELAX helped participants reduce time to get started on maintenance tasks by a factor of 5.38 or more. While most participants were unable to finish their tasks within the allotted time when they did not have recovery results available, all of them finished them successfully when they did. Additionally, participants reported that these views were easy to understand, helped them to learn the system's structure and enabled them to compare different versions of the system.

Conclusions. Through the speedup to the start of maintenance experienced by the participants as well as in their formed opinions, RELAX has shown itself to be a valuable help that could provide the basis of further tools that specifically support the development process with a focus on maintenance.

CCS CONCEPTS

• **Software and its engineering** → **Software architectures**; *Software development process management*; **Programming teams**; **Maintaining software**.

KEYWORDS

Software Architecture Recovery, Text Classification, Software Maintenance

ACM Reference Format:

Daniel Link, Kamonphop Srisopha, and Barry Boehm. 2021. Study of the Utility of Text Classification Based Software Architecture Recovery Method RELAX for Maintenance. In *ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '21)*, October 11–15, 2021, Bari, Italy. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3475716.3484194>

1 INTRODUCTION

While several definitions of what a software architecture is exist [11], e.g., the set of design decisions about a software system [13], they all refer to the structure of a software system and the reasoning process that led to that structure. A software architecture can be represented through many views that follow different paradigms, such as program comprehension and subsystem patterns [14], optimized clustering [9], dependencies, or concerns [4, 6]. (Note that the term *concern* is used with the meaning *something the system needs to have* and not *something to worry about*.)

Having an actionable view of a software system is useful for its stakeholders for a variety of reasons, such as usability [1], security [2], and maintenance, as we are about to show. However, many reasons exist why an architectural view of a system that reflects its current state may not be available. These include that such a view may never have existed or that the system may have evolved away from it over time [13]. Therefore, an interest exists to produce such a view in an efficient and expedient manner. This is where software architecture recovery comes in. It produces an architectural view of a system from its implementation artifacts, such as its source code. RELAX [6] is a software architecture recovery method that follows a concern-oriented paradigm. It produces a software architecture by running text classification on the code entities of a system and building a view of the architecture from the results. The results are represented textually by a list of concern clusters (source code entities grouped by concerns) as shown in Figure 1 as well as graphically via a directory graph [6]. Additional information for each source entity includes SLOC and dependencies.

From the outset, this method was designed with various groups of its end users in mind, such as system architects, developers, integrators and software engineering researchers, whether they would be interested in only one version of a system or plotting its architectural evolution over time. One important subgroup of developers is those who are new to a system, do not have access to current documentation (or any at all) and would like to get started with their work on the system as soon as possible, either because of approaching deadlines or in order to be seen as valuable contributors.

While several of RELAX's attributes, such as its determinism, efficiency, and scalability were able to be shown purely through

This work is licensed under a Creative Commons Attribution International 4.0 License.

ESEM '21, October 11–15, 2021, Bari, Italy

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8665-4/21/10.

<https://doi.org/10.1145/3475716.3484194>

```

contain io org.apache.jackrabbit.core.data.LazyFileInputStream
contain io org.apache.jackrabbit.core.data.LazyFileInputStreamTest
contain io org.apache.jackrabbit.core.fs.mem.MemoryFile
contain io org.apache.jackrabbit.core.fs.mem.MemoryFileSystem
contain io org.apache.jackrabbit.core.fs.mem.MemoryFileSystemEntry
contain io org.apache.jackrabbit.core.fs.mem.MemoryFileSystemTest

contain security org.apache.jackrabbit.core.security.authorization.principalbased.TestAll
contain security org.apache.jackrabbit.core.security.principal.DefaultPrincipalProvider
contain security org.apache.jackrabbit.core.security.principal.GroupPrincipals
contain security org.apache.jackrabbit.core.security.principal.PrincipalIteratorAdapter
contain security org.apache.jackrabbit.core.security.principal.PrincipalManagerImpl
contain security org.apache.jackrabbit.core.security.principal.PrincipalProviderRegistry

contain sql org.apache.jackrabbit.core.fs.db.DB2Filesystem
contain sql org.apache.jackrabbit.core.fs.db.DerbyFilesystem
contain sql org.apache.jackrabbit.core.fs.db.MSSQLFilesystem
contain sql org.apache.jackrabbit.core.journal.DatabaseJournal
contain sql org.apache.jackrabbit.core.journal.JNDIDatabaseJournal
contain sql org.apache.jackrabbit.core.journal.MSSQLDatabaseJournal

```

Figure 1: Partial Example of RELAX Textual Architecture Recovery Concern Clusters

tools that measured its runtime or sameness of results [6], its utility had not been studied yet. Specifically, we were interested in learning whether the task of getting started on modifying a system that was unknown to the study participants could be accelerated by having RELAX recovery results available.

The remainder of this paper is organized as follows: Section 2 describes our approach. Section 3 details the results of our study. Section 4 examines the meaning of the results. Section 5 addresses threats to its validity. Section 6 sums up our conclusions, while Section 7 outlines our future work.

2 STUDY APPROACH

The main objective of our study was to determine the utility of RELAX in a work environment. A task that maintainers working with unfamiliar code encounter repeatedly and regularly is that of searching for relevant code to perform maintenance on [5]. We concluded from this that if we could show that the availability of RELAX recovery results significantly reduces this time, it would hold utility for maintainers. Accordingly, we simulated a situation in which individuals with at least some Java experience are assigned maintenance work on projects they are not familiar with and need to become productive as soon as possible.

2.1 Study Questions

Our study sought to answer two research questions:

SQ1. Does using RELAX architecture recovery results reduce the time to find the location in the code where maintenance needs to be performed?

SQ2. What are the perceptions of new maintainers who work with RELAX architecture recovery results?

2.2 Participants

The software engineering group at our faculty holds several directed research sessions for graduate students of computer science which take place over the full duration of the regular semesters and Summer sessions. We sent out a call for participation in our forthcoming study to enrolled directed research students via email and held a presentation. In it, we outlined the project and made the main criterion outside of interest in the project that they should have Java programming skills commensurate with a bachelor's degree in computer science. Prospective participants were offered extra credit. Nine participants joined our study.

2.3 Duration

Two considerations drove our selection of a duration for the study. On one hand, enough time needed to be available to instruct participants about software architecture, its recovery and RELAX. On the other hand, since the emphasis of the project was on studying performance and not educating the students, we could not expect to work with the participants for an entire semester. Our personal experience over several directed research projects had also shown us that the interest and with it the motivation wears out over time when tasks are repetitive and regarded as work. This would have risked influencing the outcome.

We felt that an overall duration of eight hours distributed over four weeks suitably reconciled the two opposing factors.

2.4 Parameters

During the four weeks our study ran, each participant contributed up to a maximum of two hours per week, depending on whether and how fast they completed the tasks that were part of the study. All nine participants from our convenience sample stayed for the entire duration. The first week was spent on introducing them to or refreshing them on the basics of software architecture and instructing them on software architecture recovery with RELAX, followed by an initial survey on their experience levels going into the project. The second week served to install, try out and validate the participants' hardware and software setups and perform dry runs with warm-up tasks that were not part of the study. The third week was spent exclusively on the first experimental task, while the fourth week was spent on the second task as well as the exit survey.

2.5 Tools and Instrumentation

We gathered our data through surveys and observations. To hold online surveys, the Qualtrics¹ platform was used. Live programming sessions were held over Zoom².

The participants had the option to choose between two maintenance environment setups, depending on their comfort levels:

- (1) A downloadable virtual machine that was pre-configured by us and contained a current desktop installation of Linux Ubuntu³ with current versions of three major Java IDEs (Eclipse⁴, NetBeans⁵ and IntelliJ IDEA⁶ as well as the current source code of Apache Jackrabbit⁷ (version 2.20) and Apache Ant⁸ (version 1.9.15), which were the two systems that were chosen for the experiment.
- (2) Their own system with an IDE of their choice.

The time to download source code and to set up the chosen IDEs was not counted in the second case, while in the first one this had already been done. Time was recorded by us. The participants had 60 minutes to finish each task. A timer was started when the URL of the task description was sent to the participant via Zoom

¹qualtrics.com

²zoom.us

³ubuntu.com

⁴eclipse.org

⁵netbeans.apache.org

⁶jetbrains.com/idea/

⁷jackrabbit.apache.org

⁸ant.apache.org

chat. If the participants finished a task successfully, time-taking was stopped and the actual time taken in minutes was recorded. If they reached the 60 minute mark without having finished the task, *DNF* (standing for *Did Not Finish*) was recorded and counted as 60 minutes. In order to make timekeeping as accurate as possible, all tasks needed to be worked on live during a Zoom session in which the participants shared their screen and sound with the observer. This part of the setup aimed to prevent consultations of outside sources to the extent possible over a remote connection. The observers were authorized to help the participants with their Zoom or Virtual Machine setups. However, such a need did not arise.

The open source Java projects Apache Jackrabbit and Apache Ant were chosen for the following principal reasons:

- No dependencies on older libraries or Java versions that may be hard to obtain
- Thousands of source files make it unlikely that participants can find maintenance locations by guessing or browsing
- Mature systems that are still updated and not likely to exhibit artifacts of antiquity or instability

2.6 Tasks

There were two warm-up tasks that were observed, but not counted, followed by two experimental tasks that were used for the study.

2.6.1 Warm-Up Tasks. The warm-up tasks had the goal to establish the participants' programming skills and validate the infrastructure and setup of the experiment without using RELAX or its architecture recovery results. We estimated these tasks to be feasible in less than ten minutes without knowing anything about the respective systems. The only items that were provided to the participants were the source code of the systems and instructions on how to run the command-line version of the systems in a shell.

The first, easier task was to change the textual output of Apache Jackrabbit on command-line startup from a shell. The second, more difficult task was to do the same with Apache Ant. While in the first case, this was achievable by making changes within the Java source code, it required modification of an external XML file in the second.

2.6.2 Experimental Tasks. Since there were 9 participants, the split between the control group (participants not given the recovery results) and the experimental group (participants given the recovery results) could not be even, but was made 4 to 5 for the first task and 5 to 4 for the second (compare Table 2 below). Whichever participant was in the control group for the first task was in the experimental group for the second and vice-versa.

Both tasks were to be performed on Apache Jackrabbit 2.20 and had a common setting in which the participants were to assume that they know nothing about the architecture of Apache Jackrabbit and are working alone in a closed and sealed bunker with no internet access. As surveyed at the outset of the study (Section 3.1), none of the participants had any experience with the architecture or use of Jackrabbit and none had contributed any code to it. The participants had 60 minutes to finish each task. Timing was capped at 60 minutes as described in Section 2.5.

Task 1:

- *You are tasked with making fixes to a filesystem that interfaces with a database. In order to start your work, you think of adding a few print statements - but where?*

Task 2:

- *You are tasked to clean up system after an unreachable former contributor has made modifications to one source file so it does not handle security anymore, but SQL instead. The source file was not renamed. You want to make the system maintainable again and change the name of that file – but where do you look?*

While the first experimental task was oriented toward an individual system version, the second one focused on the evolution of a system.

2.7 Surveys

The participants were given an initial survey in the first week, directly after the instruction phase (so that there would be no confusion over terms used) and an exit survey at the end of the study period.

The surveys consisted of a variety of question types, including text fields (e.g., for email addresses), multiple-choice (e.g., whether textual results, the diagram or both together gave participants the most complete information about the software system), yes-no (e.g., whether participants had contributed to Apache Jackrabbit), with rating questions shown as rating scales or Likert scales.

To reduce fatigue and hold the attention of participants, the survey varied question types where possible. For example, some questions that could have been put on a Likert scale were put on ratings scales. For the same reason, some ratings questions used sliders and some asked participants to enter their ratings as numbers in text fields.

Questions are shown along with their results in Section 3.3.

2.7.1 Initial Survey (Participants' Experience). This asked participants about the levels of experience they had before the experiment started, regarding

- Programming
- Text classification
- Software architecture and its recovery
- Apache Jackrabbit

Some questions, such as one question that asked participants whether they had contributed to Apache Jackrabbit or another that asked about preexisting experience with software architecture recovery opened further branches of questions on the relevant details. However, this did not apply to any of our participants.

2.7.2 Exit Survey (Evaluations of RELAX). This asked the participants utility-related questions regarding RELAX recovery results and working with them on individual and multiple system versions.

3 RESULTS

The full results of our study including the surveys are shared publicly [7].

Table 1: Survey Results (Scales from 1-5)

Question	Average	σ
Familiarity with...		
Text Classification	2.56	1.13
Software Architecture Recovery	1.00	0.00
Apache Jackrabbit Source Code	1.11	0.33
Apache Jackrabbit Architecture	1.00	0.00
Apache Jackrabbit Documentation	1.11	0.33
Apache Jackrabbit Use	1.00	0.00
RELAX view of Apache Jackrabbit:		
Easy to understand	4.22	0.42
I can learn the structure of the system from it	4.78	0.42
It can support new contributors	4.67	0.47
Real-time RELAX version valuable	4.33	0.47
RELAX views of several Apache Jackrabbit versions:		
Easy to compare	4.89	0.31
Explain differences between versions	4.89	0.31
Describe system evolution	4.67	0.47

On the question whether the textual results, the diagram or both gave the most complete information on the architecture of a system, 3 participants chose the textual ones and 6 the combination of both.

3.1 Initial Survey on Preexisting Experience

The top of Table 1 shows results of our survey on the experience the participants had before they began the study. Out of the fields surveyed, the only field they tended to have some experience in is that of text classification. They were new to Apache Jackrabbit when it came to its architecture and use and had next to no experience with its code and documentation. They had no experience with Software architecture recovery and none had contributed to Apache Jackrabbit.

Out of our nine study participants, eight had 1-5 years of programming experience, only one had 5-10 years of experience. Typically, they had some text classification experience, with the average being 2.56 on a scale of 1-5, with 5 being the highest.

Interpretation: This is in line with expectations for graduate students in computer science and was not likely to distort any results. It is particularly agreeable for our study that they did not have any significant experience with Apache Jackrabbit.

3.2 Timings of Tasks

Table 2 shows how much time in minutes each participant took for each individual task. As mentioned, if they did not finish a task within 60 minutes, the time is shown as *DNF* for “Did Not Finish” and counted as 60 minutes. The columns refer to the warm-up tasks described in Section 2.6.1 and experimental tasks 1 and 2 outlined in Section 2.6.2.

The time RELAX took to recover an architectural view of Jackrabbit 2.20 on a 3.6 GHz Intel Core i9 using macOS was 235 seconds. For both experimental tasks, the results are shown without and with counting the time the architectural view recovery took, rounded up to 4 minutes and added for each case in which a participant had the recovery results. Timings including the recovery runtime are shown in parentheses. The addition of this time is discussed in Section 4 below.

Before we look at the average speedups, two observations are notable: First, that of the nine participants, three were unable to

Table 2: Task Timing Results (in Minutes)

Participant	Warm-Up		Experimental	
	1	2	Task 1	Task 2
1	DNF	DNF	DNF	3.00 (7.00)
2	DNF	27.00	DNF	2.00 (6.00)
3	11.50	8.00	6.00	2.00 (6.00)
4	20.75	4.00	DNF	1.00 (5.00)
5	DNF	DNF	DNF	3.00 (7.00)
6	12.50	DNF	7.08 (11.08)	DNF
7	14.50	DNF	2.50 (6.50)	4.50
8	16.50	DNF	5.00 (9.00)	DNF
9	12.00	DNF	6.00 (10.00)	DNF
Averages				
	29.75	44.33		
Without recovery results			49.20	46.13
With recovery results			5.15 (9.15)	2.20 (6.20)
Speedup factor			9.56 (5.38)	20.90 (7.44)

Numbers in parentheses include RELAX runtime of 4 minutes.
DNF (Did Not Finish) counts as 60 minutes.

complete the first warm-up task and six the second. Second, while for each experimental task, only one participant was able to complete it without having recovery results available, all participants were able to complete their tasks when they had them. Having recovery results has sped up the maintenance tasks by a factor of at least 5.38 for the first task or 7.44 for the second. These factors rise to 9.56 and 20.97 if we exclude the recovery time. It needs to be considered that in those cases where participants did not finish, the speedup factors are influenced by the timing cap of 60 minutes. Higher caps could have resulted in higher speedup factors.

3.3 Exit Survey

Table 1 shows the results of our exit survey in its lower two sections regarding working with either architectural views of one system version or views produced for several versions, respectively. The participants have formed overall favorable opinions from the view of one Jackrabbit version as well as the views of several versions. The highest rankings are held by the ease of understanding the view of a single version as well as comparing and explaining the differences between versions. One result that will be discussed below is that they found a hypothetical version of RELAX that would recover architectural views in the background valuable.

4 DISCUSSION

The basis of our discussion can be summed up as follows:

- The start of maintenance is sped up by at least a factor of 5.38
- Across the board, participants strongly considered architectural views produced by RELAX helpful in maintenance tasks
- This applies to both individual architectures and evolution

To which extent the recovery runtime for the current system version needs to be counted as part of the time the participant is working on a task depends on the specific setup regarding the number of maintainers and other factors. As discussed below in

Section 4.2.2, we do not envision this to be a major timing factor. However, we present our results with and without counting the runtime and conservatively report the longer time as our overall result.

4.1 Answers to our Study Questions

4.1.1 SQ1 (Does using RELAX architecture recovery results reduce the time to find the location in the code where maintenance needs to be performed?) The start of maintenance is sped up by factors of 5.38 and 7.44 for both tasks, respectively. Based on these results, SQ1 can therefore be answered affirmatively.

4.1.2 SQ2 (What are the perceptions of new maintainers who work with RELAX architecture recovery results?) As reported in Section 3.3, the participants strongly considered the architectural views produced by RELAX to be helpful in their maintenance tasks. Based on this result, SQ2 can be answered affirmatively.

4.2 Implications

4.2.1 Enabling Maintenance. Two standout observations come from Section 3.2. The first is that without having an architectural view of the system, for each of the two respective experimental tasks, only one participant was able to finish it. The second is that with an architectural view of the system, all 9 participants, including the 7 that had not successfully finished one or both warm-up tasks, were able to successfully finish their experimental tasks. This means that even if one does not consider any speedups, a benefit of RELAX lies in making timely maintenance of a system possible.

4.2.2 RELAX IDE Plugin. Based on the positive results of our study, in which we have demonstrated the utility of RELAX for maintenance tasks, we considered new tools for software maintainers, whether they are new to a system or know it well. A possible tool would be a plugin for major IDEs. One feature that sets RELAX apart from other recovery methods is that it assembles its architectural view in an additive manner from the text classification of individual source code entities of a system. This has the important consequences that individual source code files can be classified while they are being edited and that in order to see changes in the system, there is no need to recover the architecture of the whole system again because classifying the changed entities is sufficient.

We have considered how an IDE plugin that takes all this into account would be designed with regards to its internal and visual workflow. Some of the details of RELAX that the IDE would reference are beyond the scope of this paper and are described in [6]. Four considerations would be that the plugin would (1) recover the system by request or automatically in the background after each successful compilation, (2) show the maintainer the current status of the architecture visually at a glance, (3) allow the kind of searches that helped the participants in our study, and (4) warn if the current architecture has issues, such as architectural smells [3].

Figure 2 shows the plugin status view that a maintainer could leave open in an IDE. The upper left shows information related to the current code entity being edited, such as its name, file size, SLOC, color-coded outgoing and incoming dependencies as described in [6]. The right half shows a graphical or textual view of the most recently recovered architectural view, allowing zooming in

Figure 2: Plugin Status

Figure 3: Plugin Configuration Dialog

Figure 4: Plugin Layer Dependency Warning

and panning. Finally, the lower left shows how many architectural issues have been found in the last recovery, along with a button that will open up a dialog showing issue details. Figure 3 shows the configuration dialog of the plugin. This would allow the maintainer to select or deselect auto-recovery after each compilation, select a trained classifier and to open a dialog that allows them to train a new classifier. The latter option could become helpful if the concerns of interest for the maintainer have changed. Figure 4 shows an example of an architectural warning that is enabled by the concern-orientation of RELAX. Maintainers could assign layers to some or all concerns in their system and the IDE plugin could analyze the architecture for layer breaches in which a lower-level entity calls a higher one. The maintainer could then check the details and take action by opening the file in the editor or view the next warning.

5 THREATS TO VALIDITY

5.1 Task Selection

The experimental tasks focused on finding code locations and not on other parts of the maintenance process, such as following dependencies, reading documentation, or more. Such activities would have been outside of the scope of our study, however.

5.2 Subject System

The subject system was selected after dismissing several other ones for not being a good fit due to lack of compilability, trivial size, not being open, not having any documentation (needed to come up with tasks) and other reasons.

5.3 Number of Systems

While a larger amount of systems might have provided more insights, this was not possible due to the limitations in the availability of participants. In order to select a system that is as representative of a project that could conceivably take on new contributors, Apache Jackrabbit was chosen as project from the Apache Software Foundation⁹, which hosts a large number of open source projects.

5.4 Open Source vs Closed Source

The programming style or quality of closed source systems may be significantly different from that of open source systems in ways that would lead to different results. However, studies [10, 12] have found no significant differences between the code quality of open- and closed-source systems.

5.5 Java Programming Language

A Java-based subject system was selected because the current version of RELAX only allows to recover architectural views from Java source code. The experimental tasks could be easier or more difficult if posed in different programming languages. Java, however is one of the most commonly used programming languages among developers worldwide [8] and therefore a typical environment.

5.6 Participants

The participants were self-selected from computer science graduate students with at least some Java programming experience. While this group does not match up fully with all possible stakeholders in a system that can be served by RELAX recoveries, some of which can be non-technical, it does match up with the group of maintainers who would want to get started on their tasks.

6 CONCLUSION

We conducted a user study with nine participants to assess the utility of the architectural recovery method RELAX. The participants were subjected to a controlled experiment in which speed and other metrics with and without access to RELAX recovery results were compared to each other on Apache Jackrabbit 2.20.

Through speeding up the start of maintenance by factors of at least 5.38 and other benefits experienced by the participants in our study, RELAX has shown itself to be a valuable help in the software

maintenance process that can make the difference that allows new maintainers to finish a maintenance task on time.

7 FUTURE WORK

7.1 Larger-Scale Study

The encouraging emerging results from our study provide motivation for us to hold a larger-scale study with a broader scope.

Under consideration are several ways to achieve this.

- *Wider*: More systems and systems based on other programming languages than Java could be studied.
- *Deeper*: More tasks could be given to the participants, either several on one system or on a number of systems.
- *Larger*: A larger number of participants could be recruited and drawn from different pools, such as from the software industry and software engineers with the experience using other recovery methods.

7.2 Software Support

The envisioned IDE plugin that shows architectural changes as the system evolves (see Section 4.2.2) could be developed together with former participants.

REFERENCES

- [1] Len Bass and Bonnie E John. 2003. Linking usability to software architecture patterns through general scenarios. *Journal of Systems and Software* 66, 3 (2003), 187–197.
- [2] Christophe Bidan and Valérie Issarny. 1997. Security benefits from software architecture. In *International Conference on Coordination Languages and Models*. Springer, Berlin, 64–80.
- [3] J Garcia, D Popescu, G Edwards, and N Medvidovic. 2009. Identifying Architectural Bad Smells. *2009 13th European Conference on Software Maintenance and Reengineering* 1, 1 (2009), 255–258. <https://doi.org/10.1109/CSMR.2009.59>
- [4] Joshua Garcia, Daniel Popescu, Chris Mattmann, Nenad Medvidovic, and Yuanfang Cai. 2011. Enhancing architectural recovery using concerns. In *2011 26th IEEE/ACM International Conference on Automated Software Engineering (ASE 2011)*. IEEE, Melbourne, 552–555.
- [5] Andrew J. Ko, Brad A. Myers, Michael J. Coblenz, and Htet Htet Aung. 2006. An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks. *IEEE Transactions on Software Engineering* 32, 12 (2006), 971–987. <https://doi.org/10.1109/TSE.2006.116>
- [6] Daniel Link, Pooyan Behnamghader, Ramin Moazeni, and Barry Boehm. 2019. Recover and RELAX: concern-oriented software architecture recovery for systems development and maintenance. In *Proceedings of the International Conference on Software and System Processes*. IEEE, Montréal, 64–73.
- [7] Daniel Link, Kamonphop Srisopha, and Barry Boehm. 2021. RELAX ESEM 2021 DR User Study Survey Results Shared Document. <https://doi.org/10.5281/ZENODO.5323958>
- [8] Shanhong Liu. 2021. Most used languages among software developers globally 2020. <https://www.statista.com/statistics/793628/worldwide-developer-survey-most-used-languages/>
- [9] S. Mancoridis, B.S. Mitchell, Y. Chen, and E.R. Gansner. 1999. Bunch: a clustering tool for the recovery and maintenance of software system structures. *Proceedings IEEE International Conference on Software Maintenance - 1999 (ICSM'99)*. 'Software Maintenance for Business Change' (Cat. No.99CB36360) 1, 1 (1999), 50–59. <https://doi.org/10.1109/ICSM.1999.792498>
- [10] Srinivasan Raghunathan, Ashutosh Prasad, Birendra K Mishra, and Hsihui Chang. 2005. Open source versus closed source: software quality in monopoly and competitive markets. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 35, 6 (2005), 903–918.
- [11] Fritz Solms. 2012. What is software architecture?. In *Proceedings of the south african institute for computer scientists and information technologists conference*. ACM, Pretoria, 363–373.
- [12] Diomidis Spinellis. 2008. A tale of four kernels. In *Proceedings of the 30th international conference on Software engineering*. ACM, Leipzig, 381–390.
- [13] Richard N. Taylor, Nenad Medvidovic, and Eric M. Dashofy. 2010. *Software Architecture - Foundations, Theory, and Practice*. Wiley, Hoboken.
- [14] Vassilios Tzerpos and RC Holt. 2000. ACDC: An Algorithm for Comprehension-Driven Clustering. In *Proceedings Seventh Working Conference on Reverse Engineering*. IEEE Computer Society, Brisbane, 258–258.

⁹apache.org