

Evidence Profiles for Validity Threats in Program Comprehension Experiments

Marvin Muñoz Barón, Marvin Wyrich, Daniel Graziotin and Stefan Wagner
 Institute of Software Engineering, University of Stuttgart
 Stuttgart, Germany
 {firstname.lastname}@iste.uni-stuttgart.de

Abstract—Searching for clues, gathering evidence, and reviewing case files are all techniques used by criminal investigators to draw sound conclusions and avoid wrongful convictions. Medicine, too, has a long tradition of evidence-based practice, in which administering a treatment without evidence of its efficacy is considered malpractice. Similarly, in software engineering (SE) research, we can develop sound methodologies and mitigate threats to validity by basing study design decisions on evidence.

Echoing a recent call for the empirical evaluation of design decisions in program comprehension experiments, we conducted a 2-phases study consisting of systematic literature searches, snowballing, and thematic synthesis. We found out (1) which validity threat categories are most often discussed in primary studies of code comprehension, and we collected evidence to build (2) the evidence profiles for the three most commonly reported threats to validity.

We discovered that few mentions of validity threats in primary studies (31 of 409) included a reference to supporting evidence. For the three most commonly mentioned threats, namely the influence of programming experience, program length, and the selected comprehension measures, almost all cited studies (17 of 18) did not meet our criteria for evidence. We show that for many threats to validity that are currently assumed to be influential across all studies, their actual impact may depend on the design and context of each specific study.

Researchers should discuss threats to validity within the context of their particular study and support their discussions with evidence. The present paper can be one resource for evidence, and we call for more meta-studies of this type to be conducted, which will then inform design decisions in primary studies. Further, although we have applied our methodology in the context of program comprehension, our approach can also be used in other SE research areas to enable evidence-based experiment design decisions and meaningful discussions of threats to validity.

Index Terms—program comprehension, threats to validity, empirical software engineering

I. INTRODUCTION

Between 2008 and 2013, drivers on German highways had reasons to be concerned. Criminal investigators were on the trail of *The Highway Shooter*, as the media called them. Over the years, the Highway Shooter had fired a gun at passenger cars more than 700 times, with a predilection for trucks [1]. As it became clear that the shooter was mobile and that cases had spread across the country, the German Federal Criminal Police Office took over.

Several attempts were made to apprehend the criminal, most of which were covered intensively by the media. These attempts culminated in a public manhunt with a reward of 100 000 €. In the end, however, what led to success was a

combination of two measures: firstly, carriers were sensitized to immediately report bullet holes in their vehicles. Secondly, several license plate recording devices were set up along the highways. April 2013 was a hectic month: six shootings were reported within five days. The culprit's route started showing patterns. These patterns, crossed with the captured license plate numbers, ultimately resulted in a single revealing license plate number. The final evidence led to the arrest of The Highway Shooter in June 2013.

In the same year that the criminal was caught, Wohlin published a paper on evidence profiles for software engineering research and practice [2]. The overlap in time is likely coincidental, but some parallels between the case and Wohlin's proposed methodology are not.

Wohlin [2] motivates the work with the increasing importance of evidence-based software engineering (EBSE). For example, critical decisions such as the introduction of a new tool that could affect software quality or developer productivity should be based on scientific evidence. Wohlin proposes a model by which evidence from different studies could be evaluated in a manner similar to criminal law investigations. Practical conclusions could be drawn by putting the individual pieces together in an evidence profile. Just as with license plates and reports from drivers in the case of The Highway Shooter, we can more effectively provide evidence in software engineering investigations.

Just as investigators were left in the dark for several years in the case of The Highway Shooter, researchers have been uncertain about the consequences of design decisions in code comprehension experiments. In the corresponding papers, researchers regularly discuss and speculate about the threats to validity of their experiments. Two meta-studies have categorized these threats, coming up with more than 50 different threat categories [3], [4]. Moreover, papers that do not discuss threats to validity in code comprehension experiments are rather the exception today [4], [5]. At the same time, hardly anyone is sure about the actual extent of the discussed threats, and almost no paper cites evidence on the assumed threats [4]. This makes it difficult to evaluate study designs in an evidence-based manner. Researchers have to decide which of the more than 50 potential threats do, in fact, threaten the validity of a study design, execution, and interpretation, and to which extent they should be disclosed and elaborated on.

In this paper, we apply Wohlin's methodology of evidence

profiles. We echo a recent call by Wyrich et al. [4] to provide more evidence in the design of empirical program comprehension studies. To that end, we examined the threats to validity in program comprehension experiments to collect evidence of their existence, to understand the context and nature in which they occur, and to ultimately assist researchers in designing controlled experiments with high validity. Specifically, we extracted the threats to validity reported in 95 code comprehension experiments through thematic synthesis. Then, focusing on the three most frequently mentioned threat categories, we collected evidence that contradicted or supported the influence of the threat, using systematic literature searches and snowballing. Finally, we individually scored the evidence that passed our filtering criteria to create an evidence profile, serving as an overview of the evidence for each of the three threat categories.

Our work shows that the extent of commonly assumed threats to validity can be very dependent on the particular design and context factors of a primary study. Nevertheless, certain threats have so far been generalized to all code comprehension experiments without restriction. With three evidence profiles, we show how evidence-based experiment design decisions can be made and discussed. We will illustrate, throughout the paper, how such an approach can improve the validity of experiments even beyond the context of program comprehension for many other experiments in software engineering.

The paper is structured as follows. In Section II, we define the context of our investigation regarding the current handling of threats to validity in software engineering research and explain their specific meaning in the program comprehension field. In Section III, we introduce concrete research questions based on the introductory motivation and describe in detail our methodological approach to answer these questions. In Section IV, we present the results of this work and discuss them in Section V, before concluding the paper in Section VI.

II. BACKGROUND AND RELATED WORK

Medicine has a long tradition of evidence-based practice, in which administering a treatment without evidence of its efficacy is considered malpractice. Sackett et al. describe evidence-based medicine as “the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients” [6, p. 71].

In software engineering, attempts have been made to establish similar approaches to help practitioners make informed decisions in their daily work [7]. Proponents of evidence-based methods emphasize the importance of using evidence when adopting new technologies and when understanding and identifying problems in existing development processes. Making uninformed decisions may lead practitioners to favor ineffective solutions over better alternatives, resulting in financial losses or even human costs. We consider evidence to be the available body of empirical knowledge “indicating whether a belief or proposition is true or valid” [8, p. 607]. Throughout our work, this definition and our refinement of the term with

different levels of evidence strength in Section III-C guided us in identifying how empirical knowledge is documented and how we can use it to assist researchers.

Since the first calls for evidence-based software engineering in the early 2000s [7], [9], investigative methods from EBSE have found widespread usage. There has been a significant increase in the number of secondary software engineering studies [10], [11] and educators are actively incorporating EBSE in their university curricula [12]–[14]. However, recent discussions highlight difficulties in the application of evidence gained from primary and secondary research [15], [16]. In response, more tools and structured approaches are being introduced in an attempt to address the slow transfer of knowledge from research to practice (e.g., evidence profiles [2], evidence briefings [17], and rapid reviews [18]).

While evidence-based approaches may assist practitioners in making decisions, Kitchenham et al. [9] also describe how these approaches place additional requirements on researchers when developing experimental protocols. Ideally, researchers can maximize the range of application of a study’s results while minimizing potential threats to validity.

Validity refers to the degree to which we can trust the results of an empirical study [19]. In our study, a threat to validity refers to deliberate design decisions and uncontrolled extraneous factors that may impair the validity of experimental results. When readers are explicitly informed about validity threats, they can better assess the context in which experiment results may be applied. Consequently, they are empowered to understand what difficulties may arise when they attempt to replicate the study design or plan a similar study of their own.

There have been several meta-studies to summarize and categorize common threats to validity in software engineering research. Petersen and Gencel [20] compared validity threats with different worldviews. They assumed that researchers have a subconscious tendency to choose methods based on their worldview. Likewise, Devanbu et al. [21] found in a case study that programmers tend to hold strong beliefs based on personal experience rather than empirical evidence. The main implication of these studies is that researchers and practitioners make biased decisions based on their intuition, which is at odds with the main goals of evidence-based methods. Rather, they should address the discrepancy between the evidence and their perceptions and reconsider their decisions accordingly. Peterson and Gencel [20] suggest that the literature needs to be further analyzed regarding the threats and mitigation techniques mentioned therein and that inquiries need to be made into what worldviews dominate in the various sub-disciplines of software engineering.

Biffel et al. [22] followed this suggestion and created a knowledge base of threats to validity in software engineering experiments to assist researchers in planning their studies. They found that only a small fraction of validity threats are reported in most studies and that, instead, the vast majority of threats are too specific to be generalized outside the particular research area they occur in. These findings highlight the complexity of managing threats to validity, as even switch-

ing between different sub-disciplines of software engineering introduces a whole new set of potential threats that must be considered. They conclude that there is a need for an overview of threats to validity as they are reported in specific areas of software engineering research.

Managing threats to validity is particularly difficult in code comprehension experiments, where researchers seek to uncover the underlying processes of how developers understand code and evaluate ways to support that comprehension process scientifically [23]. This complexity is reflected in the wide variety of different methodologies researchers use to address the potential validity threats in their experiments [4].

For example, Siegmund and Schumann [3] surveyed the literature to obtain information on confounding parameters in studies of program comprehension. The main insights they gained were that each paper reported only a small subset of all confounding factors, and that researchers used different methods to mitigate the same factors. They recommend that other researchers include the identification of confounding parameters in their experimental design and explicitly report the relevant parameters and how their influence is controlled.

Building on the findings of Siegmund and Schumann, Wyrich et al. [4] analyzed the characteristics of 95 source code comprehension experiments and identified several shortcomings and inconsistencies in their experimental designs and in how they reported the threats to validity. They noted that, currently, researchers tend to rely on intuition when designing their experiments because of how difficult it is to reliably measure a person's understanding of code. Researchers have to deal with potential threats to validity from over 50 categories and sometimes question whether their measure of understanding is adequate at all [4].

These meta-studies illustrate the complexity of designing valid code comprehension experiments, considering potential confounding factors and other threats to validity [3], [4]. Rather than choosing methods based on intuition, we argue that researchers should make informed decisions based on empirical evidence. Previous research has outlined a need for common knowledge bases of validity threats and collections of evidence backing up both the existence of threats and the effectiveness of their mitigation techniques. The first step in solving these issues is to identify where evidence is lacking.

Surveying the evidence landscape on any particular issue can be quite daunting. For this purpose, Wohlin [2] proposes the creation of evidence profiles to gain an overview of both the amount and the direction of evidence. By scoring each individual piece of evidence in the context of an explicit research question, researchers are assisted in identifying sufficient or missing evidence. This approach shares similarities with other meta-studies, such as meta-analyses, in that it aims to summarize and combine the results of multiple studies. Evidence profiles differ in their method of synthesis, as meta-analyses take a quantitative approach in composing effect sizes, whereas evidence profiles use multiple reviewers to qualitatively evaluate the studies by scoring them.

III. METHOD

We first investigate which threats to validity are most frequently discussed in program comprehension literature (III-A). The frequency of discussion, however, should only be considered as an indicator of what researchers are most often concerned about, not of the threats' evidence. Therefore, in a second step, we will examine the scientific evidence for the most frequently discussed threats (III-B and III-C).

Fig. 1 provides a schematic overview of the research methodology. The research questions that guided us in our endeavor are as follows:

- RQ1** Which validity threat categories are most often discussed in primary studies of code comprehension?
- RQ2** What are the evidence profiles for the three most commonly reported threats to validity in code comprehension experiments?

A. Scoping the Relevant Threats

To answer **RQ1**, we surveyed existing studies of code comprehension and identified which threats were most frequently mentioned by researchers. In particular, we extracted the threats to validity reported in 95 source code comprehension experiments found in the openly available dataset by Wyrich et al. [24]. Their systematic search protocol included empirical studies of bottom-up code comprehension with human participants, published in a peer-reviewed journal, conference, or workshop before 2020. While Wyrich et al. [4] reviewed the threats to validity to some extent, we opted to repeat the coding and categorization activities to enable a more fine-grained classification of threats with the explicit goal of gaining insights into the reporting of validity threats as a whole. Moreover, we performed this bottom-up coding activity to stay closer to the data, building a foundation for the subsequent evidence searches.

We examined the list of primary papers, categorizing and summarizing each threat to validity reported in the full text. Specifically, we adopted a thematic synthesis [25] approach to identify the individual threats to validity and, if given, the mitigation techniques mentioned by the study. For this, we extracted relevant text areas using inductive coding [26] to find and describe the passage. In inductive coding, codes are formulated during the review process as the corresponding concepts become evident. Throughout the coding process, these codes are refined and reapplied, improving their quality through iteration. Upon the completion of code assignment, we categorized threat codes into high-level categories and themes. Further, in some cases, papers already contained evidence of a threat in the form of references to other studies and consequently, those were also documented. Finally, we counted the number of papers that cite a threat for each threat category, resulting in an overview of the most frequently discussed threat categories.

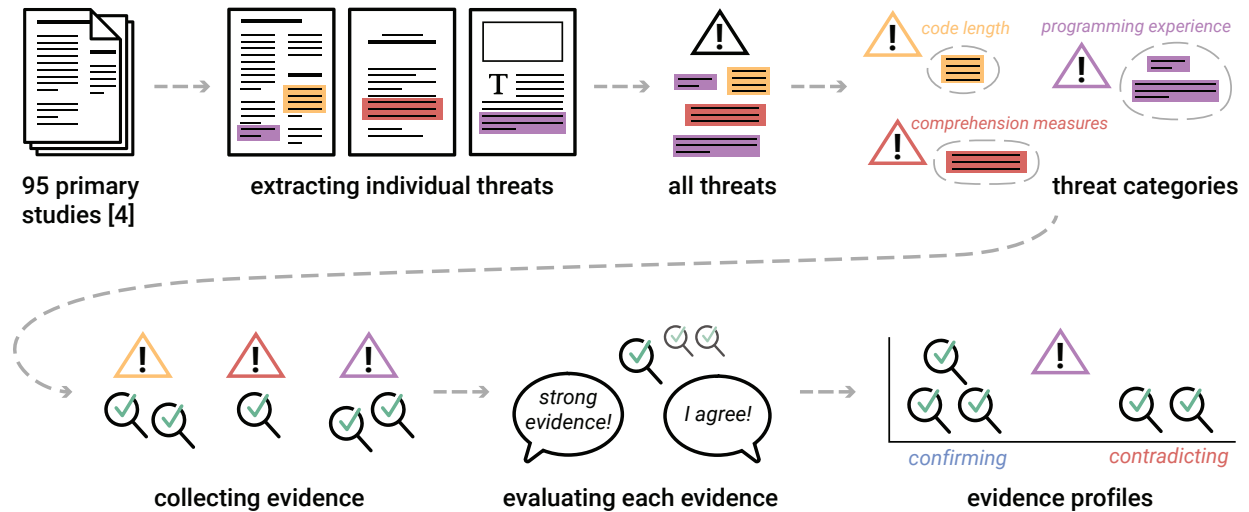


Fig. 1. Schematic representation of the research methodology.

In summary, reviewing threats reported in primary studies consisted of the following individual steps:

- 1) Extraction of relevant text passages on threats to validity, mitigation techniques, and evidence.
- 2) Inductive coding, in which each threat is assigned a descriptive code.
- 3) Categorization of threat codes and composition into higher-order themes.
- 4) Counting the number of papers mentioning each threat category.

B. Evidence Collection

At this point, we had a list of threat categories, how frequently they were reported, and in some cases, a pool of starting evidence. Due to the large number of different threat categories, we could not collect evidence for each individual one. Instead, we focused on the three categories that were mentioned most frequently. These threat categories were the level of programming experience of participants, the length of the programs used in the experiment, and the measures that were used as a proxy for the concept of comprehension.

To answer **RQ2**, we conducted systematic searches for each of the three threat categories, collecting evidence on their influence. The steps described in this section and in Section III-C were repeated for each individual category.

Search Protocol: Our search protocol utilized four different sources to search for potentially relevant papers. We describe each of these sources and list the filtering criteria applied to the literature found within them. Furthermore, we describe how we used snowballing [27] as a technique to further extend the search. Backward snowballing in this context means including the reference list of a paper, while forward snowballing means including papers that cite the paper in question.

a) Primary Papers: The 95 papers with primary research on program comprehension may already examine the threat in

question as part of their study. Consequently, we evaluated them as potential evidence. Forward snowballing was less likely to yield relevant results, as investigating the threat in question was not the main focus of these papers. Unless the threat was the main subject of a primary study, other studies examining the threat are unlikely to cite the primary study in the context of discussing the threat. Backward snowballing for this set of papers was not required because the evidence cited in the primary papers was already a separate source, as we describe in the next paragraph.

b) Evidence Cited in the Primary Papers: As part of the analysis described in Section III-A, we identified and documented all references in the primary papers that were cited to support assertions made about a threat to validity. We analyzed and filtered this list of references in the same manner as the primary papers. The evidence from this source did not necessarily focus on the threat as their primary object of research. Forward snowballing was therefore not used, as it was less likely to yield relevant results. However, we performed backward snowballing, as the papers may refer to similar evidence when comparing their results with other works.

c) Evidence Found Through the Title Search: To further enrich the dataset with research from sources independent of the primary papers, we also conducted additional systematic searches. While the literature search by Wyrich et al. [4] already captured code comprehension experiments up to 2020, we were able to extend the range of our search to additionally include literature published between 2020 and 2022. Furthermore, we focused this search on studies that mention the particular threat to validity in their title using the following search strings in Google Scholar:

TABLE I
 TYPES OF EVIDENCE ACCORDING TO THE EVIDENCE PROFILE.
 HIGHER VALUES MEAN STRONGER EVIDENCE. POSITIVE VALUES SUPPORT THE THREAT, NEGATIVE VALUES CONTRADICT IT.

Score	Level of Strength	Description
+5/-5	Strong evidence	Studies that focused on the threat in question as the main subject of their investigation or conducted an in-depth analysis of the threat as part of their overall approach and show significant results. Systematic reviews that examined the threat in question and provided a conclusion.
+4/-4	Evidence	Studies that did not have the threat as their main focus but still included it in their analysis. These studies may have more uncontrolled confounding factors, which is why they should be considered separate from strong evidence.
+3/-3	Circumstantial evidence	Similar to evidence, studies that are considered circumstantial evidence did not have the threat as the main focus of their study. Furthermore, they showed additional methodological shortcomings that reduced the reliability of the study and decreased its strength as evidence in our evaluation.
+2/-2	Third-party claim	Studies that made claims about the threat in question but only provided a slim level of empirical backing for said claim. They may have deferred to other sources of information or given general impressions about the influence of the threat in their study, but provided no dedicated statistical analysis of their own to support their claims.
+1/-1	First- or second-party claim	Studies that made a claim about the threat in question but did not provide any empirical backing. This could have been, for example, references to “common knowledge” or speculation.

- **Programming Experience** - allintitle: (experience OR novice OR expert) (code OR software OR program) (understandability OR comprehension OR comprehensibility OR readability OR analyzability OR “cognitive load”)
- **Program Length** - allintitle: (size OR length OR short OR long OR LOC OR “lines of code”) (code OR software OR program) (understandability OR comprehension OR comprehensibility OR readability OR analyzability OR “cognitive load”)
- **Comprehension Measures** - allintitle: (measure OR measures OR measurement) (code OR software OR program) (understandability OR comprehension OR comprehensibility OR readability OR analyzability OR “cognitive load”)

Each search string was composed by combining terms describing a threat (e.g. experience, novice, expert) with terms related to program comprehension experiments (e.g. code, understandability, comprehension). We used both backward and forward snowballing, as the studies found in the title search were likely to have the threat in question as the main subject of their research.

d) *Evidence Found Through Snowballing*: By exploiting clusters of related research, snowballing can identify relevant literature with high precision, but tends to miss some papers [28]. Therefore, we use a hybrid approach to complement disadvantages of snowballing with a database search and vice versa. Backward snowballing was applied to the evidence cited in the primary papers, while backward and forward snowballing were used on the evidence from the title search. In all three cases, we did one iteration of snowballing.

All papers found in the sources a) to d) were filtered according to the following four inclusion criteria and were only included when all of them were fulfilled:

- The paper reports a primary study measuring program comprehension.
- The paper reports an analysis of the threat in question.
- The paper is published in a peer-reviewed journal or conference proceeding.
- The paper’s full text is available in English.

C. Evidence Profiles

After collecting all available evidence of a particular threat, we evaluated the evidence itself. To this end, we employed the evidence profile proposed by Wohlin [2]. This profile is a model for evaluating evidence based on criminal law. Each piece of evidence is judged individually and rated with a corresponding level of strength, from 1 (lowest) to 5 (highest). This way, evidence strength is represented by an ordinal scale with no zero. In addition, the profile distinguishes between positive and negative evidence, with positive evidence supporting the theory in question and negative evidence contradicting it. It is important to stress that evidence of low quality is not synonymous with negative evidence. Negative evidence can be of high quality, but it opposes the notion that a threat has an influence on the validity of a primary study. For example, negative evidence with a score of -5 has the same evidence strength as positive evidence with a score of 5. Thus, a score close to zero indicates the strength of a piece of evidence is low, while a negative or positive score indicates the outcome of the study.

Table I describes the different levels of strength for evidence. Due to the flexible nature of the evidence profile, the descriptions do not match the ones provided by Wohlin word-for-word. Wohlin emphasizes that the evidence profile should be adapted to the context in which it is used. For example, quality aspects such as *vested interest* and the *aging of evidence* do not play a major role when collecting evidence on threats to validity, unless the threat pertains to a specific technology or approach that may influence experiment results. By contrast, the methodological rigor captured in the *quality of evidence* as well as the *relevance of the evidence* is of utmost importance and strongly influenced the descriptions of the different levels of evidence strength.

The placement of a study on a particular level is based on its adherence to the factors described in the level description, as well as on the previously mentioned quality aspects. Therefore, a study may be placed in a lower category if it has significant shortcomings regarding any of the quality aspects. The concrete definitions of what each level of strength means

were established before the evaluation process began. Because the evaluation of evidence is a largely subjective process, we involved two researchers in this step. Once both researchers had scored a piece of evidence, they compared their scores and discussed possible disagreements. In these discussions, each researcher explained their reasoning for their score, and, together, they decided on a single final score. Where the initial scores were identical, the final score was set to the same value without discussion. We documented the score of each study and the motivation behind its placement, which can be used to paint an overall picture of the evidence landscape for each threat. We documented scoring and agreement in our replication package (Section VII). Based on each evidence profile, we provide conclusive recommendations to researchers.

IV. RESULTS

We first present the results of extracting and categorizing validity threats from 95 primary studies. This provides us with an answer to **RQ1**. We then answer **RQ2** by presenting the evidence profiles for the three most frequently reported threats to validity in code comprehension experiments.

RQ1 Which validity threat categories are most often discussed in primary studies of code comprehension?

Among the 95 papers, 81 (85%) mentioned at least one threat to validity, with 45 (47%) reporting them in a dedicated section and 33 (35%) differentiating between different types of validity, such as internal or external validity. The largest concentration of studies was published in the period from 2012 to 2019, with a total of 63 of the 95 studies (66%), while only covering 7 years (18%) of the entire 40-year span. The first study in our set of papers to include a dedicated “threats to validity” section was a study published in 2005 [29].

In total, we found 409 individual threat mentions. Out of the 409 threat mentions, 198 (48%) included an explanation of a possible mitigation technique, but only 31 (8%) were reported with supporting evidence. The 31 references to supporting evidence were found in 20 out of 95 studies (21%). The 409 threat mentions were then assigned 215 unique threat codes, which captured the different nuances of how a threat was mentioned in a study. Multiple threat mentions could receive the same code, which is why the number of unique codes is lower than the number of total threat mentions. To better analyze related threat codes, we additionally categorized them, which resulted in 81 unique threat categories. For example, the threat category *Programming Experience* included threat codes such as “Missing diversity in participants’ programming experience leads to limited generalizability” and the opposite code “Diversity in participants’ programming experience confounds treatment effects.” The two individual threat codes differed in nuance, but they both emphasized the importance of programming experience as a threat to validity. When we counted the number of threat mentions to prioritize the three most common threat categories, each occurrence of either threat code would increase the category’s count.

Table II highlights the threat categories with more than five reported threat mentions and shows the themes to which

TABLE II
NUMBER OF THREAT MENTIONS PER THEME AND CATEGORY FOR CATEGORIES WITH MORE THAN FIVE THREAT MENTIONS.

Theme and Category	Count
Theme: Code Snippets	112
<i>Program Length</i>	26
Complexity	16
Code Selection	13
Programming Language	9
Synthetic Samples	9
Familiarity	6
Other Categories	33
Theme: Participant Factors	101
<i>Programming Experience</i>	44
Number of Participants	16
Programming Skills	10
Other Categories	31
Theme: Experimentation	89
Learning Effect	18
Lab Experiment	11
Fatigue	9
Code Presentation	7
Cheating	6
Other Categories	38
Theme: Measurement	67
<i>Comprehension Measures</i>	22
Eye-Tracking	20
Instrumentation	9
Other Categories	16
Theme: Comprehension Tasks	21
Type of Comprehension Task	7
Task Difficulty	6
Other Categories	8
Theme: Data Analysis	14
Statistics	10
Other Categories	4
Theme: Other	5
Total	409

the categories were assigned. Both categories and themes are presented with the total number of threat mentions. All mentions from categories with less than five mentions are counted in “Other Categories” under their respective themes. “Theme: Other” contains mentions for threat codes that did not fit into any given theme. The three most common threat categories (i.e., *Programming Experience*, *Program Length*, and *Comprehension Measures*) are highlighted in blue. Overall, most threats were related to the characteristics of the code snippets, the individual factors of the participants, or general threats in experimentation.

RQ1: Main Findings

- 85% of code comprehension experiments report at least one threat to validity.
- Only 8% of threat mentions are supported with referenced evidence.
- The three most commonly reported threat categories are **programming experience**, **program length**, and **comprehension measures**.

TABLE III
OVERVIEW OF THE EVIDENCE ANALYZED AND FILTERED IN EACH STEP OF THE EVIDENCE COLLECTION FOR PROGRAMMING EXPERIENCE.

Source	Analyzed	Excluded	Final
a) Primary Papers	95	-52	43
b) Cited in Primary Papers	13	-12	1
c) Title Search	54	-52	2
d) Snowballing	276	-262	14
Evidence Profile	60	-11	49
Total	438	-389	49




Fig. 2. Evidence Profile 1: Programming experience influences code comprehension.

RQ2 What are the evidence profiles for the three most commonly reported threats to validity in code comprehension experiments?

We selected the three threat categories that were most frequently reported and collected evidence on their influence on different facets of validity. We provide the complete evidence lists, including rationales for the placement of studies in the different evidence categories of an evidence profile, in our replication package (Section VII).

1) *Programming Experience*: In this analysis, we examine the effect a participant's programming experience has on their code understanding in comprehension experiments. We collected evidence in the form of studies that measured the programming experience of participants and determined whether it had a significant influence on program comprehension. Table III shows how many papers we found in each step of the evidence collection and how many we excluded because they did not meet the inclusion criteria. Overall, most evidence was found in the primary papers, followed by snowballing. We excluded 12 out of 13 (92%) documents cited in the references of the primary papers, as they did not meet the filtering criteria. The reasons for exclusion varied and are described in more detail in Section V.

After filtering, 60 papers remained and were evaluated as potential pieces of evidence to be used in the evidence profile.

TABLE IV
OVERVIEW OF THE EVIDENCE ANALYZED AND FILTERED IN EACH STEP OF THE EVIDENCE COLLECTION FOR PROGRAM LENGTH.

Source	Analyzed	Excluded	Final
a) Primary Papers	95	-78	17
b) Cited in Primary Papers	3	-3	0
c) Title Search	29	-29	0
d) Snowballing	0	0	0
Evidence Profile	17	-4	13
Total	127	-114	13




Fig. 3. Evidence Profile 2: Program length influences code comprehension.

In this evaluation, 11 additional papers were discarded as they did not meet the criteria to be considered evidence. For example, one paper was excluded because it was not peer-reviewed and multiple papers were excluded because they either did not measure programming experience or they did not use the collected experience data in their analysis. Fig. 2 presents the final evidence profile. The result largely indicates that programming experience influences code comprehension. In total, 37 (76%) pieces of evidence were rated as positive evidence and 12 (24%) were rated as negative evidence. Furthermore, there were 11 pieces of strong positive evidence and no strong negative evidence.

2) *Program Length*: In this analysis, we examine how the length of a program affects program comprehension. We gather evidence in the form of studies that measure the length of a program and examine its impact on program comprehension. Table IV shows how many papers we found in each step of the evidence collection and how many we excluded because they did not meet the inclusion criteria. The only evidence was found in the list of primary papers. All documents cited in the references of the primary papers were excluded as they did not meet the filtering criteria. Moreover, as we found no relevant papers in b) and c), no snowballing was performed.

After filtering, 17 papers remained and were evaluated as potential evidence using the evidence profile. In this evaluation, 4 more papers were discarded as they did not meet the

TABLE V
OVERVIEW OF THE EVIDENCE ANALYZED AND FILTERED IN EACH STEP OF THE EVIDENCE COLLECTION FOR COMPREHENSION MEASURES.

Source	Analyzed	Excluded	Final
a) Primary Papers	95	-88	7
b) Cited in Primary Papers	3	-3	0
c) Title Search	53	-51	2
d) Snowballing	134	-131	3
Evidence Profile	12	-5	7
Total	285	-278	7




Fig. 4. Evidence Profile 3: Different comprehension measures do not correlate with each other.

criteria to be considered evidence. The final evidence profile is shown in Fig. 3. We found conflicting results regarding the influence of program length on program comprehension: 6 (46%) pieces of evidence were rated as positive evidence and 7 (54%) were rated as negative evidence. There were 2 pieces of strong positive evidence and no strong negative evidence.

3) *Comprehension Measures*: Finally, we examine whether common comprehension measures are associated with distinct aspects of comprehension. We gather evidence in the form of comparative studies that analyze correlations between commonly used comprehension measures. Positive evidence describes that different comprehension measures are not correlated and if different measures were used in a primary study, one may arrive at a different conclusion. Table V shows how many papers were found in each step of the evidence collection and how many were excluded because they did not meet the inclusion criteria. Overall, we found most evidence in the primary papers, with slightly less evidence found in both the title search and through snowballing. All documents cited in the references of the primary papers were excluded as they did not meet the filtering criteria.

After filtering, 12 papers remained and were evaluated as potential pieces of evidence using the evidence profile. In this evaluation, 5 more papers were discarded as they did not meet the criteria to be considered evidence. The final evidence profile is shown in Fig. 4. Most evidence supported that the

commonly used comprehension measures do measure distinct aspects of program comprehension and are not correlated. But overall, only a small amount of evidence was found: 5 (71%) pieces of evidence were rated as positive evidence and 2 (29%) were rated as negative evidence. We found 3 pieces of strong positive evidence, but no strong negative evidence.

RQ2: Main Findings

- **Programming experience:**
37 positive and 12 negative evidence
- **Program length:**
6 positive and 7 negative evidence
- **Comprehension measures:**
5 positive and 2 negative evidence
- 94% of cited evidence in the primary papers did not meet our criteria for evidence and was excluded.
- No evidence was categorized as strong negative evidence (-5).

V. DISCUSSION

We will first consider the results of the evidence profiles on the three most frequently discussed threats to validity.

Overall, the **first evidence profile** confirms that, often, programming experience does influence the comprehension performance of programmers. In multiple cases, experienced programmers showed different comprehension behavior when compared to novices [30], [31] and this difference could be measured in their performance [32]–[34]. In contrast, however, we also found credible evidence contradicting those assertions. Twelve studies reached the conclusion that programming experience does not influence code comprehension. One explanation for this contradiction could be the specific contextual factors of each study. Siegmund et al. [35] found that depending on how programming experience is measured and operationalized, its predictive power varies. Moreover, some negative evidence still found correlations with very specific types of experience measures such as self-estimated Java knowledge [36] or correlations for only specific comprehension measures such as the number of eye fixations [37]. In other cases, the range of programming experience was quite limited, for example due to only including students as participants [36], [38]. These results imply that mentioning programming experience as potential threat alone is insufficient when publishing experiment results. Researchers should instead discuss the threat within the context of their study and discuss how and why they adapted procedures, measures, and artifacts to mitigate it.

The **second evidence profile** indicates that an influence of program length on code comprehension behavior or performance should not be assumed blindly in every context. We found conflicting evidence for this potential threat to validity. Ribeiro et al. [39] discuss similarly conflicting evidence on program length in their study. Even after conducting a follow-up study, they were unable to reach a clear conclusion. Novices

and experts had differing opinions regarding program length and its influence on comprehensibility and readability. They also mention that differences in the procedures and tools used to measure lines of code may affect the comparability of results from different studies. While the overall number of samples is small, we can identify some patterns when comparing positive and negative evidence. All studies rated as positive evidence included students as participants, and one of them included both students and experts. Every study that only included experts was rated negative. This pattern suggests that while program length may influence the comprehension performance of novices, experts appear to perform consistently, regardless of program length.

The results of the **third evidence profile** mostly support the notion that the commonly reported comprehension measures are not correlated. The three supporting pieces of evidence found no correlations between the time and accuracy of comprehension task performances, suggesting that they measure different effects, aspects, or dimensions of comprehension task difficulty [40]–[42]. However, both pieces of negative evidence found correlations between the time and accuracy of comprehension tasks [43], [44]. Furthermore, the two studies that compared physiological measures with other measures found no correlation with subjective ratings [45] and time and accuracy of comprehension tasks [45], [46]. The significance of this finding becomes more apparent when one considers that of the 95 primary papers, more than one third (37) used only a single measure to assess the comprehension performance of participants. Depending on the study context, this may cause them to miss aspects of code understanding that would have been uncovered if they had analyzed more than one measure. If different proxy measures are associated with different aspects of program comprehension, this poses further difficulties for meta-studies of comprehension experiments. For example, two studies might obtain different results because they measured comprehension performance in different ways and not due to other intrinsic factors. Using the results from this work, we were unable to identify concrete patterns in the context factors which would suggest why comprehension measures correlate in some instances and not in others. A dedicated follow-up study, solely focused on comparing different comprehension measures, may shed more light on this question. For now, combining studies with different comprehension measures should be done with care until we better understand the way they measure different aspects of comprehension.

When compared, we found much less evidence overall for the influence of program length and the comparability of comprehension measures than for the influence of programming experience. Even though program length and comprehension measures are the second and third most discussed threats to validity in code comprehension experiments, there are few studies examining their influence. With the little evidence we found, we nevertheless arrived at similar results to the first evidence search: The extent of all three validity threats is context-dependent. The influence of a validity threat varies for each individual study and must, as such, be interpreted in the

light of the contextual factors surrounding it. When seeking patterns to explain the contradiction of evidence, we found that the way a threat impacts the result of a study can depend on how a confounding factor is measured, what the sample population is, and how the experiment is designed in general.

While this conclusion might sound obvious at first, we often find generic statements about a potential threat in primary studies. For example, a study might solely mention that their experiment sample consisted of students without further elaboration. Discussing a study characteristic this way, however, does not clarify why that design decision might constitute a threat to the validity of that specific study. This issue is problematic regardless of whether references are used to back up assumptions, as generic threats are used in place of more nuanced discussions that take the study-specific context factors into account. Evidence should be cited as an additional layer to further support the explanations. For example, Jbara and Feitelson [47] show how evidence from past studies can be used when discussing the threats to validity of their study. They describe how their use of undergraduate students was adequate because the students were at a high enough skill level to complete the experimental task and cite evidence to support their assertion. In this way, they use evidence to guide the threat discussion within the context of their study.

Minor Observations

Besides this main finding, we also made some further observations. First, the studies used in the primary papers to support claims about threats to validity were, with one exception, almost entirely dismissed as evidence. Studies were mostly excluded because they either did not relate to program comprehension, solely cited another mention of the threat, or they were not relevant at all. This pattern was consistent over all three evidence searches.

Second, the results of our review of how threats to validity are reported match the investigation by Wyrich et al. [4]. In general, researchers in program comprehension tend to focus on threats relating to the code snippets and the study participants the most. The factors we investigated in our evidence search, programming experience, program length, and comprehension measures are among the most frequently mentioned threats in both studies. The terms used to describe these threats differ between the works. For example, the category named “comprehension measures” can be found in the category “instrumentation” in Wyrich et al. and “program length” is named “program size”. Looking at the bigger picture, the results for program comprehension studies are in line with reporting in SE in general. Siegmund et al. found that in SE, 51% (47% in our study) of papers discussed threats to validity in a dedicated section and 23% (35% in our study) differentiated between different types of validity [48].

Third, we found that, overall, there was less evidence for less common threats to validity. This is in line with what would be expected intuitively, if fewer researchers deem that a threat poses a danger to a study’s validity then corollary, fewer will investigate whether that assumption holds. Meta-studies

such as this one can highlight which threats are less frequently examined and provide directions for further experimentation.

Fourth, we observed that throughout the entire evidence collection process, not a single study was classified as a “-5”, which would represent strong negative evidence. While we do not have a conclusive reason for this, it is common for researchers to experience difficulties publishing negative results. Obtaining negative results can be disheartening, but publishing them is nonetheless crucial and provides value to the scientific community [49]–[51].

A. On the Desire for Evidence

Looking back at The Highway Shooter case, we also want to acknowledge some key differences between the criminal case and a scientific experiment. First, in the media coverage following the case, the recording and usage of license plates as evidence was criticized as a breach of privacy. In the context of academic research, similar issues with data privacy mainly arise when collecting data for primary research. When synthesizing results of different independent studies, however, applying their research findings as evidence with proper citation is usually encouraged. Furthermore, when creating evidence profiles, synthesis is even possible without disclosing the data sets of the studies. The information contained in the paper itself is often sufficient to assess its weight as evidence. Another important difference is that there was only one perpetrator responsible for the crimes, whereas in any given experiment, there may be more than one validity threat.

When designing an experiment, it is infeasible to conduct systematic reviews to collect evidence for each of the dozens of potential threats. Rather, researchers may look into existing research summarizing evidence and then contextualize it by comparing the underlying studies with the characteristics of their own methodology. The present paper can be one such resource for evidence collection, future meta-studies on common research questions another. Evidence is then used both during the design stages of an experiment and when interpreting and discussing its results.

Conversely, one may discuss the nature and role of the threats to validity section in scientific literature itself. While we examined and reported on various descriptive aspects of the threats reported in existing papers, we did not inquire into why researchers chose to report and discuss specific threats in a certain way. Furthermore, our research does not provide explicit guidance on how a paper author ought to write a *threats to validity* section. While the position presented in this paper suggests that reported threats should be supported with evidence whenever possible, this is not necessarily a sentiment shared by all members of the scientific community. A section on validity threats may also be a place where researchers should be able to speculate without concrete evidence and point out potential shortcomings as directions for future research. This debate should be conducted within the respective research communities, and agreements should then be recorded and incorporated into existing guidelines for the reporting of validity threats.

B. Implications

The results presented in this work lead us to concrete suggestions for the larger research community. First, threats to validity are dependent on individual context factors of a study. Researchers should therefore explicitly discuss the influence of each threat within the context of their study. Merely listing off potential threats alone is not sufficient. Second, context discussions for validity threats should be supported by existing evidence, rather than relying on intuition or speculation. Meta-studies provide the appropriate evidence base for this endeavor, as they analyze the influence of a threat in multiple different study contexts.

C. Limitations

Evidence profiles are created by qualitative evidence evaluation from human reviewers, and thus may contain bias. To mitigate this threat, two researchers independently rated the evidence and then compared their results to reach an agreement. The threat extraction, categorization, and subsequent evidence collection was done by only one researcher, which again might incur a bias and threaten the internal validity of our study. However, when comparing the threat categories and number of occurrences with a previous systematic review on confounding factors in program comprehension studies by Siegmund and Schumann [3], we find mostly similar results regarding threat codes and their frequencies.

Throughout this work, we used the terms program comprehension and code comprehension interchangeably for comprehension of source code, as is currently the norm. However, in principle, we agree with Wyrich et al.’s observation that code comprehension studies form a subset of program comprehension studies [4]. The list of primary papers we built on was a result of a previous search by Wyrich et al. and its search parameters were not fully aligned to the goals of our work. In their study, Wyrich et al. focused on a subset of program comprehension experiments, namely experiments on bottom-up code comprehension. This means that one might obtain a different ranking for the most common threats in the broader area of program comprehension experiments, limiting the external validity of our results. In the evidence collection, this limitation was mitigated as we added additional search sources.

When analyzing our data, we also found a rather strong bias in the number of studies with quantitative data from experiments over qualitative data. Even though we did not explicitly exclude any papers in the evidence search on the grounds of them being a qualitative study, the results still heavily favored controlled experiments. In numbers, 68 out of 69 (99%) pieces of evidence in the evidence profiles were from experiments. To avoid such bias in future studies, we recommend formulating descriptions of evidence profiles in a more method-neutral terminology (e.g., avoiding terms such as confounding factors, which are primarily indicative of experiments). On a related note, new evidence profiles for validity threats from code comprehension studies using other research methodologies than experimentation can, provided

equal or similar epistemological stances, help similarly well in making informed design decisions. It depends on one's epistemological standpoint and understanding of the concept of evidence to determine whether it is appropriate to use evidence in the design of non-experimental studies.

When developing the evidence profiles, we decided to focus on threat categories instead of individual codes. For example, rather than searching for evidence of programming experience as threat to internal or external validity, we combined the two into the category of programming experience. A finer distinction was not necessary for our purpose, as evidence for the influence of programming experience on code comprehension is equally relevant for all validity types. Further, a finer distinction would have been difficult to make. As previously noted, validity threats in primary studies are currently not discussed in a way that makes clear what type of validity the authors believe their study characteristics would affect.

Another limitation in the external validity of our study is that due to the amount of effort required to systematically collect evidence, we were only able to analyze three of dozens of different validity threats. We focused on the most frequently mentioned threats, as we expected them to provide the most value for the widest range of program comprehension studies. It is possible that repeating the same methods for the remaining threats will uncover new idiosyncrasies of program comprehension. As such, generalizing the results obtained in this work to other threats should be done with care.

VI. CONCLUSION

In this work, we investigated threats to validity in program comprehension experiments. First, we analyzed the state of the art by reviewing the reported validity threats in 95 papers of primary research. We found that while most studies mentioned threats to validity, few supported them with corresponding evidence. Furthermore, only in one case did the cited evidence support the validity of a threat and meet our quality criteria to be included in an evidence profile. Next, we searched for evidence regarding the three most frequently mentioned threat categories. Our evidence collection yielded both positive and negative evidence for each threat category, seemingly leading to divergent results.

After looking closer into the collected evidence and comparing the characteristics of the different studies, we concluded that validity threats are highly context-dependent. Even the threats that are intuitively expected to affect program comprehension, such as programming experience, depend on how they are measured, the sample population, comprehension tasks, and other context factors. Therefore, we must consider all individual characteristics of a study when assessing potential validity threats to develop methodologies that use evidence as a basis for implementing appropriate mitigation techniques. Furthermore, discussions of validity threats in papers need to explicitly address context factors and researchers should use evidence to support these discussions. Finally, we encourage the usage of threat catalogs and recommend the adherence

to reporting guidelines for threats to validity to improve the reproducibility and comparability of study results.

Structured guidelines for reporting threats to validity need to be established further to inform researchers on how they can incorporate evidence into their validity assessments. We need more knowledge documentation on which threats exist, the evidence supporting them, the context in which they occur, and which mitigation techniques can be used to address them. Previous works laid the groundwork in this endeavor by documenting threats in software engineering studies in a knowledge base and providing guidelines for controlling the influence of confounding factors in program comprehension experiments, respectively [3], [22]. These works can be further extended by incorporating evidence and documenting threats to validity in a common database.

We envision a future in which scientists regularly apply our methodology to collect and summarize the evidence for additional threats to validity and for experiments far beyond code comprehension. Even when looking at just three of the most common threats, we found less evidence for less popular threats, which might indicate even bigger gaps in evidence for those threats that have yet to be analyzed. Moreover, our approach may be applied to other domains to investigate how threats to validity affect experiments there.

Taking a final look back at the past, in 2014, The Highway Shooter was sentenced to 10 years in prison. Good investigative work led to the necessary evidence and ultimately to the apprehension of the criminal. Evidence profiles in software engineering research equally lead to necessary data points, not to catch criminals, but, in our case, to gain more certainty about the actual impact of threats to validity in program comprehension experiments.

VII. DATA AVAILABILITY

We publish our dataset on Zenodo for transparency and reproducibility of our approach [52]. The dataset comprises data that emerged in intermediate steps of the validity threat analysis. This includes extracted text passages from the primary studies and complete lists of threat codes, threat categories and threat themes. Further, the dataset contains artifacts produced during the evidence search and the creation of the evidence profiles. We document comprehensibly, for example, how individual reviewers evaluated individual pieces of evidence and reached a final agreement. Finally, R scripts are included, that mainly deal with quantitative analyses and plotting.

ACKNOWLEDGMENTS

We thank members of the German Federal Criminal Police Office (BKA), who wish to remain anonymous, for their information and guidance on The Highway Shooter case, which inspired us when presenting our findings. We are very grateful to three anonymous reviewers, whose feedback helped to improve the paper. We also thank Katharina Plett for proofreading to enhance the text quality.

REFERENCES

- [1] Pressestelle des Bundesgerichtshofs, "Press release of the federal court of justice of october 8, 2015," 2015. [Online]. Available: <https://www.bundesgerichtshof.de/SharedDocs/Pressemitteilungen/DE/2015/2015171.html>
- [2] C. Wohlin, *An Evidence Profile for Software Engineering Research and Practice*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 145–157.
- [3] J. Siegmund and J. Schumann, "Confounding parameters on program comprehension: a literature survey," *Empirical Software Engineering*, vol. 20, no. 4, pp. 1159–1192, 2015.
- [4] M. Wyrich, J. Bogner, and S. Wagner, "40 years of designing code comprehension experiments: A systematic mapping study," *arXiv*, 2022. [Online]. Available: <https://arxiv.org/abs/2206.11102>
- [5] I. Schröter, J. Krüger, J. Siegmund, and T. Leich, "Comprehending studies on program comprehension," in *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*, 2017, pp. 308–311.
- [6] D. L. Sackett, W. M. Rosenberg, J. M. Gray, R. B. Haynes, and W. S. Richardson, "Evidence based medicine: what it is and what it isn't," pp. 71–72, 1996.
- [7] T. Dyba, B. A. Kitchenham, and M. Jorgensen, "Evidence-based software engineering for practitioners," *IEEE software*, vol. 22, no. 1, pp. 58–65, 2005.
- [8] A. Stevenson, *Oxford dictionary of English*, 3rd ed. Oxford University Press, USA, 2010.
- [9] B. A. Kitchenham, T. Dyba, and M. Jorgensen, "Evidence-based software engineering," in *Proceedings. 26th International Conference on Software Engineering*. IEEE, 2004, pp. 273–281.
- [10] B. Kitchenham, R. Pretorius, D. Budgen, O. P. Brereton, M. Turner, M. Niazi, and S. Linkman, "Systematic literature reviews in software engineering—a tertiary study," *Information and software technology*, vol. 52, no. 8, pp. 792–805, 2010.
- [11] D. Budgen and P. Brereton, "Evolution of secondary studies in software engineering," *Information and Software Technology*, vol. 145, p. 106840, 2022.
- [12] M. Jorgensen, T. Dyba, and B. Kitchenham, "Teaching evidence-based software engineering to university students," in *11th IEEE International Software Metrics Symposium (METRICS'05)*. IEEE, 2005, pp. 8–pp.
- [13] A. Rainer and S. Beecham, "A follow-up empirical evaluation of evidence based software engineering by undergraduate students," in *12th International Conference on Evaluation and Assessment in Software Engineering (EASE) 12*, 2008, pp. 1–10.
- [14] B. J. Oates and G. Capper, "Using systematic reviews and evidence-based software engineering with masters students," in *Proceedings of the 13th international conference on Evaluation and Assessment in Software Engineering*. BCS Learning & Development Ltd., 2009, pp. 79–87.
- [15] E. Hassler, J. C. Carver, N. A. Kraft, and D. Hale, "Outcomes of a community workshop to identify and rank barriers to the systematic literature review process," in *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*, ser. EASE '14. New York, NY, USA: Association for Computing Machinery, 2014. [Online]. Available: <https://doi.org/10.1145/2601248.2601274>
- [16] R. E. Santos and F. Q. d. Silva, "Motivation to perform systematic reviews and their impact on software engineering practice," in *2013 ACM / IEEE International Symposium on Empirical Software Engineering and Measurement*, 2013, pp. 292–295.
- [17] B. Cartaxo, G. Pinto, E. Vieira, and S. Soares, "Evidence briefings: Towards a medium to transfer knowledge from systematic reviews to practitioners," in *Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ser. ESEM '16. New York, NY, USA: Association for Computing Machinery, 2016. [Online]. Available: <https://doi.org/10.1145/2961111.2962603>
- [18] B. Cartaxo, G. Pinto, and S. Soares, "Rapid reviews in software engineering," in *Contemporary Empirical Methods in Software Engineering*. Springer, 2020, pp. 357–384.
- [19] B. A. Kitchenham, *Evidence-Based Software Engineering and Systematic Reviews*, ser. Chapman & Hall / CRC Innovations in Software Engineering and Software Development Series. Boca Raton: CRC Press, 2015, vol. v.4.
- [20] K. Petersen and C. Gencel, "Worldviews, research methods, and their relationship to validity in empirical software engineering research," in *2013 Joint Conference of the 23rd International Workshop on Software Measurement and the 8th International Conference on Software Process and Product Measurement*, 2013, pp. 81–89.
- [21] P. Devanbu, T. Zimmermann, and C. Bird, "Belief & evidence in empirical software engineering," in *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*, 2016, pp. 108–119.
- [22] S. Biffel, M. Kalinowski, F. Ekaputra, A. A. Neto, T. Conte, and D. Winkler, "Towards a semantic knowledge base on threats to validity and control actions in controlled experiments," in *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ser. ACM Digital Library, M. Morisio, Ed. New York, NY: ACM, 2014, pp. 1–4.
- [23] V. Rajlich and G. Cowan, "Towards standard for experiments in program comprehension," in *Proceedings Fifth International Workshop on Program Comprehension. IWPC'97*, 1997, pp. 160–161.
- [24] M. Wyrich, J. Bogner, and S. Wagner, "Replication package: 40 Years of Designing Code Comprehension Experiments: A Systematic Mapping Study (v1.0)," Jun. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6657640>
- [25] D. S. Cruzes and T. Dyba, "Recommended steps for thematic synthesis in software engineering," in *2011 International Symposium on Empirical Software Engineering and Measurement*. IEEE, 2011, pp. 275–284.
- [26] J. Corbin and A. Strauss, *Basics of Qualitative Research (3rd ed.): Techniques and Procedures for Developing Grounded Theory*. 2455 Teller Road, Thousand Oaks California 91320 United States: SAGE Publications, Inc, 2008.
- [27] C. Wohlin, "Guidelines for snowballing in systematic literature studies and a replication in software engineering," in *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering - EASE '14*. ACM Press, 2014.
- [28] D. Badampudi, C. Wohlin, and K. Petersen, "Experiences from using snowballing and database searches in systematic literature studies," in *Proceedings of the 19th International Conference on Evaluation and Assessment in Software Engineering*. ACM, 2015.
- [29] B. Du Bois, S. Demeyer, and J. Verelst, "Does the "refactor to understand" reverse engineering pattern improve program comprehension?" in *Ninth European Conference on Software Maintenance and Reengineering*. IEEE, 2005, pp. 334–343.
- [30] B. Adelson, "When novices surpass experts: The difficulty of a task may increase with expertise," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 10, no. 3, pp. 483–495, 1984.
- [31] C. S. Yu, C. Treude, and M. Aniche, "Comprehending test code: An empirical study," in *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2019, pp. 501–512.
- [32] S. Wiedenbeck, "Novice/expert differences in programming skills," *International Journal of Man-Machine Studies*, vol. 23, no. 4, pp. 383–390, 1985.
- [33] P. A. Orlov and R. Bednarik, "The role of extrafoveal vision in source code comprehension," *Perception*, vol. 46, no. 5, pp. 541–565, 2017.
- [34] F. Medeiros, G. Lima, G. Amaral, S. Apel, C. Kästner, M. Ribeiro, and R. Gheyi, "An investigation of misunderstanding code patterns in c open-source software projects," *Empirical Software Engineering*, vol. 24, no. 4, pp. 1693–1726, 2019.
- [35] J. Siegmund, C. Kästner, J. Liebig, S. Apel, and S. Hanenberg, "Measuring and modeling programming experience," *Empirical Software Engineering*, vol. 19, no. 5, pp. 1299–1334, 2014.
- [36] N. Peitek, J. Siegmund, S. Apel, C. Kastner, C. Parnin, A. Bethmann, T. Leich, G. Saake, and A. Brechmann, "A look into programmers' heads," *IEEE Transactions on Software Engineering*, vol. 46, no. 4, pp. 442–462, 2020.
- [37] S. Jessup, S. M. Willis, G. Alarcon, and M. Lee, "Using eye-tracking data to compare differences in code comprehension and code perceptions between expert and novice programmers," in *Proceedings of the 54th Hawaii International Conference on System Sciences*, ser. Proceedings of the Annual Hawaii International Conference on System Sciences, T. Bui, Ed. Hawaii International Conference on System Sciences, 2021.
- [38] D. Hendrix, J. H. Cross, and S. Maghsoodloo, "The effectiveness of control structure diagrams in source code comprehension activities," *IEEE Transactions on Software Engineering*, vol. 28, no. 5, pp. 463–477, 2002.

- [39] T. V. Ribeiro and G. H. Travassos, "Attributes influencing the reading and comprehension of source code – discussing contradictory evidence," *CLEI Electronic Journal*, vol. 21, no. 1, 2018.
- [40] E. R. Iselin, "Conditional statements, looping constructs, and program comprehension: an experimental study," *International Journal of Man-Machine Studies*, vol. 28, no. 1, pp. 45–66, 1988.
- [41] D. Binkley, D. Lawrie, S. Maex, and C. Morrell, "Identifier length and limited programmer memory," *Science of Computer Programming*, vol. 74, no. 7, pp. 430–445, 2009.
- [42] S. Ajami, Y. Woodbridge, and D. G. Feitelson, "Syntax, predicates, idioms — what really affects code complexity?" *Empirical Software Engineering*, vol. 24, no. 1, pp. 287–328, 2019.
- [43] D. J. Gilmore and T. Green, "Comprehension and recall of miniature programs," *International Journal of Man-Machine Studies*, vol. 21, no. 1, pp. 31–48, 1984.
- [44] W. E. Hall and S. H. Zweben, "The cloze procedure and software comprehensibility measurement," *IEEE Transactions on Software Engineering*, vol. SE-12, no. 5, pp. 608–623, 1986.
- [45] M. K.-C. Yeh, Y. Yan, Y. Zhuang, and L. A. DeLong, "Identifying program confusion using electroencephalogram measurements," *Behaviour & Information Technology*, pp. 1–18, 2021.
- [46] S. Fakhoury, D. Roy, Y. Ma, V. Arnaoudova, and O. Adesope, "Measuring the impact of lexical and structural inconsistencies on developers' cognitive load during bug localization," *Empirical Software Engineering*, vol. 25, no. 3, pp. 2140–2178, 2020.
- [47] A. Jbara and D. G. Feitelson, "How programmers read regular code: a controlled experiment using eye tracking," *Empirical Software Engineering*, vol. 22, no. 3, pp. 1440–1477, 2017.
- [48] J. Siegmund, N. Siegmund, and S. Apel, "Views on internal and external validity in empirical software engineering," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*. IEEE, 2015, pp. 9–19.
- [49] P. G. Weintraub, "The importance of publishing negative results," *Journal of Insect Science*, vol. 16, no. 1, p. 109, 2016.
- [50] R. F. Paige, J. Cabot, and N. A. Ernst, "Foreword to the special section on negative results in software engineering," *Empirical Software Engineering*, vol. 22, no. 5, pp. 2453–2456, 2017.
- [51] A. Borji, "Negative results in computer vision: A perspective," *Image and Vision Computing*, vol. 69, pp. 1–8, 2018.
- [52] M. M. Barón, M. Wyrich, D. Graziotin, and S. Wagner, "Replication package: Evidence Profiles for Validity Threats in Program Comprehension Experiments," Aug. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.7038907>