

# Evidence Tetris in the Pixelated World of Validity Threats

Marvin Wyrich  
Saarland University  
Saarbrücken, Germany  
wyrich@cs.uni-saarland.de

Sven Apel  
Saarland University  
Saarbrücken, Germany  
apel@cs.uni-saarland.de

## ABSTRACT

Valid empirical studies build confidence in scientific findings. Fortunately, it is now common for software engineering researchers to consider threats to validity when designing their studies and to discuss them as part of their publication. Yet, in complex experiments with human participants, there is often an overwhelming number of intuitively plausible threats to validity – more than a researcher can feasibly cover. Therefore, prioritizing potential threats to validity becomes crucial. We suggest moving away from relying solely on intuition for prioritizing validity threats, and propose that evidence on the actual impact of suspected threats to validity should complement intuition.

## CCS CONCEPTS

• **General and reference** → **Empirical studies; Experimentation.**

## KEYWORDS

Validity Threats, Limitations, Study Design, Research Evaluation

### ACM Reference Format:

Marvin Wyrich and Sven Apel. 2024. Evidence Tetris in the Pixelated World of Validity Threats. In *International Workshop on Methodological Issues with Empirical Studies in Software Engineering (WSESE '24)*, April 16, 2024, Lisbon, Portugal. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3643664.3648203>

## 1 INTRODUCTION

When designing an empirical study, researchers make sure that the study design is as valid as possible. Validity is a multifaceted construct, and we will cover it in more detail later. For the moment, we use Kitchenham et al.'s brief summary that *validity refers to the degree to which we can trust the outcomes of an empirical study* [7]. Assessing the validity of a study design, that is, whether we can trust the results, requires expert knowledge. No one could make this assessment better than the researchers themselves, which is why we consider it a good development that nowadays threats to validity are discussed in most empirical software engineering studies [1, 3, 14, 23].

For example, one of the most frequently discussed threats to validity in empirical studies with human participants is the sampling of students as a substitute for professional software engineers.

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WSESE '24, April 16, 2024, Lisbon, Portugal

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0567-0/24/04

<https://doi.org/10.1145/3643664.3648203>

To researchers who have commented publicly on this topic, the assumption that a certain treatment influences novice and expert programmers differently makes intuitive sense [4]. As a result, the authors of a study often devote a paragraph in the discussion section of their paper to this potential influencing factor and mention, for example, that their sample consisted solely of students and that therefore the study results cannot be applied readily to more experienced developers.

However, most of these discussions about the validity of a study are purely based on the researcher's intuition and therefore speculative. Hardly anyone is sure about the actual extent of the discussed threats, and almost no paper cites evidence on the assumed threats. Verdecchia et al. [18] describe the current situation as working through a "laundry list" of threats to validity that are only vaguely discussed. We prefer to think of it as a pixelated world of validity threats, with no clarity about their actual impact.

This leads to uncertainty for the researchers when they design their studies, and at the very least, it leads to potential conflicts in peer review, when the reviewer critiques the study design based on their own, different intuition. That researchers have very different views on the assessment of validity was found, for example, by Siegmund et al. [17] in a survey of 79 program committee and editorial board members. These views included even those that would reject papers in principle if they attempted to prioritize internal validity [17]. Note that researchers should generally have their own opinions. The fundamental issue is that personal views currently determine which scientific findings are published and which not – a decision that, apart from ethical considerations, should instead be based on an informed validity assessment. Otherwise, we run the risk that a few researchers will decide the scientific discourse based on their views, while valid and potentially influential minority views will be rejected unfairly.

What is needed is greater consideration of existing evidence on the influences of assumed threats to validity. To stay with the example of student samples: whether certain characteristics of a student, such as limited programming experience, influence that person's performance and behavior in programming activities has been investigated in numerous primary studies. And we will see that the evidence does not justify speaking of a validity threat across the board in every study context. We therefore suggest in this position paper that synthesized evidence on the most commonly assumed threats to validity can and should guide researchers in discussing the validity of their empirical software engineering study.

## 2 EVIDENCE TETRIS

We introduce the concept of *Evidence Tetris* as a systematic approach for navigating validity threats in empirical studies in software engineering research. The approach consists of the following three

steps, which build on each other and must therefore be carried out one after the other within a research field:

- (1) **Threat collection:** Summarize the discussed threats to validity in a subfield of software engineering research.
- (2) **Evidence Synthesis:** Synthesize available evidence on the most discussed threats.
- (3) **Evidence-Based Study Designs:** Consider the synthesized evidence in the design and discussion of each primary study.

The first two steps represent a community effort, on which each individual primary researcher can then build in the third step. We describe the three steps in detail below. We illustrate the procedure and potential outcomes using the efforts of recent years within the code comprehension research community as an example.

## 2.1 Threat Collection

In everyday language, *validity* can denote ‘the quality of being well-grounded, sound, or correct’ [8]. In this work, validity is considered a multi-faceted construct, and the degree of validity of a study design can be assessed separately for each facet. The facets into which validity is usually divided depend on the underlying research method of a study [10]. In controlled experiments in software engineering, for instance, validity classifications into internal, external, construct, and conclusion have gained widespread acceptance in the reporting of validity threats. This classification was particularly successfully promoted by Wohlin et al. [21] and Jedlitschka et al. [6], but in both cases the idea is based on the work of Cook and Campbell [2]. In principle, researchers should aim for study designs that have a high degree of validity in those facets that are most relevant to the researcher’s particular principles and goals. “Some ways of increasing one kind of validity will probably decrease another kind” [2], making this form of prioritization often necessary.

Yet, there still is an overwhelming number of potential threats to validity that researchers can consider and discuss. Two secondary studies have categorized threats to validity discussed in code comprehension studies, coming up with over fifty different threat categories [16, 23]. The categories cover all areas of a study design, from the selection of participants and study materials to task design, research procedure, and data analysis.

It is secondary studies such as these that make a valuable contribution within the first step of Evidence Tetris. Based on already published literature, we can find out which validity threats a research community is concerned with and which are discussed most frequently.<sup>1</sup> Furthermore, such an investigation can show what proportion of papers in a research field actually report threats to validity. In code comprehension research, papers that do not discuss threats to validity are rather the exception today [14, 23].

## 2.2 Evidence Synthesis

Knowing which threats are discussed most frequently helps us to prioritize them in this second step of Evidence Tetris. Our goal here is to synthesize existing evidence on validity threats. In an ideal world, we would synthesize evidence for all potential validity threats. The solely pragmatic rationale for focusing on the most

frequently discussed threats is that the process of evidence synthesis can be time-consuming. Keep in mind that the most discussed threats are not necessarily the ones that have the biggest impact. We only count their mentions to see what concerns the community the most. Then we use evidence synthesis to see if it rightfully concerns them.

There are several ways to review literature in a secondary study [12]. In our context, methods that consider the influence of a threat to validity as a hypothesis and bring research findings on this hypothesis into a comprehensible and manageable structure are suitable. These methods range from classic vote counting to statistical meta-analyses, which in extreme cases reduce research data from several studies to a single effect size.

Muñoz Barón et al. [9] synthesized evidence for the three most frequently discussed threats to validity in code comprehension experiments: the influence of programming experience, program length, and the selected comprehension measures. They created *evidence profiles* [20], a form of data synthesis in which research findings are subdivided according to whether they support or refute the hypothesis, but in which each study is also assigned a qualitative strength of evidence. In an evidence profile, each individual primary study can still be recognized, which later in the third step will become advantageous to identify those studies that come closest to one’s own study context.

Muñoz Barón et al. [9] found that the evidence for the most frequently discussed threats in code comprehension experiments is not as clear-cut as previously assumed. It may well be that in one particular study context, for example, programming experience plays a role as a confounding variable, but in another it does not. Furthermore, for some suspected threats to validity, there is not enough evidence to even draw a conclusion. In such cases, step two of the Evidence Tetris shows where studies are still missing to explicitly investigate the significance of suspected threats.<sup>2</sup>

## 2.3 Evidence-Based Study Designs

Reaching this third and final step, a research community has already matured considerably at the meta-scientific level. Not only has a community systematically reviewed the threats to validity that are frequently discussed within its own field of research; evidence has also been compiled for the most relevant threats to form the basis for an evidence-based discussion of study designs. When designing and reporting a primary study, it would now be up to each researcher to consider the synthesized evidence and cite it in their reasoning for certain study design decisions.

Staying with the example of the frequently discussed influence of programming experience on the performance or behavior of developers in programming tasks. If one intuitively assumes that this influence can also play a role in one’s own study, one should validate this assumption based on an already available evidence synthesis. If the assumption is confirmed, programming experience should be considered accordingly in one’s own study design. In the

<sup>1</sup>This can also be the result of a working group in the case of an emerging subfield.

<sup>2</sup>See, e.g., Wagner and Wyrich’s study [19] on the influence of intelligence and personality on code comprehension, motivated by the sole intention of investigating the actual influence of suspected confounders. Such studies provide targeted evidence for suspected threats to validity, as they make the threat an actual hypothesis of their investigation. This is not the case in all studies that provide potential evidence for suspected threats to validity.

case of a controlled experiment, this could mean that programming experience is controlled as a confounder or covariate<sup>3</sup> (if internal validity is a concern) or that participants with a particular level of experience participate in the study (if external validity is a concern). If this cannot be accomplished, it should be discussed in the paper, citing the evidence for such a threat to validity.

In the case of a frequently suspected threat that evidently plays no role in a study context, a discussion of the threat citing other studies can also be in the researchers' own interest. The evidence then protects them from unjustified criticism. In any case, citing evidence will contribute to a more informed validity assessment. And if authors of primary studies already cite evidence in their discussions of threats to validity, secondary studies can also benefit: contradictory findings can be better explained, and more informed and nuanced conclusions can be made.

### 3 REMARKS

Each of the three steps of Evidence Tetris offer sufficient discussion material to engage in productive community conversations about the specific implementation and implications of evidence-based study designs. We would like to preface these discussions with a few remarks.

#### 3.1 Intuition Remains With Us

The pleasure of science is not the least driven by the challenge of finding innovative solutions to complex problems. Many research questions require creative and ingenious study designs, and human intuition should not be underestimated in this process. While the goal of this paper is to bring evidence into the design and evaluation of study designs, we need to clarify two points.

First, intuition has and must have its place also in the design and discussion of future software engineering studies. Intuition is what leads us to the hypothesis for which we collect evidence, ideas for future directions to look into, and sometimes the one methodological idea that nobody has thought about before. Evidence will not replace intuition, but it can help a researcher defend their intuition against unwarranted criticism.

Second, the collection of evidence that this paper aims for is *not* driven by a goal of coming up with the one and only true study design for a specific research question. This line of reasoning may be intuitive, as one may at some point discern which design decisions are most logical based on empirical evidence. And indeed, evidence will allow researchers to have meaningful discussions about their study design and how valid it is. However, we should not forget that the validity of a study design is not binary: There is no evidence that would suggest that a study design is valid or invalid. Evidence might tell us that a research design might be limited in some ways, and sometimes the researcher has to choose which facet of validity is important to them. This is an intuitive process of compromise.

#### 3.2 Evidence and Truth

Many years of philosophical discussions about the nature of knowledge lie behind us. If they have taught us one thing, it is that different people have different views on what evidence is and when

<sup>3</sup>Ideally, a validated instrument should be used. See, e.g., the work by Siegmund et al. [15] on measuring programming experience.

we have enough of it to ascribe a certain temporary truth to a theory [5].

For our purpose of supporting researchers in an evidence-based discussion of their empirical studies, we take a pragmatic approach and consider pieces of evidence as any information that increases our confidence in a theory. Whether one strictly follows Popper's critical rationalism [11] or allows any observation to inductively contribute to this confidence in the truth of a theory, as the logical empiricists once did [5], is left to the respective researcher. Any conscious consideration of potential threats to validity in the design and reporting of studies based on past findings and observations would already represent progress compared to the current situation, in which threats are only superficially considered in a checklist-like manner [9, 18].

Regarding the question of when we have enough evidence to accept something as sufficiently supported, we do not necessarily need an answer whether something is *actually* as we assume or not. It is sufficient to know what speaks for and against to then be able to take an informed position. This approach aligns with the thinking of Richard Rudner, who believed that evidence alone is insufficient, but that it takes a decision to accept the underlying theory [13]. Thus, while we strive to draw on more evidence when evaluating study designs, we should keep these aspects in mind and, with the evidence in hand, never argue for a single absolute truth.

#### 3.3 The Nature of 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 dedicated to validity threats may also serve as a venue for researchers to speculate without concrete evidence and identify potential deficiencies as avenues 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.

We speculate that views will vary as much as they did in Siegmund et al.'s survey of 79 program committee and editorial board members about their views on the importance and trade-off between internal and external validity [17]. Future work could proceed and similarly investigate, for example, why researchers would (not) use evidence for the evaluation of empirical studies. Such an investigation will potentially strike a chord and bring to light different perspectives on what constitutes relevant evidence and how much the research community may want to rely on its own past contributions. If there are extreme views in our community, we could potentially gain insights into the causes of inconsistency in peer review decisions that have not been explored before. And then we need to consider how to deal with the possibility that individual philosophical views may determine the publication of research findings.

### 4 CONCLUSION

The consideration and discussion of validity threats is not an end in itself or, worse still, like working through a laundry list. Over the past few years, we and other researchers in the software engineering

research community have recognized that there is greater potential for more informed validity assessments.

With our Evidence Tetris approach, we intend to help prioritize the overwhelming number of potential validity threats in a meaningful way. It is a mix of community effort and the individual efforts of each researcher. We have illustrated that this can work using the example of the code comprehension research community, which has already successfully completed the first two of three steps. We are confident that in the third step we will see a change in the literature towards validity discussions being more evidence-based than before. And of course, we envision a similar momentum in other software engineering research communities.

## ACKNOWLEDGMENTS

We would like to thank three anonymous reviewers for their constructive criticism. Wyrich's and Apel's work is supported by the European Union as part of the ERC Advanced Grant 101052182. Parts of this paper are drawn from Wyrich's doctoral thesis [22].

## REFERENCES

- [1] Apostolos Ampatzoglou, Stamatia Bibi, Paris Avgeriou, Marijn Verbeek, and Alexander Chatzigeorgiou. 2019. Identifying, categorizing and mitigating threats to validity in software engineering secondary studies. *Information and Software Technology* 106 (2019), 201–230. <https://doi.org/10.1016/j.infsof.2018.10.006>
- [2] Thomas D Cook and Donald Thomas Campbell. 1979. *Quasi-experimentation: Design & analysis issues for field settings*. Vol. 351. Houghton Mifflin Boston.
- [3] Robert Feldt and Ana Magazinius. 2010. Validity threats in empirical software engineering research—an initial survey. In *Seke*. 374–379.
- [4] Robert Feldt, Thomas Zimmermann, Gunnar R Bergersen, Davide Falessi, Andreas Jedlitschka, Natalia Juristo, Jürgen Münch, Markku Oivo, Per Runeson, Martin Shepperd, et al. 2018. Four commentaries on the use of students and professionals in empirical software engineering experiments. *Empirical Software Engineering* 23, 6 (2018), 3801–3820.
- [5] Peter Godfrey-Smith. 2021. *Theory and Reality: An Introduction to the Philosophy of Science, Second Edition*. University of Chicago Press.
- [6] Andreas Jedlitschka, Marcus Ciolkowski, and Dietmar Pfahl. 2008. Reporting experiments in software engineering. In *Guide to advanced empirical software engineering*. Springer, 201–228.
- [7] Barbara Ann Kitchenham, David Budgen, and Pearl Brereton. 2015. *Evidence-Based Software Engineering and Systematic Reviews*. Chapman & Hall / CRC Innovations in Software Engineering and Software Development Series, Vol. 4. CRC Press, Boca Raton.
- [8] Merriam-Webster. [n. d.]. Validity. In *Merriam-Webster.com dictionary*. <https://www.merriam-webster.com/dictionary/validity>
- [9] Marvin Muñoz Barón, Marvin Wyrich, Daniel Graziotin, and Stefan Wagner. 2023. Evidence Profiles for Validity Threats in Program Comprehension Experiments. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, IEEE.
- [10] Kai Petersen and Cigdem Gencel. 2013. 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*. 81–89. <https://doi.org/10.1109/IWSM-Mensura.2013.22>
- [11] K. Popper. 2005. *The Logic of Scientific Discovery*. Taylor & Francis.
- [12] Paul Ralph and Sebastian Baltes. 2022. Paving the Way for Mature Secondary Research: The Seven Types of Literature Review. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Singapore) (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 1632–1636. <https://doi.org/10.1145/3540250.3560877>
- [13] Richard Rudner. 1953. The Scientist Qua Scientist Makes Value Judgments. *Philosophy of Science* 20, 1 (1953), 1–6.
- [14] Ivonne Schröter, Jacob Krüger, Janet Siegmund, and Thomas Leich. 2017. Comprehending Studies on Program Comprehension. In *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*. 308–311. <https://doi.org/10.1109/ICPC.2017.9>
- [15] Janet Siegmund, Christian Kästner, Jörg Liebig, Sven Apel, and Stefan Hanenberg. 2014. Measuring and modeling programming experience. *Empirical Software Engineering* 19 (2014), 1299–1334.
- [16] Janet Siegmund and Jana Schumann. 2015. Confounding parameters on program comprehension: a literature survey. *Empirical Software Engineering* 20, 4 (2015), 1159–1192.
- [17] Janet Siegmund, Norbert Siegmund, and Sven Apel. 2015. Views on internal and external validity in empirical software engineering. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, IEEE, 9–19.
- [18] Roberto Verdecchia, Emelie Engström, Patricia Lago, Per Runeson, and Qunying Song. 2023. Threats to validity in software engineering research: A critical reflection. *Information and Software Technology* 164 (2023), 107329. <https://doi.org/10.1016/j.infsof.2023.107329>
- [19] Stefan Wagner and Marvin Wyrich. 2021. Code Comprehension Confounders: A Study of Intelligence and Personality. *IEEE Transactions on Software Engineering* 48, 12 (2021), 4789–4801.
- [20] Claes Wohlin. 2013. *An Evidence Profile for Software Engineering Research and Practice*. Springer Berlin Heidelberg, Berlin, Heidelberg, 145–157. [https://doi.org/10.1007/978-3-642-37395-4\\_10](https://doi.org/10.1007/978-3-642-37395-4_10)
- [21] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in Software Engineering*. Vol. 9783642290. Springer Berlin Heidelberg, Berlin, Heidelberg. 1–236 pages. <https://doi.org/10.1007/978-3-642-29044-2>
- [22] Marvin Wyrich. 2023. *Evidence for the design of code comprehension experiments*. Ph. D. Dissertation. University of Stuttgart.
- [23] Marvin Wyrich, Justus Bogner, and Stefan Wagner. 2023. 40 Years of Designing Code Comprehension Experiments: A Systematic Mapping Study. *ACM Comput. Surv.* (2023). <https://doi.org/10.1145/3626522>