

The Silent Scientist

When Software Research Fails to Reach Its Audience

Marvin Wyrich¹, Christof Tinnes^{1,2}, Sebastian Baltes³, and Sven Apel¹

¹ Saarland University, Saarland Informatics Campus, Germany

wyrich@cs.uni-saarland.de

² Siemens AG, Germany

³ University of Bayreuth, Germany

If software research were a performance, it would be a thoughtful theater play—full of rich content but confined to the traditional stage of academic publishing. Meanwhile, its potential audience is immersed in engaging on-demand experiences, leaving the theater half-empty, and the research findings lost in the wings. As long as this remains the case, discussions about research relevance and impact lack meaningful context.

An identity crisis

For researchers, reflecting on their own work is part of good scientific practice. This is common across all scientific disciplines, including software research.⁴ Though critical self-reflection is a staple in research, the software research community seems to take it further, grappling almost yearly with feelings of insignificance and entirely questioning the relevance of their work. In the past two years, such discussions have been nearly impossible to track. Two prominent examples illustrate this: In 2022, Lionel Briand delivered a keynote at ICSE, the largest software engineering conference, critically discussing the impact we can expect from software research.⁵ Two years later, one of the leading software journals, the Journal of Systems and Software, launched a new column inviting software practitioners to share their perspective on making software research more relevant [1]. Apparently, the software research community has reached a new peak in its ongoing quest to doubt its own relevance and impact.

To some extent, we understand the sentiment. Significant resources go into software research, to improve collaboration or drive innovation. It is reasonable to critically question how many of these research findings will be implemented

⁴ By “software research,” we refer to the broader research field that includes software engineering and related disciplines focused on the theory and practice of building and maintaining software systems. We consider this field as a case study within computer science, where much of the research is application-oriented and should therefore have a relatively receptive target audience. Similar calls to improve research communication can be found in related fields, such as in information systems, where recent discussions have highlighted the potential of citizen science [12].

⁵ Lionel C. Briand’s ICSE 2022 keynote on *the split minds of software engineering researchers*.

in the near or distant future. In fact, considerable resources have already been devoted to addressing this very question, and analyses have identified concrete examples showing how software research contributed to advancements in development tools and methods, such as configuration management and programming languages.⁶

However, in the ongoing discussions surrounding the relevance and impact of software research, there is a notable oversight: the crucial role of *science communication*. It is surprising how often this aspect is disregarded. Like a *silent scientist*, quietly publishing papers and expecting their work to speak for itself, many researchers assume their findings will naturally find their audience. This assumption can lead to an overly self-critical misconception: that if research lacks impact, it must be irrelevant. In this essay, we set out to challenge this view. Research can only have impact if it reaches its target audience in the first place, which requires making findings accessible through various communication channels [8,15]. The tricky thing about this situation is that, so far, active science communication has been sparse in the area of software research, and those who have tried often find their efforts unrewarded or unsuccessful [13].

Understanding the relevance of software research

To appreciate the relevance of software research, we need to consider how different topics resonate with different stakeholder groups—what’s impactful to one may be irrelevant to another.

First, *software research* covers diverse content, which can be roughly divided into technical advancement and (human-centered) empirical understanding. It seeks technical improvements, such as methods to automatically find and fix bugs, and empirical insights, such as understanding what factors influence the productivity of software developers. Corresponding studies of the two types of research use entirely different research methodologies, with some focusing on technical evaluations without involving humans, while others include human participants as primary subjects. This alone can create varied perceptions of relevance, as some studies directly involve the people they aim to help. Additionally, measuring impact and implementing findings differ significantly. For technological progress, impact is measured through adoption or metrics on quality and performance. Findings on practitioner collaboration require integration into sociotechnical processes, with impact measured through changes in behavior or improved satisfaction.

Second, different phases of the research process are of interest to different stakeholders. A research project can span several years and progress through various phases, at the end of which one or more publications may appear. Judging the relevance of research based on a single publication inevitably leads to a large proportion of the readership not (yet) feeling addressed. Let us assume

⁶ See, for example, *The Impact Project: Determining the Impact of Software Engineering Research Upon Practice*, as well as international workshops on *adoption-centric software engineering*.

that we are looking for a source code quality metric that indicates the comprehensibility of source code. Thanks to software research, we know that most code comprehensibility metrics do not, in practice, reflect what they are supposed to measure [9,10]. So, we are well advised not to repeat the mistake and simply design another metric based on our gut feeling. We start a little earlier, define the notion of code comprehension, conduct basic research, and understand the neuropsychological correlates of code comprehension in the brain of developers to develop a meaningful metric in a subsequent step [11]. Will that conceptual definition of code comprehension or neurophysiological lab studies with small code snippets hold actionable insights for practice? Probably not. Are these studies relevant at all? Absolutely, because software research has more than just the target group of software practitioners. In this case, these early findings may help educators better support novice programmers with code comprehension difficulties, guide experimenters to design better studies and ask more targeted questions, and they eventually aid other researchers build their application-oriented research on a solid theoretical foundation.

The point is that relevance is much more complex than is often portrayed. What is irrelevant to one person today may be very relevant for another tomorrow. We assure the reader that this is not just a convenient excuse that allows researchers to retreat into an ivory tower. We ourselves see systemic difficulties that reinforce the impression of a lack of relevance of software research. For example, the very people we want to help with our research at any given time are often involved too late in the research process. Researchers who have not spoken and, in particular, *listened* to their target group risk creating an artificially constructed problem space. Listening can take many forms, such as attending industry meetups, conducting targeted surveys, or monitoring developer forums and social media. Without these insights, it becomes difficult to argue for the relevance of one's research, and even harder to attract anyone as an audience for science communication. Another challenge is certainly that software researchers are incentivized by the scientific review and funding process to present solutions to be as generalizable as possible. As a result, we hear of software practitioners who fail to adapt published research findings because research has raised false expectations of applicability.

Note that, taking this viewpoint, we will neither succeed in fundamentally rethinking the academic system nor in convincing every reader of the relevance of every single research paper. We do not even want to. There are already enough opinions on why some software research is theoretical in nature and why other research should be more application-oriented. There is a multitude of opinions on what role industry-academia collaborations could play for the relevance of research topics. And there are certainly enough opinions on how the impact of software research could be evaluated. Our point is different: *As long as the research community does not manage to make its research accessible to the respective target group at all, we do not even need to talk about relevance and impact.*

No impact without science communication

The feeling of irrelevance and missing impact can often be traced back to a communication problem: The reality is that nobody cares about your research unless you make them care. This does not happen by itself with the publication of a research paper. The notion that software practitioners will dive into the academic world, eagerly browse online libraries for papers, and read them with sparkling eyes is as romantic as it is unlikely. A far more probable scenario for success is when researchers and practitioners build a bridge between these two worlds through actively engaging in dialogue.

The fact that software practitioners are indeed interested in contemporary software research is shown by a recent study that examines how these research findings are disseminated and discussed on LinkedIn [14]. These data reveal that the majority of individuals who post content and comment on the posts are software practitioners. The authors of that study conclude that researchers are not doing enough themselves. They note that some software research is so engaging that practitioners take on the role of science communicators—a role traditionally reserved for researchers, but inadequately filled by them in the case of software research [14].

The good news is that researchers can do a lot to improve science communication without having to revamp the publication system or redesign research processes. If we follow Burns et al. [4] in their definition of science communication “as the use of appropriate skills, media, activities, and dialogue to produce one or more of the following personal responses to science: Awareness, Enjoyment, Interest, Opinion-forming, and Understanding”, it is difficult to find a software research paper that cannot be brought to the target group with, at least, one of these intentions. Therefore, post-publication science communication is something every researcher can and should engage in.

Why isn’t it happening? We see two main reasons. First, publications are the currency of academia. Career advancement largely depends on the number of top-tier publications, while outreach and the practical impact of research are often secondary in university selection processes. Considering such systemic incentives, it is understandable that researchers quickly move on to the next project after publishing.

Second, there is a lack of evidence on the effectiveness of science communication in software research, so it may be useful to look at other sciences. Bauer et al.’s introduction to implementation science, i.e., “a science of implementation”, is only 10 years old, but has since been cited over 2000 times, many times by successful field reports [2]. In the context of clinical research, “the relatively new field of implementation science has developed to enhance the uptake of evidence-based practices and thereby increase their public health impact” [3]. Software research needs similar initiatives to learn how to bring research findings into software practice. In contrast, empirical insights on how research transfer works in software engineering and how science communication actually affects software engineering practice are still missing. As long as this gap persists, it is challenging to convince researchers that the extra effort is worthwhile.

In any case, what is hard to dispute is that research is unlikely to have much impact without *any* communication. This can take the form of a blog post, a social media discussion, a workshop or Dagstuhl seminar involving the target audience, or any other channel that fits the audience [5,6]. ACM itself offers several channels for this very purpose, such as the CACM blog and the ACM Queue. Even something as simple as submitting a ticket in an open-source project that you have studied in your research can help connect your work to those who would benefit. While the software research community may not yet know whether these efforts will lead to new collaborations, more citations, or industry-wide change, there is reason to assume that effective communication can make a difference. To provide, at least, anecdotal encouragement, we quote software researcher Marcos Kalinowski, who wrote on social media [7]: “Recently I shared the result of a PhD thesis on LinkedIn and it reached 4,000+ reactions and 270,000+ impressions. 95% of my network is from industry. We are well equipped to burst the academic bubble!”

Looking ahead, researchers are well advised to ensure their work reaches the right audience by complementing traditional, high-quality academic paper publishing with efforts to make research more accessible to relevant stakeholders. Identify your target audience to tailor your message! Use diverse communication channels beyond papers, and actively engage with practitioners to foster dialogue rather than broadcasting information! This approach can spark feedback, opening doors to new ideas and collaborations that shape future research.

If you are in software research, or any other research area for that matter, remember why you began your research journey. The *silent scientist* may publish papers, but without reaching the right audience, this work risks going unnoticed and unappreciated. We therefore urge researchers to break the silence and actively approach those communities that stand to benefit most from the findings.

Acknowledgments

We thank two reviewers and the associate editor for their constructive comments, which helped to enrich the article with additional perspectives. This work has been supported by the European Union under ERC Advanced Grant “Brains On Code” (101052182).

Author Information

Marvin Wyrich is a Postdoctoral Researcher at the Saarland University, Saarland Informatics Campus, Germany. Christof Tinnes is a Senior Key Expert at Siemens, Germany, and a PhD student at Saarland University, Germany. Sebastian Baltes is a Professor of Software Engineering at the University of Bayreuth, Germany. Sven Apel is a Professor of Software Engineering at the Saarland University, Saarland Informatics Campus, Germany.

References

1. Avgeriou, P., Shepherd, D.C.: Dear researchers - a new column sharing the perspective of software practitioners. *Journal of Systems and Software* **213**, 112044 (2024). <https://doi.org/https://doi.org/10.1016/j.jss.2024.112044>
2. Bauer, M.S., Damschroder, L., Hagedorn, H., Smith, J., Kilbourne, A.M.: An introduction to implementation science for the non-specialist. *BMC psychology* **3**, 1–12 (2015)
3. Bauer, M.S., Kirchner, J.: Implementation science: What is it and why should i care? *Psychiatry Research* **283**, 112376 (2020). <https://doi.org/https://doi.org/10.1016/j.psychres.2019.04.025>, vSI:Implementation Science
4. Burns, T.W., O'Connor, D.J., Stockmayer, S.M.: Science communication: a contemporary definition. *Public understanding of science* **12**(2), 183–202 (2003)
5. Cooke, S.J., Gallagher, A.J., Sopinka, N.M., Nguyen, V.M., Skubel, R.A., Hammerschlag, N., Boon, S., Young, N., Danylchuk, A.J.: Considerations for effective science communication (2017)
6. Illingworth, S.: Delivering effective science communication: advice from a professional science communicator. *Seminars in Cell & Developmental Biology* **70**, 10–16 (2017). <https://doi.org/https://doi.org/10.1016/j.semcdb.2017.04.002>, science communication in the field of fundamental biomedical research *Biology of aging: new models, new methods*
7. Kalinowski, M.: Tweet by Marcos Kalinowski from jun 22, 2023 (2023), https://web.archive.org/web/20230929100303/https://twitter.com/prof_kalinowski/status/1671812044146737152
8. Nisbet, M.C., Scheufele, D.A.: What's next for science communication? promising directions and lingering distractions. *American journal of botany* **96**(10), 1767–1778 (2009)
9. Peitek, N., Apel, S., Parnin, C., Brechmann, A., Siegmund, J.: Program comprehension and code complexity metrics: An fmri study. In: 43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22–30 May 2021. pp. 524–536. IEEE (2021). <https://doi.org/10.1109/ICSE43902.2021.00056>, <https://doi.org/10.1109/ICSE43902.2021.00056>
10. Scalabrino, S., Bavota, G., Vendome, C., Linares-Vasquez, M., Poshyvanyk, D., Oliveto, R.: Automatically assessing code understandability. *IEEE Transactions on Software Engineering* (2019). <https://doi.org/10.1109/TSE.2019.2901468>
11. Siegmund, J., Peitek, N., Brechmann, A., Parnin, C., Apel, S.: Studying programming in the Neuroage: Just a crazy idea? *Commun. ACM* **63**(6), 30–34 (2020)
12. Weinhardt, C., Klocker, S., Hinz, O., van der Aalst, W.M.: Citizen science in information systems research (2020)
13. Wilson, G., Aranda, J., Hoye, M., Johnson, B.: Experience report: It will never work in theory. *IEEE Software* **41**(03), 80–82 (may 2024). <https://doi.org/10.1109/MS.2024.3362649>
14. Wyrich, M., Bogner, J.: Beyond self-promotion: How software engineering research is discussed on linkedin. In: Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Society. p. 85–95. ICSE-SEIS'24, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3639475.3640113>
15. Wyrich, M., Wagner, S.: Teaching computer science students to communicate scientific findings more effectively. In: 2023 IEEE/ACM 45th International Conference

on Software Engineering: Software Engineering Education and Training (ICSE-SEET). IEEE, IEEE (2023). <https://doi.org/10.1109/ICSE-SEET58685.2023.00017>