Experience Paper: "*This Paper is Quite Unusual:*" Metascience and Structural Misalignment in the Security & Privacy Research Community

Henry Hosseini*†, Christian Böttger*, Nurullah Demir‡, Tobias Urban*

* *Institute for Internet Security*
Westphalian University of Applied Sciences
{*lastname*}*@internet-sicherheit.de*

† *Department of Information Systems*
University of Münster
henry.hosseini@uni-muenster.de

‡ *Stanford University*
nurullah@cs.stanford.edu

***Abstract*—Metascience examines the practices, evaluation mechanisms, and incentive structures that shape the production and validation of scientific knowledge. While metascientific research is increasingly institutionalized across disciplines, including parts of computer science, its presence in the Security and Privacy (S&P) community remains limited, particularly in flagship venues. This paper analyzes the current state of metascience in S&P and identifies structural factors, notably evaluation criteria that equate novelty with technical innovation, that constrain its integration into the field's core publication ecosystem. We (i) review existing S&P metascientific work and map its publication landscape, (ii) report on our experience submitting metascience studies to major S&P conferences, and (iii) examine the methodological and institutional barriers that arise when a predominantly technical research culture engages in self-reflective research using social-science methods. We argue that prevailing evaluative norms create a structural mismatch with metascientific contributions and propose pathways for integrating metascience into the S&P community.**

1. Introduction

Metascience, the systematic study of how research is conducted, evaluated, and incentivized, provides a valuable lens for understanding and improving scientific communities. By analyzing publication trends, collaboration networks, methodological norms, and evaluative standards, metascientific research helps to identify biases, inefficiencies, and opportunities for reform. In disciplines such as medicine, psychology, and part of computer science, this line of work has led to measurable improvements in reproducibility, transparency, and research quality [1], [2], [3].

In the Security and Privacy (S&P) research community, however, metascience remains underexplored. Throughout this paper, we use "S&P community" to refer to the research ecosystem organized around the major security and privacy venues, together with the researchers, reviewers, and institutional norms that sustain them. Based on the CORE ranking (renamed ICORE in 2026) [4], we compiled a list of the top venues considered as the "golden core" of the S&P community (see Table 1). This table does not represent the entire S&P community but captures its most prominent segment. Despite this community's rapid growth [5] and increasing societal impact [6], reflective studies of its own practices are rare, and top-tier S&P venues seldom recognize metascience as a valuable research contribution. This gap limits the community's ability to assess its research culture, including collaboration patterns and structural dynamics, potentially slowing progress and reinforcing existing biases [6], [7].

This paper examines the current state of metascience within S&P research and emphasizes structural barriers to its acceptance. We make the following contributions:

1) **Mapping S&P metascience.** We review published metascientific research across the S&P field and demonstrate that such work is largely absent from major S&P venues, appearing instead in workshops, bibliometric journals, or adjacent research communities (Section 2).
2) **Experiences submitting metascience to S&P venues.** We report on our own attempts to submit metascientific studies to top-tier S&P conferences, including cases where venue scope precluded submission entirely, and analyze the reviewer feedback we received (Section 3).
3) **Challenges of social-science-informed research in a technical field.** We discuss the methodological challenges and institutional barriers that arise when a predominantly technical research community attempts self-reflection using methods drawn from the social sciences, and identify pathways toward integrating metascientific approaches into the S&P community (Section 4).

2. Metascience in Computational Sciences

Having introduced metascience as a lens for understanding and improving scientific communities, we now position it conceptually within the broader scientific landscape and then examine its status within the S&P community.

2.1. Metascience Research in General

Rather than introducing new domain solutions, metascience evaluates existing scientific practices and outputs, i.e., their correctness, impact, and incentives, to help research communities adjust course. The roots of metascience lie primarily in biomedical and psychological research,

#	Venue	Abbr.	Rank
1	IEEE Security and Privacy	IEEE S&P	A*
2	Annual Computer Security Applications Conference	ACSAC	A
3	IEEE Computer Security Foundations Symposium	CSF	A
4	European Symposium on Research in Computer Security	ESORICS	A
5	USENIX Security	USENIX Sec	A*
6	Network and Distributed System Security	NDSS	A*
7	ACM Conference on Computer and Communications Security	CCS	A*
8	International Symposium on Research in Attacks, Intrusions and Defenses	RAID	A
9	Privacy Enhancing Technologies Symposium	PETS	A
10	ACM ASIA Conference on Computer and Communications Security	AsiaCCS	A
11	IEEE European Symposium on Security and Privacy	EuroS&P	A

TABLE 1: Top security and privacy venues, sorted by first appearance. The table lists the venues' current names.

where reflection on experiments and their results is essential for assessing the reliability of medical findings. One of the most widely cited definitions of metascience is provided by Ioannidis et al., who characterize it as "*the science of evaluating how to do, report, verify, correct, and reward scientific investigation*" [1]. In 2018, Ioannidis refined this definition as "*the study of research itself: its methods, reporting, reproducibility, evaluation, and incentives.*" [2].

These five pillars form the foundation of metascientific inquiry. (1) *Methods* focus on the design, execution, analysis, and interpretation of scientific studies. This pillar investigates whether the chosen methodologies are appropriate, statistically sound, and aligned with the research objectives. (2) *Reporting* refers to the clarity, transparency, and completeness of scientific communication. High-quality reporting ensures that research can be properly understood, critiqued, and extended by others. (3) *Reproducibility* addresses whether independent researchers can replicate results using the same data and methods. It emphasizes the availability of data, analysis code, and detailed documentation. (4) *Evaluation* investigates how research outputs and researchers themselves are assessed by peers, venues, funding bodies, and institutions. This includes the study of peer review processes, citation metrics, and journal rankings. (5) *Incentives* explore the structures that motivate or discourage certain research behaviors, including publication pressure, funding competition, and career advancement criteria. This pillar of metascience critically examines how these incentive systems can promote questionable research practices and proposes reforms to align incentives with research integrity.

This framework guides domain-agnostic examination of how science is conducted and communicated, emphasizing empirical investigation and systematic improvement of research practices. The computer science community lacks metascientific frameworks. Thus, researchers studying research practices in the computing field reference Ioannidis' framework rather than field-specific alternatives [8].

Within the computer science community, several initiatives and venues exist that institutionalize metascientific inquiry. Ziemann et al. [9] adopted Ioannidis's definition and its five-pillar framework to support computational reproducibility. D'Elia et al. [10] reported on their experience with reproducibility and artifact-evaluation processes in the computer science community, highlighting the variety of criteria that venues use to award artifact-evaluation badges. Dedicated venues have followed: The ACM Conference on Reproducibility and Replicability (ACM REP), launched in 2023, focuses entirely on reproducibility and replicability [11]. In 2025, TU Delft launched a workshop on metascience for machine learning (MS4ML), aiming to shift focus from chasing bold numbers to challenges such as reproducibility and reliability [12]. The Meta-HCI workshop was launched at ACM CHI '25 to provide a forum for metascientific topics beyond reproducibility in Human-Computer-Interaction (HCI), such as bias identification [13]. In the information systems domain, the AIS Transactions on Replication Research (AIS TRR) has published reproducibility research since 2014. The Americas' Conference on Information Systems (AMCIS) introduced a dedicated metascience track in 2018, encompassing topics such as field structure and boundaries, ranking methods, research culture, and the evaluation of scholarship [14]. These trends show that metascience is gaining institutional footing across multiple computer science communities. Next, we assess if similar developments have emerged in the S&P community.

2.2. Metascience Research in the S&P Community

Metascientific research in the S&P community is less common than in related disciplines, even though the few works that exist pursue retrospective analyses that map only partially onto Ioannidis's five-pillar framework. We organize existing S&P metascience research along four strands: (i) topical and framework-driven analysis, (ii) scientometric and bibliometric mapping, (iii) reproducibility and artifact evaluation, and (iv) peer-review and transparency studies. Based on these insights, we seek to identify why metascience remains peripheral to the S&P field's prestige economy.

Topical and Framework-Driven Analyses. Le Pochat and Joosen used Ioannidis's five-area framework to categorize cybersecurity meta-research [15], echoing the broader goal of improving research quality in the S&P community. Baset and Denning [5] examined authorship and topic trends in the top four cybersecurity conferences (USENIX, CCS, NDSS, and S&P) through 2015. Motivating their work, they note that "*meta-research [. . .] allows us, as a community, to examine trends in our research and make informed decisions regarding the course of our future research activities.*" Notably, this work was published in the Cyber Security Experimentation and Test (CSET) workshop, which has been held regularly since approximately 2008, making it one of the few venues that explicitly focuses on research

methodology in the S&P field. While this workshop was originally sponsored by USENIX Security [16], its topics ultimately may not have found their way into USENIX Security itself. Similar work has traced thematic evolution through co-word analysis [17] or mapped topical clusters in specific subareas such as malware detection [18].

Scientometric and Bibliometric Mapping. Several studies have mapped the S&P field's growth, international collaboration patterns, and disciplinary structure [19], [20], [21], [22]. Related work has examined authorship and citation dynamics, revealing cross-country differences in publication output [23] and the influence of factors such as paper length, venue, and reference count on citation rates across S&P subareas [24]. Much of this scientometric work has appeared in bibliometrics and interdisciplinary journals (e.g., Scientometrics [25]) rather than top-tier S&P venues, suggesting that such analyses are valued within the science-of-science community but lack a home within the S&P ecosystem.

Reproducibility and Artifact Evaluation. Concerns about the reliability of empirical results have driven an emerging trend in the S&P community toward reproducibility. Echoing the broader reproducibility crisis documented across scientific disciplines [3], several studies have analyzed the experimental practices and methods used in the S&P community. They highlighted recurring shortcomings across S&P subfields: under-documented experimental setups, missing or inaccessible datasets, and analytical choices that threaten the validity of reported findings [26], [27], [28], [29], [30]. In response, many venues now make the publication of artifacts mandatory [31], [32] or provide badges [33] when papers include a (technical) artifact that enables the experiment to be repeated, reproduced, or replicated.

Artifact evaluation (AE) has emerged as the most institutionalized metascientific mechanism in the S&P community. AE processes typically assess whether code, datasets, and experimental infrastructure are sufficiently documented and executable to reproduce reported results. By separating technical validation from the assessment of conceptual novelty, AE partially decouples epistemic verification from traditional peer review. Badging systems and reproducibility labels provide visible signals of transparency, potentially influencing author behavior and community norms. Nevertheless, AE primarily addresses the *reproducibility* pillar of metascience. It rarely evaluates the methodological appropriateness of experiments, the statistical validity of analyses, or the broader incentive structures shaping research design. Moreover, participation in AE is often optional or occurs after conditional acceptance, limiting its influence on initial publication decisions. Thus, while AE represents a significant step toward self-reflection, it does not yet constitute a comprehensive metascientific framework in the S&P community.

Peer Review and Transparency. Beyond artifact evaluation, a smaller body of work has scrutinized the S&P peer-review system and transparency practices. These studies have examined how evaluation criteria such as novelty, perceived impact, reviewer expertise, and bidding dynamics affect acceptance outcomes [7], [34], [35], [36], typically drawing on surveys, metadata, and conference-level data. Preliminary findings suggest that review processes may privilege incremental technical innovation framed as novel while disadvantaging replication studies or systemic critiques. Such dynamics are consistent with incentive-based explanations found in the broader metascience literature, in which prestige economies and publication pressure shape epistemic priorities. Additionally, scholars in ethics and governance have analyzed institutional mechanisms within S&P, including responsible-disclosure norms and community risk frameworks [37]. Systematic analyses of reviewing practices within the S&P community, however, remain limited. Key open questions include the operationalization of novelty, discrepancies in reviewer judgments, and the long-term effects of conference-centric publication models on research quality. Addressing these questions would broaden the community's metascientific agenda beyond reproducibility to include evaluation mechanisms and incentive alignment, thereby more comprehensively engaging with Ioannidis's five-pillar framework.

Structural Marginality of S&P Metascience. Despite this growing body of work, metascience remains peripheral within the S&P field. The work reviewed above is conducted by a relatively small subset of researchers embedded within the core community rather than by a distinct subfield. This "bubble" consists of insiders reflecting on their own field, and it remains fragmented across the strands described above rather than forming a consolidated research area.

Measured against Ioannidis's five pillars, the S&P community has primarily focused on reproducibility [38] but has largely neglected reflection on methodological validity, transparency in scientific communication, peer and venue evaluation, and the alignment of incentives with scientific rigor. One structural explanation lies in the field's evaluation criteria: novelty is consistently emphasized as a primary review criterion [7], yet it lacks a precise operational definition. Calls for papers in top S&P venues emphasize and prioritize "novel research contributions," typically interpreted as technical innovation (see also Section 3). Metascientific research, by contrast, centers on replication, methodological critique, or process evaluation, which creates a structural mismatch between what metascience produces and what top-tier S&P venues reward.

Overall, although metascientific reflection is gaining visibility within the S&P community, its relatively limited presence in top-tier venues appears to stem less from an absolute absence of such work than from prevailing incentive structures and evaluative norms. Of the works cited in this section, only 4 reproducibility studies and 3 review studies were published in the venues listed in Table 1. These norms decouple metascientific contributions from the prestige economy associated with primary technical research.

3. Our (Limited) Experience Trying to Publish Metascience at Top-tier S&P Venues

This section describes the authors' experience attempting to publish a large-scale metascience study [39] at top-tier S&P venues. We approach this analysis from a constructivist stance [40], [41], acknowledging that our perspective as authors is situated and that program chairs and reviewers may hold different, yet valid, conceptions of what constitutes relevance, rigor, and contribution. Accordingly, the following is an interpretive account of how technically oriented S&P venues approached and evaluated metascientific work; it is not intended as an assessment of the correctness of outcomes or as a basis for generalizable claims. Rather than critiquing specific conferences or reviewers, this section serves as an experience report on lessons learned from conducting and attempting to publish metascience research within the S&P community.

On the Paper Discussed in this Section. In this section, we discuss our experiences in submitting the paper *Four Decades of Security and Privacy Research: Evolution of Topics, Impact, and the Community* [39] to different major S&P venues. Our paper analyzes the evolution of IT security and privacy research from 1980 to 2023 using a dataset of over 13,000 papers and 500,000 authors. It applies social science theories and natural language processing to build a taxonomy of topics, showing the field's growth from just 5 topics to around 100, reflecting increasing topical diversity. The study also highlights strong collaboration networks within the community and finds that female representation has remained relatively stable despite increased participation. Additionally, it examines factors influencing citations, author relationships, and language use over time. Overall, the paper provides insights into how the field has developed and where it may be headed. The paper was eventually accepted for publication in the *Journal of Cybersecurity*, an Oxford University Press Q1-ranked journal [42]. It should be noted that the accepted version contained only minor adjustments compared to the version submitted to other venues. In particular, it did not include any new analyses or noteworthy methodological adjustments.

Identifying a Suitable Venue. The first step in publishing a paper is to identify a suitable venue (e.g., a conference or a journal). In the (Anglo-American) S&P community, which is predominantly conference-driven, top-ranked (i.e., $\mathtt{A}^*$ and $\mathtt{A}$) conferences are a natural starting point (see Table 1).

We manually inspected the Calls for Papers (CfP) of these eleven conferences between 2023 and 2026, covering a total of 33 individual calls for papers.[1] When manually reviewing the calls for papers, we focused on the entire section(s) on topics of interest and looked for terms that are related to metascience (e.g., scientometrics, science-of-science, (self-)reflection, scientific methodology, or meta-research). Of these 33 calls, only one ($1/33 \approx 3\%$) listed metascience as a topic of interest: the USENIX Security '25 call for papers. This addition was significant, as this was the first tier-one security conference to include metascience as a submission topic alongside "Replication and Reproduction."

Most calls for papers are non-exhaustive, allowing submissions on additional relevant topics (e.g., the call for papers of the 2026 IEEE Symposium on Security and Privacy (S&P) states "*This topic list is not meant to be exhaustive; S&P is interested in all aspects of computer security and privacy.*" [43]). Thus, in theory, a metascience paper could be submitted to these venues. While seeking a venue for our metascientific work, we contacted the program chairs of three conferences and one major journal that were accepting submissions at the time. The conferences and journals either did not respond to our inquiry or stated that metascience studies would not fall within the scope of their respective venues. One program chair's response suggested a conflation of metascience with systematization of knowledge: "*The paper that you propose does not seem to fit under any of the topics of interest of* [CONFERENCE NAME]*, and might be better suited for a venue that accepts* ***SoK submissions***."

These observations align with the patterns identified in Section 2.2: None of the existing S&P metascience works identified in our review was published at a top-tier S&P conference. Instead, such work appears in workshops, scientometric and bibliometric journals, or adjacent communities that are often not closely followed by the S&P research community. As a result, the community may be unaware of this stream of research and may not recognize the scientific value of retrospective analyses of its own practices. Overall, both our analysis of calls for papers and our submission experience provide additional indications that metascience is treated as peripheral to the S&P research community.

S&P Reviewers' Perspectives on Metascience Work. We now describe the reviews we received after submitting a metascience study to two top-tier S&P conferences, resulting in a total of 6 reviews. A recurring theme was the expectation that research must propose a technical mechanism. As one reviewer noted: "[...] *this is not a research paper because it does not propose any new technique to address certain technical challenges [...]*". This stance is notable given that HCI research, which also does not propose technical mechanisms, is an established part of the community (e.g., all venues named in Table 1 accept such submissions).

Some reviewers were positive about this "new" direction for S&P research (e.g., "*I found this paper quite unusual but also very interesting [...]*" and "*Overall, I enjoyed the premise of this paper, as it was something different than what we traditionally see at S&P venues.*"). Others, however, questioned whether the work constituted research at all (e.g., "[...] *this is neither a research paper nor an SoK paper* [...]"). Reviewers also questioned the significance of non-technical findings (e.g., "[...]*, which finally leads to some interesting but not technical observations,* [...]") or expressed confidence that identified issues in the S&P community would resolve themselves (e.g., "[...] *there is no doubt that we will continuously improve the community's*

1. Many of these conferences issue multiple calls per year due to rolling or multi-deadline submission cycles.

diversity."). These responses highlight differing conceptions of what constitutes a research contribution and what types of insights social-science methods can offer the community.

Overall, the received reviews highlight two main takeaways: (1) reviewers at top-ranked S&P conferences tend to assume that valuable research must be technical in nature, and (2) reviewers may not view the S&P community itself as a meaningful object of study, treating self-reflective analyses as outside the scope of S&P research. These observations raise the question of how metascience, as an interdisciplinary social aspect, can be integrated into the community.

4. Discussion: Metascience in the S&P World

Our preceding analysis suggests that metascience occupies a peripheral position within top-tier S&P venues. This observation is supported by three converging signals: the near-complete absence of metascience in calls for papers (Section 2.2), explicit statements by program chairs that such work is out of scope, and reviewer assessments that systematically privilege technical novelty over reflective or analytical contributions (Section 3). Having made these observations, we examine the underlying dynamics and their consequences, and propose potential pathways for change.

Technical Evaluation as Gatekeeping. A key reason for this marginalization lies in the S&P community's evaluative norms, which are designed for technical contributions. Top-tier venues have prioritized new systems, attacks, defenses, measurements, or formal techniques. Within this framing, research that does not introduce a novel technical mechanism risks being perceived as lacking rigor or relevance, regardless of its analytical depth. While non-technical subfields such as HCI are established within S&P, they are often justified through their direct interaction with technical systems (e.g., user studies of authentication mechanisms). Metascience, by contrast, treats the research community itself, i.e., its practices, incentives, and outputs, as the object of study. This alternative focus challenges prevailing notions of what constitutes a scholarly research contribution and creates a category that existing evaluative frameworks are not equipped to accommodate.

The Reviewer Expertise Gap. Metascience is inherently interdisciplinary, drawing on methods and theories from the social sciences, scientometrics, bibliometrics, and qualitative research. In a community whose reviewer pool is predominantly trained in technical evaluation, familiarity with metascientific methodologies, validity criteria, and epistemic goals may be limited. As a result, reviewers may default to technical evaluation criteria that are misaligned with the contribution type, assessing whether a study proposes a "solution" rather than whether its analysis yields useful insights about the field. This mismatch highlights the absence of shared community norms and review guidelines for evaluating intradisciplinary research. Other communities have addressed similar challenges by developing dedicated review criteria based on contribution type. For example, ACM CHI maintains distinct evaluation criteria for different contribution types, such as empirical and replication studies, as well as methodological and theoretical work [44]. Such models could be adapted for S&P venues to help reviewers engage with metascientific work on its own terms.

Consequences of Excluding Metascience. By systematically excluding or discouraging metascientific inquiry, the S&P community risks limiting its capacity for self-reflection and evidence-based reform. Metascience can provide insights into publication bias, evaluation practices, diversity and inclusion, incentive structures, and the long-term evolution of research agendas. Without such perspectives, discussions about community health, fairness, and scientific progress may rely on anecdote or intuition rather than empirical evidence. Moreover, the fact that related metascientific work is predominantly published in journals outside the S&P ecosystem suggests that relevant findings may go unnoticed by the very community they analyze. This separation reduces knowledge transfer and reinforces the perception that metascience lacks direct value for S&P research.

Technical Innovation vs. Social Innovation. The observed resistance illustrates a common challenge of intradisciplinary research: Studying a community from within can be perceived as less valuable than producing artifacts for that community. When a field does not recognize itself as a legitimate object of scientific inquiry, opportunities for social and organizational innovation are constrained. This may contribute to stagnation, as entrenched practices remain unexamined and structural issues persist unaddressed. In contrast, other scientific domains have demonstrated that metascience can play a vital role in improving research quality, transparency, and inclusivity.

5. Conclusion and Recommendations

This paper examines the current state of metascientific inquiry within the S&P research community and our own experience attempting to publish metascience in top-tier S&P venues. By analyzing strands of published metascientific work and calls for papers, reflecting on our submission experiences, and interpreting reviewer feedback, we identified a structural misalignment between the epistemic goals of metascience and the dominant evaluation criteria of leading S&P venues. While metascientific reflection exists within the broader ecosystem of S&P research, it remains institutionally peripheral and largely decoupled from the prestige structures that shape research agendas.

Integrating metascientific approaches into mainstream S&P venues is not a departure from rigor but an expansion of it. The following recommendations outline possible steps toward reducing the identified structural barriers:

1) **Venue-level reforms.** Top-tier S&P venues could explicitly recognize metascience (science-of-science or community studies) as a legitimate submission type, e.g., via a dedicated track or as an explicitly welcome category within existing tracks.

2) **Reviewer and process adjustments.** Since metascience draws on methods from scientometrics, bibliometrics, and qualitative research, reviewer pools and review forms may need to include relevant methodological expertise and criteria aligned with non-technical contributions.
3) **Community recognition.** S&P research routinely investigates systems, adversaries, infrastructures, and users. Extending this analytical lens to the S&P community itself is a progression rather than a conceptual departure.

Ultimately, the long-term vitality of the S&P research community depends not only on technical breakthroughs but also on the quality of its evaluative and institutional foundations. Integrating metascientific perspectives into main S&P discourse could strengthen evidence-based reform, improve transparency, and sustain the field's credibility.

Acknowledgments

The authors would like to thank the anonymous reviewers for their valuable feedback. This work was supported by the "Daimler und Benz Stiftung" grant: Ladenburger Kolleg (Project: KonCheck), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - FOR 5393, Project No. 462287308 (BE 1422/27-1), and Google's Educational Cloud funding (grant number: EDU Credit 426550496).

References

[1] J. P. A. Ioannidis, D. Fanelli, D. D. Dunne, and S. N. Goodman, "Meta-research: Evaluation and Improvement of Research Methods and Practices," *PLOS Biology*, vol. 13, no. 10, pp. 1–7, 10 2015. [Online]. Available: https://doi.org/10.1371/journal.pbio.1002264

[2] J. P. A. Ioannidis, "Meta-research: Why research on research matters," *PLOS Biology*, vol. 16, no. 3, pp. 1–6, 03 2018. [Online]. Available: https://doi.org/10.1371/journal.pbio.2005468

[3] National Academies of Sciences, Engineering, and Medicine, *Reproducibility and Replicability in Science*. Washington, DC: The National Academies Press, 2019.

[4] CORE Ranking, "ICORE Conference Portal," https://portal.core.edu.au/conf-ranks/?search=4604&by=for&source=CORE2023&sort=arank&page=1, 2026, Accessed: 2026-02-08.

[5] A. Baset and T. Denning, "A Data-Driven Reflection on 36 Years of Security and Privacy Research," in *Proceedings of the 12th USENIX Conference on Cyber Security Experimentation and Test*, ser. CSET'19. USA: USENIX Association, 2019, p. 6.

[6] J. M. Spring, T. Moore, and D. J. Pym, "Practicing a Science of Security: A Philosophy of Science Perspective," in *Proceedings of the 2017 New Security Paradigms Workshop (NSPW '17)*, 2017, pp. 1–18.

[7] A. Soneji, F. B. Kokulu, C. Rubio-Medrano, T. Bao, R. Wang, Y. Shoshitaishvili, and A. Doupe, ""Flawed, but like democracy we don't have a better system": The Experts' Insights on the Peer Review Process of Evaluating Security Papers," in *Proceedings of the 43rd IEEE Symposium on Security and Privacy (SP)*, 2022, pp. 1845–1862.

[8] N. Demir, "Better web measurements: Enhancing security and privacy research," Ph.D. dissertation, Karlsruher Institut für Technologie (KIT), 2023.

[9] M. Ziemann, P. Poulain, and A. Bora, "The five pillars of computational reproducibility: bioinformatics and beyond," *Briefings in Bioinformatics*, vol. 24, no. 6, p. bbad375, 10 2023. [Online]. Available: https://doi.org/10.1093/bib/bbad375

[10] D. C. D'Elia, T. D. Doudali, C. Giuffrida, M. Matos, M. Payer, S. Pirelli, G. Portokalidis, V. Schiavoni, S. Signorello, and A. Vahldiek-Oberwagner, "Lessons Learned from Five Years of Artifact Evaluations at EuroSys," in *Proceedings of the 3rd ACM Conference on Reproducibility and Replicability*, 2025, pp. 108–120.

[11] ACM Conference on Reproducibility and Replicability (ACM REP), "ACM REP: ACM Conference on Reproducibility and Replicability," https://acm-rep.github.io/, 2023, Accessed: 2026-02-10.

[12] H. Hung, J. van Gemert, and M. Loog, "Metascience for Machine Learning (MS4ML)," https://www.tudelft.nl/ellis-delft-unit/programmes/metascience-for-machine-learning, 2025.

[13] J. Oppenlaender, S. Malacria, X. Fang, N. van Berkel, F. Chevalier, K. Yatani, and S. Hosio, "Meta-HCI: First Workshop on Meta-Research in HCI," in *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '25. New York, NY, USA: Association for Computing Machinery, 2025. [Online]. Available: https://doi.org/10.1145/3706599.3706723

[14] *Meta-Research in Information Systems*. Cancun, Mexico: Association for Information Systems, 2019, AMCIS 2019 Proceedings Track: Meta-Research in Information Systems. [Online]. Available: https://aisel.aisnet.org/amcis2019/meta_research_is/

[15] V. Le Pochat and W. Joosen, "Analyzing Cyber Security Research Practices through a Meta-Research Framework," in *Proceedings of the 16th Cyber Security Experimentation and Test Workshop*, ser. CSET '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 64–74. [Online]. Available: https://doi.org/10.1145/3607505.3607523

[16] USC Information Sciences Institute (USC-ISI), "CSET'24: 17th Cyber Security Experimentation and Test Workshop," https://cset24.isi.edu/, 2024, Workshop held August 13, 2024, in conjunction with the USENIX Security Symposium, Philadelphia Marriott Downtown, Philadelphia, PA, USA. Accessed: 2026-02-01.

[17] W. H. Lee, "How to identify emerging research fields using scientometrics: An example in the field of Information Security," *Scientometrics*, vol. 76, no. 3, pp. 503–525, sep 2008. [Online]. Available: https://doi.org/10.1007/s11192-007-1898-2

[18] F. Alqurashi and I. Ahmad, "Scientometric Analysis and Knowledge Mapping of Cybersecurity," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 3, 2024. [Online]. Available: http://dx.doi.org/10.14569/IJACSA.2024.01503117

[19] F. A. Loan, B. Bisma, and N. Nahida, "Global research productivity in cybersecurity: a scientometric study," *Global Knowledge, Memory and Communication*, vol. 71, no. 4/5, pp. 342–354, 2022.

[20] B. G. S.M. Dhawan and B. Elango, "Global Cyber Security Research Output (1998–2019): A Scientometric Analysis," *Science & Technology Libraries*, vol. 40, no. 2, pp. 172–189, 2021. [Online]. Available: https://doi.org/10.1080/0194262X.2020.1840487

[21] M. Lang, "The Fragmented Research Space of Cybersecurity: A Map of the Territory," in *Cybersecurity*, I. Praça, S. Bernardi, and P. R. Inácio, Eds. Cham: Springer Nature Switzerland, 2025, pp. 116–132.

[22] M. Koca and S. Çiftçi, "A comprehensive bibliometric analysis of Big Data and Cyber Security: intellectual structure, trends, and global collaborations," *Knowledge and Information Systems*, vol. 67, no. 11, pp. 10 245–10 270, jul 2025. [Online]. Available: https://doi.org/10.1007/s10115-025-02531-1

[23] N. V. Olijnyk, "A quantitative examination of the intellectual profile and evolution of information security from 1965 to 2015," *Scientometrics*, vol. 105, no. 2, pp. 883–904, Nov 2015. [Online]. Available: https://doi.org/10.1007/s11192-015-1708-1

[24] S. Wendzel, C. Lévy-Bencheton, and L. Caviglione, "Not all areas are equal: analysis of citations in information security research," *Scientometrics*, vol. 122, no. 1, pp. 267–286, Jan 2020. [Online]. Available: https://doi.org/10.1007/s11192-019-03279-6

[25] S. Nature, "Scientometrics – An International Journal for all Quantitative Aspects of the Science of Science, Communication in Science and Science Policy," https://link.springer.com/journal/11192, 2026, Accessed: 2026-02-08.

[26] C. Rossow, C. J. Dietrich, C. Grier, C. Kreibich, V. Paxson, N. Pohlmann, H. Bos, and M. v. Steen, "Prudent Practices for Designing Malware Experiments: Status Quo and Outlook," in *2012 IEEE Symposium on Security and Privacy*, ser. SP '12. USA: IEEE Computer Society, 2012, pp. 65–79. [Online]. Available: https://doi.org/10.1109/SP.2012.14

[27] H. Hosseini, M. Degeling, C. Utz, and T. Hupperich, "Unifying Privacy Policy Detection," *Proceedings on Privacy Enhancing Technologies*, vol. 2021, no. 4, pp. 480–499, Jul. 2021. [Online]. Available: https://doi.org/10.2478/popets-2021-0081

[28] N. Demir, M. Große-Kampmann, T. Urban, C. Wressnegger, T. Holz, and N. Pohlmann, "Reproducibility and Replicability of Web Measurement Studies," in *Proceedings of the ACM Web Conference 2022*, ser. WWW '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 533–544. [Online]. Available: https://doi.org/10.1145/3485447.3512214

[29] D. Arp, E. Quiring, F. Pendlebury, A. Warnecke, F. Pierazzi, C. Wressnegger, L. Cavallaro, and K. Rieck, "Dos and Don'ts of Machine Learning in Computer Security," in *Proceedings of the 31st USENIX Security Symposium : August 10-12, 2022, Boston, MA, USA*. USENIX Association, 2022, pp. 3971–3988, 46.23.01; LK 01. [Online]. Available: https://www.usenix.org/system/files/sec22summer_arp.pdf

[30] J. Evertz, N. Risse, N. Neuer, A. Müller, P. Normann, G. Sapia, S. Gupta, D. Pape, S. Shaw, D. Srivastav, C. Wressnegger, E. Quiring, T. Eisenhofer, D. Arp, and L. Schönherr, "Chasing Shadows: Pitfalls in LLM Security Research," in *Proceedings of the 2026 Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2026.

[31] ACM Conference on Computer and Communications Security (CCS), "Call for papers 2026," https://www.sigsac.org/ccs/CCS2026/call-for/call-for-papers.html, 2026, Accessed: 2026-02-08.

[32] USENIX Security Symposium, "USENIX Security '26 Call for Papers," https://www.usenix.org/conference/usenixsecurity26/call-for-papers, 2026, Accessed: 2026-02-08.

[33] Privacy Enhancing Technologies Symposium, "PoPETs 2026 Artifact Evaluation," https://petsymposium.org/artifacts.php#artifact-badges/, 2026, Accessed: 2026-02-08.

[34] M. Reinhart and C. Schendzielorz, "Peer-review procedures as practice, decision, and governance—the road to theories of peer review," *Science and Public Policy*, vol. 51, no. 3, pp. 543–552, 02 2024. [Online]. Available: https://doi.org/10.1093/scipol/scad089

[35] J. H. Klemmer, J. Schmüser, F. Fischer, J. Suray, J.-U. Holtgrave, S. Lenau, B. M. Lowens, F. Schaut, and S. Fahl, "How transparent is usable privacy and security research? a meta-study on current research transparency practices," in *Proceedings of the 34th USENIX Conference on Security Symposium*, ser. SEC '25. USA: USENIX Association, 2025. [Online]. Available: https://www.usenix.org/system/files/usenixsecurity25-klemmer.pdf

[36] S. Jung, G. Pyeon, I. Heo, and H. Ahn, "What Drives Paper Acceptance? A Process-Centric Analysis of Modern Peer Review," https://arxiv.org/abs/2509.25701, 2026.

[37] K. Macnish and J. van der Ham, "Ethics in cybersecurity research and practice," *Technology in Society*, vol. 63, p. 101382, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0160791X19306840

[38] "Security Research Artifacts," https://secartifacts.github.io/, 2022, Website for research artifacts from the security community, collecting resources and results around artifact evaluation for security conferences and workshops. Accessed: 2026-02-05.

[39] N. Demir, C. Böttger, H. Hosseini, M. Demir, C. Wressnegger, N. Pohlmann, and T. Urban, "Four decades of security and privacy research: evolution of topics, impact, and the community," *Journal of Cybersecurity*, vol. 12, no. 1, p. tyag007, 04 2026. [Online]. Available: https://doi.org/10.1093/cybsec/tyag007

[40] T. A. Schwandt, "Constructivism," in *The SAGE Dictionary of Qualitative Inquiry*, 3rd ed. Thousand Oaks, CA: SAGE Publications, Inc., 2007, pp. 38–41.

[41] T. E. Costantino, C. Thompson, and S. Bales, "Constructivism," in *The SAGE Encyclopedia of Qualitative Research Methods*. Thousand Oaks: SAGE Publications, Inc., 2008.

[42] Oxford University Press, "Journal of Cybersecurity," https://academic.oup.com/cybersecurity/, 2026, Accessed: 2026-03-23.

[43] I. S. on Security and Privacy, "47th IEEE Symposium on Security and Privacy," https://sp2026.ieee-security.org/cfpapers.html, 2026, Accessed: 2026-02-08.

[44] ACM CHI Conference on Human Factors in Computing Systems, "Contributions to CHI – CHI 2025," https://chi2025.acm.org/contributions-to-chi/, 2025, Accessed: 2026-02-10.