RESEARCH-ARTICLE

Ethics in Computer Security Research: A Data-Driven Assessment of the Past, the Present, and the Possible Future

HARSHINI SRI RAMULU, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

HELEN SCHMITT, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

BOGDAN RERICH, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

RACHEL GONZALEZ RODRIGUEZ, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

TADAYOSHI KOHNO, Georgetown University, Washington, D.C., United States

YASEMIN ACAR, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

Open Access Support provided by:

Paderborn University

Georgetown University

Published: 19 November 2025

Citation in BibTeX format

CCS '25: ACM SIGSAC Conference on Computer and Communications Security
October 13 - 17, 2025
Taipei, Taiwan

Conference Sponsors:
SIGSAC

https://doi.org/10.1145/3719027.3765071
ISBN: 9798400715259

Ethics in Computer Security Research: A Data-Driven Assessment of the Past, the Present, and the Possible Future

Harshini Sri Ramulu
Paderborn University
Paderborn, Germany
harshini.sri.ramulu@uni-paderborn.de

Helen Schmitt
Paderborn University
Paderborn, Germany
helen.schmitt@uni-paderborn.de

Bogdan Rerich
Paderborn University
Paderborn, Germany
bogdan.rerich@uni-paderborn.de

Rachel Gonzalez Rodriguez
Paderborn University
Paderborn, Germany
rachel.gonzalez.rodriguez@uni-paderborn.de

Tadayoshi Kohno
Georgetown University
Washington, D.C., USA
yoshi.kohno@georgetown.edu

Yasemin Acar
Paderborn University & The George Washington University
Paderborn, Germany
yasemin.acar@uni-paderborn.de

Abstract

Ethical questions are discussed regularly in computer security. Still, researchers in computer security lack clear guidance on how to make, document, and assess ethical decisions in research when what is morally right or acceptable is not clear-cut. In this work, we give an overview of the discussion of ethical implications in current published work in computer security by reviewing all 1154 publications at top 4 security conferences published in 2024, finding inconsistent levels of ethics reporting with a strong focus of reporting institutional or ethics board approval, human subjects protection, and responsible disclosure, and a lack of discussion of balancing harms and benefits. We further report on the results of a semi-structured interview study with 24 computer security and privacy researchers (among whom were also: reviewers, ethics committee members, and/or program chairs) and their ethical decision-making both as authors and during peer review, finding a strong desire for ethical research, but a lack of consistency in considered values, ethical frameworks (if articulated), decision-making, and outcomes. We present an overview of the current state of the discussion of ethics and current de-facto standards in computer security research, contributing suggestions to improve the state of ethics in computer security research.

CCS Concepts

• **Security and privacy → Human and societal aspects of security and privacy**.

Keywords

security ethics; usable security and privacy; research ethics

ACM Reference Format:
Harshini Sri Ramulu, Helen Schmitt, Bogdan Rerich, Rachel Gonzalez Rodriguez, Tadayoshi Kohno, and Yasemin Acar. 2025. Ethics in Computer Security Research: A Data-Driven Assessment of the Past, the Present, and the Possible Future. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25), October 13–17, 2025, Taipei.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3719027.3765071

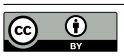

CCS '25, Taipei

ACM ISBN 979-8-4007-1525-9/2025/10
https://doi.org/10.1145/3719027.3765071

1 Introduction

Ethical reasoning in computer security has been discussed for decades [4, 18, 45, 64], and has recently, received increased attention, arguably triggered at least in part by a widely discussed study on the Linux Kernel accepted for publication, then withdrawn, from the IEEE Symposium on Security and Privacy (IEEE S&P) [55]. Here, researchers experimented on the Linux Kernel without developers' consent nor the upfront understanding that they were conducting human subjects research. Responses included the creation of ethics review committees within high-ranking S&P venues [12, 40] as well as broader discussions [68], also through workshops[1], changes in calls for papers[2] and peer reviewing fields (where flagging a submission for deeper ethics review became possible), and research targeting ethical decision-making in security [31, 43].

As evidenced by recurring discussions during peer review, in ethics committees, technical committee meetings, when creating calls for papers, and public discussions, the security research community is actively refining its understanding of the ethical challenges within the field. However, there are still no community wide-accepted standards for tackling ethical challenges in security research. The (now-archived) Menlo Report is often referred to as the de facto reference for ethical guidance, as also noted by ethical guidelines in security conferences' call for papers. Recently, the Menlo Report has been marked as archived by the Department of Homeland Security (DHS) in the US, without a replacement or a more current reference [74]. At the same time, academic research studies with ethically ambiguous methods are receiving increased attention, for example for conducting research on a Reddit forum without participant consent [53] or for sending deceptive legal requests to companies to understand their response [33].

There is, however, no current assessment of the status quo and best practices across the research field as a whole, and written

[1] https://www.ndss-symposium.org/ndss-program/ethics-2023/
[2] Since 2022, the IEEE S&P Call for Papers includes a Research Ethics Committee and encourages authors to review the Menlo Report; USENIX Security 2025 requires discussing the ethics of submitted research, and explicitly formulates an ethics guidelines document, see https://www.usenix.org/conference/usenixsecurity25/ethics-guidelines.

ethics sections published in research papers omit many of the ethical discussions authors may have had. As our research findings confirm, this situation makes it hard to establish a community standard based on written artifacts and also makes it hard for researchers newer to the field or branching out into different subdisciplines to understand the full spectrum of ethical considerations and current conventions.

In this work, we explore the de-facto status of ethical reasoning and reporting in computer security research. We assess the status quo in reporting ethical decisions, practices, and precautions discussed in research papers through a meta-analysis of 1154 published security research papers, and discuss processes, challenges, best practices, and possible improvements to handling ethical decision-making and reporting through semi-structured interviews with 24 security researchers. We answer the research questions:

RQ1: *How are ethical decisions reported in computer security research? What are gaps in reporting?*

RQ2: *How does the field of computer security research reason about ethics? When do researchers reason about ethics, what factors are considered, how are decisions made, and why are these approaches used?*

RQ3: *What are challenges to ethical reasoning in computer security research and the broader community, and what can be improved?*

RQ4: *How does ethical decision-making in computer security research interact with peer review? How does the community reason and decide about ethics in their roles as reviewers, ethics committee members? How should the security community handle work that reviewers consider ethically questionable?*

We hope that this work can serve as a foundation for the next stage in the evolution of ethics considerations within the computer security research field. While there may *not* be an unambiguously right ethical decision in all situations, as prior work demonstrated [4, 31, 43], it is our belief that an informed understanding of how the field has approached ethics in the past *can* provide a foundation for future, informed decision in challenging ethical situations.

2 Context: Background & Related Work

We discuss the development of ethics in the area of computing research and cover the related work on limitations and current work on supporting the consideration of ethics and decision-making in computer security research.

2.1 History of ethics and ethical frameworks in computing research

Published guidelines concerning research ethics have their roots in human subject protection and medical and biomedical research, as early as the 1970s following unethical research like the Tuskegee syphilis study [18, 63, 66]. Over time, ethical concepts were also adapted to apply to computing research with computer science communities devising codes of conduct in computing research, such as the 1974 IEEE Code of Ethics (updated 2006) [41, 56] and Association for Computing Machinery (ACM) code of conduct in 2018 [3]. These codes are supposed to serve as a basis for ethical decision-making and foster better ethical standards among computer science professionals. Key attributes include the contribution to humans and society, acknowledging stakeholders, avoiding harm, honesty and the respect of work, privacy, and confidentiality [3, 41].

Another relevant guideline for technological research is the 2011 Menlo Report, making the older Belmont report's guidelines for biomedical research applicable for computer and information security research [4]. The Menlo Report itself envisions ethics as a form of governance and was shaped by relying on bricolage work with available resources on past controversies in computing research [24]. In the companion to the Menlo Report, Bailey et al. argue that their re-conceptualization is needed because of the nature of information technology research; for example, it has a much larger scale and higher speed than traditional medical research. The research has the potential to harm humans, even if they are not a research entity [18]. While regulations are important to ethical research, discussions in research communities are also very important, as ethical and moral understanding is refined through discussion and conversations [9, 16].

While the ACM and IEEE Codes of Ethics aim to aid researchers and computer science professionals with ethical thinking in professional settings, McNamara et al. found that these frameworks do not significantly impact on the decision-making of professionals [50]. The security and privacy community currently has differences in opinions on ethical considerations, and no guidelines have been widely adopted [46, 68]. Further, there is criticism on the granularity and comprehensibility, and therefore, usefulness of these guidelines [51, 57, 78]. Principlist ethical frameworks, i.e., frameworks built on defined ethical principles, have been proposed, for example, with the five principles of beneficence, non-maleficence, autonomy, justice, and explicability [26]. Complementing abstract principlist frameworks, prior work provides guidelines, values, examples, and case studies for ethical challenges faced in cybersecurity in practice [15, 80]. While these works have been cited, it is unclear whether these tools are widely used in computer security research, though USENIX security has recently begun including ethics guidelines in the call for papers [75].

New development in ethical frameworks and tools for researchers and review boards. In an effort to support researchers in considering ethical issues and the implications of their work, various studies have used tools, frameworks, and resources. This includes a self-questionnaire, which is also to be used by Institutional Research Boards (IRBs) and Ethics Review Boards (ERB)[3] to address cybersecurity issues [61]. Additionally, Kohno et al. seek to support ethical discussion, moral reasoning and decision-making in computer security research through tools and insights from ethics and moral philosophy [43]. They offer moral dilemma scenarios

[3]In the US, approval by an Institutional Review Board (IRB) is mandatory for all government-funded human subjects research across research areas, and generally required by academic research institutions. This is not true for research not involving human subjects, even though there may be ethical implications. IRBs may not exist for authors in countries outside the United States. Similar boards may instead be called Ethics Reviews Boards (ERBs), and they may not be offered, applicable, or meaningful for computer security research. Outside the US, ethics review may be at the discretion of researchers, or may become mandatory when research is tied to a certain type of funding. We will write of "review boards" when we mean both IRBs and ERBs. We distinguish these from Research Ethics Committees (RECs) introduced in 2023 at IEEE S&P; these conduct ethics reviews of fully written research papers during peer review.

and analyze them following consequentialist[4] and deontological[5] frameworks. Focusing mainly on the Chinese security research community, Zhang et al. explored the ethical requirements of the most important computer security conferences and to what extent this was followed by the research community over the last few years; they find an increase in ethics discussions and urge researchers to follow Menlo Report principles, detail ethical considerations in research papers, and avoid ethically ambiguous work [82]. Similarly, prior work has also created resources for researchers and technologists to anticipate the broader ethical and societal impacts of their work like tarot cards for ethics [34], value cards [67], value sensitive design [27], real world case studies for researchers to consider [54], anticipatory ethics for new technologies [6], and several other methods and activities to anticipate ethical issues [1, 14, 21, 42, 44, 47]. Further, Reijers et al. provide recommendations for considering relevant stakeholders and ethical impacts of future technologies by reviewing published research on ethical practices of past innovations [62]. However, these tools are either not aimed at researchers [30], or they are not broadly used in practice [54, 59].

Research indicates that in computer security, ethical decision-making must consider nuances and complexities [25] and there is a lack on common agreement on what is 'ethical' in some cases [8]. Lack of organizational support, incentives, and personal precarity can impede raising and resolving ethical concerns [79].

Prior work shows that computer science researchers and professionals rarely consider potential unintended consequences of their innovations [19], due to lack of knowledge, awareness, and considering ethics as an afterthought or just as a compliance [59, 76], both within and beyond computer security and privacy. Studies point towards a lack of ethics education for computer security students at universities as one of the reasons for the missing ethical expertise in the community [46]. Additionally, prior work highlights a lack of consistency in ethics curricula for technologists [23] and advocates for more integrated ethics education for computer science students to better equip them to make ethical decisions in their work [49] and while working on 'real world' technologies [73].

2.2 Ethics in security and privacy research

Recent ethically questionable research in computer security. In the recent past, ethically questionable works have intensified ethics discussions within the security and privacy research community. In the much-discussed hypocrite commits paper, researchers wanted to test the feasibility of introducing vulnerabilities into open source projects. While the study included an experiment on a live repository, the authors were unaware that they were enrolling participants without consent, and did not obtain IRB approval [13, 68]. This study arguably violated multiple ethical principles [55]. However, when the researchers were advised to contact their IRB during peer review, their IRB deemed this research non-human-subjects. This demonstrated that relying solely on an IRB determination for ethical decision-making in security and privacy research may be insufficient [55]. More broadly, security research often involves using sensitive data, e.g., working with stolen data from illicit markets and data breaches [48], and studying censorship [22], public data [65], and vulnerable populations [69, 77]; such studies may need careful ethical considerations beyond ethics reviews [20, 29]. Ethical considerations can arise in computer security research even when people or their data are not directly studied or impacted [4, 43]. For example, research advances that can help software developers better find vulnerabilities could also empower adversaries, and cryptographic systems can be used for a diversity of purposes, not all of those purposes are "good".

Community reaction and changes in peer review. Partly due to piror ethically ambiguous incidents, interest in ethical discussion in the security research community increased. Changes were made to call for papers [38], and Research Ethics Committees (REC) were formed to review papers flagged with ethical concerns [37]. While the hypocrite commits paper was withdrawn, conferences also sometimes publish papers with ethical concerns documented [11, 22, 70]. Soneji et al. explored the peer review process in top-tier security conferences; they noted inconsistencies, randomness, and subjectivity in peer reviews and that the roles of program committee members were not well-defined [71]. The security and privacy community is working towards considering ethics in the review process, including adding ethics checkboxes [38], publishing public meta reviews [39], or, most recently, publishing dedicated ethics guidelines, requiring ethics considerations from all submissions, and dedicating extra space to ethical considerations [75].

Education within community and beyond. As humans, computer systems, and information systems move closer together, the possibility of impacting humans and society with security research grows. This is true whether or not the research directly involves human subjects. However, researchers may not always be aware that their research should consider the same ethical concerns as research involving human subjects. As a result, they may not receive help with thinking about ethics [8], or even skip (mandatory) ethics review [10]. Review boards typically operate across disciplines, and are not typically equipped to review the ethics of security research [17, 72]. Especially concerning vulnerability disclosures, review boards may decline review (as IRBs are only applicable for human subjects studies), or lack subject matter expertise [7, 17, 25, 36, 61]. Further, prior work notes that institutional review processes can be time-consuming and expensive—which may not be representative of all review boards—and can impede research progress [2, 7].

Reidsma et al. developed guidelines to better equip reviews boards for computer security ethics reviews, specifically when it comes to research involving vulnerability disclosures [61]. Like prior work, they discuss disclosing publicly and the threat of a vulnerability being exploited by adversaries [52], unintended negative consequences of research, and whether or not to conduct research in the first place [31]. Others advise to balance risk and benefits and minimize risk to parties affected by vulnerabilities [81].

Overall, prior work helps us understand the de-facto status of ethical considerations and decision-making as fragmented across sub-disciplines. We base our study on prior literature and inquire how researchers and reviewers reason about ethics. We also explore challenges, roadblocks, and wishlists for improvement.

[4]Consequential ethics centers on the consequences of decisions, e.g., considering whether a decision's net benefits benefits outweigh the net harms [43].
[5]Deontological ethics centers on one's moral obligation, duties, and rights, e.g., considering stakeholders' right to privacy and autonomy when making a decision, [43].

3 Methodology

We conducted a meta-review of all 1154 publications at the 2024 top 4 security conferences and the corresponding Calls for Papers, analyzing discussion of ethics in published security and privacy research papers, and conducted a semi-structured interview study with 24 researchers in the security and privacy research community. *Replication.* An extended version of our paper with the appendix containing the interview guide, codebooks, and meta analysis results is available online [58]; additional study materials for replication are available at https://doi.org/10.5281/zenodo.17034796.

3.1 Meta-Analysis

We conducted a meta-analysis of a year's worth of published security and privacy research, exploring of how and how frequently ethical implications were discussed in published security and privacy research, and which ethics aspects were addressed. We thus analyzed 1154[6] security papers published in 2024 at the big four security conferences: ACM Conference on Computer and Communications Security (CCS), IEEE Security and Privacy, Network and Distributed System Security Symposium (NDSS), and USENIX Security Symposium (USENIX Security). We also analyzed corresponding calls for papers through content analysis [35].

3.1.1 Data extraction. Each of the 1154 papers was manually assessed for (a) no mention of "ethics", (b) discussion of ethics without a dedicated ethics section, (c) discussion of ethics with a dedicated ethics section, and (d) mention of "responsible disclosure" or "vulnerability disclosure".

We also assessed if the paper discussed practices or considerations that directly map to principles from the Menlo Report. We used Menlo Report's companion to determine definitions of each principle (e.g., the report defines respect for persons as the presence of informed consent and protection of vulnerable persons) [74]. We then manually looked for indicators in the papers to code for each principle by first skimming through the papers and then we used keywords related to ethics to verify the presence of the principles in each paper individually. For instance, for the principle "Respect for persons", we assessed whether informed consent or protection of vulnerable persons/groups was mentioned. For the principle "Beneficence", we assessed whether confidentiality or balancing risk and benefits was discussed. For the principle "Justice", we assessed whether fairness and equity and compensation were mentioned. Finally, for the principle "Respect for law and public interest", we assessed whether compliance with laws and regulations was discussed. We also assessed *deception*, *human factors research*, and whether IRB, ERB, or any ethics board review or approval was reported (categorizing *review mentioned, no human subjects, exempt from review, review board approved, no board approval*), as well as mentions of *responsible disclosures, and vulnerability disclosures*.

3.1.2 Data analysis: Paper coding. The gathered papers were qualitatively coded by three researchers using qualitative content analysis [5]. For the first 30 papers, all three coders read the papers and categorized them, also memoing about interesting ethics content, then met to resolve questions and conflicts. Two researchers proceeded to independently code the papers in chunks of 30, noting questions and highlights, meeting in the team of three to resolve questions and concerns after each set of 30 papers. After 460 papers, the remaining papers were coded by individual authors only, who met twice weekly to resolve ambiguous situations. The lead author spot-checked a random set of 10% of total papers for consistent coding. One researcher independently analyzed all calls for papers.

[6]The number of our dataset does not match up exactly with DBLP; we think there may have been small consistency issues with indexing problems for DBLP.

3.2 Interview methodology

We describe developing and piloting our interview guide, interview process, recruitment, ethical considerations, and limitations.

3.2.1 Interview guide. We developed a semi-structured interview guide for the interviews based on our research questions and pilots of our meta-analysis of published security and privacy research. Four authors iterated over the interview guide. We piloted the interview guide with four (junior and senior) researchers from different sub-areas of security and privacy, and incrementally changed the guide for comprehension, depth through follow-up probes, and flow. After eleven interviews, we added questions for a deeper understanding of considerations on unintended consequences and research ideas or directions not followed after ethical consideration; these are marked in the interview guide [58].

First, we discussed interviewees' *background* in security and privacy research, their roles as authors, and potentially members of chairs of program and research ethics committees. Second, we asked about their *decision-making for ethics in security and privacy research*, including at which time in the research process they think about ethics, what prompts them to think about ethics, what factors they consider, how they collaborate with co-authors regarding ethics, and how they write and learn about ethics. Reviewers were also asked how they determine whether research they review is ethical, and what they expect to read about ethics in the work they review. Third, we asked how they handle potentially *unethical research*, including negative outcomes, how to mitigate ethics violations throughout the stages of the research process, and how they address this during peer review. Fourth, we ask about *personal experiences with research ethics*, past encounters with unethical research, teaching ethics to students, and changes they observed in how the research community addresses ethics over the span of their research career. Fifth, we asked about *the future of ethics research*: their assessment of the field's current handling of ethics, and desired changes and support. Sixth, we reflected on the interview, and held space for additional comments.

3.2.2 Recruitment. Four authors jointly discussed who within the security and privacy research community should be invited to be interviewed to achieve a broad sample. Through our professional network, we recruited authors of research papers and security program committee members across sub-disciplines of security and privacy research, with a broad range of experience, seniority, and geolocation. We recruited participants who had published and reviewed research papers at top venues in security and privacy (CCS, IEEE S&P, NDSS, PETS, USENIX Security) and/or had successfully published security and privacy research in HCI venues

(CHI, CSCW). We aimed for diversity in terms of geographic location, sub-discipline, gender, and seniority. We started a list of possible interviewees from a list of program committee members of these conferences, then diversified as follows. While we required successful publication in those venues as a recruitment criterion, for diversity of experience, we purposively interviewed junior participants (i.e., PhD students early in the program); we recruited participants who had previously served as PC members or chairs or REC members or chairs at security and privacy conferences. We had participants from the US and Germany (broad distribution across universities and states), UK, Switzerland, Austria, and Japan. We emailed reviewers and authors from a broader demographic range (including Hong Kong, India, Qatar, Netherlands, Switzerland, Belgium, Russia), but did not get responses from them. Participants' research areas span sub-disciplines, including cryptography, network security, systems security, usable security and privacy, privacy, and cybercrime. In addition to security venues, they publish in venues for privacy, human computer interaction, computer supported cooperative work, software engineering, cryptography, internet measurement, and criminology. We conducted qualitative data analysis concurrently with recruitment, and stopped inviting participants when responses became repetitive, which we interpreted as approaching data saturation.

3.2.3 Data collection. We invited potential participants one by one, in an email sent by the author who was most familiar with the invitees, cc-ing the interviewer (a junior researcher) and other team members. If participants opted into the interview, the lead interviewer scheduled a time with them, who conducted all the interviews, for five interviews, a second (junior) interviewer was present. For two interviews, with consent, another (senior) author was also present. We added two more interviews while this paper was under review to obtain more insights on ethical considerations while researching at-risk populations; these interviews enrich our findings, but do not otherwise influence data saturation. Participants received a consent form detailing their rights (to withdraw without loss of benefits, to revoke participation, data protection rights), and verbally consented at the start of the interview. We conducted and recorded (locally) via Zoom. Only the interviewer kept a file linking interviewee names to participant IDs. All interviews were transcribed through a GDPR compliant external service (for the first 12 interviews) or zoom's internal transcription feature. We corrected and de-identified transcript prior to analysis. De-identified transcripts and the link between participant IDs and emails were stored separately in a secure, self-hosted cloud, and only used to send participants drafts of this paper for comment.

3.2.4 Qualitative analysis of interviews. We used open coding to analyze interview transcripts qualitatively. We created the initial codebook with deductive codes based on the interview guide, inductively added codes by reading interview transcripts together,iterated over the codebook until it became stable and well-defined, then double (re)coded a total of six interviews. We achieved a high intercoder agreement (Krippendorff's $\alpha > 0.80$). The primary coder then used this codebook to code the remaining interviews, after which we affinity-diagrammed results into categories. We use these categories to inform our results section. Inspired by the results of the paper meta analysis, the second coder later re-coded all transcripts for mentions of complexity in ethics definition, and mentions of Menlo Report and Menlo principles.

3.2.5 Interview participants. We conducted the research interviews with 24 participants lasting 48 minutes on average. The participants differ in seniority and work experience: we interviewed junior and senior PhD students and research assistants, industry researchers, and professors with varying levels of seniority. We interviewed researchers from different research topics within security and privacy research, including systems and web security, cryptography, measurement, cybercrime, and human centered security and privacy. Six participants worked with at-risk populations, seven were women and 17 were men, and many participants had research experience in multiple areas.

3.3 Ethics considerations

We carefully adhered to the Menlo Report principles. Our meta-analysis of research papers was conducted on public data. As perceptions of "what is ethical" evolve over time, we acknowledge that in our analysis, we may have uncovered situations in which some authors may have made decisions that they should not have (either then, or under today's understanding of what is "right"). We do not want our data to be used as a mechanism to identify, shame, or harm other researchers, but as a way to empower future researchers to make more ethical decisions. Therefore, we believe that it is most ethical not to share the entirety of our data publicly.

For our interview study, we obtained ethics and data protection review and approval. Participant information was de-identified as much as possible while retaining meaning. Further, with two exceptions, no senior author was present and involved in the interview process to limit uncomfortable situations in case of personal connection or (previous) work relationships, or future requests for recommendations. In two exceptions, a senior researcher was present, as they personally recruited the participant and obtained their consent; however, the participant was more senior than the senior author present, and would not reasonably ever require a recommendation by the author.

As we asked questions about ethical conduct and their own previous research, we reassured participants that we were not trying to judge their opinions or handling of situations. We also made it clear to them that they could decline to answer questions or withdraw from the interview at any time. To protect participants' identities, we choose not to make full transcripts public, and only share aggregate data as well as de-identified quotes. We also hope that the positive impacts of this paper—hopefully helping the security and privacy research community better consider ethics—will outweigh the risks (e.g., the small risks of re-identification or causing negative discussions). There is a risk that writing a paper about current ethical practices will cement these, and foster "compliance"-thinking [60]. We hope that our research will instead foster communication, and also observe that, independently of our research, the field moves forward, for example with USENIX Security's new ethics guidelines [75] and other security venues' considerations of research ethics. We note that none of the authors who reviewed the Calls for Papers had a role in their creation.

3.4 Limitations

The paper analysis is limited by the venues we chose, and by the year we chose and therefore cannot generalize to other venues or other points in time. Our choice of venues further excludes specialized sub-discipline venues in security, workshops and other more informal published work. Therefore, the review is not an overview of the entire published works in the field, but more of work accepted at primary venues, with a heavy focus on security. Both from our own observations and also from interview data, we think that expectations for similar work, accepted at a high-quality venue, cross over specific venues, both with reviewer and author overlap. Though we read through all the papers, it is possible that we missed discussions on ethics in papers that used a different language than what we would associate with ethical decisionmaking. Further, the aspects of ethical considerations that we explored may be most applicable for research papers presenting studies, measurements and security vulnerabilities. Studies presenting new systems or tools (e.g., code vulnerability analysis tools) might have an in-depth discussion on beneficence and harms, but do not discuss aspects such as informed consent. Our analysis therefore does not entirely represent the depth and nuance of ethical discussions in papers.

The results from the semi-structured interviews have limited generalizability due to their qualitative nature. We also cannot represent the full diversity of how ethical processes are handled globally. All our participants agreed to be interviewed on ethics in security and privacy research, therefore they likely have a personal interest in the topic. Participants might present their ethical considerations and their work more favorably in the interviews, as breaking community standards or laws could have repercussions for involved researchers; we however felt that participants were speaking openly, and also reported thoughts and experiences that would not unambiguously be perceived as "correct". Our sampling strategy for research papers and interviewees was chosen to answer our specific research inquiry, and should not be understood as a blueprint for further meta-research.

4 Results

We discuss results from our meta-analysis of 1154 research papers published in 2024 at the big four security conferences: CCS, IEEE S&P, NDSS, and USENIX Security as well as our interview study with 24 security researchers, and present them aligned with our research questions. We include additional data from our meta-analysis in the extended version of our paper [58].

4.1 RQ1: Reporting and discussing ethics in published Security research

We find that most papers—especially technical security papers—lack discussions about their ethical considerations, and in the majority of technical cases with discussions of research ethics, the focus is on responsibly disclosing vulnerabilities to affected parties. Additionally, we find that human-centered security papers generally present detailed ethics discussions—sometimes with dedicated ethics sections—and some papers discuss ethics at a high level without a dedicated ethics section. Below, we provide a brief analysis of how ethics are discussed in calls for papers, and an analysis of how ethical discussions are presented, and which ethical aspects are discussed often in security research.

Ethics in calls for papers. The calls for papers corresponding to the publications we analyzed (all from 2024) vary greatly, with detailed guidance (IEEE S&P, NDSS, USENIX Security), and minimal guidance (CCS). In 2024, discussing ethics was not mandatory, and recommended when relevant. Unethical research can lead to rejection for all venues. Calls for papers link to Ethics guidelines ACM Code of Ethics and Professional Conduct (CCS) and the Menlo Report (IEEE S&P, USENIX Security; NDSS). The role of research ethics committees are outlined (IEEE S&P, USENIX Security, NDSS), with varying clarity and detail. The calls for papers mention expectations for vulnerability disclosures (IEEE S&P, USENIX Security, NDSS), human subjects research (IEEE S&P, USENIX Security, NDSS), and sensitive data (IEEE S&P), and discuss IRBs (IEEE S&P, USENIX Security, NDSS). They mandate integrity in the peer review process including conflicts of interests (CCS, IEEE S&P). They additionally explain how to consider stakeholders in the research, and address potential harms (USENIX Security, NDSS). Overall, we see different explicitly described expectations across venues, which may contribute to uncertainty in research paper preparation for authors, and also to uncertainty for reviewers during peer review.

Presence and absence of ethics discussions in security research papers. An overwhelming majority (839) of the papers we analyzed did not contain any discussion about ethics. A majority of these papers were making theoretical contributions, contained proofs, suggested better system designs, proposed vulnerability detection methods, and did not report on specific vulnerabilities or human subjects research.

The rest of the papers (315) had of some form of ethics discussions, with a subset of the papers (270) having a dedicated section in the paper with headline "ethics" (or similar). A subset of these papers (45) mentioned ethical considerations in the body of the paper without a dedicated ethics section, often in the context of methodological decisions, limitations, and presentation of data in the paper. For instance, one study mentioned ethical concerns as their reasons not to name brands of vulnerable devices in the paper. Another study justified the use of their dataset by stating that they are publicly available and do not contain harmful material (e.g., not depicting abuse). Another study briefly explained that they had institutional ethical clearance, obtained no sensitive data from participants, and participants consented to the study. While some papers explained ethical considerations briefly but concisely, a few were vague and only mentioned adherence to ethical guidelines without any mention of *which* ethical guidelines they followed.

We also observed that a few papers without dedicated ethics sections provided detailed explanations of ethical considerations throughout the paper, explaining ethical decisions they took in every step of the process, often woven into decisions in the methods section and study design.

Balancing harm and benefits are commonly mentioned ethical aspects. To better understand ethical principles presented in papers, we analyzed ethical consideration using principles from the Menlo Report, which serves as a guiding principle in Information communication technology research (see Section 2). While only 36 papers explicitly referred to using the Menlo Report by name.

Using explanations that we were able to map to the Menlo Report principles, namely, respect for persons, beneficence, justice, and respect for law and public interest, we found explanations of these principles appeared in 251 papers.

The most common discussion of ethics was on the principle of *beneficence* through balancing/ mitigating risks and benefits (159) (i.e., through discussion of harm and how the benefits outweigh those harms), and confidentiality (110) (i.e., protecting identity of users and stakeholders through data protection, anonymization, etc.) For example, studies reported that advantages of disclosing a vulnerability outweigh harms, or that a paper's research does not inflict any form of harm and that research advances benefit users.

The principle of *justice*, through fairness and equity was mentioned in 18 papers, which could be any discussion of burden placed on vulnerable and marginalized population, and 79 papers mentioned compensating study subjects (fairly). For the principle of *respect for persons*, 89 papers mentioned informed consent by debriefing their participants and 11 papers explicitly mentioned the protection of vulnerable persons. Finally, 56 papers mentioned *compliance with laws*, often data protection laws (specifically: GDPR). Generally, we find that published research seems to view the Menlo Report principles as a useful framework to report ethical thinking.

The focus lies on 'responsible' vulnerability disclosure. 257 papers mentioned disclosing vulnerabilities, and most of these papers use the term *responsible disclosure or responsibly disclosed*[7]. Some papers also included *responsible* disclosure under the ethics section. Some of these papers had dedicated responsible disclosure sections where they detailed their ethical considerations and decision making processes and specified details of their disclosures.

Human-centered security papers engage with ethics explicitly. Of the 1154 papers we analyzed, 116 reported on studies with or involving human subjects. Of these, 102 papers include ethical considerations, 90 include dedicated ethics sections. Human-centered papers often require and specifically mentioned obtaining IRB approval or some sort of ethics review.They also often provided detailed explanations of ethics, including the presence of informed consent (64 papers), confidentiality (72 papers), and minimizing harm to participants (52 papers). All but 17 papers mentioned IRB or review board approval, 9 papers mentioned wanting but not having access to review board approval. Only 52 papers without human subjects research mentioned review boards approvals.

We generally find that, while ethics discussions are present in published work, in more or less detail, we can rarely understand deeper ethical considerations from the published research papers. And, of course, research that was never published (or never conducted) for ethical reasons, is not present in this data set.

4.2 RQ2: Authors' reasoning and decision-making around ethics

We find that researchers describe ethics as hard and complex to refine, and context-dependent. Ethical reasoning is often guided by the principle of beneficence from the Menlo Report—however, harms and benefits can be complex to determine in security research, and harms can extend to people, systems, companies, the environment, and the researchers themselves. Further, there are different ethical considerations across sub-disciplines and industries, calling for a nuanced approach to research ethics.

4.2.1 *Ethical reasoning and definitions.*

Defining ethics is complex and context dependent. A majority of the participants stated that a definition for ethics is not easy to formulate, and often implicitly communicated from a consequentialist stance, namely that ethics in research should consider the (potential) harm of the research—specifically "*balancing potential harm with a potential positive impact*" (P12). Some researchers explained that they approach ethics through the lens of avoiding causing harm to humans, and stated: "*Is my methodology harming someone or is my outcome or my results? Are they in any way hurting someone?*" (P18) and "*If you consider ethics you are considering human well-being in your studies, or at least try to avoid harm to other humans or avoiding harm to humans in general*" (P02).

Researchers consider their work's benefits, describing their definition of ethics in research as "beneficence". They want to balance the harm and benefits of their research to have overall positive impacts. Therefore, in some contexts, a limited amount of harm might be within the bounds of their definition. Participants also described that in security research there is always a harm, and it is important to balance the harm out with potential benefits, as P12 elaborated: "*There can be harm in what we release. […] There are a lot of ways for there to be potential for harm, but there's also this question, what is the potential positive in doing this research and understanding problems?*" (P12). And one participant also explicitly talked about their ethical reasoning directed by the impact of the study: "*I'm not sure if I would say this is always necessarily ethically driven, but I think there's this question, what impact do you want to have? What is the point of going after a certain research paper*" (P14)? Additionally, one researcher even mentioned that it may be *unethical* "*to not do a little harm if there is a lot of good coming out [of the research]*" (P10).

In general, participants also indicated that balancing harm and benefit is complicated, not only in assessing certain boundaries but also in determining what level of is acceptable. One researcher mentioned that there are "*certain red lines I would not cross that are part of good scientific practice*" (P18) and further elaborated that they would for instance never misinform and avoid getting consent from participants because "*I think everybody has the right to know if they participate in a study*" (P18).

Beyond causing harm, one industry researcher considers ethics as necessary to navigate conflicts of interest between all the involved parties, "*whether it's the person doing some work, the employee, the researcher, perhaps the person who's paying for it, perhaps the person who might be impacted by it*" (P08). They discuss balancing those different interests with ethical considerations.

Context dependency was mentioned as a struggle, and also as a source of disagreement with other members of the research community: "*One of the challenges I find in debating with people on PCs about ethics is that the people want it to be simple. They want it to be black and white. That there is an easy rule and they don't like the context and they don't like the fact that we actually end up having different standards*" (P06). The context dependency can make

[7]Though we find papers using the term responsible disclosure, the community in the industry prefers to use the term coordinated disclosures [43].

decision-making regarding ethical considerations more difficult, especially for researchers with limited experience and guidance.

Harm can extend beyond humans to businesses, systems, and environment. Most participants defined ethics as majorly avoiding harm, they mostly focused on causing harm to humans—through study participation or resulting harm through publishing research. Echoing findings in earlier works, some participants elaborated that harm can extend to companies, systems, and critical infrastructure [31]. Some participants also elaborated that they can harm systems, by stating that they often have to take "*the role of an attacker*" (P22) and research with a live system. In this case, one researcher elaborated their definition of ethics as: "*I find it unethical research if you cause detrimental side effects to running systems. So I would extend the definition of ethical research beyond human-centered research at this point*" (P22). Another researcher altered the methodology of their measurement study to avoid harm and disruption to older network infrastructure: "*Because they [old network] might actually hit systems that are really not robust enough to withstand even standard scans, that might fall over*" (P21).

One researcher who worked with publicly available questionable marketing claims from companies reflected on whether publishing their findings will harm the reputation of these organizations: "*It is harming companies and their businesses. The question is, is that ethical?[…] It's causing damage for those companies but it's true. I don't know. My solution is I don't do this[publishing findings]*" (P03). Participant P21 also mentioned that certain studies around blockchain can harm the environment, and they intentionally avoid this type of research: "*[…] the harm that those [crypto and proof of work schemes research] do to the environment is much too big and from an ethical point of view I cannot condone and I will not myself support this kind of research*" (P21). The researchers themselves can experience harm in various forms: psychological when dealing with deeply sensitive topics like intimate partner violence or child sexual exploitation material (P23, P24), legal issues or encounters with law enforcement (P06, P23), or harassment due to published results (P23). Therefore, it is vital to consider stakeholders and systems that maybe harmed via the research process.

Researchers often rely on their *instincts* and experience to guide their ethical reasoning. Though the Menlo Report is commonly used in ICT research to guide ethical reasoning, we found no common agreement on a universal standard that researchers followed for ethical guidance. Most often researchers mentioned that they rely on their instincts, for instance, based on experiential learning—through life and work experience—as P08 explains their reasoning of fairness: "*I think this question of fairness is a thing we built as kids. […] It's not something I've ever studied. There's no process framework, formalism, the question of what makes an instinct theorem*" (P08). One senior researcher in particular, who mentioned that they pioneered some of the first ethics sections in security research mentioned that they "*have no idea. We just starting doing it. I think we did a little bit of research on ethical philosophy and picked the structure that matched how we looked at the world*" (P06). Other researchers—especially during their early career—mentioned that they developed their reasoning, for instance, through learning from their advisor (P09, P24), interacting with stakeholders (P12, P23), and through colleagues (P21, P23).

Ethics differs across research sub-areas and academic and industry research. Ethical concerns and IRBs are expected in peer review for human-subjects work and researchers in this area have guidance—sometimes checklists—to ethical considerations: "*For human subjects research, there is a lot of groundwork that we can just rely on, which is why I also think for, for some of those cases we can do checklist-based approaches of bit of self certification*" (P21). This is also reflected in our paper analysis (Section 4.1), where most human-subjects papers discussed ethics in some form.

Researchers from a cryptography and mathematics background less commonly discuss ethical concerns in their research. A participant lamented what they describe as a lack of direct connection between cryptography research papers and the real world: "*I actually wish we had more papers being published where there were ethical conundrums involved. A lot of the papers end up being just so pure math. What somebody did to a bunch of algebraic variables is not really an ethical question. The disconnect between a lot of the research results in the real world, I would argue, is a problem rather than a feature. I would look forward to seeing more people asking questions about that*" (P08). Nevertheless, cryptographic research impacts society down the line, so the motivations for the cryptographic explorations or the advancements made can still receive ethical considerations [64].

Participants who have experience working in or with industry perceive the standard for ethical research as quite different from academia. "*I think the academic world tends to be overly cautious […] whereas industry tends to be dramatically under-cautious. There's certainly a double standard there*" (P08). This divide may be connected to a lack of support for ethical research in the industry. "*But there's not like an ethics department usually in the industry that you could contact or even a smaller subset of part of the company that you can contact. […] In this case we managed to do the IRB because we were partnering with researchers from a university and they did it through their university*" (P19).

Research affecting human subjects and at-risk populations call for additional care. A few participants mentioned that they treated all their research subjects as being vulnerable and always take precautions (e.g., following appropriate de-identification practices), with one participant specifically stating: "*I treat in many ways all my research projects as potentially sensitive, and all of my research participants are potentially vulnerable because you never know where the issue might rise*" (P24). However, participants (P19, P24) who work with at-risk populations like incarcerated people and intimate partner violence survivors suggested that they take additional precautions in their research process. They mentioned reaching out to NGOs to approve their research protocol and involving them in the research process, paying special attention to de-identify participant data, and enforcing strict policies to prevent sharing data with third parties like law enforcement that could endanger participants.

Researchers think writing about ethical decisions is important as security research can have negative consequences. None of the participants were indifferent to the topic of ethics in research.[8] Many had examples of research they had encountered

[8]This result may be strongly influenced by sample bias, as all participants agreed to participate in an interview about ethics in security research.

in the field of security and privacy that they thought was unethical. P14 criticizes that not all researchers consider and present the ethics of their work in submitted papers: "*And that's something [...] that I do ask as reviewer [...] Can you write a part about what is good about the research? And what is bad, right? How it can be misused. How it can be like this computer science idea that we're computer scientist[s] and therefore we're not guilty of anything. We're past that time. We actually, the world, everything we do, mediates the world, creates interventions, changes things, including all of these privacy technologies*" (P14).

Further, many participants mentioned that ethical considerations are a core part of a study's design, and it is important to write about their ethics considerations, especially if the research is critical: "*I expect to see some discussion of it. I don't feel strongly about whether it's an explicitly labeled section or if it's in line with the text, but I do expect to see something written about it*" (P10).

4.2.2 Decision making with ethics in research.

Most consider ethical considerations during their research project design. Considerations depend on the nature of the research topic: user studies as well as measurements and attacks on live systems prompt researchers to consider ethics during the initial stages of research design. Participants reported thinking about ethics during the research project design as part of the methodology, as soon as ethical implications are apparent. "*We think about it a lot in the context of methodology. What can we do? What can't we do? That's early on before experiments get done*" (P06). Some participants working on technical security topics reported that ethical implications do not come up in early study design and therefore ethical considerations might start later in the writing process. "*They don't involve human subjects or attacking systems. The main place ethics comes up regularly is in terms of evaluation and how to present results*" (P10). Ethical concerns also lead some participants to not pursue research questions, due to a lack of ethical research methods to answer these questions. "*For us, there are a lot of questions we'd love to understand about the internet, but we think would be unethical to answer. These are questions of how far you interrogate a system*" (P12).

A few researchers emphasized the need for deliberating ethics in the early stages of research design, especially by consulting with stakeholders and colleagues who can provide input on the study early on: "*My advice is just think about things as early as possible and it's still possible to make changes or do things differently. It's very hard after research has been done or after you're locked into some research methodology too*" (P13).

Ethically ambiguous research ideas are discarded before implementation. Participants reported deeply discussing ethical implications during the inception stages of research ideas, and multiple participants expressed not pursuing research ideas due to ethical concerns: "*There have certainly been lots of things that we just simply don't do because we can't figure out a way to do that in a way that we feel ethically comfortable with*" (P11). One researcher in particular even mentioned that they did not pursue replicating attacks on live systems citing ethical reasons: "*One thing I really wanted to do was to kind of replicate various sophisticated cyber attacks and that's very challenging and that is also an ethical concept like how would you do it on a university network, right?*" (P17). Further, one participant highlighted that in their line of work some researchers work with cybercrime offenders using deception, and they state that "*it's[the research idea] been something that I've kind of shot down pretty quickly[based on ethical grounds]*" (P23).

Ethical decisions are mostly made through collaborative processes. Ethical decisions are usually not made in isolation, but through collaborative processes. One researcher expressed that when planning for a large-scale internet scanning study, they had many ethical questions due to which they had to involve various stakeholders in their study design: "*A project I did in graduate school had to do with doing very large-scale internet scanning, which had a lot of ethical implications immediately. It's something that we thought about from the very, very beginning, and we talked to a lot of folks across the community. We talked to CSOs, operators, attorneys, and senior members of the community to get their thoughts*" (P04). Sometime collaborative decision making can also help researchers to not have a "*myopic view*" (P05) of for instance, the industry. Collaborations help junior researchers thoroughly think through study designs: "*When I have felt unclear, I've been pretty comfortable talking to colleagues and getting feedback*" (P06).

Ethical decisions shaped by the research community and prior work—The Menlo Report and vulnerability disclosures. Researchers first bring up the ethical considerations with which they are familiar, broadly participants mentioned principles from Menlo Report and vulnerability disclosures. Researchers who are involved with user studies usually first discuss general human subject protection, including informed consent, confidentiality, and fair treatment of participants, which are also prominently mentioned in the Menlo Report: "*respect for persons, for people, which involves trying to obviously minimize harm to people, give them informed consent, let them know the purpose of the research where possible*" (P13). These researchers discussed that working with vulnerable populations required in-depth ethics considerations.

For researchers dealing with security vulnerability research, coordinated disclosure and informing affected parties is directly linked to ethics: "*Thinking through the disclosure process is an important part of the ethics*" (P05). In addition to disclosures, participants were also concerned with fixing security vulnerabilities, to minimize harm: "*So immediately as we kind of realized that there was really a vulnerability, we [...] interactively think of a fix that we could work together with the people who maintain the system that has the vulnerability*" (P19). Those not directly working with human subjects or vulnerabilities discussed research ethics as truthful representation of the collected data and not altering results. Many participants were also concerned with the potential broader societal impact of security and privacy research, which some mentioned learning from prior work and conversing with peers: "*What's the impact on the environment of the participant or a broader scale?*" (P03). This can include the misuse of developed tools by malicious actors or the introduction of technology that could affect societal changes.

4.3 RQ3: Challenges to ethical reasoning in security and privacy research

Review boards were discussed as the most common prompt to considering ethics during research, but they were rarely mentioned as a prompt to *deeply think* through ethics—depending on the specific

instantiation of the board, researchers appreciated sharing responsibility or finding (small) flaws in study designs. However, compliance with review boards might frustrate or slow down researchers.

Review boards can be (un)helpful and can slow down the project. Many participants reported contact with their review board. For research involving human subjects, participants report that review board approval is common and often important for paper acceptance. However, most review boards do not have the expertise to advise researchers with more technical study designs; these studies are usually not reviewed (in detail).

Without in-depth understanding, review boards might not be able to make a proper assessment: "*always [it] is really on authors to explain properly to the IRB why they might need to care and not trying to provide as little information to get an easy IRB [approval], that would be wrong by authors and I've seen that*" (P22).

Review board approval can be difficult to obtain, as researchers outside the U.S. reported not having access to a review board: "*Often they didn't have an ethics board review them […] To that time, or sometimes even now, I don't have access to a board. I don't have money to pay an external board*" (P03). Researchers who are working within the industry can have the same problem. Sometimes this is solved by collaborating with researchers from (U.S.) universities.

Obtaining board approval can often—but not always—be a lengthy process, involving bureaucratic hurdles, and back and forth interactions that frustrate researchers and slow down their research: "*[…] it's a common problem that if you submit stuff to an IRB, you now have to wait like three months until they get back to you*" (P02). These insights echo the critiques describing institutional review boards as overly bureaucratic and ill equipped to critically assess ethical risks in security research [2].

Review board approval does not mean that the approved research is ethically in line with regulatory research requirements. "*[…]it's compliance with the common rule for federal funding. It doesn't necessarily mean that,[…]we as a community might view that particular work as ethical. There are steps beyond it*" (P11).

Some participants find their review board's recommendations helpful: "*[…]sometimes you hear back from IRB and they have some recommendations for improving things, or they ask for stuff you haven't thought about[…]These give additional pushes towards more ethical thinking*" (P02). It can also be reassuring for researchers to share some responsibility of making decisions, and to receive support in how to protect their potential human subjects.

Identification of stakeholders and affected people can be difficult and complicate the decision-making process. Finding and establishing contact with appropriate stakeholders, especially for vulnerability disclosures and sensitive research, is often difficult. While some participants feel comfortable with what they perceive as community guidelines or best practices for vulnerability disclosure as "*rather straightforward on how or what you do*" (P07), some participants also reported that it can be difficult to identify the affected parties—especially for open-source projects—and inform them properly within acceptable deadlines: "*Sometimes you don't even know the library which is affected, and finding the developers who are responsible for this is also hard. Finding the email addresses of the responsible persons seems to be obvious but it's not always that easy*" (P02). The identification of affected parties can also be difficult for measurement research, where it is often hard to find who is (in)directly affected by research activities and network interventions: "*If you interfere with a network in the course of your research, you don't know who's using it and what impact it might have on them, so there's a lot of projection*" (P13).

There is no formal guidance for what is (un)ethical. Most participants did not receive formal ethics education nor formal ethics guidance and learned by doing research: "*Not explicitly. You learn along the way*" (P16). Some participants reported receiving ethical training later in their career, e.g., training by their review board[9], to be able to conduct user studies, or training to serve on a review board. Most participants reported learning about ethics through advisors, colleagues, and personal interest. "*It's something that was top of mind for my advisors, and I think I absorbed a lot of it from them*" (P12). Participants see room for improvement in ethics education: "*As a community, what I would wish for is training integrated as early as possible, repeated from time to time and clearly have processes maybe established at institutions*" (P22).

4.4 RQ4: Reasoning about ethics in peer review and as a community

The last step for research, once written up, is peer review—and in some cases, this is where ethical concerns are discussed, possibly for the first time. Participants described their experiences and concerns with assessing ethical research practices in papers submitted for peer review, pitfalls of requiring ethics discussions that may yield boilerplate text but no deep considerations, and the desire for more guidance, constructive discussions, and community standards.

4.4.1 Factors that help reviewers during peer review.

Ethics peer review strongly depends on information provided by the authors. Information about ethics in papers helps paper reviewers assess papers' ethics; more information is frequently requested during rebuttal phases, which allows authors to explain aspects missing in the initial submitted paper to clear up ethical concerns. Sometimes reviewers mentioned assessing papers solely based on what authors describe: "*I don't think you can determine if the research was done ethically. You can determine if they address the ethics issues in their write-up*" (P04). When reviewing papers not in their expertise, reviewers rely on authors' explanations to assess ethics: "*Some crypto person tells me, this will change the world in such and such a way, because this crypto, this math is good. I'm like, okay, I have to believe that. I'm not necessarily always in a position to understand what that means and how that's going to impact the benefits*" (P11). Relying solely on what authors write may even be a *flaw* in the reviewing process: "*I think the entire review process is fundamentally flawed. It's very hard to determine if the research was done ethically*" (P04). Sometimes, papers are rejected due to insufficient explanations and lack of detail for ethical decisions, for example P17 stated rejecting papers that "*didn't do a good job explaining all the [ethics] details*" (P17).

Written ethics sections as the de-facto community standard. The de-facto community standard is evidenced by research that ends up published at the major conferences, and the ethics considerations

[9] particularly for U.S. IRB approval, researchers are usually required to complete online ethics coursework.

are therefore reflected in the publications' ethics sections. Writing about ethics in research papers is considered as a requirement when preparing a paper for submission. Some participants mentioned an ethics and/or responsible/coordinated disclosure section as part of their "paper template". "*It's always in our introduction template*" (P01). Reviewers stated generally expecting ethical reasoning in papers: "*[…] part of ethical research is having a good explanation and careful documentation of what happened. There have been a couple of cases where authors were missing details, couldn't really explain and […] I'm like, 'If you can't clearly explain why the research was done ethically, then that's an ethical issue'*" (P13).

Some mentioned that it is now an expected community standard to have an ethics section as conferences require them, the decision is more about what to include in the sections. "*It's just best practice and there's been increased standards in committees. It's not really a discussion of whether we're going to have an ethics and disclosure section nowadays. It's just what's going to be in it*" (P05). Some mentioned having no feelings about separate sections but think that ethical explanations help with educating people: "*I have no strong feelings. I think that it is easier to educate people if you put it in a separate section though*" (P06).

Not all papers may need an ethics section. Some participants note that indiscriminately always requiring ethics sections can devalue ethics discussions into *"boilerplate"* ethics sections. "*I'm not a big fan of the fact that everyone has to have an ethics section. I think it diminishes the value of ethics sections if it becomes boilerplate. I think that papers either have the potential to have real side effects or have touched people in a way that any reasonable person might find harmful, that it's the right thing to do*" (P06).

Research ethics committees (RECs) support reviewers. RECs reflect the change in the importance of ethics, as they support conference chairs, and support reviewers unsure about ethics. Most participants who interacted with RECs described them as helpful in the reviewing process and improving the discussion on ethics: "*I also think they're good. […] They help. Keep the discussion going and we can only deal with this through discussion*" (P14). However, the same participant also criticizes that REC members might not always have the necessary knowledge to review all the research they are assigned. Others suggest that RECs could do more to educate the community on how to conduct research ethically. "*Instead, I think there's an opportunity there to teach and to help fill the skills gap. I don't think that those ethical review boards are doing that*" (P11).

4.4.2 *Challenges for researchers with peer review.*

Overly strict and inconsistent approach to ethics in peer review across sub-disciplines and geographies. Program committee members may sometimes lack domain and subject matter expertise for papers they review (for ethics), some participants mentioned that this could even be an issue in program committee meetings. Some members of the community might have a very clear-cut, inflexible approach to ethics, which can lead to disagreement with researchers who value contextual discussions of ethics. One participant said "*criticizing that others may, for example, consider attacking a US-based company's firewall unethical, and consider breaching the Great Firewall of China ethical, at the same time insisting that ethics are clear-cut*" (P06). Similarly, one participant stated frustration that sometimes reviewers lack cultural understanding in other locations, as P23's study was misunderstood as unethical, highlighting a very US-centric approach to ethics: "*we had a reviewer who was obviously very American, and had a problem with what we had done based on their assumption that students share rooms, and in our [university], they don't*" (P23). Several other participants noted that review process needs to account for non-Western and non-US contexts (e.g., non-US institutes do not have IRBs).

Review processes are reactive. Some participants criticize that the ethical review process is presently mostly reactive. Oftentimes research is only checked for ethical considerations after submission to a major conference; harm may already be done at that point. "*I think we have to figure out how we get to the point where we're being proactive instead of retrospective. That's pretty fundamentally broken right now, and it's not something that we're trying. I think that that's probably not fair. I don't think we have yet successfully figured out how to be more proactive and preventative of causing harm*" (P12).

One solution for this supported by many participants could be a community wide ethics body which would be available for questions during study design and the entire research process. P10 states that ethics are handled " *[…] reasonably well, but maybe still slightly ad hoc, which is why I think this community-wide board would be beneficial. It'd be a central repository of expertise, and knowledge and to some degree create a more universal standard*" (P10). This would support researchers in their decision-making process, possibly add new perspectives, and give authors more security when submitting to a conference. Some of the participants expressed ideas of how such a body could work. One suggestion is pre-approving study designs, mirroring psychology research, another is a community-wide advisory board, which could support both authors and reviewers.

It is unclear how such an ethics body could be initiated, who would staff it, and how it would be funded. "*Yes. If it existed, it would be helpful. I think everybody agrees with that, but don't really know how to start it, how to provide incentives for it to operate, and for people to spend time on it* " (P13).

4.4.3 *'Unethical' research shapes community's ethical experiences.*

Unethical research: unpublishable or educational? Many of the participants express that research conducted in an unethical and potentially harmful way, for example by involving human subject without proper consent, should not be published[10]. P05 states: "*I also think an instant rejection. Other reviewers on this paper don't feel the same way, but I think that we should not be publishing papers that have ethical concerns, nontrivial ethical concerns*" (P05). Some participants think publishing papers with ethical concerns could set a precedent for future, similarly problematic research.

Research assessment is dependent on the context and how much harm the research has done or would do when published. Researchers agree that if the harm of the paper lies in publishing it, e.g., by publishing confidential information, it should not be published. Some participants think other research with ethical concerns where the research has already happened could be published if the results provide benefit. "*Oh, yes. I think everything should be published. Meeting some sort of regulations, but it should be published even if something goes wrong. It's a lesson learned*" (P16).

[10]This is also discussed in Scenario E* in Kohno et al.'s work [43].

Publishing research with ethical concerns could also create the possibility to discuss these considerations in public, otherwise the reviewing process is bound to confidentiality. P06 explains: "*There is some value in having some mechanism for public signaling. I don't think we've found a great one that both prevents the harm from getting out there and does not unduly penalize researchers who acted out of ignorance as opposed to malice*" (P06). Participants rejects the option of not publishing ethically questionable work: "*It's too late at that point. What are you going to do? Tell them to publish it somewhere else. No, it's almost always too late at that point. […] I think that this practice of banning research, because the underlying ethics were problematic, is not helping the community*" (P04).

The discussion of past problematic/unethical research could help in strengthening community standards. Many participants think that it is normal for older published papers to not conform to current community standards, as they "*change over time*" (P04). They emphasize that, if a paper with unethical methods causes no harm by existing online, they can serve as case studies for the community. Moreover, it would be very difficult to effectively remove research published online: "*I think it'd be one thing if the research is doing ongoing harm. If you know someone who is actually being victimized by virtue of this thing being out there, that's a different situation […] For example, we used to do tons of research that involved just having wide open taps on the full content of traffic, which now I think is much harder to do. I don't think it solves any purpose to say, "Let's omit those papers"*" (P06). Availability of past case studies might help researchers and reviewers when they assess ethical concerns: "*If there are some corner cases, and how to handle them. Maybe a paper that summarizes specific problems that happened in recent years and how to handle them*" (P01).

However, participants draw the line at retracting work with fraudulent data. P15 encountered such a situation and explains:"*[…] we tried to reproduce this with a tool that we had written. And then we saw that the benchmark was really very carefully crafted to make the author's tool look good*" (P15).

5 Discussion & Recommendations

Ethics norms and ethical understanding are commonly shaped through discussions [9, 68]. We find that the identification of ethical problems and ethical decision-making in security and privacy research are also formed through discussions within research teams (see Section 4.2.2); likewise, (changes to) the approach to ethics in the community at large are also mediated through discussions, including discussions of past research.

5.1 Determining what is (un)ethical in S&P research

Our study shows that defining what is ethical is nuanced and intricate. Interview participants often defined ethical research as reducing harm and maximizing benefits. However, delineating harm and benefit is not simple and is contextual—leading researchers to assess their research or the ones they review on a case-by-case level. For instance, whether internet scans that cause issues for operators, servers, and services may be acceptable can depend on the proposed research benefit, and the target of the scans [31, 43]. To confidently determine, for example, when benefits outweigh harms, researchers may still need support from the community, peers, examples of prior work, and some level of self-determination. We make recommendations for research ethics below.

Institutional Review Boards (IRBs) do not equate ethical research. IRB processes, though sometimes helpful in bringing up ethical issues with research, have been critiqued for cumbersome bureaucracy and slowing down research [2, 28, 61]—also echoed by participants in our interview study. IRBs can also be treated as compliance check, not necessarily prompting researchers to carefully think through other ethical implications. IRBs may not be equipped to assess security research appropriately (as they may not understand the data sources, methodology, or risks), as demonstrated by past security research not flagged by IRB, but during peer review. Additionally, research may have consequences that the researchers may not have accounted for like data re-identification, psychological harm to participants and researchers, data misuse by adversaries, legal issues and so on [72]. Therefore, we encourage researchers to proactively engage with local review boards, but also to go beyond IRB requirements in ethical decisionmaking. Similarly, we recommend including specific guidelines in conferences' call for papers to encourage both authors and reviewers to consider ethical issues beyond IRB approval.

Using the Menlo Report as a starting point for ethics. As a starting point to support researchers with ethical deliberation, many report and recommend using the Menlo Report as a *self-governance* tool for conducting ethical research—especially for researchers writing papers—as it offers concrete guidance on what is morally right to do [43]—and the Menlo Report has become a widely used starting point to contemplate ethical decisions [68, 75]. While the Menlo Report is a good starting point, as discussed in Section 4.2.1, it has its own limitations like lacking actionable guidance for researchers, which may help them reason through challenging new scenarios—especially for security research. Additionally, there is a lack of focus clear guidelines for working with vulnerable populations and stakeholders who may be disproportionately affected by the research and the outcome of the research. Therefore, we recommend the research community to work on providing more practical guidance or expanding on the Menlo Report, for more challenging ethical dilemmas that security researchers may encounter.

5.2 How conferences can shape ethical thinking

Conference ethics guidelines in calls for papers play a significant role in shaping the security research community's approach to research ethics. For instance, they can mandate ethics sections in all research papers, and include a research ethics committee to review research that is flagged as unethical during peer review. These approaches, however, may lead to ethics considerations around the time of or after paper submission. Therefore, we implore conferences to provide clear guidance for proactively approaching ethics through resources like offering prompts for ethics sections, adding examples for ethics sections from prior work, and clarifying that frameworks like the Menlo Report should be considered in-depth.

Reflecting on ethical decisions and documenting them in research papers for the community. Including ethics consideration in papers has become common, and, for some venues, mandatory [75]. Writing about ethics and addressing review board approval or exemption may be required for paper acceptance. We highlight the criticism that this approach can lead to generic and perfunctory ethics sections (4.4.1) without in-depth discussion and reasoning through ethical considerations [32], which may be suitable for research with low ethical impact. While some works may have low ethical concerns, we deem it critical for authors not to *a priori* assume that their work has low ethical concerns but to instead to a full analysis before making such a determination. Interviews show that researchers use consequentialist reasoning about beneficence, balancing harm and benefits (4.2.1). Their reasoning is not reflected in the literature meta-analysis and this underlines that publications do not presently reflect the depth of researchers' ethical reasoning (4.1), possibly due to a limited amount of space, a lack of incentives, and the invisibility of research abandoned for ethics concerns. USENIX Security's new approach to offer an extra page for mandatory ethics considerations may prove beneficial, and may help bring the "hidden" discussions into the published body of literature. In addition to requiring ethics considerations for publications, research specific ethics guidance and considerations should also be supported, similarly to Hantke et al.'s collection of practical guidance in server-side scans [31]. Further, we recommend sharing ethics prompts and practical guidance for sub-disciplines, e.g., through practical guidelines, non-exhaustive lists of key considerations, and examples of prior work. This may make it easier for those new to sub-disciplines to create ethical study designs, without often-scarce direct advice from more experienced peers. However, these prompts should be treated as starting points to guide ethical thinking, not as a checklist to be completed. We recommend approaching ethical research as a fundamental commitment to the well-being of affected stakeholders rather than a compliance checklist.

Encouraging a proactive approach to ethics through early guidance in call for papers. Handling of ethics in security research has evolved over the last decade, with conferences including research ethics committees and prioritizing ethics in their call for papers [12, 40]. If research ethics problems are only discovered during peer review, harm may have already been done, or, at best, resources squandered. We note that USENIX Security's current ethics guidance encourages thinking about ethics early in the research process, and potentially abandoning unethical research [75]. It may be interesting to observe the impact of this change in ethics requirements on published work, to identify if this helps with thorough and proactive ethical deliberation.

Research Ethics Committees and conferences can shape the discussion of ethics. Currently, research ethics committees make decisions during peer review, i.e., after research has already been conducted. This raises the question about the dual role of RECs: Do they simply gatekeep unethical research from being published, or do they also play a role in educating researchers? Hantke et al. suggest guidance through anonymized and published REC decisions [31]. Ideally, RECs can play a role in educating researchers—as a last line of defense—providing guidance on mitigating ethical issues before publishing. This mechanism to ensure that published work is ethical does not fill the gap of an active community outreach to encourage proactive ethical thinking. This outreach might entail up-to-date case studies, current events, and guidance for researchers in different sub-fields.

5.3 Encouraging continuous community discussions and learning

As technology changes, as researchers incorporate new research methods, and as researchers and the community learn more, community discussions and standards on ethics should evolve. We propose continuous ongoing ethical discussions involving new technologies and new community understandings and discussions of identifying ways to handle ethically ambiguous research.

Adapting ethical guidelines and discussions to new technologies. Ethical considerations may require flexibility to adapt to new technologies (e.g., generative AI) and emerging challenges (e.g., large-scale misinformation that threatens democracies, security tools used for surveillance of civilians); as new technology evolves, community norms for what is morally right have to adapt.

Therefore, we recommend ongoing ethics discussions relating to new technologies and research methods. We propose to anchor them in the community by giving them space in community events (e.g.,ethics discussions and workshops in conferences). Likewise, as the community learns more, what might have been considered ethical in the past may no longer seem ethical, thereby warranting the revisiting of past strategies for considering ethics in new contexts. Our interviews show that participants reported informal and private discussions about ethics, which is promising, but excludes researchers without easy access to more experienced persons in their field, or not on program committees. Discussions and decisions made by research teams and during peer review might be documented and made available for the broader community.

Handling of ethically ambiguous research at conferences—publish or reject? We found that researchers were not in agreement on how to handle research with ethical concerns at the reviewing stage in conferences (see 4.4.3). This question is also addressed in Scenario E* by Kohno et al. [43]. The arguments for publishing research with ethical concerns are encouraging ethical discussions.

An argument against publishing the research is that committees want to avoid setting a precedent for other researchers. There is no standard approach to this. One approach has been to publish research with ethical concerns and adding a public statement by the program committee about ethical concerns and authors' reflection on their research methodology. As a community, we invite discussions on whether the approach to publish ethically questionable work with a disclaimer is the optimal way forward. IEEE S&P's public meta reviews and USENIX Security's required ethics sections may serve similar purposes.

Security ethics education—through conference sessions, peer discussions, and ethics curricula. While review boards act as a checkbox bar to clear and rarely inspire deeper ethical conversations, these conversations do happen: within research teams, and during peer review. If the latter are not informally shared, this leaves out junior researchers not yet on program committees, and also researchers who for other reasons may not be on program committees. Based on participants' wishes for community ethics

conversations, we propose embedding "ethics post mortems" into security conferences, where these discussions can be had and attended by a larger share of the community. Similarly, conferences can provide consultation sessions with research ethics committee members with feedback and insights into ethics decision making.

The approach to ethics and the importance placed on it by researchers depended on their personal experiences. Many participants reported learning about ethics through their (past) advisors (see 4.3) and exert a similar influence on their own students. Researchers with less ethics advising may be less attuned to possible ethical issues in their research. Although efforts are being made [23], ethical education is not yet an established and standardized component of computer security education [23, 46], and we invite progress towards ethics in computer security curricula.

6 Conclusion

We explored researchers' perspectives on ethics in their research process, the publishing process and the security and privacy community through 24 semi-structured interviews and a meta-analysis of 1154 security research papers published in 2024.

We find that, while authors may deeply care about ethics, few considerations are reflected in published research; researchers struggle with how to consider ethics during research inception, execution, writing, peer review, and ethics review. We hope an informed understanding of how the security and privacy community has approached ethics in the past can help towards guidance for future challenging ethical situations. We look forward to exploring how new directions in encouraging authors to consider ethics (e.g., [75]) will shape the future of the research field.

Acknowledgments

We are grateful to the researchers for their participation and early feedback on this paper. We thank Dr. Wulf Loh, Prof. Dr. Dominik Wermke, Alex Cooper, Cora Sula, Sara Olschar, and Anna Lena Rotthaler for feedback and help with early drafts. This project was supported in part by NSF grant CNS-2206865. Tadayoshi Kohno is supported by the McDevitt Chair in Computer Science, Ethics, and Society at Georgetown University. The majority of this research was conducted while Tadayoshi Kohno was at the University of Washington. Tadayoshi Kohno had a chairing role with the USENIX Security 2025 and 2026 Research Ethics Committees and additionally thanks the conference chairs (Lujo Bauer and Giancarlo Pellegrino for 2025 and Ben Stock and Elissa M. Redmiles for 2026) and his REC co-chair for 2026 (Eleanor Birrell) for related discussions.

References

[1] Mamia Agbese, Rahul Mohanani, Arif Khan, and Pekka Abrahamsson. 2023. Implementing AI ethics: Making sense of the ethical requirements. *ACM Proceedings of the Evaluation and Assessment in Software Engineering Conference* 1, 1 (2023).

[2] Ross Anderson. 2012. Ethics Committees and IRBs: Boon, or Bane, or More Research Needed? *Financial Cryptography and Data Security* (2012), 133–135.

[3] Association for Computing Machinery. 2022. ACM Code of Ethics and Professional Conduct. https://www.acm.org/code-of-ethics.

[4] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo Report. *IEEE Security & Privacy Magazine* 10, 2 (2012), 71–75.

[5] Virginia Braun and Victoria Clarke. 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and psychotherapy research* (2021).

[6] Philip AE Brey. 2012. Anticipatory Ethics for Emerging Technologies. *NanoEthics* 6, 1 (2012), 1–13.

[7] Barry Brown, Alexandra Weilenmann, Donald McMillan, and Airi Lampinen. 2016. Five Provocations for Ethical HCI Research. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, Jofish Kaye, Allison Druin, Cliff Lampe, Dan Morris, and Juan Pablo Hourcade (Eds.). ACM, New York, NY, USA, 852–863.

[8] Noelle Brown, Benjamin Xie, Ella Sarder, Casey Fiesler, and Eliane S Wiese. 2024. Teaching ethics in computing: a systematic literature review of ACM computer science education publications. *ACM Transactions on Computing Education* 24, 1 (2024), 1–36.

[9] Amy Bruckman. 2020. 'Have you thought about. . . '. *Commun. ACM* 63, 9 (2020), 38–40. https://doi.org/10.1145/3377405

[10] Elizabeth Buchanan, John Aycock, Scott Dexter, David Dittrich, and Erin Hvizdak. 2011. Computer science security research and human subjects: emerging considerations for research ethics boards. *Journal of empirical research on human research ethics : JERHRE* 6, 2 (2011), 71–83.

[11] Sam Burnett and Nick Feamster. 2015. Encore: Lightweight measurement of web censorship with cross-origin requests. In *Proceedings of the 2015 ACM conference on special interest group on data communication*. 653–667.

[12] Kevin Butler and Kurt Thomas. 2022. Message from the USENIX Security '22 Program Co-Chairs. (2022). https://www.usenix.org/sites/default/files/sec22_message.pdf

[13] Monica Chin. 2021. How a university got itself banned from the Linux kernel. Accessed: 2024-04-09.

[14] Shruthi Sai Chivukula, Colin Gray, Ziqing Li, Anne C Pivonka, and Jingning Chen. 2021. Surveying a Landscape of Ethics-Focused Design Methods. *ACM Journal on Responsible Computing* (2021).

[15] Markus Christen, Bert Gordijn, and Michele Loi. 2020. *The ethics of cybersecurity*. Springer Nature.

[16] David B. Resnik. 2022. What Is Ethics in Research & Why Is It Important? https://www.niehs.nih.gov/research/resources/bioethics/whatis/index.cfm. Accessed: 2024-09-04.

[17] Alexandra Dirksen, Sebastian Giessler, Hendrik Erz, Martin Johns, and Tobias Fiebig. 2024. Don't Patch the Researcher, Patch the Game: A Systematic Approach for Responsible Research via Federated Ethics Boards. In *Proceedings of the New Security Paradigms Workshop*. 126–141.

[18] David Dittrich, Erin Kenneally, and Michael Bailey. 2013. Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report. *SSRN Electronic Journal* (2013).

[19] Kimberly Do, Rock Yuren Pang, Jiachen Jiang, and Katharina Reinecke. 2023. "That's important, but...": How Computer Science Researchers Anticipate Unintended Consequences of Their Research Innovations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM.

[20] Serge Egelman, Joseph Bonneau, Sonia Chiasson, David Dittrich, and Stuart Schechter. 2012. It's not stealing if you need it: A panel on the ethics of performing research using public data of illicit origin. In *Financial Cryptography and Data Security: FC 2012 Workshops, USEC and WECSR 2012, Kralendijk, Bonaire, March 2, 2012, Revised Selected Papers 16*. Springer, 124–132.

[21] Felix Anand Epp, Tim Moesgen, Antti Salovaara, Emmi Pouta, and İdil Gaziulusoy. 2022. Reinventing the wheel: The Future Ripples method for activating anticipatory capacities in innovation teams. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference*. 387–399.

[22] Shencha Fan, Jackson Sippe, Sakamoto San, Jade Sheffey, David Fifield, Amir Houmansadr, Elson Wedwards, and Eric Wustrow. 2025. Wallbleed: A Memory Disclosure Vulnerability in the Great Firewall of China. In *32nd Annual Network and Distributed System Security Symposium, NDSS 2025, San Diego, California, USA, February 24-28, 2025*. The Internet Society.

[23] Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. ACM Technical Symposium on Computer Science Education.

[24] Megan Finn and Katie Shilton. 2023. Ethics governance development: The case of the Menlo Report. *Social Studies of Science* 53, 3 (2023), 315–340.

[25] Ivan Flechais and George Chalhoub. 2023. Practical Cybersecurity Ethics: Mapping CyBOK to Ethical Concerns. In *Proceedings of the 2023 New Security Paradigms Workshop*. 62–75.

[26] Paul Formosa, Michael Wilson, and Deborah Richards. 2021. A principlist framework for cybersecurity ethics. *Computers & Security* 109 (2021), 102382.

[27] Batya Friedman, Peter H Kahn, Alan Borning, and Alina Huldtgren. 2013. Value sensitive design and information systems. *Early engagement and new technologies: Opening up the laboratory* (2013), 55–95.

[28] Simson L Garfinkel. 2008. IRBs and Security Research: Myths, Facts and Mission Creep. *UPSEC* 8 (2008), 1–5.

[29] Martyna Gliniecka. 2023. The ethics of publicly available data research: A situated ethics framework for Reddit. *Social Media+ Society* 9, 3 (2023).

[30] Colin M Gray and Shruthi Sai Chivukula. 2019. Ethical mediation in UX practice. In *Proceedings of the 2019 CHI conference on human factors in computing systems*.

[31] Florian Hantke, Sebastian Roth, Rafael Mrowczynski, Christine Utz, and Ben Stock. 2024. Where are the red lines? Towards Ethical Server-side Scans in

Security and Privacy Research. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 103–103.
[32] Brent Hecht, Lauren Wilcox, Jeffrey P. Bigham, Johannes Schöning, Ehsan Hoque, Jason Ernst, Yonatan Bisk, Luigi de Russis, Lana Yarosh, Bushra Anjum, Danish Contractor, and Cathy Wu. 2021. It's Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process.
[33] Jon Henshaw. 2021. *Princeton privacy study halts GDPR/CCPA research over ethics concerns and industry blowback*. Last updated December 27, 2021.
[34] Sarah Beth Hopton. 2021. The Tarot of Tech. *Equipping Technical Communicators for Social Justice Work: Theories, Methodologies, and Pedagogies* (2021), 158.
[35] Hsiu-Fang Hsieh and Sarah E Shannon. 2005. Three Approaches to Qualitative Content Analysis. *Qualitative health research* 15, 9 (2005), 1277–1288.
[36] Jina Huh-Yoo and Emilee Rader. 2020. It's the Wild, Wild West: Lessons Learned From IRB Members' Risk Perceptions Toward Digital Research Data. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–22.
[37] IEEE Security and Privacy. 2022. REC Annual Summary. https://docs.google.com/document/d/15x5Qd1UTaoMSRZgRRvurPbgRg4SWWLQKhG41ouYV0TY/edit. Accessed: 2024-06-06.
[38] IEEE Security and Privacy Symposium. 2022. Call for Papers. https://www.ieee-security.org/TC/SP2022/cfpapers.html. Accessed: 2024-09-04.
[39] IEEE Security and Privacy Symposium 2024. 2024. Call for Papers. https://sp2024.ieee-security.org/changes-cfp.html. Accessed: 05-06-2024.
[40] IEEE Symposium on Security and Privacy. 2021. Call for Papers. https://www.ieee-security.org/TC/SP2022/cfpapers.html.
[41] Institute of Electrical & Electronics Engineers. 2023. IEEE Code of Ethics. https://www.ieee.org/about/corporate/governance/p7-8.html. Accessed: 2023-06-03.
[42] Kai-Kristian Kemell, Ville Vakkuri, and Erika Halme. 2022. Utilizing user stories to bring AI ethics into practice in software engineering. In *International Conference on Product-Focused Software Process Improvement*. Springer, 553–558.
[43] Tadayoshi Kohno, Yasemin Acar, and Wulf Loh. 2023. Ethical Frameworks and Computer Security Trolley Problems: Foundations for Conversations. In *32nd USENIX Security Symposium (USENIX Security 23)*. 5145–5162.
[44] Daria Korobenko, Anastasija Nikiforova, and Rajesh Sharma. 2024. Towards a Privacy and Security-Aware Framework for Ethical AI: Guiding the Development and Assessment of AI Systems. In *Proceedings of the 25th Annual International Conference on Digital Government Research*. 740–753.
[45] Steven Levy. 2001. *Crypto: How the Code Rebels Beat the Government–Saving Privacy in the Digital Age*. Penguin.
[46] Kevin Macnish and Jeroen van der Ham. 2020. Ethics in cybersecurity research and practice. *Technology in Society* 63 (2020), 101382.
[47] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–14.
[48] Tina Marjanov and Alice Hutchings. 2025. SoK: Digging into the Digital Underworld of Stolen Data Markets. In *2025 IEEE Symposium on Security and Privacy*.
[49] Nora McDonald, Adegboyega Akinsiku, Jonathan Hunter-Cevera, Maria Sanchez, Kerrie Kephart, Mark Berczynski, and Helena M Mentis. 2022. Responsible computing: A Longitudinal Study of a Peer-led Ethics Learning Framework. *ACM Transactions on Computing Education (TOCE)* 22, 4 (2022), 1–21.
[50] Andrew McNamara, Justin Smith, and Emerson Murphy-Hill. 2018. Does ACM's Code of ethics change ethical decision making in software development? *(ESEC/FSE 2018)*. Association for Computing Machinery.
[51] Brent Mittelstadt. 2019. Principles alone cannot guarantee ethical AI. *Nature machine intelligence* 1, 11 (2019), 501–507.
[52] Giovane Moura and John Heidemann. 2023. Vulnerability Disclosure Considered Stressful. *ACM SIGCOMM Computer Communication Review* 53 (07 2023), 2–10.
[53] Kathleen O'Grady. 2025. *'Unethical' AI Research on Reddit Under Fire*. https://www.science.org/content/article/unethical-ai-research-reddit-under-fire Accessed: 2025-07-22.
[54] Rock Yuren Pang, Sebastin Santy, René Just, and Katharina Reinecke. 2024. BLIP: Facilitating the Exploration of Undesirable Consequences of Digital Technologies. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.
[55] Program Committee Chairs IEEE Symposium on Security and Privacy. 2021. Program Committee Statement Regarding the "Hypocrite Commits" Paper. https://www.ieee-security.org/TC/SP2021/downloads/2021_PC_Statement.pdf. Accessed: 05-06-2024.
[56] Emerson W. Pugh. 2009. Creating the IEEE Code of Ethics. In *2009 IEEE Conference on the History of Technical Societies*.
[57] Robert Ramirez, Shun Inagaki, Masaki Shimaoka, and Kenichi Magata. 2020. A cybersecurity research ethics decision support UI. *USENIX Association* (2020).
[58] Harshini Sri Ramulu, Helen Schmitt, Bogdan Rerich, Rachel Gonzalez Rodriguez, Tadayoshi Kohno, and Yasemin Acar. 2025. [Extended] Ethics in Computer Security Research: A Data-Driven Assessment of the Past, the Present, and the Possible Future. arXiv:2509.09351 [cs.CR] https://arxiv.org/abs/2509.09351
[59] Harshini Sri Ramulu, Helen Schmitt, Dominik Wermke, and Yasemin Acar. 2024. Security and Privacy Software Creators' Perspectives on Unintended Consequences. In *2024 33rd USENIX Security Symposium*.
[60] Paul Rehren and Hanno Sauer. 2024. Another brick in the wall? moral education, social learning, and moral progress. *Ethical Theory and Moral Practice* (2024).
[61] Dennis Reidsma, Jeroen van der Ham, and Andrea Continella. 2023. Operationalizing Cybersecurity Research Ethics Review: From principles and guidelines to practice. *Proceedings of the 2nd International Workshop on Binary Analysis Research* (2023).
[62] Wessel Reijers, David Wright, Philip Brey, Karsten Weber, Rowena Rodrigues, Declan O'Sullivan, and Bert Gordijn. 2017. Methods for practising ethics in Research and Innovation: A Literature Review, critical analysis and recommendations. *Science and Engineering Ethics* 24, 5 (09 2017), 1437–1481.
[63] Todd W. Rice. 2008. The historical, ethical, and legal background of human-subjects research. *Respiratory care* 53, 10 (2008), 1325–1329.
[64] Phillip Rogaway. 2015. The moral character of cryptographic work. *Cryptology ePrint Archive* (2015).
[65] Juliane Schmüser, Harshini Sri Ramulu, Noah Wöhler, Christian Stransky, Felix Bensmann, Dimitar Dimitrov, Sebastian Schellhammer, Dominik Wermke, Stefan Dietze, Yasemin Acar, et al. 2024. Analyzing Security and Privacy Advice During the 2022 Russian Invasion of Ukraine on Twitter. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
[66] Karen J. Schwenzer. 2011. Best practice & research in anaesthesiology issue on new approaches in clinical research ethics in clinical research. *Best practice & research. Clinical anaesthesiology* (2011).
[67] Hong Shen, Wesley H Deng, Aditi Chattopadhyay, Zhiwei Steven Wu, Xu Wang, and Haiyi Zhu. 2021. Value cards: An educational toolkit for teaching social impacts of machine learning through deliberation. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 850–861.
[68] Katie Shilton, Megan Finn, and Quinn DuPont. 2021. Shaping ethical computing cultures. *Commun. ACM* 64, 11 (2021), 26–29. https://doi.org/10.1145/3486639
[69] Lucy Simko, Ada Lerner, Samia Ibtasam, Franziska Roesner, and Tadayoshi Kohno. 2018. Computer security and privacy for refugees in the United States. In *2018 IEEE symposium on security and privacy (SP)*. IEEE, 409–423.
[70] Konstantinos Solomos, John Kristoff, Chris Kanich, and Jason Polakis. 2021. Tales of favicons and caches: Persistent tracking in modern browsers. In *Network and Distributed System Security Symposium*.
[71] Ananta Soneji, Faris Bugra Kokulu, Carlos Rubio-Medrano, Tiffany Bao, Ruoyu Wang, Yan Shoshitaishvili, and Adam Doupé. [n. d.]. "Flawed, but like democracy we don't have a better system": The Experts' Insights on the Peer Review Process of Evaluating Security Papers. In *2022 IEEE Symposium on Security and Privacy*.
[72] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. 2017. Ethical issues in research using datasets of illicit origin. In *Proceedings of the 2017 Internet Measurement Conference*. Association for Computing Machinery.
[73] Michelle Tran and Casey Fiesler. 2024. " It's Not Exactly Meant to Be Realistic": Student Perspectives on the Role of Ethics In Computing Group Projects. In *Proceedings of the 2024 ACM Conference on International Computing Education Research-Volume 1*. 517–526.
[74] U.S. Department of Homeland Security. 2012. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf
[75] USENIX Security. 2024. USENIX Security Ethics Guidelines. https://www.usenix.org/conference/usenixsecurity25/ethics-guidelines. Accessed: 2024-09-04.
[76] Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. 2017. Ethics regulation in social computing research: Examining the role of institutional review boards. *Journal of Empirical Research on Human Research Ethics* (2017).
[77] Miranda Wei, Sunny Consolvo, Patrick Gage Kelley, Tadayoshi Kohno, Tara Matthews, Sarah Meiklejohn, Franziska Roesner, Renee Shelby, Kurt Thomas, and Rebecca Umbach. 2024. Understanding Help-Seeking and Help-Giving on Social Media for Image-Based Sexual Abuse. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 4391–4408.
[78] Jess Whittlestone, Rune Nyrup, Anna Alexandrova, and Stephen Cave. 2019. The role and limits of principles in AI ethics: Towards a focus on tensions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 195–200.
[79] David Gray Widder, Derrick Zhen, Laura Dabbish, and James Herbsleb. 2023. It's about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them?. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 467–479.
[80] Emad Yaghmaei, Ibo van de Poel, Markus Christen, Bert Gordijn, Nadine Kleine, Michele Loi, Gwenyth Morgan, and Karsten Weber. 2017. Canvas white paper 1–cybersecurity and ethics. *Available at SSRN 3091909* (2017).
[81] Mojtaba Zaheri, Yossi Oren, and Reza Curtmola. 2022. Targeted deanonymization via the cache side channel: Attacks and defenses. In *31st USENIX Security Symposium (USENIX Security 22)*. 1505–1523.
[82] Yiming Zhang, Mingxuan Liu, Mingming Zhang, Chaoyi Lu, and Haixin Duan. 2022. Ethics in Security Research: Visions, Reality, and Paths Forward. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*.