Cybersecurity has become inseparable from everyday life; from hospitals and universities to the infrastructures that power science and government. Yet despite massive investments in automation, analytics, and defense technologies, major security failures continue to occur. These failures rarely result from missing tools alone; they emerge when people, processes, and technologies fall out of alignment or when the people responsible for security are constrained by the very systems meant to support them.

My research highlights these gaps through a human-centered lens. I ask two complementary questions: *(1) How do people shape security from the bottom up, through their expertise, collaboration, and everyday interactions? and (2) How does security shape people from the top down, through the infrastructures and policies that govern their work?*

I study cybersecurity as a socio-technical system that integrates human reasoning, organizational processes, and computational infrastructure. Methodologically, my work blends social science and security practice by using qualitative and mixed-methods techniques including interviews, surveys, content analysis, and process tracing.

Together, these questions structure my research across two levels, *practice and structure*, linking human reasoning with institutional security norms and structure:

At the practice level, my work examines how humans construct and sustain resilience across security operations, research infrastructures, and policy environments. My studies of Security Operations Centers [1], hacker reasoning [2], Research Computing Infrastructures [3], and the implementation of Zero Trust principles (under review) reveal how organizational goals, mental models, and culture jointly shape the effectiveness of technical defense.

At the structural level, my research investigates how governance and accountability operate within complex ecosystems such as higher education and cybersecurity peer review system. My studies on educational technology procurement and adoption [4, 5] expose how risk and responsibility are distributed across institutions and vendors, while my meta-research on security peer review [6]) examines how the field adapts its reviewing processes to the growing scale of submissions and evolving expectations of rigor, transparency, and sustainability.

My research collectively underscores *security as alignment work*, the continuous process of reconciling human reasoning, organizational practices, and institutional systems. This human-centered perspective informs practical improvements within critical systems and infrastructures, strengthening organizational resilience, supporting the people who rely on them, and advancing the public good.

Main Topic 1: How do People Shape Security in Practice?

Humans form a critical layer within organizational cybersecurity. They may be analysts triaging and escalating incidents, managers prioritizing security investments and overseeing operations, white-hat hackers experimenting with new exploits, researchers collaborating across institutions, or leaders translating policy into practice. Each plays a distinct yet interconnected role to ensure overall security posture. My work in this area examines how resilience is constructed and sometimes compromised through human and organizational practices at different layers.

At the operational layer, our CCS 2019 paper [1] examined *how Security Operations Centers (SOCs) function from the perspective of the people who run them: analysts and managers.* Despite SOCs being the core of enterprise defense, serious incidents continue to occur even in well-resourced organizations. While prior work has focused on improving technology, less is known about what SOC professionals perceive as issues resulting into inefficiencies. To fill this gap, we turned our investigative lens to the everyday human factors of SOC operations. Using iterative open coding to analyze the interview data of 18 SOC members, we found that managers and analysts valued certain key aspects of SOC operations differently: analysts emphasized visibility and log quality, whereas managers prioritized automation, response speed, and performance metrics. These mismatched priorities were sometimes sharply opposed. The study reframes SOC performance as a problem of differing mental models and shows how analysts and managers value and make sense of the same systems in divergent ways.

At the individual level, our IEEE S&P 2024 paper [2] examined *how security expertise unfolds in practice; specifically, how hackers make and recover from mistakes while solving binary exploitation challenges.* Although exploitation is typically an offensive activity, understanding it is essential for improving security education and preparedness. We hypothesized that much of this effort stems from the mistakes hackers while exploiting the targets; understanding those matter because they reveal how experts reason through uncertainty and adapt when things go wrong. To observe authentic real-world hacking behavior, we analyzed 30 screencasts from 11 hackers uploaded to YouTube. Mistakes consumed nearly 40% of total work time; most stemming from cognitive (memory lapses and missteps) rather than technical causes. Yet these errors also revealed adaptive expertise: hackers debugged collaboratively and adjusted strategies. The findings highlight that learning to recognize and recover from mistakes is central to developing expertise, suggesting practical opportunities for hacker education and debugging support.

At the institutional level, our IEEE S&P 2025 work [3] examined how research institutions manage security and collaboration within Research Computing Infrastructures (RCIs)—the high-performance systems that now underpin data-driven science and enable large-scale, collaborative research across disciplines and institutions. Ensuring their security extends beyond technical safeguards to the people who coordinate access and share resources across institutions. Through interviews with 24 researchers and system administrators across 12 institutions, we found that

everyday research often relies on informal, trust-based sharing practices (temporary access, shared credentials, and inherited permissions) that keep projects agile but create lasting security exposures. Administrators described difficulty balancing openness and compliance, while researchers viewed rigid controls as barriers to productivity. Our findings highlight that securing RCIs requires designing systems that make security a shared and sustainable practice across technical and human boundaries, revealing that *trust, while convenient, is flawed and a vulnerability in itself.*

At the strategic level, in my ongoing research (under review), I explore how Zero Trust principles are being interpreted and implemented across academic, industry, and government organizations. Across the broader cybersecurity community, the move away from implicit trust has crystallized in the adoption of Zero Trust architectures; an approach that redefines security by assuming no implicit trust and ensures continuous verification. Through interviews with 27 cybersecurity professionals, we found that while Zero Trust has become a policy mandate for government organizations, its meaning and application vary widely across sectors. For some, it represents a genuine rethinking of security culture; for others, it is a branding exercise driven by vendors. Many grapple with how to reconcile its ideals with existing organizational norms, legacy systems, and human workflows. Rather than treating Zero Trust as a purely technical or policy directive, my work examines how its implementation interacts with organizational culture, existing infrastructures, and human constraints and what it takes to make "continuous verification" viable in practice. This work extends the notion of resilience beyond defense and recovery, toward the capacity to continuously evolve security culture in response to shifting technologies, mandates, and threats.

Main Topic 2: How does Security Shape People and Institutions?

At the structural level, my work examines how organizations govern security and privacy, particularly within education sector. Over the last decade, education has become a major target for cyberattacks,[1] driven by the combination of valuable personal and research data, under-resourced security infrastructures, and a large, decentralized attack surface. In parallel, classrooms and campuses have become increasingly dependent on interconnected educational technologies (EdTechs), for teaching, learning, and administration-related activities, that collect and exchange vast amounts of student and institutional data. This dependence has widened the threat landscape, making governance both essential and precarious. An investigation of those governance decisions can help understand how institutions uphold data privacy, regulatory compliance, and end-user trust in an era of pervasive datafication.[2]

To investigate how such decisions are made in practice, our CCS 2024 work [4] examined the *EdTech acquisition process* through interviews with people in leadership positions for learning enterprise across seven U.S. universities. We found that procurement decisions were driven by pedagogical, accessibility, and cost considerations, not the security and safety of vendor tools. Though regulations such as FERPA and HIPAA help education institutions establish baseline data protections, ambiguities surrounding data breach liability limit their ability to hold vendors accountable. Prior to acquisition, risk assessment happens yet depends heavily on vendor self-attestations. Once adopted, these tools became deeply embedded in institutional networks with little visibility into data exchanges among vendors and their sub-vendors. While contracts can restrict vendors' use of institutional data, the lack of visibility leaves HEIs unaware of potential misuse until after the fact, creating a governance structure that exists but struggles to act. Offboarding processes are similarly fraught with uncertainty: institutions rarely have the power or technical means to verify data deletion. This raised a troubling question: if institutions already struggle to govern the tools they formally acquire, *what happens when technologies enter classrooms through informal channels?*

Building on this, a follow-up study [5] surveyed 375 educators across U.S. K-12 and higher education institutions to investigate the landscape of tools adopted by instructors outside formal procurement. Educators reported 494 distinct tools, chosen largely for instructional utility and student engagement rather than its data protection and security features, and many described knowingly bypassing institutional policy to better support their classrooms. These choices, while well-intentioned, create invisible data flows between unvetted vendors, expanding exposure to breaches and non-compliance with FERPA and HIPAA. Together, these studies reveal that governance gaps in education are not isolated policy failures but systemic conditions embedded in everyday institutional practice.

Main Topic 3: Meta Research and Science of Security

Human-factors research strengthens how people perceive and use security systems. Hence, another dimension of my structural arm focuses *inward* on how the security community itself produces, evaluates, and sustains its own knowledge. My IEEE S&P 2022 paper [6] on the cybersecurity peer-review process was the first to investigate how the security research community decides what counts as good science and how those decisions shape what becomes accepted knowledge. Through interviews with 21 program committee members and chairs across top-tier conferences, we examined how reviewers interpret core evaluation criteria, assume reviewing responsibilities, and define review quality. We found that novelty was universally valued; however, its meaning varied widely—from methodological innovation to problem selection to just the subjective "wow" factor—revealing the lack of shared standards for evaluating research.

[1] https://www.highereddive.com/news/data-breaches-cost-higher-education-colleges/689499/
[2] https://wonkhe.com/blogs/universities-continue-to-struggle-with-datafication/

Reviewers provided richer justifications for rejection than for acceptance, fueling anecdotal perceptions of "randomness" in outcomes. Efforts by authors to game the system added further complexity into this notion which resulted into extra reviewing load while authors try to *get lucky*. These dynamics were further complicated by systemic pressures such as record submission volumes over rolling deadlines and overlapping service-related commitments. By empirically examining the peer-review process, this study helped scientifically and systematically voice a range of long-standing community concerns, transforming anecdotal frustrations into documented evidence that the field could analyze and build upon. Moreover, many advisors include this paper in graduate courses and lab reading groups as an accessible introduction to how scientific cybersecurity research and its peer-review process operate; a system that is often a black box to students entering the field.

Since this study, subsequent work in the security community, including investigations at IEEE S&P, USENIX Security, ACM CCS, and ACM REP, has begun to explore related aspects of the review process, such as transparency in usable security studies [7, 8], ethical oversight [9], and artifact evaluation [10]. This growing body of research reflects a broader recognition that understanding the *science of security* itself is essential to its progress.

Other Research

Alongside my primary work on organizational security, governance, and meta-research, my work on the adoption of OnlyFans [11] examined why creators without prior experience in sex work joined the platform and how its affordances such as boundary setting, privacy, and self-governance, enabled agency within stigmatized digital labor. My USENIX Security 2024 paper [12] on digital safety and privacy investigated how OnlyFans creators experience and mitigate risks such as harassment, censorship, and deplatforming, revealing how informal security practices evolve in response to stigma, prominence, and uncertainties in platform policies. Together, these studies extend my broader research on security governance to the context of digital platforms where individuals, rather than institutions, bear the responsibility of managing risk and trust.

Future Work

In the near term, I plan to deepen my human-centered cybersecurity research across three directions. First, I will study transparency in security research, examining how new artifact-evaluation practices shape perceptions of rigor, reproducibility, and trust in published results. Second, I will investigate human–process alignment in incident response, using multi-modal analyses to understand how analysts coordinate across detection, response, and recovery, and how reflective practices can improve preparedness. Third, I will explore the lived consequences of educational data breaches, analyzing user experiences and online discussions to uncover how affected communities interpret, cope with, and rebuild trust after institutional failures.

Long term goals: The rapid rise of artificial intelligence (AI) is reshaping the cybersecurity landscape—fueling a digital arms race in which both attackers and defenders harness the same technologies. This convergence magnifies a central paradox: *to strengthen security, we must trust AI systems to act intelligently, yet also mistrust them enough to retain control, accountability, and human judgment.* Modern cybersecurity paradigms such as *Zero Trust* already aim to minimize implicit trust by continuously validating identities and access requests. However, as organizations adopt autonomous AI agents that learn, reason, and act across boundaries, the assumptions underlying these architectures begin to strain. Existing models of verification and access control presume static identities, predictable behaviors, and clearly defined system boundaries; all of which dissolve when AI systems can reinterpret goals, delegate authority, and make context-dependent decisions.

At the same time, humans across different skills, expertise, and organizational roles must learn to collaborate with systems whose reasoning is partially opaque and whose speed outpaces deliberate reflection. This combination of *technical autonomy* and *human cognitive vulnerability* introduces new risks: misplaced trust, automation bias, accountability gaps, emotional strain, and fatigue.

These tensions define the space my research seeks to address: *how trust in AI can be systematically designed, calibrated, and governed in cybersecurity operations.* I begin by viewing access control not as a fixed gate-keeping mechanism but as part of a broader trust governance architecture. And in the long-term, my investigation aims to answer following questions: How can systems quantify and recalibrate trust as AI behavior evolves? What mechanisms allow agents to earn, lose, or regain privilege based on performance and context? How do executives, policymakers, employees, and end-users align expectations of AI accountability and acceptable autonomy? What governance structures foster transparency and shared responsibility without constraining innovation? How do cognitive load, fatigue, and stress shape security decision-making in AI-assisted operations? Can human well-being and mental resilience become integral components of cyber defense posture?

By examining multi-stakeholder aspects including interactions, trust, and emotional burden, I aim to make the human layer more resilient, and by extension, strengthen the resilience of organizations and nations. This reflects my broader research ethos: to redefine *the human not as the weakest link, but as an indispensable part of the solution.*

Collaborative Interests & Funding Potential: My future research plans naturally support collaborations across cybersecurity, organizational defenses, AI, education, public policy, and research on critical infrastructures, and I look forward to contributing to cross-departmental and institute-level initiatives. My vision aligns well with funding priorities at NSF, particularly Security, Privacy, and Trust in Cyberspace (SaTC 2.0), DARPA programs on human–AI collaboration, and other initiatives supporting socio-technical and organizational cybersecurity research. These directions position my research vision well for multi-PI collaborations and sustained external support.

References

[1] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. Matched and mismatched socs: A qualitative study on security operations center issues. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 1955–1970, 2019.

[2] Irina Ford, Ananta Soneji, Faris Bugra Kokulu, Jayakrishna Vadayath, Zion Leonahenahe Basque, Gaurav Vipat, Adam Doupé, Ruoyu Wang, Gail-Joon Ahn, Tiffany Bao, et al. "watching over the shoulder of a professional": Why hackers make mistakes and how they fix them. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 350–368. IEEE, 2024.

[3] Souradip Nath, Ananta Soneji, Jaejong Baek, Tiffany Bao, Adam Doupé, Carlos Rubio-Medrano, and Gail-Joon Ahn. "It's almost like Frankenstein"': Investigating the Complexities of Scientific Collaboration and Privilege Management within Research Computing Infrastructures. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 2995–3013. IEEE Computer Society, 2025.

[4] Easton Kelso, Ananta Soneji, Sazzadur Rahaman, Yan Shoshitaishvili, and Rakibul Hasan. Trust, Because You Can't Verify: Privacy and Security Hurdles in Education Technology Acquisition Practices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1656–1670, 2024.

[5] Easton Kelso, Ananta Soneji, Syed Zami-Ul-Haque Navid, Yan Shoshitaishvili, Sazzadur Rahaman, and Rakibul Hasan. Investigating the Security & Privacy Risks from Unsanctioned Technology Use by Educators. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2025.

[6] Ananta Soneji, Faris Bugra Kokulu, Carlos Rubio-Medrano, Tiffany Bao, Ruoyu Wang, Yan Shoshitaishvili, and Adam Doupé. "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 (SP)*, pages 1845–1862. IEEE, 2022.

[7] Jan H Klemmer, Juliane Schmüser, Byron M Lowens, Fabian Fischer, Lea Schmüser, Florian Schaub, and Sascha Fahl. Transparency in usable privacy and security research: Scholars' perspectives, practices, and recommendations. In *In 46th IEEE Symposium on Security and Privacy*, 2025.

[8] Jan H Klemmer, Juliane Schmüser, Fabian Fischer, Jacques Suray, Jan-Ulrich Holtgrave, Simon Lenau, Byron M Lowens, Florian Schaub, and Sascha Fahl. How transparent is usable privacy and security research? a {Meta-Study} on current research transparency practices. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 5967–5986, 2025.

[9] Harshini Sri Ramulu, Helen Schmitt, Bogdan Rerich, Rachel Gonzalez Rodriguez, Tadayoshi Kohno, and Yasemin Acar. [extended] ethics in computer security research: A data-driven assessment of the past, the present, and the possible future. *arXiv preprint arXiv:2509.09351*, 2025.

[10] Daniel Olszewski, Allison Lu, Anna Crowder, Nathaniel Bennett, Seth Layton, Sri Hrushikesh Varma Bhupathiraju, Tyler Tucker, Siddhant Kalgutkar, Hunter Ver Helst, Carson Stillman, et al. Reproducibility in applied security conferences: An 11-year review on artifacts and evaluation committees. In *Proceedings of the 3rd ACM Conference on Reproducibility and Replicability*, pages 96–107, 2025.

[11] Vaughn Hamilton, Ananta Soneji, Allison Mcdonald, and Elissa M Redmiles. "Nudes? Shouldn't I charge for these?": Motivations of New Sexual Content Creators on OnlyFans. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, pages 1–14, 2023.

[12] Ananta Soneji, Vaughn Hamilton, Adam Doupé, Allison McDonald, and Elissa M Redmiles. "I feel physically safe but not politically safe"': Understanding the Digital Threats and Safety Practices of {OnlyFans} Creators. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 1–18, 2024.