

Cyber-security Research Ethics Dialogue & Strategy Workshop

Erin Kenneally
CAIDA/UC, San Diego
La Jolla, CA, USA
erin@caida.org

Michael Bailey
University of Michigan
Ann Arbor, MI, USA
mibailey@umich.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The inaugural Cyber-security Research Ethics Dialogue & Strategy Workshop was held on May 23, 2013, in conjunction with the IEEE Security Privacy Symposium in San Francisco, California. CREDS embraced the theme of "ethics-by-design" in the context of cyber security research, and aimed to:

- Educate participants about underlying ethics principles and applications;
- Discuss ethical frameworks and how they are applied across the various stakeholders and respective communities who are involved;
- Impart recommendations about how ethical frameworks can be used to inform policymakers in evaluating the ethical underpinning of critical policy decisions;
- Explore cyber security research ethics techniques, tools, standards and practices so researchers can apply ethical principles within their research methodologies; and
- Discuss specific case vignettes and explore the ethical implications of common research acts and omissions.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; K.4.1 [Public Policy Issues]: EthicsPrivacyRegulationUse/abuse of Power

General Terms

Management, Measurement, Documentation, Design, Reliability, Experimentation, Security, Human Factors, Theory, Legal Aspects.

Keywords

Ethics, Law, Trust, Cyber security, Network measurement

1. INTRODUCTION

The future of online trust, innovation and self-regulation is threatened by a widening gap between users' expectations informed by laws and norms, and new capacities for benefits and harms generated by technological advances. As this gap widens so too does ambiguity between asserted rights, interests, and threats. As a result society perceives heightened tensions and risks when engaging

the Web. How do we narrow this gap and thereby lower risks of actions online in a manner that instills trust, safeguards autonomy, and promotes ingenuity? One part of this solution is to embrace the fundamental principles of ethics to guide our decisions in the midst of information uncertainty.

One context where this solution is germinating is cyber security research. Commercial and academic researchers and policymakers are tackling novel ethical challenges that exert a strong influence on online trust. These challenges are not exceptional, but increasingly the norm. For example: (i) to recognize significant Internet threats and develop effective defenses researchers infiltrate malicious botnets; (ii) to understand Internet fraud (phishing) studies researchers surreptitiously observe users in order to ascertain typical behaviors; and (iii) to empirically measure network usage and characteristics researchers require access to users' nonpublic traffic.

These research activities are prerequisite for evidence-based policymaking that impacts us individually and collectively, such as infrastructure security, cyber crime, network neutrality, free market competition, spectrum application and broadband deployment, censorship, technology transfer, and intellectual property rights. Therefore, in the wake of struggles to resolve the aforementioned mounting tensions, ethics has re-emerged as a crucial ordering force. For this reason, ethics underpins the debate among cyber security researchers, oversight entities, industrial organizations, the government and end users about the acceptability of Internet research activities.

Motivated by this context, the Cyber security Research Ethics Dialogue Strategy (CREDS)¹ workshop embraced the theme of "ethics-by-design", and aimed to:

- Educate participants about underlying ethics principles and applications;
- Discuss ethical frameworks and how they are applied across the various stakeholders and respective communities who are involved;
- Impart recommendations about how ethical frameworks can be used to inform policymakers in evaluating the ethical underpinning of critical policy decisions;
- Explore cyber security research ethics techniques, tools, standards and practices so researchers can apply ethical principles within their research methodologies; and
- Discuss specific case vignettes and explore the ethical implications of common research acts and omissions.

2. THEMATIC DISCUSSIONS

Our goal was to create a set of targeted discussions among relevant stakeholders whose actions impact cyber security research

¹<http://www.caida.org/workshops/creds/1305/>

ethics policy and practice, rather than to conduct a peer reviewed mini-conference. Therefore, submissions in the form of position statements or research experiences were organized into panels according to thematic similarity. Group discussions were anchored around the following three meta themes, which were kicked off by brief panel statements derived from the submissions.

In terms of stakeholder demographics the workshop participants were representative of primarily the researcher community, with a smaller proportion comprising oversight entities and policymakers.

2.1 Brave New World - Ethical Research Amidst Expanding Opportunities

The presentations and discussions for this theme focused on the lines being drawn between ethical and unethical research in the Information and Communication Technology Research (ICTR) community.

In their paper, *Welcome to the World of Human Rights: Please Make Yourself Uncomfortable*, Henry Corrigan-Gibbs posited that considering Internet access as a human right does not sufficiently address all of the ethical questions around anonymity, firewalling, and censorship-circumvention research. He suggested that computer scientists might look to the humanitarian aid community for guidance on how to address and reason about the ethical dilemmas that arise when promoting human rights throughout the global Internet. Whereas Gibbs proposed a humanitarian model to guide ethical actions in the context of censorship, ensuing discussion highlighted the broader issue of the proper role of researchers as interventionists in social, civil, economic, political policy areas. For example, how valid is the analogy between humanitarian aid and Internet access? How far can such analogies be applied in cyber security? Are bot herders the equivalent of despotic regimes, and, if so, what rights should they be afforded? What about when researchers are studying criminal underground and must engage with nefarious actors to study activities such as trading currencies, drugs, and credentials? What role do researchers play when the areas being studied may also be of interest to law enforcement?

Another paper following this theme was presented by Sebastian Schrittwieser: *Ethics in Security Research: Which lines should not be crossed?* Sebastian focused on how we might ensure that research activities do not harm others by trying to motivate “a discussion on how research activities in the field of information security can be evaluated from an ethical point of view and how, we as a community, can establish ethical standards similar to other sciences such as medical research.” This work hit upon the need for ethical guidelines or frameworks based on common principles using four controversial research cases. There is a recurring call for researchers to articulate how their actions and the underlying principles relate to existing ethical framework advanced by the Belmont Report² or the IEEE Code of Ethics³. There is a related need to understand what frameworks exist and what are the pros and cons of the frameworks researchers are currently aware of.

The final paper in this theme involved John Aycock discussing *Why “No Worse Off” is Worse Off*, wherein he examined the various ethical justifications made by researchers in the context of infiltrating the activities of miscreants, criminals, and organized crime—a domain traditionally under the exclusive purview of law enforcement. One prominent discussion point that arose, in addition to the aforementioned role of researchers as interventionists in computer security and law enforcement, was whether novel or controversial computer security research (e.g., botnets and phishing) is

increasingly justified by a utilitarian bias in the community. This spurred questions as to whether it is reasonable to work toward a unifying philosophy, and how do we move forward in reaching agreeable actions amidst fractured mindsets (e.g., utilitarian v. deontological ethical approaches)? From a more pragmatic perspective, how do we identify what entity should provide oversight and set standards? Finally, do the challenges posed by ICTR demand changes in methodologies or does this research context demand changes in expectations regarding how to apply ethical principles (e.g., respect for persons)?

2.1.1 Discussion

The overall participant discussion on this theme highlighted more specific examples and research experiences that were helpful to ground the larger concept. For instance:

- One analogy proposed that taking over an illegal botnet is similar to invading a criminal’s home. This was challenged and countered with the proposition that the better analogy for reasoning purposes is more like the criminal breaking into your home and then law enforcement breaks into the criminal’s homes.
- Specific to the UCSB case study mentioned by Schrittwieser, it was noted that researchers obtained access to command and control, took all the credit cards, and repatriated them to the banks from where they were issued. The claim then, from an ethical standpoint, was that any fraud that would have been perpetrated was actually prevented by the researchers – in essence demonstrating that there was affirmative action to reduce financial harm. Significantly, this point was understated in the paper.
- Institutional oversight (e.g., Institutional Review Boards) and legal compliance do not necessarily identify or resolve ethical issues.
- All of the case studies mentioned have a lot of subtleties. For example, most of the Botnet researchers allow themselves to be infected so that they can study the data, raising the question as to whether that poses ethical issues in and of itself?
- There is no single “right” ethical approach. We need to come to some collective understanding of the common principles and their application in cyber security research. One participant was unconvinced that a universal framework is possible, expressing skepticism that we can actually articulate it in a standard way. Rather, the various nuances of fact-specific cases relegate the approach to that of “I-know-it-when-I-see-it”.
- The role of researchers is not to be the judge and jury when they encounter purported “bad guys”, leading to the question of how can Ethical Research Boards (ERB) facilitate these types of challenges for researchers? Research papers that explicitly or implicitly address these issues of first impression will undoubtedly serve as exemplars for ethical standards, so we must be cautious about establishing norms in this way. There are gaps and loopholes between the ERBs’ mandated responsibilities (i.e., protecting human research subjects) and what they are capable of providing in terms of technical expertise and advisement regarding larger ethical issues. Who fills those gaps and how? Similarly, do researchers have a duty to act as proverbial Good Samaritans when they discover compromised users or fruits of malicious acts that would ease the harm that arises from bad actors?
- When encountering challenging ethical issues there should be more acknowledgements about how we can explore eth-

²<http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html>

³<http://www.ieee.org/about/corporate/governance/p7-8.html>

ical alternatives as opposed to outright abandoning the research. In other words, the path forward consists of reasoned decision-making along a spectrum of options rather than assuming a binary, take it or leave it approach to risky research.

- One co-Chair noted that the community’s challenges are not unlike what we see in the realm of cyber crime and technology governance in general. Researchers can wield the same tool to both serve beneficial research purposes as well as to take advantage of the new capabilities. For example, anonymity works both ways—it can shield bad actors and “bad researchers.” A fundamental question, then, is whether it is the researcher’s role to decide who gets to use research results? If not, who is in a better position?
- The point was made that we have far to go in developing performance measurements regarding stated researcher benefits or impact metrics.
- Another participant questioned how much of the issue space regarding ethics principles is awareness-raising or education versus needing better fitting ethics models?

2.2 Checking Our Collective Assumptions—Risks and Benefits at the Frontline of ICT Research.

The second theme revolved around evaluating and balancing risks and benefits, and in general, how to find common ground.

Tadayoshi Kohno and Stefan Savage kicked off the topic with a foray into *Vulnerability Research in the Cyber-Physical World*. They described some real scenarios, their own reasoning in evaluating the ethical concerns of the scenarios, and how this work has affected the way in which they conduct, describe, and publish work related to cyber-physical embedded systems. Their testimonial describing cyber physical research vulnerability disclosure upended two assumptions underlying the familiar responsible full disclosure debate: (i) it distinguished scope of harm and (ii) ability to minimize harm. As such, it raised questions such as whether responsible disclosure is an issue of degree or of kind. Is there is a distinction between responsible disclosure in cyber-physical research compared and more traditional software systems vulnerability research?

Mark Allman supplemented this theme with his position paper on *Traffic Monitoring Considered Reasonable* wherein he contended that network traffic monitoring research activities fit well within the networking and security research community’s accepted behavioral norms. The proposition is that of capturing norms of acceptable behavior and developing more nuanced characterization of types of research that should have rebuttable presumptions about ethical propriety. This begs the question of whether there are “ethics-free” types of research, as well as whether, in such broad categories of research, the community can or should self govern?

Jean Camp questioned whether participating in privacy research benefits the participant, and if so, under what conditions in her paper *I Just Want Your Anonymized Contacts! Benefits and Education in Security Privacy Research*. This topic hit upon a range of recurring issues such as how to measure risks and benefits to privacy research subjects, what is the educational value that can be harnessed from such studies, and consistency of review. Her presentation explored the interesting question of how does a subject-oriented philosophy toward research benefits differ from a societal-focused one.

2.2.1 Discussion

The overall participant discussion on this theme was as follows:

- There are nontrivial difficulties with following the traditional responsible disclosure model in the context of interventions

at scale, especially where researchers are dealing with a large diverse population. There is a need for greater adoption of security and privacy practices with regarding to cyber-physical embedded technologies.

- As with the earlier discussions related to the proper role of researchers, it is not clear when and where it is ethically-appropriate to monitor and observe vulnerabilities and harmful actions, or to take some level of action to reduce, mitigate, or preempt further harm. This decision is confused when it is difficult to identify the impacts of one’s actions. So, in order to know that you can do more than just avoiding harm and actually reduce expected harm, researchers need to have more confidence in the consequential impact of their actions.
- We need to better document standards, but standards vary depending on the community and context (e.g. Anthropology or Journalism), two factors that are not necessarily amenable to clear delineation here.
- One participant astutely asked whether we are skating to the puck or to where we think it will be? In other words, should we not think about how the context (e.g., technology capabilities) will be different in the future, and how that will impact our ethical reasoning?
- We should be cautious when delving into “degrees of ethicalness” of activities or casting judgment. Referencing slavery as a metaphor, just because it has been done does not mean it is an appropriate or desirable model to follow.
- There is a trade-off between transparency of methodology and data reliability, and effective mitigation of harm in the cyber physical research.
- A recurrent theme was the relatively poor performance of researchers in disclosing their ethically defensible decision-making in concert with their published research.
- Similarly, researchers need to better articulate research benefits in the context of community and individual norms. A feedback loop would be quite helpful to preempt inaccurate conclusions and assumptions about the risks taken by researchers in terms of how they can be justified by benefits to individuals and society.
- Ethics are not a strict liability regime, but risk-based. Our approach to developing and socializing ethics norms is not about a priori, binary, right versus wrong action, but about agreeing on a common framework to reach those decisions.

2.3 Teaching Researchers to Fish - Tools to Implement Ethics Principles and Applications

The presentations and discussions for the third theme explored techniques, tools, standards, and practices to facilitate the application of ethical principles in practice.

Stuart Schechter introduced an approach focused on *Reusable Ethics-Compliance Infrastructure for Human Subjects Research*. He contended that shared tooling may assist in three different research functions related to ethical compliance: obtaining informed consent, debriefing, and the surveying of surrogate participants when consent cannot be obtained from actual participants. This spurred some interesting questions such as what are potential drawbacks to commoditizing ethics compliance tools in the context of informed consent, debriefing, surrogate subjects? What entities are in the best position to design, implement, and assess such a tool?

Conducting Ethical yet Realistic Usable Security Studies was

the focus for Ronen Margulis, wherein he contended that usable security research necessarily focuses on the behavior of the user that poses challenges by way of ensuring realistic scenarios and user behaviors, and ensuring ethical conduct. One prominent application space where this challenge is manifested is when conducting deception research. Specifically, is the knowledge produced from such deception studies influencing effective approaches to detection, prevention and/or response to phishing-type attacks? And, addressing the question of how is that currently measured raises a more fundamental question about whether there is a gap between knowledge creation and transfer to impactful security tools/public awareness.

Rula Sayaf rounded out this theme by tackling the question, *Can Users Control their Data in Social Software? An Ethical Analysis of Control Systems*. She explored whether full control is needed and should be granted in the context of social software users including the related ethical issues. She posited a data control approach (access control model) to privacy legal issues that argues for a balance between policy and technical mechanisms to respect user privacy. It prompts the question as to whether ACM has increasingly less relevance for facilitating respect for persons in the online ICTR context of big data analytics and resulting inference risk. Participant discussion was very limited at the end of this panel.

3. PATHS FORWARD

The workshop's capstone theme catalyzed dialogue around the question, *Who's Driving the Train?* where workshop chairs encouraged discussions about the shifting roles, responsibilities, and relationships between Researchers, Ethics Review Boards (ERB), Government, Professional Societies, and Program Committees in incentivizing and overseeing ethical research. Specifically, the organizers posited that if building a more effective research ethics culture is a prerequisite for balancing research innovation (i.e., academic freedom, reduced burdens, and ambiguities) with public trust (i.e., respect for privacy and confidentiality, accountability, data quality), we should focus on the pillars of such a culture and what strategies might be adopted to incorporate them into research operations. To that end we explored:

1. What leadership should be engaged (i.e., institutional, government, peer groups), and what should their respective roles and responsibilities be?
2. What education and awareness is needed?
3. What information sharing/coordination needs to be improved: between researchers, between oversight entities, and between researchers and oversight entities?
4. What knowledge and technology-transfer mechanisms can meet stated needs?

Faced with such far-reaching questions, unsurprisingly the discussion touched upon only the first two pillars during the course of the several hour allotted time slot. Regarding the leadership pillar, there was explicit skepticism about the extent to which government or ERBs can lead regarding the stated proposition. There was, however, more support for the likelihood that peer groups, like conference program committees, can influence and/or lead.

There was general backing for the strategy that attempts to educate and raise awareness of information and communications technology (ICT) research ethics issues should: target all stakeholder oversight entities rather than take a serial approach, be open to the variable appetites toward knowledge transfer among such overseers, and avoid taking an "all or nothing" strategy. There was a related thread suggesting that perhaps what is needed is a new oversight entity to fill the gap created by ERB's relatively poor comprehension of ICT issues. However, it was contended that we should find a way to educate them rather than look to replace them. Further, because cyber security research spans multiple domains (e.g., computer science, engineering, computer-human interface), it would be more effective and prudent to target ERBs rather than domain-specific program committees.

In support of the peer-centric approach, discussion turned to incentive mechanisms such as leveraging program committees to require in its call for papers that authors include an ethics statement. This approach may be deficient insofar as the reviewing PC members would be called to proxy community norms in a context where ethical norms are embryonic and heterogeneous. Although authoritative guidance upon which formal and informal norms are built certainly exist, there is still an interpretation and/or application gap in computer security research. Furthermore, PC members are part of this emergent process themselves, and likely lacking in ethics subject matter expertise. Designing the process so that paper vetting involves a larger conversation amongst the entire PC and/or is informed by appropriate expertise might both better represent and inform broader community consensus.

Notable participant comments on this final theme were as follows:

- The Common Rule provides ill-fitting guidance for the activities in ICT research, and as a result, ERBs are asking the wrong questions.
- The community needs to have a place to engage in ethics-related discussions. This would be valuable to interpreting and applying Common Rule requirements.
- We should find ways to educate and inform ERBs about what ICT researchers are doing so they can better understand and provide more thoughtful and effective reviews and guidance.
- The community should make available case studies of both ethically commendable and regrettable studies.

Somewhat orthogonal to the core goal of this final session, a nontrivial discussion organically emerged around research using "found data"- data that is purportedly publicly-available on the Internet. Specifically, the conversation centered on the Carna botnet and the *Internet Census 2012* data that was made available by the author. Notably, the Carna botnet used hundreds of thousands of insecure embedded devices to scan the entire Internet. Conversation focused on the acceptable use of data collected in this fashion. An informal survey of participants indicated that an ample majority supported

the use of such data for research, while a minority rejected its use for research on ethical grounds.

A detailed account of the grounds for this disproportionate split in opinions is beyond scope, but nevertheless it offers concluding insight and motivation going forward. With both positions on the Carna issue, the confluence of: (a) the relatively evident difficulty in articulating reasoned ethical justifications; and (b) the unflinching drive to innovate and enhance knowledge, promises that controversies will continue. Responsibility dictates that we address the issues and solutions articulated herein in a proactive manner.

ACKNOWLEDGMENTS.

The authors would like to thank Jenny McNeill, SRI International, and KC Claffy, Cooperative Association for Internet Data Analysis for their workshop assistance and with capturing the sentiments expressed herein. Support for this work was provided by the U.S. Department of Homeland Security, Science & Technology Directorate, in conjunction with the PREDICT project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors or originators and do not necessarily reflect the views of this organization.