

Human Subjects, Agents, or Bots: Current Issues in Ethics and Computer Security Research

John Aycock¹, Elizabeth Buchanan², Scott Dexter³, and David Dittrich⁴ *

¹ University of Calgary

² University of Wisconsin-Stout

³ Brooklyn College

⁴ University of Washington

Abstract. In this panel, we explore some of the issues surrounding the ethical review of computer security research by institutional review boards (IRBs) and other ethical review bodies. These issues include interpretation of legal language defining how ethical review is to be performed, the impact of information and communication technologies (ICT) on research methods and ethical analysis, how terms like “risk” and “harm” must be interpreted in the light of ICT. We examine two case studies in which these issues surface, and conclude by providing some ideas on the path forward.

1 Introduction

This statement addresses issues of human research ethics boards/institutional review boards and the concept of computer security research from the perspectives of four researchers, including computer scientists, ethicists, and computer security experts. We frame this discussion from the extant regulations, particularly those from the United States Code of Federal Regulations at 45/46 [1], and the Canadian Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans [3]. As computer security research has grown in such forms of bot research, malicious software, and denial of service attacks [4], attention from and interest by review boards to these forms of research has correspondingly grown [7]. However, on such review boards, there is a lack of expertise and representation on boards, with anecdotal evidence from 2009 showing that of computer security researchers, only 4 out of 200 at the Network and Distributed System Security Symposium reported serving on review boards, while a lack of knowledge on US institutional review boards around computer science and computer security research has been reported. Buchanan and Ess [2] found that in their respondents, 75% of 700 boards did not have a technical expert to review internet-computer related research protocols, and, 75% of boards did not provide training for their boards in this area. While some years earlier, Hall and Flynn [5]

* Copyright © 2011, IFCA. Primary source of publication is <http://www.springer.de/comp/lncs/index.html>

conducted a survey of Computer Science departments in the UK regarding human subjects research ethics in software engineering (SE) research with a response rate of 47% (44 department heads). At that time, they found several trends that point to a lack of culture of support surrounding human subjects research ethics. Few other empirical studies have been conducted to disprove what Hall and Flynn, and Buchanan and Ess, have found. Unfortunately, it appears that both the current state of CS departments and IRBs are not in sync around human subjects and ethics awareness just yet.

2 Specific Computer Security Issues: Risk

The traditional principles of research ethics include respect for persons, beneficence, and justice. IRBs have the mission to protect *human subjects of research* and serve as the advocate for research subjects in evaluation and review of research. A particular area of concern for computer security revolves around risk. According to the Office of Human Research Protections [9], “risks to research subjects posed by participation in research should be justified by the anticipated benefits to the subjects or society. This requirement is clearly stated in all codes of research ethics, and is central to the federal regulations. One of the major responsibilities of the IRB, therefore, is to assess the risks and benefits of proposed research.” At 45 CFR 46.102(i) [1], the regulations require researchers only to *minimize*, not *eliminate*, risks to research subjects. In IRB discourse, many computer security protocols would fall into the realm of minimal risk – CS research presents many different possibilities for reconceptualizing this regulatory concept of minimal risk. The concepts of universal precautions, individual precautions and responsibility are key. Researchers and boards must balance presenting risks related to the specific research with risks related to the technologies in use.

With computer security research, major issues around risk arise, for society at large especially. The risk may not seem evident to an individual but in the scope of security research, larger populations may be vulnerable. There is a significant difficulty in quantifying risks and benefits, in the traditional sense of research ethics, and an ultimate question that is emerging is, can the computer security researcher articulate those in terms that an IRB understands and can quantify appropriately? The goal of the IRB in general is to protect human subjects, while ensuring appropriate methods and ethics in research. Meaningful assessment of computer security research may involve understanding technical details well outside the area in which IRB members are trained; effectively, the methods are concealed. However, if computer security research is opaque as to methods, it is opaque as to ethics, as methods and ethics are inherently intertwined. This is a challenge for both security researchers and for IRBs.

Moreover, in computer security research, the distance between researcher and “subject” or participant influences how an IRB will review the risk-benefit. For instance, as the “distance” between the researcher and subject/author/participant decreases, we are more likely to define the re-

search scenario as one that involves “humans” [8]. As the distance increases, we are more likely to define the research scenario as one that does not involve “humans.” For instance, an aggregation of surfing behaviors collected by a bot presents greater distance between researcher and respondent than an interview done in a virtual world between avatars. This distance leads us to suggest that computer security research focus less concern around human subjects research in the traditional sense and more concern with human harming research. As an IRB reviews a computer security research protocol, it is imperative that they consider the larger picture beyond the individual. This is in keeping with novel review forms, such as those corresponding to community-based participatory research [10].

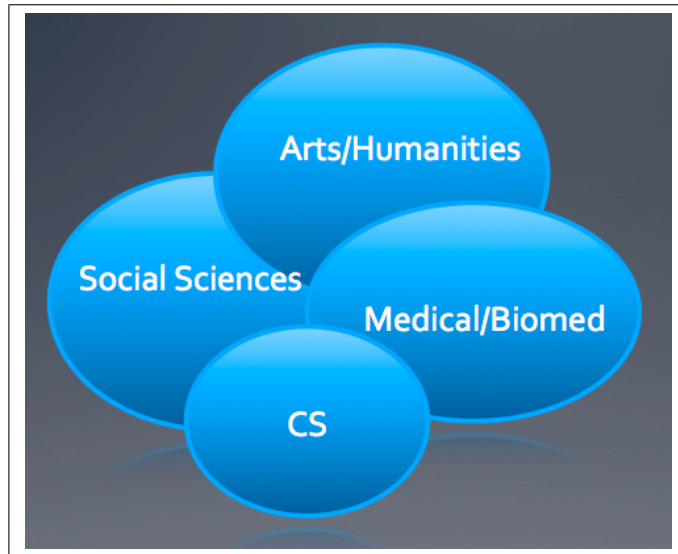


Fig. 1. Where does CS fit in?

2.1 CS Specificity

Computer security is a sub discipline, not analogous with biomedical or social sciences, and thus, fails to fit smoothly within regulatory and practical IRB language. IRBs have little expertise available to them to know all of these potential risks to protected data, let alone how to identify when a researcher is not adequately explaining the full range of protective measures that are necessary to prevent data breach, or other risks involved with computer security research that go beyond just data. In the realm of IRB discourse, specific language has been used to review all disciplinary research, though, as noted, novel interpretations are emerging [10]. It is common to see in an IRB application for a research

study involving collection of personally identifiable information phrases such as:

- “No others will have access to the data.”
- “Anonymous identifiers will be used during all data collection and analysis and the link to the subject identifiers will be stored in a secure manner.”

These phrases are boilerplate, in that they inherit the standard means of protecting data stored on pieces of paper physically located in a locked cabinet in a locked researcher’s office in a building with a guard at the front desk. They give very little detail as to the methods used to secure data, but an IRB committee that includes no experts in computer security may be satisfied with these assurances. Only after a breach would questions arise.

More and more, researchers are learning about the policies and practices of data security being enforced as part of the compliance regime at university medical centers. These policies follow years of records breaches that seem to be growing in scale and scope.⁵ IRB applications now contain language such as:

- “Data files that contain summaries of chart reviews and surveys will only have study numbers but no data to identify the subjects. The key [linking] subject names and study identifiers will be kept in a locked file.”
- “Electronic data will be stored on a password protected and secure computer that will be kept in a locked office. The software ‘File Vault’ will be used to protect all study data loaded to portable laptops, flash drives or other storage media. This will encode all data using Advanced Encryption Standard with 128-bit keys (AES-128).”

These statements are more explicit about protection of *data*, because the risk of disclosure (or exposure) of subject identities is a well-known risk. But what does it mean to store data in a “locked file?” Is the lock the username/password required to disable the screen saver or log in to the user account? Or does it mean a password protected ZIP archive file or Excel spreadsheet? What is the quality of the password (i.e., is it the word “password” or is it the the same string as the user name?) or the encryption algorithm that is used? Is the password written on a note stuck to the monitor? Who else has access to the directory in which the file is kept, have they undergone the same ethics training and signed the same confidentiality agreements other research staff have signed? Is that directory configured so as to be openly readable from any computer on the local network or anywhere on the internet? Is the data copied to a USB thumb drive that requires no password at all to mount and read from another computer? Even when specifying the details of using whole-disk encryption with File Vault, if the password is trivially guessable, even military-grade encryption does little good in terms of protecting stored electronic data.

Securing data is about more than just a password, or using a particular encryption algorithm. Data must be protected not only at rest, but in

⁵ <http://www.privacyrights.org/data-breach>

transit over the network, and while it is being processed (when it must exist in decrypted form in order to be of use to a researcher.) It must also be securely destroyed when no longer needed. Simply deleting a file does not, by itself, over-write the sensitive data contained in the file: it simply makes the space on disk that the data occupied available for potential re-use when new files are created and can often be trivially recovered by anyone with rudimentary knowledge of computer forensics and the proper software tools to recover deleted files.

But the risks go even farther than just data breach. Some computer security research puts at risk the integrity and/or the availability of information and information systems. Altering or destroying data, or causing a computer system to crash, can have just as serious harms as disclosure of patient records could have (perhaps even more). There is a long way to go before current IRB review mechanisms will adequately handle the breadth and depth of risks and benefits associated with computer security research.

To assist researchers and IRBs as they struggle with the complexity of data security, Harvard University has recently devised a tiered system of data management [6]. Specifically, Harvard outlines responsibilities for researchers, IRBs, and the IT departments within the university. In addition, they have designated different types of data along a continuum and apply appropriate standards of protection to those various forms of research data. For example, in Harvard's Information Security Categories range from Level Five: "Extremely sensitive information about individually identifiable people;" to Level Four: "Very sensitive information about individually identifiable people;" to Level Three: "Sensitive information about individually identifiable people;" to Level Two: "Benign information about individually identifiable people;" to finally, Level One: "De-identified research information about people and other non-confidential research information." Harvard's approach is a meaningful and practical approach that stands to assist researchers and IRBs in their understanding and appropriate evaluation of data security. Yet, there remain obstacles to the computer security field moving more readily to a research ethics awareness, or, towards an suitable model that best fits the specificity of the discipline, while ensuring the basic principles of research protections.

3 Best Practices and Suggestions

As this current WECSR Workshop has shown, the community of computer security researchers themselves need to identify and promulgate best practices, develop curriculum standards, self-regulate, perhaps through the model of an extra-institutional review board, and develop and apply "ethical clean bill of health" standards for publications. For instance, the Journal of the American Medical Association has an explicit statement for ethical considerations⁶, asserting: "For all manuscripts reporting data

⁶ <http://jama.ama-assn.org/site/misc/ifora.xhtml#EthicalApprovalofStudiesandInformedConsent>

from studies involving human participants or animals, formal review and approval, or formal review and waiver, by an appropriate institutional review board or ethics committee is required and should be described in the Methods section. For those investigators who do not have formal ethics review committees, the principles outlined in the Declaration of Helsinki should be followed. For investigations of humans, state in the Methods section the manner in which informed consent was obtained from the study participants (ie, oral or written). Editors may request that authors provide documentation of the formal review and recommendation from the institutional review board or ethics committee responsible for oversight of the study.”

Similarly, and in a positive directive, the Symposium on Usable Privacy and Security (SOUPS)⁷ has recently moved in a similar fashion and adopted an ethics statement for papers and publications: “Papers should mention how the authors addressed any ethical considerations applicable to the research and user studies, such as passing an IRB review.” Specific examples from SOUPS include⁸:

“Example descriptions of ethical considerations:

- This study was approved as a minimal risk study by our university’s IRB.
- Study participants were debriefed after the study to make them aware of the deception used in the study and to inform them of how they could protect themselves had this been an actual phishing attack. This study, including our use of deception and subsequent debriefing procedure, was approved by our university’s IRB.
- Our organization does not require human subjects review.
- According to the rules of our institution, this study did not require IRB approval because all human subjects data was gathered from previously-published publicly available data sets.”

In order to expose the next generation of researchers to broader, more appropriate research ethics models, we advocate for a number of pathways. These include pedagogical, professional, and regulatory. In the realm of the pedagogical, we recommend CS classes with a “research methods” component allow instructors to serve as ex officio members of an IRB in order to strengthen faculty-IRB connections and expose students to IRB motivations and methods. Academic environments, in conjunction with professional societies should develop short module on research ethics which could be used in variety of courses (security, computer ethics, software engineering, HCI, etc), that expand in scope and detail the current Collaborative Institutional Training Initiative (CITI) models, which are used internationally as an online research ethics training program. In addition, the field of CS needs to explore the importance and centrality of ethics in such initiatives as NSA and the Department of Homeland Security (DHS) jointly sponsored National Centers of Academic Excellence in IA Education (CAE/IAE) and CAE-Research (CAE-R)

⁷ <http://cups.cs.cmu.edu/pipermail/soups-announce/2011/000055.html>

⁸ <http://cups.cs.cmu.edu/soups/2011/ethics-examples.html>

programs (CAE/IAE or CAE-R criteria), and clearly comprehend and articulate the standards for federally-funded research.

In terms of industry connections, as Shou [11] has shown, the opportunity for industry-academy partnership around ethics is possible. His work on the ethics of data sharing shows that industry and academia can share in their ethics frameworks for the betterment of research in general. In terms of regulation, we need to engage as a discipline in significant risk assessment evaluation: Does security research carry the possibility of harm at the levels of, say, pharmaceutical research or research involving infectious diseases/nuclear materials? As such, regulations will be formed accordingly. Current laws around computer security are mixed, and ambiguous, and for researchers, this complexity in law contributes to confounding ethics.

4 Conclusion

We have seen some of the issues surrounding the ways computer security researchers communicate with IRBs, how well IRBs do or do not understand the risks inherent in computer security research, and some ways in which the computer security research community can move forward towards improved ethical evaluation capacity.

References

1. 45 CFR 46. <http://www.hhs.gov/ohrp/humansubjects/guidance/45cfr46.htm>.
2. E. Buchanan and C. Ess. Internet Research Ethics and the Institutional Review Board: Current Practices and Issues. In *ACM SIGCAS Computers and Society*, volume 39, 2009.
3. Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, Social Sciences and Humanities Research Council of Canada. Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans. http://www.pre.ethics.gc.ca/pdf/eng/tcps2/TCPS_2_FINAL_Web.pdf, December 2010.
4. D. Dittrich, M. Bailey, and S. Dietrich. Towards Community Standards for Ethical Behavior in Computer Security Research. Technical Report CS 2009-01, Stevens Institute of Technology, April 2009.
5. T. Hall and V. Flynn. Ethical issues in software engineering research: A survey of current practice. *Empirical Software Engineering*, pages 305–317, 2001.
6. Harvard University Information Security. Harvard Research Data Security Policy Protection Memo. <http://security.harvard.edu/harvard-research-data-security-policy-protection-memo>, October 2010.
7. E. Kenneally, M. Bailey, and D. Maughan. A framework for understanding and applying ethical principles in network and security research. In *Proceedings of the 14th international conference on Financial cryptography and data security*, FC'10, pages 240–246, Berlin, Heidelberg, 2010. Springer-Verlag.

8. A. Markham and E. Buchanan. The Distance Principle in Internet Research Ethics. (Forthcoming) *International Journal of Internet Research Ethics*, 2011.
9. Office for Human Research Protections (OHRP). Institutional review board guidebook. http://www.hhs.gov/ohrp/archive/irb/irb_chapter3.htm, 1993.
10. L. Ross, A. Loup, R. M. Nelson, J. Botkin, R. Kost, G. Smith, and S. Gehlert. Human subjects protections in collaborative community-engaged research: A research ethics framework. *Journal of Empirical Research on Human Research Ethics*, 5(1):5–17, 2010.
11. D. Shou. Ethical considerations of sharing data for cybersecurity research. In *Proceedings of the 15th international conference on Financial cryptography and data security*, FC'11, Berlin, Heidelberg, 2011. Springer-Verlag.