

Moving Forward, Building An Ethics Community (Panel Statements)

Erin Kenneally¹, Angelos Stavrou², John McHugh³, and Nicolas Christin⁴ *

¹ Cooperative Association for Internet Data Analysis; Elchemy

² George Mason University

³ Redjack

⁴ INI/CyLab, Carnegie Mellon University

Abstract. The organizing question around which this panel at WECSR 2011 rallied was how to move toward building a nation-state-agnostic ethics community in computer security research.

1 Ethics As a Three-Legged Stool

To jumpstart the discourse, panel moderator Erin Kenneally framed the issue of ethics in computer security research as a metaphorical three-legged stool consisting of principles, the applications of principles, and implementation of those applications. Accepting that model, the problems that define the current state of affairs of ethics in computer security research expose frailties along each of the three appendages, as well as that of a domain-agnostic yet nebulous fourth limb.

Specifically, the security research community and the larger domain of information and communication technology research (e.g., network measurement, computer-human interface, software engineering) lack shared community values - guiding principles around which 'right and wrong' research conduct can be assessed, systematized, influenced, and defended[3]. The growth and persistence of debate among relevant conference program committees over the ethical propriety of certain research offers a glimpse of this disharmony. Arguably, the problem may be less one of disagreement over principles than a failure to galvanize principles into a coherent delivery vehicle.

Moving on to the second leg of the ethics stool, the community is faced with a dearth of domain guidance and technical enablers to translate the abstract and theoretical ethics principles into practicable actions. Specifically, there is a lack of both formal institutional and ad hoc peer guidance in ethical decision management, thereby reinforcing the vacuum within which first order ethics principles are embraced at the community-level.

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

Further, assuming the existence of guidance, there are nary few tools that embed, consistently reproduce, and scale such expert ethics advice. Together, these deficiencies all but relegate an ethics-by-design goal for the computer security community fantastical.

Finally, there is a shortage of forcing functions that would carry the weight of the third leg of the stool, implementing the applications of ethics principles. Specifically, while Institutional Review Boards (IRB) have carried the mandate to ensure ethics in research involving human subjects, their relevance and capabilities in computer security research is under debate. Furthermore, it is unclear the extent to which other institutions, such as conference program committees or funding agencies, can or are willing to provide the oversight and quality control to ensure that ethics are identified, applied, and evaluated in research endeavors. Incentives are the implicit fourth element of the structure that directly relate to implementation. Currently, the community of researchers are neither presented with carrots - e.g., accolades, competitive advantage by way of funding or publication, nor faced with sticks - e.g., termination of funding, conference rejection.

Lest the panel end before it got started, Kenneally segued the discussion by highlighting a path forward paved by promising mechanisms to shore up the three-legged stool. Specifically, the Menlo Report is a multi-year work in progress by a collection of community stakeholders to galvanize ethics principles and their applications. The document is modeled after the Belmont Report, a bellwether guide for biomedical and behavioral research, which roots U.S. federal regulations governing ethical protections of human subjects in research. As for tools to help elucidate and systematize the application of ethics principles, an emerging solution is the Ethical Impact Assessment (EIA)[4]. Modeled after security requirements documents (if you are a techie) or privacy impact assessments (if you are a policy wonk), the design goals of the EIA are to lower the barrier of entry for researchers and oversight or advisory entities to operationalize the application of the Menlo principles into their research design, implementation, and assessment activities. These mechanisms are a path forward for the community to embrace a self-regulatory approach to embedding ethics in their respective research so as to evolve a more mature and community-built notion of what is ethically defensible. One alternative is to wait for an unfortunate event to trigger hasty, top-down forcing functions that will likely not bear the input of this community that will shoulder much of the consequences.

2 Computer Security Ethics, Quo Vadis?

Panelist Angelos Stavrou rhetorically imparted the question, "CS Ethics, Quo Vadis?" to jumpstart his commentary.

Research Ethics has been a subject of active debate in health and human-related sciences including medicine, biology, and behavioral sciences. In

those fields, researchers have to submit their research plans to an Institutional Review Board (IRB) or Ethical Review Board (ERB). Such committees are formed within the researcher's institution to approve, monitor, and regulate conducted research that involves human subjects. The mission of these committees is to provide an independent mechanism to protect the rights and well being of the participating subjects from the effects of the conducted research.

Although the mission of the IRB and ERBs encompasses the entire research that is conducted in an institution, its role has been limited to sciences that involve human or living subjects. Their design and requirements for fields such as computer science has been vastly inadequate to capture the "essence" of what needs to be protected and how. Researchers find themselves in a conundrum when requesting and IRB approval for research that does not involve direct human interaction but involves human activity (for instance human generated network traffic). The IRB committee either provides a "carte blanche" to the researcher or denies the request based on unspecified concerns for harming the rights of the human subjects. In the first case, the CS researcher is compelled to explain the risks and potential harm to the human subjects only to find out that her research plan has been denied because the IRB committee does not have the mechanisms and expertise to apply medical and behavioral protocols to the new brave world of computers and computer generated information.

Moving forward, it is the duty of the computer science researchers to discuss and take action on the Ethical issues, risks, and mitigating factors for collecting, processing, and storing human generated information. We, as a community, are responsible to form the right mechanisms that will allow unequivocally and without bias experimentation in the CS field. Indeed, it seems that now is the right time to analyze what older and more mature scientific fields have done regarding ethics rules and adapt them to the Computer Science research.

3 Be Careful What You Wish For

John McHugh further enhanced the dialogue by cautioning, "Be Careful What You Wish For."

Over the course of the last few years, there has been a movement to draft a set of ethical standards for the conduct of research in computer security. While there is a clear need for such standards, the effort and its resulting guidelines, commonly known as the "Menlo Report" are primarily directed at the academic community.

To a large extent, the Menlo Report is an attempt to adapt the earlier "Belmont Report" which provides ethical guidance primarily for medical research involving human subjects. The Belmont report was a result

of widely publicized abuses of human subjects by researchers in the period leading up to, during, and after the Second World War. The report and regulations stemming from it place restrictions on research involving human subjects funded by the US department of Health and Human Services. Identical regulations have been adopted by some 14 other U.S. government agencies. The regulations effectively cover any research being performed at an institution receiving funds from one of these agencies, whether government funded or not. Most academic researchers have learned to accommodate the requirements, factoring into their research plans the time needed to obtain Institutional Review Board approval and documenting their research approach and process accordingly. Since most of the medical and pharmaceutical research in the U.S. involves academic participation, the regulations also affect substantial industrial research programs, as well. In recent years though the pharmaceutical industry has turned to the third world to conduct clinical trials under conditions that would not pass IRB scrutiny in the U.S.⁵.

Academic computer security research (and academic computer science research, as well) is already in a state of crisis, largely due to pressure to publish early and often. When the author obtained his PhD in the early 1980s it was often the case that a new graduate's first conference or journal publication resulted from the work that led to the degree and was excerpted from the dissertation. Today, it is not uncommon for graduate students (and their advisers) to amass several publications per year out of work leading to the degree. For several years, I have been involved in an effort to raise the quality of academic research in the field by insisting that experimental papers contain an explicit description of the research question or hypothesis being investigated and detailed description of the experimental setup and methodology used to conduct the experiment. At a recent IFIP workshop, these suggestions were met with substantial resistance by a number of well known researchers in the fault tolerant and dependable systems area, largely on the grounds that the effort involved would slow the pace of the student's publication, jeopardizing employment prospects upon graduation.

I note that, to a large extent, the process that I advocate for research in general would be required for IRB approval in cases where human subjects are directly involved, and that much of the effort might be required in building a case that IRB approval should be waived for research with only a tenuous connection to real human subjects. One of my concerns is that imposing such conditions might result in driving students away from meaningful research questions requiring IRB interactions on the grounds that other research will produce more publications with less effort. Another concern is that this will exacerbate the current trend towards rapid (though trivial and largely useless) research leading to quick publications.

⁵ See "Deadly Medicine", Donald L. Barlett and James B. Steele, Vanity Fair, January 2011 on line at <http://www.vanityfair.com/politics/features/2011/01/deadly-medicine-201101?printable=true¤tPage=3ixzz18NY8yGh9>.

Unlike the medical area, a substantial amount of computer security research is conducted outside academia. Much of this work receives no government funding whatsoever and is largely beyond the reach of the processes proposed by the Menlo report. In a keynote address at the 6th European Conference on Computer Network Defence, Felix 'FX' Lindner of Recurity Labs gave a talk entitled "On Hackers and Academia" in which he took the academic community to task for concentrating on largely irrelevant approaches in an area that desperately needs useful results to help solve real problems. Recurity Labs is but one of hundreds of organizations that conduct research in computer security. These organizations have been largely left out of the ethics discussion although their actions in areas such as vulnerability disclosure and the development of both attack and defense techniques have the potential to cause serious societal harm on a broader scale the work of many academic researchers and they should be brought into the discussions.

There is a need for an open dialog on ethical issues in the community. Insofar as I can tell, the topic is completely ignored in most academic training programs at both the graduate and undergraduate levels. When it is approached, it is often couched in a legalistic rather than in an ethical framework. The Menlo report is, perhaps, too human subjects centric in its emphasis on IRB involvement. The issue is much broader than that and needs to be placed in a context of societal expectations for ethical behavior that apply inside and outside of the research arena. Although imposing an ethics review process on academic research sounds like a good idea initially, we need to be careful to ensure that it does not alter the research landscape so that valuable lines of research are avoided or pushed underground, out of academia, because the approval process is viewed as too onerous or time consuming.

4 Incorporating Cultural Differences

Finally, Nicolas Christin rounded out the topic with yet another insightful angle.

The 2011 edition of the WECSR workshop has focused a number of discussions on legal liabilities, and how transnational studies and research could result in interesting legal problems. In particular, a fair amount of time was devoted to arguing about United States vs. international law. Yet, this legal focus leads me to believe that we have ignored a more important point related to ethics: the need to be sensitive to cultural differences.

Specifically, the very definition of ethics varies depending on the culture considered. While I am not an ethicist, I have done research both in the United States and in Japan, and have, as is common for information security specialists, interacted with a large number of scientists from different cultures.

The West usually distinguishes between three different types of ethics. Utilitarian ethics, where the criterion to decide on whether or not a given action is ethical is whether society as whole would be better off – even though the action itself may hurt some individuals. Deontological ethics decide on whether an action is ethical or not, based on its consequences. Virtue ethics, on the other hand, use the character of the agent performing the action as a decision criterion.

From the discussions that preceded, it seems that a fair number of computer security experts use such a utilitarian view. In particular, the paper by Moore and Clayton presented earlier in the workshop uses this utilitarian argument to justify certain experiments that were conducted.

Yet, it is interesting to note that, in Asia for instance, the notions of ethics are completely different. Buddhist ethics can be construed to some extent as a combination of deontological and utilitarian ethics (“anatta”), while some (e.g., Gier⁶) have compared them to virtue ethics. In addition, a certain amount of modesty would be considered as an ethical necessity. The author that uses a large dataset, potentially hurting a large number of people in the process may be viewed as unethical if s/he does so to publish a research paper to further his/her reputation, even though, from a purely utilitarian standpoint, the action would be ethical if the benefits to society are considerable.

In another ethical puzzle, in Japan, it is often the case that, when exposed to a scandal, top management of a company resigns *even if they are (and are believed to be) personally innocent*.⁷ In the West, this would amount to an admission of guilt. Again, our ethical frames are colored by our cultural backgrounds.

Where does this leave us for Computer Security research? My thesis is that, when dealing with data coming from geographically diverse origins, we need to adopt ethical frames of reference that match the culture or ethnic groups we are considering rather than ours. For instance, when a large number of Mechanical Turk users participating in online behavioral experiments (e.g., [1]) are from India, we need to apply ethical notions relevant to our Indian users; if we study frauds or online scams prevalent in a single country, like One Click Fraud[2] we need to adopt a definition of ethics consistent with the predominant culture in that country.

References

1. N. Christin, S. Egelman, T. Vidas, and J. Grossklags. It’s all about the Benjamins: An empirical study on incentivizing users to ignore security advice. In *Proceedings of IFCA Financial Cryptography’11*, Saint Lucia, February–March 2011. To appear.

⁶ See <http://www.class.uidaho.edu/ngier/307/buddve.htm>.

⁷ See <http://ccbs.ntu.edu.tw/FULLTEXT/JR-PHIL/wargo.htm> for anecdotes.

2. N. Christin, S. Yanagihara, and K. Kamataki. Dissecting one click frauds, Oct. 2010.
3. 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.
4. 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.