Ethical Behavior in Cyberspace Research

JOHN MURRAY

John Murray is a Program Director in the Computer Science Laboratory at SRI International, Silicon Valley, CA. His research interests include interactive human-machine technologies, collaborative intelligence, multi-player game systems, and cognitive engineering. Prior to joining SRI, Dr. Murray held executive and technical leadership positions at several international information systems firms. He holds advanced engineering degrees from Stanford, the University of Michigan, and Dublin Institute of Technology in Ireland. jxm@sri.com

Traditional principles of scientific ethics for studies involving people are primarily designed to address direct *human-centered* research. However, online research in cybersecurity is inherently *data-centered* in nature. Consequently, cyber-researchers often operate with limited awareness or oversight of the potential human risks and effects of their activities. Newly proposed changes in the US policies and regulations governing human studies are likely to have a significant impact on the online research community. Furthermore, given the inherently transnational nature of online studies, there is a pressing need for harmonizing ethics observance regulations and guidelines for security research on virtual worlds, social network systems, and other cyber-environments across multiple jurisdictions.

At the USENIX 2015 Security Symposium, a panel of academic and industry experts focused on cybersecurity research ethics. The discussions specifically centered upon the dilemmas facing those involved in academic publishing—editors, reviewers, etc.—when confronted by articles that discuss cybersecurity explorations, which may reveal potential or real exposures or vulnerabilities.

The underlying questions of research integrity concern the beliefs and justifications that system developers, investigators, and experimenters use as the basis for undertaking their explorations, what types of impacts are considered and when, what benefits vs. harms tradeoffs are made, and so on.

While the moral dilemmas of revealing system vulnerabilities in academic publications are indeed important, they generally come towards the end of a (potentially lengthy) research effort, well after other damage may already have been done. In reality, the actual ethical challenges should be considered early in the process, when the research team is designing their initial investigations.

For example, suppose that cyber-investigators are exploring aspects of real-time online censorship in various countries. Their strategy is to find ways to initiate download requests across national boundaries for various forms of potentially controversial material. In order to accomplish this, the researchers gain unauthorized access to some individual devices in a targeted jurisdiction, which they use as proxy platforms for issuing their exploratory requests.

However, in their zeal to deploy their probes, they neglect to consider the possible adverse effects that their study might have on the owners or operators of the compromised systems. Such individuals or groups may be put at risk vis-a-vis their own government authorities, as a result of the investigators' actions. A tech-savvy ethics review of the research plans would probably have drawn attention to the potential problems and ensured that appropriate safeguards were put in place.

Many academic institutions and larger corporate organizations have ethics oversight panels or institutional review boards (IRBs), which are responsible for monitoring the human health and safety of experimental subjects, and attending to the potential for exploitation or coercion, especially among vulnerable populations. The historical background for IRBs grew out of the revelation of numerous misguided and abusive scientific studies during the twentieth century. The result was the introduction of policies and standards based on the 1979 Belmont Report [1], which is still used today to guide scientific ethics reviews across the US and beyond. These guidelines translate into government regulations for several categories of human subject research, in particular those that are supported by US federal funding.

Although these guidelines and regulations have continued applicability to classical human research laboratory work in fields like medicine and psychology, they have limited practical relevance to modern behavioral studies that involve highly networked information and communications technology (ICT) systems. These impracticalities are exacerbated by the pervasive need to undertake comprehensive, transnational experimental projects, where much of the human data collection and analysis is undertaken remotely across varied, and often incompatible, legal regimes and social norms. Yet such is the case for numerous researchers nowadays, who are working not just in ubiquitous social networks and popular gaming worlds, but also with online educational environments, cybersecurity applications, and monitoring/surveillance systems.

In consideration of these challenges, a 2011 update to the earlier guidelines, called the Menlo Report [2], was specifically developed to address issues of online security, privacy, anonymity, and other personal identifiable information (PII) concerns. The report's authors recognized that the broad cyber-research community needs a more rational and coordinated strategy for managing ethics observance, which particularly considers the scope and needs of ICT research. Such a tailored approach should emphasize studies of human behavior and community activity online, and apply across multiple jurisdictions in interactive professional and social environments.

This transition of some of these concerns into formal policies and regulations recently moved forward with the publication of a Notice of Proposed Rulemaking (NPRM) in the US Federal Register [3]. This serves to promote conversation and comment from parties affected by the proposed changes. The latest comment period is open until December 2015, after which revisions to the proposal will be considered in light of comments received.

As they currently stand, some of the proposed changes may have significant implications for transnational cyber-research. One key concern is the extent that they might exacerbate the differences between human subjects research requirements in the US and elsewhere, while at the same time relaxing some of the more stringent requirements that currently apply to the US research community.

Traditional ethics reviewers try to ensure equitable distributions of burdens and benefits among the human subjects actually involved in the study. However, as noted in the example earlier, online research activity may adversely affect innocent bystanders and neutral nonparticipants. Given the risks associated with real-time data-intensive experiments, such studies might better be reviewed in terms of *human-harming* research rather than human subjects research.

For example, solid contingency and response plans are needed for mitigation of realized harms, especially for low-probability/high-impact events. These types of safety monitoring procedures are standard in traditional biomedical studies, but are rarely considered in ICT research. Furthermore, when research involves surveillance, profiling, or monitoring, additional vulnerability protections are needed to prevent the misuse of findings and results. This is particularly the case when novel mergers of partial data from several public sources may produce PII that is not individually available from just one of them. Other concerns arise from the potential for abuse of data for social discrimination, especially by non-investigators.

Provisions are required to ensure conformance with international regulations on transborder data flow that include personal information. In this regard, the current oversight policies and data handling processes for multi-jurisdictional ethics approvals are primarily centered upon the requirements of pharmaceutical drug trials, medical device tests, etc., rather than on the research needs in global-scale social science, human-machine systems, and ICT.

To address this gap, an international ethics observance organization is needed, which would coordinate/oversee regulations and guidelines for research in online systems and other cyber-environments across multiple jurisdictions. This could be a consortium of nonprofit organizations in several domains, which would ensure smooth transnational processing of approvals. It seems appropriate that such a consortium would need to have the backing of a recognized international entity such as UNESCO.

The first steps toward such harmonization could be merely a matter of coordinating and making available the critical features of each local research context, or it could extend to negotiating safe harbors for compliance with a local research context. Thus, if a study complies with certain key components, then it is deemed to satisfy local research context requirements for specific countries. Another, further step might be to aim for legislative harmonization on the topic of research protection.

The bottom line is that almost any form of standardized ethical framework would help cyberspace researchers worldwide become more aware of the challenges and know when they have addressed some required basic considerations. This must be better than the current haphazard obstacle course, which generally leaves everyone guessing as to what they still need to do to work through this ethical minefield.

References

[1] Belmont Report: www.hhs.gov/ohrp/humansubjects/guidance/belmont.html.

[2] Menlo Principles: www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCOMPANION-20120103-r731_0.pdf.

[3] Notice of Proposed Rulemaking, Protection of Human Subjects: www.federalregister.gov/articles/2015/09/08/2015-21756/federal-policy-for-the-protection-of-human-subjects.