A Cybersecurity Research Ethics Decision Support UI

Robert B. Ramirez Shun Inagaki Masaki Shimaoka
Kenichi Magata
Intelligent Systems Laboratory, SECOM CO., LTD.

Abstract

Evaluating cyber security ethics gives researchers and ethics committees a tough time the world over. There are currently no easily usable, granular, or comprehensive benchmarks for evaluating the ethics of cyber security research. Existing abstract frameworks for ICT do not give concrete recommendations for most dilemmas faced by security researchers, and instead either focus on the ethical assessment process itself, give only general advice, or do not focus on security. There is also currently no method for evaluating security research ethics in a truly systematic or reproducible manner.

In this work we present a decision support tool with a web-based UI that can be used to prospectively or retrospectively evaluate the ethics of specific cyber security research actions. We constructed this tool by analyzing the ethical and research practices described in 101 conference papers published between 2013 and 2017, selected from a corpus of 943 papers from the top cyber security conferences.

1 Introduction

Research ethics in Information and Communications Technology (ICT) has seen a resurgence in popularity in recent years, spurred in part by AI [1, 2]. Although a number of general standards have been issued in the past decade for ICT, there are currently no easily usable, granular, or comprehensive benchmarks for evaluating the ethics of cyber security research projects. There is also currently no method for evaluating security research ethics in a truly systematic or reproducible manner [3–8].

USENIX Symposium on Usable Privacy and Security (SOUPS) 2020.
August 9–11, 2020, Boston, MA, USA.

In reality, researchers have to deal with concrete ethical dilemmas on a variety of topics, as evidenced by the prevalence of 'ethical issues' sections in research papers [9].

As a result, despite ethical analyses being demanded by many top conferences [10], traditionally, committees or Internal Review Boards (IRBs) composed of experts have been necessary to comprehensively review the ethics of papers, which has been seen as a domain requiring significant expertise, often from senior members of the research community [11, 12]. However, many IRBs lack experience specifically evaluating ICT research, particularly in a cyber security context [5, 13, 14]. This frequently leads to ICT research being exempted from IRB review [15–17]. Thus, researchers need to be able to evaluate their own research.

In this paper we present a UI for a knowledge base (KB) of concrete cyber security research ethics best practices, which we compiled from a large semi-random survey of research papers. We sorted papers using a topic model, performed systematic reviews, and, using Deontic Logic, organized the results into a decision tree with 210 branches, spanning Data Privacy, Human Subjects Testing, Autonomy, Software Examination, and General Research.

2 Related Work

There are a number of existing guidelines relevant to cyber security research ethics. Here we describe some of the most well known ones and their relation to our work.

The Menlo Report was released in mid-2012 as 'guidance for ICT researchers' in 'the context of ... information security research' [3]. It was inspired by the Belmont Report developed in the 1970s for medical ethics. When formal ethical discussions are included in cyber security research papers, the Menlo Report or its principles are sometimes referenced [18].

The Menlo Report gives examples of its *principles* applied to security, including phishing, vulnerability disclosure, and handling sensitive information; but it does not explain the ethics of those *actions* themselves.

Table 1: Conferences sourced for our Knowledge Base.

Publication	Years	Docs Collected
USENIX	2013-2016	248
IEEE S & P	2013-2017	187
ACM CCS	2016	138
SOUPS	2014-2016	65
USEC and NDSS	2013-2016	253
CREDS	2013-2014	8
PETS	2015-2017	93
All	2013-2017	992

The *Companion to the Menlo Report* [14] does, however, include a synthetic case study showing how to apply its guidelines, but it is limited in scope, organized as prose, lacks specific references to research articles, and does not describe what conditions determine the ethics at play. At the same time, the Menlo Report's guidelines were invaluable for vetting and structuring the practices we compiled.

We also referred to the ACM Code of Ethics (CoE) [19] and the Association of Internet Researchers Recommendations [5] when evaluating practices. The ACM CoE, like the Menlo Report, is written at an abstract level without much guidance on specific research project types, so our reference to these benchmarks was largely a form of due-diligence.

A tool called 'CREDS' is apparently in development by researchers affiliated with the authors of the Menlo Report and the U.S. Department of Homeland Security [20]. The stated goal of this tool, first proposed in 2015, and at one point intended for release in 2017, seems similar to our research in that it seeks to somehow analyze best practices, as well as laws, to create an online ethics tool for the community to use [21]. In 2018 a private alpha version of a checkbox-style questionnaire tool focusing on evaluating data issues for ICT research from a legal perspective could be found online [22], but was soon made private. We have found no other recent evidence of development. We intend to reach out to the CREDS team about combining our efforts.

3 Decision Support Tool

3.1 Cyber Security Ethics Knowledge Base

We extracted descriptions of ethical practices from 101 relevant papers out of a collection of 992 published in the top conferences in cyber security between 2013-2017 [23].

Our work is founded on the assumption that there is a core set of shared ethical issues that covers most cyber security-related questions researchers grapple with. We posit that these can be discovered by sampling from the same set of resources that such research ethics experts use – namely, research papers and ethics standards. The methodology for developing our ethical assessment framework can be divided into the five steps given below.

Step 1. Collect Past Research As the most advanced research often deals with subjects that would serve as interesting ethical case studies, first, we collected the corpus comprising the prior literature published by the top cyber security conferences (Table 1).

Step 2. Identify Grey-Area Papers and Papers Mentioning Ethics We used titles and abstracts to manually identify papers with "possible negative externalities," i.e. research that could potentially impact entities beyond the researcher. Following [9] we also used a regular expression to locate papers mentioning ethics in the corpus.

Step 3. Sort Papers with a Topic Model Since different technological areas give rise to their own sets of ethical problems [9, 24], we used an LDA topic model to classify documents by technological area to pinpoint these areas.

Step 4. Examine Ethics Content To aid researchers in determining best practices or the lack thereof during their ethical assessments, we created a granular KB of grey-area case studies in cyber security research. To create this KB we systematically manually extracted ethical discussions and relevant research details from the grey-area and ethics papers and recorded their ethical issues and arguments (implicit or explicit), randomly selecting at least one paper from each conference represented in a given topic.

Step 5. Organize Best Practices Best practices extracted by the above and those detailed in the Menlo Report and the ACM CoE were compiled into a decision tree (DT) to create a tool based on ethics decisions researchers have made. A DT is an efficient data structure commonly used by IRBs [25–27]. Anticipating future modification due to changes in technology and culture [28], this tool was organized by research area (not by topic label). Based on the content analysis delineated above, the nodes of the DT describe the details of various actions researchers have taken in the course of their security research, while the leaves indicate the ethical consensus (if any) we identified from the research corpus.

3.2 Decision Tree Overview

Our decision support tool consists of our case-study knowledge, cross-referenced with a decision tree containing five main branches, or classes, of rules, grouped and named for organizational convenience. The five classes and their immediate subclasses are shown in table 2. The current version of the tree has 150 branches of more specific activities stemming from these, excluding the Menlo Report and ACM Code of Ethics, which combined have 57.

Table 2: Ethical Guidelines Decision Tree, Main Classes

Main Class	Immediate Subclasses
Software Examination (Related to research involving understanding programs made by others)	Vulnerability Research Reverse Engineering Malware Disclosure
Privacy (Related to administering third parties' information about persons or technology systems)	Collecting Data Handling Data Publishing Data Transferring Data To Third Parties
Autonomy (Related to interactions with others' systems)	Web Scraping Accessing others' systems (Accessing, Trespassing)
Human and Animal Subjects Testing (Related to experiments directly or indirectly involving subjects)	Deceiving human or animal test subjects Misleading, false, or deceptive advertising Honeypots Criminal and Unethical Services Consulting with REB or IRB
General Rules (Rules or ethical processes separated from the rest of the tree due to their universality)	Terms of Service Ethical consistency Documentation and Accountability Menlo Report ACM Code of Ethics

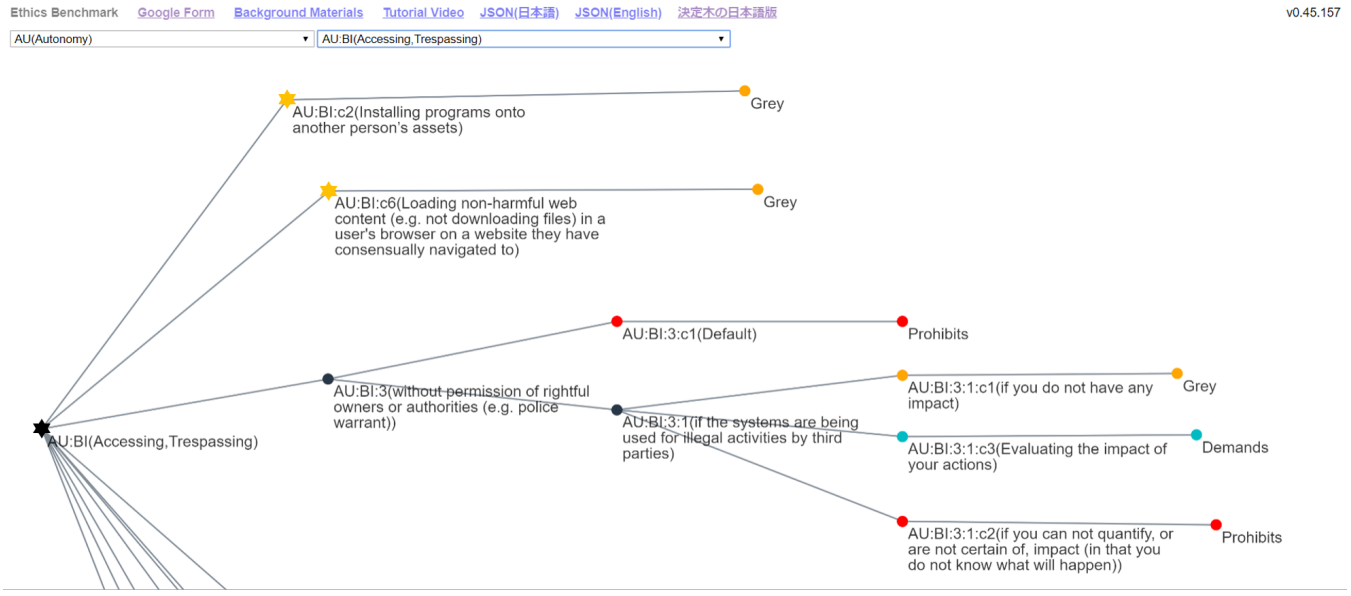

Figure 1: Decision Tree UI

3.3 Structure

Each branch of the tree terminates with leaves that, along with the branch's nodes, specify the ethics of actions a researcher might consider taking. A path may have multiple nodes seemingly referring to different actions by the researcher (as opposed to "conditions"), but only one such action-node is the true referent of any given path's leaf (see Section 3.3.2).

3.3.1 Leaf Nodes

A natural choice for labeling leaf nodes is Deontic Logic, a formal predicate logic that specifies relationships between the ethics of an action given certain conditions [29–32].

Deontic Logic in its most cited and studied form, is based on the so-called "Traditional Scheme" (TS) of normative statuses [33]. For ease of interpretation we instead use our own 8th-grade-English terms for these statuses, as recommended by the Menlo Report [3].

The five normative statuses of the TS, which we use for leaves' labels, [34, 35] are (formal names in parentheses):

1. *Permitted* (**OP**, optional that): performing the action is not in itself unethical
2. *Prohibited* (**IM**, impermissible that): performing the action is in itself unethical
3. *Demanded* (**OB**, obligatory that): not performing the action is itself unethical
4. *Grey* (**OM**, omissible that) and 5. *Recommended* (**PE**, permissible that) are "*TBD*" placeholders indicating a lack of clear consensus on the ethics of the action. *Grey* means the action could be either *Prohibited* (**IM**) or *Permitted* (**OP**); *Recommended* indicates *Permitted* (**OP**) or *Demanded* (**OB**).

All but the last status can be used as a base operator to define the rest. For example, if **IM** is taken as base:

$$\mathbf{OB}p \leftrightarrow \mathbf{IM}\neg p \quad (1)$$

$$\mathbf{PE}p \leftrightarrow \neg\mathbf{IM}p \quad (2)$$

$$\mathbf{OM}p \leftrightarrow \neg\mathbf{IM}\neg p \quad (3)$$

$$\mathbf{OP}p \leftrightarrow (\neg\mathbf{IM}\neg p \,\&\, \neg\mathbf{IM}p) \leftrightarrow (\mathbf{OM}p \,\&\, \mathbf{PE}p) \quad (4)$$

Through the lens of the TS, we see in 1 that any *Demanded* action can be rewritten as *Prohibited* simply by negating it.

3.3.2 Leaf Referents

The natural way of structuring the decision tree is from the perspective of the circumstances or desired actions of the user (that is, the researcher). That is,

- ☆ A primary, internally desired action of the user/researcher's agency
- ○ A secondary, externally motivated requirement, obligation, or condition in response to a star

We model this difference as metadata by labeling actions implying agency with ☆ , as opposed to conditions, circumstances, or incidental actions (labeled with ○). A *Prohibited*, *Permitted*, or *Grey* leaf refers to the most recent ☆ in its branch. As a rule the tree is structured such that *Demanded* and *Recommended* nodes refer to the immediately-preceding node, a ○ node. See Figure 1.

4 Discussion

4.1 Early User Testing

As an early user test, in Fall 2018 our tool was used by the inaugural ethics committee of the largest cyber security conference in Japan, the Computer Security Symposium (CSS). On top of informing usability goals, the testing results motivated the use of Deontic Logic, to separate the (originally combined) "TBD" node categories into *Grey* and *Recommended*, as users ignored such nodes under the TBD label.

4.2 General Recommendations

To make our findings more accessible, we conclude this section with a few very general ethics lessons learned:

- Make use of a secure-data-handling pipeline designed to minimize data retention and burdens on websites
- For research proposals, create a standard review process and/or statement of ethics format
- Use formal software testing methods whenever networks, products, or others' systems are involved [19]
- Keep thorough documentation especially when doing vulnerability or human subjects research, and when handling data (especially personal data)

4.3 Future Work

Below we summarize additional steps we are taking to make our decision support tool robust and usable:

- Set up a way for researchers to submit evidence and opinions on *Grey* and *Recommended* nodes and new practices, to resolve open questions as a community,
- A check-box-style branching web form version of the tool for quickly identifying areas of ethical concern
- Expand the analyzed research to 2008-2020.

5 Conclusion

In our research we demonstrated that we can reveal a number of ethical best practices relevant to human subjects testing, privacy, and other areas of concern to security researchers by systematically examining prior work. We compiled and distilled these practices into a navigable, maintainable, and granular UI with a formal decision logic.

References

[1] "The Ethics in AI Institute." [Online]. Available: https://www.schwarzmancentre.ox.ac.uk/Page/ethicsinai

[2] "Ethically Aligned Design, First Edition." [Online]. Available: https://ethicsinaction.ieee.org/

[3] M. Bailey, E. Kenneally, D. Dittrich, and D. Maughan, "The Menlo Report," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 2145676, Mar. 2012. [Online]. Available: https://papers.ssrn.com/abstract=2145676

[4] "ACM Ethics." [Online]. Available: https://ethics.acm.org/

[5] A. Markham and E. Buchanan, "Ethical decision-making and Internet research: Recommendations from the AoIR ethics working committee (Version 2.0)." AoIR, 2012. [Online]. Available: http://aoir.org/reports/ethics2.pdf

[6] R. Chatila and J. C. Havens, "The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems," in *Robotics and Well-Being*, M. I. Aldinhas Ferreira, J. Silva Sequeira, G. Singh Virk, M. O. Tokhi, and E. E. Kadar, Eds. Cham: Springer International Publishing, 2019, vol. 95, pp. 11–16. [Online]. Available: http://link.springer.com/10.1007/978-3-030-12524-0_2

[7] "EC Council Code Of Ethics," Mar. 2016. [Online]. Available: https://www.eccouncil.org/code-of-ethics/

[8] T. Jim, "There is no standard of ethics in computer security research," Aug. 2014. [Online]. Available: http://trevorjim.com/there-is-no-standard-of-ethics-in-computer-security-research/

[9] M. Akiyama, "研究倫理に関して我々の置かれている状況," in *SCIS 2017*. SCIS, 2017.

[10] "USENIX Security '20 Call for Papers," Oct. 2019. [Online]. Available: https://www.usenix.org/sites/default/files/sec20_cfp_101519.pdf

[11] "Committee Members and Staff | Welcome to COUHES." [Online]. Available: https://couhes.mit.edu/committee-members-and-staff

[12] "45 CFR § 46.107 IRB membership." Jul. 2018, library Catalog: www.law.cornell.edu. [Online]. Available: https://www.law.cornell.edu/cfr/text/45/46.107

[13] D. R. Thomas, S. Pastrana, A. Hutchings, R. Clayton, and A. R. Beresford, "Ethical issues in research using datasets of illicit origin," in *Proceedings of the 2017 Internet Measurement Conference on - IMC '17*. London, United Kingdom: ACM Press, 2017, pp. 445–462. [Online]. Available: http://dl.acm.org/citation.cfm?doid=3131365.3131389

[14] M. Bailey, D. Dittrich, and E. Kenneally, "Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report," 2013.

[15] D. Dittrich, "The ethics of social honeypots," *Research Ethics*, vol. 11, no. 4, pp. 192–210, Dec. 2015. [Online]. Available: https://doi.org/10.1177/1747016115583380

[16] M. Jackman and L. Kanerva, "Evolving the IRB: Building Robust Review for Industry Research," p. 17, 2016.

[17] A. Narayanan and B. Zevenbergen, "No Encore for Encore? Ethical Questions for Web-Based Censorship Measurement," *SSRN Electronic Journal*, 2015. [Online]. Available: http://www.ssrn.com/abstract=2665148

[18] N. Carlini Pratyush Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D. Wagner, and W. Zhou, "Hidden Voice Commands," 2016.

[19] "ACM Code of Ethics," Jun. 2016. [Online]. Available: https://ethics.acm.org/code-of-ethics/

[20] "IMPACT Cyber Trust Ethos," Feb. 2018. [Online]. Available: https://www.impactcybertrust.org/ethos

[21] E. Kenneally and M. Fomenkov, "Cyber Research Ethics Decision Support (CREDS) Tool." ACM Press, 2015, pp. 21–21. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2793013.2793017

[22] "Cyber-risk Research Ethics Decision Support Tool," Sep. 2017. [Online]. Available: http://creds.impactcybertrust.org/

[23] S. Inagaki, R. Ramirez, M. Shimaoka, and K. Magata, "Investigation on Research Ethics and Building a Benchmark." Niigata, Japan: The Institute of Electronics, Information and Communication Engineers, Jan. 2018. [Online]. Available: https://www.iwsec.org/scis/2018/program.html

[24] L. Winner, "Do Artifacts Have Politics?" *Daedalus*, vol. 109, no. 1, pp. 121–136, 1980. [Online]. Available: http://www.jstor.org/stable/20024652

[25] "Ethics Decision Tree." [Online]. Available: https://www.cisco.com/c/en/us/about/corporate-social-responsibility/ethics-office/decision-tree.html

[26] "IRB Decision Tree," Aug. 2019. [Online]. Available: https://funding.yale.edu/applying/fellowships-institutional-review-board

[27] hhs.gov, "Human Subject Regulations Decision Charts," Feb. 2016. [Online]. Available: https://www.hhs.gov/ohrp/regulations-and-policy/decision-charts/index.html

[28] L. E. McCray, K. A. Oye, and A. C. Petersen, "Planned adaptation in risk regulation: An initial survey of US environmental, health, and safety regulation," *Technological Forecasting and Social Change*, vol. 77, no. 6, pp. 951–959, Jul. 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0040162509001942

[29] P. McNamara, "Deontic logic," in *The Handbook of the History of Logic, vol. 7: Logic and the Modalities in the Twentieth Century*, D. Gabbay and J. Woods, Eds. Elsevier Press, 2006, pp. 197–288.

[30] E. Mally, "Grundgesetze des Sollens : Elemente der Logik des Willens." Graz: Leuschner & Leubensky, 1926.

[31] G. H. von Wright, "Deontic logic," *Mind*, vol. 60, no. 237, pp. 1–15, 1951.

[32] S. Knuuttila, "The Emergence of Deontic Logic in the Fourteenth Century," in *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, ser. Synthese Library, R. Hilpinen, Ed. Dordrecht: Springer Netherlands, 1981, pp. 225–248. [Online]. Available: https://doi.org/10.1007/978-94-009-8484-4_10

[33] P. McNamara, "Deontic logic," in *The Stanford Encyclopedia of Philosophy*, summer 2019 ed., E. N. Zalta, Ed. Metaphysics Research Lab, Stanford University, 2019. [Online]. Available: https://plato.stanford.edu/archives/sum2019/entries/logic-deontic/

[34] H. Leblanc, "Prior A. N.. Formal logic. Oxford at the Clarendon Press, London 1955, ix + 329 pp.Prior A. N.. Formal logic. Second edition. Oxford at the Clarendon Press, London 1962, xi + 341 pp," *Journal of Symbolic Logic*, vol. 27, no. 2, pp. 218–220, 1962.

[35] P. McNamara, "Making room for going beyond the call," *Mind*, vol. 105, no. 419, pp. 415–450, 1996.