The National Coordination Office (NCO) for Networking and Information Technology Research and Development (NITRD)

Via email

Request for Information (RFI)—National Privacy Research Strategy
Comments of Deirdre K. Mulligan and Helen Nissenbaum

Thank you for the opportunity to contribute to the development of a National Privacy Research Strategy. We appreciate the Administration's efforts to spur public dialogue, research, and policy development on security and privacy.

As a lawyer and philosopher who have spent many years working with the technical community to: analyze the privacy implications of emerging technical systems, devices, and design choices; build privacy protections into protocols and technical standards, as well as add-ons and free standing systems to inject privacy properties into the existing landscape; design legal and other incentives for the adoption of privacy and security protective technology and practice; and, develop and refine ways of conceptualizing and working with privacy in the technical arena we offer feedback on the pay-off and ongoing *promise of multi-disciplinary approaches* and the related *assessment capabilities they bring to the mix*. For privacy to be embedded appropriately into technical practice and systems, we need methods to identify contextually appropriate understandings of privacy, and methods and tools that simplify the provisioning of appropriate affordances. Multi-disciplinary research is essential to this goal.

The Value of Multi-disciplinary Research

Multi-disciplinary approaches contribute to our knowledge of what ought to be incorporated into the scientific study of privacy to address the rising societal calls for privacy protection in technologies of computing, information, and networking. Multi-disciplinary research plays a role in identifying the concepts, factors and methods necessary to integrate privacy into scientific theory, design practice, and system development to facilitate robust and systematic inclusion of privacy properties into components and systems of the information infrastructure. It extends the scientific boundary of privacy research to include the consideration of human, socio-technical, and conceptual factors. Technological design intended to incorporate privacy—not an attenuated proxy—must take such factors, in addition to the technological, into consideration.

Human factors science or engineering is the study and development of technology that systematically accounts for individual human capabilities and characteristics. A field with a relatively long history, it includes the study of psychological, emotional, cognitive, behavioral, physiological, and anthropometric characteristics relating to artifacts, as well as design and engineering practices. In computing, it has manifested in studies of usability and HCI. There is ample evidence that systems, mechanisms, devices, etc. failing to take human factors into consideration may also fail in their aim to protect privacy. Human factors of particular relevance to privacy include cognitive competencies,

preferences and desires. Human responses to technology turn on interpretation and derived meaning. For example, a technical constraint can be understood as a barrier or as a challenge. Understanding the signals sent by technical constraints is important to privacy by design but technical rules without moral legitimacy are merely practical constraints, deserving workarounds, not respect. Human factors researchers use a diverse range of methodologies including ethnographic and other qualitative (e.g. interviews and focus groups) approaches, quantitative (e.g. surveys), and lab-based as well as "in the wild" experimental studies.

Conceptual work is necessary to ensure that our privacy efforts are appropriate aimed and tailored to address privacy concerns. Systems, mechanisms, or metrics aiming to protect, measure, foster and promote privacy might strike a target, but to succeed it must be the *correct* target. Given the ambiguity of privacy and its contextual variance, developing technical tools that support both conceptual work, and conceptually informed practical work is a priority. Or, put another way, modeling privacy, and not an approximation of it or some other value altogether, requires a sound conception of privacy. Conceptions of privacy as control over information, or as secrecy dominate the computer science and policy communities. To the extent that critics of these conceptions are correct, the failure of these systems represents a theoretical or conceptual error not a scientific or technological one. Adherence to these flawed conceptual targets is an important source of failure in both technical and policy design. Advancing privacy in technical work requires us to understand how software engineers and participants in technical standard-setting processes consider privacy during design and how those concepts compare to user perspectives and the proposed technologies, as well as other sources of privacy norms. Identifying and reconciling dissonance between users' and developers' mental models of privacy is key to addressing privacy in the software development process. Privacy will both mean and require different things in different contexts, however the process of identifying both meaning and action should be stable at some level of generalizability. Articulating the specific concepts of privacy at issue in different contexts is essential to ground engineering and other efforts to protect privacy. Such work provides a detailed basis for understanding what specific mechanisms can support a given theory of privacy within a given context and socio-technical system. One example from the 2013 PCAST "Designing and Digital Future" Report is instructive. The Report notes that for NIT research, "The area of social computing augments the study of individual human-computer interaction with research in social collaboration and problem-solving in a networked, online environment." Likewise, in the increasing networked, online environments, assessment of privacy should not focus solely on individual risk but networked risks encompassing groups and interpersonal connections. Privacy is a social and political concept, not merely an individual preference or desire, or technical affordance. In order to enhance privacy assessment capabilities, privacy should not only be construed as a property or decision of individuals, or systems, but as a property of interpersonal and group connections, specific contexts, and communities. The importance of conceptual fidelity to this effort cannot be over stated.

Exploring the **legal, institutional, and economic aspects** that shape privacy in various contexts is essential to meaningful privacy work. The telephone sitting on your desk is able to mediate communication at a distance not only because of a particular

configuration of hardware and software but a complex telecommunications system coupled with institutional, economic, and political arrangements and cultural and societal norms. Properties, such as privacy, security, and so forth, may be due not to features in the artifact alone, but emerge from their situation within the larger network of dependencies—"socio-technical" systems, as they are sometimes called by a growing community of social scientists and philosophers in their studies of computer and information systems. Thus, whether technical systems protect privacy or undermine it can be a function of their situation within larger systems. Taking human factors into account means designing with individual human capabilities in mind; taking societal factors into account means considering the ways the larger societal systems shape, or ought to shape, properties of technical artifacts embedded in them. Privacy requires both. Giving due consideration to sociotechnical factors may be decisive to whether technical design achieves its aims.

These three lines of research—human factors, conceptual, and sociotechnical—combine with technical research aimed at improving the ability of designers to express privacy through technical design choices. This work requires tools and methodologies suited to the design process, preferably mapped to concepts of privacy, and/or specific contexts in which privacy arises. Ideally work in this area would aim to reduce the complexity of deploying the conceptual complexity of privacy introduced by work in the other domains through the development of common software design techniques such as vocabularies, patterns and threat models.

While social and legal solutions were determined to be out of scope in this RFI, it is essential to note that as with security in networks, and cybersecurity in particular, deploying whatever models, techniques, and tools developed through this initiative, and others, may well require incentives beyond those that exist today. While recent events have fueled renewed interest in protecting privacy across many domains, in practice, there are a range of actors that work to limit the adoption of privacy protective systems and technologies. To understand the power of technical progress on privacy to effect genuine protection it is crucial to situate it within the context of a broader social objectives and incentives.

Assembling a Field and Community: Privacy in and by Design

The concept of privacy-by-design (also discussed as privacy in design) has begun to gel a previously disparate and fragmented group of researchers who share the commitment to technological progress through multidisciplinary collaborations. (etc.)Today, there is a limited, disparate, and fragmented body of research attempting to pull privacy fully into the technical domain, some affirmatively positioned as privacy-by-design. One finds a broader range of related work in privacy-by-design looking across various disciplines. However, the work remains relatively disconnected and siloed within various sub disciplines of computer science, law, and ethics. Within and across those disciplines and sub-disciplines there are divergent views both about objectives regarding the appropriate concepts of privacy and related privacy requirements and the scope of the problem space.

To build the multi-disciplinary community required to advance privacy through technology it is necessary to bring together these diverse researchers, as well as the community of privacy practitioners. The *Computing Community Consortium (CCC*) has funded an upcoming Workshop Series on Advancing Privacy by Design toward this end.

The series includes four workshops: 1) documenting the current privacy design ecosystem of practice, and work in a range of academic disciplines to make privacy concepts more tractable in design; 2) a deep dive on engineering research and practice, focusing on who uses (or is anticipated to use) various technical approaches and instruments, what training or skills are viewed as essential for sound use, how their utility is (or could be) evaluated, and how well these results respond to the concepts and open problems identified in the prior workshop; 3) will cover the latest research results in user interface design, usability and human factors, again mapping current approaches, tools, motivations, methodologies—looking at practice and research—to consider how well they address real-world problems as framed by various stakeholders; and, 4) examining how existing regulatory models, along with other factors, shape organizations' understanding of privacy problems, approaches, and solutions, and how it might foster privacy-by-design approaches.

The workshop series will broaden the lens through which privacy-by-design is viewed by the research community—positioning technical design along side theoretical/conceptual, organizational, and regulatory design questions. Building on the insights gain through the workshops we will identify open research questions and offer a map to orient and facilitate communication, collaboration, and integration across the diverse privacy research community. The results of each workshop and the ultimate roadmap will be widely disseminated.

The current organizing committee includes:

- Deirdre K. Mulligan (Chair), University of California, Berkeley
- Annie Antón, Georgia Institute of Technology
- Kenneth A. Bamberger, University of California, Berkeley
- Travis Breaux, Carnegie Mellon University
- Nathan Good, Good Research
- Susan Landau, Worcester Polytechnic Institute
- Helen Nissenbaum, New York University
- Ira Rubinstein, New York University
- Fred B. Schneider, Cornell University
- Peter Swire, Georgia Institute of Technology
- David C. Vladeck, Georgetown University Law Center

* * *

In sum, a multi-disciplinary approach can ensure that technical design promote critical societal ends and values. Practices and research in the social sciences and humanities, as well as interdisciplinary fields such as human-computer interaction and science and technology studies can focus privacy research on the myriad potential futures represented by tomorrow's technologies. Social science and humanities research can contribute to an appreciation of future social and technological contexts where privacy concerns will occur. They can investigate how social and technological futures are mutually constructed, through sociotechnical imaginaries in science and technology studies, speculative design in human-computer interaction, and other refined approaches and methods in the social and humanistic studies of technology. Finally interdisciplinary approaches ensure that we consider a range of methods for protecting privacy as they reveal that threats to privacy as well as its protection can take form at many layers including: the user interface layer, the software layer, the hardware layer, and the Internet infrastructure layer. Even if we discover that protection necessarily requires policy, practice, law, and norms, a multi-disciplinary approach, investigating multiple layers of technology can help us to pinpoint the most promising and efficacious technical interventions, particularly if considered early on in development.

Sincerely,

Deirdre K. Mulligan
Associate Professor
School of Information
University of California, Berkeley
Faculty Director Berkeley Center for
Law and Technology

Helen Nissenbaum
Professor, New York University
Media, Culture, and
Communication & Computer Science
Director, Information Law Institute

With contributions from

Nicholas P. Doty and Richmond Wong, PhD Candidates, School of Information, UC Berkeley