Data & Society

Meeting Minutes: March 4, 2016

Attendees

Council Members:

Alessandro Acquisti, Carnegie Mellon University
Solon Barocas, Princeton University
danah boyd, Data & Society / Microsoft Research (co-PI)
Kate Crawford, Microsoft Research / New York University (co-PI)
Seeta Peña Gangadharan, London School of Economics and Political Science
Alyssa Goodman, Harvard University
Rachelle Hollander, National Academy of Engineering
Barbara Koenig, University of California, San Francisco
Arvind Narayanan, Princeton University
Paul Ohm, University of Colorado Law School
Frank Pasquale, University of Maryland
Matt Zook, University of Kentucky

Staff:

Emily F. Keller, Project Coordinator
Jacob Metcalf, Postdoc

Council members unable to attend: Mark Andrejevic, Geoffrey Bowker, Edward Felten, Eric Meslin, Alondra Nelson, Helen Nissenbaum, Latanya Sweeney, Sharon Traweek

1. Industry & Research Discussion – led by danah and Kate

- What can those at the intersection of industry and research teach us about ethics?

The Council brought in and engaged a panel of industry leaders to discuss how the data ethics challenges identified by the Council play out, or do not play out, at companies. The guests spoke to the Council anonymously. The issues discussed ranged from current and potential ethics oversight mechanisms to the risks and benefits of data sharing and collaborations.

Ethical oversight varies by company. However, product and business decisions, as well as research projects on social issues, are primarily focused on building the best product rather than being centered on ethics or public benefit. Industry research on sensitive populations is often reviewed by legal, public policy, or engineering leaders, though defining what constitutes sensitive material can be a moving target based on current events. Some data ethicists and legal scholars have proposed going further by creating ethics review bodies in industry to serve a function similar to IRBs in academia. Challenges to introducing ethical oversight mechanisms at companies include the difficulty of fitting emerging research into predetermined categories, resistance to the application of one strict set of standards to all departments or teams, and the need for flexibility to accommodate the fluid and fast-moving nature of research projects and methods. Enforcement of voluntary processes could also be difficult. Without external or legal repercussions, internal ethics review mechanisms could be difficult to enforce and may end up morphing to support new research projects or ultimately breaking their own rules.

Collaboration between academia and industry takes various forms, including data sharing. Academics frequently request corporate data for research, leading to company discussions about potential ways to release their data to contribute to the scientific community. Discussions include looking for ways to allow people to run queries and get answers without seeing the data. Ethics training in academic programs - or lack thereof - also influences industry, as many researchers have noted that a dearth of

RESEARCH ☩ INSTITUTE

ethics education in computer science programs may affect professional decisions down the road. However, caution prevails, as companies have faced public scrutiny in the past for decisions to share anonymized data that unexpectedly led to controversy. Even after being scrubbed, sensitive corporate data is uncontrollable after release and can lead to identifiable information and legal problems following combination with other data sets. Benefits must be weighed against risks.

At some companies, the only way to access their data is to work there. Students frequently work as interns at companies, where they have strict rules limiting their data access and restricting sharing with others. However, professors have cited instances in which students overlooked or misunderstood these requirements, which underlines the importance of oversight for responsibility and accountability in cross-sector data sharing. Restricting the data release to an IRB-reviewed academic purpose and prohibiting combination with other data sets is one possibility allowing greater control, but concerns remain about feasibility or potential backlash against the research community if a perceived violation of trust or privacy results. Vetting researchers for individual data release can also lead to perceived bias if elite institutions have greater leeway in gaining access to data sets than other researchers or members of the public.

2. Demo of Online Ethics Center – led by Rachelle Hollander

Council member Rachelle Hollander demonstrated the redesigned website of the [National Online Ethics Center](#), a collection of case studies and other pedagogical resources hosted by the National Academies of Engineering (NAE). Hollander's group at the NAE recently completed an overhaul of the website to include searchable metadata. She encouraged Council members to submit case studies, and the site will host the case studies solicited by the Council.

3. Future Research Questions – led by Kate and danah

- What are the big open research questions that need to be asked?

The meeting moved to a discussion of the collaborative white paper currently under draft. The white paper is an output promised in the NSF award as a capstone. It will present the Council's top-level recommendations for the future of data ethics scholarship, pedagogy, and policy.

First the Council discussed what they believed to be open research questions that should be taken up by data ethics scholarship. The white paper draft in circulation listed the following areas:

- Is human data science human-subjects research? There are some perspectives out there, but this problem requires more sustained conversation.
- What are the quantifiable risks posed by correlative data research? Does it differ with public datasets?
- How should data/computer/security scientists approach illicitly gained, publicly available data?
- What are the options for self-regulation in data science? What needs to be done with paper committees?
- What resources are needed in the university context to encourage engagement with data ethics issues, particularly outside of the IRB? (Again, reference Bonnie & Emily's work.)

Kate Crawford asked Council members to discuss these points and add to the list. She raised the first question: How do we think about power imbalances in ethics? Frank Pasquale responded that it seems in order to protect privacy you have to restrict the data to just one or a few places, which limits the utility of data reuse and repurposing. This tradeoff suggests the need for a more robust form of self-regulation. He also cited a study that showed the more you make it difficult to transfer or share data, the more companies will merge. This has already occurred with hospitals to reduce the costs of managing electronic medical records. Seeta Peña Gangadharan suggested that we should be looking at informal practices of institutions regarding how they are making ethical decisions.

Solon Barocas responded that there's a big difference between debating the right way to achieve an

objective (means) and debating the objective (ends). Is the discourse we call "ethics" the right tool to reason about the objective (justice, equitable distribution of power, etc.)? How do we get to this objective in a way that's defensible, particularly when the means-ends distinction is internally structuring the debate? Because privacy has been perceived as inadequate to reach such deeper objectives, there appears to be an evolving, widely held perception that we can pick up the slack with "research ethics." Research ethics may just be a means, not an end. Alyssa Goodman responded that it sounds like a two-step process: one is ethics, the second is research data ethics.

Rachelle Hollander asked: what are the drivers of the development of research in big data? We should expect those drivers to be at least partially a matter of engineering culture and be limiting factors in the development of ethical processes. Crawford noted that a discussant in the earlier industry panel pointed out that we didn't talk about money, yet people have to show that ethics work will improve the bottom line. We don't have a lot of literature that talks about the money and ethics conundrum. Matt Zook said we need more research on the political economy of big data. People could make a lot more money working for a hedge fund and yet find other endeavors more satisfying. Pasquale suggested that we see a 'conservation of ethical energy' where something like ethics that seems totally off the table as bad for the bottom line will become the norm over time. Goodman likened that to the rise of carbon offsets.

Pasquale also pointed to the power of metaphors, such as "the privacy dinosaur will go extinct" as an important driver in ethical thinking. Crawford asked: what are the messages and metaphors built into design changes? She also noted that as a group we have not had much discussion on the ethics of user experience (UX) design. Hollander noted that concern about reputational risk is very powerful driver alongside financial bottom lines. Crawford noted that public opinion includes outrage. Barocas asked whether UX design decisions are driven by research, and not politics and/or subjective preferences? Research ends up covering over the politics, which is how a veneer of scientific research makes it hard to talk about ethics. Presumed neutrality makes deep ethical questions appear to be no longer up for debate.

danah boyd said that 'slippage' between presumed neutrality and political contestation is part of a broader story, implicated in the easy movement between academic and industry research. Zook likened this to post-war cartography that emphasized the 'neutrality of map.' Crawford pointed out the continuing problem of 'data' being presented as neutral and beyond critique. boyd called back to the earlier conversation with industry scientists, in which one speaker said that engineers are aware of these philosophical-level problems that get lost in layers of interpretation and reuse. Obsession with control over interpretation may or may not be warranted.

Goodman discussed how IRBs have become ossified things, despite their origin as peer groups. She indicated there is a danger of that happening again as we pursue IRB-like options for data ethics in industry. Crawford noted there is a lot of interest in ethics in tech but there is a conflict between the drive for a more formal model and the drift toward just having long discussions. boyd noted that researchers are surprised when you can actually call the IRB and talk to them in alternative and independent IRBs. Gangadharan emphasized the need for alternative processes for engaging in ethics and accountability conversations outside of a traditional IRB that consists largely of professional peers. Community IRBs place researchers in a different community of practice. Barbara Koenig pointed to efforts that link data between genomics and health records in the "pseudo-not-for-profit" health care sector, which has different constraints and attention to bottom line. What alternative public governance could we think about in an age of democratic collapse? Crawford asked: why would people trust alternative mechanisms? Who gets to be consulted matters significantly, like making something for the disability community. How do you bring in effective communities, and what mechanism do you use? There are established models for situations like this, such as the Negotiated Rule-Making that is adequately effective at the federal level. These new companies are like the new governance agencies. Gangadharan stated that we should also consider the local level, separate from the more expansive model of deliberation. Koenig noted the extensive discussions about who gets to sit at the table in discussions about tribal health research.

Hollander noted Donna Reilly, a professor of engineering and public policy, is doing research on "shadow codes," a potentially useful concept here. Core concepts that are absent from explicit codes

may be absent because they are intentionally or unintentionally left out as implicit knowledge. How do you encourage an alternative community who might have important information on what should be included and why?

Chaitan Baru raised the question of how we should account for missing data and how that affects the social justice consequences of research. It is often imagined that research done with Twitter and Facebook data is broadly representative, but it simply is not because they do not reach all demographics evenly. If we think of this as a formal problem in data science, we can see people have their own networks. Crawford framed this as how we should account for missing data and communities.

Barocas suggested conducting a closer examination of slippery concepts such as 'sensitive' and 'creepiness' in order to develop actionable or measurable definitions. 'Sensitive' in the context of political issues could be seen as meddling, or refer to certain information items. Without care, such terms risk conveniently collapsing all problems into heuristics and replacing more robust concepts of autonomy or beneficence. Crawford asked why 'creepy' has become a go-to term? Koenig likened it to the so-called 'yuck factor' in bioethics, a loose concept about instinctual dislike of certain technological alterations of human bodies or behaviors.

Paul Ohm discussed a recent conference on the history of professional ethics. He said we need room to study data panics and their sociologies and discourse. Companies want to survive data panics while advocates' goals are unclear.

Pasquale raised the matter of professions that are said to be heading towards obsolescence because even complex interpretive work will be done by big data analytics, and he pointed to a start-up that tries to predict how courts will rule. This suggests a viable research question: are predictive analytics having a positive or negative effect on professional ethics? Crawford said that analytics can improve or erode ethics, or both.

Barocas raised Michelle Mayer's argument that we also have to ask about the flipside of unethical experimentation—sometimes not experimenting is ethically troublesome, which she argues is often the case in big data. Crawford noted that experimentation appears to be increasingly detached from the familiar model of the scientific method. boyd responded that most machine learning engineers will say their discipline is not about scientific method, it is all observational. Crawford noted that this may be the case, except for automated experimentation.

Goodman noted that even for science that has nothing to do with the public, like her research in theoretical physics, she doesn't agree with extreme views about a radical change in scientific method being a license to mess around more just because we can. There appears to be an erosion in ethical safeguards even on a yearly basis. Barocas stated that those who use experiments generally want to tease out causality. Goodman noted that figuring out causality can raise ethical issues.

Alessandro Acquisti noted he is interested in how IRBs make incredibly different decisions in different contexts. Even with straightforward academic IRBs, some proposals have no problem in one IRB but get stuck in another. This creates a sense of uncertainty about how the process really works. Jacob Metcalf noted that an IRB is a unique institutional structure, part of a governmental mandate for self-regulation within institutions that leaves a lot of discretion to the institutions themselves. There are some meaningful moves towards having greater consistency between IRBs in different universities, such as allowing researchers with cross-institution proposals to choose a single IRB.

The conversation closed around a discussion of what collective practices are best for sustaining ethics conversations. Examples included debates, guidelines, and consulting.

4. Paper, Network, and Final Council Outputs – led by Jacob and Emily

The next session began with Chaitan Baru notifying the Council about the BIGDATA program PI meeting at the NSF on 4/20 and 4/21. There will be several hundred PIs, including those in Big Data Innovation Hubs. There will be a slot on the agenda for the Council to do an ethics presentation.

The conversation turned to the remaining aspects of the white paper and network expansion effort.

First the Council discussed core lessons from the morning's conversation with industry actors. boyd noted the informal and haphazard way ethics is being dealt with in industry. People often learn about ethics from talking to people in their community, not through formal processes. Goodman observed that ethics becomes a luxury component once companies have adequate resources and more prominent vulnerabilities. Hollander suggested that we can see signs of significant interest from industry, and so collaborations might be more doable than previously expected.

boyd asked: who are the bridging translators out there who can think across ethical and technical issues, and between academia and industry? The Council brought in some excellent bridges, but how many others like them are out there? The conversation indicated that it takes some institutional knowledge to know where and how ethics conversations are happening. It may be the case that some people in major organizations are completely unaware of ethics processes occurring elsewhere in the organization, but there can be significantly different cultures within teams and clusters. How one cluster of employees acts is not indicative of the entire organization. In some cases, people may even be violating the explicit or implicit rules of their employer in order to establish a community with other professionals interested in ethics inquiry. There is also clearly some value in using external pressure through journalism and critique to encourage ethics leadership within companies.

Gangadharan observed that there can be major differences between big and small companies, and that the available resources affect the frame of the ethics conversations that can happen. Crawford said she worried that this conversation is about elite companies with money to burn, and most companies don't have time to interface with an ethics advisory board. boyd noted that general counsels play a heavy role in these conversations. To what degree are they the right point of intervention instead of engineers themselves? Institutional culture makes it hard to go through legal oversight. In New York, a handful of counsel manages most start-ups and this may be a route into more encounters with multiple companies at once. Ohm noted that there is a wave of different approaches in law school. There are 500 privacy lawyers in D.C., and a lot of lawyers have had privacy exposure. Crawford noted that this is a nice augmentation of our pedagogy model.

Gangadharan said we should be thinking about educational spaces non-traditionally. We should also think more broadly about the different kinds of publics that can be engaged—if we exclude communities outside of the normal scholarly interlocutors them then we haven't been true to the flexibility of a radically changed system.

boyd raised the matter of how to leverage privacy conversations into conversations about ethics. Privacy issues that started with big tech moved into the health sector fast. If we couldn't scaffold ethics into moving fast, can we tie them into privacy? Crawford considered the relationship between existing privacy structures and where ethics might fit in, which might be a little hard. Privacy has become a legal cover in many cases, but there are lessons to be learned. boyd pointed to the maturation of Chief Information Officer (CIO) as a possible point of entry within many companies. Baru stated that CTOs may get more of that power because in industry they are becoming the infrastructure people. Maybe a group like this could engage with them. Security may also be a feasible route into industry conversations. Barocas noted that theme recently came up at a workshop he organized at a data mining and search conference, in part because security researchers and professionals have a long-running ethics discourse with fairly mature norms.

5. Future of the Council and Ongoing Collaborations – led by Kate

The meeting closed with a brief discussion of the Network expansion effort. boyd asked to what extent such a network should be a goal as the Council moves to new efforts. Crawford and Geof Bowker are pursuing some leads for continuing the Council, and a Network may be useful for that. It may also be useful for establishing a community interested in an annual or bi-annual conference.