

# Data&Society

Meeting Minutes: October 16, 2015

## Attendees

### Council Members:

Solon Barocas, Princeton University  
danah boyd, Data & Society / Microsoft Research (co-PI)  
Kate Crawford, Microsoft Research / New York University (co-PI)  
Alyssa Goodman, Harvard University  
Rachelle Hollander, National Academy of Engineering  
Barbara Koenig, University of California, San Francisco  
Arvind Narayanan, Princeton University  
Helen Nissenbaum, New York University (co-PI)  
Paul Ohm, University of Colorado Law School  
Frank Pasquale, University of Maryland  
Sharon Traweek, University of California at Los Angeles  
Matt Zook, University of Kentucky

### Staff:

Emily F. Keller, Project Coordinator  
Jacob Metcalf, Postdoc

*Council members unable to attend: Alessandro Acquisti, Mark Andrejevic, Geof Bowker, Edward Felten, Eric Meslin, Alondra Nelson, Seeta Peña Gangadharan, Latanya Sweeney*

## 1. Status Update & Q&A – led by danah boyd

danah boyd led a conversation about how the discussion of ethics in computer science can be perceived as a judgment or assertion of control from outside, and the importance of trying to understand practices within the field, such as the elements of teaching and tradeoffs. This was particularly acute following the Facebook emotional contagion study, when researchers felt some critiques were not based on an informed understanding.

Members said that ethics conversations have to come from within the field and be based on an appreciation of what the work entails and meeting people on their own terms. Historical ethics standards from law and medicine emerged from within those professional communities. In computer science, there is resistance to “ethical experts” from outside the field, particularly in communities that work closer to people than machines. Computer science is a relatively young community. Pressure from outside can help spur ethical discussions in the field, but it is a mistake to assume that computer scientists do not want to discuss ethics.

“Ethics” is often discussed under other terms particular to each field. “Reputational risk” is a term related to ethics that is understood by legal clients. Journalists think about broader impacts. Developing ethics codes internally can help prevent impositions from outside. Addressing ethical issues as a community issue versus pointing the finger at others to find “rule-breakers” will advance the ethics conversation further. However, when someone violates community norms, are there situations in which they should be shunned?

## 2. Workshop based on Matt Zook’s draft PLOS article: “Ten Simple Rules for Ethical Big Data Studies” – led by Kate Crawford and Matt Zook

Matt Zook asked the Council to come up with examples or citations to support the concepts represented in his draft PLOS article. He summarized the rules as:

1. Putting big data in a larger context.
2. Lots of things about data are human.
3. If you're not tending to your own ethics in your research, someone might come along and impose from outside.
4. Argue that ethics makes better research, such as the reputational risk issue noted earlier.
5. Bring voice to these decisions and make sure you have these conversations.
6. A lot of big data work is being done outside the academy.
7. A call to have people talk across disciplines.
8. When does public become private, especially when combining data sources?
9. Once collected, big data sticks around a long time.
10. Ethical standards are coming. If communities don't come up with their own from within, regulations will be imposed from outside.

Council members suggested providing a definition of big data while inviting others to provide their own, and offering a complex definition of ethics, including by adding the words: "consider whether or how this research should be done." Additional suggestions were made to add more examples rather than scholarly links, to reach commercial actors, including data monopolies, and to provide precise, practical advice that assumes the audience is thinking about ethics already. Microsoft is running an internal group on ethics and machine learning and Kate Crawford will look into the AT&T archive. Alyssa Goodman noted that her PLOS article received 30,000 views but no citations.

Next, the Council members drafted new versions of the rules in three breakout groups with minimal use of the word "ethics." They considered changing the title to "Ten Simple Rules for Responsible Big Data Research." Ideas were discussed about identifying individuals or groups who are impacted or represented by data, knowing when to break a rule, considering what is counted as data and why, distinguishing group harm (such as environmental impacts) from the traditional focus of human subjects protections for individuals, and recognizing that attending to ethics is good business sense and can help avoid reputational harm. Council members noted that ethics are about conversations and processes rather than clear right/wrong outcomes. The members discussed replacing general principles with actionable, prescriptive items for use in the field.

Council members provided specific editorial ideas, such as replacing, "Acknowledge that big data and human subjects are deeply interlinked" with "Identify how your research and its results can help or harm individuals or communities." They suggested noting the importance of diversity in the section on inter-disciplinarity, suggesting that people talk to others with different values or backgrounds. The piece could also acknowledge sustainability conversations in industry, such as the serious environmental costs of doing data, as measured in energy and water usage; and the impact of Google allowing people to store data indefinitely. Regarding the conceptual approach, the piece could highlight that the act of interpretation, beyond the collection and processing of big data analytics, is a human process rather than an automatic one, and there are implicit judgments in data analytics. The Council should determine whether the rules should target groups, individuals, or types of people.

Several Council members asked the group to think about when it is appropriate to violate the rules, such as to save lives in a crisis or public health emergency, to provide educational opportunities to children, or to give a bribe to get meals to people in corrupt countries. The members said perhaps rule #10 should communicate the need to think about when to break the other rules or how to not mistake them for a checklist. The rules should be tactical and usable in the moment for people to think about when working with a dataset or starting work in industry or academia, and the tone should focus on helping researchers improve their research versus saying they are doing something wrong. Rule #2, "acknowledge that big data and human subjects are deeply interlinked" is very abstract and could be changed to: "remember that your data can be used in ways you didn't anticipate." The opening paragraph should specify that the rules are targeted at human-centered data, but some data, such as urban informatics and sensors in the built environment, have human impacts even though they are not about humans. The rules should be somewhat provocative, punchy, and controversial.

### 3. Revisions to the Common Rule – led by Jacob Metcalf

#### Overview:

The Department of Health & Human Services released a "Notice of Proposed Rulemaking" for revising the Common Rule. The Council aims to submit an official public statement to the HHS on the proposed changes by late November. See HHS's official summary:

<http://www.hhs.gov/ohrp/humansubjects/regulations/nprm2015summary.html>

#### Discussion:

This portion of the meeting concerned the proposed revisions to the Common Rule. After publishing a Notice of Proposed Rulemaking in September, the HHS opened a 90-day window for public comment. The Council discussed submitting a collective public comment regarding the proposed revisions, and decided that it was worthwhile to draft a narrowly tailored letter to be collectively edited and signed individually. Jacob Metcalf wrote a pre-meeting briefing on the proposed revisions and hosted the conversation along with presenting slides.

A number of changes are big data-related without directly addressing the issues most specific to big data research methods.

For example, Council members first discussed the proposal to reclassify biospecimens as human subjects data. Until recently, biospecimens taken in research or routine medical care were treated as de-identified material that did not require consent to use in research. Responding to research demonstrating that techniques to de-identify biospecimens are not as reliable as once assumed, the revisions propose to newly treat biospecimens as identifiable material from human subjects. This change is paired with the proposal to require biobanks to receive consent for research re-use, but allows them to receive blanket long-term consent with requiring re-consent for each use. Together, these are arguably the most significant and contentious revisions. Despite resulting from big data research techniques that make re-identification increasingly plausible, the Council decided to not directly comment on this issue because many other biomedical groups will be offering input.

The discussion moved to the new rules concerning exclusion and exemption from the Common Rule. Currently much non-biomedical research falls under the category of "Exempted," which means that the project is only reviewed by IRB staff (not a full review panel) and typically does not have continuing supervision. In order to address long-standing complaints from social scientists that their work is a poor fit for regulatory practices oriented toward the biomedical sciences, the HHS is proposing to create a new category of "Excluded" to encompass research that poses such minimal risk to human subjects that it does not require any oversight. Additionally, the HHS proposes to significantly expand the "Exempted" category (minimal oversight) and conduct the exemption determination via an online portal, resulting in most non-biomedical research never being seen by an independent reviewer.

Conversation turned largely toward how the new Excluded and Exempted rules address big data research techniques. In particular, the revisions expressly assume that research which makes secondary use of pre-existing research data or personally identifiable non-research data is inherently low-risk to human subjects. Council members expressed significant concern about that assumption, which may have once been sound but is much more questionable in the age of big data. Whereas the primary risk of exposure of private data was the exposure itself—and therefore once made public, any harm is assumed to already be done—the utility of big data research and practice is that existing data can be put to inherently unpredictable secondary uses. Therefore, the public status of datasets is no longer an accurate proxy for the risks posed to human subjects.

Council members agreed this assumption should be the primary concern to focus on in the public letter. Kate Crawford pointed out that although there will be some confusion and disagreement about how these policies will play out, it is important to flag an erroneous foundational assumption as troublesome regardless of the downstream outcomes.

Jacob Metcalf will write the first draft of the letter and circulate it for comments among the Council. Members will sign the letter individually.

#### **4. White paper and future of the Council – led by danah boyd**

##### **Part I**

The Council's long-term goals and responsibilities include the creation of a white paper summarizing its work and findings. Council members said that one key audience for the white paper will be NSF, but it should also be written for use by academics and practitioners. Any recommendations that the Council provides to the NSF must be made public.

Members suggested the following topics and sections:

- Questions that have come out of big actionable research.
- Suggested revisions to NSF data management plans (DMPs) and associated guidance for researchers.
- How big data techniques affect ethical responsibilities of research.
- Injustice and environmental issues.
- Historical framing of how the Council came together, inspired by emerging problems such as scandals of the last five years and the areas they affect.
- What should individual data scientists be thinking about?
- Flesh out what ethical challenges look like, such as stigmatization (being in or out of research) and privacy.
- How do we judge the ethics of algorithms or hold them accountable (example: who is responsible when a Google Search algorithm shows an anti-Semitic website?)

Council members brainstormed their collective contributions as members to issues of data ethics in the field. They suggested a diverse group like themselves should convene every few months, like the FCC of data science, to discuss issues with no precedent and no clear oversight body on a case-by-case basis, and communicate the results. The RAC was put into effect to review protocols for recombinant DNA, and committees of experts study cloning. What are the carrots and sticks, questions that can help program officers, and levers at the organizational level that can have a role in ethical guidance, such as funding for research and convening powers for things like special committees? Possible structures include that of the Council and the NSF Research Coordination Networks. What do program officers supporting big data need to recognize about ethical questions and how can NSF make a difference in ways that go beyond proclamations, lists, and regulations, such as proposal requirements and ways that PIs come together?

Barbara Koenig asked if the NSF should fund experiments in alternative forms of governance or research practices for a different way of getting participants' voices in. Jacob suggested drawing on analogs from biomedicine. Arvind Narayanan suggested thinking about the mechanics of the funding process and opportunities at NSF and other funding agencies to insert ethical review and guidance. NSF annual reviews are one potential method to raise consciousness retrospectively by providing feedback on published papers, and annual PI meetings are another opportunity for discussion. Alyssa Goodman, who has been on the National Academy Board of Research Data and Information, noted that boards are supposed to spur committees. Matt Zook suggested putting together a series of questions to the effect of: "what are the ethical implications of your study?" with a goal of discussion rather than the pro forma nature of DMPs. Helen Nissenbaum and Jacob Metcalf suggested trying to encourage the NSF to include ethics in its Broader Impacts criteria.

Regarding white paper content, Paul Ohm offered to work on a section about interdisciplinarity, which raises time-consuming vocabulary issues. Sharon Traweek suggested having a section with one paragraph per member on how serving on the Council has altered their work and perspectives. Solon Barocas noted there is a dearth of writing oriented towards practitioners in ethical scholarship and this could be well received by people who are teaching.

Council members said there is a lot of money available for teaching data ethics and data science but varying opinions or turf wars over who should teach it – computer science, engineering, information studies, etc. Arvind Narayanan suggested that the NSF could assist with building a technical infrastructure at Princeton to detect online discrimination. Solon Barocas said an interdisciplinary grant that brings social scientists and engineers together could work.

## **Part II**

danah boyd asked the Council to think about possible future models and who would want to take on a leadership role when her time as lead PI ends in March 2016. Kate Crawford said that she and Geof Bowker are interested in serving as the lead on a rotating basis, beginning in 12 months. Paul Ohm suggested that a colleague of his might be interested in taking it on. Kate Crawford suggested that foundation funding could supplement NSF funds to lighten the administrative load. Alyssa Goodman suggested looking at the option to serve as a committee, possibly through the Research Data Alliance. Rachele Hollander suggested that something like the National Research Council at the National Academies could be a good model, as agencies can create boards like these. Staff can also raise money to form consensus studies, which make recommendations and publish reports. Barbara Koenig noted that she has been a part of the Roundtable on Translating Genomics-based Research at the National Institutes of Health, and the institute is also pushing the Big Data To Knowledge (BD2K) program and a committee to re-envision the future of medicine, chaired by NHGRI Director Eric Green, that could potentially have some overlap. Discussions will continue between members offline.

*\*Funding for the Council for Big Data, Ethics, and Society was provided by the National Science Foundation (# IIS-1413864).*