

## Opinion

## Who Should Do Data Ethics?

Caitlin D. Wylie<sup>1,\*</sup><sup>1</sup>School of Engineering and Applied Science, University of Virginia, Charlottesville, VA, USA\*Correspondence: [cdw9y@virginia.edu](mailto:cdw9y@virginia.edu)<https://doi.org/10.1016/j.patter.2020.100015>

**Who decides what good data science looks like? And who gets to decide what “data ethics” means? The answer is all of us. Good data science should incorporate the perspectives of people who create and work with data, people who study the interactions between science and society, and people whose lives are affected by data science.**

The launch of a new journal, like this inaugural issue of *Patterns*, can be part of a defining characteristic of a scientific field. For example, the papers in *Patterns* may help to define data science. As such, a new journal offers an important opportunity for practitioners to consider what it means to do their work well. So, who should decide what *good* data science—i.e., that is beneficial for society and produces useful, reliable knowledge—should look like? The short answer is all of us. Good data science should incorporate the perspectives of people who create and work with data, people who study the interactions between science and society (as I do), and people whose lives are affected by data science.

One might think that it's up to the government to define good data science. After all, governments regulate research in many ways, including by licensing practitioners, accrediting training programs, limiting research on people and animals, and defining intellectual ownership through patents and copyrights. However, regulation is usually reactive rather than proactive. As the classic Collingridge dilemma explains, the social effects of a technology (or, in the case of data science, a method and mindset of doing research) cannot be known until that technology is widely used. But once a technology has become widely used, regulation is less effective at curbing its negative outcomes.<sup>1</sup> This conundrum limits the power of regulation to minimize harm from emerging tools and techniques. A government must either try to predict harms through anticipatory governance or wait until the harms become clear and then try to limit them. Clearly, we can't rely on these

problematic approaches alone to define or achieve good data science.

Furthermore, few people in government are experts in science or technology. The embarrassingly ignorant questions that some US senators asked Facebook CEO Mark Zuckerberg during congressional hearings in 2018 demonstrate the potentially enormous gap in knowledge between elected officials and practitioners of science and technology. Without guidance from experts (specifically, experts who are pursuing collective good rather than indiscriminately trying to block regulation of their field), government has little chance of successfully anticipating or reducing harm from emerging technologies.

Unlike the government's actions before or after the implementation of a new technology, data experts face ethical issues *throughout* their work of managing data and designing methodologies. How they respond to these everyday dilemmas influences the design, function, and social impact of their products. Because data practitioners understand the knowledge and systems they are producing better than anyone else, the primary responsibility for doing good data work must fall on them.

Traditionally, experts' conceptions of good work have been informed by discipline-specific professional societies, which issue codes of ethics. For example, in the United States, computer scientists are expected to follow the Association for Computing Machinery's (ACM) code, electrical engineers the Institute of Electrical and Electronics Engineers' (IEEE) code, and statisticians the American Statistical Association's (ASA) code. These codes alone are not

the solution because they are not laws, are rarely enforced, are specific to their disciplines, and are vague enough to be broadly interpretable. However, they serve as a professional community's declaration of what it means to do good work.<sup>2</sup>

But in the case of research—and journals—that claim to stand apart from any specific discipline or professional society, who decides what good work means? For example, *Patterns* Editor-in-Chief Sarah Callaghan celebrates data science's ability to exist between fields: “Domain boundaries are difficult to cross, but there are many exciting and fruitful collaborations and developments that can be found and nurtured at those edges. This is why *Patterns* is data agnostic and focused on the commonalities across fields. It's about building a community.”<sup>3</sup>

By “data agnostic,” I understand Callaghan to mean that *Patterns* welcomes datasets and analyses regardless of which domain—i.e., discipline, field—produced them. Similarly, in the founding editorial of the *Harvard Data Science Review*, Editor-in-Chief Xiao-Li Meng defines data science as “a collection of disciplines with complementary foundations, perspectives, approaches, and aims, but with a shared grand mission. That is, to use digital technologies and information of any kind to advance human society.”<sup>4</sup> I interpret these two editors' admirable visions to mean that data science draws from and contributes to many domains while remaining distinct from them. These editors specify that responsible research and social benefit are integral components of data science. But who will define these ideals when data practitioners lack a shared



disciplinary society, as the traditional arbiter of professional ethics? I worry that claiming that data science is agnostic to disciplines could invite assumptions that it is also agnostic to *ethics*.

First, what might it mean to be domain agnostic? (See Ribes<sup>5</sup> and Ribes et al.<sup>6</sup> for more.)

- (1) It can mean that ways of working with data can be applied to data from any discipline, such that the domain is not relevant to the success of these methods. In this sense, domain doesn't matter.
- (2) It can mean that data work spans domains, such that members of any domain can be equally capable of using and developing these methods. In this sense, data science welcomes all domains.

Ascribing to the first sense could invite the idea that ethics, as a component part of each domain, is irrelevant to data work, like the domains themselves. This is obviously dangerous. It also implies that data work—including datasets, analytic techniques, and infrastructure—is neutral and free of social values. This misconception obscures the rich social world in which people define, organize, manipulate, and interpret data in ways that serve their cultural, political, and professional context.<sup>7</sup>

Adopting the second sense of domain agnosticism could mean that anything goes in data science. Each domain is valued for its potential to improve mutually useful ways of working with data; therefore, each field assumedly should also import their ethics to this shared space of data science. But inviting all ethical views into professional practice could create a Wild West of unclear, conflicting values that practitioners do not share or consider important. Having too many conceptions of ethics is not as dangerous as considering ethics irrelevant, but it nonetheless offers insufficient guidance and opportunities for the community to reflect on what good data science should be.

Could there be, then, domain-agnostic ethics? Are there notions of good work that span fields and can guide how all practitioners work with

data? Perhaps, but it would be difficult to unite data practitioners enough to begin defining this kind of ethics due to their wide variety of training, job titles, and everyday work. People who share a field already share a set of knowledge and values, which is an important starting point for defining their community's beliefs about good work. Thus, a domain-agnostic space, with diverse kinds of experts, arguably requires *even more* ethical awareness and reflection than a domain space does, so that practitioners can figure out how to align their many co-existing ways of addressing ethical dilemmas.

I propose that the responsibility for serving society through good data work lies with practitioners—all of them, including people who produce, curate, analyze, and/or interpret data. Every data worker is an expert about how they handle data, which gives them crucial skills for thinking carefully about how their work affects the world. This makes them the best option we have to entrust with our social good. Of course, individual practitioners' judgments vary based on their personal values, experiences, and cultural context; nonetheless, their sense of responsibility and moral duty to society is the foundation for all professional ethics. If this were not true, then these experts would be merely "guns for hire,"<sup>8</sup> meaning immoral thugs with enormous power for harm, rather than professionals striving for social good.

But data practitioners should not bear this responsibility alone. The definition and achievement of good data work requires the input of a wide variety of experts, including those in social science, history, ethics, policy, law, and education. These experts can help data practitioners align their work with broader, more evidence-based conceptions of social benefit, while also gaining insight into their own areas of expertise. Just as data experts hope to learn from and inform various domains, many people in those domains hope to collaborate with data experts. Furthermore, data science deserves the input of those we might not consider experts, such as people who generate data by using the internet or just strolling past a CCTV camera and people whose lives are

influenced by data science, such as through algorithm-informed policing and banking and through new insights from interpreting data about climate, health, and social inequality. Including a wide range of stakeholders in discussions of good data work is the right thing to do, and it will also improve data science. After all, if data science strives to serve society, then it must welcome the wisdom of everyone it affects.

Once individual data practitioners have accepted the responsibility that goes with their expertise, the next step is to publicly explain their ethical worries and judgments, thereby initiating discussion and eventually developing a shared sense of how to define dilemmas, weigh possible responses, enact the best response, and evaluate the outcomes. Crucially, this communication would create the expectation that practitioners be transparent about (1) their concerns about a dataset or technique or product, (2) how they think through these concerns, and (3) how they decide how to respond. Transforming individual experts' worries and decision-making processes into open, valued, inclusive discussions would help build a common understanding among practitioners—and the rest of us—of what good data work is and who is responsible for its practice.

As a journal that welcomes data work from all domains, *Patterns* could be a powerful forum for these discussions. *Patterns* could deploy its domain agnosticism as a way to unite data practitioners as they embrace the responsibility to do good data work and to decide together, with help from the rest of us, what that means. After all, the world desperately needs good data science. Let's work together to achieve it.

## REFERENCES

1. Collingridge, D. (1980). *The Social Control of Technology* (St. Martin's Press).
2. Herkert, J.R. (2001). Future directions in engineering ethics research: microethics, macroethics and the role of professional societies. *Sci. Eng. Ethics* 7, 403–414.
3. Callaghan, S. (2020). Starting a New Pattern. Cell Press. <https://www.cell.com/patterns/editorial>.

4. Meng, X.-L. (2019). Data Science: An Artificial Ecosystem. *Harvard Data Science Review* 1, <https://doi.org/10.1162/99608f92.ba20f892>.
5. Ribes, D. (2019). STS, Meet Data Science, Once Again. *Sci. Technol. Human Values* 44, 514–539.
6. Ribes, D., Hoffman, A.S., Slota, S.C., and Bowker, G.C. (2019). The logic of domains. *Soc. Stud. Sci.* 49, 281–309.
7. Leonelli, S. (2019). Data - from objects to assets. *Nature* 574, 317–320.
8. Johnson, D.G. (2008). Computer experts: guns-for-hire or professionals? *Commun. ACM* 51, 24–26.

**About the Authors**

**Caitlin Donahue Wylie** is an assistant professor of science, technology, and society in the University of Virginia's School of Engineering and Applied

Science. She studies how technicians, students, and volunteers contribute to research in science and engineering. This topic includes who works in laboratories and what they do, who receives recognition for research work, how students learn to conduct research, and how researchers define expertise. Wylie teaches engineering students how to assess the social and ethical dimensions of technology so that they can design safe, equitable, and successful sociotechnical systems.