

A non-expert's introduction to data ethics for mathematicians

Mason A. Porter

*This chapter is dedicated to my current and former PhD students.
They mean the world to me.*

ABSTRACT. I give a short introduction to data ethics. I begin with some background information and societal context for data ethics. I then discuss data ethics in mathematical-science education and indicate some available course material. I briefly highlight a few efforts—at my home institution and elsewhere—on data ethics, society, and social good. I then discuss open data in research, research replicability and some other ethical issues in research, the tension between privacy and open data and code, and a few controversial studies and reactions to studies. I also discuss ethical principles, institutional review boards, and a few other considerations in the scientific use of human data. I then briefly survey a variety of research and lay articles that are relevant to data ethics and data privacy. I conclude with a brief summary and some closing remarks.

My focal audience is mathematicians, but I hope that this chapter will also be useful to others. I am not an expert about data ethics, and this chapter provides only a starting point on this wide-ranging topic. I encourage you to examine the resources that I discuss and to reflect carefully on data ethics, its role in mathematics education, and the societal implications of data and data analysis. As data and technology continue to evolve, I hope that such careful reflection will continue throughout your life.

*Don't say that he's hypocritical
Say rather that he's apolitical
"Once the rockets are up, who cares where they come down?
That's not my department!" says Wernher von Braun.*
(Tom Lehrer, *Wernher von Braun*, 1975)

1. Introduction

The use of digital data to examine and help understand human behavior is both powerful and dangerous [Edi21]. Every day, it seems like there is a new nightmare with problematic uses of data and algorithms. The use of predictive policing to identify criminal activity can exacerbate existing racial and ethnic

2020 *Mathematics Subject Classification.* Primary 97M70, 97B70; Secondary 97P80, 91C99, 91-11.

Key words and phrases. Data, ethics, social systems, human society, complex systems, mathematics, the mathematics profession.

inequities [Nat18], algorithmic social-credit ratings of individuals have frightening dystopian uses [Wikc], and other manifestations of “algocracy” (i.e., algorithmic government) have many human ramifications [ET19, OBS⁺19, Has21, SMBM23].

The ever-increasing use of and reliance on “big data” and the accelerating application of tools such as machine learning and artificial intelligence (AI) to make societally consequential decisions can cause very significant harm and exacerbate societal inequities [O’N16, Ree18]. Problems like algorithmic bias and the collection, measurement, and use of enormous amounts of data about humans and their behavior have significant societal consequences [WSAO⁺21, LHF⁺21, SVW21, Bir21, WJ23]. New technologies like “deep fakes” [IPA20], synthetic media in which a person in an existing image or video is replaced with somebody else’s likeness,¹ have terrifying potential to cause harm. There are significant risks and potential nefarious uses of tools like generative AI and large language models (LLMs) [BGMMS21, EMFG⁺22, Fer24]. Modern data analysis also has been accompanied by the reincarnation of pseudosciences, such as physiognomy, which can now be performed at a large scale [SH22]. Along with such dangers come enormous potential benefits, and there are nascent efforts to create documents like an “AI Bill of Rights” to help guide the design and deployment of automated systems in a way that protects people [Off22].

The situation is already scary. In 2016, Microsoft released a Twitter bot that “learned” from its interactions with other Twitter accounts; in less than a day, it was regularly producing racist tweets [Vin16]. In December 2020, Stanford University’s use of an algorithm to determine which people would receive the first batches of a COVID-19 vaccine resulted in an inequitable vaccination rollout that prioritized high-ranking doctors over frontline human-facing medical personnel [Che20]. The output of algorithms have also led to decisions to cut off medical care [RH23]. In 2021, the outgoing Editor-in-Chief of a scientific journal publicly released refereeing data, including the numbers of decisions to accept or reject papers by each of the journal’s referees during his tenure [Sch21]. In 2023, the LLM ChatGPT invented² a sexual-harassment scandal and named a real professor in it [VO23] and Microsoft limited its Bing AI chatbot after its unsettling conversations [Les23]. In today’s world, it often seems necessary to laugh to keep from crying [Wei21].

The mathematical, statistical, and computational sciences are interconnected with human communities and society at large [Cam, Eth, MC23, Mos21]. Neither our education nor our research occur in a vacuum, and many of our algorithms and other methods are now applied to social systems [WSAO⁺21, CM23]. However, it is not traditional in mathematical-science education to discuss data ethics and other ethical considerations [RSMCM24, Mü24]. What we choose to teach (and choose not to teach) impacts what our mentees do with their education. Consequently, it is crucial for mathematical scientists (and others) to think carefully about ethics and engage with it throughout their careers [CM24]. See [LMP18] for a short book on ethics and data science, [FT16] for a theme journal issue on data ethics, [Vé21] for a wide-ranging handbook of digital ethics, and [MCF22] for a recent

¹A light-hearted example is Gollum’s precious cover of the song “Nothing Compares to U” [Dor20].

²It seems that LLMs and generative AI tools often simply make things up.

discussion of the ethics landscape in mathematics and advocacy of a “Hippocratic Oath” for mathematics.

The present chapter, which I intend as an introductory resource about data ethics, is a companion to my oral presentation at the American Mathematical Society Short Course on Mathematical and Computational Methods for Complex Social Systems in January 2021 [BFPV21]. My slides and video presentation are available at [Por21]. I discuss many of the same key points in my presentation as in the present chapter, although my emphases often differ. I have also learned about more data-ethics resources in the last four years, and certain ethical issues and relevant technologies have gained increased prominence in that time. I encourage you to look at my slides and watch my presentation, and I especially encourage you to look at the resources that I discuss in those slides and in the present chapter. I am not an expert on data ethics—and, to be frank, writing this article has been accompanied by my most serious case of “imposter syndrome” in a long time—and it is important that you look at what actual experts have to say. I will attempt to give some helpful thoughts and pointers to begin a journey in data ethics. It is vital to continue to reflect on data ethics and the societal impact of data and data analysis throughout your life.

In this chapter, I cover a diverse variety of topics. I start with data ethics and education, which is important for almost the entire mathematical-science community. In Section 2, I discuss data ethics in the mathematics community and mathematical-science education. I highlight the importance of curricular data ethics in Section 2.1, and I point to existing educational material in Section 2.2. In Section 3, I describe some of our efforts at University of California, Los Angeles (UCLA) on data ethics and society. I also briefly mention a few efforts by others and point out an important caveat in efforts in “data for social good”. In Section 4, I discuss various ethical issues and tensions in research. Some of these topics, such as research replication and appropriately acknowledging others, are directly relevant to everybody in the mathematical-science community. Other topics are most directly relevant to people who use data in their research, although it is still important for others in the mathematical sciences to have some familiarity with them. I discuss research replication and open data and code in Section 4.1, acknowledgements and licensing material for open use in Section 4.2, privacy concerns and its tension with open data and code in Section 4.3, and some controversies in studies and in reactions to studies in Section 4.4. In Section 4.5, I summarize the key ideas of Section 4. In Section 5, I discuss some ethical principles and other considerations in the use of human data in scientific research. In Section 6, I highlight a few research articles about issues in data ethics. I conclude in Section 7.

2. Data ethics, the mathematics community, and mathematical-science education

In this section, I advocate for data ethics in mathematical-science education and point to some existing course materials. Many students who obtain undergraduate or graduate degrees in the mathematical sciences will become data scientists or otherwise will work extensively with data. Therefore, data ethics needs to be a central part of their professional lives, and it is our responsibility as mathematics educators to mentor them to face these ethical challenges with careful reflection, wisdom, and humility.

2.1. The importance of data ethics in the mathematics community and mathematical-science curricula. The mathematical, statistical, and computational sciences are intertwined with human society [Cam, Eth, MC23]. This notion is not new. Throughout his life, educator and civil-rights activist Bob Moses championed mathematical literacy as a civil right for all public-school children, with particular advocacy for those who are most vulnerable [MWW⁺23]. In a 2021 article [Mos21] that was published a month after his death, Moses wrote

In the 1960s, voting was our organizing tool to demolish Jim Crow and achieve political impact. Since then, for me, it has been algebra. What's math got to do with it?—you ask. Everything, I say.

Amidst the planetwide transformation we are undergoing, from industrial to information-age economies and culture, math performance has emerged as a critical measure of equal opportunity.

Mathematical research also does not occur in a vacuum, especially when we apply our algorithms and other methods to social systems [CM23]. It is not a traditional part of curricular education or research education in the mathematical sciences to teach our students and other junior community members about data ethics and other ethical considerations [RSMCM24, Mü24], but this needs to change.³ What we choose to teach (and choose not to teach, or choose to mention in only a cursory way) impacts what our students and other mentees do with their education. Other disciplines (e.g., the social, behavioral, and medical sciences) have thought *a lot* more about ethics than mathematics (and allied disciplines), and we should be guided by the best practices that they have developed. Many of these best practices arose in the aftermath of ethically problematic studies, and ethical guidelines have developed as people have tried to learn from past mistakes. As in other arenas, the mathematics community has had major ethical problems (e.g., through various forms of social toxicity), but these problems traditionally have arisen in issues other than the scientific use and abuse of data. This situation has changed; mathematicians are now also facing data-related issues directly in their research. What we learn, teach, and do needs to catch up to this reality.

It is important for mathematical scientists (and others) to think carefully about ethics and engage with it throughout their careers [CM24]. We need to incorporate data ethics into the core education in the mathematical, statistical, and computational sciences. We are increasingly using social, animal, and human data (including potentially personal data) in our research. We need to think carefully about when it is appropriate to use such data and when it is not appropriate, and there needs to be systematic training to help mathematical scientists confront these issues. See [LMP18] for a short book on ethics and data science, [FT16] for a theme journal issue on data ethics, [Vé21] for a wide-ranging handbook of digital ethics, [MCF22] for a discussion of the ethics landscape in mathematics and advocacy of a “Hippocratic Oath” for mathematics (also see [D’A22]), and [Sku21] for an article with a useful discussion and pointers to several relevant resources. See [Ass18] for the Code of Ethics and Professional Conduct of the Association for Computing

³One positive development is that some computer-science conferences now require authors to include ethics statements in their submitted papers.

Machinery (ACM). Buell et al. [BPT22] leveraged the ethical-practice standards of the ACM and the American Statistical Association (ASA), which both represent disciplines that are relevant to mathematics, in a survey of mathematicians about ethical standards. See [TPB24a, TPB24b] for further discussions of the guidelines that they developed for ethical mathematical practice, and see [WT24] for associated teaching materials. To help acknowledge issues like data ethics in the mathematical sciences, it is also relevant to include pertinent information on departmental websites. For an example website, which I helped produce at University of Oxford in the aftermath of a study [TMP12] by my collaborators and me that used Facebook data, see [Oxf].

2.2. Course material on data ethics and society. Many existing courses discuss data ethics and the societal impact of data. Casey Fiesler (Department of Information Science, University of Colorado Boulder) has collected the syllabi of many such courses at [Fie18]. A good book on ethics and data science to use as a starting point is [LMP18], which the students in UCLA’s course on societal impacts of data are asked to read.

Several courses have websites with a lot of useful information about data ethics and related topics. I will highlight a few examples. Matt Salganik (Department of Sociology, Princeton University) has taught a graduate course (from which I drew some material for the present chapter) called *Computational Social Science: Social Research in the Digital Age*. The material for the course’s Fall 2016 edition is available at [Sal16]. It includes material that appeared later in book form [Sal17]. Johan Ugander (Department of Management Science and Engineering, Stanford University) has taught a graduate course called *Data Ethics and Privacy*. See [Uga20] for the course website, from which I drew several articles that I mention in Section 6. Johan Ugander’s graduate course *Social Algorithms* [Uga23] also has excellent resources. Chris Bail (Department of Sociology, Duke University) has taught a graduate course called *Data Science & Society*. See [Bai22] for its course materials, schedule, and YouTube videos. A recent book by Chris Wiggins (Department of Applied Physics and Applied Mathematics, Columbia University) and Matthew L. Jones (Department of History, Columbia University) [WJ23] is an adaptation of their course *Data: Past, Present, and Future* [WW24].

Rachel Thomas (co-founder of FAST.AI and Professor of Practice in the Center for Data Science at Queensland University of Technology) has posted a wealth of resources at [Tho]. This material includes a data-ethics course, a data-science blog, a diversity blog, and more. For example, take a look at her collection of short videos on ethics in machine learning [Tho21]. The course *Calling Bullshit* by Carl Bergstrom (Department of Biology) and Jevin West (Information School) at University of Washington includes a section on ethics [BW19]. This course also includes lots of other valuable material and reading suggestions.

The Cambridge University Ethics in Mathematics Project [Cam] also has relevant course material, including a description and recordings of an 8-lecture course called *Ethics for the Working Mathematician*. Other useful material for courses on data ethics are research papers (see Section 6) and discussions of controversies (see Section 4.4).

3. A few efforts related to data ethics, society, and social good

In this section, I briefly discuss a few efforts that are related to data ethics and society. I highlight two recent efforts at University of California, Los Angeles (UCLA), and I briefly mention a few of the many efforts of others. I also highlight the importance of being careful and conscientious in these efforts.

3.1. Two recent initiatives at UCLA.

3.1.1. *An undergraduate course: Societal impacts of data.* I helped design a new undergraduate major in “Data Theory”⁴ at UCLA. All students in our Data Theory major must take a new upper-division course, which was designed by Mark Handcock of the Department of Statistics & Data Science, called “Societal Impacts of Data”. This course covers a variety of topics—including privacy, algorithmic bias, and many others—through an ethical lens. UCLA’s course catalog has the following description of the course:

Consideration of impacts that data collected today have upon individuals and society. Rapid increase in scale and types of data collected has impacted commerce and society in new ways. Consideration of economic, social and ethical, legal and political impacts of data, especially that collected on human behavior. Topics include privacy and data protection, intellectual property and confidentiality, sample selection and algorithms, equality and anti-discrimination.

In a recent offering of the course, [LMP18] was used as reading material.

In my view, such courses should be mandatory not only for undergraduates who are majoring in data science (and similar topics), but also for all other students. Because data ethics is crucial for any member of human society, taking a course in it should be a core requirement for literally all students to obtain an undergraduate degree. However, it is especially important for the many students (in mathematics, statistics, computer science, and other subjects) who become data scientists or otherwise work with human data (and nonhuman animal data) in their careers.

3.1.2. *Social-justice data-science postdoctoral scholars.* A new UCLA academic position, which I designed along with Deanna Needell (Department of Mathematics) and Mark Handcock, is a Social-Justice Data-Science (SJDS) postdoctoral scholar. This innovative postdoctoral position, which I hope to see in various forms at academic and other institutions worldwide, is a joint venture of UCLA’s mathematics and statistics departments. An SJDS postdoc has two mentors: (1) a faculty member from the Department of Mathematics or the Department of Statistics & Data Science and (2) a faculty member who is a social-justice scholar. We hope to continue to hire SJDS postdocs.

Mathematical, statistical, and computational scientists can accelerate the quantitative study of social-justice issues by harnessing data. They can also leverage recent advances in data science into social justice and activism. Individuals who are trained in other fields (such as sociology) have a long tradition of such involvement in activism. Importantly, mathematicians, statisticians, and computer scientists also have a lot to contribute. Such contributions are not simply a matter

⁴See [UCLa] for a description of UCLA’s undergraduate major in Data Theory.

of studying abstract problems in mathematical sociology or allied topics. Critically, it requires engagement with social-justice scholars and the communities and other stakeholders that we seek to help. This is what we expect SJDS postdocs to do. Moreover, by sponsoring SJDS scholars as postdocs in mathematics and statistics departments in positions of comparable prestige to our usual postdocs, the mathematical-science community (along with our colleagues in statistics, computer science, and other disciplines) can show students that these paths—whether in academia, in industry, as a data scientist for a nonprofit organization that serves communities, or elsewhere—are available and viable career pathways. It is crucial that we send this message.

3.2. Some other noteworthy efforts. There are many other efforts (which take a variety of forms) on data and society. In Section 2.2, I discussed several courses and resources on data ethics. In this section, I briefly highlight a few initiatives on data and society.

One of these efforts is Mechanism Design for Social Good (MD4SG) [AG18, MD4], which uses techniques from algorithms, optimization, mechanism design, and disciplinary insights to improve the equity, social welfare, and access to opportunities for historically underserved, disadvantaged, and marginalized communities. Another effort, which was launched by Timnit Gebru, is the Distributed AI Research Institute (DAIR) [DAI]. This institute works on community-based and interdisciplinary AI research. AI4All, which was founded by Fei-Fei Li and Olga Russakovsky in 2015, is a nonprofit organization that aims to increase diversity and inclusion in AI education, research, development, and policy [AI4].

3.3. An important caveat. There are a diverse variety of both research efforts and practical efforts that fall under the auspices of “data for social good” [ACX23]. Harnessing data for social good is both important and laudable, but there are many challenges [ABK⁺20, BBS22]. For example, it is vital to engage meaningfully and respectfully with communities, rather than attempting to help people by trying to be a “new sheriff in town”, which can be very harmful.

4. Research replication and ethics, open data and code, the tension between privacy and open data and code, and some controversy

4.1. Replication of research. It is crucial to be able to replicate scientific research, and improvement in current practices is necessary to produce reproducible and reliable computational science [CGH21]. Obviously, it is necessary to be honest about data and other aspects of research, but professional obligations go far beyond mere honesty. Scholarship needs to be transparent. It is important to precisely explain all details of analysis, implementation, and data cleaning in scholarly works; it is also important to openly provide source code and data. See [ATG24] for a set of recommendations for sharing code, and see [S⁺25] for a discussion of guidelines for the management of scientific data. The inclusion of precise explanations and source material is necessary for research dissemination. When a mathematician publishes a theorem and its proof, they also give other people a license to use it freely in their own work. The sharing and portability of knowledge lies at the core of both science and mathematics. Accordingly, it is a professional obligation to share research in a usable form, including by providing

source code and data. Depriving readers of such material is analogous to publishing a theorem statement without a proof or publishing a theorem without permission to use it.

To the extent possible, it is very important to publish the relevant data and code (including code to reproduce all figures) that accompanies a scholarly manuscript.⁵ In a manuscript, one also should explicitly and carefully discuss each step in the procedures for data anonymization, cleaning, sampling, and transformation. It is important to be explicit about anything that one does with data so that readers know precisely what choices have been made and can then evaluate whether or not they think that those choices are good ones for the analysis in a manuscript. For example, sampling biases can change the properties of data in fundamental ways [SWM05]. Additionally, by providing the original data when possible, others can analyze that data with procedures that involve different choices. There are many choices that scientists make in data analysis—it is impossible not to make such choices—but these choices are a fundamental part of the scientific procedure when conducting research, so it is imperative to inform others of exactly what one has done (with particular highlighting of choices) in any scientific work. They may want to make different choices.

In making data publicly available, posting the output of synthetic models is safer than posting even the safest real-world data (see Section 4.3). When using synthetic data, such as the output of numerical simulations of a differential equation or data that one generates from a random-network model, it is good to publish code to generate the output data (e.g., the examples) that one presents in a manuscript. Publishing user-friendly and open-source code is important for a paper’s readers, and it also facilitates the fair evaluation of methods and results. One way to publish code is as supplementary material on a journal website; another way is through repositories such as Bitbucket, GitLab, and GitHub. Posting synthetic data is relevant not only for the output of numerical computations and any other data that one generates, but also even for examples (such as adjacency matrices in a paper about networks) that one constructs by hand in a definition–theorem–proof paper. How relevant it is to include such data depends on the sizes of the examples. Even posting the entries of a 10×10 matrix in a repository saves time for others and reduces transcription errors.

We are all human, and it is easy to forget something or to inadvertently be insufficiently precise about a procedure, so gaps often occur. If somebody e-mails you to ask for a clarification, copy of code (even if poorly commented), or something else, it is important to respond and provide it to them (assuming that it is something that you have the legal and ethical right to provide).

4.2. Acknowledgements, giving credit, and licenses to use and share material. Another part of doing scientific research and presenting it in scholarly works is acknowledging the contributions of others. Naturally, acknowledging contributions includes things like coauthorship and citations of prior work. It also includes things like acknowledging all sources of data, all sources of funding, and thanking people for their useful comments and ideas. In the acknowledgements section of a manuscript, one should include precise details of how one obtained data

⁵In the interest of admitting my own flaws, I note that I have been imperfect in my career about publishing source code with my papers. I am doing this increasingly often, and I seek to improve further.

and how others can also obtain that data (especially if one cannot publish it oneself, perhaps because of privacy considerations or because it is not one’s own data to share). It is important to be generous when acknowledging others in manuscripts. If somebody gives useful comments, one should thank them for it (assuming that they want to be thanked).

One should be fair, appropriate, and precise when discussing prior work in a manuscript. It is crucial to give credit where it is due. The research in prior work has heterogeneous levels of mathematical rigor, scientific rigor, and even correctness. How one writes about such work is affected by such things, including some facets that are factual and others that may reflect a variety of opinions. For example, there is a difference in writing that something was “shown” versus “reported” in a prior work. The former wording has a built-in claim, by the author(s) of a manuscript, of the validity of that aspect of that prior work. By contrast, the latter is merely a historical fact (assuming that what one writes is itself accurate).

The increasing prominence of LLMs has brought new ethical concerns to the attribution of proper credit and to the creative process more generally.⁶ Scientific journals, funding agencies, professors, and universities are scrambling to develop policies that govern the use of LLMs and other text-generation tools in submissions [Bra23]. See [Con] for the LLM and AI policy that the Consortium for Mathematics and its Applications (COMAP) has started using for the Mathematical Contest in Modeling (MCM) and their other contests. For UCLA’s guidelines for the ethical and “safe” use of AI, see UCLA’s new AI website [UCLb], which also discusses various AI resources.⁷

Another aspect of open science is the different types of licensing that are available through Creative Commons. See [Wika] for a discussion of the different types of Creative Commons licenses. Some licenses allow work to be duplicated for any purpose, and other licenses enforce a variety of use restrictions. I advocate making one’s work as open as possible, as others can then use it readily for purposes such as teaching and explaining ideas.

4.3. Privacy concerns and practical considerations for open data and code. There are various tensions and practical considerations with the lofty ideals of openly publishing data and code. For example, one may not be allowed to publish data for privacy reasons or because of nondisclosure agreements. For empirical data, if you have permission to post something (e.g., does the data “belong” to somebody else?) and it does not pose privacy concerns, then it makes sense to post it because doing so promotes good science. Because of privacy issues, one may choose not to publish certain data that one is technically allowed to publish. It is crucial that researchers navigate these issues in a conscientious way. Another issue is that publicly posting usable code and data takes time and energy, and key participants (such as students and other junior researchers) in a project move on to other things,

⁶The use generative AI and LLMs also brings significant ethical concerns about data ownership and copyright infringement, such as through the training data that is used for AI-generated art and text [Fra22, GM23]. It also can lead to comically disturbing situations, as illustrated by the recent discovery of child sexual-abuse material in a prominent data set (which has been downloaded by many researchers) in the AI community [Col23] and by the AI-generated image of a rat with a gigantic penis in a (subsequently retracted) published scientific paper [Pea24].

⁷Amidst the very serious concerns about AI, it is also important not to lose sight of its immense promise. For example, see the curated list [RT24] (and the announcement of it at [Tao24]) of resources for AI in mathematics (e.g., to use AI to assist in mathematical reasoning).

so the team members who are best equipped to do this effectively may no longer be available. In other words, there is sometimes a practical tension between the well-being of one's collaborators and the publishing of code and data.

In Section 4.1, I mentioned that it is important to indicate how one has anonymized data in scholarly works. When is data actually “anonymous”, and is it ever possible to “fully” anonymize data? Consider the following scenario, which is discussed in [Sal16, Sal17]. Suppose that we have a data set of medical records of individuals that includes their names, their home addresses, the zip codes of these addresses, their birth dates, their sexes, their ethnicities, the dates that each individual visited a doctor, the medical diagnoses, the medical procedures, and the prescribed medications. We can try to “anonymize” this data set by removing the names and home addresses of all individuals. We now have a data set of “anonymized” medical records. Suppose that we obtain a data set of voting records and that this data set includes names, home addresses, political-party affiliations, voter-registration dates, zip codes, birth dates, and sexes. Both data sets include the zip codes, birth dates, and sexes of the individuals that are common to the two data sets. One can use such data—namely, data that is common to these two data sets—to deanonymize people in the supposedly “anonymized” data set of medical records [NS08]. An infamous example of data deanonymization by combining data sets led to the cancellation of the sequel to the Netflix Prize [Wikb].

Given the simultaneous presence of privacy concerns and the desire to produce replicable scientific research, what should one do if the employed data, an algorithm (or part of an algorithm), or something else needs to remain private? This was one key topic of discussion following a publication by Bakshy et al. [BMA15], who examined the exposure of different individuals to heterogeneous news and opinions in their Facebook feeds. The paper's authors, who were all Facebook employees at the time, could not reveal how Facebook determined the feeds that individuals saw, so how can others replicate their work to try to evaluate and verify their observations and insights? Which of the insights in [BMA15] apply exclusively to Facebook, and which of them also apply to other social-media platforms? In principle, it should be possible to attempt a weaker form of replication of the study's most interesting qualitative results, which are not merely a property of something that is specific to Facebook.

4.4. Controversies in studies and in reactions to studies. Unsurprisingly, some studies that involve human data have been controversial. In other cases, there has been controversy in how authors of studies were treated in the aftermath of their work.

One controversial study, with much ensuing public discussion (see, e.g., [Luc14] and many other sources), was an examination of emotional contagions using experiments with Facebook in which user feeds were altered [KGH14]. There were angry accusations that the researchers manipulated people's emotions, with additional questioning of subsequent actions by the journal that published the paper. There were also discussions of the procedure to obtain permission to undertake the study in the first place. One key consideration is that there are crucial differences between the ethical procedures for academic and commercial researchers [Gri16].

Academic researchers need to obtain approval from an Institutional Review Board (see Section 5.3) before undertaking a study that involves humans, whereas companies like Facebook have publication review boards to approve the publication of a study after it has already been done. Therefore, we know that this study occurred because Facebook concluded that it could be published. By contrast, we do not know about what research is done with our data by Facebook and other entities when an associated document is not placed in the public domain. This leads to an important question: Should academic researchers and companies follow the same rules?

Many technology companies have units that do research on data ethics and related subjects, although that too can lead to controversy. One example is the departure from Google of researcher Timnit Gebru and others who study data ethics [Wikd] in the aftermath of a paper that Gebru and her collaborators wrote about the significant risks (including environmental costs, unknown and dangerous biases, and potential uses to deceive people) of LLMs [Hao20].

Other research about online social networks has also been controversial. I will mention two examples that have been influential scientifically because of their research findings. One of these examples involved experimental manipulation of feeds by seeding posts on social media with a small number of fake initial upvotes or fake initial downvotes [AW12]. In this study, the researchers found that initial upvotes had a persistent effect on the overall positivity of the voting on posts, whereas the initial downvotes were overturned. The other example is the “Tastes, Ties, and Time” study of several waves of Facebook data from students at an Ivy League university in the United States [LKG⁺08]. See [Per11] for one discussion of the data-privacy controversy of this study and its associated data set.

4.5. Summary. To summarize some key ideas in Section 4, here are a few things to think about:

- There is tension between open data and personal privacy.
- The use and publication of data, code, and anything else that one reports in a manuscript or discusses with others can be subject to terms-of-service agreements and nondisclosure agreements.
- In what sense can one make one’s research replicable if one cannot make all of the associated data (or algorithms or something else) publicly available? There are weaker notions of replication, such as whether or not others can observe similar phenomena in circumstances that are similar but not precisely the same. For instance, in studying human behavior on social media, perhaps certain phenomena are very similar on Facebook and X (the social-media platform formerly known as Twitter), but other phenomena are specific to only one of these two social-media platforms.
- Are you comfortable doing research in collaboration with private companies or government entities? Maybe there are some entities with which you are willing to collaborate (perhaps depending on their purposes, goals, and history), but there are others with which you are not willing to collaborate or use data from? If you work with or for such an entity, what is permissible to include in a publication or post online?

5. Ethical principles and other considerations in the scientific use of human data

In this section, I discuss some ethical principles and other considerations in the scientific use of human data.⁸ I also discuss Institutional Review Boards (IRBs).⁹ Much of my discussion also applies to data from nonhuman animals, but I am focusing on human data in this chapter, so I typically phrase my exposition in human terms.

5.1. Online courses for ethics training. In studies that use human data, it is important to think carefully about ethics and to have formal training in it. A popular choice is the Collaborative Institutional Training Initiative (CITI) program [CIT], which offers a variety of courses. For more information, see the website of the UCLA Office of the Human Research Protection Program (OHRPP) [OHR]. See Section 5.3 for further discussion of these courses and of IRBs.

5.2. Four key principles. I now enumerate a few critical ethical principles, which are discussed in detail in [Sal17, Chapter 6].

In scientific pursuits that involve human data, it is important to do the following:

- be honest and fair (obviously);
- design ethically thoughtful research;
- explain your decisions to others.

Four key principles in research that involves humans are

- respect for persons;
- beneficence;
- justice;
- respect for law and public interest.

These four principles can come into tension with each other, so how do we balance them?

In conducting research with human data, there is a sliding scale: the more your research has the potential to violate personal privacy, the more helpful for humanity its outcome needs to have the potential to be. Four things to ponder with research that involves personal data are the following:

- informed consent;
- understanding and managing informational risk;
- privacy;
- making decisions in the face of uncertainty.

As you design and conduct research, put yourself in the shoes of other people. Think of research ethics as continuous (i.e., there is a sliding scale), rather than as discrete.

⁸Parts of my discussion draw heavily from material in Matt Salganik's course on computational social science [Sal16]. See his associated book [Sal17] for further discussion.

⁹At some institutions, the backronym for IRB is Internal Review Board.

5.3. Institutional Review Boards (IRBs).

“Yeah, yeah, but your scientists were so preoccupied over whether or not they could, they didn’t stop to think if they *should*.”

(stated by the character Ian Malcolm in *Jurassic Park*)

In many situations, it is a professional requirement to obtain permission to undertake a study in the first place. Such permission gives an ethical floor to satisfy; it is not a ceiling. It may be legally and professionally permissible to do something, but it is important to hold oneself to higher standards when it comes to whether or not it is actually the right decision to do it. For example, the limits to informed consent with human data [LAH⁺22] may influence such a decision. Whatever you decide for your own work, make sure that you think carefully about it.

In universities in the United States, a researcher that is working with personal data needs to check with their university’s IRB to ensure that they are conducting research in an ethical way. Universities in other countries, private companies, government laboratories, and other organizations often have bodies that are similar to IRBs, but the procedures and especially the specific details are very different [Gri16]. A university IRB may inform you that you do not need to submit a formal application for a research project to be approved, or they may inform you that a formal application is necessary. Let your IRB know what data you have (or what data you plan to acquire and how you plan to acquire it) and what you plan to do with it. Different IRBs can rule differently. When an IRB grants permission to undertake a study, they have decided that a proposed project is above the ethical floor and hence that it is permissible to do that project. One’s own standards should be higher. (Again, ethical approval is a floor, rather than a ceiling.) In this light, it is worth examining the discussion following a controversial IRB-approved study of “emotional manipulation” in which researchers adjusted user feeds on Facebook [Luc14, KGH14].

For a more concrete idea about IRBs and conducting ethical research in a university setting, see the materials at [OHR]. The requirements to conduct research with human data (and nonhuman animal data) include taking various online training courses, such as those in the CITI program. These courses, which are available online at [CIT], are in common use in the United States. The training that is required, expected, and available for research projects that involve human data and other sensitive data is rather different in different countries. Some of the topics that are covered in courses on ethical research are animal care and use, biosafety and biosecurity, human-subject research, information privacy and security, and responsible conduct of research. It is useful to take a variety of these courses even when they are not required.

5.4. Another salient warning. It is important to be cognizant that your research and data can potentially be “weaponized” by other people—who may not care about nuances in research findings and who may interpret unfortunate wording or insufficiently careful exposition in nefarious ways—so you need to be conscientious about the precise wording in your publications and other media. This is particularly relevant in research on social systems and on tools that are applicable to social systems [Ste22]. This possibility also makes it particularly crucial to be

open-minded about potential sources of bias in your research. As always, it is important to be ethically thoughtful.

6. Some research and lay articles that are relevant to data ethics and data privacy

There is a wealth of research about data ethics and related topics in computer science, sociology, and other fields [JIV19, LHF⁺21, WSAO⁺21, ABK⁺20]. There is no way that I can possibly be exhaustive,¹⁰ so I will highlight a selection of studies to give a taste of existing research. I will also briefly discuss a few articles in blogs, news websites, and other nontechnical venues. Unsurprisingly, much of this research is simultaneously fascinating and concerning [Exp18].

6.1. Algorithmic biases and other biases. A key research area is the analysis and mitigation of biases in computational techniques and data analysis.

One central topic is algorithmic bias, in which systematic and repeatable errors result in unfair outcomes, such as privileging one group of people over others [Bir21]. For example, such biases have systematically hurt certain racial and demographic groups in predictive policing [Nat18]. Algorithmic biases also reinforce negative racial and gender biases in online search-engine results [Nob18]. As is now evident, algorithmic biases can have very serious consequences. For example, in 2020, an algorithm for prioritizing COVID-19 vaccinations severely disadvantaged human-facing medical workers [Che20].

There is much research on the mitigation of algorithmic biases, such as in the algorithmic hiring of people for jobs [RBKL20] and the algorithmic classification of individuals (e.g., for admission to a university) [DHP⁺12]. To encourage transparent reporting, clarify intended use cases, and minimize usage in inappropriate contexts, some researchers have proposed the inclusion of “model cards” to accompany trained machine-learning models [MWZ⁺19]. Such model cards are short documents that give intended use cases; benchmarked evaluations across cultural, demographic, phenotypic, and intersectional groups; and other salient information.

Biases in computer systems, machine learning, and AI go far beyond only algorithmic biases [FN96, HV24]. See [MMS⁺21, BHN23, CDNSG23] for reviews of bias and “fairness” in machine learning, and see [N⁺20] for an introductory survey of bias in data-driven AI systems. As advocated in [KA21], it is important to go beyond notions of mere algorithmic “fairness” (which focuses on intra-group versus inter-group differences). One must also analyze (1) inequality and the causal impact of algorithms and (2) the distribution of power. It is also important to be cognizant of traps that can beset work on “fair” machine learning in sociotechnical systems [SBF⁺19]. In “fair” machine learning, it is necessary to carefully examine context-specific consequences [CDNSG23].

To mitigate biases, it is imperative for researchers to carefully justify their choices of data sets. It is crucial to consider the social context (country, gender, race, and so on) of data. Why is one using a particular data set, and is it appropriately representative for the problem that one is studying [KDHF21]? It is common and often convenient to import data sets that were used originally to study one problem for investigations of many other problems, and that can lead to several

¹⁰I selected some of the highlighted articles, including both research studies and lay articles, from the website for Johan Ugander’s course on data ethics and data privacy [Uga20].

issues. This is particularly relevant for human data and social data, and it is important to think carefully about whether a data set is appropriate for a given study. Particular facets of a setting can interfere with attempts to generalize a study’s insights from its specific use case to other potentially similar situations.

6.2. Human privacy, personal characteristics, and personalization.

Other salient research focuses on human privacy. As has been studied thoroughly [EN16], people are tracked extensively when they visit websites. User profiles, which encode both characteristics and behavior, are generated through interactions with websites and other digital systems. User profiling is ubiquitous in daily life, and it is also a vast area of research [PBL24]. It is possible to infer private traits and attributes from digital records of human behavior [KSG13]. One can also steal the identities of visitors to websites [Nar10], and trackers can use social-network structure to deanonymize data from browsing the World Wide Web [SSGN17]. Companies can exploit the information that they observe (or infer) from website visits. For example, about a decade ago, the online travel agent Orbitz showed higher prices for flights to users of Macintosh computers than to users of other types of computers [Whi12]. It is imperative that individuals and other stakeholders have agency in how their personal data are used [SVW21], although this can come into tension with the scholarly desire to promote open data (see Section 4). For example, the sharing of data from the continent of Africa has often been driven by non-African stakeholders [AAB⁺21]. Additionally, conventional studies of algorithmic fairness are centered on Western culture; data proxies are different in different cultures, and it is important to consider local context when building data models [SAH⁺21]. See [LHF⁺21, WSAO⁺21] for discussions of access, ethics, and best practices in the algorithmic measurement of human data.

There are many ways to infer the characteristics of people and communities, as well as social and other ties between people, using data analysis. For example, it is possible to infer social ties by examining the geographical proximity of individuals in time and space using offline or online data [CBC⁺10]. Additionally, researchers have used machine learning and Google Street View to estimate the demographic composition of neighborhoods across the United States [GKW⁺17]. Unsurprisingly, there are significant gaps and biases (and associated exclusion issues) in such geographic data [GD22]. Moreover, as I discussed in Section 4, the fact that individuals appear in multiple social networks (e.g., multiple social-media platforms and multiple databases) gives considerable ability to identify them or information about them even when they appear to be “anonymous” in those individual networks [NS08]. For example, the inference of the sexual orientation of some individuals by combining Netflix data—which was released as part of a public competition to improve Netflix’s algorithmic ability to infer user ratings of movies—of movie rentals with movie rating data from the Internet Movie Database (IMDB) led to legal troubles, and Netflix ultimately cancelled a planned sequel competition [Wikb]. More recently, researchers have successfully inferred gender using mobile-payment data [SM23]. Other researchers have illustrated that one can use interactions to identify people even across long time periods [CMM⁺22]. Tools like data analysis and machine learning can yield crucial and actionable insights (e.g., of racial disparities in police stops [PSO⁺20]), but they need to be applied in a careful and respectful way.

A related issue is personalization [Uga17], such as in the targeted advertisements that inundate people on social-media platforms. These advertisements can inform us of desirable products (such as t-shirts, dice, or plushies in my case), but the online personalization of commercial and political advertising can cause serious problems. Infamously, Donald Trump’s 2016 presidential campaign employed the company Cambridge Analytica for targeted online political advertising. Additionally, the movie *Straight Outta Compton* was advertised differently to people based on their inferred demographic characteristics [McA16]. There is much pertinent research on personalization. For example, there are studies of the effectiveness of spam-based marketing [KKL⁺09], the detection of spam in social-bookmarking websites (such as Pinterest and Digg) [MCM09], and other aspects of spam and related online marketing. Users on social-media platforms like Facebook have personalized feeds, and a study that involved the adjustment of such feeds [KGH14] led to a major controversy, including with the issue of manipulating people’s emotions [Luc14].

6.3. AI ethics guidelines, tradeoffs in different ethical values, and theoretical barriers. Guidelines for AI ethics appear to have converged on five key values:¹¹ transparency, fairness, safety, accountability, and privacy [JIV19]. There are inevitable tensions between these ethical values. Academics, public-sector organizations, and private companies emphasize different ethical values, with particularly systematic differences between practitioners and the general public [JBAO22].

It is mathematically impossible for an AI system to simultaneously include many parameters, be robust to poisoning (e.g., with fake data) by an adversarial actor, and preserve privacy [EMFG⁺22], so the use of AI will always include tradeoffs between different ethical values even when everything works perfectly. Of course, AI systems don’t work perfectly. For example, it was proven recently that one can plant undetectable backdoors into machine-learning models and thereby confuse “adversarially robust” classifiers, demonstrating a major theoretical barrier to the certification of adversarial robustness [GKVZ22].

7. Conclusion

The use of digital data to study and analyze humans, create technologies, and enact policies that involve humans is both powerful and dangerous [Edi21, Off22]. Mathematical scientists have a younger tradition of studying human data than researchers in many other disciplines, such as the social, behavioral, and medical sciences. Most researchers in the mathematical, statistical, and computational sciences have not had research-ethics and data-ethics training. This needs to change. One recent paper has even proposed a “Hippocratic Oath” for the mathematical sciences [MCF22]. The mathematical-science community needs to learn from the best practices of other disciplines to ensure that our research is ethical. Other scholarly communities have been considering research ethics a lot longer than we

¹¹These values overlap with (but are organized differently from) the five principles to guide the design, use, and deployment of automated systems in the recent blueprint for an AI Bill of Rights from the United States government’s Office of Science and Technology Policy (OSTP) [Off22]. The OSTP’s five principles are (1) safe and effective systems, (2) algorithmic-discrimination protections, (3) data privacy, (4) notice and explanation, and (5) human alternatives, consideration, and fallback.

have, and it is important that we learn from them. As in the social and medical sciences, the mathematical and computational sciences need a robust program of ethics training.

In this chapter, I have discussed many aspects of data ethics in education, publishing, and research. In some of my discussions, I have also touched on issues of data and justice. It is important to distinguish between ethics and justice, including when it comes to data [Kit14]. Ethics presupposes laws and social norms, and then one considers what is morally appropriate within those frameworks. Justice is concerned with how to change laws and modify social norms to attain broader equity. Both ethics and justice are core aspects of our interactions with data.

I encourage you to read widely, think about, and discuss how to do research ethically, especially for studies of—or with consequences for—social systems and human data. It is also valuable to read about past research and societal controversies. There have been mistakes in the past (and there continue to be mistakes), and we need to learn from them. We may all set our ethical bars in different places and have different views on different issues, but our scholarship needs to be conscientious and ethically thoughtful. As a reminder, official approval (e.g., from an IRB) to undertake a study is only a lower bound. The ethical bar that one needs to surpass in the design and performance of research is a sliding bar: the more potential for invasion of human privacy (or other potential harm), the more valuable to humanity the potential outcome of a research project has to be.

Be ethically thoughtful.

Acknowledgements

In parts of this chapter, I drew on materials from a class [Sal16] that Matt Salganik has taught on computational social science at Princeton University. Johan Ugander suggested several excellent resources and research studies. I thank Jacob Foster, Mark Handcock, Andrei Kleshin, Peter Mucha, Dennis Müller, Deanna Needell, Matt Salganik, Massimo Stella, Alexandria Volkening, Chris Wiggins, Richard Yeh, and an anonymous referee for helpful comments. Finally, I thank my fellow Short Course organizers (Heather Zinn Brooks, Michelle Feng, and Alexandria Volkening) for their continuing support on this project and elsewhere, their collaboration, and their friendship.

References

- [AAB⁺21] Rediet Abebe, Kehinde Aruleba, Abeba Birhane, Sara Kingsley, George Obaido, Sekou L Remy, and Swathi Sadagopan, *Narratives and counternarratives on data sharing in Africa*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), Association for Computing Machinery, New York, NY, USA, 2021, pp. 329–341.
- [ABK⁺20] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G. Robinson, *Roles for computing in social change*, Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20), Association for Computing Machinery, New York, NY, USA, 2020, pp. 252–260.
- [ACX23] Ahmed Abbasi, Roger H. L. Chiang, and Jennifer Xu, *Data science for social good*, J. Assoc. Inf. Syst. **24** (2023), no. 6, 1439–1458.
- [AG18] Rediet Abebe and Kira Goldner, *Mechanism design for social good*, AI Matters **4** (2018), no. 3, 27–34.
- [AI4] AI4All, accessed December 26, 2023. <https://ai-4-all.org>
- [Ass18] Association for Computing Machinery, *ACM code of ethics and professional conduct, 2018*, accessed December 26, 2023. <https://ethics.acm.org>

- [ATG24] Richard J. Abdill, Emma Talarico, and Laura Grieneisen, *A how-to guide for code sharing in biology*, PLoS Biology **22** (2024), no. 9, e3002815.
- [AW12] Sinan Aral and Dylan Walker, *Identifying influential and susceptible members of social networks*, Science **337** (2012), no. 6092, 337–341. MR2986000
- [Bai22] Christopher A. Bail, *Data science & society*, 2022, Sociology 367S, Spring 2022, Duke University. <https://dssoc.github.io>
- [BBS22] Caroline Buckee, Satchit Balsari, and Andrew Schroeder, *Making data for good better*, PLoS Digital Health **1** (2022), no. 1, e0000010.
- [BFPV21] Heather Z. Brooks, Michelle Feng, Mason A. Porter, and Alexandria Volkening, *Short course: Mathematical and computational methods for complex social systems*, 2021. https://zerodivzero.com/short_course/aaac8c66007a4d23a7aa14857a3b778c/titles
- [BGMMS21] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, *On the dangers of stochastic parrots: Can language models be too big?* 🐦, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), Association for Computing Machinery, New York, NY, USA, 2021, pp. 610–623.
- [BHN23] Solon Barocas, Moritz Hardt, and Arvind Narayanan, *Fairness and machine learning: Limitations and opportunities*, MIT Press, Cambridge, MA, USA, 2023. <http://www.fairmlbook.org>
- [Bir21] Abeba Birhane, *Algorithmic injustice: A relational ethics approach*, Patterns **2** (2021), no. 2, 100205.
- [BMA15] Eytan Bakshy, Solomon Messing, and Lada A. Adamic, *Exposure to ideologically diverse news and opinion on Facebook*, Science **348** (2015), no. 6239, 1130–1132. MR3380630
- [BPT22] Catherine A. Buell, Victor I. Piercey, and Rochelle E. Tractenberg, *Leveraging guidelines for ethical practice of statistics and computing to develop standards for ethical mathematical practice: A white paper*, arXiv:2209.09311, 2022.
- [Bra23] Jeffrey Brainard, *Journals take up arms against AI-written text*, Science **379** (2023), no. 6634, 740–741.
- [BW19] Carl T. Bergstrom and Jevin West, *Calling bullshit: Data reasoning in a digital world. Week 10. The ethics of calling bullshit*, 2019, Informatics 270 / Biology 270, University of Washington. <https://www.callingbullshit.org/syllabus.html#Ethics>
- [Cam] The Cambridge University Ethics in Mathematics Project, accessed December 26, 2023. <https://www.ethics.maths.cam.ac.uk>
- [CBC+10] David J. Crandall, Lars Backstrom, Dan Cosley, Siddharth Suri, Daniel Huttenlocher, and Jon Kleinberg, *Inferring social ties from geographic coincidences*, Proceedings of the National Academy of Sciences of the United States of America **107** (2010), no. 52, 22436–22441.
- [CDNSG23] Sam Corbett-Davies, Johann D. Gaebler, Hamed Nilforoshan, Ravi Shroff, and Sharad Goel, *The measure and mismeasure of fairness*, Journal of Machine Learning Research **24** (2023), 312. MR4664749
- [CGH21] Peter V. Coveney, Derek Groen, and Afons G. Hoekstra, *Reliability and reproducibility in computational science: Implementing validation, verification and uncertainty quantification in silico*, Philosophical Transactions of the Royal Society A **379** (2021), 20200409.
- [Che20] Caroline Chen, *Only seven of Stanford's first 5,000 vaccines were designated for medical residents*, ProPublica, December 18, 2020. <https://www.propublica.org/article/only-seven-of-stanfords-first-5-000-vaccines-were-designated-for-medical-residents>
- [CIT] CITI Program, *Explore our courses*, accessed December 26, 2023. <https://about.citiprogram.org>
- [CM23] Maurice Chiodo and Dennis Müller, *Manifesto for the responsible development of mathematical works—A tool for practitioners and for management*, arXiv:2306.09131, 2023.
- [CM24] Maurice Chiodo and Dennis Müller, *A field guide to ethics in mathematics*, Notices of the American Mathematical Society **71** (2024), no. 7, 939–947. MR4776111

- [CMM⁺22] Ana-Maria Crețu, Federico Monti, Stefano Marrone, Xiaowen Dong, Michael Bronstein, and Yves-Alexandre de Montjoye, *Interaction data are identifiable even across long periods of time*, *Nature Communications* **3** (2022), 313.
- [Col23] Samantha Cole, *Largest dataset powering AI images removed after discovery of child sexual abuse material*, 404 Media, December 20, 2023. <https://www.404media.co/laion-datasets-removed-stanford-csam-child-abuse/>
- [Con] Consortium for Mathematics and its Applications, *Use of large language models and generative AI tools in COMAP contests*, accessed April 21, 2024. https://www.contest.comap.com/undergraduate/contests/mcm/flyer/Contest_AI_Policy.pdf
- [D'A22] Susan D'Agostino, *Do scientists need an AI Hippocratic oath? Maybe. Maybe not.*, *Bulletin of the Atomic Scientists*, June 9, 2022. <https://thebulletin.org/2022/06/do-scientists-need-an-ai-hippocratic-oath-maybe-maybe-not/>
- [DAI] *Distributed AI Research Institute (DAIR)*, accessed December 26, 2023. <https://www.dair-institute.org>
- [DHP⁺12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel, *Fairness through awareness*, *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS '12)*, Association for Computing Machinery, New York, NY, USA, 2012, pp. 214–226. MR3388391
- [Dor20] Lori Dorn, *A deepfaked Gollum performs a 'precious' cover of Sinead O'Connor's version of 'Nothing Compares 2 U'*, *Laughing Squid*, December 18, 2020. <https://laughingsquid.com/gollum-sinead-oconnor-nothing-compares-2-u/>. The music video itself is available at <https://www.youtube.com/watch?v=9d9pZi7uZQs>.
- [Edi21] Editors, *The powers and perils of using digital data to understand human behaviour*, *Nature* **595** (2021), 149–150.
- [EMFG⁺22] El-Mahdi El-Mhamdi, Sadegh Farhadkhani, Rachid Guerraoui, Nirupam Gupta, Lê-Nguyễn Hoàng, Rafael Pinot, Sébastien Rouault, and John Stephan, *On the impossible safety of large AI models*, *arXiv:2209.15259*, 2022.
- [EN16] Steven Englehardt and Arvind Narayanan, *Online tracking: A 1-million-site measurement and analysis*, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, Association for Computing Machinery, New York, NY, USA, 2016, pp. 1388–1401.
- [ET19] Zeynep Engin and Philip Treleaven, *Algorithmic government: Automating public services and supporting civil servants in using data science technologies*, *The Computer Journal* **62** (2019), no. 3, 448–460.
- [Eth] Ethical Mathematics, accessed December 26, 2023. <https://ethicalmath.com>
- [Exp18] Explain xkcd, *2072: Evaluating Tech Things*, 2018. https://www.explainxkcd.com/wiki/index.php/2072:_Evaluating_Tech_Things. The original comic (by Randall Munroe) is available at <https://xkcd.com/2072/>.
- [Fer24] Emilio Ferrara, *GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models*, *Journal of Computational Social Science* **7** (2024), 549–569.
- [Fie18] Casey Fiesler, *Tech ethics curricula: A collection of syllabi*, *Medium*, July 5, 2018. <https://cfiesler.medium.com/tech-ethics-curricula-a-collection-of-syllabi-3eedfb76be18>
- [FN96] Batya Friedman and Helen Nissenbaum, *Bias in computer systems*, *ACM Transactions on Information Systems* **14** (1996), no. 3, 330–347.
- [Fra22] Matthew R. Francis, *The ethics of artificial intelligence-generated art*, *SIAM News* **56** (2022), no. 09. <https://sinews.siam.org/Details-Page/the-ethics-of-artificial-intelligence-generated-art>
- [FT16] L. Floridi and M. Taddeo, *What is data ethics?*, *Philosophical Transactions of the Royal Society A* **374** (2016), no. 2083, 20160360.
- [GD22] Mark Graham and Martin Dittus, *Geographies of digital exclusion: Data and inequality*, Pluto Press, London, UK, 2022.
- [GKVZ22] Shafi Goldwasser, Michael P. Kim, Vinod Vaikuntanathan, and Or Zamir, *Planting undetectable backdoors in machine learning models*, 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS 2022), IEEE Computer Society, Los Alamitos, CA, USA, 2022, pp. 931–942. MR4537269

- [GKW⁺17] Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei, *Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States*, Proceedings of the National Academy of Sciences of the United States of America **114** (2017), no. 50, 13108–13113.
- [GM23] Michael M. Grynbaum and Ryan Mac, *The Times sues OpenAI and Microsoft over A.I. use of copyrighted work*, The New York Times, December 27, 2023. <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>
- [Gri16] Peter Grindrod, *Beyond privacy and exposure: Ethical issues within citizen-facing analytics*, Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **374** (2016), no. 2083, 20160132.
- [Hao20] Karen Hao, *We read the paper that forced Timnit Gebru out of Google. Here's what it says.*, MIT Technology Review, December 4, 2020. <https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/>
- [Has21] Gry Hasselbalch, *Data ethics of power: A human approach in the big data and AI era*, Edward Elgar Publishing, Northampton, MA, USA, 2021.
- [HV24] David W. Hogg and Soledad Villar, *Position: Is machine learning good or bad for the natural sciences?*, Proceedings of the 41st International Conference on Machine Learning (ICML '24), Proceedings of Machine Learning Research **235** (2024), 18439–18453.
- [IPA20] Institute for Pure and Applied Mathematics, *White paper: Deep fakery*, January 7, 2020. <http://www.ipam.ucla.edu/news/white-paper-deep-fakery/>
- [JBAO22] Maurice Jakesch, Zana Bućinca, Saleema Amershi, and Alexandra Olteanu, *How different groups prioritize ethical values for responsible AI*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), Association for Computing Machinery, New York, NY, USA, 2022, pp. 310–323.
- [JIV19] Anna Jobin, Marcello Ienca, and Effy Vayena, *The global landscape of AI ethics guidelines*, Nature Machine Intelligence **1** (2019), 389–399.
- [KA21] Maximilian Kasy and Rediet Abebe, *Fairness, equality, and power in algorithmic decision-making*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), Association for Computing Machinery, New York, NY, USA, 2021, pp. 576–586.
- [KDHf21] Bernard Koch, Emily Denton, Alex Hanna, and Jacob G. Foster, *Reduced, reused and recycled: The life of a dataset in machine learning research*, NeurIPS 2021: Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, 2021.
- [KGH14] Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, *Experimental evidence of massive-scale emotional contagion through social networks*, Proceedings of the National Academy of Sciences of the United States of America **111** (2014), no. 24, 8788–8790.
- [Kit14] Rob Kitchin, *The data revolution: Big data, open data, data infrastructures & their consequences*, Sage Publishing, Thousand Oaks, CA, USA, 2014.
- [KKL⁺09] Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage, *Spamalytics: An empirical analysis of spam marketing conversion*, Communications of the ACM **52** (2009), no. 9, 99–107.
- [KSG13] Michal Kosinski, David Stillwell, and Thore Graepel, *Private traits and attributes are predictable from digital records of human behavior*, Proceedings of the National Academy of Sciences of the United States of America **110** (2013), no. 15, 5802–5805.
- [LAH⁺22] Juniper L. Lovato, Antoine Allard, Randall Harp, Jeremiah Onaolapo, and Laurent Hébert-Dufresne, *Limits of individual consent and models of distributed consent in online social networks*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), Association for Computing Machinery, New York, NY, USA, 2022, pp. 2251–2262.

- [Les23] Kif Leswing, *Microsoft limits Bing A.I. chats after the chatbot had some unsettling conversations*, CNBC, February 17, 2023. <https://www.cnn.com/2023/02/17/microsoft-limits-bing-ai-chats-after-the-chatbot-had-some-unsettling-conversations.html>
- [LHF+21] David Lazer, Eszter Hargittai, Deen Freelon, Sandra González-Bailón, Kevin Munger, Katherine Ognyanova, and Jason Radford, *Meaningful measures of human society in the twenty-first century*, *Nature* **595** (2021), 189–196.
- [LKG⁺08] Kevin Lewis, Jason Kaufman, Marco Gonzalez, Andreas Wimmer, and Nicholas Christakis, *Tastes, ties, and time: A new social network dataset using Facebook.com*, *Social Networks* **30** (2008), no. 4, 330–342.
- [LMP18] Mike Loukides, Hilary Mason, and DJ Patil, *Ethics and data science*, O’Reilly Media, Inc., Sebastopol, CA, USA, 2018.
- [Luc14] Michael Luca, *Were OkCupid’s and Facebook’s experiments unethical?*, *Harvard Business Review*, July 29, 2014. <https://hbr.org/2014/07/were-okcupids-and-facebooks-experiments-unethical>
- [Mü24] Dennis Müller, *Situating “ethics in mathematics” as a philosophy of mathematics ethics education*, *Ethics and Mathematics Education*, Springer, Cham, Switzerland, 2024, pp. 71–87. MR4830879
- [MC23] Dennis Müller and Maurice Chiodo, *Mathematical artifacts have politics: The journey from examples to embedded ethics*, [arXiv:2308.04871](https://arxiv.org/abs/2308.04871), 2023.
- [McA16] Nathan McAlone, *Why ‘Straight Outta Compton’ had different Facebook trailers for people of different races*, *Business Insider*, March 16, 2016. <https://www.businessinsider.com/why-straight-outta-compton-had-different-trailers-for-people-of-different-races>
- [MCF22] Dennis Müller, Maurice Chiodo, and James Franklin, *A Hippocratic Oath for mathematicians? Mapping the landscape of ethics in mathematics*, *Science and Engineering Ethics* **28** (2022), 41.
- [MCM09] Benjamin Markines, Ciro Cattuto, and Filippo Menczer, *Social spam detection*, *Proceedings of the 5th International Workshop on Adversarial Information Retrieval on the Web (AIRWeb ’09)*, Association for Computing Machinery, New York, NY, USA, 2009, pp. 41–48.
- [MD4] Mechanism Design for Social Good, accessed December 31, 2023. <https://www.md4sg.com>
- [MMS⁺21] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan, *A survey on bias and fairness in machine learning*, *ACM Computing Surveys* **54** (2021), no. 6, 115.
- [Mos21] Bob Moses, *Returning to ‘normal’ in education is not good enough*, *The Imprint*, August 24, 2021. <https://imprintnews.org/opinion/returning-to-normal-in-education-is-not-good-enough/58069>
- [MWW⁺23] Benjamin Moynihan, Robin T. Wilson, Joan Wynne, Lee J. McEwan, Mary M. West, Frank E. Davis, Herbert Clemens, Greg Budzban, Edith Aurora Graf, and Aidan Soguero, *What’s math got to do with it?: Bob Moses, algebra, and the movement for civil rights, January 23, 1935–July 25, 2021*, *Notices of the American Mathematical Society* **70** (2023), no. 2, 258–277. MR4537170
- [MWZ⁺19] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru, *Model cards for model reporting*, *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* ’19)*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 220–229.
- [N⁺20] Eirini Ntoutsi et al., *Bias in data-driven artificial intelligence systems—An introductory survey*, *WIREs Data Mining and Knowledge Discovery* **10** (2020), e1356.
- [Nar10] Arvind Narayanan, *Cookies, supercookies and ubercookies: Stealing the identity of Web visitors*, *33 Bits of Entropy: The End of Anonymous Data and What to Do About It*, February 18, 2010. <https://33bits.wordpress.com/2010/02/18/cookies-supercookies-and-ubercookies-stealing-the-identity-of-web-visitors/>

- [Nat18] National Academies of Sciences, Engineering, and Medicine, *Proactive policing: Effects on crime and communities*, The National Academies Press, Washington, DC, USA, 2018.
- [Nob18] Safiya Umoja Noble, *Algorithms of oppression: How search engines reinforce racism*, NYU Press, New York, NY, USA, 2018.
- [NS08] Arvind Narayanan and Vitaly Shmatikov, *Robust de-anonymization of large sparse datasets*, 2008 IEEE Symposium on Security and Privacy (sp 2008), 2008, pp. 111–125.
- [OBS⁺19] Osonde A. Osoba, Benjamin Boudreaux, Jessica Saunders, J. Luke Irwin, Pam A. Mueller, and Samantha Cherney, *Algorithmic equity: A framework for social applications*, RAND Corporation, Santa Monica, CA, USA, July 11, 2019. https://www.rand.org/pubs/research_reports/RR2708.html
- [Off22] Office of Science and Technology Policy, The White House, *Blueprint for an AI bill of rights: Making automated systems work for the American people*, 2022, accessed March 1, 2025. <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/>
- [OHR] OHRPP, *UCLA Office of the Human Research Protection Program*, accessed December 26, 2023. <https://ohrpp.research.ucla.edu>
- [O’N16] Cathy O’Neil, *Weapons of math destruction: How big data increases inequality and threatens democracy*, The Crown Publishing Group, New York, NY, USA, 2016. MR3561130
- [Oxf] University of Oxford Mathematical Institute, *Research using data involving humans*, accessed December 26, 2023. <https://www.maths.ox.ac.uk/members/policies/data-protection/research-using-data-involving-humans>
- [PBL24] Erasmo Purificato, Ludovico Boratto, and Ernesto William De Luca, *User modeling and user profiling: A comprehensive survey*, arXiv:2402.09660, 2024.
- [Pea24] Jordan Pearson, *Scientific journal publishes AI-generated rat with gigantic penis in worrying incident*, VICE, February 15, 2024. <https://www.vice.com/en/article/dy3jbz/scientific-journal-frontiers-publishes-ai-generated-rat-with-gigantic-penis-in-worrying-incident>
- [Per11] Marc Perry, *Harvard researchers accused of breaching students’ privacy*, The Chronicle of Higher Education, July 10, 2011. <https://www.chronicle.com/article/harvard-researchers-accused-of-breaching-students-privacy/>
- [Por21] Mason A. Porter, *Data ethics for mathematicians*, 2021. https://zerodivzero.com/short_course/aaac8c66007a4d23a7aa14857a3b778c/ttitle/181b207c7d2941278be4641ea5fe0e21
- [PSO⁺20] Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, and Sharad Goel, *A large-scale analysis of racial disparities in police stops across the United States*, Nature Human Behaviour **4** (2020), 736–745.
- [RBKL20] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy, *Mitigating bias in algorithmic hiring: Evaluating claims and practices*, Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* ’20), Association for Computing Machinery, New York, NY, USA, 2020, pp. 469–481.
- [Ree18] Chris Reed, *How should we regulate artificial intelligence?*, Philosophical Transactions of the Royal Society A **376** (2018), no. 2128, 20170360.
- [RH23] Casey Ross and Bob Herman, *Denied by AI: How Medicare Advantage plans use algorithms to cut off care for seniors in need*, STAT, March 13, 2023. <https://www.statnews.com/2023/03/13/medicare-advantage-plans-denial-artificial-intelligence/>
- [RSMCM24] Lucy Rycroft-Smith, Dennis Müller, Maurice Chiodo, and Darren Macey, *A useful ethics framework for mathematics teachers*, Ethics and mathematics education: The good, the bad and the ugly (Paul Ernest, ed.), Springer, Cham, Switzerland, 2024, pp. 359–394. MR4830892
- [RT24] Talia Ringer and Terrence Tao, *AI for Math resources*, accessed April 21, 2024. <https://docs.google.com/document/d/1kD7H4E28656ua8j0GZ934nbH2HcBLyxcRgFDduH5iQO/edit>

- [S⁺25] Herbert M. Sauro et al., *From FAIR to CURE: Guidelines for computational models of biological systems*, [arXiv:2502.15597](https://arxiv.org/abs/2502.15597), 2025.
- [SAH⁺21] Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Tulsee Doshi, and Vinodkumar Prabhakaran, *Re-imagining algorithmic fairness in India and beyond*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), Association for Computing Machinery, New York, NY, USA, 2021, pp. 315–328.
- [Sal16] Matthew Salganik, *Computational social science: Social research in the digital age*, Sociology 596, Princeton University, 2016. https://www.princeton.edu/~mjs3/soc596_f2016/
- [Sal17] Matthew Salganik, *Bit by bit: Social research in the digital age*, illustrated edition, Princeton University Press, Princeton, NJ, USA, 2017. <https://www.bitbybitbook.com>
- [SBF⁺19] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi, *Fairness and abstraction in sociotechnical systems*, Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19), Association for Computing Machinery, New York, NY, USA, 2019, pp. 59–68.
- [Sch21] G. William Schwert, *The remarkable growth in financial economics, 1974–2020*, Journal of Financial Economics **140** (2021), no. 3, 1008–1046.
- [SH22] Luke Stark and Jevan Hutson, *Physiognomic artificial intelligence*, Fordham Intellectual Property, Media & Entertainment Law Journal **32** (2022), no. 4, 2.
- [Sku21] Joe Skufca, *Incorporating ethical discussions in the mathematics classroom*, SIAM News **54** (2021), no. 06. <https://www.siam.org/publications/siam-news/articles/incorporating-ethical-discussions-in-the-mathematics-classroom/>
- [SM23] Ben Stobaugh and Dhiraj Murthy, *Predicting gender and political affiliation using mobile payment data*, [arXiv:2302.08026](https://arxiv.org/abs/2302.08026), 2023.
- [SMBM23] Vincent J. Straub, Deborah Morgan, Jonathan Bright, and Helen Margetts, *Artificial intelligence in government: Concepts, standards, and a unified framework*, Government Information Quarterly **40** (2023), no. 4, 101881.
- [SSGN17] Jessica Su, Ansh Shukla, Sharad Goel, and Arvind Narayanan, *De-anonymizing Web browsing data with social networks*, Proceedings of the 26th International Conference on World Wide Web (WWW '17), International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2017, pp. 1261–1269.
- [Ste22] Janet D. Stemwedel, *Science must not be used to foster white supremacy*, Scientific American, May 24, 2022. <https://www.scientificamerican.com/article/science-must-not-be-used-to-foster-white-supremacy/>
- [SVW21] Jathan Sadowski, Salomé Viljoen, and Meredith Whittaker, *Everyone should decide how their digital data are used—Not just tech companies*, Nature **595** (2021), 169–171.
- [SWM05] Michael P. H. Stumpf, Carsten Wiuf, and Robert M. May, *Subnets of scale-free networks are not scale-free: Sampling properties of networks*, Proceedings of the National Academy of Sciences of the United States of America **102** (2005), no. 12, 4221–4224.
- [Tao24] Terrence Tao, *Two announcements: AI for Math resources, and erdosproblems.com*, April 19, 2024. <https://terrytao.wordpress.com/2024/04/19/two-announcements-ai-for-math-resources-and-erdosproblems-com/>
- [Tho] Rachel Thomas, Twitter post, August 22, 2020, accessed March 15, 2025. https://twitter.com/math_rachel/status/1297255819965169664
- [Tho21] Rachel Thomas, *11 short videos about AI ethics*, fast.ai, August 16, 2021. <https://rachel.fast.ai/posts/2021-08-17-eleven-ethics-videos/>
- [TMP12] Amanda L. Traud, Peter J. Mucha, and Mason A. Porter, *Social structure of Facebook networks*, Physica A **391** (2012), no. 16, 4165–4180.
- [TPB24a] Rochelle E. Tractenberg, Victor Piercey, and Catherine A. Buell, *Defining “ethical mathematical practice” through engagement with discipline-adjacent practice standards and the mathematical community*, Science and Engineering Ethics **30** (2024), 15.
- [TPB24b] Rochelle E. Tractenberg, Victor Piercey, and Catherine A. Buell, *Proto ethical guidelines for mathematical practice*, 2024. <https://osf.io/x5ur9/>

- [UCLa] University of California, Los Angeles, *Data Theory at UCLA*, accessed December 26, 2023. <https://datatheory.ucla.edu>
- [UCLb] University of California, Los Angeles, *UCLA: Generative AI*, accessed April 21, 2024. <https://genai.ucla.edu/>
- [Uga17] Johan Ugander, *Truth, lies, and an ethics of personalization*, Medium, January 23, 2017. <https://medium.com/@jugander/truth-lies-and-an-ethics-of-personalization-e4ccfa7f2b84#.rzap3hm70>
- [Uga20] Johan Ugander, *Data privacy and data ethics*, Management Science and Engineering 234, Stanford University, 2020. <https://web.stanford.edu/group/msande234/cgi-bin/wordpress/>
- [Uga23] Johan Ugander, *Social algorithms*, Management Science and Engineering 231, Stanford University, 2023. <https://msande231.github.io/syllabus>
- [Vé21] Carissa Véliz (ed.), *The Oxford handbook of digital ethics*, Oxford University Press, Oxford, UK, 2021.
- [Vin16] James Vincent, *Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day*, The Verge, March 24, 2016. <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>
- [VO23] Pranshu Verma and Will Oremus, *ChatGPT invented a sexual harassment scandal and named a real law prof as the accused*, The Washington Post, April 5, 2023. <https://www.washingtonpost.com/technology/2023/04/05/chatgpt-lies/>
- [Wei21] Zach Weinersmith, *Software*, Saturday Morning Breakfast Cereal, September 3, 2021. <https://www.smbc-comics.com/comic/software>
- [Whi12] Martha C. White, *Orbitz shows higher prices to Mac users*, Time, June 26, 2012. <https://business.time.com/2012/06/26/orbitz-shows-higher-prices-to-mac-users/>
- [Wika] Wikipedia, *Creative Commons license*, accessed March 15, 2025. https://en.wikipedia.org/wiki/Creative_Commons_license
- [Wikb] Wikipedia, *Netflix Prize: Cancelled sequel*, accessed December 26, 2023. https://en.wikipedia.org/wiki/Netflix_Prize#Cancelled_sequel
- [Wikc] Wikipedia, *Social Credit System*, accessed December 26, 2023. https://en.wikipedia.org/wiki/Social_Credit_System
- [Wikd] Wikipedia, *Timnit Gebru: Exit from Google*, accessed December 26, 2023. https://en.wikipedia.org/wiki/Timnit_Gebru#Exit_from_Google
- [WJ23] Chris Wiggins and Matthew L. Jones, *How data happened: A history from the age of reason to the age of algorithms*, W. W. Norton & Company, Inc., New York, NY, USA, 2023. MR4795998
- [WSAO⁺21] Claudia Wagner, Markus Strohmaier, Emre Kıcıman Alexandra Olteanu, Noshir Contractor, and Tina Eliassi-Rad, *Measuring algorithmically infused societies*, Nature **595** (2021), 197–204.
- [WT24] Stephen M. Walk and Rochelle E. Tractenberg, *Helping students deal with ethical reasoning: The proto-Guidelines for Ethical Practice in Mathematics as a deck of cards*, arXiv:2403.16849, 2024.
- [WW24] Madi Whitmann and Chris Wiggins, *Data: Past, present, and future*, History-APMA UN2901, Columbia University, 2024, accessed December 31, 2023. <https://data-ppf.github.io>

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, LOS ANGELES, CALIFORNIA 90095, USA; DEPARTMENT OF SOCIOLOGY, UNIVERSITY OF CALIFORNIA, LOS ANGELES, LOS ANGELES, CALIFORNIA 90095, USA; AND SANTA FE INSTITUTE, SANTA FE, NEW MEXICO, 87501, USA

Email address: mason@math.ucla.edu