

---

7-31-2020

## Integrating Data Science Ethics into an Undergraduate Major

Benjamin Baumer  
*Smith College*, [bbaumer@smith.edu](mailto:bbaumer@smith.edu)


Randi L. Garcia  
*Smith College*, [rgarcia@smith.edu](mailto:rgarcia@smith.edu)

Albert Y. Kim  
*Smith College*, [akim04@smith.edu](mailto:akim04@smith.edu)

Katherine M. Kinnaird  
*Smith College*, [kkinnaird@smith.edu](mailto:kkinnaird@smith.edu)

Miles Q. Ott  
*Smith College*, [mott@smith.edu](mailto:mott@smith.edu)

Follow this and additional works at: [https://scholarworks.smith.edu/sds\\_facpubs](https://scholarworks.smith.edu/sds_facpubs)

 Part of the [Categorical Data Analysis Commons](#), [Computer Sciences Commons](#), and the [Other Mathematics Commons](#)

---

### Recommended Citation

Baumer, Benjamin; Garcia, Randi L.; Kim, Albert Y.; Kinnaird, Katherine M.; and Ott, Miles Q., "Integrating Data Science Ethics into an Undergraduate Major" (2020). Statistical and Data Sciences: Faculty Publications, Smith College, Northampton, MA.  
[https://scholarworks.smith.edu/sds\\_facpubs/20](https://scholarworks.smith.edu/sds_facpubs/20)

This Article has been accepted for inclusion in Statistical and Data Sciences: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact [scholarworks@smith.edu](mailto:scholarworks@smith.edu)

# Integrating data science ethics into an undergraduate major

Benjamin S. Baumer \*

Statistical & Data Sciences, Smith College

and

Randi L. Garcia

Psychology and Statistical & Data Sciences, Smith College

and

Albert Y. Kim

Statistical & Data Sciences, Smith College

and

Katherine M. Kinnaird

Computer Science and Statistical & Data Sciences, Smith College

and

Miles Q. Ott

Statistical & Data Sciences, Smith College

August 4, 2020

## Abstract

We present a programmatic approach to incorporating ethics into an undergraduate major in statistical and data sciences. We discuss departmental-level initiatives designed to meet the National Academy of Sciences recommendation for weaving ethics into the curriculum from top-to-bottom as our majors progress from our introductory courses to our senior capstone course, as well as from side-to-side through co-curricular programming. We also provide six examples of data science ethics modules used in five different courses at our liberal arts college, each focusing on a different ethical consideration. The modules are designed to be portable such that they can be flexibly incorporated into existing courses at different levels of instruction with minimal disruption to syllabi. We conclude with next steps and preliminary assessments.

*Keywords:* data ethics, education, case studies, undergraduate curriculum

---

\*Benjamin S. Baumer is Associate Professor, Statistical & Data Sciences, Smith College, Northampton, MA 01063 (e-mail: [bbaumer@smith.edu](mailto:bbaumer@smith.edu)). This work was not supported by any grant. The authors thank numerous colleagues and students for their support.

“The potential consequences of the ethical implications of data science cannot be overstated.”

—National Academies of Sciences, Engineering, and Medicine (2018)

## 1 Introduction

Data ethics is a rapidly-developing yet inchoate subfield of research within the discipline of data science, which is itself rapidly-developing (Wender & Kloefkorn 2017). Within the past two years, awareness that ethical concerns are of paramount importance has grown. In the public sphere, the Cambridge Analytica episode revealed how the large scale harvesting of Facebook user data without user consent was not only possible, but permissible and weaponized for political advantage (Davies 2015). Facebook CEO Mark Zuckerberg initially characterized “the idea that fake news on Facebook influenced the [2016 United States Presidential] election in any way” as “pretty crazy”—comments he later regretted (Levin 2017). Nevertheless, the subsequent tongue-lashing and hand-wringing has led to substantive changes in the policies of several large social media platforms, including several prominent public figures being banned. Popular books like O’Neil (2016) and Eubanks (2018) have highlighted how algorithmic bias can steer even well-intentioned data science products into profoundly destructive forces. These incidents have revived a sense among tech professionals, and the public at-large that ethical considerations are of vital importance.

As academics, it is our responsibility to educate our students about ethical considerations in statistics and data science *before* they graduate. To that end, recent work by Elliott et al. (2018) addresses *how* to teach data science ethics. The machine learning community convenes the ACM Conference on Fairness, Accountability, and Transparency (which includes Twitter as a sponsor), which focuses on ethical considerations in machine learning research and development. Some of the first wave of data science textbooks include chapters on ethics (Baumer et al. 2017).

Most specifically, the National Academies of Sciences, Engineering, and Medicine Roundtable on Data Science Postsecondary Education devoted one of its twelve discussions to “Inte-

grating Ethics and Privacy Concerns into Data Science Education” (Wender & Kloefkorn 2017). National Academies of Sciences, Engineering, and Medicine (2018) includes the following recommendations for undergraduate programs in data science:

Ethics is a topic that, given the nature of data science, students should learn and practice throughout their education. Academic institutions should ensure that ethics is woven into the data science curriculum from the beginning and throughout.

The data science community should adopt a code of ethics; such a code should be affirmed by members of professional societies, included in professional development programs and curricula, and conveyed through educational programs. The code should be reevaluated often in light of new developments.

In light of this, it seems clear that indifference to ethics in data science is not an informed position, and in fact, the default position of indifference prevalent in the tech community is exactly the problem we are trying to help our students solve. In this sense, indifference to ethics in data science is counter to the mission of our program, and in a larger sense to our profession.

In the major in statistical and data sciences at Smith College, we have incorporated discussions of ethics (in one form or another) into all of our classes, including the senior capstone, in which about 25% of the content concerns data science ethics. Especially in light of concerns about academic freedom, we wish to stress that this treatment is not about indoctrinating students about *what* to think, but rather to force students to grapple with the often not-so-obvious ramifications of their data science work and to develop their own compasses for navigating these waters (Heggeseth 2019). It is not a political stance—it is an educational imperative, as stressed by recommendations 2.4 and 2.5 in National Academies of Sciences, Engineering, and Medicine (2018).

In this paper, we present a programmatic approach to incorporating ethics into an undergraduate major in statistical and data sciences. In Section 2, we review and delineate notions of ethics in data science. We discuss departmental-level initiatives designed to meet the NAS recommendation for weaving ethics into the curriculum from top to bottom, and

from side-to-side as well through co-curricular programming in Section 3. In Section 4 we provide six different modules that focus on data science ethics that have been incorporated into five different courses. The modules are designed for portability and are publicly available at our website <sup>1</sup>. We review evidence of our progress in Section 5.

## 2 Ethical considerations in statistics and data science

Ethical considerations in statistics have been taught for decades, going back to the classic treatment of misleading data visualization techniques in Huff (1954). In this section, we review the literature on teaching data science ethics with an eye towards explicating different notions of what falls under that umbrella.

### 2.1 Types and scope of ethical considerations

From a legal perspective, the General Data Protection Regulation (European Parliament 2018)—which became enforceable in 2018—provides Europeans with greater legal protection for personal data stored online than is present in the United States. This discrepancy highlights the distinction between ethical and legal considerations—the former should be universal, but the latter are patently local. At some level, laws reflect the ethical values of a country, but a profession cannot abdicate its ethical responsibilities to lawmakers. As O’Neil notes: “it is unreasonable to expect the legal system to keep pace with advances in data science.” (Wender & Kloefkorn 2017)

Within statistics, a major ethical focus has been on human subjects research. *The Belmont Report* is still required reading by institutional review boards (IRBs) (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1978). It posits three major ethical principles (respect for persons, beneficence, and justice) and outlines three major applications (informed consent, assessment of risks and benefits, and selection of subject). Yet just as we reject the argument that all legal data science projects are ethical, we question the supposition that all IRB-approved data science projects are ethical. IRBs have not been able to keep pace with the rapid development of data science

---

<sup>1</sup><https://bit.ly/2v2cf8n>

research, and have little authority over research fueled by data collected by corporations.

For example, Facebook data scientists manipulated the news feeds of 689,003 users in order to study their “emotional contagion” (Kramer et al. 2014). While Facebook did not break the law because users relinquished the use of their data for “data analysis, testing, [and] research” when they agreed to the terms of service, many ethical questions were subsequently raised, notably whether informed consent was legitimately obtained. Moreover, Cornell University IRB approval was obtained only *after* the data had been collected, meaning that the approval covered the analysis of the data, not the collection or the design of the experiment. This example illustrates how university IRBs are ill-equipped regulate “big data” studies (Meyer 2014).

Major professional societies, including the American Statistical Association (ASA) (Committee on Professional Ethics 2018*b*), the Association for Computing Machinery (ACM) (Committee on Professional Ethics 2018*a*), and the National Academy of Sciences (NAS) (Committee on Science, Engineering, and Public Policy 2009), publish guidelines for conducting research. These documents focus on topics like professionalism, proper treatment of data, negligence, and conflicts of interest. Similarly Tractenberg (2019*a*), Tractenberg (2019*b*), and Gunaratna & Tractenberg (2016) explore ethics in statistical practice but don’t mention newer concepts like *algorithmic bias*. Loukides et al. (2018) focuses on industry and identifies five framing guidelines for building data products: consent, clarity, consistency, control, and consequences. For oversight, Germany is considering recommendations for a data science ethics review board (Tarran 2019). Canney & Bielefeldt (2015) present a framework for evaluating ethical development in engineers.

A broader discussion of professional ethics in statistics and data science would include issues surrounding reproducibility and replicability, which would include concepts like transparency, version control, and p-hacking (Wasserstein et al. 2016, Wasserstein et al. (2019)). Inappropriate analysis remains a problem in many fields, including biostatistics (Wang et al. 2018). The machine learning community is having intense debates about the extent to which data or algorithms are ultimately most responsible for bias in facial recognition and other AI-driven products (Cai 2020).

While these areas remain crucially important—and continue to play a role in our

curriculum—we focus here on more modern manifestations of data science ethics brought on by “big data.” These include primarily algorithmic bias, but also ethical concerns when scraping data from the web, storing your personal data online, de-identifying and re-identifying personal data, and large-scale experimentation by internet companies in what Zuboff (2018) terms “surveillance capitalism.” These ethical areas are obviously informed by longstanding ethical principles, but are distinct in the way that computers, the Internet, and databases have transformed the way we live (Hand 2018).

Our focus areas mostly intersect with those identified by National Academies of Sciences, Engineering, and Medicine (2018) as needed by data scientists:

- Ethical precepts for data science and codes of conduct,
- Privacy and confidentiality,
- Responsible conduct of research,
- Ability to identify “junk” science, and
- Ability to detect algorithmic bias.

This paper offers examples for implementing these focus areas. For example, Section 4.6 contains a module that has students apply ethical codes in context. The modules in Sections 4.1 and 4.3 explore notions of privacy and confidentiality. Sections 4.5 and 4.3 provide modules that illuminate notions of responsibility when conducting research. Sections 4.2 and 4.6 present modules that encourage students to detect algorithmic bias in action. Yet we also go beyond these key areas. The module in Section 4.4 explores boundaries between legal and ethical considerations. In other activities not presented here, we engage students in our senior capstone and machine learning courses with deep questions about the impact that actions by large-scale Internet companies have on our lives.

## **2.2 Approaches to teaching data science ethics**

While discussion about data science ethics abounds, there are few successful models for how statisticians and data scientists can teach it. Indeed, relevant work on teaching data science by Donoho (2017), Hicks & Irizarry (2018), Baumer (2015), Hardin et al. (2015), and Kaplan (2018) either barely mentions ethics or doesn’t mention them at all. Despite

recommending the inclusion of ethics into data science curricula, even the the National Academies of Sciences, Engineering, and Medicine (2018) report does not include explicit recommendations for *how* to do so. One of the primary challenges is that while educators are typically well-trained in the ethics of human subjects research, few have explicit training in, say, algorithmic bias, or even general ethical philosophy. But why should a lack of training prevent us from teaching our students? As Bruce (2018) points out, ethical issues are not really a technical problem, but rather “a general issue with the impact of technology on society,” to which we all belong. We might make up for our lack of training by partnering with philosophers and ethicists to develop a robust ethical curriculum (Bruce 2018).

Echoing Bruce (2018) that “there is a long history of scholars and practitioners becoming interested in ethics when faced with new technologies,” Gotterbarn et al. (2018) argue forcefully that the recent uptick in interest in “computing ethics” is merely the most recent star turn for a longstanding and valued component of the computer science curriculum. While this is surely true at some level and important to keep in mind, it hardly seems like the renewed attention on ethics is unwarranted. Moreover, Gotterbarn et al. (2018)’s focus is on artificial intelligence driven systems like self-driving cars, whereas our focus is on *data*.

Several examples of how to teach ethics in statistics, data science, and (mostly) computer science exist. Neff et al. (2017) takes a broad view of data science ethics, bringing tools from critical data studies to bear on the practice of actually doing data science. Burton et al. (2018) outlines a strategy for teaching computer science ethics through the use of science fiction literature. Elliott et al. (2018) provides a framework for reasoning about ethical questions through the dual prisms of Eastern (mainly Confucianism) and Western ethical philosophies. We found this inclusive approach to be particularly valuable given the large presence of international (particularly Chinese) students in our classes. Perhaps presaging many recent scandals, Zimmer (2010) analyzes a Facebook data release through an ethical lens. Fiesler analyzes ethical topics in a variety of computer science courses (Saltz et al. 2019, Fiesler et al. 2020, Skirpan et al. 2018). Grosz et al. (2019) describes how ethics education is integrated into the computer science curriculum at Harvard. Barocas teaches an undergraduate elective course on data science ethics at Cornell (Wender & Kloefkorn



2017).

These articles offer guidance on how to teach ethics in data science, but leave many stones unturned. In this paper, we present six additional concrete modules for teaching data science ethics, as well as outlining departmental initiatives for fully integrating ethics into a data science curriculum and culture.

### 3 Department-level initiatives

At Smith, every department periodically reviews and updates a list of learning goals for their major. The major in statistical and data sciences (SDS) is designed to cover a broad range of topics to produce versatile future statisticians and data scientists. Our learning goals include skills like: fitting and interpreting statistical models, programming in a high-level language, working with a wide variety of data types, understanding the role of uncertainty in inference, and communicating quantitative information in written, oral, and graphical forms. Most recently, we added the following learning goal:

Assess the ethical implications to society of data-based research, analyses, and technology in an informed manner. Use resources, such as professional guidelines, institutional review boards, and published research, to inform ethical responsibilities.

In support of this learning goal, we have taken measures to:

- incorporate ethics into all of our classes, culminating in a thorough treatment in the senior capstone course.
- support student engagement in extra-curricular and co-curricular events that touch on data science ethics.
- bring a diverse group of speakers to campus to give public lectures that often focus on ethical questions.
- include a candidate's ability to engage with data science ethics as a criterion in hiring.
- increase inclusion at every level of our program.

We discuss six specific modules for courses in Section 4. In this section we discuss approaches for the other measures. We recognize that not every institution has the curricular flexibility and resources that we have at Smith, nor is our student body representative of those at different types of institutions (e.g., R1’s or two-year colleges). Nevertheless, most of the modules we present can be stitched into a single class period, which should provide instructors at any institution with a reasonable opportunity to incorporate some of this material.

Our students are very interested in ethical questions in data science (see Section 5.2). As digital natives, they bring an importantly different perspective to questions about, for example, sharing one’s personal data online. Many of them have never seriously considered the ramifications of this. The notion that “if you’re not paying for the product, then you are the product” is new, scary, challenging, relevant, personal, and engaging to them in a way that helps them see data science as more than just a battery of technical skills (Fitzpatrick 2010). Thus, teaching ethics in data science is another way to foster student interest in the *discipline*. Framing ethical questions in data science as unsolved problems helps students imagine themselves making meaningful contributions to the field in a way that may seem too remote of a possibility in, say, estimation theory.

In particular, algorithmic bias intersects with questions about inclusion and diversity with which students are already grappling on a daily basis. During the past two years, we have applied for (and received) funds from the community engagement center and the Provost’s office to support student engagement with the *Data for Black Lives* conference (Milner 2019). In 2018, the first year of the Data for Black Lives conference, we hosted a remote viewing party on campus. In 2019, one of us attended the conference with five students. This experience led to *a student* inviting Data for Black Lives founder Yeshimabeit Milner to campus for a public lecture entitled “Abolish Big Data”. These experiences help students connect what they are learning in the classroom to larger movements in the real world, and give them the sense that their skills might be used to affect positive change in the world—a powerful motivator.

The SDS major at Smith includes an “application domain” requirement. One of the purposes of this requirement is to ensure that students understand that all data and analyses

have a context. Conducting ethical data analysis requires knowledge of the context in which the data is being used. For example, only through having some understanding of the history of racial/ethnic groups in the United States can data scientists hope to code and use race appropriately in their analyses (see Section 4.5).

The major at Smith requires every student to take one course that focuses explicitly on communication. Another simple initiative was to allow students to fulfill this requirement by taking the “Statistical Ethics and Institutions” course taught at nearby Amherst College by Andreas V. Georgiou, the former President of the Hellenic Statistical Authority (Langkjær-Bain 2017). Although the course did not explicitly focus on communication, we made an exception to our policy to allow students to have this unique opportunity to learn about statistical ethics from the person at the center of a world-famous episode. Moreover, ethics and communication are intertwined, in that conveying ethical subtleties requires a different skill set than say, explaining a statistical model.

We are fortunate that our institution provides generous funding for bringing outside speakers to campus, and we have taken full advantage of their largesse over the past two years. We welcomed BlackRock data scientist Dr. Rachel Schutt to give a talk titled “A Humanist Approach to Data Science,” in which she underscored the importance of recognizing the people behind the numbers, and highlighted examples of recently published research that raised profound ethical dilemmas. Dr. Terry-Ann Craigie of Connecticut College came to talk about the intersections of race, data science, and public policy. Dr. Emma Benn of Mount Sinai discussed how her intersectional social identity has informed her work as a biostatistician. Alumna Gina DelCorazon spoke about her experiences as Director of Data & Analytics at the National Math and Science Initiative in her talk “From Interesting to Actionable: Why good context matters as much as good code.” At the invitation of a student group, Dr. Alisa Ainbinder, an alumna working locally in data science, discussed ethical considerations in her work in non-profit accelerator programs. Hearing from professionals about the ethical considerations in their work helps reinforce the messaging we give them class.

Finally, we take small steps to ensure that incoming faculty are capable of supporting our program in meeting this newest learning goal. They cannot be dismissive of ethical concerns

in data science. In the same way that a candidate who didn't understand correlation would not be hireable, we consider whether a candidate who seemed ignorant of data science ethics would be hireable. To assess this, we might ask a question about data science ethics during a first round or on-campus interview. We might ask candidates to submit a separate statement on data science ethics as part of their application, or to discuss ethical considerations in their teaching and/or research statement. To be clear, we cannot and do not infringe upon the candidate's academic freedom by assessing *what* they think about data science ethics. Rather, we are merely trying to assess *how deeply* they have thought about data science ethics and thus whether they are sufficiently prepared to help the program meet our learning goals.

## 4 Modules for teaching data science ethics

In this section we present six modules for teaching ethics in data science that are used in a variety of courses. Here, we give a brief description of each module, its learning goals, and the context of the course in which it is delivered. In our supplementary materials, we provide more complete teaching materials.

### 4.1 Three uses of OkCupid data

OkCupid is a free online dating service whose data has been scraped on at least three known occasions. Kim & Escobedo-Land (2015) presented scraped data on nearly 60,000 OkCupid users in the greater San Francisco area in the early 2010's for use in the classroom and subsequently released the data as the R package `okcupiddata`. Around that same time, Chris McKinlay created 12 fake OkCupid accounts and wrote a Python script that harvested data from around 20,000 women from all over the country (Poulsen 2014). In 2016, Kirkegaard & Bjerrekær (2016) published a paper in an open-access psychology journal investigating a variety of hypotheses about OkCupid users—along with the corresponding data from 70,000 users. From the same underlying data source, these three incidents provide fertile ground for substantive discussions about the corresponding ethical considerations.

Some further detail reveals fascinating disparities:

- Kim & Escobedo-Land (2015) obtained explicit permission from OkCupid CEO Christian Rudder before publishing the data in a statistics education journal. Their goal was to illuminate statistical phenomenon using data that was relevant to undergraduate students. In addition, the authors removed usernames from the data as a modest attempt at de-identifying the users. Only later were the authors alerted to the fact that even though usernames had been stripped, the full-text of the `essay` field often contains personally-identifying information like Facebook and Instagram handles.
- McKinlay did not publish the data he collected—his goal was personal. Essentially, he trained his own models on the data he collected to find his own match. It worked—he is now engaged to the woman he met. Only after his story was published were questions raised about whether he had violated the Computer Fraud and Abuse Act.
- Kirkegaard & Bjerrekær (2016) included username, age, gender, and sexual orientation in the data set. This meant that users were easily identifiable and particularly vulnerable. While the blowback in this case was immediate, Kirkegaard insisted that the data were already public and his actions were legal.

Collectively, these episodes raise issues about informed consent, data privacy, terms of use, and the distinction between laws and ethics. One could use these incidents to motivate coverage of technical concepts such as  $k$ -anonymity (Sweeney 2002) and differential privacy (Dwork et al. 2006). In our senior capstone course (see Section 4.6), we ask students to break into three groups and discuss the relevant ethical issues involved in each case. Then, we bring students together to write a coherent response. Some students elect to use these incidents as the subject of a longer essay, as described in Section 4.6.

## 4.2 Algorithmic bias in machine learning

Discussions on the perniciousness of “algorithmic bias” in machine learning and artificial intelligence have become more prevalent of late, both in the news media as well as in academic circles (Noble 2018, Eubanks (2018) O’Neil (2016)). However, few of these

ideas have been incorporated into the classroom. For example, in James et al. (2013)—a popular introductory textbook on machine learning—the `Credit` dataset is often used as an example (it is available in the companion `ISLR R` package). Readers are encouraged to apply various predictive algorithms predict the credit card debt of 400 individuals using demographic predictors like `Age`, `Gender` (encoded as binary), and `Ethnicity` with levels `African American`, `Asian`, and `Caucasian`. While the data is simulated, one must still wonder what kind of thinking are we tacitly encouraging to students by using ethnicity to predict debt and thus perhaps credit score. This is especially fraught in light of existing inequalities to access to credit that fall on demographic lines. In other words, to quote Milner (2019), “What are we optimizing?”





In this module, we propose a hands-on in-class activity to help students question the supposed objectivity of machine learning algorithms and serve as a gateway to discussions on algorithmic bias. The activity centers around StitchFix, an online clothing subscription service that uses machine learning to predict which clothes consumers will purchase. New users are asked to complete either a men’s or women’s “Style Profile” quiz, whose responses are then used as predictor information for the company’s predictive algorithms. However, both quizzes differ significantly in the types of questions asked, how the questions are asked, in which order they are asked, and what information and visual cues are provided.

Figure 1 (current as of December 16, 2019) presents one example relating to clothing style preferences, specifically jean cut. The prompt in the men’s quiz shows photographs of an individual actually wearing jeans, whereas the women’s quiz presents the options in a much more abstract fashion. On top of differences relating to clothing style and fit, many differences exist in how demographic information is collected. Figure 2 presents an example of a question pertaining to age. While both groups are asked the same question of “When is your birthday?” individuals completing the women’s quiz are primed with a “We won’t tell! We need this for legal reasons!” statement, whereas those completing the men’s are not. One has to suspect such a difference was not coincidental, but rather reflects a prior belief of the quiz designers as to the manner in which one should ask about age. Other differences include questions pertaining to parenting and occupation.

Many of these differences can be attributed to prevailing biases and beliefs on the nature

How do you like your jeans to fit?

Select all that apply.

<b>SKINNY</b>	<b>SLIM</b>	<b>STRAIGHT</b>	<b>RELAXED</b>
			
Cut narrow through the thigh, tight at the ankle	Cut close through the thigh, slightly narrow at the ankle	Comfortable through the thigh, straight below the knee	Relaxed through the thigh, straight below the knee

Choose all the **styles of jeans** you would wear






			
<input type="radio"/> Skinny	<input checked="" type="radio"/> Straight	<input type="radio"/> Bootcut	<input type="radio"/> Wide

Figure 1: Example difference between men's (left) and women's (right) StitchFix Style Quizzes: Question on jean preferences. Contrast the abstract presentation of jeans shown to women with a picture of someone actually wearing jeans shown to men.

When is your birthday?

mm/dd/yyyy



When is your birthday?

We won't tell! We need this for legal reasons.

1	,	1	,	1970
MM		DD		YYYY

Figure 2: Example difference between men's (left) and women's (right) StitchFix Style Quizzes: Question about age. We note the disclaimer present for women is omitted for men.

of gender and thus can serve as fertile ground for student discussions on algorithmic bias.

Specifically, this module can satisfy three goals. First, it provides students with an example of algorithmic bias to which they can directly relate. This stands in contrast to more abstract and much less accessible examples discussed in academic readings and news media, such as facial recognition software. Second, it asks students to view the statistical, mathematical, and machine learning topics covered in class through a sociological lens, in particular relating to the nature of gender. Third, it gives students the opportunity to think about statistical models in a rich, real, and realistic setting, in particular what predictor variables are being collected and what modeling method/technique is being used.

### 4.3 Social networks

Perhaps in part thanks to the aptly-named Facebook movie (*The Social Network*), social networks are intuitive to students. The relatively simple mathematical formulation of networks (i.e., graphs) makes them easy to understand, but the complex relationships and behaviors in such networks lead to profound research problems. Moreover, analyzing social network data leads to thorny ethical questions.

A 300-level course on statistical analysis of social network data has as its primary objective for students to “learn how to answer questions by manipulating, summarizing, visualizing, and modeling network data while being vigilant to protect the people who are represented in that data.” Thus, ethical concerns surrounding privacy and confidentiality are woven directly into the main course objective.

The primary textbook is Kolaczyk & Csárdi (2014), which provides a thorough treatment of both the theoretical and applied aspects of social network analysis. However, supplementary readings are especially important, since Kolaczyk & Csárdi (2014) fails to address the many complex ethical issues that arise for these data. We employ supplemental readings to address data ethics on topics including:

- collecting social network data
- informed consent for social network surveys
- data identifiability and privacy in social networks
- link prediction



- data ethics specific to social networks

In our supplementary materials we present a module applied during the first week of class in which we use an example from popular culture (the television show *Grey's Anatomy*) to motivate ethical issues in social network analysis. It has several goals:

- Prime students to always think about how the data were collected
- Prime students to think about the benefits of and risks of each data collection / analysis / visualization, etc.
- Encourage students to create their own understanding of how data ethics pertain to social network data as opposed to being provided with data ethics rules. This encourages critical thinking which can then be transferred to other topics and types of data.

It is especially important to introduce ethical considerations on the first day of the course to set the tone and give students the message that data ethics is inextricable from the rest of the content of the course.

#### 4.4 Copywriting music

Ethical usage of data can come into conflict with copyright law. Music usage, for example, is heavily protected by copyright laws. The field of Music Information Retrieval (MIR) seeks to address questions about music, such as finding all covers of a particular song or detecting the genre of a song. In MIR, access to music is critical to conducting research, and that access is governed by copyright laws.

Music is also a medium that has a fraught history navigating the line between sharing and violating copyright. This history is complicated by the power dynamics at play between recording companies and artists, and recording companies and listeners. Today, music is often consumed through streaming services, distorting our understanding of music ownership. Since music is heavily protected by copyright but remains omnipresent in our lives, conversations about data access require nuance about ownership, sharing, and the subtleties of ethical vs. legal considerations.

Understanding that the goal of copyright is to protect artists, and then contrasting students' experiences of accessing and digesting music, this debate's overarching goal is to have students navigate legal considerations (i.e., copyright) and ethical considerations (i.e., when to share or not share data) in the contexts of pushing research forward and of the capitalist motivations of the music industry. The legal restrictions of copyright and the ethical responsibilities of a researcher to protect and appropriately use (and share) their data provide a fascinating grey area for this debate. The generational experience of our current students informs their notions of morality and access, which in turn leads them to confront legal restrictions in an interesting way.

To explore data access and copyright, we provide a module in which students have a debate about whether the music copyright laws should be softened for those conducting MIR research. This debate is not as simple as whether to relax these laws, instead one side is defending the role and purpose of copyright laws for music while the other side not only advocates for relaxing these laws but also for *how* to accomplish this. This requirement of proposing a solution required students to hold the responsibilities of a researcher who has broad access to data in contrast with the ease at which we can share music (and data).

This debate activity was originally part of a senior seminar introducing students to MIR, but this activity could be done in any course where data provenance, data usage, or data access is discussed. For this activity, students were randomly assigned to one side of the debate. In preparation for the debate, students were required to submit a position paper (due just before the debate) that presented a coherent argument that is well supported by the literature. Students were also barred from sharing arguments with each other (even if assigned to the same side of the debate). However, they could share resources with each other (just not their opinion of these resources).

This structure of a preparatory paper followed by a debate required students to engage with the research process at a deep level. For the actual debate, each side was given opportunities to present their ideas and offer rebuttals to the other side. This meant that not only did they have to find resources and digest them, they had to discuss the ideas both in written text and orally in a debate setting.

## 4.5 Teaching about race and ethnicity data

In an upper-level research seminar on intergroup relationships cross-listed in the psychology department, students learn the psychology of close relationships between people who have differing social group identities (e.g., racial/ethnic and gender group identities). In addition, students learn to analyze dyadic data through multilevel modeling (i.e., mixed linear modeling), and write reproducible research reports in APA format with the R package `papaja` (Aust & Barth 2018). This course attracts a diverse group of students in terms of majors, professional goals, interests, statistical preparation, and personal identities. In this ethics module, we describe a discussion and data cleaning activity used to get students thinking in a more careful and nuanced way about the use of race and ethnicity data in their analyses.

The instructor provides psychological data from her own research program, and the overarching focus of the course is to form research questions answerable through the analysis of data that has already been collected. Since the focus is on analyzing existing data (in addition to talking about race), we also discuss:

- how to communicate transparently one's use of confirmatory versus exploratory analyses
- the philosophical differences between inductive and deductive reasoning
- the prevention of p-hacking (Wasserstein et al. 2016, Wasserstein et al. (2019)) and HARKing (Hypothesizing After the Results are Known; Kerr 1998)

On the first day of this course, we have a class discussion about how we will try to create a climate of psychological safety (Edmondson 1999) together. This initial discussion helps to set the tone of respect and generosity that we will need in order to have fruitful discussions about race and ethnicity data. In the first half of the course, class sessions alternate between discussions about assigned readings (from psychology) and the statistical and data science instruction they need to complete their projects. In the second half of the course, class sessions are mainly used for actively working on their projects. The two parts of this ethics module (discussion and data cleaning) might be split across two class sessions.

The activity described in this module consists of a class discussion about race and a race/ethnicity data cleaning activity in the context of a psychology article about interracial

roommate contact (Shook & Fazio 2008). The structure of this activity invites students to discuss the article first in small groups, and then as a class. The larger class discussion portion of this activity is designed to evolve into a broader discussion about the coding and use of race and ethnicity data in quantitative research. Some important revelations that might be pulled from the discussion include:

- Researchers studying interracial interactions make choices about who to focus on, and, in the past, this choice has often been to focus on white participants only. An acknowledgement of white privilege and who, historically, has been asking the research questions might come out as well.
- A person's personal racial/ethnic identity may be different from how they are perceived by another person (roommate).
- The choice to use a person's own racial/ethnic identity data or someone's perception of their race depends, in part, on the research question. When is identity or perception more important for the specific research context?
- Race is not as clear of a categorical variable as we think it is. Can we think of other instances of this, for example, with gender categorization?
- Are there times when it could serve a social good to use race in our analyses and, in contrast, are there ways in which using race and ethnicity data in analyses might reify socially constructed racial categories?
- If you decide to use race in your analyses, what might you do in smaller samples if there are very small numbers of ethnic minority groups relative to White/European-Americans? Is it ever OK to collapse racial/ethnic categories? What immediate consequences do these choices have for the interpretation of your analysis and what broader consequences might these choices have when your results are consumed by your intended audience?

The second part of this activity asks students to code raw race/ethnicity data into a new categorical variable called `race_clean`. They do this part in pairs. Then, in small groups, they discuss the decisions they made when completing this task and also any feelings they had during the task, as those feelings reflect the hard realities that researchers must confront in their work. The raw data comes in check-all-that-apply and free response

formats. Students will find this task quite difficult, and perhaps uncomfortable. The goal is not to have them finish, but to get them to recognize the ambiguity inherent the construction of categorical race/ethnicity variables. They may have used the clean version of race/ethnicity variables in the past without thinking much of it.

Lastly, the module also contains notes on closing thoughts the instructor might offer their students after this activity. It is very important not to skip the wrap-up for this activity. Let students know that this is not the end of the discussion. As future data scientists, they can play an active role in creating ethical guidelines for moving towards more appropriate use of race and ethnicity data.

#### **4.6 *Weapons of Math Destruction* in the senior capstone**

In the senior capstone course, roughly 25% of the course is devoted to learning about data science ethics. During the first half of the semester, we spend *every other class period* discussing ethical considerations that arise from weekly readings of O’Neil (2016). These readings introduce students to episodes in which often well-intentioned data science products (e.g., criminal sentencing algorithms, public school teacher evaluations, US News and World report college rankings, etc.) have had harmful effects on society. These episodes are accessible to students and provide many opportunities to engage students in thoughtful conversation.

The material in O’Neil (2016) also intersects with a wide variety of statistical topics, such as modeling, validation, optimization, Bayesian statistics, A/B testing, Type I/II errors, sensitivity and specificity, reliability and accuracy, Simpson’s paradox, multicollinearity, confounding, and decision trees. A clever instructor could probably build a successful course entirely around these topics.

Moreover, the ethical considerations that O’Neil (2016) raises about algorithmic bias, informed consent, transparency, and privacy, also touch on hot-button social questions surrounding structural racism, gender equity, software licensing, cheating, income inequality, propaganda, fake news, scams, fraud, pseudoscience, and policing bias. Situated in the fallout from the 2008 global financial crisis, but presaging Cambridge Analytica and fake news, the book feels simultaneously dated and relevant. Our students lived through the

global financial crisis but most were too young to understand it—for many of them the book allows them to grapple with these events for the first time as adults.

The first major goal of the module is to raise awareness about the manifold ethical considerations in data science. Reading O’Neil (2016) and having class discussions about the material will help accomplish this learning goal. We employ a variety of techniques, including think-pair-share, breakout groups, student-led discussions, and even lecturing to keep students engaged in class.

However, the second major goal is to have students write something constructive about data ethics. To this end, more structured readings are needed. We present students with two frameworks for thinking critically about data science ethics: Data Values and Principles (Gershkoff et al. 2019) and the Hippocratic Oath for Data Science (National Academies of Sciences, Engineering, and Medicine 2018). The former defines four values (inclusion, experimentation, accountability, and impact) and twelve principles that “taken together, describe the most effective, ethical, and modern approach to data teamwork.” The latter provides a data science analog to the oath that medical doctors have taken for centuries. We then ask students to write an essay in which they analyze a data science episode—perhaps drawn from O’Neil (2016)—in the context of one of these frameworks.

During the course, students write four papers of varying length on data science ethics. Together, these assignments not only impress upon students the importance of ethics in data science, but also give them tools and experience to reason constructively about data science ethics in the future. The goal is to produce students who have fully integrated ethical questions into their understanding of statistics and data science.

## **5 Assessment of our ethical curriculum**

Early returns suggest that our emphasis on teaching data science ethics is having an impact. To support this claim we relate two concrete anecdotes, analyze results from an anonymous student survey, and provide several free responses from students.

## 5.1 Data science ethics in action

One student used her experience with data science ethics directly in a summer internship with an anonymous company to help draft the company’s heretofore non-existent policies around ethical data use (Conway Center for Innovation and Entrepreneurship 2019).

“[She] was also the first data scientist to work in the [company] space. Until her arrival, [company]’s businesses lacked clear guidelines for collecting data and ways for using that data to generate insights. Surprised by this, [she] first initiated conversations with the [company] team around ethical concerns in data collection. **Drawing on lessons from her academic work**, and discussions with her Smith mentors, she helped to develop policies for [company] businesses to ethically collect, manage, and act on customer data moving forward.”

While this is clearly just one example, we note that the connection between data science ethics in practice and her academic coursework was made explicit *by the student*.

Another anecdote involves a student group supporting students in our major that held their annual “Data Science Day” on November 9th, 2019. At the open house portion of the event, in addition to operating booths on data visualization and machine learning, students set up a “data ethics” booth with handouts posing ethical and philosophical questions about the use of data (see Figure 3). While this event was sponsored by the program, programming for the event was entirely determined by students. The inclusion of the booth suggests that students see ethics as an integral component of data science, on par with data visualization and machine learning. We interpret this as an early sign of our program’s success at emphasizing the importance of ethical thinking in data science.

Furthermore, in the wake of discussions on racism and white supremacy spurred by the death of George Floyd in May 2020, two students created a Data Science Resources for Change website. They state: “In order to be thoughtful, effective, and inclusive data scientists, we believe it is important to understand the ways in which bias can play a dangerous role within our field, to understand the ways in which data can be used to either reinforce/exacerbate or fight oppression, and to support the inclusion of voices of color within the community.” To this end this website includes numerous resources such as reading lists, videos and podcasts, organizations to support, and notable people to follow.



Figure 3: The SDS student group chose to staff a 'data ethics' booth at Data Science Day 2019.

## 5.2 Analysis of survey responses

We conducted an anonymous online survey during the summer of 2019, in which 23 students participated<sup>2</sup>. The results in Figure 4 reveal that students are interested in learning more about data science ethics and feel that it is an important part of their education. However, they are less certain that they have achieved our stated learning goal. Unfortunately, none of the respondents had taken the capstone course (see Section 4.6), and so these results almost certainly undersell the effectiveness of our ethical curriculum.

The first panel in Figure 4 reflects self-assessments from students about two aspects of our major learning goal. The questions reflect both the ability of a student to assess the ethical implications of data science work, as well as their ability to draw on published materials to inform their thinking. These ideas are most explicitly and thoroughly tackled in the senior capstone, and so the lack of respondents with that course under their belt renders this picture incomplete.

The second panel addresses the importance of ethics to a student's data science edu-

---

<sup>2</sup>This survey was approved by the Smith College IRB, protocol 18-111.



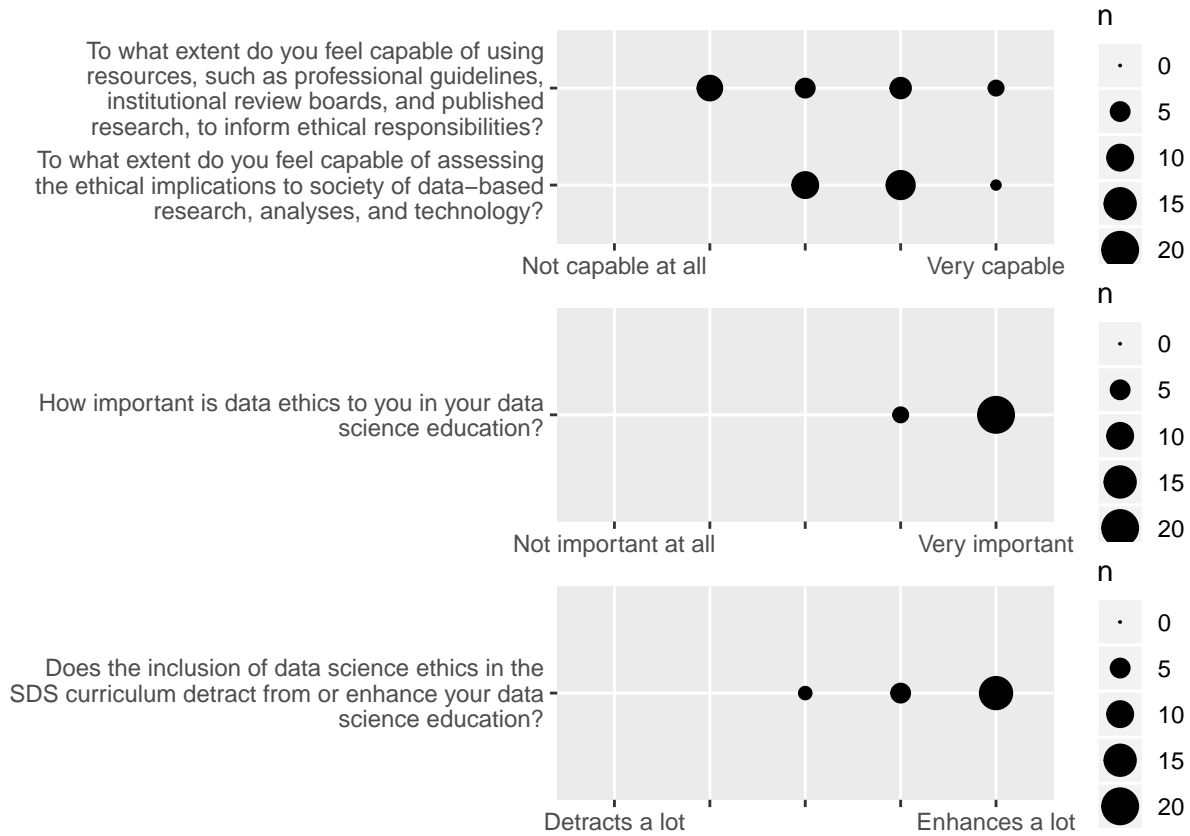


Figure 4: Student self-assessment of their ethical capabilities, and the importance of data science ethics in their education, from an anonymous survey of 23 students. We note that nearly all respondents saw the inclusion of data science ethics as an important enhancement to their education, although they were less certain of their own capabilities in analyzing ethical concerns.

cation. Here, students *universally* believe that data science ethics is important *to them* in their education, with most responding that it is “very important.” This finding supports the recommendation of National Academies of Sciences, Engineering, and Medicine (2018).

Finally, the third panel in Figure 4 makes plain that no students feel that the inclusion of data science ethics detracts from their data science education, with most students seeing the inclusion as an enhancement. We encourage data science programs contemplating adding ethical content to consider this point particularly. That is, the respondents to this survey did not see the inclusion of data science ethics as a distraction from more important, interesting, technical, or valuable content. Rather, learning about data science ethics enhances that curriculum.

### **5.3 Anonymous text responses from the survey**

In Appendix ?? we present selected quotes in their full context. Here, we highlight a few of the most relevant thoughts and connect them to broader themes.

First, faculty should not assume that students know about data science ethics just because it is often in the news. To the contrary, learning about data science ethics can be revelatory for students.

“It was the first time I had ever thought that data science had ethical implications and it really changed the way I thought about the work that I do.”

Second, as noted by National Academies of Sciences, Engineering, and Medicine (2018), teaching data science ethics as a one-off topic is not likely to be sufficient, and students notice the tangential nature of this approach.

“I’d like to see more data ethics integrated with in-class work. In my experience, data ethics has been presented as an additional topic as opposed to something that is an intrinsic part of data science work itself.”

Third, far from being off-putting, these students found ethics in data science to be a topic likely to engage a broader set of students with data science.

“I also think that providing... a push for students outside of the major [to] have access to more resources about data ethics, (i.e. data talks which include ethics being more widely broadcast to the rest of the student body) should be seriously considered.”

## 6 Conclusion

The long-term health of data science as a discipline relies on public trust. Ethical lapses, or gross indifference to ethics, has resulted in the deployment of data science products that are harmful to society, due to biases that we now recognize. Our students are part of the generation of data scientists that will address these issues and restore faith in data-driven applications. In order to do this, they need to see weighing ethical considerations as an integral part of the process of doing data science. We present our approach to achieving this in the hopes that others will emulate and refine what we have started.

## References

- Aust, F. & Barth, M. (2018), *papaja: Create APA manuscripts with R Markdown*. R package version 0.1.0.9842.  
**URL:** <https://github.com/crsh/papaja>
- Baumer, B. S. (2015), ‘A data science course for undergraduates: Thinking with data’, *The American Statistician* **69**(4), 334–342.  
**URL:** <http://dx.doi.org/10.1080/00031305.2015.1081105>
- Baumer, B. S., Kaplan, D. T. & Horton, N. J. (2017), *Modern Data Science with R*, Chapman and Hall/CRC Press: Boca Raton.  
**URL:** <https://www.crcpress.com/Modern-Data-Science-with-R/Baumer-Kaplan-Horton/9781498724487>
- Bruce, K. B. (2018), ‘Five big open questions in computing education’, *ACM Inroads* **9**(4), 77–80.  
**URL:** <https://dl.acm.org/citation.cfm?id=3230697>

Burton, E., Goldsmith, J. & Mattei, N. (2018), ‘How to teach computer ethics through science fiction’, *Communications of the ACM* **61**(8), 54–64.

**URL:** <https://dl.acm.org/citation.cfm?id=3154485>

Cai, F. (2020), ‘Yann LeCun quits Twitter amid acrimonious exchanges on AI bias’, Synced: AI Technology & Industry Review.

**URL:** <https://syncedreview.com/2020/06/30/yann-lecun-quits-twitter-amid-acrimonious-exchanges-on-ai-bias/>

Canney, N. & Bielefeldt, A. (2015), ‘A framework for the development of social responsibility in engineers’, *International Journal of Engineering Education* **31**(1B), 414–424.

Committee on Professional Ethics (2018a), *ACM Code of Ethics and Professional Conduct*, Association for Computing Machinery, Inc.

**URL:** <https://www.acm.org/binaries/content/assets/about/acm-code-of-ethics-booklet.pdf>

Committee on Professional Ethics (2018b), Ethical guidelines for statistical practice, Technical report, American Statistical Association.

**URL:** <http://www.amstat.org/asa/files/pdfs/EthicalGuidelines.pdf>

Committee on Science, Engineering, and Public Policy (2009), *On being a scientist: a guide to responsible conduct in research*, 3 edn, Washington, DC: National Academies Press.

**URL:** <https://www.ncbi.nlm.nih.gov/pubmed/25009901>

Conway Center for Innovation and Entrepreneurship (2019), ‘One data scientist’s experience innovating at [company]’, The Jill Ker Conway Innovation and Entrepreneurship Center. Article is no longer available online.

Davies, H. (2015), ‘Ted Cruz campaign using firm that harvested data on millions of unwitting Facebook users’, The Guardian.

**URL:** <https://www.theguardian.com/us-news/2015/dec/11/senator-ted-cruz-president-campaign-facebook-user-data>

Donoho, D. (2017), ‘50 years of data science’, *Journal of Computational and Graphical Statistics* **26**(4).

**URL:** <https://amstat.tandfonline.com/doi/full/10.1080/10618600.2017.1384734>

Dwork, C., McSherry, F., Nissim, K. & Smith, A. (2006), Calibrating noise to sensitivity in private data analysis, in S. Halevi & T. Rabin, eds, ‘Theory of cryptography’, Springer, pp. 265–284.

**URL:** [https://link.springer.com/chapter/10.1007/11681878\\_14](https://link.springer.com/chapter/10.1007/11681878_14)

Edmondson, A. (1999), ‘Psychological safety and learning behavior in work teams’, *Administrative science quarterly* **44**(2), 350–383.

Elliott, A. C., Stokes, S. L. & Cao, J. (2018), ‘Teaching ethics in a statistics curriculum with a cross-cultural emphasis’, *The American Statistician* **72**(4), 359–367.

**URL:** <https://www.tandfonline.com/doi/abs/10.1080/00031305.2017.1307140>

Eubanks, V. (2018), *Automating inequality: How high-tech tools profile, police, and punish the poor*, St. Martin’s Press.

European Parliament (2018), Regulation on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (data protection directive), Technical report, European Union.

**URL:** <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>

Fiesler, C., Garrett, N. & Beard, N. (2020), What do we teach when we teach tech ethics? a syllabi analysis, in ‘Proceedings of the 51st ACM Technical Symposium on Computer Science Education’, pp. 289–295.

**URL:** <https://dl.acm.org/doi/10.1145/3328778.3366825>

Fitzpatrick, J. (2010), ‘If you’re not paying for it; you’re the product’.

**URL:** <https://lifesacker.com/if-youre-not-paying-for-it-youre-the-product-5697167>

Gershkoff, A., Therriault, A., Satyanarayan, A., Jones, B., Burg, B., Hurt, B., Granger, B., Jacob, B., Doig, C., Fryar, C., Ramanan, D., Bhargava, D., Perez, F., Greenleigh, I., Feng, J., Loyens, J., Morgan, J., Ram, K., Green, L., Barba, L., Colaco, M., Rocklin,

- M., Jamei, M., Horn, M., Harris, N. E., Elprin, N., Kaldero, N., Chopra, N., McGarry, P., Todkar, R., Journey, R., Brener, S., Couture, T., Thibodeaux, T. & McKinney, W. (2019), 'Data values and principles'.  
**URL:** <https://datapractices.org/manifesto/>
- Gotterbarn, D., Wolf, M. J., Flick, C. & Miller, K. (2018), 'Thinking professionally: The continual evolution of interest in computing ethics', *ACM Inroads* **9**(2), 10–12.  
**URL:** <https://dl.acm.org/citation.cfm?id=3204466>
- Grosz, B. J., Grant, D. G., Vredenburg, K., Behrends, J., Hu, L., Simmons, A. & Waldo, J. (2019), 'Embedded EthiCS: Integrating ethics across CS education', *Commun. ACM* **62**(8), 54–61.  
**URL:** <https://doi.org/10.1145/3330794>
- Gunaratna, N. S. & Tractenberg, R. E. (2016), Ethical reasoning with the 2016 revised ASA ethical guidelines for statistical practice, American Statistical Association.  
**URL:** [https://www.researchgate.net/publication/313309250\\_Ethical\\_Reasoning\\_with\\_the\\_2016\\_Revised\\_A](https://www.researchgate.net/publication/313309250_Ethical_Reasoning_with_the_2016_Revised_A)
- Hand, D. J. (2018), 'Aspects of data ethics in a changing world: where are we now?', *Big Data* **6**(3), 176–190.  
**URL:** <https://doi.org/10.1089/big.2018.0083>
- Hardin, J., Hoerl, R., Horton, N. J., Nolan, D., Baumer, B. S., Hall-Holt, O., Murrell, P., Peng, R., Roback, P., Temple Lang, D. & Ward, M. D. (2015), 'Data science in statistics curricula: Preparing students to 'think with data'', *The American Statistician* **69**(4), 343–353.  
**URL:** <https://doi.org/10.1080/00031305.2015.1077729>
- Heggeseth, B. (2019), 'Intertwining data ethics in intro stats', Symposium on Data Science and Statistics.  
**URL:** <https://drive.google.com/file/d/1GXzVMpb6GVNfWPS6bd9jggtqq1C77Wsc/view>
- Hicks, S. C. & Irizarry, R. A. (2018), 'A guide to teaching data science', *The American Statistician* **72**(4), 382–391.  
**URL:** <https://amstat.tandfonline.com/doi/abs/10.1080/00031305.2017.1356747>

- Huff, D. (1954), *How to lie with statistics*, WW Norton & Company, Inc.
- James, G., Witten, D., Hastie, T. & Tibshirani, R. (2013), *An Introduction to Statistical Learning: with Applications in R*, Springer.  
**URL:** <https://faculty.marshall.usc.edu/gareth-james/ISL/>
- Kaplan, D. (2018), ‘Teaching stats for data science’, *The American Statistician* **72**(1), 89–96.  
**URL:** <https://amstat.tandfonline.com/doi/full/10.1080/00031305.2017.1398107>
- Kerr, N. L. (1998), ‘Harking: Hypothesizing after the results are known’, *Personality and Social Psychology Review* **2**(3), 196–217.
- Kim, A. Y. & Escobedo-Land, A. (2015), ‘OKCupid data for introductory statistics and data science courses’, *Journal of Statistics Education* **23**(2).  
**URL:** <https://amstat.tandfonline.com/doi/abs/10.1080/10691898.2015.11889737>
- Kirkegaard, E. O. & Bjerrekær, J. D. (2016), ‘The OKCupid dataset: A very large public dataset of dating site users’, *Open Differential Psychology* **46**.
- Kolaczyk, E. D. & Csárdi, G. (2014), *Statistical analysis of network data with R*, Vol. 65, Springer.
- Kramer, A. D. I., Guillory, J. E. & Hancock, J. T. (2014), ‘Experimental evidence of massive-scale emotional contagion through social networks’, *Proceedings of the National Academy of Sciences* **111**(24), 8788–8790.  
**URL:** <https://www.pnas.org/content/111/24/8788>
- Langkjær-Bain, R. (2017), ‘Trials of a statistician’, *Significance* **14**(4), 14–19.  
**URL:** <https://rss.onlinelibrary.wiley.com/doi/full/10.1111/j.1740-9713.2017.01052.x>
- Levin, S. (2017), ‘Mark Zuckerberg: I regret ridiculing fears over Facebook’s effect on election’, *The Guardian*.  
**URL:** <https://www.theguardian.com/technology/2017/sep/27/mark-zuckerberg-facebook-2016-election-fake-news>

- Loukides, M., Mason, H. & Patil, D. (2018), *Ethics and data science*, Sebastopol, CA: O'Reilly Media.  
**URL:** <https://www.oreilly.com/library/view/ethics-and-data/9781492043898/>
- Meyer, R. (2014), 'Everything we know about Facebook's secret mood manipulation experiment', The Atlantic.  
**URL:** <https://www.theatlantic.com/technology/archive/2014/06/everything-we-know-about-facebooks-secret-mood-manipulation-experiment/373648/>
- Milner, Y. (2019), Data for Black Lives II, Data for Black Lives, MIT Media Lab, 75 Amherst St., Cambridge, MA 02139.  
**URL:** <http://d4bl.org/conference.html>
- National Academies of Sciences, Engineering, and Medicine (2018), *Data science for undergraduates: opportunities and options*, National Academies Press.  
**URL:** [http://sites.nationalacademies.org/cstb/currentprojects/cstb\\_175246](http://sites.nationalacademies.org/cstb/currentprojects/cstb_175246)
- National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1978), The Belmont report: Ethical principles and guidelines for the protection of human subjects of research, Technical Report 0012.  
**URL:** [https://videocast.nih.gov/pdf/ohrp-belmont\\_report.pdf](https://videocast.nih.gov/pdf/ohrp-belmont_report.pdf)
- Neff, G., Tanweer, A., Fiore-Gartland, B. & Osburn, L. (2017), 'Critique and contribute: A practice-based framework for improving critical data studies and data science', *Big Data* **5**(2), 85–97.  
**URL:** <https://www.liebertpub.com/doi/full/10.1089/big.2016.0050>
- Noble, S. U. (2018), *Algorithms of Oppression: How Search Engines Reinforce Racism*, NYU Press.  
**URL:** <http://www.jstor.org/stable/j.ctt1pwt9w5>
- O'Neil, C. (2016), *Weapons of math destruction: How big data increases inequality and threatens democracy*, New York: Crown.  
**URL:** <https://weaponsofmathdestructionbook.com/>



- Poulsen, K. (2014), ‘How a math genius hacked OKCupid to find true love’.  
**URL:** <https://www.wired.com/2014/01/how-to-hack-okcupid/>
- Saltz, J., Skirpan, M., Fiesler, C., Gorelick, M., Yeh, T., Heckman, R., Dewar, N. & Beard, N. (2019), ‘Integrating ethics within machine learning courses’, *ACM Transactions on Computing Education (TOCE)* **19**(4), 1–26.  
**URL:** <https://dl.acm.org/doi/10.1145/3341164>
- Shook, N. J. & Fazio, R. H. (2008), ‘Interracial roommate relationships: An experimental field test of the contact hypothesis’, *Psychological Science* **19**(7), 717–723.
- Skirpan, M., Beard, N., Bhaduri, S., Fiesler, C. & Yeh, T. (2018), Ethics education in context: A case study of novel ethics activities for the CS classroom, *in* ‘Proceedings of the 49th ACM Technical Symposium on Computer Science Education’, pp. 940–945.  
**URL:** <https://dl.acm.org/doi/10.1145/3159450.3159573>
- Sweeney, L. (2002), ‘k-anonymity: A model for protecting privacy’, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **10**(5), 557–570.  
**URL:** <https://www.worldscientific.com/doi/abs/10.1142/S0218488502001648>
- Tarran, B. (2019), ‘German commission calls for risk-based regulation of algorithmic systems’, *Significance* **16**(6), 4–5.  
**URL:** <https://doi.org/10.1111/j.1740-9713.2019.01329.x>
- Tractenberg, R. E. (2019a), ‘Strengthening the practice and profession of statistics and data science using ethical guidelines’.  
**URL:** <https://osf.io/preprints/socarxiv/93wuk>
- Tractenberg, R. E. (2019b), ‘Teaching and learning about ethical practice: The case analysis’.  
**URL:** <https://osf.io/preprints/socarxiv/58umw/download>
- Wang, M. Q., Yan, A. F. & Katz, R. V. (2018), ‘Researcher requests for inappropriate analysis and reporting: A US survey of consulting biostatisticians’, *Annals of Internal*

*Medicine* **169**(8), 554–558.

**URL:** <https://doi.org/10.7326/M18-1230>

Wasserstein, R. L., Lazar, N. A. et al. (2016), ‘The ASA’s statement on p-values: context, process, and purpose’, *The American Statistician* **70**(2), 129–133.

**URL:** <https://doi.org/10.1080/00031305.2016.1154108>

Wasserstein, R. L., Schirm, A. L. & Lazar, N. A. (2019), ‘Moving to a world beyond “ $p < 0.05$ ”’, *The American Statistician* **73**(sup1), 1–19.

**URL:** <https://doi.org/10.1080/00031305.2019.1583913>

Wender, B. & Kloefkorn, T. (2017), ‘Roundtable on data science postsecondary education, Meeting #5 highlights’, The National Academies of Sciences, Engineering, and Medicine.

**URL:** <https://www.nationalacademies.org/event/12-08-2017/docs/D8EE65EFC7F4B0C368D267EDAL>

Zimmer, M. (2010), ‘“but the data is already public”: on the ethics of research in Facebook’, *Ethics and information technology* **12**(4), 313–325.

**URL:** <https://link.springer.com/article/10.1007/s10676-010-9227-5>

Zuboff, S. (2018), *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, 1st edn.