

Readings in Ethical Data Science

An Incomplete Annotated Reading List

Introduction

This reading list gives an overview of the ethical concerns specific to data analysis, data science, and artificial intelligence. Ethics is used broadly here to mean concerns related to racial and economic equity, justice, fairness, and the protection of democratic and human rights.

This list is intended to spark new ideas and prompt critical thinking about data system design and integration into business processes in an organization. This is not an endorsement of all viewpoints represented in the readings below – except to say that each of the readings raise questions, put forward ideas, and make critiques that are worthy of your deep consideration.

All links last accessed July 11th, 2019. This guide was last updated July 11th, 2019. An unannotated version of this reading list is available [on GitHub](#). You can get an updated version of this list as well as suggest additions to the reading list there.

Books

Eubanks, Virginia. 2018. [Automating Inequality](#). St. Martin's Press.

Noble, Safiya. 2018. [Algorithms of Oppression: How Search Engines Reinforce Racism](#). NYU Press.

O'Neil, Cathy. 2016. [Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy](#). Broadway Books.

These are the “big three” books uncovering the ways that algorithms affect our lives invisibly and sometimes visibly. Of the three I am partial to Eubanks because of the in-depth way she centers the voices of those affected by algorithms and her focus on algorithms in the social services sector. Noble is one of the most important voices in technology today – especially for thinking about the impact of major technology companies on our lives. I prefer the article below to the book length treatment by O'Neil. The intersection of expertise and democracy has been studied by social scientists for decades and that literature is better summarized elsewhere.

Broussard, Meredith. 2018. [Artificial Unintelligence: How Computers Misunderstand the World](#). MIT Press.

Benjamin, Ruha. 2019. [Race After Technology: Abolitionist Tools for the New Jim Code](#). Polity.

These are two newer texts which offer updated takes on the above themes. Broussard is specifically tackling artificial intelligence which is related to, but adjacent to most applications of data science in fields like education and social services (for now). Benjamin's book is one I anticipate greatly as it brings a much needed critical race perspective to the conversation about the affect of data analysis and collection on our society. Even better it is intended to teach the reader how to critically review the promises of technologies like algorithms and automated decision support systems.

Loukides, Mike, Hilary Mason, and DJ Patil. 2018. Ethics and Data Science. O'Reilly.

This is a pragmatic and brief overview of the major ethical concerns with data science. This text focuses on practical steps that a data science team can take to be more ethical. This practical approach is different than the above readings which is why I recommend it as a supplementary reading – but is very helpful for answering the “what do I do now?” question.

brown, adrienne maree. 2017. Emergent Strategy: Shaping Change, Changing Worlds. AK Press.

This book isn't about ethics, data science, or technology explicitly at all. It is about how to work together with a large inclusive set of stakeholders to build something that reflects the voices of a diverse community. This is, in fact, the main solution proposed by almost all the authors above – inclusive design done together with a wider community. This book will stimulate your thinking on how to go about that.

Articles

Wallach, Hanna. 2014. “Big Data, Machine Learning, and the Social Sciences: Fairness, Accountability, and Transparency.” Medium. [Online](#). 12.19.2014

O'Neil, Cathy. 2016. “How to Bring Better Ethics to Data Science.” Slate. [Online](#). 2.4.2016

Broussard, Meredith. 2019. “Letting Go of Technochauvinism.” in Public Books. [Online](#). 6.17.2019.

Together these three articles provide a great overview of the limits of data science, the limits of our ability to “technology” our way out of social problems, and the intersection of the data systems we design and the world we live in. The Broussard article challenges the reader with the proposition that automation is perhaps not the best answer to each and every problem. The O'Neil article is a great overview of her critically acclaimed book – clearly presenting the arguments and the implications. The Wallach piece is maybe the best of the bunch – it gives a comprehensive tour of the ethical concerns of data science starting from the questions we ask all the way through how we use the answers our algorithms provide.

Dash, Anil. 2018. Humane Tech. Medium. [Online](#).

Dash is one of the most important voices in the tech industry. While not strictly about data science, this series of articles provides a great overview for thinking about how to build technology tools for society as it is in ways that make society better – instead of exploiting the flaws in our society for profit. You should also check out his podcast – [Function](#).

Fischer, Frank. 1993. “Citizen participation and democratization of policy expertise: From theoretical inquiry to practical cases.” Policy Sciences. v. 26 pp. 165-187.

Diakopoulos, Nicholas. 2016. “How to Hold Governments Accountable for the Algorithms They Use.” Slate. [Online](#). 2.11.2016

Angwin, Julia. 2016. “Making Algorithms Accountable.” ProPublica. [Online](#). 2.1.2016

Government uses of data science tools are a special case and merit their own discussion. Dakopoulos and Angwin both present good overviews for how to make algorithms in government accountable and how to enforce accountability of algorithms in general. For me, though, the take on this topic that expanded my mind the most was an older article on the role of expertise in governing a democratic society by Fischer. This article is heavy on the academic side but takes a look at the unique challenges that a reliance on expertise poses to a democratically governed society.

Patil, DJ. 2016. “A Code of Ethics for Data Science.” Medium. [Online](#). 2.1.2018

Wheeler, Shaun. 2018. “An ethical code can’t be about ethics.” Towards Data Science. [Online](#). 2.6.2018

Eubanks, Virginia. 2018. “A Hippocratic Oath for Data Science.” [Online](#). 2.21.2018

There has been a healthy debate about a “Hippocratic Oath” for Data Science or a “Data Science Code of Ethics”. These articles provide different viewpoints on that debate and help think about what it means to ethically do data science and what role a professional code of ethics may play.

Further Reading Lists

Venkatasubramanian, Suresh and Katie Shelef. 2017. “Ethics of Data Science Course Syllabus.” University of Utah. [Online](#).

This syllabus contains a lot of foundational texts in the ethics of social science as well as a wonderful set of examples of the ethical challenges posed by data science.

Malliaraki, Eirini. 2018. "Toward ethical, transparent and fair AI/ML: a critical reading list." Medium. [Online](#).

This is the closest thing to a comprehensive current reading list on transparency and fairness in machine learning and artificial intelligence. This thorough and well-organized reading list has plenty of great further reading to extend on any of the readings covered here.

Wickham, Hadley. 2018. "Readings in Applied Data Science." [Online](#).

A wide-ranging reading list of applied data science topics. Some would make great case studies for ethical dilemmas in data science, others are critical analyses of the ethics of particular applications of data science.

Various. 2018. Readings in Data Ethics. O'Reilly. [Online](#).

Five short articles that will give you a practical and pragmatic overview for how to implement some ethical safeguards into your data science team and products.