4. Data Science Ethics

The Need for Ethics

In most empirical disciplines, **ethics** are introduced early in the educational process and end up playing a crucial role in researchers' activities.

Data scientists who come to the field by way of mathematics, statistics, computer science, economics, or engineering, however, are less likely to have encountered ethical research boards or **formal ethics training**.

Discussions on ethical matters are often tabled in favour of pressing technical or administrative considerations when faced with hard deadlines.

But the current deadline is replaced by another deadline, and then by another one, with the end result being that the conversation may never take place.

The Need for Ethics

When large-scale data collection first became possible, there was a ‘Wild West’ mentality to its use: **everything was allowed as long as it was feasible**.

Modern data science has **professional codes of conduct**

- outlining **responsible** ways to practice data science
- legitimate rather than fraudulent, ethical rather than unethical.

This shifts **added responsibility** to data scientists, but provides **protection** from clients/employers who want them to carry analysis in questionable ways.

The Need for Ethics

Recent focus on data ethics does not seem to have slowed breaches:

- Volkswagen
- Whole Foods Markets
- General Motors
- Cambridge Analytica
- Amazon
- Ashley Madison

What is/are Ethics?

Ethics refers to the study and definition of **right** and **wrong** conduct:

- in general
- applied in specific circumstances

Ethics is not (necessarily) the same as:

- social convention
- religious beliefs
- laws

What is/are Ethics?

In the West, ethical theories are used to frame debates around ethical issues:

- **Golden rule:** do unto others as you would have them do unto you
- **Consequentialism:** the end justifies the means
- **Utilitarianism:** act in order to maximize positive effect
- **Moral Rights:** act to maintain and protect the fundamental rights and privileges of the people affected by actions
- **Justice:** distribute benefits and harm among stakeholders in a fair, equitable, impartial way

What is/are Ethics?

But humans subscribe to a wide variety of ethical codes/cultures, including:

- Confucianism
- Taoism
- Buddhism
- Shinto
- Ubuntu
- Te Ara Tika (Maori)
- etc.

It is easy to imagine contexts in which any of these would be better-suited to the task at hand –remember to **inquire** and to **heed the answers**.

Ethics and Data Science

How might these ethical theories apply to data analysis?

- who, if anyone, owns data?
- are there limits to how data can be used?
- are there value-biases built into certain analytics?
- are there categories that should never be used in analyzing personal data?
- should data be publicly available to all researchers?

The answers depend on a number of factors. To give you an idea of some of the complexities, let us the first question: *who, if anyone, owns data?*

Ethics and Data Science

Is it the **data analysts** who transform the data's potential into usable insights?

Is it the **data collectors** who have a copy and make the work possible?

Is it the **sponsors** or **employers** who made the process economically viable?

In some instances, the **law** may chime in as well. Anybody else?

This simple question is not easily answered; it's on a case-by-case basis.

Hidden truth: **there is more to data analysis than *just* data analysis**.

Ethics and Data Science

Similar challenge for **open data** (“pro” vs. “anti” both have strong arguments).

General principle of data analysis: eschew the **anecdotal** for the **general**. **Sound**, as focus on specific observations can obscure the full picture.

But data points are **not** just marks on paper or bytes on the cloud. Decisions made on the basis of data science may **affect living beings in negative ways**. It cannot be ignored that outlying individuals and minority groups often suffer disproportionately at the hands of so-called evidence-based decisions.

First Nations Principles of **OCAP** (Ownership, Control, Access, Possession).

Best Practices

"Do No Harm": data collected from an individual **should not be used to harm** the individual.

Informed Consent:

- Individuals must **agree to the collection and use** of their data
- Individuals must have a **real understanding of what they are consenting to**, and of **possible consequences** for them and others

Respect "Privacy": excessively hard to maintain in the age of constant trawling of the Internet for personal data.

Best Practices

Keep Data Public: data should be kept **public** (all? most? any?).

Opt-In/Opt-Out: Informed consent requires the ability to **opt out**.

Anonymize Data: removal of id fields from data prior to analysis.

“Let the Data Speak”:

- no cherry picking
- importance of validation (more on this later)
- correlation and causation (more on this later, too)
- repeatability

The Good, the Bad, and the Ugly

Data projects could whimsically be classified as **good**, **bad** or **ugly**, either from a technical or from an ethical standpoint (or both).

- **good** projects increase knowledge, can help uncover hidden links, etc., as harmlessly as possible
- **bad** projects can lead to bad decisions, which can in turn decrease the public's confidence and potentially harm some individuals
- **ugly** projects are, flat out, unsavoury applications; they are poorly executed from a technical perspective, or put a lot of people at risk; these (and similar approaches/studies) should be avoided

The Good, the Bad, and the Ugly

Good:

- P. A. B. Bien Nicholas AND Rajpurkar, "Deep-learning-assisted diagnosis for knee magnetic resonance imaging: Development and retrospective validation of MRNet," *PLOS Medicine*, vol. 15, no. 11, pp. 1–19, 2018, doi: 10.1371/journal.pmed.1002699.
- BeauHD, "Google AI claims 99 percent accuracy in metastatic breast cancer detection," *Slashdot.com*, Oct. 2018.
- Columbia University Irving Medical Center, "Data scientists find connections between birth month and health," *Newswire.com*, Jun. 2015.

The Good, the Bad, and the Ugly

Bad:

- Indiana University, "Scientists use Instagram data to forecast top models at New York Fashion Week," *Science Daily*, Sep. 2015.
- D. Wakabayashi, "Firm led by Google veterans uses A.I. to 'nudge' workers toward happiness," *New York Times*, Dec. 2018.
- N. Cohn, "How one 19-year-old illinois man is distorting national polling averages," *The Upshot*, 2016.

The Good, the Bad, and the Ugly

Ugly:

- J. Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," *Reuters*, Oct. 2018.
- I. Johnston, "AI robots learning racism, sexism and other prejudices from humans, study finds," *The Independent*, Apr. 2017.
- M. Judge, "Facial-recognition technology affects African-Americans more often," *The Root*, 2016.
- M. Kosinski and Y. Wang, "Deep neural networks are more accurate than humans at detecting sexual orientation from facial images," *Journal of Personality and Social Psychology*, vol. 114, no. 2, pp. 246–257, Feb. 2018.

Suggested Reading

Data Science Ethics

Data Understanding, Data Analysis, Data Science
Data Science Basics

Ethics in the Data Science Context

- The Need for Ethics
- What Is/Are Ethics?
- Ethics and Data Science
- Guiding Principles

Exercises

Data Science Ethics

1. Research the recent data ethics scandals involving Volkswagen, Amazon, Whole Foods Markets, Cambridge Analytica, Ashley Madison, General Motors, or any other organization. What transpired? Who was affected? What were the consequences to the general public, the organization, the data community? How could it have been avoided?

2. Establish a statement of ethics for your data work. Are there areas that you are unwilling to work on?