Ethics in Data Science (Part 1, Overview)

Daniel Lawson — University of Bristol

Lecture 12.1.1 (v1.0.2)

Signposting

- We've covered **how to do data science** in this course.
- This final lecture is about how to do it **responsibly**, i.e.
  - What laws govern data science?
  - What is Ethical data science?
  - What is privacy?
  - How can we design data science tools to protect it?
  - What is Statistical discosure?
  - What is algorithmic bias? How can we mitigate it?
- Thanks to Raul Santos-Rodriguez and Peter Flach for access to their slides on these topics.
- This lecture is split into three parts:
  - This is part 1, on ethics and the law,
  - Part 2 covers Privacy and disclosure,
  - Part 3 covers Fairness and interpretability.

ILOs

- ▶ ILO5 Be able to reason about and conceptually align problems involving real data to appropriate theoretical methods and available methodology to correctly make inferences and decisions

Ethical data science

- Ethical behaviour puts benefits to group or society above benefits to the individual or organisation.
  - Typically, ethical outcomes are in individuals and organisations long term interest.
  - Ethics is not **law**. But it is good when the two coincide.
  - Ethics asks what we **should** do, rather than what we **can** do.
- Data Science can be done for external ends, ethically.
  - It can also be used to improve the world by making ethical outcomes a goal.
- Not all data science is ethical. Yours can be.

Ethics and Privacy: the law

- ▶ The laws that govern data privacy in the UK (and similarly in the EU) include:
  - ▶ Human Rights Act (HRA 1998)
  - ▶ EU General Data Protection Regulation (GDPR 2018)
  - ▶ Data Protection Act 2018 (DPA 2018)

Human Rights Act

- The **HRA (Article 8)** provides the **right to privacy**. Specifically "respect for your family and private life, your home and your correspondence" including **how your personal information is held and protected**.
- The HRA sets out intentions. This right is **implemented** in various other forms of legislation.

EU GDPR / UK Data Protection Act

- The **GDPR** is an important and pervasive law implementing **data protection** rules.
- You will learn it in detail if you are employed by a company that does business in the EU.
- The **DPA (2018)** implements GDPR *in the UK* and supersedes the DPA (1998). It goes slightly further than GDPR but not in any ways that directly affect data science. These include:
  - National security,
  - Immigration,
  - Implementation.

GDPR and Data Science

- GDPR affects data science practice in three areas.
    1. GDPR imposes **limits** on data processing and consumer profiling.
    2. For organizations that use automated decision-making, GDPR creates a **"right to an explanation"** for consumers.
    3. GDPR holds firms **accountable** for bias and discrimination in automated decisions.
- The good news is that as responsible data scientists we wanted to do these things all along.
- KDNuggets has a very good description of how GDPR affects Data Science.

GDPR implications

- Some key implications include:
  - **Informed consent** for data usage is required except for "ordinary conduct of business".
  - There are **constraints on data processing and profiling**, which apply **only to identifying data** individual consumer.
  - Therefore **robust anonymisation** is vital for data mining activities.
  - Pre-anonymised data must be carefully guarded, and there are legal implications for employees and businesses that do not follow best practice.

Ethical topics in data science:

- Data **ownership**:
  - Who owns your data?
- Statistical **disclosure** attacks:
  - How can data be extracted from anonymised data?
- **Interpretable** data science:
  - How do we know why an algorithm has made a particular choice?
- Algorithmic **Fairness**:
  - How can we ensure that automated decisions are not biased against particular groups of people?

Data ownership

- Who owns data? Under the GDPR, there is a **data subject** and a **data controller**. The subject has several rights, including:
  - To **object** to specific usage, and to give **clear consent** where use is acceptable,
  - To have **easy access** to their data,
  - To **rectification** of errors and for most information on them to be forgotten,
  - To data **portability**, so that their data can be moved between providers.
- The controller has the responsibility to protect the subject's rights.

Implications of data ownership

- A Company that gathers data must provide provisions for access, rectification, and portability.
  - Auditable logs of data use are required.
  - If you are the controller of data, you will need additional training.
- For data science use:
- Data used for data science must be subject to rectification and removal!
  - This is **not** retrospective, so any **results** from this data used historically, or being used in contemporary analysis, do not need to be deleted.
  - However, you cannot store copies of personal data arbitrarily; these are subject to the right to be forgotten.
- Careful processing pipelines are therefore needed to retain critical results whilst not storing individual data except temporarily.
  - Anonymised and/or aggregated data can be retained. Provided that it truly cannot be linked to an individual. . .

Reflection

- What do the acronyms HRA, GDPR and DPA mean?
- How do they relate to you as a data scientist?
  - What specific implications can you think of?
- How does GDPR change your practice?
- By the end of the course you should:
  - Be aware of the **basic regulatory framework** of privacy in the UK,
  - Be able to reason using the **general human rights** principles underlying this.

Signposting

- Still to come:
  - 12.1.2 Protecting data from statistical disclosure,
  - 12.1.3 Ensuring algorithms are fair and interpretable.
- References: Make sure you know broadly what the legislation does!
  - Human Rights Act (HRA 1998)
  - EU General Data Protection Regulation (GDPR 2018)
  - Data Protection Act 2018 (DPA 2018)