Institute and Faculty of Actuaries

ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS

Additional information and resources: ethical data science

Contents

Introduction 1

1. Seek to enhance the value of data science for society 2

2. Avoid harm 3

3. Apply and maintain professional competence 4

4. Seek to preserve or increase trustworthiness 5

5. Maintain accountability and oversight 6

Contact details 7

DISCLAIMER

The views linked to in this publication are those of external parties and not necessarily those of the Royal Statistical Society ("the RSS"). The RSS does not endorse any of the views stated, nor any claims or representations made in this publication or external links and accepts no responsibility or liability to any person for loss or damage suffered as a consequence of their placing reliance upon any view, claim or representation made in this publication or in external links

Introduction

The Royal Statistical Society ('the RSS') and Institute and Faculty of Actuaries ('the IFoA') guide for ethical data science considers five reoccurring themes from existing ethical frameworks and working practices across a wide range of sectors and industries.

This reference document outlines the existing ethical frameworks that were considered when developing the guide and provides some additional resources to further explore and implement the ethical themes.

As this is an evolving area, the resources may not reflect most recent developments but could provide a starting point for further exploration. This is a snapshot of links at the time of developing the guidance, and we recognise that links will change and practice will move on. Nevertheless, we hope this will serve as a useful guide.

Existing ethical frameworks considered

- AI Now Institute – **AI Now 2017 Report**, **2018 Report**
- Association for Computing Machinery (ACM) **Code of Ethics and Professional Conduct**
- DataKind – **Ethical Principles for Pro Bono Data Scientists**
- **Future of Life Institute** – Asilomar AI Principles
- Institute of Electrical and Electronic Engineers (IEEE) – **Ethics in Action Standards Projects**
- MetroLab Network – **Ethical Guidance for Applying Predictive Tools**
- Machine Intelligence Garage **Ethics Framework**
- Microsoft Ethics and Law in Data and Analytics course materials
- Nuffield Foundation **Ethical and Societal Implications of Data and AI report**
- National Statisticians Data Ethics Advisory Committee **Ethical Principles**
- Open Data Institute **Data Ethics Canvas** tool
- UK Department for Digital, Culture, Media & Sport – **Data Ethics Framework**

1. Seek to enhance the value of data science for society

Being aware of the potential biases, errors, assumptions and risks inherent in predictive modelling

There can be many sources of bias to consider within statistical processes, for example Inside Big Data wrote a simple introduction to **machine learning bias**. There are examples of where automated algorithms have led to poor outcomes for certain individuals, or groups, for example algorithms that have been used to decide who will have **access to healthcare**. Pro Publica report on the bias in the **COMPAS Recidivism Algorithm**.

Fairness in data science

As with decisions that are made by humans, algorithms can be subject to bias and errors. This **tutorial on fairness in machine learning** explains some of the problems and solutions.

Some examples of data science being used for social good can be found in **The Big Data Social Justice Foundation, DataKind** and the **Data Science for Social Good Fellowship**.

A key organisation doing work in this area is the **Ada Lovelace Institute** which is an independent research and deliberative body with a mission to ensure data and AI work for people and society. The Institute was established in 2018 by the Nuffield Foundation, and the Royal Statistical Society is one of its partner organisations.

Acting in the public interest

The IFoA's policy briefing on **Data Science in Insurance** discusses public interest concerns within insurance.

The UK Government's **Data Ethics Workbook** also includes discussion on serving the public good and the **UN's Sustainable Development Goals** provides a useful framework.

2. Avoid harm

Using ethically sourced data

The Statistics Authority **Code of Practice for Statistics** provides a framework for considering data sources.

Considerations for using Big Data

Using Big Data is both technically and ethically challenging compared with more traditional sources such as statistical and operational data. The **RSS report on Big Data** and ethics helps to outline some of the key considerations. The Financial Conduct Authority (FCA) have also released guidelines on the use of **big data in general insurance**.

Data linkage

Linkage is the process of joining two data sets together, usually to combine different sources of information about individuals and provide a richer data set. There may be additional or increased risks to consider when linking data, including errors, privacy and bias. For a discussion on data linkage see The Office for Statistics Regulation review: **Joining up Data for Better Statistics**.

Using anonymisation

Anonymisation is the practice of protecting confidentiality in data, so that it can be legally and ethically shared. Although it has implications for privacy, it is primarily concerned with reducing the risk of individuals being identified. This can be challenging in data science, especially when using big data or linking data from different sources. If data is in the public domain, it can be possible to identify individuals even if their personal information is not contained in the data. It can be good practice to use a combination of anonymisation and data security to protect the confidentiality (and privacy) of individuals. The **UK Anonymisation Network** has several resources to help practitioners consider anonymisation including **An Anonymisation Decision Framework**.

Having ethical governance

Governance refers to organisational policies which describe how a business will operate and the procedures for monitoring their implementation. Embedding ethical governance can be achieved by having ethical assessment 'check points' to assess whether the project is meeting the ethical standards of the organisation. Suggestions for ethical checks are shown in the implementation checklist in the full guide.

Ethical assessment could range from simple self-assessment checklists to seeking approval from an ethical committee. Governance may also be linked to mandated governance within the industry or profession within which the data scientist works.

For example, the **National Statisticians Data Ethics Advisory Committee** is an independent committee that provides advice on the ethical use of data within government. Their **Ethics Self-Assessment process** is an example of combining informal and formal governance. The Open Data Institute have a **Data Ethics Canvas** to help organisations ask the right questions about their projects.

3. Applying and maintaining professional competence

Professional body codes

The RSS **Code of Conduct** represents best practice in doing statistical work which equally applies to data science projects.

The IFoA **Actuaries' Code** applies to all IFoA members.

There may be other professional codes in different industries. For example the **Code of Practice for Statistics** applies to anyone producing official statistics, and represents a gold standard for producing statistics for decision makers.

Legal and regulatory requirements

There are various legal and regulatory requirements that will be applicable to data science work, such as the European Union General Data Protection Regulation (GDPR). The UK Government provides a guide on how to apply **GDPR** that may be useful.

Open Source tools for data cleaning, exploration and visualisation

It is important to fully explore and clean data before starting any modelling. This can be more challenging when working with unstructured and very large data sets. Some libraries in Python which are useful include **Pandas Profiling** library and the **Missingno** library. There is also **JanitoR** and **pyJanitor** for semi-automatic data cleaning.

Open source tools that may help explain machine learning models

Some tools are becoming available to help visualise data and models that may be useful to help in explaining them to others.

For example, for text analysis practitioners may use the **lime** and **eli5** python packages.

For numerical / other features, **DALEX** (Descriptive mAchine Learning Explanations) in R or **PyDALEX** in Python can be useful.

On the python / scikitlearn side there is also **Yellowbrick** which enables visualisation and steering of python **scikit** machine learning models.

4. Seek to preserve or increase trustworthiness

Publishing and sharing work

One platform being used to share data science work is **Github**, which allows the sharing of code, libraries for open source software, data and associated documentation. There are privacy options which allow users to control who can access the information. This **article** discusses how Github can be used.

The **Open Data Institute** also provides information about sharing data.

5. Maintain accountability and oversight

Programmes of research

Several organisations have programmes of research to address the practical and ethical questions that AI poses.

A key organisation doing work in this area is the **Leverhulme Centre for the Future of Intelligence** (founded by Professor Stephen Hawking) who are running several **key programmes** to progress this area:

- AI: Futures and Responsibility (Projects "The Value Alignment Problem" and "Autonomous Weapons – Prospects for Regulation" seem particularly relevant to this guiding principle)
- AI: Trust and Society
- Kinds of Intelligence
- AI: Narratives and Justice
- Philosophy and the Ethics of AI

These Centre programmes are still ongoing since research in this area is in its infancy, but some **outputs** are becoming available and they are active on **Twitter**.

Other organisations working on this include:

- **Machine Intelligence – Research Institute**
- **Centre for the Study of Existential Risk**

Emerging views

DataKind previously published the results of a survey asking for **members views on the responsible use of AI.**

The AI Now Institute have several **reports** around the challenges of AI.

The Guardian writes about AI in the article "**Can we stop AI outsmarting humanity?**".

Institute and Faculty of Actuaries

ROYAL STATISTICAL SOCIETY

DATA | EVIDENCE | DECISIONS

Beijing
14F China World Office 1 · 1 Jianwai Avenue
Beijing · China 100004
Tel: +86 (10) 6535 0248

Edinburgh
Level 2 · Exchange Crescent · 7 Conference Square
Edinburgh · EH3 8RA
Tel: +44 (0) 131 240 1300

Hong Kong
1803 Tower One · Lippo Centre · 89 Queensway
Hong Kong
Tel: +852 2147 9418

London (registered office)
7th Floor · Holborn Gate · 326-330 High Holborn
London · WC1V 7PP
Tel: +44 (0) 20 7632 2100

Oxford
1st Floor · Park Central · 40/41 Park End Street
Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200

Singapore
163 Tras Street · #07-05 Lian Huat Building
Singapore 079024
Tel: +65 6906 0889

www.actuaries.org.uk

Royal Statistical Society
12 Errol Street
London · EC1Y 8LX
Tel: +44 (0) 20 7638 8998

www.rss.org.uk

© 2019 Institute and Faculty of Actuaries and Royal Statistical Society