




Institute
and Faculty
of Actuaries

ROYAL
STATISTICAL
SOCIETY
DATA | EVIDENCE | DECISIONS

A Guide for Ethical Data Science

A collaboration between the Royal Statistical Society
(RSS) and the Institute and Faculty of Actuaries (IFoA)



Contents

- Introduction 1
 - Background
 - Relevant themes
- 1. Seek to enhance the value of data science for society 4
- 2. Avoid harm 5
- 3. Apply and maintain professional competence 7
- 4. Seek to preserve or increase trustworthiness 9
- 5. Maintain accountability and oversight 10
- Implementation checklist 11
- Contact details 12

DISCLAIMER

This guidance imposes no new obligations upon members. Rather the Institute and Faculty of Actuaries ('the IFoA') and the Royal Statistical Society ('the RSS') hopes that the guidance will be a useful tool for members. It does not constitute legal advice, nor does it necessarily provide a defence to allegations of misconduct. While care has been taken to ensure that it is accurate, up-to-date and useful, the IFoA and RSS will not accept any legal liability in relation to its contents.



1. Introduction

- 1.1 This guide has been developed jointly by the **Royal Statistical Society (RSS) Data Science Section** and the **Institute and Faculty of Actuaries (IFoA)** for members working in the area of data science. It is intended to complement existing ethical and professional guidance and is aimed at addressing the ethical and professional challenges of working in a data science setting.
- 1.2 While there is no single definition of data science, it can be broadly thought of as scientific, computational and analytical methods used to process and extract information from data. It is synonymous with 'big data'¹, machine learning (artificial intelligence), and data pipelines (automated systems that capture and process data). Data science brings together several fields including maths, statistics and computer science.
- 1.3 The ethical significance of data science, and the implications for industries and the wider public, is constantly evolving. As data science methods become more common within statistical and actuarial fields, there are both opportunities and challenges for individuals working in data science ('practitioners').
- 1.4 For example, managing privacy, fairness and bias can be difficult and complex when using algorithmic methods. Additionally, public perceptions are still developing around many aspects of data science, including the use of artificial intelligence (AI) in systems and decision making, and 'big data' sources about people, such as social media and mobile phone data. The importance of keeping up with these challenging ethical issues led to the partnership between the **RSS** and the **IFoA** on the practical and ethical implications of data science.
- 1.5 This partnership draws on both organisations' commitment to ensuring high professional standards and upholding the public interest.

2. Background

- 2.1 The RSS and IFoA developed this guide using a bottom-up approach, starting in 2018 with four workshops in which data science professionals were consulted on four key questions affecting the data science industry. Those questions were:
 - What does a good data science workflow look like?
 - How should data science fit into the structure of an organisation?
 - What do executives and managers need to know about data science?
 - What is a data scientist's responsibility to wider society?
- 2.2 The key **findings** from the workshops included:
 - best practice for data science is dependent on the industry, organisational design, historical analytical workflows and the availability of skills within data science teams

1 | 'Big data' is commonly defined as having one or more of the 'four Vs': Variety (includes unstructured formats eg images, videos, audio files, web pages, emails, and documents), Velocity (generated/updated at speed), Volume (cannot be held/analysed with traditional databases/statistical tools), Veracity (high levels of inconsistencies and uncertainty). It is also often used to refer to data about people's behaviour, such as their digital footprint, social media use or mobile data.

- the data science professionals broadly agreed several high-level principles and practices around workflows, ethics, what executives need to know, and how to effectively place data science within the structure of an organisation

- 2.3 While agreement for good practice was found across many aspects, some of the more complex issues will require further thought and input from the professions.
- 2.4 This guide considers five reoccurring ethical themes from existing ethical frameworks² relating to data science and AI, with sources including government and industry, as well as bringing in discussion points arising from the joint workshops.
- 2.5 This guidance is non-mandatory and does not impose any obligations upon RSS or IFoA members. Members must comply with any relevant and applicable professional body codes, standards of practice and legal or regulatory obligations, which take precedence.³
- 2.6 The guide will be reviewed regularly. If members have any queries or feedback, please contact **dataethics@actuaries.org.uk**

3. Relevant themes and principles within data ethics

- 3.1 There are five relevant ethical themes associated with data science. Within each of these themes are examples of corresponding working practices which aim to help members consider data ethics. This is a condensed outline and is not intended to be a comprehensive list of ethical principles.
- 3.2 The five themes for consideration are:

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

As the impact that data science can have on society could be significant, an important ethical consideration is what the potential implications could be on society as a whole. A common theme within ethical frameworks discussing data science and AI is for practitioners to attempt to seek outcomes within their work which support the improvement of public wellbeing. This could involve practitioners seeking to share the benefits of data science and balancing this with the wellbeing of potentially affected individuals.

2. Avoid harm

Data science has the potential to cause harm and this ethical consideration therefore focuses on how practitioners can avoid this by working in a manner that respects the privacy, equality and autonomy of individuals and groups, and speaking up about potential harm or ethical violations.

Practitioners may be subject to legal and regulatory obligations in relation to the privacy of individuals, relevant to the jurisdiction in which they are working, as well as regulatory obligations to speak up about harm or violations of legal requirements.

This can also be applied to work relating to businesses, animals or the environment, with consideration of commercial rights, animal welfare and the protection of environmental resources.

2 | These are available in the additional references paper at www.rss.org.uk

3 | Members of the Institute and Faculty of Actuaries carrying out technical actuarial work within the UK geographic scope are subject to the Financial Reporting Council Technical Actuarial Standards and must ensure that they comply with them.

3. Apply and maintain professional competence

This ethical principle expects data science practitioners to apply best practice and comply with all relevant legal and regulatory requirements, as well as applicable professional body codes.

Professional competence involves fully understanding the sources of error and bias in data, using 'clean' data (eg edited for missing, inconsistent or erroneous values), and supporting work with robust statistical and algorithmic methods that are appropriate to the question being asked. Practitioners can also thoroughly assess and balance the benefits of the work versus the risks posed by it, and keep models under regular review.

4. Seek to preserve or increase trustworthiness

The public's trust and confidence in the work of data scientists can be affected by the way ethical principles are applied. Practitioners can help to increase the trustworthiness of their work by considering ethical principles throughout all stages of a project.

This is another reoccurring theme that encourages practitioners to be transparent and honest when communicating about the way data is used. Transparency can include fully explaining how algorithms are being used, if and why any decisions have been delegated, and being open about the risks and biases.

Engaging widely with a diverse range of stakeholders and considering public perceptions both from the outset, and throughout projects, can help to build trustworthiness and ensure all potential biases are understood.

5. Maintain accountability and oversight

Another key issue in data ethics around automation and AI is the question of how practitioners maintain human accountability and oversight within their work.

Being accountable can include being mindful of how and when to delegate any decision making to systems, and having governance in place to ensure systems deliver the intended objectives. When deciding to delegate any decision making, it would be useful to fully understand and explain the potential implications of doing so, as the work could lead to introducing advanced AI systems which do not have adequate governance. Practitioners should note that delegating any decisions to these systems does not remove any of their individual responsibilities.

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

Data science has the potential to be both beneficial and detrimental to individuals and/or the wider public. To help minimise any adverse effects, members can seek to understand the potential impact of their work and consider any opportunities that may deliver benefits for the public. It is recognised that not all work will have a defined societal benefit, but members could strive to seek fairness or an overall increase to well being, within commercial applications.

‘Public or societal good’ has many aspects and could include:

- economic empowerment
- employment
- education
- environment
- equality and inclusion
- the concept of fairness
- health and hunger
- information certification and validation
- infrastructure
- public and social sector
- security and justice
- cyber risk to national security.

Practice examples

Below is a table of examples of how this principle could be put into practice:

Practices	
What	How
Considering the potential impact that models have on decisions (especially in relation to people)	<ul style="list-style-type: none"> • Understanding how the models will be used, what impact there may be and for whom • Understanding who the relevant stakeholders are • Understanding the potential biases, errors, assumptions and risks inherent in predictive modelling • Inviting peer and bias review
Seeking to act in the public interest	<ul style="list-style-type: none"> • Choosing or seeking work that is considered to have a benefit to the public/society • Understanding how the benefits may be distributed across society • Considering who or what could be affected by outcomes • Thinking about what could mitigate any risks for unfairness • Leaders encouraging a culture which values social justice and fairness



2. Avoid harm

Data science can involve using potentially sensitive data, such as an individual's personal information, or have underlying sensitivity due to the volume of data and its ability to link with other sources. Additionally, newer sources of data, such as social media, have limited capacity for informed consent.

Examples of potential harm could include:

- financial loss or disadvantage
- damage to reputation, privacy or psychological wellbeing
- exclusion from benefits or services.

This applies to individuals, public bodies and private organisations. Therefore, as well as applying applicable legal and regulatory obligations, members are encouraged to think broadly about how data could cause harm, both in the context of its current purpose, and any future uses if the data is kept, linked or shared.

Practice examples

The following practices contain examples of ways to minimise and manage harm in data science work:

Practices	
What	How
Using data that is 'ethically sourced', for example where data subjects have knowingly given their data and the perceptions of data subjects have been considered	<ul style="list-style-type: none">• Understanding and asking questions about the origin of the data, such as how it was collected• Investigating the attached consents and legal uses for the data• Considering the privacy, dignity and fair treatment of individuals when selecting data• Leaders being aware of and understanding the implications of using 'big data'
Linking in privacy and ethics into work, as well as legal and regulatory requirements	<ul style="list-style-type: none">• Following good data handling practices, including data security, and considering ways to reduce the need for collecting, storing and utilising personally identifiable information• Considering the impact of deriving demographics or linking with other data• Being transparent and providing evidence of privacy considerations to the public and/or regulators



What	How
Embedding ethical risk management into work	<ul style="list-style-type: none"> • Using ethical assessment checkpoints • Seeking independent or domain-expert advice on assumptions and risks • Appropriately monitoring and communicating ethical risks and potential biases to relevant stakeholders • Regular ethical reviews of models/algorithms • Applying organisational or industrial ethical policies • Leaders putting clear ethical governance in place for data science work
Educating the workforce to recognise the risks for harm and/or unfairness	<ul style="list-style-type: none"> • Communicating to others what data science can and cannot do • Helping non-technical individuals understand the ethical, professional and technical issues relevant to a project • Leaders ensuring that their organisation as a whole understands the ethical principles and policies surrounding data science
Considering the impact on the environment and its resources	<ul style="list-style-type: none"> • Fully understanding error, bias and risks associated with models applied to environmental problems • Evaluating the energy-cost of storing and processing large volumes of data, and considering different options such as cloud computing, aggregating data and regularly reviewing the need for data



3. Applying and maintaining professional competence

Practitioners can help minimise uncertainty and risk in their work by complying with best industry and professional practices and applying analytical rigour. This links closely with the other listed ethical considerations and can be challenging in the data science setting where the use of data is often complex.

Practice examples

The following practices may help to maximise the benefits of data science and minimise potential harm.

Practices	
What	How
Complying with relevant professional and regulatory practices	<ul style="list-style-type: none">• Investigating what the relevant professional and legal/regulatory requirements are to the work in which you are involved• Ensuring that the work complies with these requirements and justifying this to relevant authorities• Being aware of any evolving legal requirements and communicating these appropriately to relevant stakeholders
Ensuring the business's ethics policies, procedures and governance are applied to data science work	<ul style="list-style-type: none">• Providing executive staff with enough information on the advantages and limitations of the work to make decisions about the use of models• Understanding and communicating to stakeholders and decision-makers the ethical risks of the project, such as bias, uncertainty, quality issues, individual/commercial harm, methodology assumptions, disadvantages of chosen methods• Leaders setting clear lines of responsibility• Educating the workforce appropriately about relevant policies and procedures
Following best practice for data analysis	<ul style="list-style-type: none">• Using appropriate statistical and algorithmic methods for the question being addressed• Leaders taking the level of uncertainty in work into account when using data science in any decision making
Monitoring and maintaining models	<ul style="list-style-type: none">• Having clear owners for models• Monitoring the performance of models• Having regular reviews of models• Having an agreed action plan for updating or retiring models that are no longer fit for purpose

What	How
Validating and improving work	<ul style="list-style-type: none"><li data-bbox="507 264 1342 331">• Thoroughly quality-assuring data and analysis, including for bias, error, uncertainty<li data-bbox="507 349 1198 383">• Encouraging challenge from others involved in the project<li data-bbox="507 400 1445 468">• Sharing appropriate code, diagnostics and results with colleagues/customers to ensure they sufficiently understand the models<li data-bbox="507 486 1007 519">• Publishing results publicly where possible<li data-bbox="507 537 1286 571">• Seeking feedback from peers, academia or subject matter experts



4. Seek to preserve or increase trustworthiness

Data science involves the analysis of varied sources of data and greater use of AI. As this is relatively new area, public perceptions are still evolving and developing. In order to build trust and understanding of their work, practitioners will need to work with the public and stakeholders. Trustworthiness can be built and maintained if organisations using data science engage regularly with stakeholders, including potential or actual critics, and have an open and honest dialogue about the use of data and AI. This engagement may also help to better understand potential sources of bias and risk.

Practice examples

Below is a table describing how this could be put into practice.

Practices	
What	How
Engaging with appropriate ethical bodies	<ul style="list-style-type: none">• Leaders in this area may wish to communicate with ethical bodies and the public about how data is used and handled in relation to data science work
Avoiding unnecessary complexity in methods to improve transparency	<ul style="list-style-type: none">• Considering simpler models and documenting the performance of different models tested• Disclosing the reasons for any improvements in accuracy if one model is favoured over another
Building trust through transparent communication with the public	<ul style="list-style-type: none">• Communicating the ways in which data is used and the benefits to customers/wider public• Taking accountability for involvement in the work• Considering the views of the public and data subjects along with associated consents and legal uses of data• Publishing and sharing methods with relevant internal stakeholders, customers and the wider public, clearly referencing limitations, bias and uncertainty• Being honest and objective in communications

5. Maintain accountability and oversight

There are various emerging views on how to manage AI and whether it poses risks to society as well as opportunities. A common ethical consideration is how accountable individuals are when introducing AI and what level of oversight is put into place when delegating any decision making to machines.

Practice examples

Below is a table describing how this could be put into practice.

Practices	
What	How
Maintaining human oversight of automated solutions	<ul style="list-style-type: none">• Implementing model governance by deciding how to monitor the model over time, such as setting review points and sign-off• Agreeing where responsibility lies for models in production (approval, reviews, longer-term quality assurance)• Agreeing a review process
Clearly defining operating constraints	<ul style="list-style-type: none">• Adding clear constraints within systems and corresponding documentation to control and communicate what systems can and cannot do
Public involvement	<ul style="list-style-type: none">• Using public engagement and involvement to feed into projects, where appropriate• Checking what public concerns may be
Governance and accountability structures	<ul style="list-style-type: none">• Having ethics oversight committees where appropriate• Having clear processes for accountability so that stakeholders have a clear route to any redress

To find further reference material and resources on how this is put into practice, please visit www.actuaries.org.uk and www.rss.org.uk.

Implementation checklist

This summary of the ethical practices highlights when they can be implemented within a project to help embed ethics into data science work.

Project planning

To embed ethics when defining, scoping and initiating projects, you can:

- Complete an ethics assessment including:
 - Is it in the public interest, and how might value be distributed fairly?
 - Can data be ethically sourced?
 - Are there risks (privacy, harm, fairness) for individuals, groups, businesses, environment?
- Engage with the public/stakeholders to gather perceptions
- Seek early feedback from domain experts
- Define the governance for the project, including data security and handling

Data management

To embed ethics when acquiring, storing and preparing data for analysis, you can:

- Fully understand the consents and legal uses of the data
- Have data security measures in place, and ensure staff are trained in data handling
- Consider impacts of data processes to privacy, bias and error; including linking data sources, estimating demographic or other omitted information

Professional competence might include:

- Thorough quality checks
- Applying best practices for data cleaning (errors, inconsistency, missing values)
- Detecting and mitigating sources of bias

Analysis and development

To embed ethics when analysing data, building models and designing systems, you can:

- Apply consents and permitted uses, professional and regulatory requirements
- Monitor risks identified at planning, assess for additional risks (harm, bias, error, privacy)

Professional competence might include:

- Best practices, analytical rigour, quality assurance of methods, including peer review
- Minimising complexity in models and algorithms, validating thoroughly
- Fully explaining outcomes and uncertainty when making recommendations

Implementation and delivery

To embed ethics when operationalising and deploying models and systems, you can:

- Be transparent about when, how and why decisions have been delegated
- Share methods, results/limitations and code with relevant parties; invite feedback

Professional competence might include:

- Applying best practice in anonymisation before sharing data or disseminating outputs
- Adequate quality assurance of systems
- Regular reviews of models, with risk assessments

Communication and oversight

To embed ethics in communications:

- Communicate with stakeholders throughout projects and help any non-technical individuals understand the technical and ethical issues
- Be honest and objective when communicating about work

Leaders can:

- Ensure the organisation has and understands ethical principles and policies
- Engage with ethical bodies, privacy groups, stakeholder groups, the public



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