Department for
Digital, Culture
Media & Sport

Data Ethics Framework

Date: 13 June 2018

Contents

Ministerial Foreword

Making better use of data offers huge benefits, in helping us provide the best possible services to the people we serve.

However, all new opportunities present new challenges. The pace of technology is changing so fast that we need to make sure we are constantly adapting our codes and standards. Those of us in the public sector need to lead the way.

As we set out to develop our National Data Strategy, getting the ethics right, particularly in the delivery of public services, is critical. To do this, it is essential that we agree collective standards and ethical frameworks.

Ethics and innovation are not mutually exclusive. Thinking carefully about how we use our data can help us be better at innovating when we use it.

Our new Data Ethics Framework sets out clear principles for how data should be used in the public sector. It will help us maximise the value of data whilst also setting the highest standards for transparency and accountability when building or buying new data technology.

We have come a long way since we published the first version of the Data Science Ethical Framework. This new version focuses on the need for technology, policy and operational specialists to work together, so we can make the most of expertise from across disciplines.

We want to work with others to develop transparent standards for using new technology in the public sector, promoting innovation in a safe and ethical way.

This framework will build the confidence in public sector data use needed to underpin a strong digital economy. I am looking forward to working with all of you to put it into practice.

The Rt Hon Matt Hancock MP, Secretary of State for Digital, Culture, Media and Sport

Introduction

What is data ethics?
Data ethics is an emerging branch of applied ethics which describes the value judgements and approaches we make when generating, analysing and disseminating data. This includes a sound knowledge of data protection law and other relevant legislation, and the appropriate use of new technologies. It requires a holistic approach incorporating good practice in computing techniques, ethics and information assurance.

What is data science?
A core aspect of data ethics is using data science appropriately. Data science describes analysis using automated methods to extract knowledge from data. It covers a range of techniques, from finding patterns in data using traditional analytics to making predictions with machine learning. It presents new opportunities for identifying factors for answering important policy questions — factors which might be difficult for people to spot on their own. Data science, therefore, offers huge public benefits in creating better evidence-based policy and in making government operations more targeted and efficient. However, we must carefully consider the social implications of the data and algorithms used, our practices and the quality assurance processes we follow to ensure this is done well.

Why do we need a framework?
Advances in computing power and techniques mean newer, more powerful, computational models or data science tools are seeing uptake across the public sector. Coupling this with an increase in skills means we now have the ability to analyse larger volumes of data more rapidly and more regularly.

Increasingly public servants from across disciplines will need to understand insights from data and emerging technologies. It is crucial that public servants are equipped to use data-informed insight responsibly and processes must be in place to support this.

How to use the Data Ethics Framework

The Data Ethics Framework guides the design of appropriate data use in government and the wider public sector. This guidance is aimed at anyone working directly or indirectly with data in the public sector, including data practitioners (statisticians, analysts and data scientists), policymakers, operational staff and those helping produce data-informed insight.

The Data Ethics Framework builds on the core values of the Civil Service Code - integrity, honesty, objectivity and impartiality - to encourage ethical data use to build better services and inform policy.

Specifically the framework will help:

- policy or operational professionals understand the uses and limits of data science, define ethics-related requirements, and develop context and domain-specific questions when planning a project or writing a tender
- data practitioners ensure that they have considered policy and subject matter when designing data science approaches, and develop questions around project-specific requirements
- information technology providers better understand the core ethical expectations for public sector data science projects, and to tailor their offerings appropriately

All public servants and especially data practitioners, working with data, have a responsibility to act appropriately with that data. This framework is intended to provide advice on how to do that. It should be used throughout public sector organisations, especially where a prior governance framework does not exist.

The framework consists of 3 parts:

- the data ethics principles
- additional guidance for each principle in the framework
- a workbook to help your team record the ethical considerations you've made about your project

Teams should work through the framework together before starting the design or discovery phase of a new project, policy or service. Use the workbook to consider legal and ethical questions to inform the best approach for your use of data.

Each part of the framework is designed to be regularly revisited throughout your project, especially when any changes are made to your data collection, storage, analysis or sharing processes.

Data Ethics Framework Principles

Your project, policy, service or procured software should be assessed against the 7 data ethics principles.

1. Start with clear user need and public benefit
Using data in more innovative ways has the potential to transform how public services are delivered. We must always be clear about what we are trying to achieve for users - both citizens and public servants.

2. Be aware of relevant legislation and codes of practice
You must have an understanding of the relevant laws and codes of practice that relate to the use of data. When in doubt, you must consult relevant experts.

3. Use data that is proportionate to the user need
The use of data must be proportionate to the user need. You must use the minimum data necessary to achieve the desired outcome.

4. Understand the limitations of the data
Data used to inform policy and service design in government must be well understood. It is essential to consider the limitations of data when assessing if it is appropriate to use it for a user need.

5. Use robust practices and work within your skillset
Insights from new technology are only as good as the data and practices used to create them. You must work within your skillset recognising where you do not have the skills or experience to use a particular approach or tool to a high standard.

6. Make your work transparent and be accountable
You should be transparent about the tools, data and algorithms you used to conduct your work, working in the open where possible. This allows other researchers to scrutinise your findings and citizens to understand the new types of work we are doing.

7. Embed data use responsibly
It is essential that there is a plan to make sure insights from data are used responsibly. This means that both development and implementation teams understand how findings and data models should be used and monitored with a robust evaluation plan.

1. Start with clear user need and public benefit

How to implement principle 1 of the Data Ethics Framework for the public sector.

Before you start working with data, you must consider the user need and expected public benefit. Understanding this is fundamental to making sure you take the right approach. Having a clear user need means that you know what problem you are trying to solve, even if you don't yet know what the solution is or even the path to the solution. Considering the Data Ethics Principles holistically should help you make a decision on whether your approach is delivering public benefit.

How data can help meet user needs

Teams in the public sector must clearly define user needs to make sure:

- service teams build the right government service
- content on GOV.UK is clear, helpful and to the point
- support and skill requirements, risks and operational needs are considered before new technology purchases

A range of public sector user needs which can be met using data analysis are:

- running and improving services
- building new services
- trialling new processes for internal operations
- testing existing and new policies

For these needs, using data can:

- help you identify themes in large volumes of text
- predict what will happen
- automatically categorise stuff
- spot something unusual
- show you how things are connected to each other
- spot patterns in large volumes of data
- spot geographic patterns in services or data

Often projects involving data analysis are requested by non-practitioners - people with an ill-defined problem they would like to understand better. Reframing their request as a user need will help you understand what they're asking for and why, or expose what you don't know yet.

Writing your user needs

Writing out the needs you've identified can help make sure the whole team understands the project objective. User needs are usually written following a set format.

Example
As a ...
I need to ...
So that I can ...

Your user need should not focus on a specific technology, data analysis technique, or dataset. If there's more than one, prioritise the needs to help your project focus on the most important thing.

Example user needs

Here are some examples of real user needs

- As a user researcher working in the Government Digital Service
- I want to understand if natural language processing can support user research
- So that I can analyse research findings more efficiently

or:

- As a policy expert at the Department for Education
- I need to identify schools with maintenance issues
- So that I can understand the effect of school investment patterns on maintenance

or:

- As a data analyst working in a fire and rescue service
- I need to identify homes which are likely to not have a smoke alarm fitted
- So that I can advise how to prioritise fire safety checks

Determining the public benefit

When determining user needs, public servants must also be confident they are acting in the public benefit. Defining public benefit means demonstrating that:

- your approach offers value for money
- the appropriate governance and decision-making oversight exist to ensure success of the project
- potential risks or negative consequences have been weighed up against the risk in not proceeding
- there is supporting evidence for each of the above

Determining the public benefit is crucial for any project, but is particularly important before any large expenditure. HM Treasury provides a series of guidance to aid public servants:

- the Green Book provides guidance for how to achieve transparent, objective and evidence-based advice for decision making in Government
- Managing public money explains how public servants can ensure value for money for citizens, while maintaining transparency and accountability

Evidence that you've understood the problem correctly

Data scientists and policy professionals should work together to make sure everyone involved understands:

- the overall problem you're trying to solve
- who the users of this data process or analysis are
- what needs they have

Try these things to help you better understand the problem:

- apply the Service Manual's guidance on learning about users and their needs and general user research guidance
- research any past similar projects
- consult more widely or get independent advice - read the Open Policy Making Toolkit guidance on consulting
- speak to operational staff to understand how data is collected and what impact any new tool might have

Having a clearly defined user need will determine the project approach.

You should work through the Data Ethics Framework principles 2 through 7 to help you decide what tools are most appropriate and proportionate for meeting your user need.

2. Be aware of relevant legislation and codes of practice

How to implement principle 2 of the Data Ethics Framework for the public sector.

You must be aware of legislation and codes of practice that apply to your use of data.

This includes knowing about:

- legislation that applies to your proposed data use
- how to produce statistics
- data protection by design
- data minimisation
- information governance

Other important pieces of central government guidance that are helpful for using data and designing projects in the public sector include:

- The Civil Service code
- HM Treasury Aqua Book: guidance on producing quality analysis for government
- HM Treasury Magenta Book: guidance for evaluation

What the law says

Here are some important pieces of legislation that typically apply to using data. If you are unsure how relevant laws might affect your work, speak to a legal adviser within your organisation.

Personal data

If you are using personal data, you must comply with the principles of the EU General Data Protection Regulation (GDPR) and the Data Protection Act 2018 (DPA 2018) which implements aspects of the GDPR and transposes the Law Enforcement Directive into UK law. It also provides separate processing regimes for activities which fall outside the scope of EU law.

Personal data is defined in Section 3(2) DPA 2018 (a wider explanation is detailed in Article 4 of the GDPR).

Equality and discrimination

Analysis or automated decision making must not result in outcomes that lead to discrimination as defined in the Equality Act 2010.

Sharing and re-use of data

When accessing or sharing personal data, you must follow the Information Commissioner's Code of Practice for Data Sharing which should be read alongside the ICO's Guide to GDPR. This Code of Practice is due to be updated to align with the new Data Protection Act 2018.

When accessing and sharing data under powers in Part 5 of the Digital Economy Act 2017, you must follow the relevant Codes of Practice.

When re-using published and unpublished information relating to public tasks, you must follow the Re-use of Public Sector Information Regulations 2015.

Copyright and intellectual property

Copyright and intellectual property are often governed by combinations of statutes.

When using data, respect copyright laws and database rights, covered in part by the Copyright and Rights in Databases Regulations 1997.

When procuring software, consider potential intellectual property constraints covered in the Intellectual Property Act 2014.

Freedom of information

Your use of data may be subject to the Freedom of Information Act 2000. You should also consider the wider publishing of datasets released following a Freedom of Information request, in accordance with the Protection of Freedoms Act 2012.

Sector specific legislation

Specific sectors like finance and health have further data use legislation and frameworks, including those relating to the use of non-personal data. Health research has its own UK Policy Framework for Health and Social Care Research drafted by the NHS Health Research Authority (HRA). The NHS HRA also provides specific guidance for health researchers on the new data protection principles being introduced by the General Data Protection Regulation.

Statistics

When using or producing statistics, you must follow the Code of Practice for Statistics.

The National Statistician's Data Ethics Advisory Committee (NSDEC) provides independent and transparent ethical assurance that the access, use and sharing of public data for research and statistical purposes is ethical and for the public good. The UK Statistics Authority can work with

statisticians and researchers to identify potential ethical issues in their research and guide them through the NSDEC application process.

Data protection by design

Data protection by design and by default is a legal requirement under the GDPR. It means taking a holistic approach to embed data protection from design through to application of any use of personal data.

GDPR requires that anyone handling personal data protects the rights of individuals by:

- using personal data for a specific task
- putting in place technical and organisational measures to implement data protection principles effectively
- integrating necessary safeguards into the processing of personal data

This includes a commitment to completing data protection impact assessments (DPIA) (also known as privacy impact assessments) throughout the lifecycle of your project or service. DPIAs are an important tool for identifying privacy risks.

It is a legal obligation under Article 35 of the GDPR to complete a DPIA when there's likely to be high risk to people's rights, particularly when using new technologies. However it is often good practice to do a DPIA for any use of personal data.

Things to think about:

- always seek the advice of your organisation's Data Protection Officer when doing a DPIA
- privacy should be considered throughout the project – although you may not be using personal data at the outset of your work, the project type and privacy considerations may change as work develops
- you should consider how often you will repeat the DPIA when using personal data and may need to change this if the project changes significantly
- when joining a new project, seek out and review the existing DPIA to familiarise yourself with any risks to rights and freedoms identified and the relevant mitigation strategies proposed
- if you discover a DPIA has *not* been completed for a project for which it is relevant, this should be flagged as soon as feasible
- see the ICO's guidance on DPIAs, including its Privacy Impact Assessments Code of Practice, and practical advice on doing DPIAs for data analytics

Accountability

An important aspect of complying with data protection law, is being able to demonstrate what measures you are taking to ensure this (see Article 5(2) of the GDPR (the accountability principle) and Article 30 on keeping records of processing activities).

Your organisation and information assurance teams will be responsible for this at a high level including ensuring policies and training are in place. However, it is essential to show how you are doing this at an individual level, through thorough documentation of things like Data Protection Impact Assessments.

Data minimisation

You must use the minimum data necessary to achieve your desired outcome.

Article 5(1)(c) of the GDPR states that personal data should be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed. You should only use as much data as is required to successfully accomplish a given task.

The law tells us the minimum we need to do to protect the rights of citizens when using data. However, when deciding if a particular data use is ethical, we need to think beyond legal compliance only. See principle 3 (Use data that is proportionate to the user need) to evaluate proportionality.

Information governance

Organisations have a responsibility to keep both personal data and non-personal data secure.

How personal data should be collected, stored, shared, processed and deleted is covered by the General Data Protection Regulation (GDPR) and the Data Protection Act 2018 (DPA 2018).

Government departments, services and public bodies set out how they use, store and share personal data including how data subjects can exercise their rights in their Personal Information Charters or service privacy notices. Personal Information Charters contain guidance on how people can access their data, as prescribed in Articles 13 and 14 of the GDPR. A useful example of a Personal Information Charter is from the Department for Work & Pensions.

The Security Policy Framework requires that risk assessments are carried out to 'identify potential threats, vulnerabilities and appropriate controls to reduce the risks to people, information and infrastructure to an acceptable level'.

Information assurance (IA) helps do this by:

- assessing the information risks
- helping to define the appropriate measures required to reduce those risks to levels acceptable to your organisation's risk appetite
- ensuring that contracts provide the required measures

You should engage as early as possible with your IA specialists so they can provide effective support through all stages of your work.

In many organisations information risk is overseen by a Senior Information Risk Owner (SIRO). Usually your organisation will have a risk appetite statement that sets out how information risk is managed.

You should consult with your information assurance team when you need to delete data.

3. Use data that is proportionate to the user need

How to implement principle 3 of the Data Ethics Framework for the public sector.

Your use of data must be proportionate. Whether your proposed data collection, storage or analysis is proportionate will depend on:

- the user need and expected public benefit, which according to Principle 1, you should be able to clearly demonstrate
- the type of data (personal or non-personal)
- whether personal data can be de-identified, also known as pseudonymising or anonymising
- how and where you got the data

If you decide your proposed data use isn't proportionate, then either:

- change your data source, collection mechanism or analysis technique so that it becomes proportionate
- consider if there's some other way to meet your user need (like qualitative user research)

You must not proceed with your project if your data use is not proportionate to the user need.

When deciding if a use of data is proportionate, you should document the process used to determine this and any supporting evidence.

Personal data and proportionality

De-identifying personal data before use

When using personal data it is good practice to work with data that has been de-identified to the greatest degree possible.

This process is often called anonymisation, although it is important to note that many forms of data can never be fully and irreversibly anonymised. Pseudonymisation is a related process of removing the most re-identifying components of the data and storing them separately.

If data is anonymised to the greatest degree possible, it is likely to be out of scope of data protection law as it is no longer considered personal data. Pseudonymised data, however, is subject to the same laws as fully identifiable personal data. Recital 26 of the GDPR provides further information.

If you plan to anonymise or pseudonymise personal data before linking or analysis, make sure you follow the ICO's Anonymisation: managing data protection risk code of practice and document your methods.

You can find more technical advice in the UK Anonymisation Networks anonymisation guidance.

Synthetic data

You may also be able to use synthetic data, replacing sensitive data with a set of plausible values, where individual citizens would not be identifiable**.** Synthetic data is becoming increasingly popular for training machine learning models while preserving individuals privacy. The model can then be used to make predictions on real data. Synthetic data is only as good as the underlying generative methods used to produce it which can impact the performance of the model on making predictions on real data. This can be a useful approach when working on very sensitive data or specific machine learning applications such as computer vision.

It is important to remember that pseudonymising or anonymising data or creating synthetic data does not make it automatically appropriate for your objective to use. It is possible to make incorrect inferences and develop potentially intrusive or damaging policies based on less identifiable data. Considering the Data Ethics Principles holistically should help you make a decision on whether your approach is proportionate and appropriate.

Determining proportionality

You should assess the proportionality of your proposed approach i.e. whether it is necessary and appropriate. The following questions must be answered in multidisciplinary teams made up of data practitioners and subject matter and operational experts. Evidence to support your decision should be recorded and accessible to any individual joining the project at a later date.

- Is the measure suitable to achieve the aim?
- Is the measure necessary to achieve the aim?
- Would the proposed use of data be deemed inappropriate by those who provided the data?
- Would the proposed use of data for secondary purposes make it less likely that people would want to give that data again for the primary purpose it was collected for?

Some other ways to understand if the data you intend to use is proportionate to the user need are:

- user research or speaking to the public about your plans
- getting advice from ethics committees and individual experts, like the National Statistician's Data Ethics Advisory Committee
- reviewing public dialogue studies on how citizens feel about data use, including the Government Data Science Partnership and Ipsos Mori review of government use of data science and the Royal Society and British Academy Data governance: Public engagement review
- speaking to colleagues in different disciplines to get a broader view of the policy issue, including your organisation's Data Protection Officer
- getting advice from the Government data science community and the Data Leaders Network

Data sources and proportionality

Data insight used to inform policy making and service design must be representative and accurate.

Usually government generally relies on four sources of data for analysis:

- repurposed operational data
- repurposed third party data e.g. social media
- statistics (derived from some other form of raw data)
- purposefully collected data (through new processes)

No one data source is in itself inappropriate to use for analysis, but the needs of the user must always be considered alongside any relevant data protection subject rights, to assess the overall suitability.

For any data being repurposed for analysis, without original individual consent, you must assess whether or not the new purpose is compatible with the original reason for collection (Article 6(4) GDPR).

Repurposed operational data

You must consider the proportionate use of departmental operational data i.e. data collected through the operation of the department or a delivery of one of its services. Although most organisations have well-established processes for using operational data to improve services, proportionality must be considered for each new piece of work.

Repurposed third party data

Repurposed personal data from third party sources

If obtaining personal data from sources, other than directly from an individual, you must be able to make this fair and transparent. This means being able to provide data subjects with the information listed in Articles 13 and 14 of the GDPR within a reasonable period, and no later than one month, after obtaining their personal data.

You should always consider the effects that using this data will have on them and what their reasonable expectations are likely to be. You also need to determine if you are obtaining any special category personal data or data relating to criminal convictions and offences which the GDPR and DPA 2018 gives more protection. In order to lawfully process special category data, you must identify both a lawful basis under Article 6 of the GDPR and a separate condition for processing special category data under Article 9, as supplemented by Section 10 and Schedule 1 of the Data Protection Act 2018 .

Social media data

Social media data must be used responsibly. Using some social media data may be considered too intrusive to use without an individuals consent. It can also be difficult to determine the representativeness of social media data when working at regional or national level.

The Government Social Research (GSR) profession has published guidelines for using social media data responsibly in research. The GSR have also published a social media ethics grid which aims to aid ethical decision making when using social media data.

Web scraped data

You need to consider if using web scraped data is appropriate for the intended data analysis. You should also ensure individuals' privacy is respected. Although information is publicly accessible, this does not automatically provide a lawful basis for processing.

Even where its use is legal, citizens may feel scraping of particular websites is not ethical. Scraping information citizens consider private can be controversial. If scraping social media sites and forums, decide whether your intended use is intrusive or breaches citizen trust.

When web scraping you must:

- always respect website terms and conditions and robots exclusion protocol (like robots.txt)
- make sure you do not breach intellectual property rights if you republish any data sourced from the web
- schedule web scraping activities to minimise the impact on target websites
- not scrape websites anonymously - make sure an identifiable IP address is visible
- have data protection processes in place to manage any personal data you unintentionally collect

Purposefully, newly collected data

If you are collecting new data specifically designed for your project through a tool or application, an End User Licensing Agreement (EULA) must be presented to users. To ensure informed consent, the EULA must fully explain the terms and conditions of their data disclosure, including the objective of the data collection and when it will be destroyed.

Consent is one of six lawful bases for processing personal data, under data protection legislation. GDPR makes it clear that if relying on consent from data subjects, it must be informed, unambiguous and involve a clear affirmative opt-in. Read the UK Data Service guidance on getting consent for research. If your data use falls under a public body carrying out its tasks, it is unlikely that any consent will be considered 'freely given'. This means consent is unlikely to provide a valid legal ground for data processing.

Statistics

Production is covered by the statutory Code of Practice for Statistics and subject to independent regulation, ensuring that these data uphold rigorous quality standards, are well documented and are free from bias. Statistics can originate from survey, census and operational data sources, but have usually undergone extensive post-processing and quality assurance.

Statistics are already aggregated and released according to strict statistical disclosure procedures so there is a lower risk of disproportionate use.

4. Understand the limitations of the data

How to implement principle 4 of the Data Ethics Framework for the public sector.

Though legal and proportionate, there may be limitations to your data that make your proposed approach inappropriate, unreliable or misleading - and therefore unethical as a basis for public sector policy making or service design.

Things to consider when deciding if a source of data is suitable include:

- provenance (for example how and why the data was collected)
- errors in the data
- bias (from historical decision making, or unrepresentative surveys or social media)
- if metadata and field names are ambiguous

Provenance

When designing a new use of data, you must understand the impact of data provenance on accuracy, reliability and representativeness.

Specifically assess the impact of:

- the source of the data, such as a transactional service, survey, administrative task in a public sector organisation, a government department, social media or open dataset
- whether the data was collected by humans or an automated system
- how well the data reflects its target population
- any likely omissions, exclusions or systematic biases
- patterns in the data and whether they are likely to stay static or change over time
- quality assurance processes when the data was collected
- the sampling strategy used to collect the data
- any other problems surrounding data collection

Errors

Errors in data are inevitable; however it can be difficult to understand how frequent they are, if they are random, the cause and ways to mitigate or remove them. Errors are not always immediately obvious, especially in large datasets. Simple data visualisations can be the best way of spotting anomalies and systematic errors.

You will need to consider and document how identified errors will impact the work.

If you find errors in the way data is collected or interpreted, report them to policy or operational staff.

The UK Statistics Authority Quality Assurance of Administrative Data framework provides useful

resources to help you understand the data that you are using, how it was collected and any likely quality impacts.

Bias

You should be aware of the types of bias that can exist in the data you are using by reviewing how the data was collected.

There are many ways in which bias can be introduced into datasets, through collection techniques, limited representativeness of a particular cohort and social bias from historical decision making.

Carefully considering potential bias and its impacts on outputs from data analysis is technically well established. When using data in more contexts, to inform policy or service design, it's critical to involve policy or subject matter experts to fully consider types of bias which might not be immediately obvious to a data practitioner.

Measurement bias

Bias in measurement is the selection of data or samples in a way that does not represent the true parameters (or distribution) of the population.

Social bias

Any data sources about citizens, collected from services, surveys, or elsewhere will contain some level of social bias as the information is based on historic decisions and actions by humans, or was shaped by laws no longer in force.

Bias in training data leads to bias in algorithms. Machine learning is a data-driven technology and the characteristics of the data are reflected in the properties of the algorithms.

Read more about social bias in algorithms in Principle 5.

Social media

Data from social media sources may give valuable real time or historic insight, but the data should be properly investigated to identify any representation or selection bias. Include metrics and caveats about who or what the data is representative of, and importantly, what you cannot determine from the data.

Practitioner bias

Data practitioners and others involved in a project may inadvertently introduce their own confirmation bias into the design of projects, analyses, or interpreting outputs. Ensuring a diverse team from a range of backgrounds is a good way to mitigate potentially damaging practitioner bias.

Survey methodology

Surveys must be carefully designed and used to ensure they cover your target population. Low response rates may mean it's inappropriate to use survey data.

If the proportion of non-respondents or 'invisible data' is too high in a survey, it may be irresponsible to describe the results of work with this data as representative. Determining your proportion of 'invisible data' is also crucial when using machine learning to spot correlation or network effects.

Response and selection bias affect the generalisability of findings. With high bias, you can't infer that patterns exist beyond the sample who responded.

Metadata and field names

Metadata and the names of fields in datasets can be misleading or inaccurate. It's critical that you work with the subject matter expert for the dataset collection to understand if it is fit for purpose for your project. Improve documentation of the metadata, if you can.

5. Use robust practices and work within your skillset

How to implement principle 5 of the Data Ethics Framework for the public sector.

To make best use of data we must ensure we have robust and consistent practices.

This involves:

- working in multidisciplinary teams
- getting help from experts outside your team
- ensuring accountability of algorithms
- avoiding outputs of analysis which could result in unfair decision making
- designing for reproducibility
- testing your model under a range of conditions
- defining acceptable model performance: false negatives and false positives

Public servants must work within their skillset, recognising where they do not have the expertise to use a particular approach, type of data or tool to a high standard.

Read the Government Aqua Book before you design any analytical quality assurance process. It gives guidance on setting up analytical process and developing the right culture for providing reliable and accurate evidence for policy making.

Within your team, you should establish and document a consistent process for delivering a data science project or new process using data. Having a consistent approach to delivering projects will improve efficiency and simplify management.

Multidisciplinary teams

Ensuring practitioners are working in multidisciplinary teams, with access to all necessary subject matter expertise, is essential to ensure robust processes are established for using data appropriately.

Questions to consider before starting any new project:

- are data scientists working in multidisciplinary teams to assess how data can be used to meet a user need?
- do people at all levels of the team understand how and why the proposed use of data is expected to deliver or aid the delivery of a solution to the user need?

Getting help from experts

When using data or designing new processes or tools to do so, it's essential that you recognise when you do not have all the necessary information or expertise to do something safely or accurately.

Always seek expert help when you do not have all the necessary skills or expertise.

Accountability of algorithms

You should always use the most simple model to achieve the desired outcome.

As machine learning algorithms are designed to learn and develop semi-independently, the processes used can become more difficult to interpret and understand. Teams need to have a reasonable understanding of how the machine learning or pipeline of machine learning models has worked to meet the user need.

You must be able to explain this to non-specialists.

You can design your machine learning process to improve accountability.

Staged process

This involves tackling prediction tasks as a pipeline of machine learning models which should facilitate the potential for overall interpretability. The staged process must be clearly documented with all necessary assumptions and caveats articulated.

Often we do not need to know exactly how a machine learning model has arrived at a particular result if we can show logical steps to achieving the particular outcome. This includes exactly what training data was used and which variables have contributed most to a particular result.

Simplification

Teams should always use the simplest model to achieve the intended measurable and quantifiable goal. This must be balanced against a potential loss in accuracy of the model. Pay extra attention to lost accuracy disproportionately affecting subgroups of the data that might not be well analysed by a simpler approach.

A more complex machine learning model is more likely to lose interpretability. This will be more or less tolerable depending on the intended outcome.

Social bias in algorithms

The Equality Act 2010 makes it illegal to discriminate against anyone based on nine protected characteristics.

The Equality Act 2010 includes the public sector equality duty which requires organisations to eliminate discrimination and support the advancement of equality.

Anyone doing analysis or making policy in the public sector must:

- make sure that any gathered evidence does not inadvertently produce discriminatory decisions
- recognise opportunities within their role to usefully flag social biases in data

This applies irrespective of the technique.

Analysis is most often done using historical data. This data might contain issues or patterns that policies are trying to mitigate, not replicate.

Proxy variables in machine learning

Some machine learning methods used to inform decisions can inadvertently base outputs on implicit proxies for variables which might be undesirable or even illegal.

Note that it is often not sufficient to remove protected characteristics from analysis to remove the possibility of discriminatory outcomes based on these variables. This journal article explains a number of statistical reasons for this.

One important reason is that if a protected feature is predictively powerful, you might be able to infer that protected feature from the other data you have. To make sure you're not implicitly using that feature in decision making, you have to remove statistical traces of it from the whole dataset, rather than just omit it.

This does not prevent analysis being carried out on protected characteristics. Analysis on protected characteristics gives important insight for policy making, but you must make sure it does not inadvertently result in discriminatory actions, for example tailoring service delivery in an unfair way.

Reproducibility

Reproducibility in data-informed services and decision making is essential to:

- demonstrate accuracy
- aid transparency and accountability
- allow others to use and share your work
- ensure consistency of analysis in an organisation

The 3 requirements for reproducibility are: applying the same logic (code, model or process), to the same data, in the same environment.

The whole workflow should be supported by high quality documentation, including ethical issues raised in this framework and ways to mitigate them.

Using the simplest model possible, with adequate documentation, makes it easier for other teams to understand and use your work.

Version control

Using version control allows an analyst to create a very clear audit trail. It can also be used to formalise a system of quality assurance (QA), for example by pull request review. There are several tools for version control.

Other literate programing tools (like Jupyter notebooks and Rmarkdown) allow researchers to

combine code and analysis with narrative, making for much more transparent analysis. The Software Sustainability Institute has done work on developing and publicising best practice in academia, which can be applicable to data science in government.

Software Carpentry and Data Carpentry have useful lessons which advise on good practice methods for version control and processes more broadly.

Reproducible Analytical Pipelines (RAP)

Government Digital Service (GDS) data scientists have trialled the use of tools for creating reproducible workflows called Reproducible Analytical Pipelines (RAP). The RAP approach draws on existing process management tools such as reproducible research, software engineering and DevOps.

RAP aims to aid practitioners during the analytical process by automatically recording:

- what changes were made
- who made those changes
- why those changes were made
- when those changes were made

It is important that you research and think about the best approach for your team. Email GDS for advice on developing a pipeline within your organisation.

Test the model

Good data science involves testing your models against reality, using appropriate model evaluation techniques.

Make sure you have a clear methodology for testing your findings from the start. This should include desired levels of accuracy and other performance measures.

Testing algorithmic systems before going live is essential.

As part of testing the model, you need to develop guidance on how often you need to update the data that trains it. How regularly do you need to check on a dynamic or live model that receives data about how it's performing? This is covered more broadly by principle 7.

Define acceptable model performance: false negatives and false positives

You must decide what is acceptable in terms of false negatives and false positives within your intended system. This will determine the type of model and metrics you choose to use.

What your model is used for will determine the threshold for potential errors. Some false positives or negatives can be disastrous, while others simply result in wasted time and resources.

The target you set will determine the type of algorithm, metrics and acceptable loss function. Model performance as measured by these metrics should be compared to a null model or competing

model. You must consider how relevant options or decisions can be explained for the purpose of accountability.

6. Make your work transparent and be accountable

How to implement principle 6 of the Data Ethics Framework for the public sector.

Your work must be accountable, which is only possible if people are aware of and can understand your work.

Being open about your work is critical to helping to make better use of data across government. When discussing your work openly, be transparent about the tools, data, algorithms and the user need (unless there are reasons not to such as fraud or counter-terrorism). Provide your explanations in plain English. Check within your department before speaking about your work openly.

Sharing your work builds trust in its quality by allowing other practitioners to learn from your methods. It can also inspire policy makers to use data science.

Peer review is an essential part of quality assurance. Get feedback from your own team or organisation's data science function. If you're working alone, you may need to look elsewhere to receive appropriate scrutiny. The GDS Data Science Community is a good way of receiving feedback from peers.

Opening up your code within public repositories on Github facilitates the use of free features such as automated testing and code coverage measurement. This encourages continuous improvement of the code and your own coding skills.

Feedback tells you what people care about and judge acceptable. Having evidence of this is useful when determining whether your approaches are proportionate (Principle 3). This is useful for your own project, but can also be shared with others looking for advice.

Good practice for making your work transparent

Documenting your work clearly is an essential part of working in the open and being accountable. Follow good development practices to make sure your work is easy to understand. This includes clearly explaining the caveats, assumptions and uncertainties in your work.

Your technology choices should support coding in the open where possible. Read GDS guidance on when code should be open or closed, how to keep open code secure and how to make your code reusable.

Discussing your work openly at events, blogging and documenting work clearly on Github helps to:

- build trust in its quality
- facilitate peer review
- get feedback

Sharing your data

If data is non-sensitive and non-personal, you should make it open and assign it a digital object identifier (DOI). For example, scientists share data when publishing a paper on Figshare and Datadryad. This gives others access to the data and the code, so the analysis can be reproduced. You can also publish data on Find open data and the UK Data Archive.

When sharing personal data, you must comply with the ICO data sharing code of practice, which will be updated for the new Data Protection Act 2018.

When accessing and sharing data under powers in Part 5 of the Digital Economy Act 2017, you must follow the relevant Codes of Practice.

Share your models for algorithmic accountability

Developed data science tools should be made available for scrutiny wherever possible.

There are 2 main types of algorithms used in data science.

The first is the algorithmic methodology used to train a model. It's often more useful and clear to share a document describing the analytical process than the code.

The second is the trained model itself (the result of applying the methodology to the dataset). Releasing this model allows others to scrutinise and test it, and may highlight issues that you can fix as part of your continual improvement.

When sharing models it's important that it does not endanger either the:

- privacy of those whose data was used to train it
- integrity of the task being undertaken

Even if the model cannot be released publicly, you may be able to release metadata about the model on a continual basis, like its performance on certain datasets. If your data science application is very sensitive, you could arrange for selected external bodies, approved by your organisation, to examine the model itself in a controlled context to provide feedback. This could be expertise from another government department, academia or public body.

Transparency and interpretability of algorithms

The more complex data science tools become, the more difficult it may be to understand or explain the decision-making process. This is a critical issue to consider when carrying out data science or any analysis in government. It is essential that government policy be based on interpretable evidence in order to provide accountability for a policy outcome.

You should also plan how you will explain your work to others, ensuring your approach can be held

to account.

7. Embed data use responsibly

How to implement principle 7 of the Data Ethics Framework for the public sector.

Policy decisions informed by data can have significant social impact.

Put appropriate long-term processes in place to monitor policies developed using data analysis. This applies to both traditional regression and more advanced techniques like machine learning.

For any data use you need to determine:

- the implementation plan, including ongoing agile iteration of your live service
- sustainable and ongoing evaluation methods
- the method for feedback into the data model, including when it's necessary to retrain using newly collected data

This applies to both one-off projects and ongoing operationalised models, i.e. models which are used in running of government internal or public facing services.

Making policy with data

Monitoring and evaluating policies is an established process in government. The HM Treasury Magenta Book gives guidance on evaluation.

You should follow the ROAMEF (rationale, objectives, appraisal, monitoring, evaluation, feedback) cycle for policy development and tailor it to fit data projects. Depending on your intended data use and objective, data science could strengthen each stage of the classic ROAMEF cycle as detailed in Chapter 7 of the Magenta Book.

Long-term collaboration across disciplines (practitioners, service design, policy and operational staff) should ensure this cycle is managed appropriately and all potential factors impacting the model or insight are considered.

Designing or delivering services with data

Make sure you iterate and improve frequently, to meet point 5 of the Digital Service Standard. This means your process can be changed in response to policy changes or other factors.

Ensuring appropriate knowledge and support when deploying to non-specialists

To make sure data science is embedded responsibly, develop a plan to manage appropriate use of an operationalised model by non-specialists.

Operational or service staff must have sufficient knowledge or training to understand how to use a new system including a full interpretation of outputs. You must provide them with sufficient support to avoid the misuse of models. They must have an easy way to report any suspected erroneous behaviour. The development team must make sure all of this information is provided.

How the efficacy of the model will be monitored once deployed

You should be confident that your model won't fail after being operationalised. To ensure this, development teams must advise on ongoing evaluation of models once deployed.

Any model may produce undesirable results under unusual conditions. Use the Futures Toolkit methods to consider potential risks and disruptions which could alter the performance of your model.

Who will be responsible for ongoing maintenance

HM Treasury's Aqua Book recommends that most models have a Senior Responsible Officer (SRO). Models are likely to be handed from a technical team to a responsible policy or service delivery team. There should be ongoing communication between teams to manage the use of the model.

When to retrain or redesign a predictive model

When thinking about how regularly a model should be reviewed, consider:

- how quickly the model or service affected by your work will scale
- the likely impact of your model on citizens

Your model will have been trained on historic data. Once the model is deployed, it is reasonable to expect it to change how a service or policy is delivered. This will lead to new outcomes and new data, different from the historic data the model was developed with.

Factors external to the policy may also change the data gathered over time. Policy and practitioner teams should work together to identify:

- which factors outside the policy are most likely to have an impact
- how often models will need to be retrained to account for this or redesigned based on the impact

Monitoring personalisation or tailored service delivery

One potential opportunity of machine learning is tailoring services for individuals or groups to make

them more effective (called personalisation). This may also fall under profiling (GDPR Article 22) and you will therefore need a lawful basis to do so.

If fewer choices are presented as a result of personalisation, monitor your model continuously to make sure it's still personalising effectively without negative consequences. You should also be prepared to be transparent about this process as it is essential you can explain clearly how any algorithm is personalising information.

Algorithms in decision making

Under Article 22 of the GDPR you can only use solely automated processes to make decisions with legal or similarly significant effect about an individual if you:

- have a specific lawful basis to do so
- follow the data protection safeguards laid out in Article 22 of the GDPR and Section 14 of the DPA 2018

When considering the role of algorithms in decision making, it's important to not only consider a final decision but any potential automated decisions which played an important role in forming the final decision-making process.

Read the Article 29 Working Party guidance on automated individual decision making and profiling for more information.

Data Ethics Workbook

How to follow the Data Ethics Framework which guides data use in government and the wider public sector.

Getting started

Before you start a new data project or workstream, the Data Ethics Workbook should be completed collectively by practitioners, data governance or information assurance specialists, and subject matter experts like service staff or policy professionals.

Teams must decide at the start how often they will revisit these questions. This will depend on the length and scale of the project, service or policy.

These questions will help you determine how you can align your work with the Data Ethics Framework principles. Answering the workbook questions should inform the design of an implementation plan for producing high quality results and mitigating risks.

The Data Ethics Workbook is available in html, ods and pdf format.

Data Ethics Framework						
	0	1	2	3	4	5
1. Start with clear user need and public benefit	User need is not well defined					User need is clearly defined
Description of the user need with supporting evidence						
2. Be aware of relevant legislation and codes of practice	Needs clarification or expert input					Relevant laws are well understood
List the pieces of legislation, codes of practice and guidance that apply to your project.						
3. Use data that is proportionate to the user need	Reuse not proportionate					Reuse of data is clearly proportionate to achieve user need
Describe how the data being used is proportional to the user need						
4. Understand the limitations of the data	Unreliable, unsuitable data					Data is representative and accurate
Identify the potential limitations of the data source(s) and how they are being mitigated						
5. Use robust practices and work within your skillset	Needs further expert input					Methodologies clearly designed and understood
Explain the relevant expertise and approaches that are being employed to maximise the efficacy of the project						
6. Make your work transparent and be accountable	No scrutiny or peer review available					Oversight built in through life cycle of project
Describe how you have considered making your work transparent and accountable						
7. Embed data use responsibly	No ongoing plan determined					Evaluation plan developed and resource in place to deliver it
Describe the steps taken to ensure any new model, policy or service is managed responsibly						

Questions for principle 1 - Start with clear user need and public benefit

Describe the user need.

- Does everyone in the team understand the user need?
- How does this benefit the public?
- What would be the harm in not using data science - what needs might not be met?
- Do you have supporting evidence for the approach being likely to meet a user need or provide public benefit?

Principle 1: Start with a clear user need and public benefit	
To consider:	
Describe the user need:	
Does everyone in the team understand the user need?	
How does this benefit the public?	
What would be the harm in not using data science - what needs might not be met?	
Do you have supporting evidence for the approach being likely to meet a user need or provide public benefit?	

Questions for principle 2 - Be aware of relevant legislation and codes of practice

List the pieces of legislation, codes of practice and guidance that apply to your project.

- Do all team members understand how relevant laws apply to the project?
- If necessary, have you consulted with relevant experts?
- Have you spoken to your information assurance team?
- If using personal data, do you understand your obligations under data protection legislation?
- Do you have plans in place to handle any potential security breach?

Principle 2: Be aware of relevant legislation and codes of practice	
To consider:	
List the pieces of legislation, codes of practice and guidance that apply to your project:	
Do all team members understand how relevant laws apply to the project?	
If necessary, have you consulted with relevant experts?	
Have you spoken to your information assurance team?	
If using personal data, do you understand obligations under data protection legislation?	

Questions for principle 3 - Use data that is proportionate to the user need

Describe how the data being used is proportionate to the user need.

- Could you clearly explain why you need to use this data to members of the public?
- Does this use of data interfere with the rights of individuals?
- If yes, is there a less intrusive way of achieving the objective?
- Is there a fair balance between the rights of individuals and the interests of the community?
- Has the data you're using been specifically provided for your analysis?

- By using data that the public has freely volunteered, would your project jeopardise people providing this again in the future?
- How can you meet the project aim using the minimum personal data possible?
- Is there a way to achieve the same aim with less identifiable data?
- Can you use synthetic data?
- If using personal data is unavoidable, have you answered the questions for determining proportionality?
- If using personal data identifying individuals, what measures are in place to control access? How widely are you searching personal data?

Principle 3: Use data that is proportionate to the user need	
To consider:	
If using personal data, have you answered the questions for determining proportionality? You must include evidence to support any decision.	
If using personal data, what measures are in place to control access? How widely are you searching personal data?	
How can you meet the project aim using the minimum personal data possible?	
Is there a way to achieve the same aim with less identifiable data?	
Can you use synthetic data?	
Has the data being used been provided for your analysis?	
By using data that the public have freely volunteered, would your project jeopardise people providing this again in the future?	
Could you clearly explain why you need to use that data to members of the public?	
Is there a fair balance between the rights of individuals and the interests of the community?	

Questions for principle 4 - Understand the limitations of the data

Identify the potential limitations of the data source(s) and how they are being mitigated.

- What data source(s) is being used?
- Are all metadata and field names clearly understood?
- What processes do you have in place to ensure and maintain data integrity?
- Is there a plan in place to identify errors and biases?
- What are the caveats?
- How will the caveats be taken into account for any future policy or service which uses this work as an evidence base?

Principle 4: Understand the limitations of the data	
To consider:	
Identify the potential limitations of the data source(s) and how they are being mitigated:	
What data source(s) is being used?	
Are all metadata and field names clearly understood?	
What processes do you have in place to ensure and maintain data integrity?	
Is there a plan in place to identify errors and biases?	
What are the caveats?	

Questions for principle 5 - Ensure robust practices and work within your skillset

Explain the relevant expertise and approaches that are being employed to maximise the efficacy of the project.

- Describe the disciplines involved and why.
- Are there expertise that the project requires that you don't currently have?
- Have you designed the approach with the policy team or a subject matter expert?
- Has all subject matter context, from policy experts or otherwise, been taken into account when determining the appropriate loss function for the model?
- If necessary, how can you (or external scrutiny) check that the algorithm is achieving the right output decision when new data is added?
- How has reproducibility been ensured? Could another analyst repeat your procedure based on your documentation?
- How confident are you that the algorithm is robust, and that any assumptions are met?
- What is the quality of the model outputs, and how does this stack up against the project objectives?
- If using data about people, is it possible that a data science technique is basing analysis on proxies for protected variables which could lead to a discriminatory policy decision?

Principle 5: Use robust practices and work within your skillset	
To consider:	
Explain the relevant expertise and approaches that are being employed to maximise the efficacy of the project	
Describe the disciplines involved and why.	
Are there expertise that the project requires that you don't currently have?	
Have you designed the approach with a policy team or subject matter expert(s)?	
Has all subject matter context, from policy experts or otherwise, been taken into account when determining the appropriate loss function for the model?	
If necessary, how can you (or with external scrutiny) check that the algorithm is achieving the right output decision when new data is added?	
How has reproducibility been ensured? Could another analyst repeat your procedure based on your documentation?	
How confident are you that the algorithm is robust, and that any assumptions are met?	
What is the quality of the model outputs, and how does this stack up against the project objectives?	
If using data about people, is it possible that a data science technique is basing analysis on proxies for protected variables which could lead to a discriminatory policy decision?	

Questions for principle 6 - Make your work transparent and be accountable

Describe how you have considered making your work transparent and your team accountable.

- Have you spoken to your organisation to find out if you can speak about your project openly?
- Have you considered how both internal and external engagement could benefit your project?
- How interpretable are the outputs of your work?
- How are you explaining how approaches were designed in plain English to other

practitioners, policy makers and if appropriate, the public?

- Can you openly publish your methodology, metadata about your model, and/or the model itself e.g. on Github?
- Can you get peers to review your Pull Requests?

Principle 6: Make your work transparent and be accountable	
To consider:	
Describe how you have considered making your work transparent and accountable	
Have you spoken to your organisation to find out if you can speak about your project openly?	
Have you considered how both internal and external engagement could benefit your project?	
How interpretable are the outputs of your work?	
How are you explaining how approaches were designed in plain English to other practitioners, policy makers and if appropriate, the public?	
Can you openly publish your methodology, metadata about your model, and/or the model itself e.g. on Github?	
Can you get peers to review your Pull Requests?	

Questions for principle 7 - Embed data use responsibly

Describe the steps taken to ensure any insight is managed responsibly.

- How many people will be affected by the new model, insight or service?
- Who are the users of the insight, model, or new service?
- Do users have the appropriate support and training to maintain the new technology?
- Have future events been planned for?
- Is your implementation plan correlated with the impact of a particular model?
- How often will you report on these plans to Senior Responsible Officers?

Principle 7: Embed data use responsibly	
To consider:	
Describe the steps taken to ensure any insight is managed responsibly:	
How many people will be affected by the new model, insight or service?	
Who are the users of the insight, model, or new service?	
Do users have the appropriate support and training to maintain the new technology?	
Have future events been planned for?	
Is your implementation plan correlated with the impact of a particular model?	
How often will you report on these plans to senior reporting officers?	

Data Ethics Workbook: working with suppliers

If you are procuring analytics rather than building them within your team you must still ensure the software and contract are conducive to the principles of this framework. In order to determine this, you should discuss the following questions with your provider.

- Can you answer the questions in the Data Ethics Workbook?
- What was the original intended use of the software?
- What are the contractual arrangements for any derived data produced through the use of the procured software?
- Can you scrutinise the performance of the software free from the constraints of intellectual property restrictions?
- How has the algorithm been trained?
    - What data was used to train the algorithm?
    - How interpretable is the algorithm?
    - How has the algorithm been tested for different failure modes relevant to the intended task?
- What are the plans for re-assessing the algorithm's performance at set times?
    - Are assessment methods, e.g. for bias, compatible with the software being procured?
    - Do assessment methods require the support of the provider, or can they be carried out independently?

Procurement	
To consider:	
Describe how the technology you are procuring is conducive to the principles of this framework:	
Can you answer the questions in the Data Ethics Workbook?	
What was the original intended use of the software?	
How has the algorithm been trained?	
What data was used to train the algorithm?	
How interpretable is the algorithm?	
How has the algorithm been tested for different failure modes relevant to the intended task?	
What are the plans for re-assessing the algorithm's performance at set times?	
Are assessment methods, e.g. for bias, compatible with the software being procured?	
Do assessment methods require the support of the provider, or can they be carried out independently?	
Can you scrutinise the performance of the software free from the constraints of intellectual property restrictions?	

Information assurance and suppliers

Working with commercial suppliers is another activity that your information assurance specialists will advise on. Whether the supplier is developing, managing or processing information, the information assurance specialists will work with the supplier to understand their approach to information security and management. If processing personal data, you must ensure the supplier's process allows you to adhere to Articles 13 and 14 of the GDPR and Section 93 of the DPA 2018. If suppliers are 'data processors', the data controller needs to comply with Article 28 of the GDPR. This requires the controller to only use processors who provide sufficient guarantees for people's data and this should be subject to a contract between the controller and the processor.