



Government
Digital Service



Data Ethics Framework

© Crown copyright 2020

This publication is licensed under the terms of the Open Government Licence v3.0 except where otherwise stated. To view this licence, visit nationalarchives.gov.uk/doc/open-government-licence/version/3

This publication is also available on our website at www.gov.uk/government/publications/data-ethics-framework

Data Ethics Framework

What is it for?

The Data Ethics Framework guides appropriate and responsible data use in government and the wider public sector. It helps public servants understand ethical considerations, address these within their projects, and encourages responsible innovation.

Who is it for?

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 to produce data-informed insight.

How to use it?

Teams should work through the framework together throughout the process of planning, implementing, and evaluating a new project. 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.

Questions in each section will guide you through various ethical considerations in your project. Record your answers to create a data ethics assessment. The self-assessment scoring scale is designed to help you identify which aspects of your project need further work.

Structure

The framework is split into **overarching principles** and **specific actions**. Overarching principles are applicable throughout the entire process and underpin all actions and all aspects of the project. Specific actions will guide you through different stages of the project and provide practical considerations.

In addition, the framework provides specific actions you can take at each stage of the project to advance **transparency**, **accountability**, and **fairness**. These are marked within each section.



Overarching principles



Transparency

Transparency means that your actions, processes and data are made open to inspection by publishing information about the project in a complete, open, understandable, easily-accessible, and free format. In your work with and on data and AI, use the available guidance, e.g. the Open Government Playbook, to ensure transparency throughout the entirety of your process.

Transparency					
Score					
0	1	2	3	4	5
Information about the project, its methods, and outcomes is not publicly available					Information about the project, its methods, and outcomes is widely available to the public

Key:



Start



During



After



Accountability

Accountability means that there are effective governance and oversight mechanisms for any project. Public accountability means that the public or its representatives are able to exercise effective oversight and control over the decisions and actions taken by the government and its officials, in order to guarantee that government initiatives meet their stated objectives and respond to the needs of the communities they are designed to benefit.

Accountability					
Score					
0	1	2	3	4	5
Mechanisms for scrutiny, governance, or peer review for the project haven't been established					Long-term oversight and public scrutiny mechanisms are built into the project cycle

Key:



Start



During



After



Fairness

Fairness — It is crucial to eliminate your project's potential to have unintended discriminatory effects on individuals and social groups. You should aim to mitigate biases which may influence your model's outcome and ensure that the project and its outcomes respect the dignity of individuals, are just, non-discriminatory, and consistent with the public interest, including human rights and democratic values.

You can read more about fairness and its different types in the [Understanding artificial intelligence ethics and safety](#) guide developed by the Government Digital Service and the Office for Artificial Intelligence.

You can read more about the standards regarding human rights in AI and machine learning in the [Toronto Declaration](#) developed by Access Now and Amnesty International.

Fairness					
Score					
0	1	2	3	4	5
There is a significant risk that the project will result in harm or detrimental and discriminatory effects for the public or certain groups					The project promotes just and equitable outcomes, has negligible detrimental effects, and is aligned with human rights considerations

Key:



Start



During



After

Self-assessment tables for the 3 overarching principles.

Please use the following self assessment to mark, track and share your overall progress over time against the following 3 overarching principles.

Key:



Start



During



After

		Score					
Principles		0	1	2	3	4	5
	Transparency						
	Accountability						
	Fairness						

Self-assessment tables for the 5 specific actions.

Please use the following self assessment to mark, track and share your overall progress over time against the following 5 specific actions.

Key:



Start



During



After

		Score					
Specific actions		0	1	2	3	4	5
1	Define public benefit and user need						
2	Involve diverse expertise						
3	Comply with the law						
4	Check the quality and limitations of the data						
4.1	Check the quality and limitations of the model						
5	Evaluate and consider wider policy implications						

Self-assessment tables for the 3 overarching principles.

Please use the following self assessment to mark, track and share your overall progress over time against the following 3 overarching principles.

Key:		
		
Start	During	After

		Score					
Principles		0	1	2	3	4	5
	Transparency						
	Accountability						
	Fairness						

Self-assessment tables for the 5 specific actions.

Please use the following self assessment to mark, track and share your overall progress over time against the following 5 specific actions.

Key:



Start



During




After

		Score					
	Specific actions	0	1	2	3	4	5
1	Define public benefit and user need						
2	Involve diverse expertise						
3	Comply with the law						
4	Check the quality and limitations of the data						
4.1	Check the quality and limitations of the model						
5	Evaluate and consider wider policy implications						

Define and understand public benefit and user need

When starting a public sector data project, you must have a clear articulation of its purpose.

This includes having clarity on what public benefit the project is trying to achieve and what are the needs of the people who will be using the service or will be most directly affected by it.



Define and understand public benefit and user need

When starting a public sector data project...

Using the circles move them to where your project is on the spectrum. Do this at the start of a project, during a project and after a project has completed. This will help you track, communicate and share your progress.

Define public benefit and user need					
Score					
0	1	2	3	4	5
Public benefit and user need are not clearly defined or understood					Public benefit and user need are well defined and understood by all team members

Key:



Start



During



After



Define and understand public benefit and user need

Public Benefit

1.1


Understand the wider public benefit

- What are direct benefits for individuals in this project? (e.g. saving time when applying for a government service)
- How does the project deliver positive social outcomes for the wider public?
- How can you measure and communicate the benefits of this project to the public?
- What are the groups that would be disadvantaged by the project/ that would not benefit from the project? What can you do about this?

1.2

Understand unintended and or negative consequences of your project

- What would be the harm in not using data? What social outcomes might not be met?
- What are the potential risks or negative consequences of the project versus the risk in not proceeding with the project?
- Could the misuse of the data/algorithm or poor design of the project contribute to reinforcing social and ethical problems and inequalities?
- What kind of mechanisms can you put in place to prevent this from happening?
- What specific groups benefit from the project? What groups can be denied opportunities/face negative consequences because of the project?



User need

1.7

Ensure there is a clear articulation of the problem before you start the project.

Describe the user need in your project:

- As a ...
- I need/want/expect to... [what does the user want to do?]
- So that... [why does the user want to do this?]

Example 1: the [Register to vote service's](#) user need:

- As a UK resident
- I want to get my details on the online electoral register
- So that I can vote

Example 2: user need when building a platform for fire safety checks:

- 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

For various user 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

1.6

Understand the user need

User needs' are the needs that a user has of a service, and which that service must satisfy for the user to get the right outcome for them. For more information about the user need, consult the [GDS Service Manual](#).

Example:

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

1.8

Check if everyone in your team understands the user need and how using data can help.

Does everyone in your team understand the user need?

- 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.

1.9

Repeatedly revisit your user need throughout the project.

Consider these questions as the project evolves:

- What is the overall problem you're trying to solve?
- Who are the users of this data process or analysis?
- What needs do they have?

Involve diverse expertise

Working in diverse, multidisciplinary teams with wide ranging skill sets contributes to the success of any data or tech project. If you do not have the sufficient skills or experience, you should involve others from your team or wider network with the right expertise.

2

Involve diverse expertise

Involve diverse expertise					
Score					
0	1	2	3	4	5
The project team is homogeneous and there is little expert input		The project team is diverse, multidisciplinary, with extensive expert input			

Key:

 Start
  During
  After

2

Involve diverse expertise

2.1


Get the right expertise

- Beyond data scientists, who are the relevant policy experts and practitioners in your team?
- What other disciplines and subject matter experts need to be involved? What steps have been taken to involve them?
- What should their roles be? Have you defined who does what in the project?
- Ask your team and experts if you have the right data for your research questions. Get their help in matching the dataset to the problem.

2.2

Ensure diversity within your team

- How have you ensured diversity in your team? Having a diverse team helps prevent biases and encourages more creativity and diversity of thought.
- Avoid forming homogenous teams, embrace diversity of lived experiences of people from different backgrounds. If you find yourself in a homogenous team, challenge it.



Comply with the law

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

Comply with the law

Comply with the law					
Score					
0	1	2	3	4	5
There is little clarity on legal requirements for the project					Relevant legal requirements have been met, compulsory assessments completed, legal experts have been consulted

Key:

 Start
  During
  After

3

Comply with the law

3.1

Get legal advice

- Have you spoken to a legal adviser within your organisation?
- Have you spoken to your information assurance team?
- Have you consulted your organisation's Data Protection Officer when doing a DPIA?
- What legal advice have you received?

3.2

It is your duty and obligation to obey the law in any data projects. You must ensure the project's compliance with GDPR and DPA 2018

If you are using personal data, you must comply with the principles of the [EU General Data Protection Regulation \(GDPR\)](#) and [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](#)).

3.3

Data protection by design and DPIA

Data protection by design and by default is a legal requirement under the GDPR (see [Article 25](#)).

It is a legal obligation under [Article 35 of the GDPR](#) to complete a [data protection impact assessments \(DPIA\)](#) (also known as privacy impact assessments) when there's likely to be high risk to people's rights, particularly when using new technologies. It is good practice to do a DPIA for any use of personal data.

3.4

Accountability

An important aspect of complying with data protection law, is being able to demonstrate what measures you are taking to ensure everything is documented, as seen in 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.

3.5

Transparency

Publish your DPIA and other related documents

3.6

Ensure the project's compliance with the Equality Act 2010

Data analysis or automated decision making must not result in outcomes that lead to discrimination as defined in the [Equality Act 2010](#).

- How can you demonstrate that your project meets the Public Sector Equality Duty?
- What was the result of the Equality Impact Assessment of the project?

3.7

Ensure effective governance of your data

Organisations have a responsibility to keep both [personal data](#) and non-personal data secure.

- How have you ensured that the project is compliant with data governance policies within your organisation?

3.8

Ensure your project's compliance with any additional regulations

Consider [additional relevant legislation and codes of practice](#).

Review the quality and limitations of the data

Insights from new technology are only as good as the data and practices used to create them. You must ensure that the data for the project is accurate, representative, proportionally used, of good quality, and that you are able to explain its limitations.

Review the quality and limitations of the data

Review the quality and limitations of the data					
Score					
0	1	2	3	4	5
Data for the project is of bad quality, unsuitable, unreliable, not-representative					Data used in the project is representative, proportionally used, accurate, and of good quality
The model is not reproducible and is likely to produce invalid outputs					The model is reproducible and able to produce valid outputs

Key:



Start



During



After

4

Review the quality and limitations of the data

4.1

Data source

- What data source(s) is being used?
- Are individuals and/or organisations providing the data aware of how it will be used? If the user is repurposing data for analysis without individual consent, how have you ensured that the new purpose is compatible with the original reason for collection (Article 6 (4) GDPR)?
- Are all metadata and field names clearly understood?
 - Do you understand how the data for the project is generated? Remember that depending on where the data came from, the field may not represent what the field name/metadata indicates.
- What processes do you have in place to ensure and maintain data integrity?
- 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?
- Would using synthetic data be appropriate for the project? Synthetic data is entirely fabricated or abstracted from real data through various processes, e.g. anonymisation or record switching. It is often created with specific features to test or train an algorithm.

4.3

Bias in data

- How has the data being used to train a model been assessed for potential bias? You should consider:
 - Whether the data might (accurately) reflect biased historical practice that you do not want to replicate in the model (historical bias)
 - The data might be a biased misrepresentation of historical practice, e.g. because only certain categories of data were properly recorded in a format accessible to the project (selection bias)
- If using data about people, is it possible that your model or analysis may be identifying proxy variables for protected characteristics which could lead to a discriminatory outcome? Such proxy variables can potentially be a cause of indirect discrimination; you should consider whether the use of these variables is appropriate in the context of your service (i.e. is there a reasonable causal link between the proxy variable and the outcome you're trying to measure?; do you assess this to be a proportionate means to achieve a legitimate aim in accordance with the Equality Act 2010?)
- What measures have you taken to mitigate bias?

4.2

Determining proportionality

You must use the minimum data necessary to achieve your desired outcome ([Article 5\(1\)\(c\)](#) of the GDPR). Personal data should be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

- 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, e.g. pseudonymised data?
- If using [personal data identifying individuals](#), what measures are in place to control access?
- Is the input data suitable and necessary to achieve the aim?
- Would the proposed use of data be deemed inappropriate by those who provided the data (individuals and/or organisations)?
- 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?
- How can you 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?

4.4

Data anonymisation

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](#).

If the assumption is that data in the project is anonymised, consider the following:

- How can you demonstrate that the data has been de-identified to the greatest degree possible?
- Can the data be matched to any other datasets that will make individuals easily identifiable? What measures have you taken to mitigate this? Have you considered any [determined intruder testing](#)?

4.5

Robust practices

- If necessary, how can you (or external scrutiny) validate that the algorithm is achieving the correct output decision when new data is added?
- How can you demonstrate that you have designed the project for reproducibility?
 - Could another analyst repeat your procedure based on your documentation? Have they tried?
 - Have you followed the 3 requirements for reproducibility? (Applying the same logic (code, model or process), to the same data, in the same environment).
 - How have you ensured that high quality documentation will be kept?
 - [Have you considered Reproducible Analytical Pipelines? \(RAP\)](#)
- 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?

Make your data open and shareable whenever possible

If data is non-sensitive, non-personal, and if Data Sharing Agreements with the supplier allow it, you should make the data 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](#).

Share your models - developed data science tools should be made available for scrutiny wherever possible.

- Can you openly publish your methodology, metadata about your model, and/or the model itself e.g. on Github?

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.

4.8

How to ensure transparency of sensitive models

- How are you planning to inform the public about the model?
- 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.

4.9

Explainability

Explainability is the extent to which the workings in a machine learning algorithm can be explained in human terms. It means expanding on the transparency of what variables are used to provide information on how the algorithm came to give an output, and how changing the inputs can change the output.

- Explain what your project does and how it was designed in plain language to a non-expert audience.
- Describe the process and the aim of your algorithm, as well as what variables are used for what outcomes without using technical terms.
- Make this explanation publicly available (e.g. on GitHub, blogs, or gov.uk).

Evaluate and consider wider policy implications

It is essential that there is a plan to continuously evaluate if 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 and effective accountability mechanisms.

5

Evaluate and consider wider policy implications

Evaluate and consider wider policy implications					
Score					
0	1	2	3	4	5
There are no long-term evaluation and maintenance structures in place					Continuous evaluation and long-term maintenance structures are in place

Key:



Start



During



After

5

Evaluate and consider wider policy implications

5.1

Evaluate the project

Continuous evaluation - ask yourself and the team:

- At the beginning of the project: 'are we doing the right thing?'
- During the project: 'have we designed it well?'
- After the project: 'is it still doing the right thing we designed it for?'
- How have you evaluated the project? Evaluation techniques you might use include holding retrospective roundtables at the end of the project; inviting an external expert or a 'critical friend' from a different team to observe and evaluate the project; request external consultations or audits.

5.2

Repeatedly revisit the user need and public benefit throughout the project

- How has the user need changed?
- Is the project still benefiting the public?
- Has there been a change of circumstances that might have affected the initial understanding of the public benefit in the project? If so, how can you adjust it?
- How sufficient is the current human oversight of the automated project?
- If any unanticipated harms emerged during the project, how have they been mitigated?

5.3

Check how your project influences policy

- How accurately are the insights from the project used in the practical policy context?
- How have you ensured if policymakers and secondary users of the tool fully understand its purpose and structure?

5.4

Ensure there are skills, training, maintenance for longevity of the project

- How have you ensured that the users have the appropriate support and training to maintain the new technology?
- What is the longevity of the project?
- How have you checked if the users understand the software they need to maintain the project? Have they been appropriately trained?

5.5

Accountability structures

- What are the governance structures in place to ensure a safe and sustainable implementation of the project?
- How often will you update the board/governing bodies?

5.6

Public scrutiny

- Do members of the public/ end users of the project have the ability to raise concerns and place complaints about the project? If yes, how? If not, why?
- What channels have you established for public engagement and scrutiny throughout the duration of the project?

5.7

Share your learnings

- How have you documented and shared the progress and case studies from your project with peers and stakeholders?

Next steps: if you have scored a 3 or less in any of the principles, this could indicate the need for additional checks and potential changes to make your project more ethical. Please explain the reason for the score and consult the outcome with your team leader/organisational ethics board/data ethics lead to advise on the specific next steps to improve the ethical standards of your project.