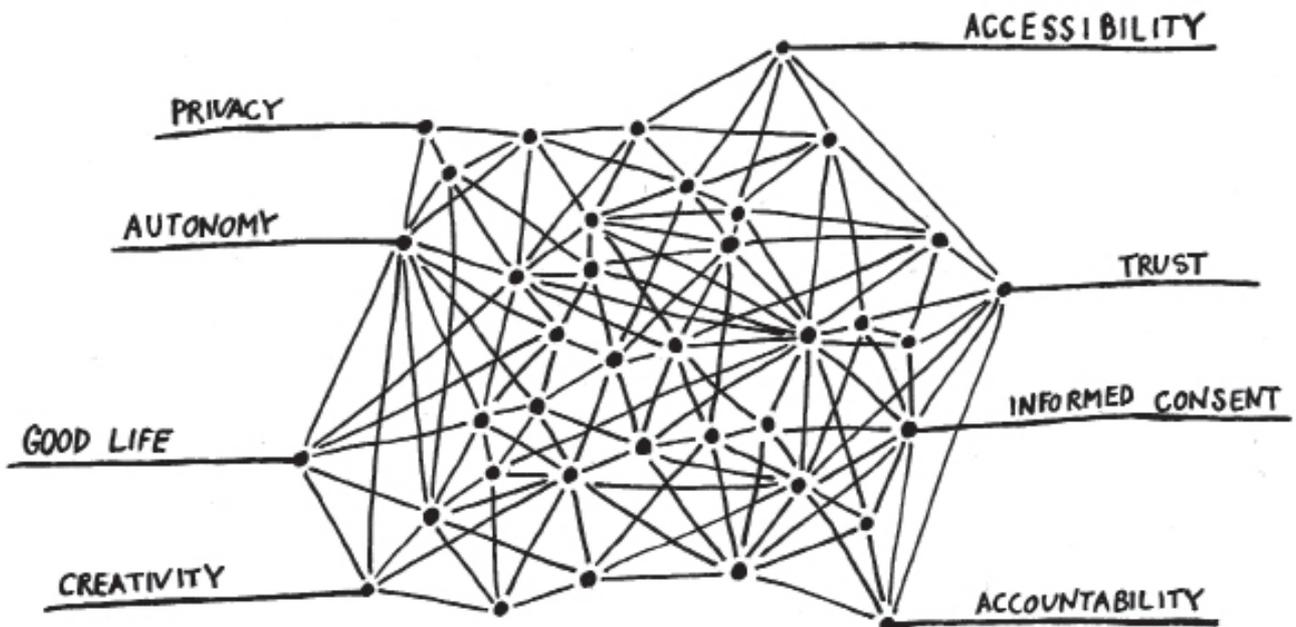


D *a t a*

E *t h i c s*

D *e c i s i o n*

A *i d*



Handbook

Assessing **ethical issues** with regard
to governmental **data projects**

Disclaimer

The 'Data Ethics Decision Aid' is developed by Data School, Utrecht University.
© Data School, Utrecht University
2026

DEDA - Version 4.0
February 2026
For updates, see:

deda.dataschool.nl

Neither Utrecht University, nor Data School, is liable for (possible) damage in connection with, or arising out of, the use of material or content of the 'Data Ethics Decision Aid'.

Utrecht, 2026.

As DEDA is used, we learn to improve it. These improvements will be implemented in future versions. If you are using DEDA and have any comments, please do not hesitate to share them with us. You can always mail to::

dataschool@uu.nl

Concept :
Mirko Tobias Schäfer, Ph.D.
Aline Franzke, MA

Editors:
Julia Straatman, MA
Iris Muis, MA
Hajo Bijleveld, BA
Ydo Bergstra, BSc
Lysa Ngouateu, BSc
Mirko Tobias Schäfer, Ph.D.
Bart Anthony Kamphorst, Ph.D.

Illustration:
Sammy Hemerik

© Data School



**Utrecht
University**

DATA SCHOOL

Table of contents

Introduction

6 - 9

Questions

10 - 31

TECHNICAL CONSIDERATIONS

10 - 17



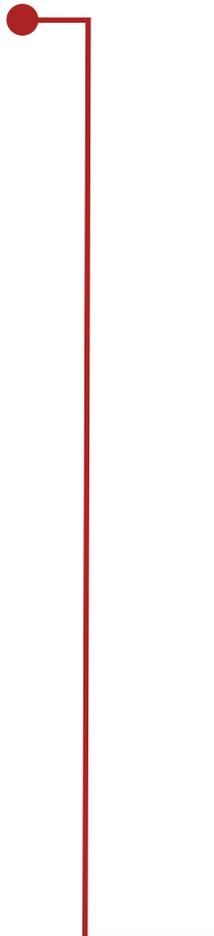
SOCIETAL CONSIDERATIONS

18 - 25



OUTCOME

26 - 31



DATA SOURCES	12
ALGORITHMS	13
ANONYMIZATION	15
VISUALISATION	16
ACCESS AND REUSE	17
PRIVACY	20
BIAS	21
TRANSPARENT COMMUNICATION	23
ORGANIZATION	24
FUTURE SCENARIOS	25
VALUES	28
CONCLUSION	29
NEXT STEPS	30
FURTHER INFORMATION	31

Introduction

Why DEDA?

Data, algorithms, and AI promise great benefits for the public sector, from facilitating smart cities to increasing economic prosperity. However, these opportunities are not without difficulties. These challenges are easy to ignore, but in the long term, they can cause good intentions to lead to bad outcomes. For example, problems can arise around liability when something goes wrong - which can lead to a lawsuit - or a project may conflict with the guidelines of good governance.

Companies and governments are increasingly criticised for the way they use data and algorithms. In response to this, laws and regulations have been amended and the use of AI is regulated by the European AI Regulation. Increased fines for violating privacy are an example of the EU's attempts to enforce the responsible use of personal information.

In addition, the regulations in the AI Regulation are intended to ensure that AI systems used in the EU are safe, transparent, traceable, non-discriminatory, and environmentally friendly. Algorithms that are not classified as AI by the AI

Act can be just as impactful.

This is because there are all kinds of problems that can arise from data projects.

That is why the Dutch government's Algorithm Register uses the term "impactful algorithms." How these relate to the definition of "high-risk AI" included in the AI Regulation can be seen in Figure 1.

For example, datasets may have questionable origins, or they might be taken out of context. There may also be bias in the datasets, models, and algorithms. In addition, questions might arise regarding (potential) conflicts of interest between commercial companies and public institutions. Finally, there can be a lack of critical evaluation of the social impact of data-driven policies.

The Data Ethics Decision Aid (DEDA) is a tool that helps identify ethical issues and develop a sense for detecting value conflicts within a data project or AI project. DEDA can also help provide insight into the (public) values that are affected, or document an ethical decision-making process.

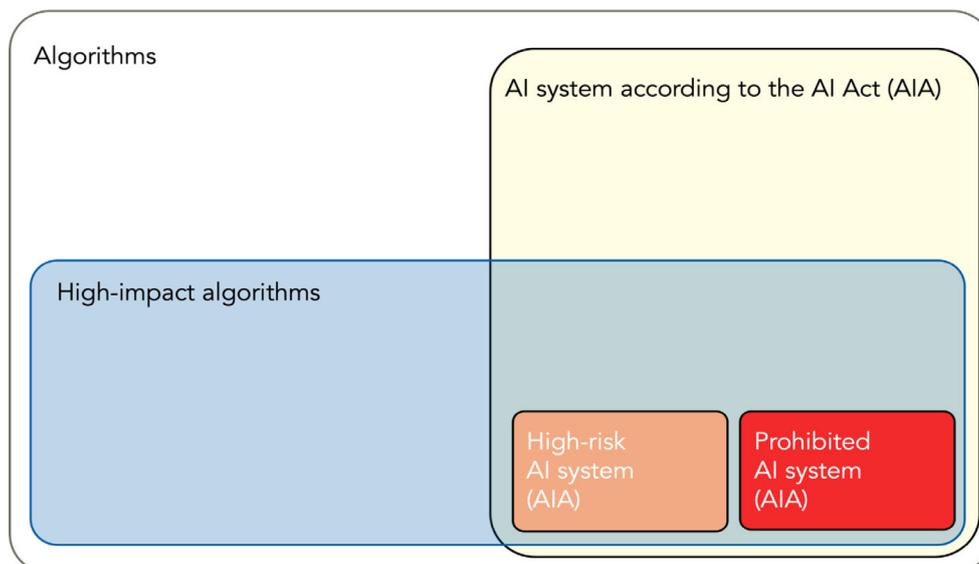


Figure 1: White paper Implementation of the AI Act – Definition of an AI system, Algorithm Audit (Feb 2025), retrieved from <https://algorithmaudit.eu/pdf-files/technical-tools/AI-AQT/20250228%20White%20paper%20Implementation%20AI%20>

It enables organizations to determine whether a data or AI project fits the ambitions, goals and context of the organization, or whether it undermines them.

Additionally, the process provides insight into the necessary capabilities and preconditions for responsibly realizing the project. DEDA helps with the responsible handling of data, models, algorithms, and more.

Goal

The DEDA handbook can be used together with the DEDA poster. The handbook aims to provide further explanation of certain concepts and to delve deeper into specific questions.

For some questions, examples are provided as well.

How to use this handbook

1 You can use this handbook as support for the DEDA poster. The poster contains coloured clusters of questions. The questions are structured in a process-oriented manner so that they highlight technical considerations (blue) and social considerations (green) in succession. You will find explanations for each cluster of questions in this handbook. You can use sticky notes to write down the answers to the questions, note down action points, and collect responses on the poster.

2 In the section with technical questions, you can skip any questions that are not (yet) relevant to the phase the project is in. However, we recommend that you revisit these questions during later phases of the project. The social considerations focus on questions relating to privacy, communication, bias, transparent communication, organization, and future scenarios.

Note: *If you are unable to answer a question because you need additional details, it is advisable to make the collection of such details an action point for answering that question. You can also note specific points of attention for that question.*

3 The last section on the DEDA poster calls for extra attention to values. Based on the values that are important within your organization and to the project members, you can make decisions about points of attention for the project and possible obstacles that need to be overcome. For this purpose, the conclusion offers a few closing questions.

DEDA provides a general overview of everything related to your data or AI project. If the conclusion from completing DEDA is that personal data is being processed that also poses a high privacy risk for the individuals concerned, then a DPIA (Data Protection Impact Assessment) must be completed. Completing the DPIA has been mandatory since May 25, 2018, based on the GDPR (General Data Protection Regulation). Assessment) must be completed.

The GDPR is in force throughout Europe and has direct effects in all member states. More information about the DPIA can be found in the DEDA questions about Privacy (p.20)

Preparing for a DEDA session

As the organizer and/or moderator of a session with DEDA, there are a number of things to consider in advance. First of all, it is important to have a suitable case study from your own organization. This can be anything, as long as it concerns a data project with potential ethical pitfalls. Secondly, it is crucial to invite an interdisciplinary group of participants. This allows everyone to complement each other's expertise.

As a guideline, the following factors can be considered. Each DEDA poster can accommodate four to eight people. Allow three to four hours for the session and invite participants well in advance to a spacious meeting room such that, if necessary, groups can be formed.

It is advisable to start the session with a brief summary of the case so that all participants are well informed about the subject matter. The session can be concluded in plenary by sharing all findings from the participants.

The conclusion on the DEDA poster provides a foundation for further follow-up steps in the project.

Questions

TECHNICAL CONSIDERATIONS



DATA SOURCES

Explanation

It is important to check where the data you want to use comes from. When datasets are purchased or data collection is outsourced, it can sometimes be difficult to understand the context in which the data was originally collected, the limitations of the dataset and whether this data is actually suitable for the purpose.

Sometimes there is a “the more, the better” mentality when it comes to collecting data. This can be problematic when certain data is not relevant to the context of the project. It is important to reflect on this carefully. A larger dataset does not always lead to better quality research.

In addition, the “expiry date” of data is important. The GDPR contains regulations on data retention periods, but it is also necessary to be aware of the limited usability of some datasets. How long will the data remain suitable and relevant for the intended project? For this reason, ethical reflections on the origin of the dataset must begin with a reflection on what kind of data is actually needed for the project and how long this data will remain relevant.

Finally, it is important to reflect on the representation of different groups of people in the dataset. For example, over-representation or under-representation of certain social groups can lead to differences in outcomes for these groups. When only one type of data is present, the model does not learn to predict well with other types of data, which can result in bias. A balanced dataset, with a representative distribution across different variables, ensures that the model can make predictions that are as accurate and fair as possible.

In addition, consideration can be given to whether the dataset contains certain proxy

characteristics. These are indirect characteristics in the data that appear to be neutral, but can nevertheless indirectly say something about certain protected characteristics. For example, the seemingly neutral characteristic “low literacy” could serve as a proxy for the protected characteristic “origin”. It is therefore necessary to be well informed about what the dataset contains and what it tells us.

When both the social groups and the outcome labels (e.g. low or high risk) are known, it is then possible to test mathematically whether there is any bias. This can be done using so-called bias and fairness measurements, which show whether the model systematically disadvantages certain groups.

Please note:

The section on Bias will discuss possible bias in the project in more detail later on. If a discussion arises that takes too much time and goes beyond bias in the data, stop the discussion here and inform the group that this topic will be discussed in greater detail at a later stage.

ALGORITHMS

Explanation

Algorithms can process data and develop insights based on that data.

An example of this is an algorithm that determines when a traffic light should turn green or an algorithm that assesses who is eligible for benefits. Algorithms use mathematical models. However, these models are rarely value-free. Models are expressed numerically and calculations often appear neutral. However, algorithms are often designed in such a way that a numerical value is linked to a normative judgment, such as “risk of fraud” and “unusual event.” Therefore, it is important to think carefully about why and how a type of algorithm will be used, as well as its impact on these normative judgements.

In-depth explanation questions 13 and 14: types of algorithms

There are many different types of algorithms. Some algorithms are very simple and often easy to explain, but there are also many more complex and difficult to explain algorithms that offer more accurate results. A simple algorithm could, for example, be a decision tree in which policy rules are programmed. These rules are therefore devised by humans and translated into an algorithmic system. A machine learning model is more complex and makes choices based on data. How these models arrive at a particular choice is more difficult to ascertain because the model trains itself using a large amount of data. The rules for arriving at a specific result are not programmed in advance by humans. The model does this itself, by adjusting its internal weights in a learning process based on the training data, until it can make accurate predictions. The advantage is that a machine learning model can create a certain outcome without people having to come up with decision rules themselves. The

disadvantage is that with these models, it is often more difficult to explain how and why a particular outcome has been reached.

In-depth explanation question 15

Algorithms will increasingly inform decision-making. That is why it is important to understand how the output of an algorithm is generated. Algorithms must be accessible for review by external experts and for verification of the results. This is one of the pillars of transparency. (Government) organizations must be able to explain how their models and algorithms work. What does not help is that it can be unclear who exactly owns an algorithm and who is responsible for its results. In addition, models and algorithms are not always made public, which can make transparent communication about them difficult or even impossible.

In-depth question 16, 17, and 18

The output of an algorithm can be numerical, such as the value of a house, but it can also be a label or category, such as ‘high risk’. It is important to consider the extent to which the result of an algorithm is used as a guide or supplement and the extent to which a human being can deviate from the output. For example, there is a risk in ‘automation bias’, the phenomenon whereby people too easily accept the answer provided by a system because the system appears to be automatically correct. To counteract automation bias, in some cases training is therefore needed for the employees who will be working with (the results of) the algorithm.

In-depth explanation questions 19 and 20: evaluation criteria and errors

Certain criteria are often used to evaluate the functioning of an algorithm. These are also referred to as evaluation criteria. Which evaluation criteria are relevant for your algorithm depends on the exact outcome of the algorithm. Is it a binary outcome (yes/no, 1/0, positive/negative)? Or is it a non-binary outcome (such as predicting values between 0 and 1)?

If the algorithm produces a binary outcome, such as the predictions of binary classification models, a so-called 'confusion matrix' is often used. Based on the predictions, there are a number of important statistics: 'precision', 'recall', and 'accuracy'. Precision refers to the proportion of positive predictions ($TP/(TP+FP)$), recall refers to the actual number of positive cases ($TP/(TP+FN)$), and accuracy refers to all correct predictions relative to the entire dataset ($(TN+TP)/N$). A false positive is what happens when the model predicts positive (1), and the actual label is negative (0). A false negative is the opposite.

To give a concrete example: if a pregnancy test indicates that someone is pregnant, when in fact they are not, we refer to this as a false positive. The reverse situation, in which the pregnancy test indicates that someone is not pregnant, when in fact they are, is called a false negative. When the pregnancy test indicates the correct state of being pregnant (or not), we refer to this as a true positive or true negative.

However, not every model classifies in such categories or in a binary manner at all. Non-binary outcomes are also possible. For example, a regression model predicts continuous values between 0 and 1, which makes discrete categorization into classes unworkable. In this case, an "error term" is often used, a function that can be used to determine how correct or incorrect an output is. The most commonly used terms are 'Mean Squared Error (MSE)' or ' R^2 (R-squared)'. The first measures the distance to the correct outcome and penalizes larger errors more heavily. The second mainly explains the variance within the dataset.

		Real Label	
		Positive	Negative
Predicted Label	Positive	True Positive (TP)	False Positive (FP)
	Negative	False Negative (FN)	True Negative (TN)

ANONYMIZATION

Explanation

Anonymization is the process of making personal data in data sets untraceable. This is irreversible. One technique that can be used for this is generalization, for example by reducing all personal identification data to age groups. An example of anonymization is changing social security numbers within a data set to a random number, after which the original social security numbers are deleted. Of course, it is important that the new random number is not linked to traceable personal data such as first and last names. Please note: by combining overlapping data points (such as zip code, age, gender, etc.) with other datasets, it may still be possible to identify individuals in the anonymized dataset.

Pseudonymization makes a dataset untraceable in a similar way, but the process can still be reversed. Returning to the above example, in the case of pseudonymization, the social security number is not completely removed. In that case, at least one person will be needed who can potentially reverse the pseudonymization. That person must have the key to reverse the process. When this happens, the random number is replaced by the social security number, possibly even including personal data such as first and last name.

VISUALISATION

Explanation

For some projects, it is necessary to visualize data or the results of data projects. This can be done in many different ways, such as through infographics, data visualizations, or dashboards. Keep in mind that visualization always influences the interpretation of the data and is therefore not neutral. Consider, for example, the use of color (red versus green) or how the scale of a graph can influence the message being conveyed.

The main purpose of this question is to reflect on the method used to present the data or results. Is this the best method? What are the reasons for choosing a (less) obvious way of visualizing? Can a visualization be misinterpreted? What visualization style and techniques are used, and do they express prejudice or bias? Keep in mind that the same dataset visualized in different ways can also lead to different “readings” of the data. In addition, it is important to note that not all results based on certain datasets are suitable for visualization.

ACCESS AND REUSE

Explanation

‘Access’ can mean many different things, depending on the context. Here, we are referring to the issue of access to collected and archived datasets within your organization. Making data accessible within a company or municipality, or making data accessible to everyone (open access), involves various considerations. On the one hand, open access can increase transparency and trust, but on the other hand, trust can be damaged if the data is not suitable for sharing with everyone. Some datasets contain confidential information and/or personal data. The confidentiality of this information must be guaranteed.

In addition, careful consideration must be given to whether the results will be reused in any way and what pitfalls this may present. When answering the questions on this topic, it may be helpful to weigh the pros and cons of reusing data.

Sometimes datasets can be beneficial not only for one project, but can be reused for other purposes and projects. However, it is quite possible that data collected for a specific project may lose its validity when reused for another project. There are also restrictions in the GDPR on the use of data for purposes other than those specified at the time of data collection. This also applies to the reuse of the dataset by third parties. Furthermore, an external party with access to the data could undermine the organization’s knowledge proposition and/or reliability.

SOCIETAL CONSIDERATIONS



PRIVACY

Explanation

The privacy of every living person is protected by the GDPR. This law strictly regulates the infringement of privacy or the leakage of personal information, and violations of this law can be punished with heavy fines. Even if there is a general perception that people are careless with their right to privacy by signing up for social media or openly discussing intimate details of their personal lives, this does not mean that the right to privacy is losing its urgency. Privacy remains essential to democracy. After all, it is up to individuals themselves to decide what information they share.

About the Data Protection Impact Assessment (DPIA)

- *The DPIA is a tool for identifying privacy risks in advance. Measures can then be taken to reduce these risks.*
- *Completing the DPIA is not always mandatory. The DPIA is mandatory under the GDPR if there is a high privacy risk for the individuals whose data is being processed.*
- *Not every data project involves the processing of personal data.*
- *Data from DEDA can be helpful when completing a DPIA.*

BIAS

Explanation

Bias is a major problem in data projects. A dataset, model, algorithm, or organization with bias produces unequal results for different (social) groups. If these biases are not recognized and subsequently prevented or resolved, discrimination can arise. This can affect the value of equality and non-discrimination. This happens, for example, when certain biases are (unintentionally) incorporated into the data and therefore do not reflect fair expectations, when the algorithm contains biased decision rules or indicators, or when the results are often wrongly assumed to be correct.

In-depth explanation question 32

In addition to concrete signs of discrimination, you may also have a gut feeling about your project. Gut feelings can be very useful in detecting possible bias, by recognizing early on that 'something' in the project is not quite right. It is important to listen to these kinds of feelings and investigate them. If these feelings are indeed based on an error in the data or in the project, valuable solutions can be found early on and potential harm can be prevented. Moreover, even if these feelings turn out to be unfounded, they are often a good indication of how the algorithm is perceived by third parties.

In-depth explanation question 33

There are various ways in which bias can reappear in your data project. All these manifestations of bias can contribute in their own way to the unequal treatment of social groups. This page contains a number of descriptions of bias for different phases of

the project. Keep in mind that these are just a few examples of many. The list of examples is therefore not exhaustive.

Bias in data

The section on data sources already discussed the under- or over-representation of certain groups of people in the data and any proxy variables in the datasets that could cause bias in the model. In addition, consideration can also be given to bias in the selection of the datasets. For example, 'selection bias' can occur when the choice of your data source is a biased choice in itself. Selection bias can occur, for example, if an algorithm is used to regulate traffic planning, but the data source only uses data from smartphone users. In this case, groups that use smartphones less or not at all (the elderly, the digitally illiterate) are not included in the algorithm. Policy decisions based on this are then based on an incomplete picture of mobility. What may appear to be an objective fact ('mobility in the city') may be influenced by the type of data used (smartphone users only). Random sampling, control groups (where possible), and discussions with your team can minimize the risk of biased selection.

Bias in algorithms

Bias can also be hidden in the way the algorithm works. A certain group may not be able to use the algorithm, or decision rules or measurement methods may have been used that could cause bias.

Bias during use of algorithms

Bias can also arise when interpreting the results. For example, in the form of confirmation bias or automation bias. Automation bias has already been discussed under the heading Algorithms. Confirmation bias refers to our human tendency to seek confirmation of pre-existing ideas and to give preference to statements that confirm our beliefs (as opposed to statements that challenge our beliefs). Confirmation bias is often reinforced by tunnel vision. This means that people tend to ignore opinions that do not correspond with their own point of view, even though these differing opinions may be important and valid. This tendency can cause problems in the use of data, because important external points of view, differing interpretations, and/or other concerns are missing or not being heard..

In-depth explanation question 34

A feedback loop is also a way in which bias can arise. A feedback loop in data projects means that the results of the project are used in some way in the project again and are deployed as new data. This can be a conscious choice, but it can also happen less deliberately. In any case, one must be vigilant for a self-reinforcing feedback loop, in which the outcomes of a previous iteration impact future decisions.

A classic example of a feedback loop is an algorithmic system that helps determine the route of police patrols. If more crime is reported or found in Neighborhood A, the system will indicate that patrols should be increased in Neighborhood A. This will result in a greater police presence in Neighborhood A, which in turn will lead to more crime being found in that neighborhood. Based on these increasing reports, even more patrols will be carried out in the neighborhood. This can also lead to reactions from residents, which in turn can be used as a reason for increased police presence. This can create a self-reinforcing effect, whereby Neighborhood A is over-patrolled and other neighborhoods are overlooked.

TRANSPARENT COMMUNICATION

Explanation

It is often the case that communication and stakeholder involvement are only considered when things go wrong and critical questions are asked. It is essential to communicate clearly not only within the organization, but also externally. This helps to convey responsibility for the project and can contribute to building trust among parties who are not directly involved in the project, but who are affected by it. Remember that political parties, citizens, lawyers, or activists can use their rights to inquire about your data projects. It is therefore important to consider the extent to which stakeholders are involved and the degree of external transparency.

In-depth explanation question 36

It is not always easy to be transparent. Models and algorithms can be very complex. A high level of knowledge of statistics and data science is often required to understand them. In such situations, transparency does not always mean that models and algorithms must be translated into understandable language, but that they must be accessible for critical questioning. Finally, it is also possible to be too transparent in data projects, for example with regard to the dataset. If this happens, an unintended data breach may occur, resulting in too much information being provided to people with malicious intentions.

ORGANISATION

Explanation

Within an organisation, models and algorithms are also subject to accountability and good governance. Government agencies are increasingly being asked to account for their models and algorithms. In this context, it is essential that responsibilities are clearly defined, not least because it can sometimes be unclear who owns an algorithm. Data projects also often influence daily practice within an organisation. They can change work processes, shift tasks or influence the way employees experience their work.

refers to employees' understanding of how systems work, the risks involved and their responsibilities when using them. However, AI literacy is more broadly applicable than just to AI systems and is relevant to any technology introduced within an organisation.

In-depth explanation question 42

Data projects can have a direct impact on employees within the organisation. An important point of attention here is “deskilling”. This is the phenomenon whereby specialised work is gradually taken over by technology, which can lead to a decline in employees' skills, with adverse consequences for quality. This can also have consequences for job satisfaction or employment, among other things. It is therefore important to also consider the impact of your project on your employees. Keep in mind that AI projects often get stuck because too little attention has been paid to implementation within the organisation and getting employees on board.

In-depth explanation question 43

Responsible and proper governance requires careful and expert use of data and AI. It is therefore important that employees are sufficiently “AI literate”. This term, also mentioned in the European AI Regulation,

FUTURE SCENARIOS

Explanation

Although some data projects are only relevant or in use for a short period of time, it is often the case that projects have medium to long-term effects. It is therefore important to consider these long-term effects now, and whether they would change the current state of affairs. Issues such as “function creep”, data sovereignty and climate impact play an important role in this. It is therefore necessary to conduct interim evaluations to ensure that the data project is still doing what it is intended to do and that there are no (unintended) negative effects.

In-depth explanation question 44

“Function creep” means that information is used for a purpose other than the original intended purpose. An example of function creep is as follows: a security system is installed in the office, requiring employees to check in and out for the purpose of facilitating secure access. However, if this information is then used to track or monitor individual employees, we refer to this as function creep. This is a violation of privacy and dignity that goes beyond the stated purpose of creating secure access. By exploring alternative applications at an early stage, clear limits can be set on the use of data. This helps to prevent misuse, promotes transparency and protects those involved from unexpected negative consequences.

In-depth explanation question 45

Data sovereignty is about maintaining control over data. Where is the data stored? Which parties have access to the data? Which laws apply to data retention? Do third parties obtain

the data and can they process it further? Is there a risk of dependence on a particular supplier? As digital environments become increasingly dependent on external (international) suppliers and infrastructures, it is important to strive for data sovereignty. This contributes to safeguarding values such as autonomy and security.

In-depth explanation question 46

Digital projects have an impact on the climate through their energy consumption, data centre load and the ecological footprint of hardware. At the same time, data projects can also help to promote sustainability, for example by making processes more efficient. It is therefore important to consider the environmental impact of a data project in order to make climate-responsible choices.

OUTCOME



OUTCOME

VALUES

Explanation

At the beginning of the workshop, you wrote down a number of values. Take these out again and check for yourself whether these values are represented in the project, and if so, where.

Go through all the members of the project group and ask them to say whether, and where, the values are represented in the project. The person whose turn it is may stick their values on the poster in a place where they think they are most positively reflected. For example, the value “ownership” could be stuck next to the questions about Organisation or next to the questions about Data Sources. When it is your turn to stick your values on the poster, explain to the rest of your group why you think the values belong in that place in the project. Discussing this as a group can lead to new insights.

You may also find that (one of) the values you have noted fit everywhere in the project. In that case, you can stick that value in the middle.

Finally, it is also possible that (one of) the values you have written down are not represented in the project at all. In that case, discuss with each other whether this value could be important for the project and whether it is necessary to change certain aspects of the project so that the value can be represented. It may happen that a value is less relevant to a project. For example, the value “loyalty” is less relevant to a data project that anonymously counts passers-by in a certain area. However, it may also happen that a value is relevant to the project, but is not reflected positively. This may be the case, for example, when “transparency” is written down as an important value, but when answering the

questions it turns out that the project is not transparent at all. In such a case, discuss as a group what needs to be changed in the project in order to respond to that value.

OUTCOME

CONCLUSION

Explanation

Once all values have been discussed, the group can move on to the final, concluding section. The Conclusion offers four concluding questions that summarise the entire session and form the basis for further follow-up steps.

OUTCOME

NEXT STEPS

Explanation

You can use the DEDA poster, including all answers and action points, as a basis for or to support a report on the data project.

The answers and action points can serve as tools that make it easier to account for the project. For example, the answers can help explain what choices were made and, just as importantly, why certain things were not done.

OUTCOME

FURTHER INFORMATION

Explanation

For further information about DEDA, we refer to the following publications:

Franzke, A., Muis, I. M., & Schaefer, M. T. (2021). Data Ethics Decision Aid (DEDA): a dialogical framework for ethical inquiry of AI and data projects in the Netherlands. *Ethics and Information Technology*, 23(3), 551-567. Article 3. <https://link.springer.com/article/10.1007/s10676-020-09577-5>

Siffels, L., van den Berg, D., Schaefer, M., & Muis, I. (2022). Public Values and Technological Change: Mapping how Municipalities Grapple with Data Ethics. In *New Perspectives in Critical Data Studies: The Ambivalences of Data Power* (pp. 243-266). (Transforming Communications – Studies in Cross-Media Research). Palgrave Macmillan. https://link.springer.com/chapter/10.1007/978-3-030-96180-0_11

Colophon

Concept:

Mirko Tobias Schäfer, Ph.D.
Aline Franzke, MA

Editors:

Julia Straatman, MA
Iris Muis, MA
Hajo Bijleveld, BA
Ydo Bergstra, BSc
Lysa Ngouateu, BSc
Mirko Tobias Schäfer, Ph.D.
Bart Anthony Kamphorst, Ph.D.

Cover illustration:

Sammy Hemerik

Utrecht University
Data School
Drift 13, kamer 0.01
3512 BR Utrecht
Netherlands

'The Data Ethics Decision Aid' was developed by the Utrecht Data School and Utrecht University on behalf of the Municipality of Utrecht. © Data School, Utrecht University 2026

DEDA - Version 4.0
March 2026
For updates, visit:

deda.dataschool.nl

Neither Utrecht University, nor the Data School, nor the Municipality of Utrecht are liable for any damage related to or arising from the use of material or content from 'The Ethical Data Assistant'.

Utrecht, 2026.

As DEDA is used, we learn how to improve it. These improvements will be implemented in future versions.
If you use DEDA and have any comments or suggestions, please do not hesitate to share them with us. You can always email us at:

dataschool@uu.nl





**Utrecht
University**

DATA SCHOOL