A Citizen's Guide to Data

Ethical, Social and Legal Issues

The Alan Turing Institute

Contents

A Citizen's Guide to Data
Author & Design: Dr Christopher Burr

Illustrations: Johnathan Lambert & Adobe Stock

This document has been produced to support specific activities for the Resident's Panel being run by Camden City Council. It is not a public document in its current version, but will be revised and developed following the workshop for public release.

01

An Introduction to Data

A Citizen's Guide

Data is everywhere!

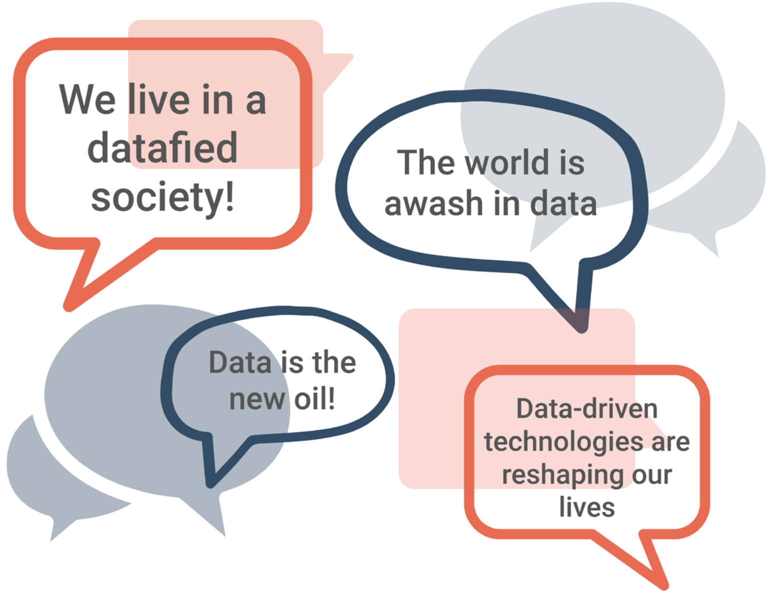

You have probably heard one or more of these phrases from the many technology entrepreneurs or political pundits eager to extol the virtues and value of data. Whether it's mentioned in relation to mundane projects, such as improving administration for a company, or addresssing global priorities, such as climate change, data seems to be everywhere.

The Value of Data

There is no question about the positive impact that data has had on many areas of society. Data has been essential to the COVID-19 recovery effort. Data helps individuals connect online or track down long-lost relatives. Data can even help our species peer into the farthest reaches of our universe or the deepest parts of our history on this planet.

While data is essential for many parts of society and scientific knowledge, there are nevertheless questions that must be asked about the use of data in public services and for the governance of society more broadly. However, before we jump in and look at these ethical, legal, and social questions there are a few preliminary remarks that ought to be made.

For instance, you may be wondering what all the fuss is about? Haven't we relied on data to manage our societies for centuries? Why are we still asking questions about data today?

We have come a long way since some of the earliest forms of data recording. Whereas our ancestors back in 18,000 BCE stored data on tally sticks, such as the Ishango Bone depicted here on the right, our present methods for data collection and the purposes for which we use it have changed drastically.

A significant reason for the renewed attention that data is receiving is the increasingly widespread use of data-driven technologies. Personal devices such as your smartphone, television, smart lights and speakers, or maybe even your fridge and toaster are now sources of diverse types of data. These devices make use of embedded sensors that can automatically measure and monitor data about your location, viewing habits, and dietary preferences, and then send this information to many interested parties via the Internet.

If the increase in data usage was restricted to personal devices, maybe there wouldn't be that much interest. Or, perhaps it would just be restricted to legal issues regarding data privacy and protection However, it is not just individual consumers that are having to adjust and adapt to the new environment.

Public and private sector organisations are also experiencing drastic shifts in how they operate as a direct result of data-driven technologies. And the effects of how these organisations are using novel, data-driven technologies often spills over into the public domain. The consequences of this can be severe for specific individuals or groups of people.

The Ishango Bone—discovered in 1950 and believed to be used as a tally stick. Image from WikiMedia Commons.

In 2020, for example, all secondary education examinations in the UK were cancelled because of the COVID-19 pandemic. In a failed attempt to prevent disruption to the affected students, Ofqual—the regulator of qualifications, exams and tests in England—produced an algorithmic system that was supposed to moderate teacher-predicted grades and prevent grade inflation. However, the algorithm, which made use of educational data from schools, was heavily criticised because of its perceived unfairness in the way that it penalised smart students from disadvantaged schools.

Part of the failure in this instance can be attributed to poor decision-making and implementation by the regulator and government ministers, and specifically their choice about which data to have the algorithm consider. For instance, the algorithm worked by assessing data about the historical performance of the school and previous year groups when moderating an individual student's grade. The choice to include this data meant that students who went to the most priviliged schools were favoured by the algorithm over and above students of similar intelligence at disadvantaged schools.

Unfortunately, it is not always easy to identify the reasons why specific uses of technology or data cause problems. And even when there is some clear fault that can be identified, it can also be challenging to know what the best course of action is to resolve the issue.

We have come a long way since some of the earliest forms of data recording.

Haven't we relied on data to manage our societies for centuries?

Why are we still asking questions about data today?

For example, should we scrap the idea of algorithms in grading altogether? Or should there just be additional oversight mechanisms to ensure that they are designed responsibly and used fairly?

Or, to take another example, do we have an obligation to share our sensitive health data with public health authoirities to benefit the rest of society in tackling the effects of the global pandemic?

While these questions are about data and data-driven technologies, they are not only about data and technology. They are also about ethics, law, and society, and any attempt to answer them must have some awareness of these wider topics.

This guide is intended to help you in answering such questions, but it is not about giving simple routes to pre-determined answers. As we will see, there is often no straightforward answer to these questions. Rather, getting to an answer often demands deliberate and inclusive forms of discussion and debate. Therefore, this guide is intended to provide you with the critical tools necessary to understand these questions and contribute to the discussion as an active and informed citizen.

This guide is intended to help you answer questions about the ethical, legal, and social issues. There is not always a straightforward answer, but by having a better understanding of the topics and issues incolved you will have the critical tools necessary to contribute to the discussion as a more active and informed citizen.

02

Understanding Data

A Data Journey

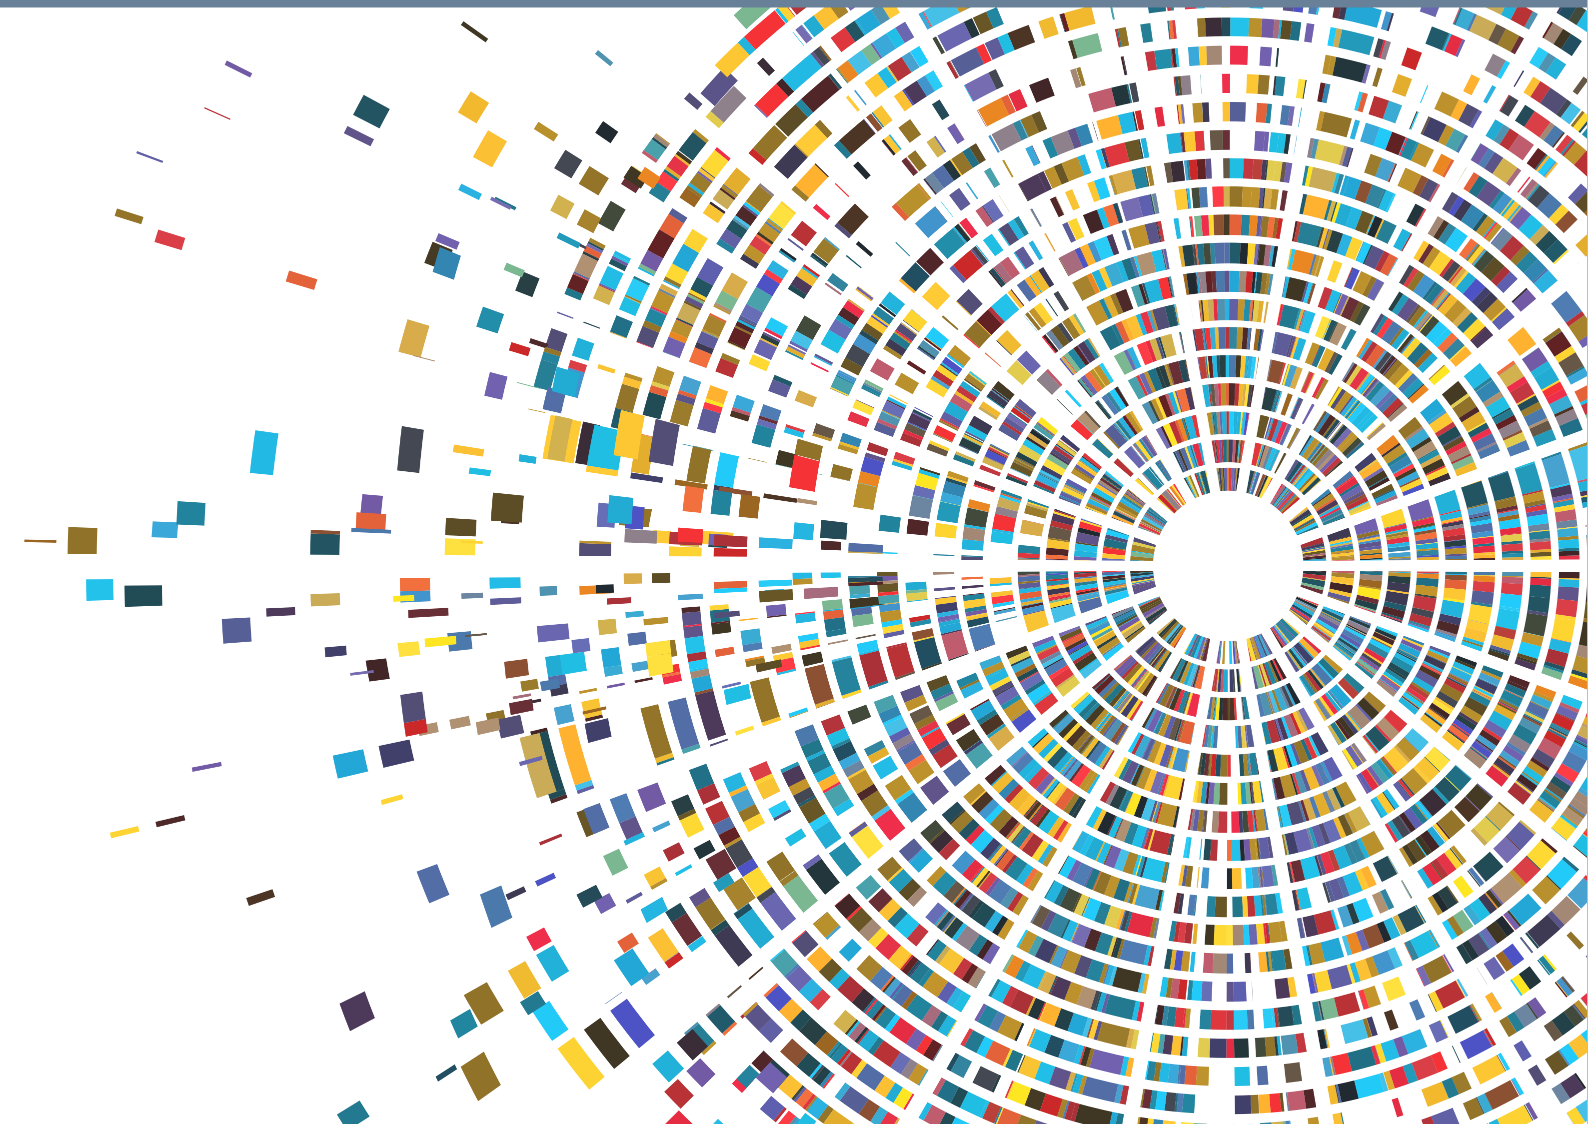

Planning our Journey

We are going to explore a data journey to help demonstrate why your data matters. We will start with an initial stage of project planning and end with the eventual deletion of data.

Alongside this data journey we will refer to an illustrative case study involving a made-up project at a local police force who are building an automated surveillance system for facial recognition.

Project Planning

There is a vast amount of data that can be collected about an object or person. Only some of it will be useful.

Planning what data needs to be collected, and how the data will be used to solve some problem is an important first step to good project management.

If you visit your doctor they may wish to measure your blood pressure, record your weight, take details about your current medication, or document any allergies you may have. These are some examples of data about your health.

The type of data that is collected will depend on the specific illness being assessed. It would be strange to go to your doctor complaining of a pain in your shoulder, only to have the doctor collect data about your eyesight.

So, before data is collected it is important to answer questions such as,

- What is the problem that is being addressed?
- What data is needed to address this problem?
- Who will or could have access to the data?
- How will they use the data to address the problem?

Without a rough answer to these types of questions, time is likely to be wasted by recording unnecessary data. In some cases, this could even result in the violation of certain laws about data privacy and protection.

When a public sector organisations starts a data-driven project it is also vital that they consider the social benefits and harms. This is especially important when groups of citizens have different views about whether the project will be beneficial or harmful. Engaging members of the public in the earliest stages of project planning is, therefore, a sensible idea.

Let's start exploring our case study (see box on the right) that we will return to for each stage in this data journey.

Case Study

Our case study involves a local police force who want to improve their ability to identify wanted suspects. A senior team member suggests that they build an artificial intelligence (AI) system to identify wanted suspects in public spaces because there are not enough officers to patrol the streets at all times.

The proposal is to develop an automated facial recognition system, which can gather images from a surveillance system, and automatically match these images to a police database of wanted suspects and known criminals. This requires collecting lots of data about individual faces.

They believe that this system will be well received by the public because it will help them reduce crime. From their perspective this is an obvious benefit of the project.

Topics for Reflection

1. Is an automated facial recognition system the best proposal for the problem of identifying wanted suspects?
2. Will all members of society see the use of this system as a public good?
3. Are there any risks involved with building an AI system to detect wanted suspects or known criminals?

Data Collection & Processing

Data collection is a process of observation and measurement to record data about a property of an object or person (e.g. price, weight).

Let's assume you visit a doctor and they make a note of the following data:

Patient #	Age	Weight (Kg)	Blood Pressure
1883652	36 > 50	72	115/75

When this data is collected it creates a small record of data about you. We can call the collection of records a **dataset**.

Your doctor can use this dataset to help assess or diagnose illnesses or conditions. But, a health insurance company could also use this data to decide on which policy to issue you. You may be happy with your doctor having access to the data but unhappy with an insurance company using the data.

A human can manually collect data (as above). But data can also be collected automatically, such as when you use a computer or smartphone. Automatic data collection is now a common feature of modern life because of increased connectivity and advances in computer technology.

Did you know?

When you use the internet, data is automatically collected about your usage. This goes beyond the obvious data such as which websites you visit. It can also include data such as how long you stay on a certain site, where you are accessing the site from (e.g. your postcode), whether you log in to a social media account, and even how you move your mouse cursor across the screen. Some of this data may seem strange. For instance, why would your mouse cursor movements be useful to anyone?

These types of data offer valuable insights for advertising companies that want to learn about your preferences. If they can use data to learn about what grabs your attention, they can show your adverts that are more likely to engage you. But whether the data leads to useful insights depends on how well it is collected and processed.

If lots of data is missing it can be difficult to make a reliable decision. It is also necessary to ensure the data is stored in the right format. It is unlikely to be useful for the doctor if her patient's data looks like this:

Patient #	Age	?	Blood Pressure
1883652	36 > 50	72	115/75
unknown	38	missing	missing
1992627	missing	116/72	78 kg

Case Study

The police force decide to go ahead and develop the automated facial recognition system. To do this they will need to collect two datasets.

First, they need to have a dataset with the prior records of known criminals or wanted suspects, including images of their faces. This is because the facial recognition system cannot identify a wanted suspect or known criminal if it has no record of who they are.

Second, they need to collect images of individual members of the public. They decide to collect this data from CCTV cameras that are in use throughout public spaces, such as high streets or shopping centres. However, the images that the CCTV cameras record contain many faces at the same time. So, the team decide to use AI to detect and extract the individual faces from the video and try to match these images with the dataset of wanted suspects.

There is no way of determining whether an individual matches a wanted target before collecting visual data of their face. Because of this limitation, the system will have to temporarily record images of all members of the public and delete the ones that are not matched.

Topics for Reflection

1. Does the benefit of this system justify the potential intrusion on public privacy? If not, should the project go ahead?

Data Exploration & Analysis

Data exploration and analysis helps project teams to spot patterns in their data that may lead to important insights or discoveries

You may have come across the term 'data science' before. The term refers to the practice of using sets of data, which can be very large, to help drive scientific discovery throuygh exploration and analysis.

For example, international projects have collected massive amounts of data about Covid-19 patients. This data is used to explore and analyse which factors, such as age or pre-existing illness, increase the risks associated with Covid-19.

Although we won't work with a dataset of the same size, let's have a go at being data scientists.

In 1912, the RMS Titanic sank after colliding with an iceberg on her maiden voyage. Because there were not enough lifeboats for all the passengers and crew onboard, 1502 people drowned.

On the following page is a small sample from a larger dataset that many new data scientists like to use. We'll look at the first 10 passengers.

For each of the passengers, you have the following information:

- **Name** (including title)
- **Sex**
- **Age**
- **Class** (the class of ticket)
- **Survived** (whether they survived)

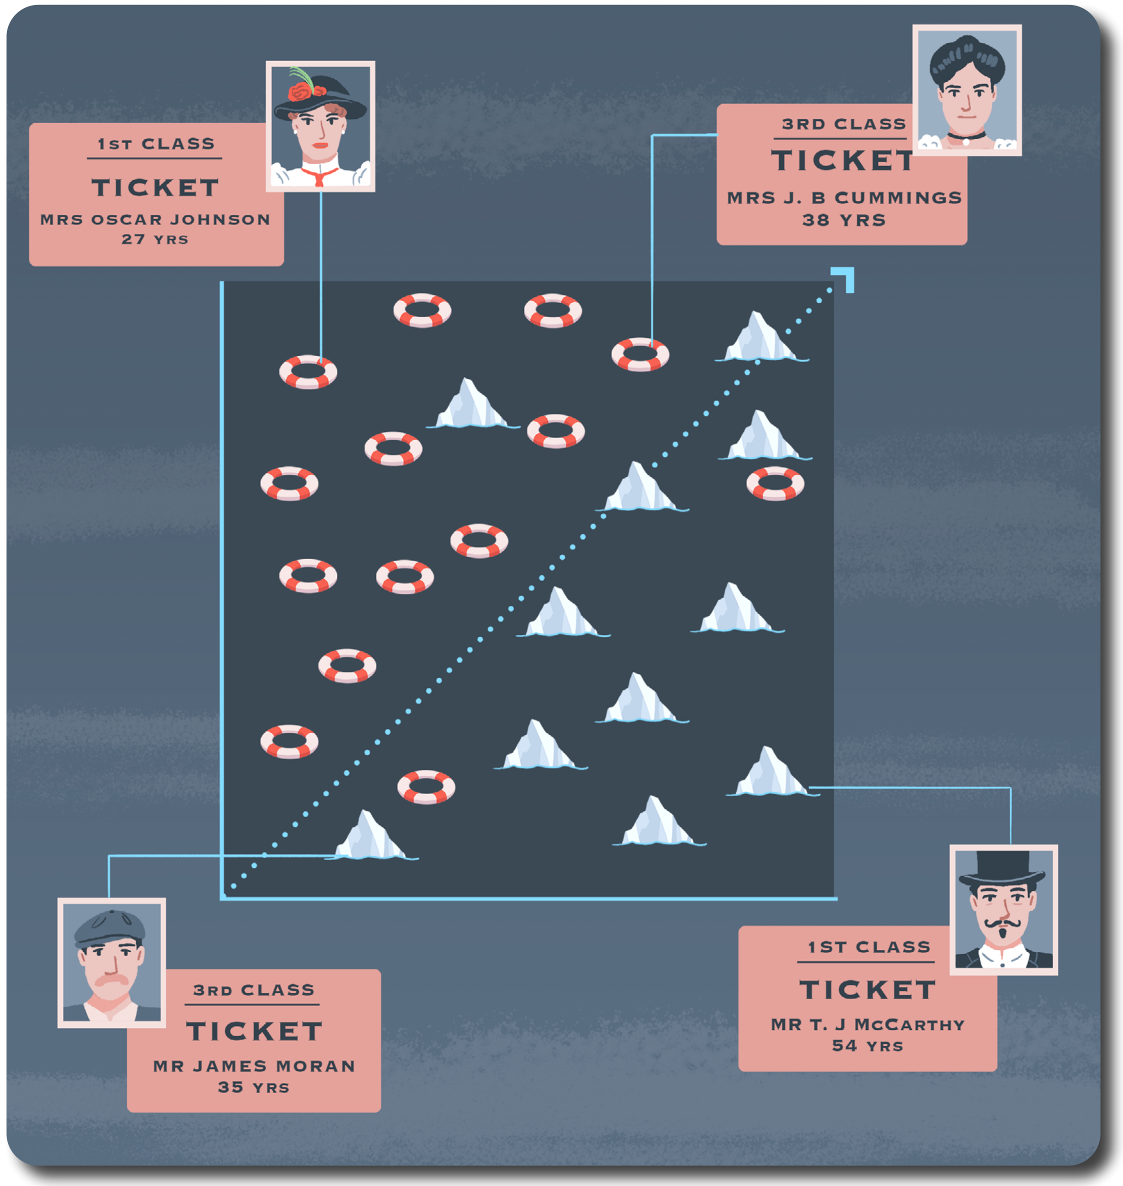

It is your task to explore this dataset and try to answer some questions. The passengers are grouped according to whether they survived, as this will be the main focus of our task. Let's start by looking at the data.

Name	Sex	Age	Class	Survived
Mr. Owen Harris Braund	Male	22	3rd	No
Mr. William Henry Allen	Male	35	3rd	No
Mr. James Moran	Male	No Data	3rd	No
Mr. Timothy J McCarthy	Male	54	1st	No
Master. Gosta Leonard Palsson	Male	2	3rd	No
Mrs. John Bradley Cummings	Female	38	1st	Yes
Miss. Laina Heikkinen	Female	26	3rd	Yes
Mrs. Jacques Heath Futrelle	Female	35	1st	Yes
Mrs. Oscar Johnson	Female	27	3rd	Yes
Mrs. Nicholas	Female	14	2nd	Yes

Table 1 - Titanic Passengers

Having looked at the data, try to answer the following questions:

- Who was more likely to survive, a male or a female?
- Does the ticket class appear to affect the chance of survival?
- Were younger or older passengers more likely to survive?

Some of these questions may be easier to answer than others. For example, in this small sample all the male passengers died and all the female passengers survived. This is not the case in the full dataset, but females did have a greater chance of survival. If you know about the sinking of the titanic, you may remember that this is because females and children had priority with the lifeboats. You may also be able to work out that the higher the ticket class of the passenger, the greater the chance of survival.

But what about age? Let's look at one more passenger without information about whether they survived:

Name	Sex	Age	Class	Survived
Miss. Constance Mirium West	Female	5	2	?

Now, if you have to predict whether this passenger survived or not, what would you say?

If we put aside our prior knowledge about the Titanic and only look at the first 10 examples it is not easy to answer this question. This is because some of the data are missing and there appears to be conflicting evidence.

For example, if we only consider sex, we would be more likely to predict that Miss West survived. But we also have 'age' and 'class' to consider. And, here, there doesn't appear to be enough data to make a reliable prediction.

If we had a look at the full dataset we would see that age does appear to affect survival rate. And the younger passengers also have a greater chance of survival than the older passengers. But, as data scientists, we may not always have access to all the data we need or want. So, we have to take care to ensure that the data we analyse is reliable and representative, and we make predictions on the basis of careful analysis. Otherwise, our results may be unreliable. For example, how can we be sure that the other 1299 passengers in the full dataset are like the 10 we have looked at?

To support this process of data exploration and analysis, data scientists are increasingly turning to machine learning algorithms to help them navigate large and complex datasets. We will discuss some of these technologies in the next section.

Case Study

The police force decide to collect data from surveillance systems already installed in several public locations: a shopping centre, a busy high street, and a park. Next, they analyse the data to determine which location proves best for detecting wanted suspects.

By exploring and analysing the data, the police force determine that a local shopping centre would be the best location for their system.

Topics for Reflection

1. Have the police force gathered data from a diverse enough set of locations to reliably determine that the shopping centre is the best location?
2. If not, what problems could arise from their analysis of the limited dataset they have collected?

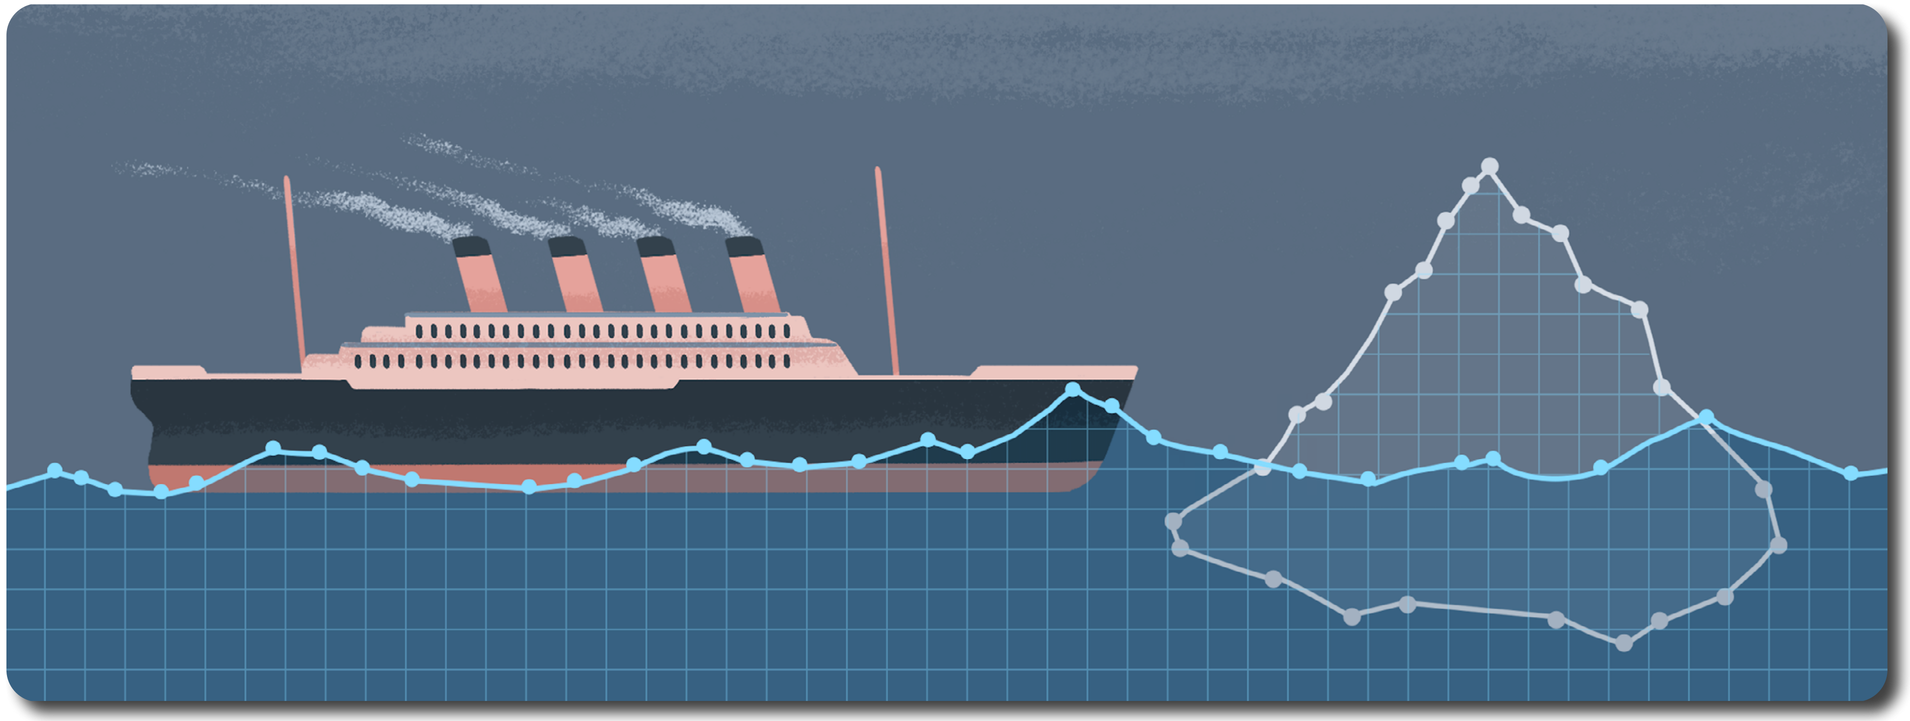

Data Use

An increasingly common use for collecting data is to develop an algorithm or train a machine learning model

It can be expensive to collect, analyse, and store data, so there is often a good reason for doing it in the first place. As we saw in the ‘project planning’ stage, this relates to the use for the data, which can be very varied. Furthermore, there may not be only one use for the collected data.

What is an algorithm?

An algorithm is a procedure or series of steps that provide instructions on how to take a series of inputs and produce an output. For example, a recipe is like an algorithm that provides instructions for taking a series of inputs (i.e. the ingredients) and creating an output (e.g. a cake).

In the case of data-driven technology, an algorithm is a series of instructions for a computer program. The algorithm takes the input data and transforms it, according to the instructions, to achieve some desired output. This could be as simple as taking a set of numbers, and adding them all together to get the sum of the inputs. But it can also be as complex as taking a list of all the stock prices of the top 100 companies over the last 12 months and trying to predict whether a stock price is likely to rise or fall.

Algorithms can be very complex—far beyond the difficulty of planning and cooking a three-tiered cake. Modern uses include machine learning algorithms, which get better at some task as they are fed more training data. Over time, they can learn to perform complex tasks, such as medical diagnosis or environmental monitoring. This requires a lot of data, and the quality of the dataset is also very important.

Case Study

The police’s facial recognition system needs to learn to match images of members of the public with images from the police database. This requires using a type of machine learning known as supervised learning.

Here, the system takes an image from a training set and has to find a possible match in the police database. The matching process can give one of two answers: a positive match, or no match. Each of the images used in the training of the system has a label. The label says whether the image belongs to a wanted suspect or an ordinary member of the public. During training, the system cannot access the labels. But once the system has made its decision, the label reveals whether the decision was correct or incorrect.

The systems uses this feedback to improve its algorithm over time. The system’s goal is to get good enough at matching images from the labelled dataset so it can work in a real-world setting.

Topics for Reflection

1. The accuracy of a machine learning system depends on whether its training data is similar to the data that the system collect in the real-world.
2. If the training data does not include pictures of elderly people or people with darker skin, how will this affect the facial recognition system’s real-world performance?

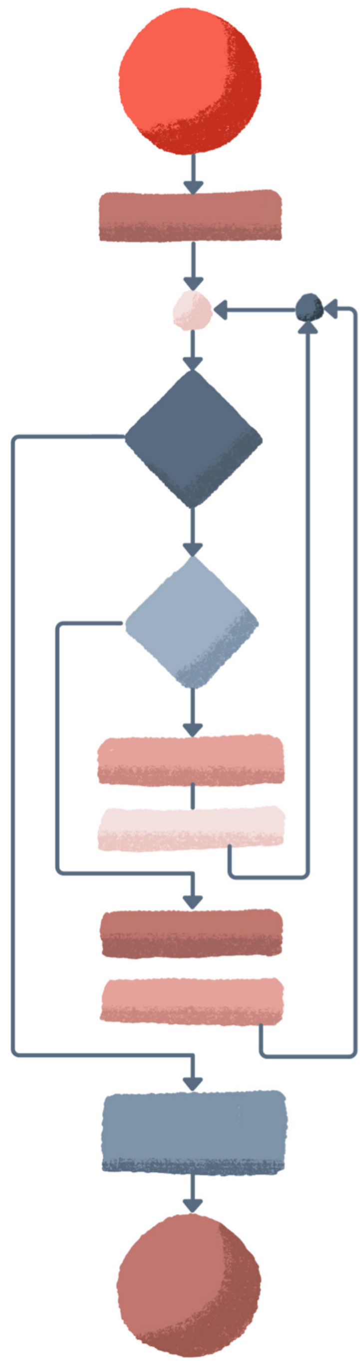

Project Documentation

There can be a lot of steps in data projects, especially those that use complex forms of AI.

So, keeping up-to-date and accurate documentation about the project and the data involved in each stage is important.

By now, we have seen several stages of the data journey. In practice, it is unlikely that a single person would be responsible for all these stages.

Good documentation helps members of the project understand what has been done by others and not repeat unecessary work. And if there is a mistake or issue, accurate and up-to-date records can also help find out what went wrong and how to fix it. For example, a project member may have forgotten to record an important detail, such as the addition of new data at a later stage.

Also, when it comes to using a system, it is vital that the project owners can prove to any affected people or regulators that their system was developed in line with best practices. If they have not kept accurate and up-to-date records, this will not be possible. As a result, users may not trust the claims made by the developers.

Good documentation, therefore, is necessary to ensure accountability and transparency.

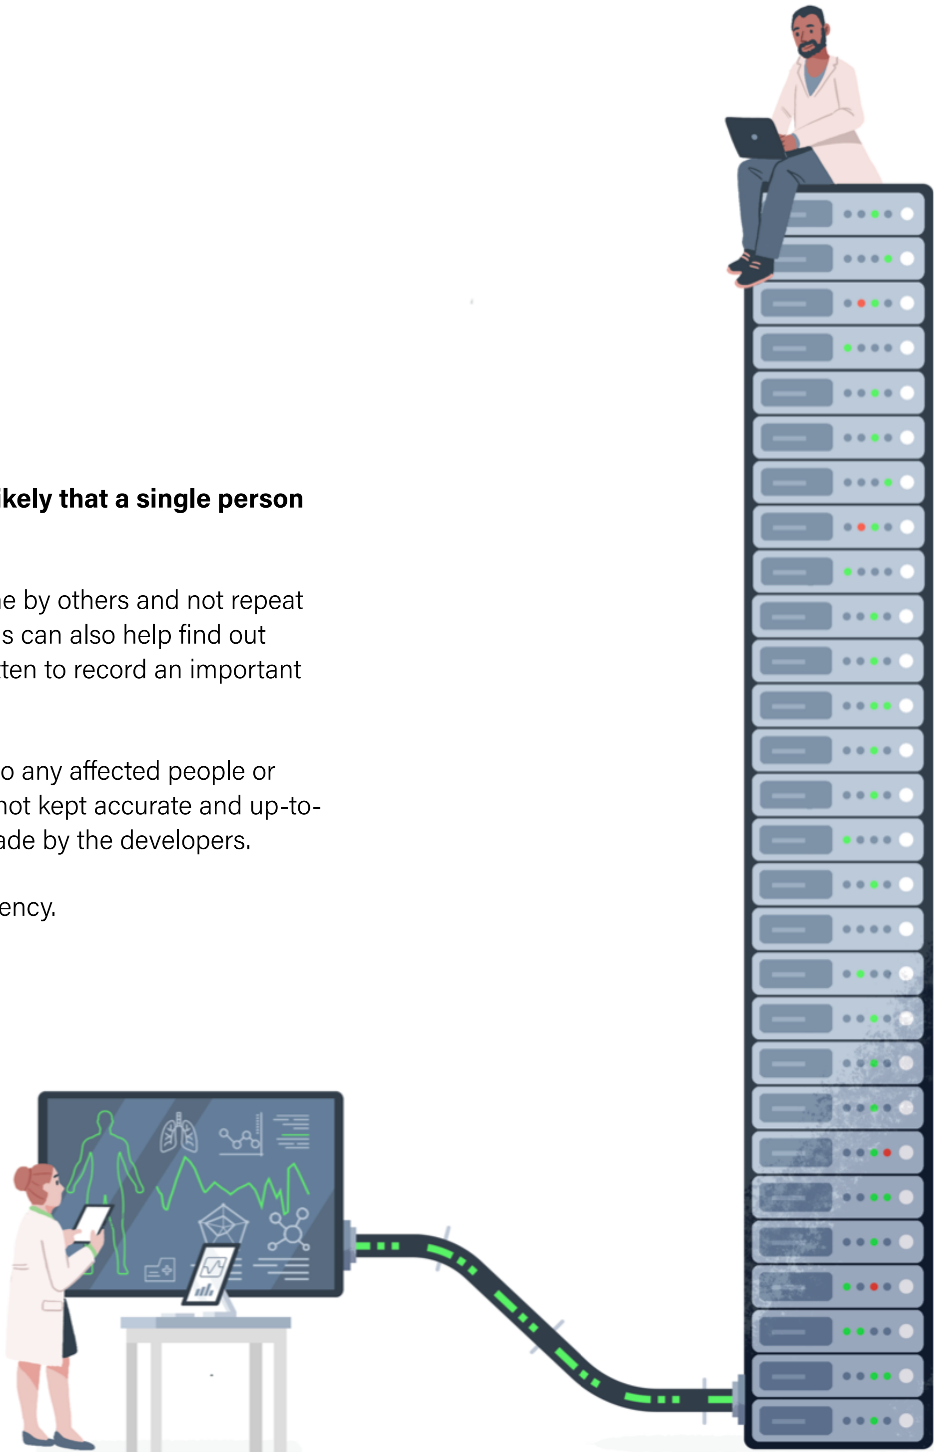

Case Study

The police force have documented all the stages of their project.

Their documentation includes the following information:

- who they consulted during the project planning,
- where they collected the data from, the steps they took to analyse the data,
- how they trained their algorithm, and
- how accurate the system is for detecting wanted suspects.

But when they deploy the system in the real world the system does not act as expected. The accuracy of the system is a lot worse than it was during their testing.

They return to their original documentation and realise that there are gaps in their original data. For example, when they use the system in the real world, it encounters more diverse groups of people than are present in their limited dataset. Specifically, it does not have many images of elderly people or people with darker skin colours. By consulting their documentation they realise this is the reason for the unexpected inaccuracy.

Topics for Reflection

1. What do you think the police force should do when they realise their system is not as accurate as initially expected?
2. Should they keep the system in use and collect more reliable data over a longer period?
3. Should they stop using the system for detection and just gather more data?
4. Should they stop the project altogether?
5. Or, is there something else they could do?

Data Reuse

Data may be collected and stored for an initial purpose, but this doesn't mean that it is only useful for this original purpose

We saw before that data is typically collected for a specific use or purpose. For example, helping your doctor make a diagnosis and prescription. But the same data could also be useful for other organisations. For instance, an insurance provider could use your health data to adjust the price of your policy.

Likewise, in principle, a public sector organisation may reuse data that was initially collected for a single project. For example, data that is used to determine which residents can access social care services may also prove useful for decisions about public spending.

An important principle is to collect only the smallest amount of data necessary and for a specific purpose. This is not just a matter of efficiency; it is also a legal rule (as we will see later). This principle of data minimisation can help prevent data from being reused for uses that fall outside of the scope of the original project.

Why does this matter? One possibility is that the data may have been collected from individuals who gave consent for a specific purpose. This does not guarantee they will consent for it to be reused for a different purpose. Again, you may consent to your doctor knowing sensitive information about your health, but feel unhappy about this data being shared with other organisations.

Case Study

Our police force have managed to improve the accuracy of the system and now want to expand its use to cover more areas.

This requires reusing the same data and system in different locations. This means they do not have to repeat the data collection stage. But because of the previous inaccuracy of the system they recognise that the (re)use of data may also create similar issues. So they decide to collect more data to supplement their existing data.

It is important to note here that the collection of citizen data for policing is rarely based on the consent of the individuals. Instead, police rely on a principle known as 'legitimate interest.' That is, they can collect data if doing so serves a legitimate interest that allows them to fulfil their duty as a public authority.

Other organisations can also appeal to 'legitimate interests' instead of consent. But this does not mean they can ignore the rights or freedoms of the people whom the data represents. Citizens also have their own interests, which may not align with the (legitimate) interests of the organisation.

Topics for Reflection

1. Let's assume you upload personal data to a social media platform to organise an event. How would you know that the data is only used for this purpose?

Data Monitoring & Deletion

Digital data doesn’t gather dust but there are still good reasons to monitor and delete it if it’s no longer serving an important function

We are now at the final stage of the data journey. At the start, we saw how data represents a property of an object or person. But many objects, and all people, change over time—especially the most fickle among us! As objects and people change, so do the properties that describe them.

Because data represent a snapshot in time it is important to keep track of the changes. If a system uses data to make decisions about people, old and inaccurate data can cause a person inconvenience or harm.

For example, let’s return to the doctor in our first example. If she reused the original data about your blood pressure they would be acting in an irresponsible manner. This is because a new measurement is necessary to know if your prescription is having the desired impact.

So too must public sector organisations monitor their data, and ensure it is still up-to-date. If it is not, they may have to delete the data—assuming it has served its original function—or store it in a secure manner for a fixed period of time. Depending on the system or project this will bring us back to the stage of project planning where we can consider what new data (or system) may be needed to replace the old one.

Case Study

The police collect images from public areas for the specific purpose of identifying wanted suspects. To do this the automated facial recognition system also has to collect images of members of the public who are not on a police database. Because the police have no legitimate use for this data they should delete it after they have determined that it does not match a wanted suspect.

The police also establish mechanisms to track the accuracy of their system. If it continues to be good enough they will redeploy it in new areas. But if it starts to perform poorly they will have to stop using it and plan a new project.

Topics for Reflection

1. The police force find out that holding on to all the images of the public improves the overall accuracy of their system. Many of these images capture innocent members of the public. The image are anonymous, but do you think the police should use and store these images after they have determined that the face does not match a suspect or criminal in their database? If not, why?

03

Automating Public Services

Benefits and Risks

Data-Driven Services

You have now been introduced to a case study that involved the development of an automated facial recognition system. The value of the data in this example is connected with the project goal of improving public safety. In this section we will explore whether there is a good reason to use technologies such as automated facial recognition in the first place.

We will consider other examples of data-driven public services across a variety of different domains.

Classification

Detecting Fraud

The type of domain in which data is used affects the types of risks and benefits that may occur. For example, inaccurate data records in healthcare are likely to cause more harm than in retail. In this section, we will focus on the function of the system though, instead of th domain.

We will explore three ways that data can support the automation of public services: classification, prediction, and recommendation. For each of these functions we will also consider several benefits and risks.

Classification is a process of assigning an object or person to a particular group. For example, does a particular animal belong to the group 'cats' or 'dogs.'

There are many ways to classify objects or people and many reasons why we may wish to classify them. For example, when you move house you may group items according to rooms (e.g. lounge, kitchen, bedroom). Or, a local council may use catchment areas to help classify school applications into 'approved' or 'rejected' groups.

Data can help automate and support classification tasks. To see how, let's continue to use the school application example.

To ensure school places are distributed fairly, a local council has to consider many features. These can include 'catchment area,' 'current siblings at school,' and 'medical needs.' The council also has to ensure that no fraudulent applications get approval.

Detecting fraudulent applications is a classification problem. The council has to classify applications according to two categories: 'fraudulent' and 'legitimate.'

Data can support this process. For instance, if an application claims that a child has many siblings at a school but the council's existing data suggests otherwise, this is one reason for suspecting a fraudulent application.

There are benefits and risk from automating this process. One benefit is that the application process is more efficient because of the speed at which the automated system runs. But there is also a risk that a small number of valid applications will be incorrectly rejected.

For instance, as pictured below, a council may have old records about a family who have recently moved or missing records about existing siblings at the school. Rejecting the applicaton because of inaccurate data could be a large inconvenience

Application

- Catchment Area: NW3
- Siblings: Two
- Medical Needs: None

Council Data

- Catchment Area: NW1
- Siblings: None
- Medical Needs: None

Some machine learning algorithms, known as 'neural networks' can classify objects in images.

However, it is not always easy to understand how they function.

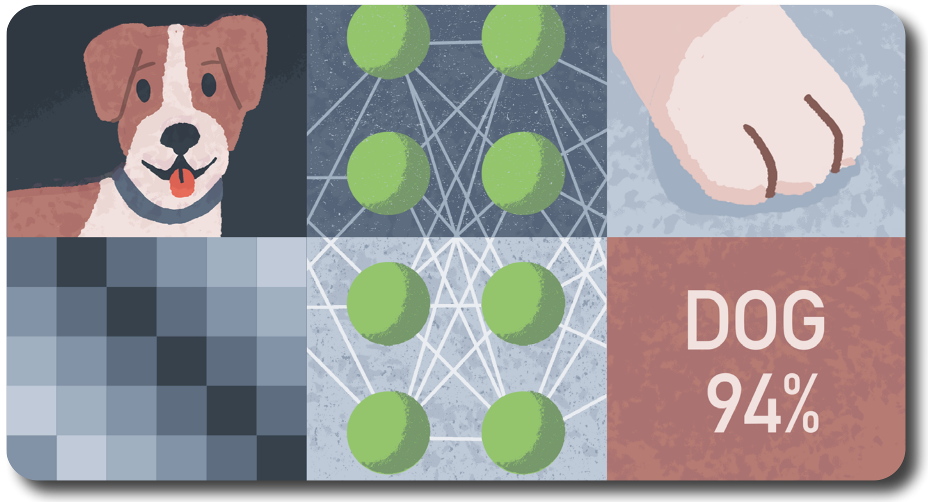

Risks and Benefits of Automated Classification

Risks

- Incorrect classification can cause inconvenience or harm to people.
- Classification systems may be unfair and lead to discrimination if they are more accurate for certain groups of people.
- Automation of tasks that are currently performed by humans can result in loss of jobs.

Benefits

- Automating simple classification tasks can improve the efficiency and effectiveness of local services.
- Automation of repetitive tasks can free people up to perform more rewarding jobs.
- When they are well designed, automated classification systems can reduce the impact of biased human decision-making.

or even affect the quality of education their child receives.

In this example, it would be sensible to ensure that applications that are classified as 'fraudulent' receive human review. This would provide an extra layer of accountability.

This type of human review is also important with the use of more advanced classification techniques. For instance, in the data journey example, classifying individuals as 'matches' or 'non matches' relied on complex automated facial recognition. The features used by such systems can be hard to understand. To see why this is the case, ask yourself how you recognise a friend in the street.

What features of their face do you use to identify them?

Automated facial recognition systems do not work the same way as our brains or visual systems. But both are very difficult to understand.

We will return to this issue in the next section.

Prediction

Supporting Vulnerable Residents

Data is a snapshot of an object or person at a moment in time. With enough relevant data it is possible to build up an intricate picture of the object or person and even predict future trends.

For example, with enough data about your past purchases a retailer could predict what you will buy on your next shopping trip.

Prediction is not an exact science though. Although the type and quantity of data will affect the accuracy of a prediction, there is always a chance that the prediction will be wrong. But informed predictions are sometimes necessary to deal with an uncertain future in a responsible way.

Policy-makers have to make predictions all the time. When they do, it is important they use the best data available. For example, predicting the effects of climate change means collecting data about the environment, such as traffic levels, water pollution, land use, and local and global temperatures. Local councils and service providers must also make predictions to ensure that they are meeting the needs of residents.

Let's consider an example of a local council trying to prioritise social care services.

Due to an increased demand for social care services and a decrease in public funding, cash-strapped local councils often have to prioritise their most vulnerable residents. To do this, they need to collect data about their residents. In particular, they will have to collect data that relates to so-called "risk factors".

What is a risk factor?

A risk factor is a specific feature or property that increases the risk of an outcome. For example, smoking is a risk factor for lung cancer and red meat consumption is a risk factor for cardiovascular disease. Collecting data about these risk factors can be helpful for making predictions about how likely an individual is to develop the associated condition.

Unfortunately, there is not always a straightforward link between a risk factor such as smoking and an outcome like lung cancer. For instance, poverty is a risk factor for many different outcomes, but is also an outcome with many other risk factors (see Image below). Complex relationships such as these make it hard to form reliable predictions based only on whether an individual falls below the poverty threshold.

It is a general principle of social care that the earlier an individual receives help the better their outcome will be. This applies to social support for both children and adults who, for example, may suffer from learning disabilities or physical impairments. Making predictions about who is vulnerable (or "at risk") is a good way to ensure that those most in need get help.

New technologies can help local councils and service providers make better predictions based on data about risk factors. Whereas a human may only be able to look at a small number of risk factors, a machine learning algorithm could analyse hundreds of risk factors for millions of people. The algorithm could use this data to form a prediction about who the most vulnerable individuals are. And a local council could use these predictions in various ways.

There are risks and benefits for automating predictions though. One benefit is that vulnerable individuals will receive timely and vital support. This could also happen without the need for the individual to come forward. For children who do not know how to get support this could be crucial.

But it can also be risky to approach an individual who has not come forward or asked for help. It is possible that they do not want to be identified and may feel as though their privacy has been violated. Or, worse, they could be a victim of domestic abuse and their identification could make their situation worse.

As with classification, it is often essential for a human to review automated predictions with care before making a decision.

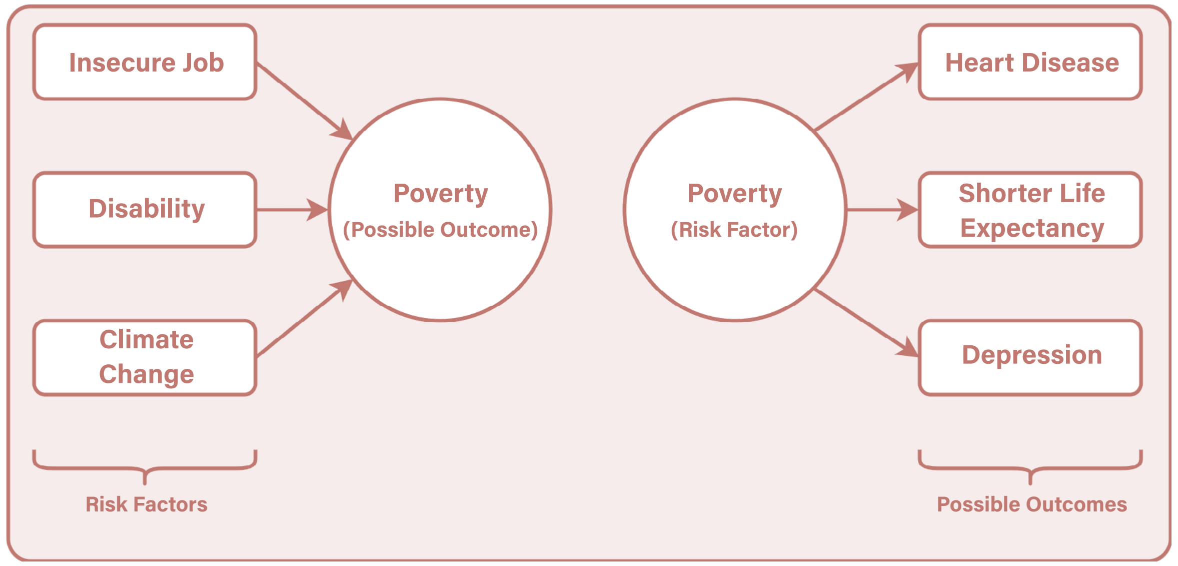

Poverty can be both a possible outcome of other risk factors and also a risk factor for other possible outcomes.

Risks and Benefits of Automated Prediction

Risks

- **Inaccurate predictions may lead to discriminatory outcomes if certain groups tend to be more affected by the errors than other groups.**
- **A prediction is based on past trends in the data. However, people typically want the opportunity to act differently in the future and not be shaped by their past behaviour.**

Benefits

- **Predictions can be used to design better systems that identify risk factors before they lead to serious harms.**
- **Machine learning algorithms can identify complicated patterns that lead to more accurate predictions that would be beyond the ability of most, if not all, humans.**

Recommendation

Electric Vehicle Recharging

Data has enabled many useful systems that automate the task of recommendation. Online shops can now make recommendations about which products you may like based on your past purchases. Media platforms can suggest new movies or television shows by comparing your viewing history to other users. And smart home speakers can play personalised music playlists in response to simple spoken requests, such as "play me some music."

Data has enabled many useful systems that automate the task of recommendation. Online shops can now make recommendations about which products you may like based on your past purchases. Media platforms can suggest new movies or television shows by comparing your viewing history to other users. And smart home speakers can play personalised music playlists in response to simple spoken requests, such as "play me some music."

Recommendation is a process of taking a large set of items and sorting or filtering them according to how well they meet some criteria.

We can refer to any technology that automates this process as a 'recommender system.'

For example, a social media site could use a recommender system to filter through all the posts of users that you follow and only show the ten posts that you are most likely to engage with. Machine learning algorithms are behind this recommendation process. The aim is to learn what interests you to keep you on the site, where you are more likely to see advertisements that generate revenue for the company.

Local councils can also use recommendation systems to help them decide which services they should invest in. To better understand this, let's use an example of installing new battery charging stations for electric vehicles.

A local council has a small budget to install 10 new charging stations but has received requests from over 100 residents. The council would like to meet all the needs of those who have requested the charging stations. But they also want to encourage more residents to switch to electric vehicles. This means they need to figure out where the 10 charging points would be best placed to meet both of these goals.

An algorithmic recommendation system could analyse massive amounts of data about (a) which streets the requests have come from, (b) where there are existing charging stations, and (c) where there is most need for electric vehicles due to high levels of traffic pollution. The system could then recommend the 10 locations that will best serve the 100 residents who requested them, while also ensuring that future residents can use them.

As before, there are also risks and benefits to using an automated system to decide where to place new charging points.

One benefit is that the automated system may be able to recommend novel locations for the charging points that are fairly distributed across the council, rather than being installed only on streets that belong to wealthy residents who have requested them. But the risk in doing this is that some of the charging points may end up in an area where there is no demand for electric charging. This would result in limited resources go to waste.

Risks and Benefits of Automated Recommendations

Risks

- **Recommendation is a process that involves filtering a *large* and *unordered* list of options into a *smaller* and *ordered* list of options (e.g. filtering the 10 most popular films from a list of 100). If the recommendation is not accurate, a user may not see options that they prefer or would benefit from.**
- **The reasons for making a recommendation may not be aligned with the interest of a user. For example, an online shop may recommend the items that they make the most profit from, rather than the ones that are the best for the customer.**

Benefits

- **Good recommendations can save people a lot of time. Imagine how long it would take to look through every film on Netflix or every television show on BBC iPlayer. The aim of the recommendation systems these platforms use is to show you the media that you are most likely to enjoy.**
- **The term 'serendipity' refers to the occurrence and development of events by chance in a happy or beneficial way. Automated recommendations can often be serendipitous, as they may expose us to new ideas, items, or options that we find rewarding, but which we would not have otherwise encountered.**

04

Evaluating Data Projects

Ethical, Legal, and Social Issues

Asking the right questions

This section will introduce you to a range of specific ethical, social, and legal issues that ought to be considered when evaluating the possible design, development, and deployment of a data-driven technology.

Data Ethics

Designing Fair Systems

When we discussed the data journey, the issue of varying levels of accuracy for different groups was raised. Different groups being subject to varying levels of accuracy can give rise to important questions about fairness. Specifically, does the variation in accuracy disproportiantely benefit or harm specific groups of individuals?

Why does fairness matter?

As a range of scientific studies and reports by journalists have documented, automated facial recognition systems are less accurate when used for individuals with darker skin than those with lighter skin. One reason for this is because the data that are used to train the systems are often insufficiently representative of the individuals that the system is actually used to classify. That is, if there is a lack of diversity in the initial set of images, then the final system may be very good at identifying light-skinned, young, males—who are well represented in the training data—but perform worse on other groups that are not well represented in the original dataset.

This means that individuals outside of this group could be disproportionately affected by their use. This scenario was actually documented in the case of passport photo checking system used by the Home Office. As the BBC reported,

> "Women with darker skin are more than twice as likely to be told their photos fail UK passport rules when they submit them online than lighter-skinned men, according to a BBC investigation." [1]

Biases with data-driven systems such as this one can lead to a range of social injustices, whereby already marginalised groups are further affected by the use of poorly designed systems.

To be clear, fairness is not just a matter of improving data representativeness—although this is vital. In addition, it is also a matter of making informed decisions about which projects to pursue in the first place. For example, if limited public resources are spent on a system that benefits a small sub-group of society that are already advantaged in various ways, then it could be argued that this is also an unfairness, especially if the money could have been spent in ways that helped disadvantaged or vulnerable members of society.

To help us answer this question, we must introduce four related concepts, which are set out in the following table.

True **Positive**	***True*** **Negative**
False Positive	*False* Negative

Each of these four options refers to a possible outcome of a data-driven technology that is used to classify objects or individuals, like the fraud detection system we saw in the previous section.

To make this clearer, let's assume that Naomi is at border control, and is using an e-passport gate that scans her face and attempts to match it with the photo on her passport. She is also in possession of her own, up-to-date passport. In contrast, Elliot is also at border control, but is using a passport that he has stolen.

Here, the positive/negative distinction refers to whether Naomi and Elliot are matched (positive) or not-matched (negative) with the respective passport. The true/false distinction refers to whether the outcome is correct or incorrect.

If Naomi is correctly matched with her passport, because she is the proper owner of the passport, then she will be let through border control. This is a *true positive*. However, if she is incorrectly flagged as not matching her passport, this could result in an unnecessary and inconvenient encounter with border security. This would be a false negative.

If Elliot is correctly identified as not-matching the passport, which he is in posession of, he will be stopped by security and, hopefully, arrested. This is a *true negative*. That is, it is true that he does not match the passport. However, if he is incorrectly flagged as matching the stolen passport, then he will be let through border control and not caught. This is a *false positive*.

These terms can sometimes be difficult to remember, and so they are listed in the following table to help:

	Positive (Match)	Negative (No Match)
Naomi	True	False
Elliot	False	True

We will often talk about the "accuracy" of a data-driven system, such as it's accuracy at classifying objects or making good recommendations. However, this term can often disguise inequalities in how the accuracy is distributed accross groups. In other words, a system can be 95% accurate overall, which seems like a good result, but the 5% of inaccuracies may only affect a specific group of individuals. That is, it may be 100% accurate for Group A, who make up 95% of the population, and 100% inaccurate for Group B, who make up the remaining 5%.

When framed this way, "95% accuracy" seems like a good way of hiding significant inequalities!

[1] https://www.bbc.co.uk/news/technology-54349538

Explaining Automated Decisions

> **"Why were you late?"**

It's a commonly asked question, and one which is in effect a request for an explanation. Perhaps it is asked out of concern, or perhaps because someone is fed up of the same old excuses! Regardless of the reason for asking, the ability to offer an explanation is vital to maintaining successful relationships with others.

Data-enabled systems can now be used to provide automated decisions. For example, banks can use predictive algorithms to determine if you are likely to pay back a loan; home smart speakers can automatically choose a playlist of songs for you; and AI technologies in healthcare can suggest possible treatments to doctors based on a patient's symptoms.

However, the ability to explain how these decisions were made is not always straightforward. For simplicity, let's label the ability to provide an explanation as the explainability a data-enabled system.

As we saw with the above example, explainability is a vital component of trustworthy relationships—would you, for example, be happy if your friend said that they were unable to offer you with a reason for why they were late? Would you be happy if you're bank denied you a loan and were unable to explain why? What if you're doctor assured you that the were recommending the right treatment for you, but could not back this up with detailed reasons?

In some cases, low explainability may not seem to matter too much. For example, when you ask your home speaker to play music, you probably don't want it to explain why it has recommended a particular song. But, if you're local healthcare provider denies you access to a service you desperately need then you will probably want to know why so you can seek redress. If they are unable to provide you with this explanation because they are relying on a data-enabled system that is too complex to be interpreted, this is likely to be a problem.

Let's take the example of a local healthcare provider choosing who should be prioritised for access to a service that is in high demand. Let's also assume that you are filling in an online application form, and the data you provide will be used to decide whether to grant you a loan.

The data that this form may request could include, among other things:

Personal Information

- Name
- Gender
- Date of Birth
- Postcode

Employment Status

Medical Records

- Current Prescriptions
- Existing Diagnoses

The healthcare provider's system will then use this data to sort you according to need, as a way to determine if you should receive support. But how does the system use the data to sort people, and which properties are actually important for the final decision?

Depending on type of system the healthcare provider may use, it can be very difficult to answer this question.

The details of why this is the case can be complex, but we can simplify some of the details by comparing the following three possibilities in the following table:

	System 1 (Least Accurate)	System 2	System 3 (Most Accurate)
Name	No Importance	Least Importance	No Information
Gender	Some Importance	-	No Information
Date of Birth	Important	-	No Information
Postcode	Some Importance	-	No Information
Employment Status	Important	-	No Information
Current Prescriptions	Some Importance	-	No Information
Existing Diagnoses	Very Important	Most Importance	No Information

All three of the systems pass an acceptable threshold for accuracy, but option 1 is the least accurate of the three systems. It occassionally approves people who do not require support and denies access to people who are most at risk. However, it also happens to give the most information about how it reaches this decision.

Option 3, in contrast, gives no information about how it reaches its decisions, as the algorithm it uses is too complex to be interpreted. As such, it is unable to support explanations that may be requested by customers. However, in contrast to option 1, it almost never provides the wrong decision and has also been shown to outperform human assessors.

Option 2 falls in between the other two, with some ability to support explanations and intermediate levels of accuracy.

These three options represent well-known trade-offs between accuracy and explainability in automated decision-making systems.

Here, the ethical challenge is when and where it is acceptable to lose some transparency over how a decision was reached (i.e. its explainability) in exchange for an improvement in accuracy?

This connects explainability to the previous issue of fairness, and also to an additional ethical need to support autonomy. That is, to what extent is an individual being provided with sufficient information to make a free decision, and if they are not, does this harm them in some way? We will explore this idea more in the next section.

Data and the Law

Equality and Human Rights

Human rights belong to every person in the world, regardless of specific details such as your gender, age, birth place, or beliefs. They are based on important moral principles like dignity, fairness, respect and equality.

Understand Human Rights

The Human Rights Act 1998 sets out the fundamental rights and freedoms that apply to all people in the UK. Its purpose is to ensure that the rights and freedoms of all UK citizens are legally enforceable. This means that if the rights of a citizen are violated they can take their case to court to seek justice.

The existence of fundamental rights and freedoms also places a duty on public bodies to respect your rights. Any new law or policy that is established must be compatible with the Human Rights Act 1998.

The UK's Human Rights Act 1998 incorporates the rights that are set out in the European Convention on Human Rights. This document was drafted in 1950 by the Council of Europe, following the atrocities of World War II.

> Despite the similar names and logos, the Council of Europe is not the same as the European Council. The European Council shapes the political direction and priorities of the European Union. Whereas the Council of Europe is an international organisation that upholds human rights, democracy and the rule of law in Europe.

Although the UK has left the European Union, it remains a member of the Council of Europe. This means that if a UK citizen has exhausted all options within the UK courts they can still appeal and have their case heard by the European Court of Human Rights.

In the previous section, we looked at how data-driven systems can give rise to ethical issues of fairness and explainability. These ethical issues also connect with legal issues such as human rights and freedoms.

Your rights and freedoms are backed by social institutions such as the courts system and also by charities and human-rights organisations such as the Equality and Human Rights Commission. There are numerous rights and freedoms that apply to all people. Some of these include:

- Right to life—nobody, including the Government, can try to end your life.
- Right to liberty and security—you cannot be imprisoned or detained without good reason.
- Right to respect for your private and family life—your private life, family life, and home life, must be respected and remain free from government interference.
- Freedom of expression—you have the right to express your views, such as in public protests or demonstrations

Because these rights and freedoms are designed to protect people they also affect the type of data that can be collected and how it can be used.

For instance, the collection of personal or sensitive data typically requires the informed consent of an individual. If a government were to spy on you at home in order to collect data about your activities, this would likely be a violation of your right to respect for your private and family life.

However, there are some cases where these rights and freedoms can be restricted, such as when the freedom of expression involves a hate crime or threatens national security. In these cases, a public body may not require the consent of an individual to collect data. Instead, they may appeal to alternatives such as the need to carry out a public task, comply with a legal obligation, or because of a genuine and legitimate interest in the data.

An example of a task that public bodies have to conduct, which provides them with a basis for collecting personal data is the Public Sector Equality Duty. This duty requires all public bodies to which it applies to:

- Eliminate unlawful discrimination, harassment and victimisation and other conduct prohibited by the Act.
- Advance equality of opportunity between people who share a protected characteristic and those who do not.
- Foster good relations between people who share a protected characteristic and those who do not.

What are the protected characteristics?

The Equality Act 2010 sets out the following protected characteristics:

- Age
- Disability
- Gender reassignment
- Marriage and civil partnership
- Pregnancy and maternity
- Race
- Religion or belief
- Sex
- Sexual orientation

They are protected in the sense that the law is designed to protect individuals from unfair treatment or discrimination on the basis of these characteristics.

This duty also extends to the use of data-enabled systems by public sector organisations when carrying out their functions.

For example, if a local council wishes to use a predictive algorithm to determine which schools are likely to underperform and need additional support, they are not only required to ensure that doing so does not lead to discrimination, on the basis of protected characteristics, but also requires them to consider how they could positively contribute to the advancement of social equality.

If the use of the system, for example, resulted in the withdrawal of funds from vital services that support children with learning disabilities, this may be grounds for unlawful discrimination, even if this outcome was not intended by a human decision-maker. One way to minimise the risk of such an outcome would be to carry out an Equality Impact Assessment. This is an activity that can be carried out prior to implementing a policy in order to help determine what impact the policy may have on social equality.

Data Privacy & Protection

At one end of the scale, poor data privacy practices can lead to some unfortunate but minor consequences. For example, perhaps you visit an online shop to buy a surprise present for your partner. You then suddenly find that all the computers on your shared home network begin to display adverts for this product, potentially ruining an important surprise.

At the other end of the scale, poor data privacy practices can lead to the disruption and undermining of vital democratic services. For instance, the now defunct consulting firm, Cambridge Analytica, were exposed by journalists for their abuse of data collected about users of Facebook. They used this data to develop predictive models that helped spread misinformation and influence a variety of elections or votes, including the Brexit referendum. Data privacy is increasingly important because of abuses such as this one.

Novel techniques, including some of the functions we saw in the previous section enable companies to link separate datasets together to learn new information.

A simple example can help to explain how this works. Consider the following pretend patient dataset:

Patient ID	Height	Sex	Age	Postcode	Medical Issues
#932745	5′ 7″	F	44	SW18	Asthma
#866180	**6′ 6″**	**M**	**23**	**SW6**	**Chronic Back Pain**
#039582	5′ 10″	M	55	SW11	Diabetes
#129034	5′ 3″	F	29	SW15	Depression

This is sensitive medical data. So it has been anonymised by using a patient ID instead of using the patient's name.

But one of the people that this data refers to could be identified if another dataset that contains similar properties was also obtained. Consider the following table, for example:

Name	Height	Age	Postcode	Sports Membership
Francesca Coates	5′ 6″	30	SW6	Yoga
Cameron Field	6′ 1″	41	SW7	Football
Mason Nash	**6′ 6″**	**23**	**SW6**	**Swimming**
Eva Coleman	5′ 9″	34	SW7	Climbing

In this example, we can see that Mason Nash is very likely the same individual from the other dataset, referred to using the Patient ID #866180. This is easy to do in the present case because it is highly unlikely that there will be two individuals of the same age living in the same postcode who happen to be 6′ 6″, and we are using very small datasets.

However, because of the vast amount of data that is available online, and the sophistication of modern machine learning, the risk of deanonymisation with real world datasets is also a genuine concern for data security experts. This problem has been made worse as a result of data breaches, such as the multiple breaches of NHS data in the last couple of years.

Therefore, it is important to protect private, personal, or sensitive information that relates to people. This is why an increasing amount of regulation has focused on protecting the rights of people to ensure that they can trust that their data is being used fairly and responsibly.

Although this is a part of respecting human rights in general, the UK also has specific data protection law—set out in the Data Protection Act 2018 and UK General Data Processing Regulation.

A lot of this regulation is aimed at organisations who collect, process or store personal data. In short, any information that could be used to identify a particular living individual. As we've just seen, what constitutes identifiable information can be tricky given new methods for linking data.

However, what's important for present purposes is the specific data protection rights that apply to you as an individual. These rights can summarised as follows:

- A right to be informed if your personal data is being used
- A right to access a copy of your data
- A right to have incorrect or inaccurate data corrected
- A right to have your data deleted
- A right to portable data that can be transferred to a different location
- A right to object to the use of your data
- A right to understand the reasons behind an automated decision made using your data and to not be subjected to automated decision-making
- A right to access information from a public body
- A right to raise a concern about how an organisation is using your data

All of these rights have restrictions, which is why it is important to understand their scope. The Information Commissioner's Office has an excellent introduction to these rights:

https://ico.org.uk/your-data-matters/

The term 'data subject' can seem a little impersonal. It is used to refer to the individual whom particular personal data is about. In this guide, where possible, we will use 'person', 'individual', and sometimes 'user' (where appropriate).

Unfortunately, because these rights are limited, there is no guarantee that the manner in which an organisation chooses to comply with the law will necessarily be in the best interest of you as an individual.

By now, the most notorious instance of data protection compliance that burdens individuals has got to be the annoying cookie pop-ups that appear when you visit a website (see the box on the following page). Unfortunately, without additional pressure from informed and empowered individuals, things are unlikely to change.

Cookies and Data Surveillance

By now, I'm sure we all feel the need to go on a cookie diet! I'm referring, of course, not to the delicious biscuits but to the small data files that are stored on your computer when you browse the internet that allow websites to keep track of information about you (e.g. what computer or device you are using, where you are located). Some of these cookies provide an essential service and are necessary for modern websites to function correctly.

However, other cookies allow companies such as Facebook and Google to track you across their vast online advertising networks. When you visit a website that is part of Facebook's advertising network, for example, you are sharing information with both the company that owns the website as well as Facebook.

Perhaps you are happy with your data being shared, and believe you have nothing to hide. However, it is not just criminals, journalists, or activists who may have legitimate reasons for being careful about what they share online.

Machine learning algorithms can use the vast quantities of data that are collected about us when we go online to predict surprising details about us. If enough data is collected about what we post on social media, what products we purchase, or which sites we visit frequently, fairly accurate machine predictions can be made about us, including our sexual orientation, religious beliefs, age, who we are likely to vote for, whether we have a physical or mental health disability, and who our friends and family members are. You may recall that each item on this list is a protected characteristic. And yet, machine learning algorithms can predict this information in some cases on the basis of unrelated data.

It is difficult to know what you are revealing about yourself online, or who may be watching.

Using Data for Social Good

Sustainable Data Use

Because of the ongoing climate crisis we tend to associated sustainability with care and consideration for the envrionment. The collection and use of data can cause harm to the environment. For instance, the servers that are used to store the world's data often fill vast warehouses, known as data centres. Many of these data centres have a footprint larger than 100 football fields put together (see imagees across the page).

The use of complex AI technologies can also require vast amounts of energy to train and run. Although renewable energies such as solar and wind carry the promise of being able to offset some of this usage, our reliance on data and data-driven technologies still has a sizeable impact on the environment.

There is another important sense in which the term 'sustainable' applies to data though. It can be costly to collect, analyse, and use data, and to build sophisticated data-driven technologies. When public bodies choose to carry out a data-driven project, therefore, they must consider the sustainability of the project and the technology themselves. Otherwise, they risk wasting limited public funds.

During project planning, it is important to consider whether there is sufficient funding and resources to deliver a succesful and effective project. It is also important to determine if there is longer-term support to ensure the project is maintained.

In the case of data-driven technologies like machine learning algorithms, the longer-term monitoring and support of a system has a particular importance for sustainability.

As their name implies, machine learning algorithms get better at some task (i.e., they learn) as more data is collected. When the algorithms are first being trained there is often close human monitoring to ensure they work as intended. For example, learning to classify images correctly. However, once they are put to work in the real-world, it is just as important to continue to monitor their performance.

Microsoft found this out the hard way when they released an AI, called Taybot, that could chat with people online, and learned what responses to give based on human conversations. Unfortunately, becuase the AI learned on the basis of online conversations the system ended up giving highly offensive and inflammatory responses. Microsoft quickly took the system down and apologised, but the AI system remains a stark example of how a system may work well in testing but not be sustainable when used in the real world.

As public bodies start to use data-driven technologies in more areas of society they must also consider the long-term sustainability of their projects. This means monitoring the accuracy of a system over time, protecting private data from hackers or other security vulnerabilities, and also ensuring that the technology continues to serve its original purpose. If not, then the technology may have to be removed or replaced. However, if the technology serves a vital public function, such as predicting social care services, this can only be done when there is a suitable replacement.

Google's Datacenter in Eemshaven, Netherlands. Access to solar, hydro, and wind energy is increasingly important to minimise the environmental impact of these sites.

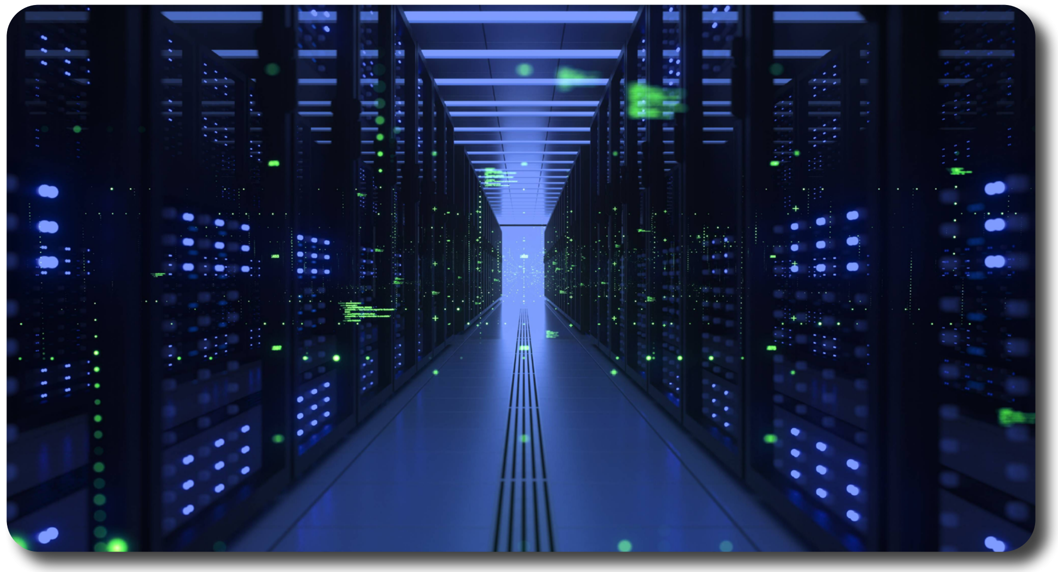

Data centres house vast numbers of computer servers that can be used for many intensive tasks involving data storage, analysis, and use. The largest data centres can cover several millions of square feet.

Inclusive and Representative Data

Using data for social good requires the active engagement and participation of citizens to help evaluate and manage the societal impact of data-driven technology (e.g. ensuring benefits are realised and harms are mitigated). However, there are often a variety of barriers that prevent or undermine the engagement and participation of individual users or communities.

Some of these barriers prevent data from being collected about certain residents or stakeholders in the first place (see 'Sourcing from the crowd'). However, even where data are collected from all relevant stakeholders, there is no guarantee that the dataset can be said to be fully representative.

Consider the following form that could be used to collect data from local residents:

Please enter/select the following information.

Name (up to 20 characters)

Age

18-25
26-30
31-40
41-50
50+

Gender

Male
Female

A dataset that was generated on the basis of this form would suffer from several representational issues. First, the limitation on characters for names may prevent individuals from specific ethnic backgrounds where longer names are common from entering their full name.

Second, the age brackets are incomplete and may also be too coarse grained, depending on what use the data are being gathered for.

And, finally, the form fails to acknowledge individuals who are non-binary from accurately recording their gender.

It can be easy to overlook the importance of ensuring that all individuals have the means to be accurately represented and included in datasets. However, the repeated failure to recognise certain groups of individuals is a social injustice that prevents people from participating eqaully in democratic practices.

In addition, even if we design inclusive forms, there is also the question about who is able to contribute and participate with the associated activity. Even the most well-designed and inclusive form would be wasted if there are significant barriers that prevent certain groups from participating (see the box on the right).

Sourcing from the Crowd

The term 'crowdsourcing' refers to projects that obtain work, opinions, financing, and data from a large group of people (i.e. the crowd). The Internet has made crowdsouring an increasingly popular means for supporting the functions or goals of various projects. This can range from the real-time feedback of traffic incidents by users of navigation apps, the submission of user-generated reviews to websites that catalogue and rank restaurants, and even the reporting of bee hive locations to help researchers and ecologists understand the effect of climate change.

For some projects, it may not matter who submits data. But for other projects, barriers to participation can result in a lack of inclusivity and representation that can severely undermine the overall goals of the project. For example, let's assume that a local council decides to partner with a technology company that provides open-source map software. The council then suggest that local residents use an app to add and review services that they find useful, as a way of better understanding how services are used by the local community (e.g. access to parking, or use of public restrooms).

Understandably, residents will tend to list and review services that they find useful. However, this can result in what is known as "missing data". That is, data about the services that serve a useful function for a particular group of residents that either do not know about this app, or perhaps cannot afford to use the app due to monetary or time constraints. The result is that their data is not shared with the council, and is, therefore, missing from the suggestions. Missing data is a big problem for crowdsourcing, and also for designing inclusive policies because it can be very difficult to know who is missing from your data if you don't have any information about them in the first place.

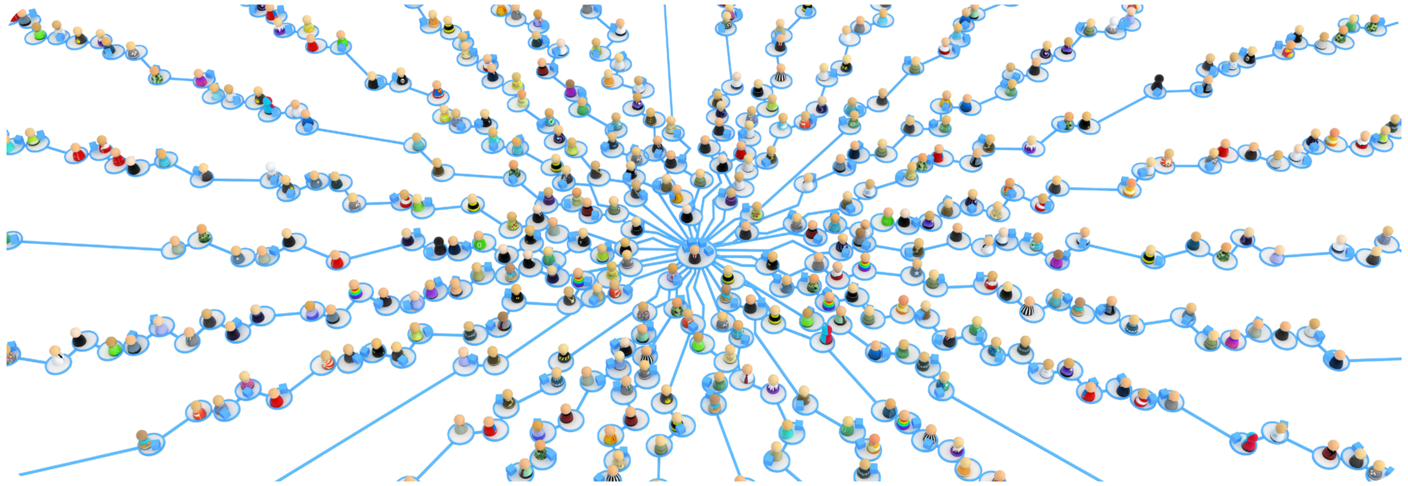

The Alan Turing Institute

A Citizen's Guide to Data