

Contents	2
About this guidance	5
Key data protection concepts	7
Biometric recognition	10
How do we demonstrate our compliance with our data protection obligations?	20
How do we process biometric data lawfully?	27
How do we process biometric data fairly?	36
How does the accuracy principle apply to biometric data?	41
How do we ensure our processing of biometric data is transparent?	43
How do we consider rights requests for biometric data?	46
How do we keep biometric data secure?	51

Contents

About this guidance

- Why have you produced this guidance?
- Who is it for?
- What does it cover?
- What doesn't it cover?
- How should we use this guidance?

Key data protection concepts

- What is personal information?
- What is biometric data?
- What is special category biometric data?
- What about other special category information?
- Are we processing personal information if we delete it quickly?

Biometric recognition

- What do you mean by "biometric recognition"?
- How do biometric recognition systems work?
 - Biometric capture, feature extraction and template creation
 - Comparison
 - Decision-making and thresholds
- What can we use biometric recognition systems for?
- Do biometric recognition systems use personal information?
- Do biometric recognition systems use biometric data?
 - The personal information is about someone's physical, physiological or behavioural characteristics
 - The personal information results from specific technical processing
 - The personal information allows or confirms someone's unique identification
- Do biometric recognition systems use special category biometric data?

How do we demonstrate our compliance with our data protection obligations?

- Who is the controller for our biometric recognition system?
- Data protection by design and default

- Do we need to do a DPIA?
- What risks to rights and freedoms should we consider?
 - Risks resulting from personal data breaches
 - Risks resulting from biometric false acceptance or rejection
 - Risk of discrimination
 - Risks resulting from systematic monitoring of public spaces

How do we process biometric data lawfully?

- How do we determine our lawful basis and special category condition?
- Can we use consent and explicit consent?
- What other lawful bases might apply?
- What other special category conditions might apply?
 - Substantial public interest
 - Research
- What if we don't have explicit consent and no other condition applies?

How do we process biometric data fairly?

- How does the statistical accuracy of biometric algorithms relate to the fairness principle?
 - How do we deal with risks resulting from errors?
- How do we deal with the risk of discrimination?

How does the accuracy principle apply to biometric data?

- Does the accuracy principle apply to our biometric recognition system?
- When do we need to collect new biometric samples?

How do we ensure our processing of biometric data is transparent?

- What information do we have to share to comply with the transparency people?
- How should we provide transparency information?

How do we consider rights requests for biometric data?

- What is the right of access?
- What is the right to rectification?
- What is the right to erasure?

- What is the right to data portability?
- What is the right to object?
- Can we use a biometric recognition system to make automated decisions about someone?

How do we keep biometric data secure?

- What are appropriate security measures?
- Can PETs help us comply with our data protection obligations?
 - Biometric template protection
 - Personal information in biometric samples and templates
 - On-device verification
- How can we comply with the data minimisation and storage limitation principles?
- Should we retain biometric samples?

About this guidance

In detail

- [Why have you produced this guidance?](#)
- [Who is it for?](#)
- [What does it cover?](#)
- [What doesn't it cover?](#)
- [How should we use this guidance?](#)

Why have you produced this guidance?

This guidance explains how data protection law applies when you use biometric data in biometric recognition systems. Read it to understand the law and our recommendations for good practice.

Who is it for?

This guidance is primarily for organisations that use or are considering using biometric recognition systems. It is also for providers of these systems (this could include vendors and developers). It therefore applies to controllers, processors and relevant third parties.

What does it cover?

This guidance looks at the definition of biometric data under the UK GDPR. It also focuses on biometric recognition uses and explains how these involve processing special category biometric data.

This guidance covers:

- what biometric data is;
- when it is considered special category data;
- its use in biometric recognition systems; and
- the data protection requirements you need to comply with.

What doesn't it cover?

This guidance does not cover requirements of the data protection regimes for law enforcement purposes or the security services. However, some of the principles explained in this guidance are relevant to these regimes too.

This guidance is intended to highlight the considerations you should give to biometric data when you use biometric recognition systems. It is not intended to be a comprehensive guide to compliance. Where this guidance refers to principles already addressed in our other guidance, we provide links to the relevant further reading.

This guidance does not cover the use of biometric classification or categorisation systems. These systems

make inferences about people based on observable characteristics and will be addressed in the next phase of our biometric technologies project through separate guidance. We will publish this guidance by the end of 2024.

This guidance also does not consider future changes to data protection legislation. We will update this guidance where necessary in response to any future changes to data protection law.

How should we use this guidance?

To help you to understand the law and good practice as clearly as possible, this guidance says what organisations **must**, **should**, and **could** do to comply.

Legislative requirements

- **Must** refers to legislative requirements.

Good practice

- **Should** does not refer to a legislative requirement, but what we expect you to do to comply effectively with the law. You should do this unless there is a good reason not to. If you choose to take a different approach, you must be able to demonstrate that this approach also complies with the law.
- **Could** refers to an option or example that you could consider to help you to comply effectively. There are likely to be various other ways you could comply.

Key data protection concepts

At a glance

- Personal information relates to an identified or identifiable person.
- If you can identify someone from the information directly or indirectly, then it is personal information.
- Biometric data is a type of personal information.
- Personal information must meet specific requirements to be biometric data. These relate to qualities of the information itself, not how you use it.
- If you use biometric data for unique identification, it is special category biometric data.

In detail

- [What is personal information?](#)
- [What is biometric data?](#)
- [What is special category biometric data?](#)
- [What about other special category information?](#)
- [Are we processing personal information if we delete it quickly?](#)

This guidance refers to the following data protection concepts. An understanding of these will help you to get the most out of this guidance.

What is personal information?

Personal information relates to an identified or identifiable person.

It is known as “personal data” in the UK GDPR and is defined in Article 4(1) as:



“any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.”

If you cannot directly identify someone from information you hold, it may still be possible to indirectly identify them. You **must** consider the information you are using, together with all the means you, or anyone else, is reasonably likely to use to identify that person.

Information that is not personal information (ie it does not relate to an identifiable person) is outside the scope of data protection law.

Personal information does not include information:

- about the deceased; or
- that has been anonymised appropriately.

Further reading

- [What is personal information: a guide](#)

What is biometric data?

Biometric data is a type of personal information. Article 4(14) of the UK GDPR defines biometric data as:



“personal data [resulting from specific technical processing](#) relating to the [physical, physiological or behavioural characteristics](#) of a natural person, which [allow or confirm someone’s unique identification](#) of that natural person, such as facial images or dactyloscopic [fingerprint] data.”

This means that personal information is only biometric data if it:

- relates to someone’s physical, physiological or behavioural characteristics (eg the way someone types, a person’s voice, fingerprints, or face);
- has been processed using specific technologies (eg an audio recording of someone talking is analysed with specific software to detect qualities like tone, pitch, accents and inflections); and
- can uniquely identify (recognise) the person it relates to.

See [Do biometric systems use biometric data?](#) for more information.

If you are using information that does not meet these criteria, you **must** still determine if you are processing personal information.

What is special category biometric data?

Article 9 of the UK GDPR singles out some types of personal information as more sensitive and gives them extra protection.

These include biometric data when used for the purpose of uniquely identifying someone. In this guidance, we use the term “special category biometric data” to refer to this.

Not all biometric data is automatically special category biometric data. It only becomes this if you use it to uniquely identify someone.

Your purpose for processing biometric data is therefore important. It defines whether you’re processing special category biometric data.

See [Do biometric recognition systems use special category biometric data?](#) for more information.

You **must** only process special category information if you can identify a valid condition for processing it.

See [How do we process biometric data lawfully?](#) for more information.

What about other special category information?

Not all biometric data is “special category biometric data”. This only applies if you use it, or intend to use it, to uniquely identify someone. However, even if this is not your purpose, the biometric data you process may still be considered another type of special category information. For example, you could use biometric data to infer someone's racial or ethnic origin or consider it as health data.

Whether you consider this to be special category information depends on if you intend to infer this information from the biometric data.

Further reading

- [Special category data](#)

Are we processing personal information if we delete it quickly?

Processing means taking any action with someone’s personal information. This includes collecting, storing and deleting personal information. It is still processing if you only briefly create, collect or store information (sometimes known as transient processing).

In many cases, technologies using biometric data process it transiently. This information may only exist for a fraction of a second. Data protection requirements still apply to you, regardless of how quickly you may delete it.

Transient processing can help you adopt a data protection by design approach by demonstrating compliance with data protection principles, such as data minimisation, storage limitation and security.

Biometric recognition

At a glance

- Biometric recognition describes when you use biometric data to uniquely identify someone.
- It is a term used in industry standards and isn't defined in data protection law.
- Biometric recognition uses personal information, biometric data and special category biometric data.
- If you are using a biometric recognition system, you are processing special category biometric data.

In detail

- [What do you mean by "biometric recognition"?](#)
- [How do biometric recognition systems work?](#)
- [What can we use biometric recognition systems for?](#)
- [Do biometric recognition systems use personal information?](#)
- [Do biometric recognition systems use biometric data?](#)
- [Do biometric recognition systems use special category biometric data?](#)

What do you mean by "biometric recognition"?

"Biometric recognition" is not a term defined in data protection law.

Biometric recognition, as defined by the International Standards Organisation (ISO) in ISO/IEC 2382-37:2022(E), refers to the automated recognition of people based on their biological or behavioural characteristics. This aligns closely with the definition of special category biometric data in the UK GDPR.

If you use a biometric recognition system, you are using biometric data to uniquely identify someone. So, "biometric recognition" encompasses all situations in which biometric data is special category biometric data.

In this guidance, we use the term biometric recognition because:

- it aligns with the definition of processing special category biometric data;
- some readers are likely to be more familiar with the ISO's definition than the scope of processing special category biometric data; and
- our previous work on biometrics found that a lack of clarity in terminology was causing confusion around how data protection law applied to biometric data.

Further reading

- [ICO biometrics reports](#)
- [Special category data](#)
- [ISO/IEC 2382-37:2022](#) [↗](#) establishes a systematic description and vocabulary for the field of

How do biometric recognition systems work?

Biometric capture, feature extraction and template creation

Biometric recognition begins with biometric capture. Biometric capture is the process of recording information relating to someone's physical, biological or behavioural characteristics. This can be either be done directly from a person, or from an existing representation of those characteristics, such as a photograph.

Biometric capture creates a biometric sample. Examples of biometric samples include:

- an image of someone's face in a digital photograph;
- a recording of someone talking; or
- a video of them walking.

The key information extracted from a biometric sample is a biometric feature. A biometric feature is a digital summary of how a person's characteristics make them unique. Biometric features often take the form of a string of numbers, and do not visually resemble the characteristics they describe. This is because they are intended to be readable by biometric algorithms, not by people.

Biometric algorithms are sets of rules that determine what automatically happens to biometric samples. For example, biometric algorithms determine how features are extracted from a sample.

When features are stored for reference, they become a biometric template. The process of creating a template for reference and associating it with a person is known as enrolment.

Both templates and samples can serve as biometric references, against which queries about someone's identity can be checked.

Comparison

Biometric recognition systems work by comparing two sets of biometric features. The intended outcome of this comparison process is to establish how likely it is the two sets of features belong to the same person.

The first set of features is extracted from a biometric reference. The second set of features is extracted from a newly created sample. This new sample is known as a biometric query or biometric probe.

Depending on the specific use-case, a biometric recognition system may compare features from a probe against those from a single reference, or from many references in a database (or watchlist). See [What can we use biometric recognition systems for?](#) for more information.

Biometric probes and biometric references are never exactly the same, even when they belong to the same person. This means that there will always be some variation between the probe and reference. Biometric recognition systems can therefore never be certain about an individual's identity.

Comparison is a statistically informed estimate of similarity (or dissimilarity) between a biometric probe and a reference. The comparison process produces an estimate based on the probability of whether the probe

and the reference belong to the same person.

Decision-making and thresholds

A threshold is the point at which a biometric recognition system considers this similarity to be statistically significant.

A threshold may be set by default by a biometric recognition system, or you may be able to vary it to meet your specific purposes.

A comparison resulting in acceptance means that the similarity between the probe and reference has met the threshold. This suggests that a probe and a reference relate to the same person.

Similarly, rejection means that the similarity between the probe and reference has not reached the threshold, which suggests that a probe and a reference do not belong to the same person.

The lower the threshold, the higher the chance that an accepted comparison may not actually relate to the same person. See [Risks resulting from biometric false acceptance or rejection](#) for more information on false biometric acceptance.

However, you **should** always take care when interpreting the acceptance and rejection decisions made by a biometric recognition system. This is because these decisions are statistically informed estimates based on probability, and there is always some degree of error that cannot be removed entirely.

Also, there are many real-world factors that can increase the probability of a rejection, such as environmental conditions. Therefore, you **should not** interpret the outcome of a comparison as a matter of fact.

What happens because of a threshold being met or not should depend on your circumstances.

For example, it depends on:

- what impact the decision could have on the person;
- what your use case and context is;
- what safeguards you have in place; and
- whether the system is making solely automated decisions.

It is important to understand:

- how the process of comparison works;
- what threshold your system works to; and
- whether your circumstances requires further steps to confirm someone's identity.

See [What risks to rights and freedoms should we consider?](#) for more information about how errors made by biometric recognition systems can result in harm.

See [How do we process biometric data fairly?](#) for more information on how to address the risks arising from the errors that biometric recognition systems can make.

Further reading

What can we use biometric recognition systems for?

Biometric recognition includes both biometric identification and biometric verification.

Identification refers to a one-to-many matching process (1:N, where N is the number of biometric references in a database). A biometric probe derived from one person is compared with many references in a database. It asks the question “Who is this person?,” or “Do we know this person?”.

For identification, the biometric reference of the person you are trying to identify must be in the database of biometric references that you are searching (ie they must have been enrolled).

While facial recognition technology is probably the best-known use of biometric recognition to identify someone, there are many other biometric approaches that are capable of identification.

Verification refers to a one-to-one matching process (1:1). A biometric probe compared against a single biometric reference (ie a biometric template, or a biometric sample such as a passport photo). It asks the question “Is this person who they claim to be?”. Passport eGates are an example of biometric verification.

The term “authentication” has historically been used about both identification and verification, sometimes interchangeably. Current industry standards move away from using this term and our guidance reflects this.

Example

An employer provides work devices to its employees. The devices include an on-device biometric recognition feature.

The employer offers its employees a free choice over whether to use this feature as an alternative to a password to access their work account.

This involves creating a biometric template that is stored on the device as a biometric reference.

Each time the employee wants to access the device, the biometric recognition system creates a biometric probe from a newly captured biometric sample (image of their face).

It then compares the biometric probe with the stored template. If the two match, the employee can then access their work account.

This means the employee’s biometric data is processed for the purpose of uniquely identifying them. Even if the comparison is unsuccessful and the employee has to enter a PIN or password instead.

Further reading

Both biometric identification and biometric verification are often used to control access to virtual or physical spaces. In these scenarios, biometric recognition systems replace a password (something you know) or a swipe card (something you have) with biometric data (something you are). Similarly, biometric verification is increasingly being used to control access to digital services, mobile devices and computers.

Biometric identification is also sometimes used to check whether someone signing up for a service has already registered under a different identity.

Both identification and verification require biometric data to uniquely identify someone.

Example

A rental company requires customers to prove that they have a valid driving licence prior to using their services. To make the check-in process quicker and more convenient for customers, the company offers customers a way to do it online before they check in.

If a customer decides to use this process, they are given the option to prove their details by uploading a scan of their driving licence and another photo of themselves. The company then uses a biometric recognition system to compare the two images and verify that they are of the same person.

This process involves processing special category biometric data to uniquely identify someone. In this case, it involves comparing biometric data generated from both photos to verify the identity of the customer.

The scalability of biometric recognition systems may be attractive in comparison to traditional access control systems that incur fixed costs (eg the need to issue new or replacement identity cards).

Biometric recognition systems may also be more secure than swipe cards or PINs from the perspective of controlling access. For example, people can't forget or lose their biometric data but can share or misuse cards or PINs.

However, unlike access control technologies requiring PINs and passwords, biometric recognition systems are reliant on special category biometric data to function.

Do biometric recognition systems use personal information?

Yes. If you use a biometric recognition system, then you are processing personal information.

The purpose of any biometric recognition system is to recognise someone. This includes checking someone is the person they claim to be (verification) and checking whether they match anyone in a database (identification).

Both processes require information about an identified or identifiable person.

Because biometric samples contain information relating to identified or identifiable people, they are personal data under data protection law.

Do biometric recognition systems use biometric data?

Yes. If you use a biometric recognition system, you are also using biometric data.

This is because biometric recognition systems process personal information that meets all three parts of the definition of biometric data in data protection law, which are listed below.

The personal information relates to someone's physical, physiological or behavioural characteristics

"Physical and physiological" means someone's biological characteristics.

These characteristics can include a person's facial features, friction ridges on their fingers (which are what create our fingerprints), iris, voice and even their ear shape.

Examples of physical or physiological biometric recognition techniques include:

- facial recognition;
- fingerprint recognition;
- iris recognition;
- voice recognition; and
- ear recognition.

Physical characteristics are used by biometric recognition systems because they can provide a lot of information that typically varies from person to person.

"Behavioural" is about biometric characteristics that relate to things like movements, gestures or motor skills. Behavioural characteristics can include a person's handwriting, how they type, their gait when walking or running and their eye movements.

Examples of behavioural biometric recognition techniques include:

- keystroke recognition;
- handwritten signature recognition;
- gait recognition; and
- gaze-based recognition.

The personal information results from specific technical processing

The term "specific technical processing" describes a processing operation – or set of operations – that can be applied to a person's physical, physiological or behavioural characteristics, which makes it possible to uniquely identify them.

This means there's a difference between things like "ordinary" digital images and how these may be used in

the context of biometrics. For example, data protection law says that photographs:



“are covered by the definition of biometric data only when processed through a specific technical means allowing the unique authentication of a natural person.”

While someone's physical characteristics may be shown in a photo, this isn't enough to make that photo biometric data. It's only when something else happens to that photo – a discrete processing operation or set of operations that result in something that allows or confirms someone's unique identification – that the result becomes biometric data.

For example, if specific techniques are applied to the photo to extract someone's facial features, then the photo has been "processed through a specific technical means" that allows the person to be uniquely identified.

The information resulting from specific technical processing not only describes the end result of the processing (ie a biometric template or a biometric probe). It also covers any information produced by these specific technical processes, regardless of how long it exists for. If this information is capable of uniquely identifying someone, it is biometric data.

The term "specific technical processing" can also refer to the main stages involved in biometric recognition systems. For example:

- biometric feature extraction, where the information in the biometric sample is extracted and transformed by an algorithm into a biometric feature; and
- biometric template generation, where a biometric feature is stored as a biometric template.

The personal information allows or confirms someone's unique identification.

This is about the properties of the information itself, not what you intend to use it for.

In data protection law, if you can distinguish someone from other people, then that person is "identified" or is "identifiable". This may be from the information itself, other information you may hold, or other information someone else may have.

However, the term "unique identification" is slightly different. Unique identification, as described in UK caselaw, refers to someone being singled out with accuracy, (ie where they are distinguished from others with a level of precision).

Unique identification for the purposes of biometric data does not consider other sources of information you may hold or might be available. It is about whether someone can be directly identified from that information, with accuracy. It therefore differs from the question of identifiability and personal information.

All biometric recognition systems process biometric data. This is the information that allows them to uniquely identify someone. Any attempt at unique identification doesn't have to be successful to meet this definition.

Even if you do not intend to use biometric data to identify someone, the properties of biometric data mean that you can use it for this purpose. The wording “allow or confirm” means that you will meet this part of the definition if it is possible to identify someone, even if this is not your intention.

Example

An employer may be able to identify a staff member from an audio recording of a meeting, even if the staff member didn't state their name. The recording therefore includes personal information about that staff member.

However, this doesn't make the audio recording biometric data, as it doesn't result from specific technical processing of the staff member's characteristics (ie their voice).


The same organisation then buys a voice recognition solution to transcribe audio recordings and attribute what was said to particular people who attend the meetings.

This first involves enrolling all meeting attendees onto the system. To do this, a sample is captured – either directly from a person or from a voice recording. Biometric features are then extracted from the sample and stored as a biometric template. It then requires the voice recognition solution to capture new biometric samples, create biometric probes and compare these against the stored templates.

All of these stages involve biometric data. The information results from specific technical processing of someone's characteristics and can be used to directly identify them with a degree of accuracy.

As the employer intends to process the biometric data for the purpose of uniquely identifying the meeting attendees, it is therefore also special category biometric data.

Further reading

- [Bridges v South Wales Police \(EWCA Civ 1058 – C1/2019/2670\)](#)  UK caselaw referring to the concept of unique identification.

Do biometric recognition systems use special category biometric data?

Yes. If you use a biometric recognition system, you are using special category biometric data. This is because the purpose of biometric recognition systems is to uniquely identify someone using biometric data.

The UK GDPR says that biometric data is special category data if it is processed:



“for the purpose of uniquely identifying a natural person.”

This is slightly different from the definition of biometric data. Instead, it is specifically about the purpose you intend to use that information for.

This makes special category biometric data different to the other special categories of information. For example, political opinions or racial origin are about the nature of the information alone, rather than any additional consideration of the purposes you are processing the information for.

This means that if your purpose is to uniquely identify someone, you are processing special category biometric data from the moment you collect the biometric data. It is not the case that you are only processing from the point that you attempt any comparison for identification or verification purposes.

However, it is also important to remember that you are still processing special category biometric data, even if:

- you do not find a match, as you are still creating and comparing biometric data for the purpose of unique identification; or
- you do not need to know who the person is to achieve your overall purpose (i.e., you do not attempt to link the comparison of biometric features to any other known information about that person’s identity like their name). This is because you are still singling someone out with accuracy (uniquely identifying them).

At any stage, if your use of biometric data requires you to uniquely identify someone, then you are processing special category information.

Example

An organisation has an area on their premises that it uses to store highly hazardous chemicals. Previous attempts to limit entrance to this area have failed because employees have shared their PINs with those not authorised to access the area. The organisation adopts a biometric recognition system instead to ensure that only approved staff can access the sensitive area.

It enrolls all authorised staff onto the system. This involves taking a digital image of their thumb. This image is processed to extract biometric features, which are then stored as a biometric template.

Every time a member of staff places their thumb on the door sensor, a probe is created and compared with the biometric reference to confirm whether they are in the database of authorised personnel.

Even if the system does not find a match, the purpose of this processing is to uniquely identify someone from their biometric data.

To implement this biometric recognition system, the organisation needs a lawful basis and an Article 9 condition.

Further reading

- [Special category data](#)

How do we demonstrate our compliance with our data protection obligations?

At a glance

- You **must** comply with data protection law when you use biometric data because it is a type of personal information.
- You **must** take a data protection by design approach when putting in place biometric recognition systems.
- You **must** complete a DPIA before you use a biometric recognition system.
- You **must** assess what impact your use of a biometric recognition system will have on the people whose information it will process (and wider society in some cases).
- You **must** be clear when you are a controller for biometric data, and if you are a joint controller with others.
- You **must** have a written contract in place with all processors.

In detail

- [Who is the controller for our biometric recognition system?](#)
- [Data protection by design and default](#)
- [Do we need to do a DPIA?](#)
- [What risks to rights and freedoms should we consider?](#)

You **must** comply with data protection law when you use biometric data as it is a type of personal information. This means you **must** demonstrate how you comply with the data protection principles.

If you are using a biometric recognition system, you are processing biometric data to uniquely identify someone. This means you are processing special category biometric data.

Processing special category information requires greater care and consideration as it is more sensitive or private than other kinds of personal information. This requirement is reflected in data protection law which prohibits you using special category information without a valid condition for processing it.

Who is the controller for our biometric recognition system?

Your use of a biometric system may involve several different organisations.

You **must** be clear about:

- when you are a controller (with the system provider as your processor); or
- whether at any stage you might be a joint controller with another organisation.

To assess whether you are a controller or a processor, you **must** consider who is determining the purposes and means of the specific processing.

If you and another organisation are joint controllers, both of you **must** ensure that people are able to exercise their rights and understand what you are doing with their information.

Whenever you, as a controller, use a processor to process personal information on your behalf, the processor **must** only process the personal information in line with your written instruction.

You **must** have a contract or other legal act in place. You **must** specify in your contract that the processor should only use the collected biometric data under your instruction.

If the processor uses this biometric data outside your instruction, it is using this information for its own purposes and is therefore considered a controller for these purposes.

If you use a provider that wants to use personal information for their own purposes, then this processing activity won't be something they're doing on your behalf or under your instruction. This means that they are likely to be a controller for this processing.

If you decide to share personal information with your provider, then you must consider and justify your reasons. You must also be clear about what data you intend to share with your provider, and what status this has under data protection law.

As a provider, you **must** also ensure that your own processing complies with the law.

Further reading

- [Controllers and processors](#) provides information about identifying your role when processing personal data.
- [Accountability and governance](#) provides more information about contracts and what to include in them.
- [Contracts and liabilities between controllers and processors](#) provides more detailed information about contracts and liabilities between controllers and processors.
- [How should we understand controller / processor relationships in AI?](#) gives specific information about this relationship in the context of AI.

Data protection by design and default

Your data protection obligations start at the point you decide to use personal information; before you begin to process it.

You **must** adopt a data protection by design and default approach. This means you **must** consider data protection and privacy issues upfront at the design stage and throughout the lifecycle of your system.

Data protection by design and default means that you **must** put in place appropriate technical and organisational measures to implement the data protection principles effectively and safeguard people's rights.

You **must** only use providers who can provide sufficient guarantees of the measures they will use for data protection by design. For example, some biometric recognition systems are designed in ways that reduce the risks associated with personal data breaches of biometric data. See [Risks resulting from personal data](#)

[breaches](#) for more information.

Your choice of biometric recognition system can help to demonstrate your compliance with the data protection by design and default principle. You **should** document your choice and the rationale for it in a DPIA.

When thinking about how your plans to use biometric data will comply with data protection law, you **must** ask yourself the following questions at the initial planning stage:

- Is it necessary for us to use biometric data (ie can we prove that using biometric data is a targeted and proportionate way to meet our needs)?
- What alternatives to biometric data have we considered?
- Could any of these reasonably meet our needs in a less intrusive way?

Further reading

- [Data protection by design and default](#)
- [Data protection principles - guidance and resources](#)
- [What does it mean if you are a controller?](#)
- [Accountability and governance](#)

Do we need to do a DPIA?

You **must** consider whether your processing is likely to result in high risk to people's rights and freedoms. If you identify such risks, you **must** complete a DPIA.

A DPIA will help you to work out the likelihood and impact of those risks, and to demonstrate your compliance with data protection law.

We have published [a list of examples of high-risk processing operations](#).

It is highly likely that you will trigger the requirement to complete a DPIA if you are using biometric recognition systems.

This is because data protection law says that you **must** do a DPIA if you plan to:

- process special category information on a large scale; or
- undertake systematic monitoring of a publicly accessible area on a large scale.

Most uses of biometric recognition systems involve one of these criteria.

And, even if you don't use special category biometric data, you may assess that your proposal to use biometric data is still likely to result in high risk, due to the context and purpose.

If you do not do a DPIA, you **must** still consider the likelihood and potential impact of specific risks that may occur, and the potential for harm that may result.

To do this effectively, you **should** understand how your chosen system works and what its capabilities are.

You may need specialist expertise, including from any providers you're considering.

You **must** also understand the AI supply chain involved in the biometric system you are using or developing, including potential third parties, and consider what risks may be involved.

You **should** also consider whether the system you intend to use involves privacy enhancing technologies (PETs) or whether you can deploy these alongside it. PETs can support you in meeting your data protection obligations. For example, by limiting the amount of personal information you use, or by providing appropriate technical and organisational measures to protect it.

See [Can PETs help us comply with our data protection obligations?](#) for more information.

If you are a provider of biometric systems, as part of your obligations to assist a controller, you **should** explain how your system works. You **must** also ensure that users understand the relevant data protection requirements associated with using your system. For example, how the system observes the security, data minimisation and storage limitation principles.

Further reading

- [Data protection impact assessments](#)
- [Examples of processing likely to result in a high risk](#)
- [Privacy-enhancing technologies \(PETs\)](#)

What risks to rights and freedoms should we consider?

In the context of biometric recognition, there are specific risks to people's rights and freedoms that you **should** consider. This is not intended to be a comprehensive guide to all these risks. See the further reading box at the end of this section for more resources.

Risks resulting from personal data breaches

A personal data breach is a security incident that has affected the confidentiality, integrity or availability of personal information. If you don't address a data breach appropriately, it can lead to a several kinds of harms. This includes those resulting from someone losing control of their personal information, psychological harms, and financial harms.

The severity of harms caused by biometric data breaches may be greater than with other types of personal information due to its sensitive nature. Biometric data represents key features of a person's physical identity that can't easily be changed (eg facial features, eye shape, the sound of their voice). Biometric data breaches can result in an indefinite (depending on which characteristics are processed) loss of control of personal information if biometric data is not appropriately protected. See [Can PETs help us comply with our data protection obligations?](#) for further information about appropriate protection of biometric data.

In addition, the properties of biometric data mean that it can pose specific risks to people.

The first is linkage. Biometric data is a unique identifier. It can, in theory, act as a link across multiple databases on which people's biometric data are stored. This could allow anyone with unauthorised access to someone's biometric data to learn more about them. This risks a serious intrusion of their privacy and

further loss of control of their personal information.

The second is reverse engineering. Biometric data can sometimes be reversed to work out what the original biometric sample (and the person in it) look like. This process could be made easier if the original biometric samples (in addition to templates) are retained and also compromised.

You **should** assess the likelihood and potential impact of reverse engineering and linkage for your use-case. You **should** address the likelihood and impact of these risks through system design and approaches like privacy-enhancing technologies.

You **must** also decide how long it is necessary to keep biometric data for and have clear retention periods. You **should** consider if it is necessary for you to keep original biometric samples once you have generated templates. If you do decide to keep these samples, you **must** protect and store them appropriately.

Biometric data breaches can lead to identity theft that can be very hard to identify as fraudulent. This may result in financial harms, in addition to continued fraudulent access to services or further sensitive information – anything that is protected by someone’s biometric data.

Further reading

- [Personal data breaches: a guide](#)
- [NCSC guidance on how common attack types against biometrics systems](#) 

Risks resulting from biometric false acceptance or rejection

Biometric recognition systems make a statistically informed judgement about whether a new input (biometric probe) and a biometric reference are sufficiently similar to each other. This determines how likely the biometric recognition system is to suggest that the two belong to the same person.

See [How do biometric recognition systems work?](#) for more information about why biometric recognition relies on probability.

The fact that the comparison process relies on probability introduces the potential for the following types of errors:

- **false biometric acceptance** is a false positive (or type I) error specific to biometric recognition systems. The rate at which these occur is sometimes described as the false match rate (FMR). False biometric acceptance occurs when a system incorrectly observes a case as positive when it shouldn’t (ie a match is suggested, but the biometric probe and biometric reference do not belong to the same person).
- **false biometric rejection** is a false negative (or type II) error specific to biometric recognition systems. The rate at which these occur is sometimes described as the false non-match rate (FNMR). False biometric rejection occurs when a system incorrectly observes a case as negative when it should be positive (ie the biometric probe and the biometric reference do belong to the same person, but the system did not recognise them).

Some degree of error is unavoidable in biometric recognition systems. However, both false acceptance and rejection can result in harms to people. For example, biometric recognition systems may be responsible for

controlling access to important resources like bank accounts. False acceptance could lead to unauthorised people gaining access to sensitive information. False rejection could deny people accessing services or opportunities, which could result in financial harms (eg gig workers unable to access apps that allow them to earn money).

For these reasons, you **should** monitor the performance of your biometric system and the rates of false acceptance and rejection. If either of these situations is happening too frequently, people may lose confidence in your system. You may also face problems in complying with the fairness principle, if your use of personal information results in people experiencing unjustified adverse effects.

See [How do we process biometric data fairly?](#) for further information.

Risk of discrimination

False acceptance and rejection are also central to the issue of bias in biometric recognition.

When algorithmic errors are systematic and repeatable, they become biases. Bias is not unique to automated processes; human decision-making is also prone to bias.

Biometric algorithms are trained on biometric data, which is reliant on people's characteristics. Differences exist between people and groups in terms of these characteristics. If training datasets are not representative of the context in which biometric recognition technologies are going to be used, they are less likely to be effective in interpreting the characteristics of certain groups. This is likely to result in certain groups being subjected to higher error rates. This can result in a bias in the biometric algorithm.

Biases may also be present in a biometric algorithm due to human biases impacting how variables are measured, labelled or aggregated.

Certain biometric technologies have previously been shown to be biased according to race, gender and age.

Bias in a biometric recognition system can give rise to discrimination. Discrimination is where people or groups are treated unjustly on the grounds of protected characteristics.

Greater rates of false positive or false negative errors when processing biometric data relating to a specific group could result in discrimination against that group. This is particularly problematic if the system is making fully automated decisions about people. The form that the discrimination takes depends on the context in which biometric recognition is deployed. You **must** ensure your use of biometric data does not result in discrimination.

Biometric recognition systems also have the potential to discriminate in other ways.

Some people may be unable to interact with a biometric recognition system (eg a fingerprint scanner) due to physical disability. This could mean that they are at a significant disadvantage when accessing a specific service, benefit or product, compared to those who can use the biometric solution.

See [How do we process biometric data fairly?](#) for more information.

Risks resulting from systematic monitoring of public spaces

Certain biometric recognition systems are capable of monitoring publicly accessible spaces. These systems often use facial recognition.

The use of this technology for overt surveillance introduces the potential for someone's identity to be determined whenever they enter the field of view of an enabled camera system.


In a practical sense, there is a risk that someone could enter the field of view without being aware that their special category biometric data is being processed.

Without being aware of the processing taking place, people are unable to exercise their right to be informed and their other data protection rights. Even when made aware of the processing, it may be difficult for someone to opt out, if they need to walk through the area under surveillance.

More generally, there are wider concerns that the use of biometric recognition systems in public spaces could result in a 'chilling effect'. This means people are less likely to exercise rights such as freedom of expression or freedom of assembly.

While some of these risks may be mitigated, many are unavoidable aspects of the technologies themselves. Therefore, you **should** consider these seriously and weigh them up against any perceived benefits. The specific circumstances of the deployment will determine whether that interference is lawful and justifiable. Unnecessary deployment of these technologies constitutes another type of harm: unwarranted intrusion.

Further reading

- [Overview of data protection harms and the ICO's taxonomy](#)
- [Opinion: The use of live facial recognition technology in public places](#) 
- [ICO video surveillance guidance](#)
- [ICO AI and data protection risk toolkit](#)

How do we process biometric data lawfully?

At a glance

- Biometric data is personal information. You **must** comply with data protection law when you process it.
- Explicit consent is likely to be the most appropriate condition available to you to process special category biometric data.
- Other conditions may apply, but these depend on the specifics of your proposal and your justification for using special category biometric data.
- If you can't identify a valid condition, you **must not** use special category biometric data.
- If you can identify a valid condition, you **must** still comply with the data protection principles for your processing to be lawful.

In detail

- [How do we determine our lawful basis and special category condition?](#)
- [Can we use consent and explicit consent?](#)
- [What other lawful bases might apply?](#)
- [What other special category conditions might apply?](#)
- [What if we don't have explicit consent and no other condition applies?](#)

How do we determine our lawful basis and special category condition?

If you are using a biometric recognition system, you **must** identify a lawful basis and a separate condition for processing special category biometric data.

Lawful bases

There are six lawful bases for processing personal information. These are:

- consent;
- contract;
- legal obligation;
- vital interests;
- public task; and
- legitimate interests.

No one basis is 'better' or more important than the others. The most appropriate depends on your purpose and relationship with those whose information you are processing.

You **must** determine your lawful basis before you begin processing. Take care to get it right first time – you **should not** swap your lawful basis without good reason.

You **must** always choose the lawful basis that most closely reflects the true nature of your relationship with people and the purpose of the processing.

Special category conditions

You **must not** use special category information without a valid condition for processing it.

There are 10 conditions for processing special category data in total in Article 9 of the UK GDPR. Further requirements for meeting certain Article 9 conditions are provided in the DPA 2018 Schedule 1.

The conditions are deliberately stringent due to potential risks to people's rights associated with processing special category information.

Not all lawful bases have an equivalent condition. There is no requirement for your chosen lawful basis and condition to 'match'.

All of the conditions require you to have a specific purpose for processing. Many of the Article 9 conditions for processing require you to demonstrate that the processing is necessary to achieve its specific purpose.

You **should** carefully consider which conditions for processing special category information are valid in your circumstances. This depends on your purpose and, in some cases, who you are.

Depending on the condition you choose, you may also have to justify why consent is not valid or appropriate in your circumstances.

Whatever condition you consider, you **must** meet all its requirements. If you are unsure, you **should** seek independent legal advice.

[This diagram in our AI guidance](#) can help you to think through what condition might be appropriate for you when processing special category biometric data.

Further reading

- [Lawful basis for processing](#)

Can we use consent and explicit consent?

When using biometric recognition systems, you **must** identify both a lawful basis and a separate condition for processing special category biometric data.

In many cases, explicit consent is likely to be the most appropriate condition. This is because there is no specific special category condition for identification or verification that can apply to the various circumstances in which you may want to process special category biometric data.

There will be circumstances where consent will not be appropriate, for example when using biometric recognition systems for systematic monitoring of public spaces, or where the requirements of consent can not be met.

However, in these circumstances it may be possible to satisfy a different special category condition.

See [What other special category conditions might apply?](#) for further information.

When using explicit consent as a special category condition, you may also use consent as your lawful basis for convenience, although there is no requirement to do so.

Consent

You **must** ensure that the consent is 'specific and informed' to be valid. This means that you **must** give people all the information they need to understand what they are agreeing to. This requires you to clearly set this information apart from other terms and conditions and legal information.

You **must** ensure that consent is also 'freely given'. This means that you **must** give people genuine choice and control about how you use their information. For example:

- You **must** give them the opportunity to refuse or easily withdraw their consent at any time without detriment.
- You **must** also offer a suitable alternative to people who choose not to consent. Otherwise, people do not have a real choice.
- Where there is an imbalance of power between you and the person, you **should** carefully consider whether relying on consent is appropriate.
- You **must** offer a suitable alternative, regardless of whether a power imbalance exists, if you are relying on consent.

This is particularly an issue for public authorities and employers. This is because anyone who depends on your services, or fears adverse consequences if they refuse, may feel they have no choice but to agree. This means people may not freely give their consent. This does not mean that public authorities and employers can never rely on consent. They do need to carefully consider the specific scenario to confirm that they can offer a genuine choice without detriment (eg deciding whether to use a PIN or password; or biometric recognition on a work device).

You **must** also make sure that the consent you receive is specific to your purpose for processing. It may last longer than a single processing operation. There is no set time limit for how long consent is valid for, this depends on the context.

If the processing operation or purposes change, then you **must** seek fresh consent.

If you are systematically monitoring a public space, it is highly unlikely that consent is appropriate. For example, you **must not** assume that someone walking into the field of view of a recognition system has given their consent. This is because consent must be informed, freely given, specific and unambiguous.

Explicit consent

You **must** make sure that the explicit consent is affirmed in a clear statement. A clear statement is not the same as a clear affirmative action. You can't infer explicit consent by what someone does – only by what they say (explicit statement).

This agreement can be given in writing or orally and doesn't require someone to say or write a consent statement in their own words. You **could** provide a statement (an express statement of consent) that they can clearly indicate their agreement to.

You **must** make your statement as clear as possible. This reduces any possibility of later complaints that any consent you received was not 'explicit.'

Example

A gym chain plans to change their swipe card access control system to one that uses facial recognition technology. The gym wants to do this to enhance the customer experience by making it easier and quicker for people to access their facilities.

The new system will use biometric recognition to process their customers' biometric data for the purpose of uniquely identifying them. The gym **must** identify a valid condition for processing this special category information.

Given its intended objective, the gym can only rely on explicit consent for the processing, as no other conditions in Article 9 apply.

To ensure the consent it obtains is valid, the gym provides its customers with all the information that it is required to (eg why they are processing their information and their right to withdraw consent at any time). It asks its customers to indicate their agreement by completing a clear and specific written statement of consent about the processing. It also gives customers the option to enter a unique PIN code if they don't want their biometric data processed.

The gym documents this in its DPIA, describing the non-biometric option as an alternative offered to customers, if they prefer not to have their biometric data used. The DPIA also details their wider compliance with data protection law, and how they have taken a data protection by design approach to the planning and roll-out of the new system.

Further reading

- [When is consent appropriate?](#)
- [When is consent valid?](#)
- [Lawful basis for processing](#)
- [Consent](#)

What other lawful bases might apply?

If you cannot offer people a genuine choice over how you use their biometric data, then you **must** identify a valid lawful basis from the remaining five (contract, compliance with a legal obligation, vital interests, public task, and legitimate interests).

To rely on any lawful basis other than consent, you **must** demonstrate that processing biometric data is

“necessary” to achieve your overall purpose.

Necessity does not mean that processing of personal information has to be absolutely essential in order for a lawful basis to be valid.

However, necessity does mean more than just useful or desirable. You cannot argue that processing is necessary just because you have chosen to operate your business in a certain way.

“Necessary” means that you **must not** rely on these lawful bases to process data unless your processing is a targeted and proportionate way of achieving your purpose.

To help you consider whether your processing would be necessary, you **should**:

- describe the purpose and the benefit you expect to get from the processing of biometric data;
- describe the wider benefits of processing this information to achieve your purpose, including to any third parties, the person themselves, and wider society;
- describe how you will achieve your purpose by processing this personal information; and
- consider the full data lifecycle, and how you have considered data minimisation and storage principles considering your necessary purpose.

You **must** also establish whether your proposed processing is a proportionate way to achieve your purpose when you consider necessity. If you could achieve your purpose in a less intrusive way, or by processing less information, then you cannot argue that your proposal is necessary.

If processing biometric data is not necessary for achieving your purpose and you are unable to rely on consent, you **must not** process biometric data.

To consider proportionality, you **should** do the following:

- Identify the risks and assess the impact the processing has on people.
- Balance the potential impact this processing may have on people against your purpose.
- Be clear about the known or expected false acceptance and rejection rates for the system you intend to use. Assess these against the benefits you are looking to get (ie reduced cost, reduced friction, increased speed). Will people experience any of these benefits?
- Are there any groups who will be unable to use this technology, or for whom the results are likely to be less accurate? If so, what are the real-world implications for people? Will your proposal reinforce or exacerbate existing exclusion or risks of harm felt by specific groups?
- Are there any wider ethical concerns with the processing?
- Describe any potential alternatives to your proposal, and why you have not chosen them.
- Consider any other relevant considerations in choosing this way to achieve your purpose.
- Demonstrate how your processing is sufficiently targeted to meet your objective. For example, a biometric recognition system for access control may be justified for sensitive areas, but not for access to the building.

If you are relying on legitimate interests as your lawful basis, then you are taking on extra responsibility for considering and protecting people’s rights. This means you **must** balance your interests against those of the people whose information you are processing.

If they would not reasonably expect the processing, or if it would cause unjustified harm, then their interests are likely to override yours, and this lawful basis won't apply.

This balancing test builds on the necessity considerations described above. The test considers the kind of information you plan to use, the context you will use it in, and the possible impact your use will have.

One possible safeguard to consider in any balancing test is whether you can offer people a choice to opt-out from their information being used in this way.

Further reading

- [Lawful basis for processing](#) provides more information on how to decide which lawful basis may apply to your proposal and what you need to consider when demonstrating how your chosen basis applies.
- Our [legitimate interest template](#) [↗](#) helps you to decide whether the legitimate interests basis is likely to apply to your processing.

What other special category conditions might apply?

If explicit consent is not an appropriate condition for processing special category information, other Article 9 conditions may still be appropriate.

We outline below two of the most likely conditions which you **could** rely on when processing special category biometric data if you are not relying on explicit consent. Again, whether these conditions are appropriate depends on your purpose and, in some cases, who you are.

Substantial public interest

If your processing is not authorised by law or does not have a basis in UK law and if explicit consent is not an appropriate condition, your only alternative is likely to be substantial public interest.

All of the substantial public interest conditions are set out in paragraphs 6 to 28 of Schedule 1 of the DPA 2018. These give you the basis in UK law to rely on the substantial public interest condition.

The public interest covers a wide range of values and principles about the public good, or what is in the best interests of society. For some of the conditions, the substantial public interest element is built in. For others, you need to be able to demonstrate that your specific processing is "necessary for reasons of substantial public interest", on a case-by-case basis.

If you rely on a condition that requires you to demonstrate that your processing is necessary for reasons of substantial public interest, you **must** make sure the "public interest" is real and of substance. It is not enough for you to make a vague or generic public interest argument to support your purpose. You **must** make a specific argument about the concrete wider benefits of your processing. This is due to the sensitive nature of processing biometric data to uniquely identify people.

The substantial public interest conditions vary in terms of whether they require you to:

- demonstrate that your specific processing is "necessary for reasons of substantial public interest" on a

case-by-case basis (for some conditions, this is assumed);

- justify why explicit consent is not appropriate; and
- have an appropriate policy document in place.

Some of these conditions may be applicable to your reasons for processing special category biometric data.

Further reading

- [What are the substantial public interest conditions?](#) provides further detail about the requirements for each of the conditions.

Prevention or detection of unlawful acts

One of the substantial public interest conditions is the prevention and detection of unlawful acts.

This condition applies if:

- you can demonstrate that it is necessary to use special category biometric data for crime prevention or detection purposes; and
- asking for people's consent means you wouldn't achieve those purposes.

This condition requires you to demonstrate that your processing is necessary for reasons of substantial public interest. This means you **must** be able to show that using special category biometric data is "necessary" both for the prevention and detection of crime and for reasons of substantial public interest. You **must** include an explanation of how the processing is likely to be effective in preventing or detecting the illegal act.

To satisfy this condition, you **must** demonstrate you are using biometric data in a targeted and proportionate way to deliver the specific purposes set out in the condition. Also, that you cannot achieve them in a less intrusive way.

You **must** also have an [appropriate policy document](#) in place at the time your processing starts.

If you are processing special category biometric data to detect or prevent a crime, it is likely that you are also processing criminal offence information.

Criminal offence information refers to 'personal data relating to criminal convictions and offences or related security measures'. This covers information about offenders or suspected offenders in the context of criminal activity, allegations, investigations and proceedings. It includes suspicions of criminal activity and evidence relating to them.

To process criminal offence information, you **must** identify a lawful basis under Article 6. You **must** also identify a condition in Schedule 1 of the DPA, or be processing under the control of official authority.

You **must** be clear about why you need criminal offence information. You can then identify the most relevant condition.

In most cases, if you have already identified your Article 9 condition for processing special category data, this may also justify the processing of criminal offence data.

Further reading

- [Criminal offence data](#)
- [What does 'under the control of official authority' mean?](#)

Research

This condition for processing special category information applies if you intend to use special category biometric data for one of the research purposes.

You **must** be able to show that using special category biometric data is “necessary” for the research purpose. This means that your use of special category biometric data is a reasonable and proportionate way to achieve your purpose.

To rely on this condition, you **must** also comply with further safeguards, including demonstrating that your use of special category biometric data is:

- not likely to cause someone substantial damage or substantial distress; and
- in the public interest.

Example

A company uses a dataset of biometric samples to assess and address the bias in its facial recognition system. The dataset is diverse in terms of ages and genders.

It uses some of the dataset to train the model so that it learns from a diverse range of inputs.

It uses the remaining information to test the performance of the model and confirm comparable statistical accuracy for all ages and genders. This will help in avoiding potential discriminatory effects when the company launches the system.

Further reading

- [What are the substantial public interest conditions?](#)
- [The research provisions](#)

What if we don't have explicit consent and no other condition applies?

If you cannot gain explicit consent, and no other condition is appropriate, then you are infringing data protection law if you process special category biometric data. This is because your processing is unlawful.

The first data protection principle requires any processing of personal information to be fair, lawful and transparent.

If you cannot identify a valid condition, then you won't be able to comply with this principle. You **must** therefore consider other options to achieve your purpose and **must not** use a biometric recognition system.

How do we process biometric data fairly?

At a glance

- Biometric recognition systems with poor statistical accuracy can result in challenges to fair processing of biometric data.
- You **must** ensure any system you use is sufficiently accurate for your purposes.
- You **should** establish a threshold that is appropriate for your circumstances.
- You **should** test for and mitigate biases in any system you use.
- If you are a provider, you **should** ensure users are aware of any biases.
- You **should** consider the real-world risks of false acceptance and rejection rates and put in safeguards to mitigate these risks accordingly.
- You **must** ensure that biases do not result in discrimination.

In detail

- [How does the statistical accuracy of biometric algorithms relate to the fairness principle?](#)
- [How do we deal with the risk of discrimination?](#)

Identifying a lawful basis for processing does not mean your processing is lawful by default. You **must** also ensure that your use of biometric systems is lawful more generally.

For processing to be fair, you **must** use information in ways that people would reasonably expect and that do not have unjustified adverse effects on them. Also, you **must not** mislead anyone when obtaining their biometric data.

How does the statistical accuracy of biometric algorithms relate to the fairness principle?

The effectiveness of a biometric recognition system (its ability to identify people correctly) is a measure of its statistical accuracy. If systems routinely fail to identify people correctly, these errors could have a real-world impact.

If you don't address this risk of inaccuracy, you could contravene the fairness principle and other equalities legislation. This may leave you exposed to further legal claims, as well as regulatory action.

How do we deal with risks resulting from errors?

Because the comparison process relies on probability, you **should not** interpret the decision of a biometric recognition system as an objective fact about someone's identity. Instead, you **should** see it as an indicator of the confidence you can have in asserting someone's identity.

Example

Traditional verification methods (eg a password) make a simple comparison between an input value (what you type) and a stored value (the password). If the input exactly matches the stored value, then access is granted.

This is a binary outcome – either the input value will match the stored password, or it won't.

Biometric recognition systems work in a different way.

Although their objective is the same, there are a range of factors which mean that no two captures of biometric data can be truly identical in the same way as a password.

Differences in environmental conditions (such as the amount of light or glare) may mean that an input value (eg an image of a face) doesn't precisely match the stored value (ie an image of their face when they first enrolled on the system).

Other issues can complicate the matching process, such as the passage of time between the original enrolment and the later re-presentation.

You can never entirely eliminate errors from biometric recognition systems. You **must** consider the potential impact such errors could have on people in real life. Your context also determines your tolerance for these errors, and what trade-offs in performance you can justify.

In a biometric recognition system, you can reflect these trade-offs in your choice for the system's threshold – the point at which the result of a comparison is considered statistically significant.

A more sensitive system has a lower threshold: this increases your likelihood of finding a 'match,' but it is more likely that the system will return false positives.

A more precise system has a higher threshold: this decreases your likelihood of finding false matches, but it increases the chance that you miss a true positive match.

You **should** use a threshold that is appropriate for your circumstances. This depends on:

- what impact the decision could have on someone;
- what your use case and context is;
- what safeguards you have in place; and
- whether the system is making solely automated decisions.

As a result, you **should not** assume that a biometric recognition system always offers the best performance in your scenario when left on any default threshold settings. You **should** understand whether you can configure your system locally (eg to optimise performance and reduce errors to an acceptable level).

If you are a provider, you **should** assist in setting up your systems to help controllers with their obligations.

In the case of biometric identification, the number of references against which you are comparing can be a factor in the rate of false positive errors observed. To help with accuracy, you **should** therefore keep the size of any reference database as small as possible. This also helps you to comply with the data minimisation principle.

While some error is unavoidable, you **should** use well-developed systems that minimise the number of errors that could occur.

For example, the National Institute of Standards and Technology (NIST), an agency of the United States Department of Commerce, tests facial recognition technologies for their error rates. This is not an accreditation process, and all tests take place under laboratory conditions. However, you may find this is a useful resource to:

- demonstrate that different facial recognition technologies have different error rates; and
- help you work out whether the service you are considering has a significant observed high error rate compared with other commercially available products.

You **should** assure yourself of the statistical accuracy of any solution. This includes where any published information on the performance of a solution came from (ie lab testing versus 'real-life' scenarios similar to yours).

Before you deploy a biometric recognition system, you **must** understand the potential implications of these sorts of errors. You **should** consider this in terms of the possible impact on both the people who will rely on the system, as well as your organisation as a whole.


You **should** have mechanisms and processes in place to:

- diagnose any quality issues or errors;
- record how these errors are resolved;
- check that your systems are working as intended; and
- highlight any inaccuracies or bias.

Maintaining a biometric recognition system is key to making sure that you process special category biometric data fairly.

You **should** consider what appropriate safeguards are required to ensure people do not experience real detriment or harm because of an error. This might require introducing a way for people to challenge the outcome, if they feel that it is inaccurate.

Further reading

- [Guidance on AI and data protection](#) – including the [section on statistical accuracy](#) and the trade-offs involved in considering terms like the precision and recall of AI systems
- [NIST - Face recognition technology evaluation \(FRTE\)](#) 

How do we deal with the risk of discrimination?

You **should** test all biometric recognition systems for bias. If you detect bias, you **should** mitigate it.

If you are a provider, you **should** inform potential users about:

- bias that exists;

- the implications of that bias; and
- how they might mitigate it.

For example, it might be possible to improve bias by training the system further on a dataset that is more representative of the context you will deploy it in.

As part of their testing process, NIST analyse the effect of demographics on the effectiveness of biometric recognition systems. As a controller, this may give you an idea of the levels of observed bias rates and how and why they can vary.

Bias in a biometric recognition system can give rise to discrimination. See [Risk of discrimination](#) for more information on the relationship between errors, bias and discrimination.

Discrimination is where people or groups are treated unjustly on the grounds of protected characteristics.

You **must** ensure that bias in biometric recognition systems does not lead to discrimination.

Biometric systems can also result in discrimination in ways not related to the statistical accuracy of the system.

For example, if a disabled person was unable to use a biometric recognition system controlling access to a room, and there was no other option for them to gain access, this would make your use of biometric recognition technology unlawful. This is both because you would violate the fairness principle, and because disability is a protected characteristic.

If you intend to use a biometric recognition system, as part of complying with the fairness principle, you **must** assess whether it is likely to have a discriminatory impact on people.

When considering a biometric recognition technology, you **should** also consider its potential for exclusion based on non-protected characteristics. For example, if your use case requires people to have a smartphone, you may disproportionately be excluding and impacting certain groups. If this happens, you are likely to find it challenging to demonstrate how your system complies with the fairness principle.

Example

Fingerprint recognition is less accurate for adults over 70 and children under 12. This is because older adults' fingerprints are less distinct and young children's fingerprints change rapidly because they are still developing.

As a result, a technology that may superficially be considered to be fair can have unfair impacts, as it will systematically perform worse for older adults and young children.

Further reading

- [AI guidance: Fairness in the AI lifecycle](#)

How does the accuracy principle apply to biometric data?

At a glance

- Biometric recognition systems do not need to be 100% statistically accurate to comply with the accuracy principle.
- You **should** ensure that your records indicate that decisions made by biometric recognition systems are statistically informed guesses rather than facts.
- The accuracy of biometric data can decrease over time because of the aging process and other life events.
- You **must** have appropriate processes in place to check the accuracy of the personal information you collect and create.

In detail

- [Does the accuracy principle apply to our biometric recognition system?](#)
- [When do we need to collect new biometric samples?](#)

Does the accuracy principle apply to our biometric recognition system?

Yes. The accuracy principle applies to the processing of all personal information. This includes any information associated with biometric data in a database (eg name or date of birth).

But there can be some confusion between the accuracy principle and the concept of statistical accuracy that features in automated systems like biometric recognition systems.

The accuracy principle is about ensuring the personal information you process is not inaccurate or misleading as to any matter of fact. As noted earlier, statistical accuracy is about how well your system performs in given conditions. For example, the rate of false biometric acceptance or rejection.

Statistical accuracy is more relevant for your considerations about the fairness principle (ie whether the decisions your system makes result in fair outcomes for people).

The outcomes of your biometric recognition system are statistically-informed judgements about someone's identity. They are based on estimates about the level of similarity between a biometric probe and a biometric reference, usually in relation to a threshold.

See [How do biometric recognition systems work?](#) for more information.

Biometric recognition systems involve the processing of personal information, and therefore the accuracy principle applies. But they don't need to be 100% statistically accurate to comply with it.

To avoid these outcomes being misinterpreted as factual, you **should** ensure that your records indicate they are statistically-informed guesses rather than facts.

At the same time, you **must** ensure that your system is sufficiently statistically accurate for your purposes. This doesn't mean every single outcome has to be correct, but you **must** factor in the possibility of errors happening and the impact they may have both on your decision-making and the people it applies to. If you don't do this, your processing may not comply with the fairness principle.

See [How do we process biometric data fairly?](#) for more information about how statistical accuracy impacts your compliance with data protection law.

Example

A company uses a fingerprint recognition system to control access to a restricted area of its premises. An employee who is permitted to access the restricted area attempts to use the system to gain access, but it falsely rejects them.

While this false rejection is a matter of fact - the system failed to allow the employee into the area - it doesn't mean that the processing of personal information is inaccurate in the context of the accuracy principle.

This is because the system's comparison process produced a similarity score rather than a statement of fact. The company clearly labels its records to note this.

When do we need to collect new biometric samples?

Our physical, physiological and behavioural characteristics change as we age. However, the biometric reference that describes these characteristics is fixed in time at the point at which this information was captured.

This means that the accuracy of biometric references can decrease over time. Therefore, the same biometric reference can result in a greater false non-match rate as someone ages.

Different characteristics change at different rates, and the rate of change of these characteristics is likely to be greatest in younger people.

If you discover that biometric data is no longer accurate, you **must** take reasonable steps to correct or erase it as soon as possible.

You **should** have a re-enrolment process that appropriately addresses these issues for your circumstances.

If a biometric reference is no longer necessary for your purposes, you **must** delete it.

How do we ensure our processing of biometric data is transparent?

At a glance

- To comply with the transparency principle, there is certain information that you **must** share.
- There is not a fixed way in which you **must** share this information.
- In deciding how you are going to share this information, you **should** consider several factors.

In detail

- [What information do we have to share to comply with the transparency principle?](#)
- [How should we provide transparency information?](#)

People have a right to know how and why you're processing their information. If you are using any biometric system that processes personal information, you **must** explain how in a way which is clear, concise, and easy to access.

What information do we have to share to comply with the transparency principle?

You **must** provide the following information:

- your retention periods for that personal information, and who you will share it with;
- all relevant contact information (eg the name and contact details of your organisation) (and your representative, if applicable) and the DPO's contact details;
- the purposes of the processing and the lawful bases (and, if applicable, the legitimate interests for the processing);
- details of all personal information that you share with other organisations and, if applicable, details of transfers to any third countries or international organisations;
- retention periods for the personal information, or if that is not possible, the criteria used to determine the period;
- details about people's rights including, if applicable, the right to withdraw consent and the right to make a complaint; and
- whether people are under a statutory or contractual obligation to provide the personal information (if applicable, and if you collect the personal information from the person it relates to).

You **must** provide privacy information to people at the time you collect their personal information from them, or ahead of time. This is because people have the right to be informed about the collection and use of their personal information.

Getting the right to be informed correct can help you to comply with other aspects of the GDPR and build trust with people. But getting it wrong can leave you open to fines and lead to reputational damage.

Further reading

- [Right to be informed](#)

How should we provide transparency information?

The best way to provide this information may differ depending on:

- your relationship to the people whose information you are processing;
- what your processing involves; and
- what your use case is.

You **should** consider how people will interact with the technology or the wider context of how you are using it when deciding how to provide this information. This will help you work out the most effective way of informing people.

You **should** also consider the potential impacts of any decisions biometric recognition systems will make and what people should know about these. It might be appropriate to provide information in different formats or levels of detail for different people, depending on their level of pre-existing knowledge.

For example, you **could** make information available in the following ways:

- using leaflets or digital techniques (eg QR codes and other local media), in advance where possible;
- members of staff on hand to discuss the processing;
- visual or audio signals; and
- making information available online and through social media, and otherwise using digital spaces that visitors are likely to use in advance of visiting the premises.

You **should** also consider who someone should ask if they have questions about the processing.

Once you have put in place mechanisms for providing transparency information, you **could** do user testing and surveys to assess whether the information is sufficiently clear and accessible.

Example

A cruise ship company decides to offer biometric facial recognition as an option for guests. The guests can use the biometric recognition system when reboarding the ship following an excursion.

Before the cruise, the company sends information to guests about the process, together with a video explaining how the system works. At the start of the cruise, a member of staff is on hand at reception to explain the process and answer guests' questions and concerns.

An alternative, non-biometric solution, is available for any guests who do not wish to use this option.

Further reading

- What privacy information should we provide?

How do we consider rights requests for biometric data?

At a glance

- Biometric data is a form of personal information. This means that you **must** carefully consider people's rights.
- Your choice of lawful basis will determine whether some rights apply in your circumstances.
- You **must** understand how people's rights apply to your processing of biometric data and have processes in place to recognise and respond to requests.
- There are strict limits to using biometric data to make solely automated decisions about people.

In detail

- [What is the right of access?](#)
- [What is the right to rectification?](#)
- [What is the right to erasure?](#)
- [What is the right to data portability?](#)
- [What is the right to object?](#)
- [Can we use a biometric recognition system to make automated decisions about someone?](#)

This section provides further context to help you decide what rights may be in scope, and what relevant factors you **should** consider. This section does not describe the right to be informed in further detail, as this is addressed in [How do we ensure our processing of biometric data is transparent?](#).

What is the right of access?

The right of access is a cornerstone right of data protection as it allows people to see what information you collect and share about them.

When you respond to a subject access request, you **must** confirm to the person whether you are processing their personal information and why (unless an exemption applies).

You **must** also provide them with other information about the processing. For example:

- your purpose for processing;
- the categories of personal information you are processing; and
- the recipients or categories of recipients you intend to disclose, or have already disclosed, the personal information to.

If your processing involves any automated decision-making, then you **must** make this clear and provide information about the logic involved and the envisaged consequences for people of the decision.

The right of access also entitles the requester to a copy of their personal information. When you use biometric recognition systems, this includes personal information, such as a biometric sample, as well as any biometric data, such as a biometric template.

You should be clear about:

- what is in scope of any subject access request; and
- whether any copy you provide includes the personal information of people other than the requester.

For example, an audio recording may include personal information of several people. In contrast, a biometric sample used to create a template just includes the voice of the person being enrolled onto the system.

At the same time, there may be some practical issues in providing someone with a copy of their biometric data.

This is because biometric data is likely to consist of complex mathematical outputs in a specific (and often proprietary) machine-readable format. Biometric data in this form is not 'readable' by people, and by design may not even be readable by other biometric recognition systems.

The format of biometric data may also mean it is not possible to provide the information in other forms (ie hard copy, or even a commonly used electronic form).

Responding to subject access requests does not require you to translate or decipher information in its entirety. However, you **must** give enough information to aid the requester's understanding (ie if records include specific technical language or acronyms).

Therefore, where there are practical considerations which prevent you providing a copy of someone's biometric data, you **should** provide additional explanations to help them understand:

- your justification for being unable to provide a copy of personal information in scope, and a summary of the specific practical issues;
- what the information consists of; and
- how you hold it.

If you are using a biometric recognition system to identify persons of interest (ie through comparison against a watchlist), then you **must** be transparent about this in your response, unless you can justify using one of the exemptions from the right of access.

Further reading

- [Right of access](#)
- [A guide to the data protection exemptions](#)
- [Explaining decisions made with AI](#)

What is the right to rectification?

The right to rectification provides the ability for anyone to rectify or complete any incorrect or incomplete

personal information about them.

Rectification requests may result from a subject access request.

If you receive a rectification request, you **must** satisfy yourself about whether the information you hold is accurate and consider what steps you have taken to assure yourself of this.

For biometric recognition systems, you **should** remember that a match suggested by a system is not subject to a rectification request. This is because a suggested match is a statistically-informed estimate and not accurate as to a matter of fact.

However, a mislabelled record or a further decision or opinion you take based on a suggested match might be subject to this right.

A biometric reference may also be subject to this right if it no longer accurately represents a person's biometric characteristics.

If you are using watchlists, then the right to rectification applies if the information you hold is inaccurate or incomplete. You **must** also inform any recipients of the watchlist about any rectifications you've made, unless this is impossible or involves disproportionate effort.

You **should** also say who those recipients of the watchlist are, if you're asked.

Further reading

- [Right to rectification](#)

What is the right to erasure?

You may receive requests to erase the biometric data you hold. The right to erasure is not absolute, so you **must** understand how the right applies in your circumstances.

If you are relying on consent to process biometric data, then people can withdraw consent at any time. If someone withdraws consent, then they can also make an erasure request. You **must** comply with this request without delay unless you have another purpose under another lawful basis to continue to hold the biometric data.

If you intend to retain biometric data for another purpose, then you **must** make this clear at the time you originally seek consent and include details of this other purpose in your transparency information.

If you are not relying on consent, then there are other reasons that people can make an erasure request. In summary, these will depend on your circumstances for processing biometric data, and whether your continued use of that information can be considered necessary.

One of these grounds depends on exercising another qualified right, the right to object, which is discussed later in this section.

If you've received an erasure request and no exemption applies, then you **must** take steps to delete the biometric data you hold – including in any backup systems.

Further reading

- [Right to erasure](#)

What is the right to data portability?

The right to data portability applies to any personal information someone provides to you where:

- your lawful basis for processing is either consent or contract; and
- you are carrying out the processing by automated means.

People have the right to receive this personal information. They can also ask you to transmit it to another organisation, and you **must** do this if it is technically feasible.

The right doesn't apply to personal information that you create based on what someone's provided to you.

This means the right doesn't apply to biometric templates. This is because a template is something you've derived from a person's characteristics. It's still personal information, and therefore other rights apply. For example, the right of access.

You **must** tell people what rights they have when you collect their personal information. These rights differ depending on your lawful basis for processing, so you **must** make sure that what you tell people reflects this.

If someone makes a request for their biometric data under the right to data portability, you **must** explain why you're not going to act in response to the request (eg because the right doesn't apply).

You **must** also clarify with the requester that biometric data falls under the right of access. See [What is the right of access?](#) for further information.

Further reading

- [Right to data portability](#)
- [Right to be informed](#)

What is the right to object?

Provided you are not using personal information for direct marketing, the right to object only applies if you are:

- processing biometric data under the public task or legitimate interest lawful bases; and
- not able to override the request by demonstrating a compelling reason that requires you to continue processing the biometric data.

If someone exercises their right to object and you don't have a strong reason to refuse their request, then you **must** also consider whether you are required to erase their personal information.

Further reading

- [Right to object](#)

Can we use a biometric recognition system to make automated decisions about someone?

People have the right not to be subject to decisions based solely on automated processing that produce legal or similarly significant effects.

Many uses of biometric recognition systems involve making solely automated-decisions about people. Depending on the specifics of your deployment, these decisions may have legal or similarly significant effects on those people (eg denial of a service).

Data protection law restricts the circumstances in which you can make solely automated-decisions with these kinds of effects. These include specifying the:

- conditions you can rely on; and
- safeguards you **must** have in place (eg the ability for a human review of any decision where required).

To determine whether your biometric recognition system makes these kinds of decisions, you **should** ask the following questions:

- What decisions do you intend the system to make?
- Who (or what) determines these decisions?
- Is the decision solely-automated, or is there any meaningful human involvement in making the decision?
- What are the potential impacts of the decisions on someone? Do the decisions affect people's legal rights or have a similarly significant effect on their circumstances or choices?

If you are using biometric recognition systems to make these decisions about people, you are making these decisions using special category information. This means that you **must not** carry out this processing unless:

- you have the person's explicit consent; or
- the processing is necessary for reasons of substantial public interest.

Further reading

- [What is the impact of Article 22 of the UK GDPR on fairness?](#)
- [Automated decision-making and profiling](#)

How do we keep biometric data secure?

At a glance

- You **must** apply appropriate security measures when you use biometric data.
- You **should** consider how privacy enhancing technologies (PETs) can help you meet your data protection requirements.
- You **must** only store information for as long as you need it.
- Your choice of provider can help you to demonstrate compliance with your data protection obligations.

In detail

- [What are appropriate security measures?](#)
- [Can PETs help us comply with our data protection obligations?](#)
- [How can we comply with the data minimisation and storage limitation principles?](#)
- [Should we retain biometric samples?](#)

What are appropriate security measures for biometric data?

The security principle is about ensuring the personal information you hold isn't accidentally or deliberately compromised.

You **must** process biometric data in a way that ensures appropriate security and protection against unauthorised or unlawful processing (amongst other things).

This means you **must** apply appropriate security measures when you use biometric data. This includes both technical and organisational measures.

You **should** determine what these measures are by carrying out a risk analysis that considers:

- the circumstances of your processing and the likely security threats you may face;
- the damage or distress that may be caused if the biometric data is compromised; and
- what forms of attack your system might be vulnerable to.

"Appropriate" is a higher bar here than for personal information more generally. This is due to the sensitive nature of biometric data and the risks associated with it falling into the wrong hands.

You **must** also conduct regular testing and reviews of your security measures to ensure they remain effective. Other examples of organisational security measures include having information security policies in place and ensuring key personnel in your organisation co-ordinate with each other.

You **must** also appropriately encrypt any biometric data that you use. This is an example of a technical security measure. There are many dimensions to security outside of cybersecurity that you **must** consider. For example, physical security measures, such as how you control access to your premises.

If you are using a processor, you **must** choose a processor that offers sufficient guarantees to implement appropriate technical and organisational measures.

If you are a processor, you **must** assess system vulnerabilities systematically and regularly and act on any findings in a timely way to ensure biometric data remains secure. You **should** also be able to respond to any identified threats quickly, to minimise the impact of any attacks.

Further reading

- [A guide to data security](#)
- The NCSC has produced some guidance on [biometric recognition systems and security](#) ↗
- [ISO/IEC 24745:2022](#) ↗ includes a list of threats and countermeasures

Can PETs help us comply with our data protection obligations?

Yes. You **should** consider using biometric recognition systems that employ PETs, as these can help you demonstrate compliance with the security principle.

But you **must** still:

- have appropriate organisational measures in place to keep information secure, such as regular testing; and
- review your security measures to ensure they remain effective.

PETs can help you meet other data protection obligations too, like the requirements around data protection by design and by default. For example, PETs that limit the amount of personal information you process can help you demonstrate compliance with the data minimisation principle.

Techniques that seek to minimise privacy risks in biometric recognition systems are sometimes called 'biometric privacy enhancing technologies', or B-PETs.

Biometric data may be impossible to anonymise completely. By definition, it must be capable of identifying someone. The general approach of all B-PETs is to strike a balance between:

- minimising the risks associated with unauthorised access to biometric data (ie the risk that an unauthorised party can identify someone); and
- ensuring the biometric data is as accurate as possible (so an authorised party can reliably identify an individual).

Not all B-PETs can be used in all use cases for biometric recognition. You **should** carefully consider which PETs are appropriate to use in your context.

Further reading

- [Reference table](#) – Our guidance on privacy-enhancing technologies (PETs) includes a non-exhaustive list of PETs and their potential use-case applications.

Biometric template protection

You **should** choose a provider that protects biometric templates appropriately.

The ISO/IEC 24745 standard on biometric information protection outlines several important characteristics for biometric templates. These include:

- **Irreversibility:** biometric templates are difficult to reverse engineer to gain information about the someone's appearance.
- **Unlinkability:** biometric templates cannot easily be attributed to the person they relate to, meaning they cannot easily be used to link between different databases based on a person's biometric data, or used in different biometric recognition systems.
- **Revocability and renewability:** biometric templates can be revoked or cancelled and replaced with a new template without needing to take a new biometric sample.

These characteristics help to mitigate the risks associated with unauthorised access to biometric data.

If your use case involves storing biometric templates, you **could** ask your provider whether their system produces templates that have these characteristics.

The first two characteristics mean that if someone gains unauthorised access to a person's biometric data, it is difficult to fraudulently use that information on its own.

Revocability and renewability mean that a single biometric sample can be used to create multiple different templates. This means that if a single biometric template is compromised, it can be revoked, and a new template created which differs from the compromised version. This approach can also mean that all templates in a database can be renewed using the revised method without requiring everyone to re-enrol.

However, the effectiveness of this approach relies on the swift detection of any data breach in the first place, underscoring the importance of having appropriate security measures in place.

There are several techniques that can help to achieve these characteristics, some of which can be used in combination with others to provide greater protection. These include:

- biometric cryptosystems;
- cancellable biometrics; and
- homomorphic encryption

All these systems work on the principle that the comparison process in biometric recognition should not use unprotected biometric data directly, as this could result in sensitive personal information being exposed. Instead, comparison happens based on a transformed or encrypted version of the biometric data.

Like any form of encryption, the effectiveness of this approach relies on the management of the keys used.

You **should** ensure that it would be difficult for an unauthorised party to undo this protection in the event of a data breach through effective key management approaches.

Further reading

- [ISO/IEC 24745:2022 - Biometric data protection](#) for further information on different methods of protecting biometric templates.

Personal information in biometric samples and templates

Biometric recognition systems do not only use biometric data. They can also capture other types of personal information, which could be used to infer several things about someone. These could include characteristics such as age, gender and hair colour.

Some B-PETs look to minimise or protect the amount of personal information in biometric samples and templates. In the case of samples, these techniques look to reduce a person's ability to recognise someone, for example, by distorting specific features of a person's face. In the case of templates, these techniques look to reduce the machine-readability of this information.

These approaches also align with the data minimisation and storage limitation principles. They can also help with the purpose limitation principle, as they can reduce the utility of this information for other purposes (such as profiling) by removing personal information about certain characteristics.

On-device verification

On-device verification is a technique that can reduce the amount of biometric data created and shared compared with other systems that verify people's identities remotely.

By configuring your systems and devices to perform on-device verification, your users only need to create and store a single biometric template. This happens entirely on the device, and no biometric data leaves it – only a token or proof of the verification. This means they can access a range of applications and services easily and securely.

Whether on-device verification is appropriate depends on your circumstances. You **should** consider the different benefits and risks involved.

For example, it reduces the potential impact of a large data breach when compared to on-server storage of biometric data. This is because the biometric data isn't all stored in one place, which can reduce the risks of harm that may arise from a data breach.

However, it may also mean that you are less likely to have control over how the biometric data is processed, because you do not have access to the biometric recognition system itself. In turn, this may make it more difficult for you to identify security threats and ensure your processing complies with the security principle.

Further reading

- [NCSC guidance on biometric recognition systems and security](#)

How can we comply with the data minimisation and storage limitation principles?

You **must** comply with both the data minimisation and storage limitation principles when processing biometric data. The less information you collect, store and retain, the less information you need to protect.

Data minimisation means you **must** limit the amount of biometric and other personal information that you process to the minimum that is adequate, relevant and necessary for your purpose.

B-PETs that minimise the amount of personal information you collect help you to comply with the data minimisation principle.

The storage limitation principle means you **must** only keep biometric (and other personal) information for as long as it is necessary for your purposes.

You **must** consider storage limitation throughout the lifecycle of personal information as it passes through a biometric recognition system. You **must** have processes in place to regularly review your database of biometric references to ensure you delete any data that you no longer need.

You **must** have clear retention periods which means you only keep this information in an identifiable form for as long as is necessary. This demonstrates compliance with the storage limitation principle and will have benefits to your organisation.

An example of a way to comply with the storage minimisation principle is to use biometric recognition systems that transiently process biometric probes and immediately delete them as soon as any comparison fails to meet the acceptance threshold.

Should we retain biometric samples?

Renewability of biometric templates relies on retaining the original biometric sample. This allows you to create new templates without creating a new sample.

To comply with the data minimisation and storage limitation principles, you **must** consider whether you need to keep the original biometric sample. If you do, you **should** ensure that you appropriately restrict access to the retained sample and template and deploy further measures like PETs.

Further reading

- [Data minimisation](#)
- [Storage limitation](#)