


Guidance on highly effective age assurance

For Part 3 services

Published 24 April 2025



Contents

Section

1. Overview.....	4
2. Introduction.....	5
3. Age assurance methods and processes	7
4. Criteria to ensure an age assurance process is highly effective	10
5. Privacy and data protection	22
6. Glossary	25

1. Overview

What this guidance covers

This is our guidance to assist providers of regulated user-to-user and search services (“Part 3 services”) in implementing highly effective age assurance for the purpose of fulfilling their regulatory obligations under the Online Safety Act 2023 (“the Act”).

This guidance is applicable for the purposes of:

- stage 1 of the children’s access assessment, as explained in the Children’s Access Assessment Guidance; and
- understanding how highly effective age assurance should be implemented where applicable to recommended measures set out in Ofcom’s Protection of Children Code of Practice for user-to-user services.

This guidance sets out additional detail to assist service providers in i) understanding whether an age assurance process is highly effective and ii) complying with the recommended measures.

Our approach to highly effective age assurance aligns, where appropriate, with the approach taken in our Guidance for service providers publishing pornographic content under Part 5 of the Act. This is to ensure that service providers in scope of Part 5 and / or Part 3 of the Act have a clear and consistent understanding of how to implement highly effective age assurance to prevent children from encountering harmful content.

This guidance was originally published on 16 January 2025. This updated version was published on 24 April 2025. We have replaced references to the draft version of the Protection of Children Code for user-to-user services with references to the final version.

2. Introduction

Background to the guidance

Children's Access Assessment

- 2.1 All providers of Part 3 services are required to carry out children's access assessments to determine whether a service, or part of a service, is likely to be accessed by children.
- 2.2 The Act says that service providers may only conclude that it is not possible for children to access a service if that service uses a form of age assurance with the result that children are not normally able to access that service or part of it.¹
- 2.3 We consider that, in order to secure the result that children are not normally able to access their service (or a part of it), service providers should deploy highly effective age assurance and implement effective **access controls** to prevent users from accessing the service (or relevant part of it) unless they have been identified as adults.²
- 2.4 As stated in the Children's Access Assessment Guidance, service providers should consult this guidance to understand what constitutes highly effective age assurance and / or to carry out an in-depth assessment of whether a particular form of age assurance is highly effective for the purpose of stage 1 of the children's access assessment.³

Protection of Children Codes

- 2.5 The Protection of Children Code of Practice for user-to-user services ("the Code"), includes recommended measures on the implementation of highly effective age assurance in certain circumstances.⁴ The Code sets out the definition of highly effective age assurance for these recommended measures, and lists the steps that service providers should take to fulfil each of the criteria.⁵
- 2.6 The Code also includes other recommended measures which may be relevant to the way that service providers implement and operate a highly effective age assurance process on their service – for example, measures relating to the clarity and accessibility of terms of service, and reporting and complaints.⁶
- 2.7 Service providers are required to keep records of (1) steps that they have taken in accordance with the Code, or (2) any alternative steps they have taken to comply with their

¹ Section 35(2) of the Act.

² We use the term "access controls" to describe a technical mechanism(s) which prevents users who have not been age assured, or having been age assured, did not meet the requirements of the age assurance process, from accessing a service (or part of it) or certain content.

³ Children's Access Assessments Guidance, April 2025.

⁴ Measures PCU B1-PCU B7 of the [Protection of Children Code for user-to-user services](#).

⁵ Paragraphs 5.10-5.15 of the Protection of Children Code for user-to-user services.

⁶ See the 'Index of Recommended Measures' in the Protection of Children Code for user-to-user services for a full list of the recommended measures and which services they apply to.

duties.⁷ Service providers should consult our Record Keeping and Review Guidance for this purpose.⁸

- 2.8 This guidance will help service providers in adopting recommended measures that relate to the implementation of highly effective age assurance, by providing additional technical detail and examples on how to meet the standard.

Navigating the guidance

- 2.9 We set out below an overview of the remaining sections of the guidance.

Section 3: Age assurance methods

- 2.10 Section 3 sets out a non-exhaustive list of age assurance methods that we consider are capable of being highly effective at correctly determining whether or not a user is a child, and those that we consider are not capable of being highly effective.

Section 4: Criteria to ensure an age assurance process is highly effective

- 2.11 Section 4 sets out the four criteria of technical accuracy, robustness, reliability, and fairness, that the age assurance process should fulfil to ensure that it is highly effective at correctly determining whether or not a user is a child.
- 2.12 We define each of the criteria, why they are important, and outline steps that service providers can take to have regard to them.
- 2.13 Section 4 sets out detail on the additional principles of accessibility and interoperability that service providers should consider alongside the criteria.

Section 5: Privacy

- 2.14 Section 5 provides some guidance on how service providers can have regard to protecting users' privacy when implementing age assurance. It includes relevant information about the data protection regime and directs service providers to ICO guidance. It also sets out examples of how service providers can have regard to privacy under the Act.

⁷ Section 23(3) and (4) of the Act.

⁸ [Record Keeping and Review Guidance](#)

3. Age assurance methods and processes

- 3.1 In this section, we set out a non-exhaustive list of the kinds of age assurance that we consider are capable of being highly effective at correctly determining whether or not a user is a child, and those that are not capable of doing so.
- 3.2 Throughout this section, we refer to age assurance **methods** and **processes**.
- An **age assurance method** refers to a particular system or technology that underpins an age assurance process.
 - An **age assurance process** refers to the end-to-end process through which the age assurance method or combination of methods are implemented to determine whether or not a particular user is or is not a child. The effectiveness of an age assurance method will depend on how it is implemented, including whether by itself or in combination with other methods. The age assurance process as a whole needs to be highly effective at correctly determining whether or not a particular user is a child.
- 3.3 We set out below a non-exhaustive list of the kinds of age assurance that we consider are capable of being highly effective at correctly determining whether or not a user is a child.⁹ We recognise that age assurance methods are developing at pace and this list may expand in time. It is for the service provider to determine which age assurance method(s) to use in order to implement an age assurance process that is appropriate to meet its duties under the Act. Implementing one of the example methods is not a guarantee that the service is acting in accordance with the requirements of the Act – service providers need to be able to demonstrate that the method has (or methods have) been implemented in such a way that ensures the overall process as a whole is highly effective.
- 3.4 We also provide examples of methods that we do not consider are capable of being highly effective at correctly determining whether or not a user is a child. Service providers should not rely on these methods alone to determine whether a user is a child in the absence of other measures.
- 3.5 All age assurance methods involve the processing of personal data. As such, service providers who are required to implement age assurance are also subject to the requirements of the UK’s data protection regime and should follow a data protection by design approach. The ICO has issued guidance on how these requirements should be met, as outlined in the ‘Privacy and data protection’ sub-section below, which will assist service providers to implement highly effective age assurance while protecting user privacy in line with the data protection regime.
- 3.6 Service providers have the flexibility to choose to build an in-house age assurance method or purchase a method from a third-party age assurance provider. Additionally, we recognise that there may be wider system-level age assurance methods that service providers could

⁹ The kinds of age assurance in this list may be referred to by different names, and each kind may be implemented in a number of ways. We have used high-level descriptions to assist service providers in understanding the options that are available to them, but it is for each provider to consider which age assurance methods and processes will be most appropriate for complying with the duties under the Act.

use to distinguish between children and adults on their service, for example, involving providers of devices, app stores, browsers operating systems, or relevant kinds of authentication systems. Regardless of where the age assurance occurs in the ecosystem or whether it is implemented by the service provider or by a third-party, it is the responsibility of the regulated user-to-user service provider to ensure that age assurance is implemented in such a way that it is highly effective at determining whether or not a user is a child. Should service providers opt to use wider system-level age assurance, they must ensure the initial age check and the process to share this information with the regulated service (e.g. through age tokens) are highly effective.¹⁰

Kinds of age assurance that are capable of being highly effective

Open banking

- 3.7 This works by accessing the information a bank has on record regarding a user's age, with the user's consent. Confirmation of whether or not the user is over 18 is shared with the relying party.¹¹ The user's date of birth is not shared with the relying party, nor is any other information.

Photo-identification (photo-ID) matching

- 3.8 This works by capturing relevant information from an uploaded photo-ID document and comparing it to an image of the user at the point of ID upload to verify that they are the same person.

Facial age estimation

- 3.9 This works by analysing the features of a user's face to estimate their age.

Mobile-network operator (MNO) age checks

- 3.10 Each of the UK's MNOs have agreed to a code of practice whereby they automatically apply a content restriction filter (CRF), which prevents children from accessing age-restricted websites over mobile internet on pay-as-you-go and contract SIMs. Users can remove the CRF by proving they are an adult.¹² MNO age checks rely on checking whether the CRF on a user's mobile phone has been removed. If the CRF has been removed, this indicates that the recorded user of the device is over 18. Confirmation of whether or not the recorded user is over 18, based on the status of the CRF, is shared with the relying party.

Credit card checks

- 3.11 In the UK, individuals must be 18 or over to obtain a credit card, therefore, credit card issuers are obliged to verify the age of applicants before providing them with a credit

¹⁰ Age tokens are reusable digital tokens that act as a digital proxy or representation of a completed age check. They can be shared by users across multiple services over a defined period of time as evidence that an age check has been completed.

¹¹ 'Relying party' refers to the service that is trying to establish the age of the user. In this context, the relying party is likely to be the regulated service.

¹² There are several ways to remove a CRF, depending on the MNO.

card.¹³ Credit-card based age checks work by asking a user to input their credit card details, after which a payment processor sends a request to check the card is valid by the issuing bank. Approval by the issuing bank can be taken as evidence that the user is over 18.¹⁴

Email-based age estimation

- 3.12 These are solutions that estimate the age of a user by analysing the other online services where that user's provided email address has been used. This could include where an email address has been associated with financial institutions such as mortgage lenders.

Digital Identity Services

- 3.13 A digital identity is a digital representation of a person which enables them to prove who they are during interactions and transactions online and in person. Reusable digital identities are those which can be used multiple times for different interactions and transactions. This includes digital identity wallets which enable users to verify and securely store their attributes (such as age) in a digital format. This verification may take place using a variety of methods, including those listed above. Once their identity or an attribute of their identity has been verified and stored in the wallet, a user may choose to share individual attributes, such as their age, or their status as an adult, with a relying party.

Kinds of age assurance that are not capable of being highly effective

Self-declaration of age

- 3.14 The Act states that measures which require users to self-declare their age (without other methods) are not to be regarded as age assurance.¹⁵ These include:
- asking a user to input their date of birth without any further evidence to confirm this information; or
 - asking a user to tick a box to confirm that they are 18 years of age or over.

Age verification through online payment methods which do not require a user to be over the age of 18

- 3.15 For example, debit cards or any other card where the card holder is not required to be 18.

General contractual restrictions on the use of the regulated service by children

- 3.16 For example:
- including as part of the terms of service a condition that prohibits users who are under 18 years old from using the service, without any additional age assurance;
 - general disclaimers asserting that all users should be 18 years of age or over; or
 - warnings on specific content that the content is only suitable for over 18s.

¹³ We are aware that in the US, the term 'credit card' can be used to refer to debit cards. For clarity, when we refer to 'credit card' we mean cards tied to an account where money is borrowed and repaid, and not debit cards tied to current or 'checking' accounts, which often do not have the same 18+ requirements.

¹⁴ Possession of credit card details is not evidence that the user is the credit card holder.

¹⁵ Section 230(4) of the Act.

4. Criteria to ensure an age assurance process is highly effective

- 4.1 Service providers need to: (a) choose an appropriate kind (or kinds) of age assurance; and (b) implement it in such a way that it is highly effective at correctly determining whether a user is a child.
- 4.2 To ensure that an age assurance process is, in practice, highly effective at correctly determining whether or not a user is a child, service providers should ensure that the process fulfils **each** of the following four criteria:
- it is technically accurate;
 - it is robust;
 - it is reliable; and
 - it is fair.
- 4.3 These criteria apply to the technical operation of the age assurance process. Table 4.1 below provides a summary of the criteria, all of which should be considered by service providers to decide their approach to age assurance. In addition to the summary in the table, we give more detail about each criterion below.

Table 4.1: Summary table of the criteria service providers should fulfil and how they can do so.

Criteria	Practical steps to fulfil criteria
Technical accuracy: the degree to which an age assurance method can correctly determine the age of a user under test lab conditions.	<p>Ensure the age assurance method(s) has been evaluated against appropriate metrics and the results indicate that the method(s) is able to correctly determine whether or not a particular user is a child under test lab conditions.</p> <p>Where the age assurance process used on the service involves the use of age estimation, the provider should use a challenge age approach.</p> <p>Periodically review whether the technical accuracy of the age assurance process for the service could be improved by making use of new technology and where appropriate, make changes to the age assurance process.</p>

Criteria	Practical steps to fulfil criteria
Robustness: the degree to which an age assurance method can correctly determine the age of a user in actual deployment contexts.	<p>Implement age assurance processes that have undergone tests in multiple environments during development.</p> <p>Identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children and where it is reasonable to assume that children may use them.</p>
Reliability: the degree to which the age output from an age assurance method is reproducible and derived from trustworthy evidence.	<p>Where age assurance methods forming part of the age assurance process rely on artificial intelligence or machine learning, take steps to ensure that:</p> <ul style="list-style-type: none"> the artificial intelligence or machine learning method(s) has been suitably tested during the development of the age assurance process to ensure it produces reproducible results; once deployed, the artificial intelligence or machine learning method(s) is regularly monitored to ensure it produces reproducible results; the outputs of the artificial intelligence or machine learning method(s) are assessed against key performance indicators designed to identify whether the artificial intelligence or machine learning produces reproducible results; in circumstances where the artificial intelligence or machine learning used is observed to be producing unreliable or unexpected results, the root cause of the issue is identified and rectified. <p>Take steps to ensure that any data relied upon as part of the age assurance process comes from a trustworthy source.</p>
Fairness: the extent to which an age assurance method avoids or minimises bias and discriminatory outcomes.	<p>Ensure that any elements of the age assurance process which rely on artificial intelligence or machine learning have been tested and trained on data sets which reflect the diversity in the target population.</p> <p>For methods reliant on artificial intelligence or machine learning, ensure the age assurance method(s) has been evaluated against the outcome / error parity and the results indicate that the method(s) does not produce significant bias or discriminatory outcomes.</p>

4.4 We recognise that different kinds of age assurance – or even the same kinds of age assurance provided by different companies – may perform more strongly in some of these

criteria than others. For example, one age assurance method could produce a highly reliable result due to limited variance, but it may provide greater opportunities for children to circumvent, therefore reducing its robustness. We expect to see that, when determining which age assurance method(s) to implement, service providers have satisfied themselves that the age assurance process as a whole fulfils each of the criteria.

- 4.5 Throughout the guidance, we refer to the importance of testing. Testing plays a key role in how service providers can evidence that they have had regard to the four criteria. Where we suggest that service providers should consider testing, in all instances, metrics and results could be derived from testing by the service provider internally (if feasible), by their third-party age assurance provider(s), or by an independent third party. Where testing has been carried out by third parties, providers should understand what tests have been conducted and what metrics have been used to measure the results.
- 4.6 Service providers may choose to implement age assurance methods provided by services that are certified against a standard or scheme, such as the UK Digital Identity and Attributes Trust Framework (“the trust framework”).¹⁶ The trust framework is a set of rules and standards governing the provision of digital verification services across the UK economy. Using a service certified against the trust framework (or any other standard or scheme) is not an automatic means of compliance, but it may help to evidence that a service provider has had regard to the four criteria to ensure that its approach is highly effective.
- 4.7 We give more details about each criterion below, including why the criteria are important, and steps service providers can take to have regard to them.

The technical accuracy criterion

What is technical accuracy?

- 4.8 Technical accuracy describes the degree to which an age assurance method can correctly determine the age or age range of a user under test lab conditions.
- 4.9 It is an indicator of the performance of an age assurance method and can be applied to methods that assess a user's age, age range, or whether a user is above a certain age.

Why is technical accuracy important?

- 4.10 An age assurance method which performs poorly in test conditions will perform worse in actual deployment contexts and is therefore unlikely to be highly effective at correctly determining whether or not a particular user is a child when deployed. This indicates that an alternative or additional age assurance method is likely to be required. Understanding the technical accuracy of the individual age assurance method(s) is therefore an important step in ensuring that the process as a whole is highly effective at correctly determining whether or not a particular user is a child.

¹⁶ Office for Digital Identities and Attributes and Department for Science, Innovation and Technology, ‘[UK digital identity and attributes trust framework](#)’. A register of certified services can be found on [GOV.UK](#).

How can service providers have regard to the technical accuracy criterion?

Ensure the method(s) has been evaluated against appropriate metrics and the results indicate that the method(s) is able to correctly establish whether or not a particular user is a child

- 4.11 To understand the technical accuracy of an age assurance method, service providers should ensure it has been evaluated against appropriate metrics to assess the extent to which they can correctly determine the age or age range of a person under test lab conditions.
- 4.12 Age assurance methods either produce:
- A binary result (for example, categorising users as either over or under the age of 18).
 - A continuous result (for example, providing an estimation of the user's age).¹⁷
- 4.13 In the case of methods that produce a binary result, examples of appropriate metrics include but are not limited to; the True Positive Rate (TPR), False Positive Rate (FPR), and False Negative Rate (FNR).¹⁸
- 4.14 In the case of methods that produce a continuous result, examples of appropriate metrics include but are not limited to the Standard Deviation, Mean Absolute Percentage Error (MAPE), and Cumulative Score (CS).¹⁹
- 4.15 Service providers should be satisfied that the results indicate that the age assurance method(s) is able to correctly establish whether or not a particular user is a child.

Use a challenge age approach for age estimation methods

- 4.16 Where the age assurance process used on the service involves the use of age estimation, the provider should use a challenge age approach.
- 4.17 A challenge age approach is widely used offline when selling age-restricted products in retail environments, for instance, through the retailing strategy 'Challenge 25.' In this approach, anyone who appears to the provider of restricted products to be under the age of 25 should be challenged to provide acceptable ID proving that they are over the age of 18 if they wish to buy alcohol. The 'challenge age' in this scenario would be 25.²⁰
- 4.18 In an online age assurance process, a challenge age approach refers to where a user who is estimated as being under a given challenge age must then undergo a second age assurance step (for example, a different age assurance method) to confirm that they are over the required age.²¹
- 4.19 The challenge age should be set at an appropriate level according to the limits of the technical accuracy of that method, for example, where system testing suggests that there is a significant risk of incorrectly estimating a 17-year-old's age by 7 years above or below. To

¹⁷ The estimation of the user's age will usually be accompanied by a confidence interval or range, which conveys the algorithm's level of uncertainty regarding the prediction. For example, where an age estimation method predicts that a user is 25 years old with a confidence interval of ± 2 years, this means that the method estimates the user's age to fall within the range of 23 to 27 years.

¹⁸ We define each of the metrics set out in the technical glossary in Annex 1 of this document.

¹⁹ We define each of the metrics set out in the technical glossary in Annex 1 of this document.

²⁰ Drink Aware, [Challenge 25](#).

²¹ ACCS, 2022. [Measurement of Age Assurance Technologies](#).

manage this risk a buffer can be set above the age by 8 years, so if the relevant age is 18 then the Challenge Age would be 25. For users estimated to be over the age of 25, no additional verification will be required. Where the method estimates that the user's age is under the challenge age, the user could be required to undergo another age check by a second method that is more technically accurate for that age group.

- 4.20 Using a challenge age approach helps to improve the overall effectiveness of the age assurance process by preventing or minimising borderline cases where the age estimation method incorrectly assesses a user as being an adult when they are a child.

Periodically review the technical accuracy of the age assurance method(s) and make changes where necessary

- 4.21 We expect that the technical accuracy and testing practices of age assurance methods will continue to improve in years to come. Providers should ensure their age assurance processes are reviewed and updated periodically to determine whether newer, more effective technologies and testing practices may provide a higher level of technical accuracy, and, where appropriate, make changes to the age assurance process.

The robustness criterion

What is robustness?

- 4.22 Robustness describes the degree to which an age assurance method can correctly determine the age of a user in actual deployment contexts. Common threats to robustness in the context of age assurance methods include:
- **Conditions that change the quality or characteristics of the input** e.g., poor lighting, blurring, brightness, contrast, or positioning of the user in the image (relevant for methods reliant on visual input e.g., facial age estimation, photo-ID matching, etc).
 - **Circumvention techniques that are easily accessible to children where it is reasonable to assume they may use them** e.g., a child user uploading an image of an ID that does not belong to them.
- 4.23 We acknowledge that it may be possible for some children to circumvent the age assurance process or access control mechanisms that the provider has put in place to meet its duties. However, risks can be mitigated by service providers taking steps to improve the robustness of their age assurance process. We therefore expect service providers to take appropriate steps to mitigate against, and refrain from promoting, any methods of circumvention which are easily accessible to children and where it is reasonable to assume they may use them.

Why is robustness important?

- 4.24 Conditions in actual deployment contexts will vary considerably to those in a test scenario.
- 4.25 If the age assurance method is not robust, there are likely to be discrepancies in how it performs across varying conditions. For example, the performance of the method might be lower where a low-quality camera is used. Therefore, such a method would not be highly effective at correctly determining whether or not a user is a child as it does not perform consistently in a varied set of conditions.

- 4.26 In addition, there may be circumvention techniques which are easily accessible to children and where it is reasonable to assume that children may use them. If the age assurance process is not robust, it will be more vulnerable to circumvention.

How can service providers have regard to robustness when implementing age assurance?

Where relevant, ensure the technology has been tested in a range of conditions

- 4.27 We expect service providers to implement age assurance processes that have undergone testing to ensure the process is highly effective in a range of conditions; e.g. poor lighting, blurring, brightness, contrast, or positioning of the user in the image.
- 4.28 Should service providers choose an age assurance method dependent on visual or audio input, they should ensure that the technology underpinning that method has been tested in multiple environments during its development, to minimise any discrepancies in the performance of the method in actual deployment contexts.

Identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children and where it is reasonable to assume that children may use them

- 4.29 We expect service providers to identify and take appropriate steps to mitigate against methods of circumvention that are easily accessible to children, and where it is reasonable to assume that children may use them.
- 4.30 Where service providers implement age assurance processes that rely on details obtained via a user's identification documents, mobile phone number, email address, or credit card, we expect providers to have means of checking that the details supplied belong to the user attempting to access the service.
- 4.31 Verifying the ownership of a user's details is one method of reducing the risk of a child user circumventing the age assurance process. For example, in the case of a mobile phone number, email address or credit card checks, this could be done through the use of a One Time Passcode (OTP) or multi-factor authentication.
- 4.32 Requiring a photo of the user at the point of ID upload when photo-ID matching helps to verify that the photo ID belongs to that user. Liveness detection provides further confidence that a child user has not uploaded a photo of an adult by ensuring that the user undergoing the age assurance process is present at the time the check is carried out.²²
- 4.33 Liveness detection can also help to ensure that children are not using still images of adults to pass through facial age estimation.
- 4.34 It is possible to obtain fake forms of identification of varying degrees of sophistication. A robust photo-ID check should not be capable of being easily circumvented, and as such, a photo-ID method should be able to detect falsified documentation or manipulation that a child could create or obtain. Government-issued guidance on how to prove and verify

²² Liveness detection is used to ensure that the face being analysed is not a photograph, video, or any other form of spoofed representation. The primary goal is to prevent attackers from using static images (print attack) or pre-recorded videos (replay attack) to trick the system into making inaccurate age estimates.

someone's identity ("GPG45") provides some useful indicators on how a document can be scored to detect certain levels of faked documentation.²³

- 4.35 Repeating an age check could also help to increase the robustness of an age assurance method, to prevent child access via device or account sharing and instances where children may be mistakenly classified as adults during the initial age check.
- 4.36 Service providers should determine whether repeated age checks are needed to secure the robustness of their solution based on the features of their service, and if so, how often it is appropriate to repeat an age check. For example, service providers may decide to age check each unique visitor to a service which does not require users to create an account. When deciding on the frequency of age checks, service providers should be mindful that data protection law requires them to assess the necessity and proportionality of the personal data processing and take a data protection by design approach to implementing the data protection principles.²⁴
- 4.37 In addition, service providers should not publish content on their service that directs or encourages UK users to circumvent the age assurance process or the access controls, for example by providing information about or links to a virtual private network (VPN) which may be used by children to circumvent the relevant processes.

The reliability criterion

What is reliability?

- 4.38 Reliability describes the degree to which the age output from an age assurance method is reproducible and derived from trustworthy evidence.
- 4.39 Reproducibility describes the ability for an age assurance method to perform in a consistent manner, producing the same or similar outputs when given the same or similar inputs.²⁵
- 4.40 Strength of evidence describes the relative weight that should be afforded to the underlying data or documents used as evidence for a user's age.²⁶ It concerns how trustworthy the documents or data are and therefore is indicative of how much reliance, or doubt, a service should place on the output of an age assurance method derived from this evidence.

Why is reliability important?

- 4.41 Without reliability, an age assurance method might correctly determine the same user to be a child in some instances, but not in others. Demonstrating that a method can account for variance and create reproducible outputs is therefore an important element of ensuring that children are prevented or protected from encountering harmful content online.

²³ Cabinet Office and Government Digital Service, 2023, [Guidance – How to prove and verify someone's identity](#). Subsequent references to this document are referred to as 'GPG45.'

²⁴ ICO, 2023. [Data protection by design and default](#)

²⁵ Gundersen OE, Kjensmo S, 2018, [State of the art: Reproducibility in artificial intelligence in Proceedings of the AAAI Conference on Artificial Intelligence](#) 32(1), p. 1645.

²⁶ 'Strength' refers to evidence being harder to forge or counterfeit, as defined in GPG45.

- 4.42 In addition, where age assurance does not rely on trustworthy age evidence, there is a risk that a service incorrectly determines a child to be an adult based on evidence that wrongly suggests they are over 18 in some instances.

How can service providers have regard to reliability when implementing age assurance?

Ensure that methods with a degree of variance have been suitably tested and that ongoing performance is measured and monitored

- 4.43 Age assurance methods that rely on machine learning or artificial intelligence, such as facial age estimation and photo-ID matching, are likely to produce outputs with a degree of variance. There are several reasons for this, including data variability and model complexity.
- 4.44 In addition, the performance of these specific methods may degrade over time due to 'model drift'. This is where the data the method has been trained on becomes less representative of the population using the age assurance method. For example, population demographics may shift over time, resulting in a greater degree of variance.
- 4.45 Where service providers implement age assurance methods that rely on artificial intelligence or machine learning, we expect them to take steps to ensure that it has been suitably tested during the development process to ensure it produces reproducible results, i.e. to ensure that outputs are consistently produced when the method is presented with the same inputs.
- 4.46 There should be regular monitoring and measurement of key performance indicators of the system once the age assurance process has been deployed.²⁷ Where necessary, root cause analysis and retraining should also be carried out where unexpected or unreliable predictions are being observed, particularly where such predictions may risk children being able to access harmful content.
- 4.47 For other kinds of age assurance methods, including credit card age checks, open banking, and MNO age checks, outputs do not generally exhibit any variance of the type described above. In these cases, identical outputs should be produced when the method is presented with the same inputs. While reliability is less relevant for these methods, the service provider should still ensure that they fulfil the criteria of technical accuracy, robustness, and fairness.

Ensure that the evidence used is derived from a trustworthy source

- 4.48 We expect service providers to have confidence in the evidence that the age assurance method is relying on by considering, for example:
- the nature and properties of any identity documents, profiles, accounts, data, etc. used as part of the age assurance process; and
 - the source of the underlying data or documents.

²⁷ For example:

- 1) Age Verification Accuracy Rate (AVAR): the percentage of users correctly identified as belonging to the appropriate age group;
- 2) Age Verification Efficiency (AVE): the time taken to complete the age verification process;
- 3) Drift Threshold: establish predefined thresholds for AVAR and AVE beyond which significant model drifting is considered to have occurred.

- 4.49 In assessing the nature and properties of the relevant evidence, service providers should identify features that they would expect to see in a trustworthy source. When deploying photo-ID matching, for example, these features might include that:
- the evidence has originated from a country or organisation that is recognised as trustworthy;
 - the positioning of the photographs on the evidence does not suggest they have been edited or replaced;
 - the layout or any logos look as expected; and/or
 - the visible security features are genuine.²⁸
- 4.50 Certification against the trust framework indicates that the evidence used by a third-party digital identity or attribute service provider should be reliable.²⁹

The fairness criterion

What is fairness?

- 4.51 Fairness describes the extent to which an age assurance method avoids or minimises bias and discriminatory outcomes.³⁰ It refers here to the internal operation of an age assurance method, rather than external factors, such as a lack of access to a particular form of identification required by the age assurance method, which are covered by the principles below.

Why is fairness important?

- 4.52 Implementing a fair age assurance process is important to avoid discriminatory outcomes for certain groups. For example, an age assurance method that provides outputs with a lower degree of technical accuracy for users of certain ethnicities when relying on facial estimation.³¹ Such an outcome might lead to children being incorrectly determined to be adult users, or adult users being incorrectly determined to be children.

How can service providers have regard to fairness when implementing age assurance?

Ensure the technology has been tested on diverse datasets

- 4.53 We expect service providers to ensure that any elements of the age assurance process which rely on artificial intelligence or machine learning have been tested and trained on data sets which reflect the diversity in the target population.

²⁸ Further examples and information on checking that evidence is genuine or valid can be found in GPG45.

²⁹ Office for Digital Identities and Attributes and Department for Science, Innovation and Technology, '[UK digital identity and attributes trust framework](#)'. A register of certified services can be found on [GOV.UK](#).

³⁰ Fairness is a separate principle in data protection law, which states that that data should be processed lawfully, fairly and transparently. For more information, see ICO, [Principle \(a\): Lawfulness, fairness and transparency](#) and ICO, 2023. [Guidance on AI and data protection](#).

³¹ There may also be obligations for service providers under the Equality Act 2010 and guidance on this can be found at [Equality and Human Rights Commission Guidance](#).

Consider the outcome / error parity

- 4.54 Outcome / error parity can help service providers to understand how an age assurance method performs across groups with different characteristics, in order to mitigate bias and discriminatory outcomes.
- 4.55 A method's outcome parity can indicate that a model is fair if it produces equal numbers of positive or negative outcomes for different groups.
- 4.56 A method's error parity can indicate that a model is fair if it produces equal numbers of errors for different groups.
- 4.57 For methods reliant on artificial intelligence or machine learning, service providers should consider and monitor the outcome / error parity across different characteristics (such as race and sex), as part of demonstrating how they have had regard to the fairness criterion.
- 4.58 Service providers should be satisfied that the results indicate that the age assurance method(s) does not produce significant bias or discriminatory outcomes.

Additional principles for providers to consider

- 4.59 Service providers should ultimately ensure that the age assurance process used is highly effective at correctly determining whether a particular user is a child.
- 4.60 Alongside fulfilling the criteria, the age assurance process should be easy to use and work for all users. Failing to do so might unduly prevent adult users from accessing legal content.
- 4.61 Service providers should therefore also consider the principles set out in Table 4.2 below when implementing age assurance methods or processes. We provide further detail about each principle below, including why they are important, and steps service providers can take to have regard to them.

Table 4.2: Summary table of the principles that services providers should consider in addition to the criteria.

Principles	Practical steps to consider principles
Accessibility: the principle that age assurance should be easy to use and work for all users, regardless of their characteristics or whether they are members of a certain group.	<p>Assess the potential impact that the chosen age assurance method(s) might have on users sharing protected characteristics.</p> <p>Consider offering a variety of age assurance methods.</p> <p>Design the user journey through the age assurance process to be accessible for a wide range of abilities.</p> <p>Make information about the age assurance process available to the user prior to completing the age check.</p>

Principles	Practical steps to consider principles
Interoperability: the ability for technological systems to communicate with each other using common and standardised formats.	Stay up to date with developments in interoperable age assurance methods and use these approaches to reduce the burden on the user where possible and appropriate for the service.

Accessibility principle

What is accessibility?

- 4.62 The Act sets out that in recommending the use of age assurance, or which kinds of age assurance to recommend, Ofcom must have regard to the principles that age assurance should:
- be easy to use, including by children of different ages and with different needs;³² and
 - work effectively for all users regardless of their characteristics or whether they are members of a certain group.³³

- 4.63 We refer to these principles collectively using the term **accessibility**.

Why is accessibility important?

- 4.64 Age assurance processes that are inaccessible either because they are complex, are less accurate for users with different characteristics, or include requirements that certain groups of users are unable to fulfil, may result in users being unable to access a service that they should otherwise be able to use.
- 4.65 To enhance accessibility, it is also important that service providers explain to users what the age assurance process is designed to do and how it works, so that users can understand why it is necessary and how to complete the process.

How can service providers have regard to accessibility when implementing age assurance?

- 4.66 It is for service providers to consider what steps are most appropriate for their service to take to ensure their age assurance process is accessible. Examples of practical steps to improve accessibility could include:
- Assessing the impact that the age assurance process might have on users sharing protected characteristics and including details of this assessment in the written record.
 - Offering more than one age assurance method to assist users who may be unable to, or may find it more difficult to, use certain kinds of age assurance.³⁴
 - Designing the user journey through the age assurance process to be accessible for a wide range of abilities. This might include, for instance, ensuring that users with visual impairments are able to use screen readers to complete the age assurance process, or

³² The Act, Schedule 4, paragraph 12(2)(e).

³³ The Act, Schedule 4, paragraph 12(2)(f).

³⁴ For example, those without credit cards will be unable to complete a credit card check. Those without a driving licence or passport will be unable to undergo a photo-ID check that relies on these documents.

ensuring that all functionality is available from a keyboard for users with limited motor control.

- Making information about the age assurance process available in the form of a pop up prior to completing the age check, for example, as a smaller, new window that appears overlaid on top of the webpage, drawing the user's attention. The text could be included in this window, or the pop up could feature a button prompting users to click for more information.³⁵

4.67 The Web Accessibility Initiative's [Web Content Accessibility Guidelines](#) provide recommendations for how service providers can make services more accessible to disabled people.³⁶

Interoperability

What is interoperability?

- 4.68 The Act sets out that in recommending the use of age assurance, or which kinds of age assurance to recommend, Ofcom must have regard to the principle of interoperability between different kinds of age assurance.³⁷
- 4.69 Interoperability describes the ability for technological systems to communicate with each other using common and standardised formats. It relies on consistent technological approaches being adopted across different methods.
- 4.70 In the context of age assurance, interoperability may involve re-using the result of an age check across multiple services allowing different providers of age assurance methods to share the information, provided this is done in line with data protection laws.

Why is interoperability important?

- 4.71 Interoperability offers a potential benefit to the user experience, as it limits the amount of information that users need to provide when accessing a new service if they have already proved their age elsewhere. This could reduce the time and effort required by users to understand, and input into, different age assurance processes.

How can regulated services have regard to interoperability?

Stay up to date with developments in interoperability

- 4.72 We recognise that the development of interoperable solutions is still at an early stage. Service providers can have regard to interoperability by staying up to date with developments in this area, and considering whether to implement interoperable solutions to age assurance where they exist and are appropriate for the service.
- 4.73 Current efforts at enabling interoperable age assurance include [the euCONSENT project](#), a non-profit non-governmental organisation that has been established with the intention of designing, testing, and implementing extensions to the eIDAS infrastructure to enable open-system, secure and certified interoperable age assurance.

³⁵ This could be part of a service provider's publicly available statement, which we provide more guidance on in Section 6 (5.28-5.5.30).

³⁶ Further guidance on businesses' legal obligations in this area can be found at [Equality and Human Rights Commission Guidance](#).

³⁷ The Act, Schedule 4, paragraph 12(2)(g).

5. Privacy and data protection

- 5.1 All age assurance methods involve the processing of personal data and should follow a data protection by design approach. As such, they are subject to the requirements of the UK's data protection regime.
- 5.2 For an understanding of how to have regard to the importance of protecting users from a breach of any statutory provision or rule of law concerning privacy (including data protection), service providers should familiarise themselves with the data protection legislation, and how to apply it to their age assurance method. This includes by consulting ICO guidance and seeking relevant independent legal advice as service providers deem appropriate.

The Data Protection Regime

- 5.3 The UK data protection regime is made up of several pieces of legislation, including the Data Protection Act ("DPA") 2018, the UK General Data Protection Regulation (UK "GDPR"), and the Privacy and Electronic Communications Regulations ("PECR") 2003.
- 5.4 Together, this legislation provides a risk-based framework for making sure the processing of personal data respects the fundamental rights and freedom of individuals. The Information Commissioner's Office ("ICO") is responsible for upholding information rights through its oversight and enforcement of the legislation.
- 5.5 Service providers should consult ICO guidance when implementing age assurance to understand how to comply with the data protection regime, including its guides to the data protection principles, identifying an appropriate lawful basis, and how to respond to users exercising their individual rights afforded by the UK GDPR.³⁸
- 5.6 The PECR will apply to anyone who stores information on or gains access to information on a user's device, for example, by using cookies or other similar technologies. Where an organisation stores, or gains access to, information on a user's device, for example, by using cookies or other similar technologies, PECR will apply. The ICO has produced [detailed guidance](#) on this topic.

ICO guidance on data protection and age assurance

- 5.7 The data protection principles are the cornerstone of the UK GDPR.³⁹ The ICO guidance includes the data protection principles for UK GDPR which are:
- Lawfulness, fairness and transparency;⁴⁰
 - Purpose limitation;⁴¹

³⁸ ICO, 2023. [A guide to the data protection principles](#); ICO, [A guide to lawful basis](#); and ICO, [Individual rights – guidance and resources](#). ICO [Guidance on controllers/ processors](#).

³⁹ For an overview of each principle, see the ICO's guide to the data protection principles.

⁴⁰ ICO, [Principle \(a\): Lawfulness, fairness and transparency](#).

⁴¹ ICO, [Principle \(b\): Purpose limitation](#).

- Data minimisation;⁴²
 - Accuracy;⁴³
 - Storage limitation;⁴⁴
 - Security;⁴⁵ and
 - Accountability.⁴⁶
- 5.8 To assist in implementing age assurance while protecting user privacy, service providers should familiarise themselves with the ICO's Children's code, and the Commissioner's Opinion on Age Assurance for the Children's code.
- 5.9 The ICO's Children's code is a statutory code of practice which sets out 15 standards that internet society services likely to be accessed by children should conform with and demonstrate that their services use children's data fairly and in compliance with data protection law. The standards include that the best interests of the child should be a primary consideration when designing and developing online services likely to be accessed by children. Services should take the standards of the Children's code into account when implementing highly effective age assurance.⁴⁷
- 5.10 The Opinion outlines how the data protection principles and other requirements can be considered in the context of age assurance. In particular, the Opinion explains how age assurance can form part of an appropriate and proportionate approach to reducing or eliminating the personal information risks children face online and enabling conformance with the Children's code. The considerations set out in the Opinion are technology neutral, making them applicable to any kind of age assurance.⁴⁸

Having regard to privacy under the Act

- 5.11 Services likely to be accessed by children have a duty when deciding on, and implementing, safety measures, to have particular regard to the importance of protecting users from a breach of any statutory provision or rule of law concerning privacy.⁴⁹ Where we have concerns that a provider has not complied with its obligations under data protection laws, we may refer the matter to the ICO.
- 5.12 To demonstrate compliance with this duty, service providers may find it helpful to include details of how they have taken privacy into account when implementing highly effective age assurance in their written record (see paragraph 2.7).
- 5.13 The examples listed below, which reflect relevant principles set out in the ICO's Children's code, are ways to demonstrate consideration of data protection law, which service providers may wish to provide details on in the written record.

⁴² ICO, [Principle \(c\): Data minimisation](#).

⁴³ ICO, [Principle \(d\): Accuracy](#).

⁴⁴ ICO, [Principle \(e\): Storage limitation](#).

⁴⁵ ICO, [Principle \(f\): Integrity and confidentiality \(security\)](#).

⁴⁶ ICO, [Accountability and governance](#).

⁴⁷ A summary of the 15 standards can be found at ICO, '[Code standards](#)' in [Age appropriate design: a code of practice for online services](#).

⁴⁸ ICO, [Children's code guidance and resources](#).

⁴⁹ Section 22(3) of the Act.

- **Conducting a Data Protection Impact Assessment (DPIA).** These are required by data protection law where processing is likely to result in a high risk to the rights and freedoms of individuals. DPIAs will assist service providers in identifying and mitigating the risks arising from their processing of personal data, which can help demonstrate that they have had regard to the importance of protecting users from a breach of any statutory provision or rule of law concerning privacy. As set out in Standard 2 of the Children's code, a DPIA can also help services to minimise and identify the specific risks to children who are likely to access the service which arise from the processing of their personal data.⁵⁰ [Detailed guidance on how to carry out a DPIA](#), and a sample template, can be found on the ICO website.
- **Providing privacy information to users.** Service providers should give users information about why they need to provide any personal data, how it will be processed, how long it will be retained, and if it will be shared with anyone else. Doing so in a child-friendly way will also help services to meet Standard 4 of the Children's code: transparency.⁵¹ More information on privacy notices can be found on the ICO website.⁵²
- **Keeping written records of processing activities.** Most organisations that process personal data must document their processing activities to some extent.⁵³
- **Having up to date data protection policies along with a record of how providers make staff aware of them.** This provides staff with clarity and consistency around their data protection obligations.⁵⁴
- **Having a record of which staff have completed any data protection training programme that is in place.** This helps to ensure all staff have adequate knowledge of data protection, as appropriate for their role.⁵⁵
- **Clearly documenting technical and organisational security measures.**⁵⁶

⁵⁰ ICO, [2. Data protection impact assessments](#)

⁵¹ ICO, [4: Transparency](#)

⁵² See ICO, [Transparency \(cookies and privacy notices\)](#) and ICO, [How to write a privacy notice and what goes in it.](#)

⁵³ ICO, [Records of processing and lawful basis](#). Also see ICO, [Governance and Accountability](#) in [Age appropriate design: a code of practice for online services](#).

⁵⁴ ICO, [Policies and procedures](#). Also see ICO, [Governance and Accountability](#) in [Age appropriate design: a code of practice for online services](#).

⁵⁵ ICO, [Training and awareness](#). Also see ICO, [Governance and Accountability](#) in [Age appropriate design: a code of practice for online services](#).

⁵⁶ ICO, [A guide to data security](#).

6. Technical glossary

Metrics used to measure the accuracy of age assurance

Term	Meaning
Absolute error (AE)	The same as the 'error,' but disregards the sign (i.e., positive or negative) thus focusing only on the magnitude (size) of the difference between the technologically-determined age and actual age.
Accuracy (ACC)	The fraction of the predictions the model got right. The formula is $ACC = (TP + TN) / (TP + TN + FP + FN)$.
Cumulative score (CS)	An aggregated score that is calculated by summing the individual score across over a period of time/category etc.
Error	The user's age determined by the technology minus the user's actual age. An overestimation yields a positive value, whereas an underestimation yields a negative value.
False negative (FN)	An outcome where a model incorrectly predicts a negative class i.e., a user is under 18 and the model predicts their age 18 or over.
False negative rate (FNR) / Miss rate	Measures the proportion of FN against all negative predictions (i.e., FN and TP). FNR highlights the performance of the model in yielding FP results and this should be minimised. The formula is $FNR = FN / (FN + TP)$.
False positives (FP)	For the purpose of age assurance, this refers to an outcome where a model incorrectly predicts a positive class i.e., a user is 18 or over and the model predicts their age as under 18.
False positive rate (FPR)	Measures the proportion of FP against all positive predictions (i.e., FP and TN). FPR highlights the performance of the model in yielding FP results and this should be minimised. The formula is $FPR = FP / (FP + TN)$.
Mean absolute error (MAE)	The central value of the absolute error. It describes the average discrepancy between a user's technology determined age and their actual age, ignoring whether it is an over- or under-estimation. It is calculated by summing the absolute errors for a given number of absolute errors, then dividing this by the number of absolute errors. The formula is $MAE = (1/n) \sum_{i=1}^n y - x $ where n = number of observations in the dataset, y = is the true value, x = is the predicted value.
Mean absolute percentage error (MAPE)	A metric that used to measure the accuracy in a regression analysis, this is useful where relative errors (age range estimations) are more meaningful than absolute errors. $M = (1/n) \sum_{t=1}^n (At - Ft) / At * 100$ Where n = number of times the summation iteration happens, At = actual value and Ft = forecast value.

Term	Meaning
Outcome / error parity	Outcome / error parity is a measure designed to compare how an age assurance process outcome impacts users in different groups, both positively and negatively, and/or how often these different groups of users are subjected to errors.
Standard deviation (SD)	A measure of variation or dispersion of the dataset relative to the mean. A low SD suggests datapoints closer to the mean, whereas a high SD suggests datapoints are more dispersed. $s = \sqrt{\sum((X - \text{MAE})^2 / (n - 1))}$ where X = is the i th point in the dataset, MAE = is the mean absolute error, and n = the number of datapoints in the dataset.
True positives (TP)	An outcome where a model correctly predicts a positive class i.e., a user is under 18 and model predicts their age as under 18.
True positive rate (TPR) / Recall	For the purpose of age assurance, this measures the proportion of TP predictions out of all actual positive instances (i.e., TP and FN). This metric highlights the model's performance in correctly identifying positive cases. The formula is $\text{TPR} = \text{TP} / (\text{TP} + \text{FN})$.

Additional terms used in this guidance

Term	Meaning
Access controls	Technical mechanism(s) which prevents users who have not been age assured, or having been age assured, did not meet the requirements of the age assurance process, from accessing a service (or part of it) or certain content.
Age assurance	Refers to both age verification and age estimation, as defined in section 230 of the Act. For these purposes, self-declaration of age is not considered to be a form of age assurance.
Age assurance method	Refers to the particular system or technology that underpins an age assurance process.
Age assurance process	Refers to the end-to-end process through which the age assurance method or combination of methods are implemented to determine whether or not a user is a child.