Yoti facial age estimation

White paper: model yoti-005 | Full version

November 2026



Executive summary

As regulators introduce stronger protections for children and young people across the world, age assurance is becoming a requirement for more organisations. New rules are placing greater responsibility on online and offline services to establish users' ages and enable age-appropriate access to a wide range of goods and services.

Many geographies in the EU, multiple states in the US, Brazil, Australia and other jurisdictions are introducing age assurance legislation.

Organisations operating across multiple jurisdictions around the world age assurance is fast becoming a practical requirement. They need ways to check that people meet the relevant age requirements without adding unnecessary friction or collecting more personal information than they need.

Over recent years, age estimation and inference approaches have evolved to offer alternatives to document-based checks. This is particularly important for people who may not have access to, or feel comfortable using, identity documents to prove their age. One widely used approach is facial age estimation. It allows organisations to estimate someone's age from an image of their face, without identifying them or comparing their face against a database. It works effectively across online and in-person settings and can be used alongside other age assurance methods. Since launching this approach in 2019, Yoti has conducted over 1.4 billion age checks with a growing average of 1.5 million checks per day.

Mean Absolute Error by age band

Yoti facial age estimation accuracy — Mean estimation error in years split by gender, skin tone and age band

Gender	Female				Male				All
Skin tone	Tone 1	Tone 2	Tone 3	All	Tone 1	Tone 2	Tone 3	All	
6-12	1.1	1.2	1.5	1.3	1.1	1.2	1.2	1.2	**1.2**
13-17	1.2	1.4	1.6	1.4	0.9	1.3	1.5	1.2	**1.3**
18-24	1.7	1.6	1.4	1.5	1.5	1.5	1.4	1.4	**1.5**
25-70	2.5	2.8	3.4	2.8	2.3	2.5	2.9	2.6	**2.7**
6-70	1.8	2.1	2.3	2.1	1.6	2.0	2.3	2.0	**2.0**

Source: Yoti internal evaluation using an independent test data set. See page 15 for further information on how we evaluate models.

© 2026 Yoti Ltd

An effective age check involves more than the model itself. It needs to fit into existing services and wider user journeys, with appropriate safety thresholds and alternative ways to prove age where needed. We believe consumers should have choice in how they assert their age and organisations should have the flexibility to offer different approaches to suit their needs.

Age assurance also needs to work reliably at scale, resist attempts to circumvent it and protect people's privacy. We focus on making it practical for organisations to deploy, with solutions that support their operational regulatory requirements.

Model naming and multiple models

In this white paper, we refer to our facial age estimation models using their model names.

We now have multiple facial age estimation models available in our production environment for organisations to use. Alongside yoti-005, we also offer:

- yoti-003: covered in our July 2025 white paper.
- yoti-004

We are doing this for two reasons. Firstly, it helps organisations clearly identify which model they are using, now that multiple models are available. Secondly, it matches the model names used on the NIST website[1], making it easier to identify our models in its evaluations. Statistics and performance results in this paper relate to yoti-005.

Key takeaways

- **Highly accurate for younger age groups.** 99.1% of 13 to 17 year olds are correctly estimated as under 21, and 98.6% of 6 to 12 year olds are correctly estimated as under 13. Accuracy across ages of interest under selected global regulations, with a two-year buffer applied:
  - Age 13: 99.8% of 6-12 year olds correctly estimated as under 15.
  - Age 15: 96.4% of 6-14 year olds correctly estimated as under 17.
  - Age 16: 97.7% of 6-15 year olds correctly estimated as under 18.
  - Age 18: 99.0% of 6-17 year olds correctly estimated as under 20.
- **Low mean absolute error (MAE).** MAE is 2.0 years across ages 6 to 70, 1.3 years for ages 13 to 17, 1.2 years for ages 6 to 12 and 1.7 for ages 18-30.
- **Consistent performance across different groups.** Gender and skin tone bias is minimised.
- **Independently tested and certified.** Tested by the National Institute of Standards and Technology (NIST) and by the Age Check Certification Scheme (ACCS).
- **Recognised for accuracy.** Ranked by NIST as the most accurate MAE for higher quality interocular distance, the most common images we see in production.
- **Strong protection against circumvention.** Facial age estimation combined with liveness detection, injection attack detection and deception detection makes attempts to fool the system much harder.
- **Proven at scale.** Over 1.4 billion age checks performed worldwide, with an average of 1.5 million checks performed per day.
- **Strong user preference.** 90% of people choose facial age estimation when it's offered as an option for age assurance.

1. NIST face analysis technology evaluation (FATE) age estimation and verification data is published here: https://pages.nist.gov/frvt/html/frvt_age_estimation.html.

Contents

The global age assurance regulatory landscape 5

Where organisations use facial age estimation 6

Facial age estimation in action 6

What is facial age estimation? 7

How does facial age estimation work? 7

How accurate is facial age estimation? 8

How do we measure the accuracy of an age check? 9

Using a waterfall approach 13

Preventing against circumvention 14

Testing and evaluation 15

Deployment and implementation 19

Data privacy 21

Data security 22

Appendix 23

Mean absolute error by year 24

Training data (data used to build the model) 27

Data used for testing 29

Accuracy across the entire data set 29

Accuracy by age, gender and skin tone 29

Standard deviation of absolute error 30

Absolute versus percentage errors 31

Improvement in accuracy as the data set grows 32

False positives 33

Live retail trials in UK 35

YOTI © 2026 Yoti Ltd

The global age assurance regulatory landscape

Age assurance requirements differ across countries and sectors, creating a complex regulatory landscape for organisations operating across multiple markets.

Across the EU, the Digital Services Act (DSA) is driving stronger protections for children and greater scrutiny of how platforms establish age.

Australia introduced new requirements in December 2025. Covered social media platforms must take reasonable steps to prevent under-16s from holding accounts.

Brazil's Digital Child and Adolescent Statute (ECA Digital) came into force in March 2026. It introduces new obligations for digital services to protect children and young people, including requirements relating to age assurance. The country's data protection regulator, the ANPD, oversees the legislation.

In the United States, a growing number of states have introduced laws addressing children's access to social media and other online services. The New York SAFE for Kids Act is the first to explicitly require accuracy standards. The ACCS have confirmed Yoti meets these standards for facial age estimation.

Across Asia, countries including Japan and South Korea have also introduced or considered measures relating to children's online safety and access to age-restricted content.

The Organisation for Economic Co-operation and Development (OECD) has reviewed age assurance laws and policies around the world. Its research highlights the differences between approaches and the speed at which requirements are developing.

In the UK, the Online Safety Act requires services covered by its child-safety duties to use highly effective age assurance in certain circumstances. In October 2026, Ofcom also identified facial age estimation as one of four methods capable of establishing whether someone is 16 or over, ahead of proposed new protections for under-16s on social media.

These examples show that there is no single regulatory approach that applies everywhere. The age thresholds, the services covered, the circumstances when an age check is required and the level of assurance expected can all differ. Organisations need to consider the requirements that apply to their service and choose an approach that provides an appropriate level of age assurance.

Facial age estimation is one way organisations can establish age without requiring a user to provide an identity document.

Where organisations use facial age estimation

Facial age estimation supports age checks across a wide range of services and settings.

Examples include:

- **Social media, live streaming and online communities** - during registration or when users try to access age-restricted features.
- **Gaming** - when users create an account or try to access age-restricted games and features.
- **Online retail and ecommerce** - at checkout, before an age-restricted purchase is completed or delivered.
- **Vaping** - for both in person and online sales
- **Adult content** - before users access age-restricted content.
- **Online dating** - during registration, to establish that users meet the minimum age requirement.
- **Gambling** - during registration and before users access gambling services.
- **Retail and supermarkets** - at self-checkout, when customers buy age-restricted products.
- **Vending machines** - before an age-restricted product is dispensed.
- **Unattended retail and self-service kiosks** - when customers use self-service terminals (supervised or unsupervised) to buy age-restricted products.
- **Live chat and user-generated content** - before users access age-restricted features or content.
- **AI services, chatbots and large language models** - before users access features or content intended for specific age groups.
- **Parental consent and child access controls -** when a service needs to establish a child's age before applying additional controls or requesting parental consent.

Facial age estimation in action

Facial age estimation is already being used across a range of industries and applications worldwide. Social media live streaming platform Yubo uses age assurance for all new users to help check whether they meet its age requirements. Because of the high accuracy of our facial age estimation across different age groups, including whether someone is under or over 13, Yubo is effectively applying age gates to group users by age. This helps to create safer, more age-appropriate experiences on the app.

Kids Web Services (KWS) use facial age estimation to support its parental consent processes, while Xbox uses age assurance to create safer, age-appropriate experiences for their players.

In gaming, Regal Gaming uses facial age estimation to check users' age before play.

OnlyFans uses age assurance to confirm that users meet its 18+ age requirement. It was also one of the first online platforms to introduce an effective age-checking solution ahead of the UK Online Safety Act.

Beyond online platforms, British American Tobacco (BAT) has explored age assurance for online and retail environments.

What is facial age estimation?

Facial age estimation uses an image of a person's face to estimate their age. It allows organisations to determine whether someone meets their chosen age threshold without asking them to provide an identity document or share personal details.

The technology analyses patterns in the facial image to estimate age. It doesn't identify the person or compare their face against a database of known faces. The image is used to estimate age and then deleted. The organisation receives the age result it needs for its use case. It never receives the facial image.

We've been developing and deploying facial age estimation since 2019 and have completed more than 1.4 billion age checks worldwide. We leverage this long-standing expertise to continually improve our technology and understand how it performs across different ages, genders, skin tones, image quality, environments and use cases.

Organisations can use facial age estimation whenever they need to check whether someone meets an age requirement. As regulatory requirements for age assurance grow, our technology is already proven in high-volume, real-world settings. Organisations can check age quickly, effectively and reliably while keeping the experience simple, practical and privacy-preserving for users.

How does facial age estimation work?

A facial age estimation check starts with a person taking a selfie. The technology analyses the image and estimates their age based on patterns in their face. It does this using a 'neural network', a type of AI that we train to recognise patterns associated with different ages.

At a technical level, the model analyses the pixels that make up the facial image. Each pixel contains information about the image and is represented by a number. The neural network processes these numbers and learns which patterns of numbers are associated with different known ages. When it receives a new image, the model applies what it has learned to estimate the person's age.

The model doesn't need to know who the person is. It doesn't look for a match with a known face or compare the image against a database of faces. Its only task is to estimate age from the information in the image.

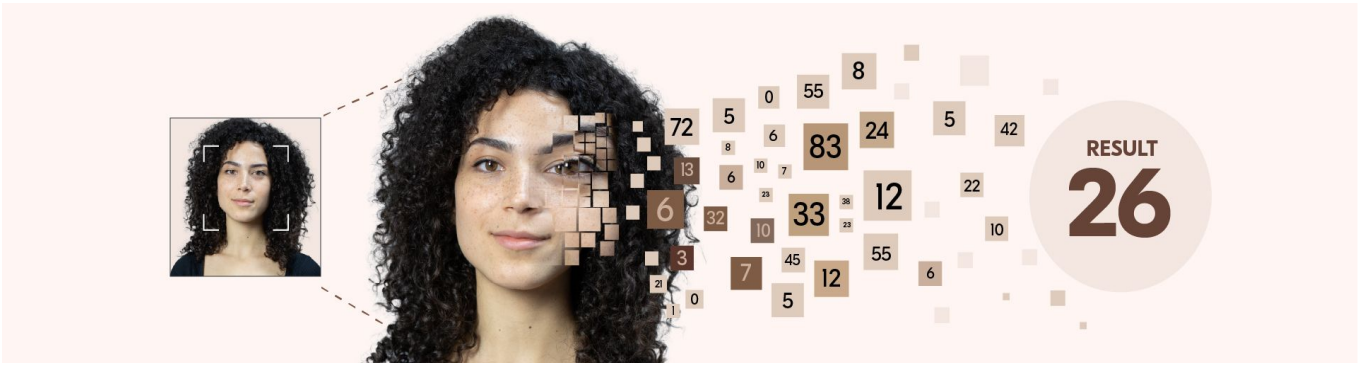

For the person taking the check, the process is straightforward. Depending on the context, they take the image using a camera on their device, such as a mobile phone, tablet or computer, or on a point-of-sale terminal. Their face needs to be clearly visible. Our testing shows that facial hair, glasses and head coverings that don't obscure the face don't prevent a successful check. In some cases, facial disfigurement can affect the accuracy of the estimate.

Image quality also plays an important role. Poor lighting, motion blur or an unclean camera lens can make it harder for the model to estimate age accurately. Our face capture technology checks the image before processing. If an image isn't suitable, it prompts the user to take another one. In most cases, they can resolve the issue by moving to better lighting, holding the device still or cleaning the camera lens.

The amount of usable facial detail in the image is also important. The size of the face within the image, its resolution and the angle at which it's captured can affect how much facial detail the model can analyse. We optimise the capture process to make sure the images contain enough facial detail for the model to produce an accurate estimate.

We train our model primarily using selfies of people looking directly at a mobile phone camera. This reflects how people most often capture facial images during age assurance checks. It also means we train the model using images that are similar to those it will assess in real-world settings.

Once the image has been captured and processed, the organisation receives the age result it needs for its chosen threshold. The captured image is then deleted.

How accurate is facial age estimation?

When an organisation uses facial age estimation, it needs to understand more than how close the estimated age is to a person's actual age.The impact of an incorrect result depends on the person's actual age, the age threshold being applied and what could happen if someone is incorrectly allowed to access age-restricted goods, content or services.

Mean Absolute Error (MAE) is one way of measuring how accurately a model estimates age. It tells us how far the estimated age is from a person's actual age. For example, a model with an MAE of 1.5 years has a mean difference of 1.5 years between its estimated ages and people's actual ages.

MAE is useful for understanding overall performance and comparing models. However, it doesn't tell an organisation how often the model will make the wrong decision at a particular age threshold.

A model's accuracy also depends on the data used to train it. The training images need to reflect the people and image conditions the model is likely to encounter in practice. Testing then helps show how accurately it estimates age across different groups of people and different image conditions.

Furthermore, organisations need to understand how the technology performs at the specific age threshold they want to apply. This involves looking at how often it correctly identifies people who meet the age requirement, and how often it incorrectly allows someone below that age to pass.

How do we measure the accuracy of an age check?

Take an 18+ service as an example. The organisation needs to know whether someone is at least 18 years old. Every age check can produce one of four outcomes:

Person's actual age	Age result	Outcome
The person is 18 or over	The person is 18 or over	Correctly allows access
The person is under 18	The person is under 18	Correctly prevents access
The person is 18 or over	The person is under 18	Incorrectly prevents access
The person is under 18	The person is 18 or over	Incorrectly allows access

These four outcomes have formal statistical names:

True Positive Rate (TPR) - The proportion of people aged 18 or over who are **correctly** estimated as being over 18.
False Positive Rate (FPR) - The proportion of people under 18 who are **incorrectly** estimated as being over 18.
True Negative Rate (TNR) - The proportion of people under 18 who are **correctly** estimated as being under 18.
False Negative Rate (FNR) - The proportion of people aged 18 or over who are **incorrectly** estimated as being under 18.

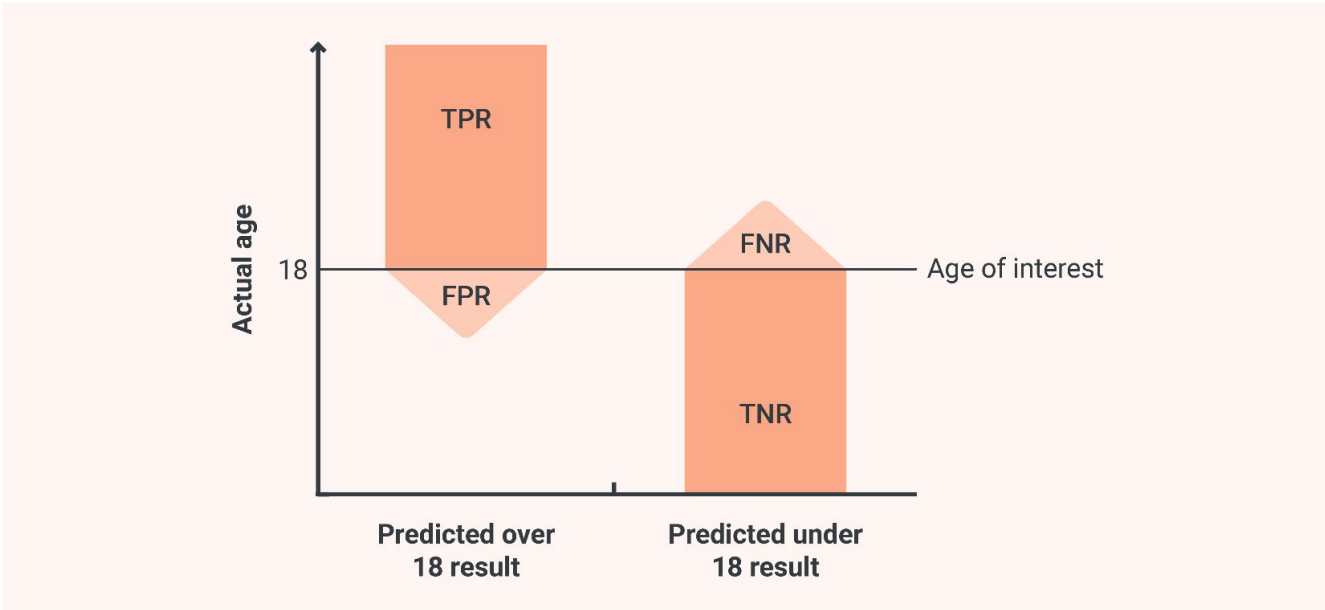

In this context, a ‘positive’ result means that the person is assessed as meeting the age requirement. A ‘negative’ result means that they are assessed as being below it.

These outcomes carry different consequences for organisations. An incorrect rejection can be frustrating for a user who is old enough to access a service. However, they are likely to be able to try again or use another method to prove their age. An incorrect approval can be more serious where access is legally restricted by age. It can allow someone under the required age to access age-restricted goods, content or services.

For this reason, the **False Positive Rate (FPR)** is a critical measure when assessing facial age estimation models. It shows how often the technology incorrectly assesses someone below the required age as old enough. For example, it could be a 14-year-old being incorrectly assessed as old enough to access a social media platform for over-16s.

However, reducing the FPR cannot come at any cost. A model could achieve a very low FPR by rejecting a large number of people who are actually old enough. This would result in a high **False Negative Rate (FNR),** creating unnecessary barriers for eligible users. The aim is to keep both rates as low as possible.

Sensitivity and specificity

Sensitivity and specificity help us understand this balance. **Sensitivity** is the True Positive Rate (TPR) - how often the model correctly identifies people who meet the age requirement. **Specificity** is the True Negative Rate (TNR) - how often it correctly identifies people who fall below it.

Youden’s J statistic combines these two measures to give an overall indication of how well a model distinguishes between two groups. A score of 1 represents a perfect classification and 0 means the model is uninformative. A negative score means its classifications are worse than random. We suggest a factor of 0.9 weight on FPR is a realistic and practical implementation level. This can be seen below for age of interest of 18.

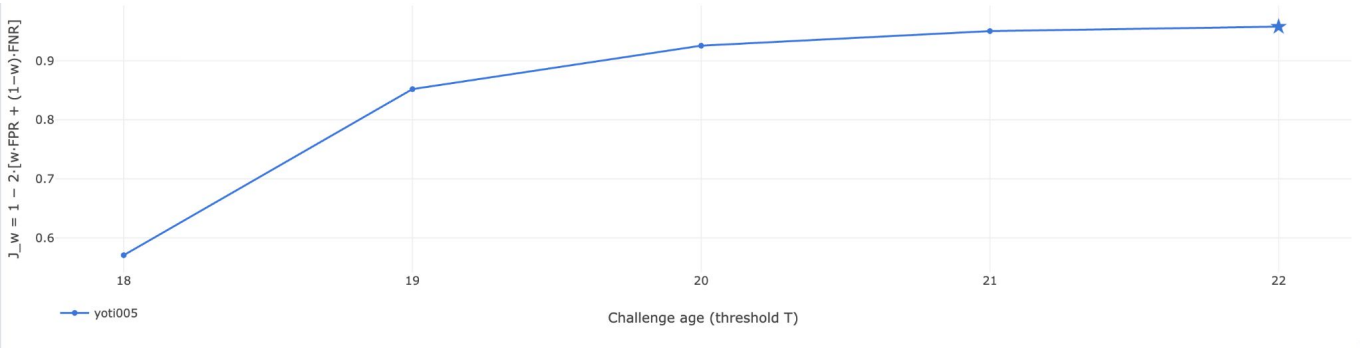

Source: Yoti R&D, yoti-005. For FPR ages 13-17, FNR ages 18-70, for an age of interest 18 - Youden’s J statistic is 0.84 at age 19, 0.91 at age 20 and 0.93 at age 21, 0.94 at age 22, where w=0.9

The level of risk an organisation is prepared to accept will depend on its use case. An organisation providing an 18+ service will have different requirements from one checking whether someone is under 13. For this reason, we recommend setting the facial age estimation threshold above the legal age, with an appropriate safety buffer. Businesses should set an appropriate safety buffer that balances the risk of underage access with the need to make age checks as inclusive as possible for those who meet the age requirement.

Using safety thresholds

We recommend setting a safety threshold above the legal age when using facial age estimation. The difference between the legal age and the safety threshold is called the safety buffer. This reduces the risk of someone below the legal age being incorrectly assessed as old enough to access an age-restricted service.

For example, if a service is for people aged 18+, an organisation could set its safety threshold at 21. This creates a three-year safety buffer. A person would need to be estimated as 21 or over to pass the age check.

For a social media service with a minimum age of 16, a platform or regulator could choose to set a lower safety threshold of 17 or 17.5 years. This would help more people who meet the age requirement, including those aged 17, 18 or 19, to quickly and easily complete the process without needing to undergo an additional age check.

Safety buffers are similar to the 'Challenge' schemes in the UK. Depending on the product and its legal age limit, retailers use policies such as Challenge 21 or Challenge 25. Customers who appear younger than 21 or 25 are asked to prove their age, even though the legal age for the product may be 18. Setting the threshold higher than the legal age gives the retailer an extra margin of safety.

Organisations can apply the same principle to facial age estimation by setting a threshold above the legal age. Unlike a human age check, we can measure how often our technology makes an error at different thresholds. This gives organisations evidence to help them choose a suitable safety buffer for their service.

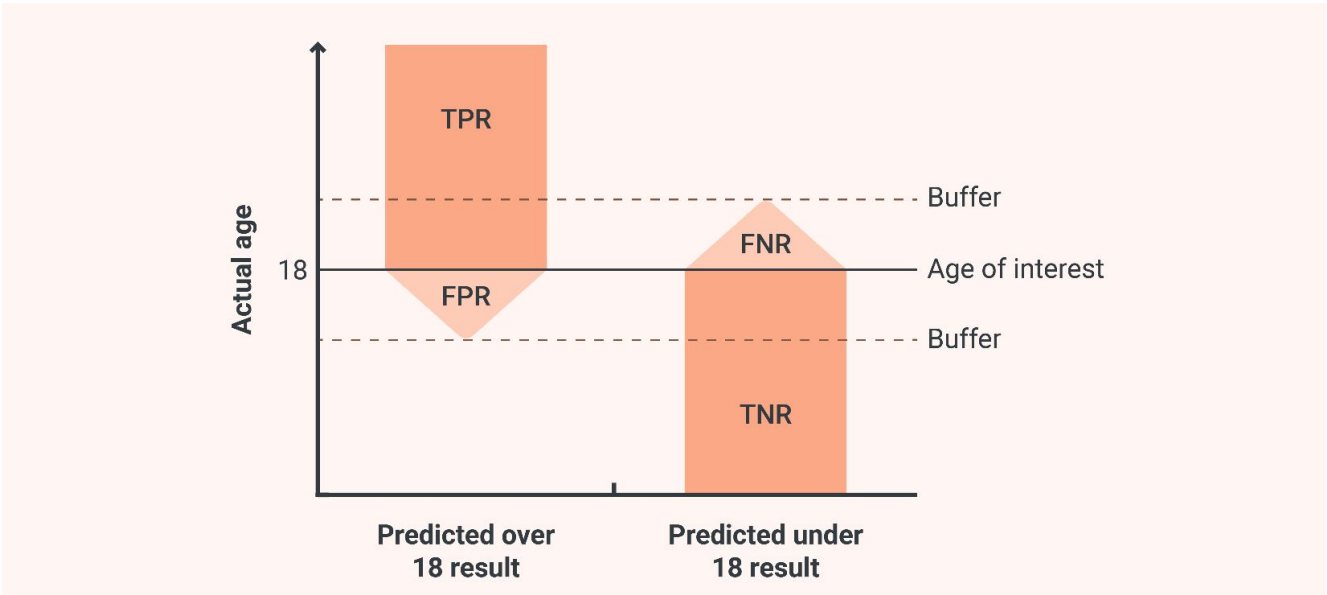

We've been assessing safety thresholds for different use cases since 2019. The appropriate threshold depends on the age being checked, how the technology is used, the risks of incorrect access and the level of protection needed.

What do our results show?

Our latest testing covers the age groups and thresholds most relevant to current age assurance requirements.

Yoti facial age estimation accuracy					Mean estimation error in years split by gender, skin tone and age band				
Gender	**Female**				**Male**				**All**
Skin tone	**Tone 1**	**Tone 2**	**Tone 3**	**All**	**Tone 1**	**Tone 2**	**Tone 3**	**All**	
6-12	1.1	1.2	1.5	1.3	1.1	1.2	1.2	1.2	**1.2**
13-17	1.2	1.4	1.6	1.4	0.9	1.3	1.5	1.2	**1.3**
18-24	1.7	1.6	1.4	1.5	1.5	1.5	1.4	1.4	**1.5**
25-70	2.5	2.8	3.4	2.8	2.3	2.5	2.9	2.6	**2.7**
6-70	1.8	2.1	2.3	2.1	1.6	2.0	2.3	2.0	**2.0**

Highly accurate for younger age groups. 99.1% of 13 to 17 year olds are correctly estimated as under 21, and 98.6% of 6 to 12 year olds are correctly estimated as under 13. Accuracy across ages of interest under selected global regulations, with a two-year buffer applied:

- Age 13: 99.8% of 6-12 year olds correctly estimated as under 15.
- Age 15: 96.4% of 6-14 year olds correctly estimated as under 17.
- Age 16: 97.7% of 6-15 year olds correctly estimated as under 18.
- Age 18: 99.0% of 6-17 year olds correctly estimated as under 20.

Low mean absolute error (MAE). MAE is 2.0 years across ages 6 to 70, 1.3 years for ages 13 to 17 and 1.2 years for ages 6 to 12 and 1.7 for ages 18-30.

Consistent performance across gender and skin tone. Accuracy remains high across both gender and skin tone groups.

- For 13 to 17 year olds, the results are:
  - 99.3% for female participants and 99.5% for male participants.
  - Across skin tones 1, 2 and 3 for this age range, the results are 99.6%, 99.0% and 98.9%, respectively.
- For 6 to 12 year olds, the results are :
  - 99.4% for female participants and 98.6% for male participants.
  - Across skin tones 1, 2 and 3 for this age range, the results are 99.0%, 98.4% and 99.4%, respectively.

Using a waterfall approach

Organisations can use facial age estimation as the first step in a waterfall approach to age assurance If someone doesn't meet the safety threshold, they can then be offered another way to prove their age.

For example, an organisation providing an 18+ service could set its facial age estimation safety threshold at 21. Someone estimated as 21 or over could access the service. Someone estimated below 21 could then be offered another way to prove they are 18 or over.

This is useful for people who fall within the safety buffer. A 19-year-old, for example, is legally old enough to use an 18+ service but wouldn't meet a safety threshold of 21. They could instead be directed to another age assurance method, such as the Yoti ID app or using an identity document.

Offering more than one way to prove age also gives people choice. Facial age estimation provides a quick and simple option for users, while alternative methods support those who cannot use it, do not have a suitable device or would prefer to prove their age another way.

The combination of checks will depend on the service and its legal requirements. Facial age estimation can therefore form part of a wider age assurance process, with different methods used at different points in the user journey. Other age assurance methods include:

Facial age estimation Global	Digital ID apps Global	ID verification Global	Age tokens Global
Credit card check Global	Mobile phone check Regional	Database check Regional	Email address check Regional
Swedish Bank ID Sweden	FTN Finland	MitID Denmark	Double anonymity France
Social Security Number check US	LA Wallet US Louisiana	CPF check Brazil	Singpass Myinfo Singapore

Yoti's age-check services incorporate patent-protected technology relating to a multi-stage age-assurance process, including EP3807794B1 in Europe and US17/13353 in the US. This process involves detecting the human characteristics of a living person at a terminal and using those characteristics to estimate their age. Where facial age estimation is unsuccessful or inconclusive, the process provides for an alternative verification of an age-related identity attribute, such as with a Digital ID or identity document.Further information about Yoti's patents can be found at www.yoti.com/patents/.

YOTI © 2026 Yoti Ltd

Preventing against circumvention

An age check is only useful if people cannot easily manipulate it into producing the wrong result. There are several ways someone might try to do this, from presenting a fake image to changing their appearance to look older. These attempts are known as spoofing attacks or circumvention attempts.

We use multiple layers of protection to address these attempts. Liveness detection checks that a real person is present in front of the camera, while Secure Image Capture and Processing (SICAP) protects the image as it is captured and processed. Our facial age estimation technology is also designed to recognise artificial features that could influence an age estimate.

We always recommend using anti-spoofing measures alongside facial age estimation. Together, these measures help organisations protect age checks against different types of manipulation.

Direct attacks

A direct attack involves presenting something to the camera instead of a person's real face. This could be a photograph, video, mask or deepfake showing someone who appears to be a different age. A bot could also attempt to interact with the check instead of a real person.

Our liveness technology, MyFace, is designed to detect these attempts and check that a real person is physically present in front of the camera. MyFace conforms to iBeta's Level 3 testing under ISO/IEC 30107-3, the highest level of testing in the iBeta presentation attack detection programme.

Indirect attacks

An indirect attack attempts to interfere with the image before it reaches the facial age estimation technology. These are also known as injection attacks. One approach is to override the device's camera feed with a pre-prepared image or video. Another is to modify the image data while it's being processed.

We developed Secure Image Capture (SICAP) to protect against these attacks. It adds security controls at the point when an image is captured and processed. This helps to ensure that the image used for the facial age estimation comes from the device's camera and has not been replaced or manipulated.

Deception detection

Someone may also try to influence their age estimate by changing their appearance to look older. For example, drawing on a fake moustache, adding glasses, or frowning. All of these could be interpreted by the model as being signs of being older. Unlike an image of a different person on a screen or image, the person presents their own face but tries to influence how old they appear.

Our facial age estimation technology is designed to recognise artificial features and account for them when estimating age. This helps reduce the risk of changes to someone's appearance distorting the result. We continue to test and improve our technology to address new attempts to influence results.

Testing and evaluation

We test our facial age estimation technology throughout its development to understand how accurately it estimates age and how consistently it performs across different groups of people and image conditions.

We have published our facial age estimation test results since 2019 through a series of white papers. Previous white papers are listed in the appendix.

A model can perform well on the images it was trained on without performing as well on new images. This is known as overfitting. It can happen when a model learns patterns specific to its training data rather than patterns that apply to other images.

To assess how well our models perform on new images, we use a separate test dataset of around 120,000 facial images, each labelled with the person's actual age. The images cover ages from 2 to 75, with representation across gender and skin tone. Because these image aren't used to train the model, we can use them to assess how accurately it estimates age when presented with images it hasn't previously seen.

We review our testing as the technology develops, using the results to identify areas where performance could improve and decide where further testing is needed.

Fairness

Age assurance needs to work consistently for different people. If performance varies between groups, some people may be more likely than others to be incorrectly allowed access or asked to prove their age another way.

We train our models using images covering a range of ages, skin tones and genders. We also test how accurately they estimate age across these groups. This helps us to identify differences in performance and focus our training and testing on areas where the model can improve.

For our testing, we use three categories to assess skin tone, ranging from lightest (1) to darkest (3). These allow us to assess differences in performance across a range of skin tones.

Our most recent testing shows minimal differences in performance across gender and skin tone for younger people, although we see greater differences among older adults. We continue to work to reduce these differences.

Performance across younger age groups is particularly relevant for age assurance because these are the ages most relevant to current regulatory requirements. For someone who is clearly above an age threshold, a small difference in estimated age is less likely to change the outcome. Around an age threshold, however, even a small difference can determine whether someone passes an age check or is asked to prove their age another way.

The National Institute of Standards and Technology (NIST) also assesses demographic performance. It uses country-level information from its border-image dataset to group people by geographic region as an indicator of ethnicity. Countries originate from regions including East and West Africa, East Europe, East Asia, South East Asia and South Asia.

Our own testing does not currently reproduce this level of geographic detail, so the results are not directly comparable. Our testing allows us to monitor performance across gender and skin tone, while NIST provides an independent assessment across a broader range of geographic groups. Looking at both gives organisations a wider view of how the technology performs across different populations.

Image quality and test conditions

The quality of the facial image can affect how accurately age is estimated. The model needs enough clear facial detail to identify the patterns it uses to estimate age. Factors such as the size of the face in the image, the camera used and the angle of the face can all affect how much detail is available.

One way of measuring this is interocular distance (IOD). This is the distance between a person's pupils in an image, measured in pixels. It gives an indication of how large the face is within the image and how much facial detail is available to the model.

For example, an image can be high resolution but still contain limited facial detail if the face is small within the frame. Conversely, a lower-resolution image can contain more usable facial detail if the face is closer to the camera. The angle of the face can also affect the amount of information available to the model. These factors can influence how accurately it estimates age, regardless of the original image resolution.



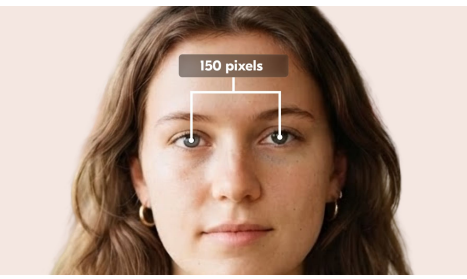

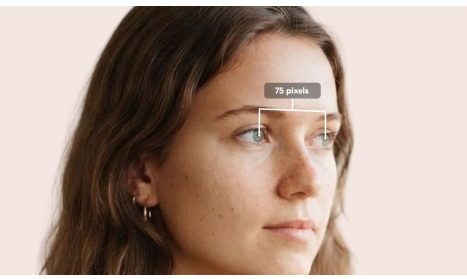

We have trained our model primarily using selfies of people looking directly at a mobile phone camera, with closely cropped faces. Our training images are optimised to a resolution of 720 × 800 pixels. This reflects the type of image the model is trained to assess and how facial images are commonly captured during real-world age assurance checks.

In live deployments, the typical IOD is around 150 pixels across both mobile and desktop devices. This means that most production checks fall within column 5 (120-199) of the table below.

Algorithm	Submission Date	Interocular Distance (IOD) Measured By Pixels					
		15-79	80-119	120-199	200-299	300-499	500-955
		Num. Img. 20448	Num. Img. 136342	Num. Img. 316807	Num. Img. 391201	Num. Img. 47935	Num. Img. 4849
yoti-005	2026-09-21	$1.92 \pm 0.09^{(3)}$	$1.73 \pm 0.03^{(1)}$	$1.70 \pm 0.02^{(1)}$	$1.73 \pm 0.02^{(1)}$	$1.61 \pm 0.05^{(1)}$	$1.5 \pm 0.1^{(1)}$
yoti-004	2026-01-07	$1.93 \pm 0.09^{(4)}$	$1.75 \pm 0.03^{(3)}$	$1.71 \pm 0.02^{(2)}$	$1.74 \pm 0.02^{(2)}$	$1.62 \pm 0.05^{(2)}$	$1.5 \pm 0.1^{(2)}$

Source: NIST FRVT age estimation public testing results. This table tests a range of mugshot images of white male participants aged 18 to 30, tested across six IOD ranges. Outside of the lowest IOD range, yoti-005 and yoti-004 rank first and second respectively.

YOTI © 2026 Yoti Ltd

Our face capture technology helps people capture suitable images by checking lighting, blur, clarity and whether enough of the face is visible. It provides live guidance to help people position their phone correctly and prompts them to retake the image if it isn't suitable.

When comparing facial age estimation performance, the conditions in which images are captured need to be considered. Results from tests using closely cropped, forward-facing images may differ from those using images where the face is smaller or turned at an angle. Results from different tests are therefore only directly comparable when the image conditions are similar.

Independent testing

Our own testing helps us monitor and improve the technology. Independent testing gives organisations an assessment from outside Yoti and shows them how the technology performs under different testing conditions.

The **Age Check Certification Scheme (ACCS)** independently evaluated yoti-004 in October 2024. It reported a Mean Absolute Error (MAE) of 1.1 years for 18-year-olds. Our own testing of the same model showed an MAE of 1.2 years.

KJR, a software quality consultancy, evaluated age assurance methods as part of the Australian Government's Age Assurance Technology Trial (AATT). It looked at how different methods could help protect minors from harmful content.

Researchers from **Princeton University, Carnegie Mellon University and the University of Melbourne** also conducted research into facial age estimation and anti-spoofing performance, involving a number of providers, including Yoti.

NIST has also independently evaluated our technology. NIST assesses the vendor models using four test datasets and specified image sizes. It considers image quality alongside accuracy, using IOD alongside MAE to assess performance at different levels of facial detail.

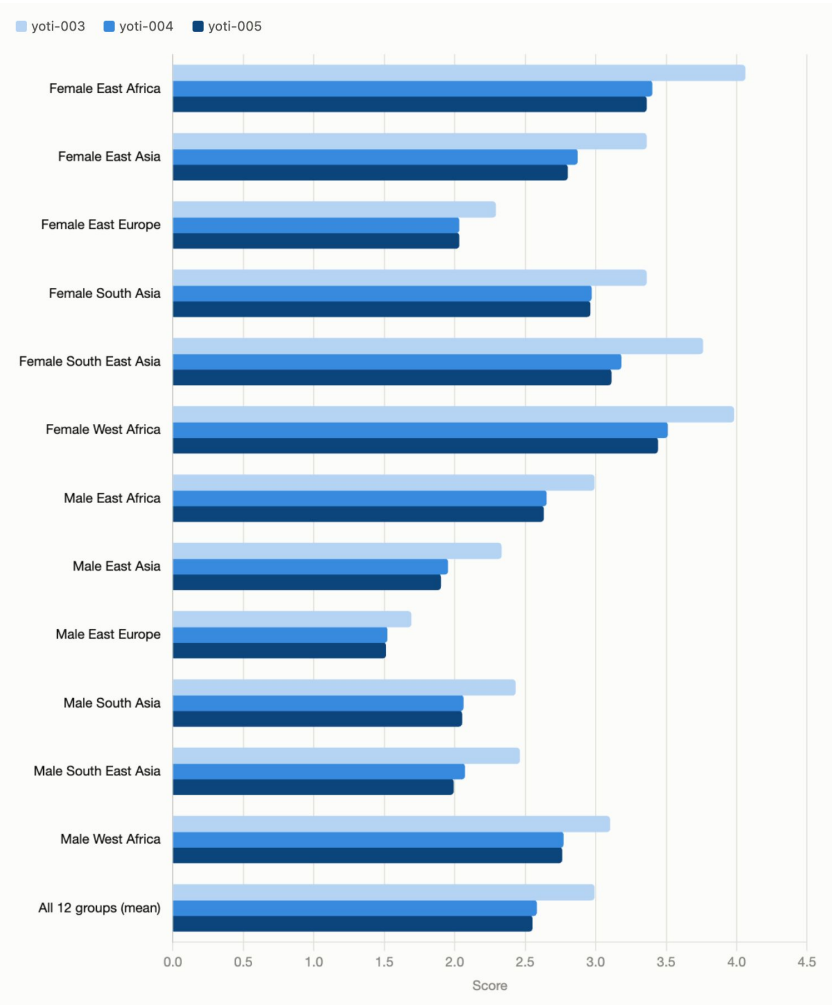

Source: NISTIR 8525 (2026-09-30 update), Tables 150–152, printed pp. 341–343 — Application images are ages 14–30 computed by this report.

Shows improvement in MAE across demographics.

The ACCS and NIST evaluations use different datasets and image conditions. ACCS uses images captured from mobile devices, while NIST's evaluation includes images from sources such as visa and mugshot datasets. These differences are important when comparing results because the conditions in which an image is captured can affect how accurately a model estimates age.

For example, NIST reported an MAE of 2.63 years for 18-year-olds using its visa images from Mexican consular offices, compared with 1.1 years in the ACCS evaluation of yoti-004. The results cannot be compared directly without considering the conditions under which the images were captured. Differences in datasets and image conditions help explain why the same model produces different results across evaluations.

Experience in the real world

Performance in controlled testing is only part of the picture. Organisations also need to know how facial age estimation performs when people use it in real services, across different ages, devices, countries and environments.

We've processed more than 1.4 billion checks worldwide and now handle an average of 1.5 million age checks each day for some of the world's largest platforms. This gives us a substantial body of real-world evidence about how people use the technology. We use this experience to monitor performance and identify issues that may not appear in controlled testing.

Real-world deployment also allows us to assess how the technology performs at scale. This includes how quickly checks can be completed, how reliably the service operates during periods of high demand and how often users need to repeat a check because the image isn't suitable.

We work with large platforms to carry out collaborative testing against their specific requirements and use cases. This helps us assess whether the proposed age assurance approach provides the level of assurance they need and how best to configure it. For example, testing might show that a particular safety threshold or combination of methods is more suitable, or highlight areas where performance could be strengthened. This allows us and our partners to refine the approach before wider deployment.

Running age assurance at this scale also means preparing for changes in demand. A new regulatory requirement, the launch of a service in a new country or a major product launch can cause the number of checks to increase quickly. We work with organisations to forecast these changes and prepare our infrastructure in advance, helping to maintain speed and reliability as demand grows.

Our solution can scale to more than 20 million checks per day, or 200 checks per second. This gives organisations the capacity to support large user bases and increases in demand without having to compromise on the speed of the age-checking experience.

A model can perform well on a test dataset, but a live service has to handle real users, changing traffic, new attempts to circumvent checks and remain available when people need it. Our technology operates in these conditions every day.

Deployment and implementation

Facial age estimation can be deployed in different ways, depending on where an organisation needs to check age and how its users access the service. We support online and in-person use cases, including websites, mobile apps, gaming platforms, social media, e-commerce, retail self-checkouts and other age-restricted environments. Deployment can be achieved at scale within 24 hours.

Online checks

For online services, facial age estimation can be integrated into a website or mobile app. Users capture an image of their face using their device's camera and the service returns the age result that the organisation needs.

The check is designed to complete quickly and fit into an existing user journey, without requiring every user to complete a document check. Organisations can use it during onboarding, when someone first accesses age-restricted content or services, periodically, or at another point where they need to check age.

Retail and in-person checks

Facial age estimation can also run on devices with a camera, such as self-checkouts, tablets, vending machines and other retail terminals. Customers use facial age estimation to prove their age without needing to wait for a member of staff to check their identity document.

The technology can return a result within a second, making it suitable for busy environments where customers need to move through a checkout quickly.

Mobile hand off

Some terminals don’t always have a suitable camera or the connectivity needed to perform a facial age estimation check. In these situations, a mobile hand-off allows the customer to complete the check on their own phone.

When a customer attempts to buy or access an age-restricted product or service, the terminal displays a unique QR code. The customer scans it with their phone and completes the age check there. The result is then returned to the terminal so they can continue their transaction.

This allows organisations to offer facial age estimation without having to install a suitable camera or provide an internet connection at every terminal.

YOTI © 2026 Yoti Ltd

NFC and Yoti ID

Users can also add an age result to their reusable Yoti ID app using facial age estimation, without first scanning a government-issued identity document such as a passport. This provides an alternative for people who don't own or have access to an identity document or would prefer not to use a document-based approach.

At a point of sale, customers can then use their Yoti ID to prove their age through an NFC-enabled payment terminal. This gives them another way to prove their age assurance in retail environments without needing to complete a new facial age estimation check for each purchase.

Server-based deployment

Our standard service processes checks using Yoti's infrastructure. This gives organisations access to our most advanced models and allows us to update the technology as it improves or as we respond to new attack methods.

Server-based deployment requires an internet connection between the device capturing the image and our service. For organisations operating large online services, this provides a straightforward way to integrate age assurance into an existing website or app.

On-device deployment

Some organisations need age assurance in places where a reliable internet connection is not available, such as certain self-checkouts or other retail terminals. In these environments, we can run facial age estimation directly on the device instead of sending the image to our servers.

This allows age checks to take place where a server-based check is not practical. It also allows devices to carry out checks without relying on a constant internet connection.

There are some trade-offs. An on-device model needs to work within the processing and storage limits of the device, so it can't be as large or computationally demanding as the models we run on our servers. A model stored on a device is also more exposed to reverse engineering than a model held within our infrastructure.

For organisations that need the highest level of performance and access to our most advanced models, our server-based service remains the stronger option. For those that need age assurance to work on devices with limited or no connectivity, on-device deployment provides a practical alternative.

We are developing an on-device workflow that combines facial age estimation, face capture, liveness and injection attack detection. Because of the limits of on-device processing, its performance may differ from our server-based service. We assess each organisation's requirements before recommending the most suitable approach.

Model architecture and efficiency

Our models are designed as small, specialised components rather than one large model responsible for every part of a check. Each component focuses on a specific task. This allows us to optimise each model for that task and update it independently as the technology develops.

This architecture also allows us to respond to new attack methods without rebuilding the entire system. If we identify a new attack vector, we can update, optimise or replace the relevant model while keeping the rest of the system in place.

Each model is designed to be efficient as well as accurate. A typical check takes less than a second to process and our models require relatively little computing power. We can also configure their size for different deployment environments, including mobile phones, self-checkout terminals and vending machines.

This gives organisations flexibility over where and how they deploy the technology. The appropriate approach depends on factors such as connectivity, device capability, expected volumes and the level of performance required.

Data privacy

Facial age estimation only needs an image of a person's face to estimate their age. It doesn't need personal details such as their name, address, date of birth or identity document.

We've designed the technology in accordance with the GDPR principles of data minimisation and privacy by design. The image is collected specifically to estimate age and - for server-based checks - is encrypted when sent to our service for processing.

Once the estimation is complete, the image is permanently deleted. We don't keep it, share it with other organisations, sell it or reuse it for other purposes, such as training our models.

The organisation only receives the age result it needs. It doesn't receive the facial image itself. For example, an organisation checking whether someone is over 18 receives an 'over 18' or 'under 18' result without needing to receive or store the person's image. This reduces the amount of personal data the organisation needs to handle. It also means the organisation doesn't need to collect and retain an identity document simply to check whether someone meets an age requirement.

Training data used to develop our models is collected separately. We obtain this data during the onboarding process for the Yoti ID app, in accordance with data protection requirements. We also conduct various ad hoc data collection exercises including from trusted third-party suppliers, where additional training data is needed. This helps us improve accuracy and support consistent performance across different demographic groups.

Facial age estimation is not facial recognition

Facial age estimation and facial recognition use facial images for different purposes. Facial recognition is used to determine whether two images show the same person. NIST describes facial recognition as comparing information from two photographs to determine whether they are from the same person.

Facial age estimation doesn't identify the person in the image or compare their face with a database of known faces. It only analyses the facial image to estimate the person's age. The technology doesn't need to know the person's name or identity to make that estimate.

The UK's Information Commissioner's Office (ICO) has recognised the distinction between facial age estimation and facial recognition. In its 2021 assessment of Yoti's facial age estimation technology through the ICO Sandbox, ahead of the Age Appropriate Design Code, the ICO concluded that the technology did not result in the processing of special category data because it was not used to uniquely identify individuals.

Data security

Our security practices are independently assessed against recognised standards, including ISO 27001, the international standard for information security management, and SOC 2 Type II. These are an independent review of our data handling and security practices to ensure they meet recognised standards and remain effective as our services develop.

In September 2026, we also gained certification against ISO 27566, the international standard for online age assurance. The standard covers privacy and safety requirements for age assurance services.

For server-based checks, images are encrypted when transmitted to our service for processing. They are processed within our service and permanently deleted once the facial age estimation is complete.

We regularly assess our security controls to identify and address risks as our services develop and new threats emerge.

Appendix

This appendix provides further context and information on the accuracy of yoti-005.

© 2026 Yoti Ltd

Mean Absolute Error by year

Age	Gender								
	Female				Male				All
	Skin Tone								
	Type 1	Type 2	Type 3	All	Type 1	Type 2	Type 3	All	
	MAE	MAE	MAE	Average MAE	MAE	MAE	MAE	Average MAE	Average MAE
6	1.2	1.7	1.7	1.5	1.2	1.8	1.5	1.5	1.5
7	1.0	1.1	1.1	1.1	0.9	1.0	1.8	1.3	1.2
8	1.3	1.1	1.5	1.3	1.3	1.6	0.9	1.3	1.3
9	1.1	1.1	1.3	1.1	1.5	1.1	0.9	1.2	1.1
10	0.9	1.1	0.8	1.0	0.8	1.1	0.9	0.9	0.9
11	0.9	1.1	2.0	1.3	1.0	1.0	1.1	1.0	1.2
12	1.2	1.1	1.8	1.4	1.0	1.3	1.2	1.2	1.3
13	2.0	2.2	2.5	2.2	1.4	1.4	2.0	1.6	1.9
14	1.5	1.9	2.4	1.9	1.1	1.5	2.2	1.6	1.8
15	1.2	1.5	1.8	1.5	1.0	1.3	1.8	1.3	1.4
16	1.0	1.3	1.5	1.2	0.8	1.2	1.4	1.1	1.2
17	1.0	1.0	1.0	1.0	0.9	1.2	1.1	1.0	1.0
18	1.2	1.2	0.9	1.1	1.1	1.2	1.0	1.1	1.1
19	1.6	1.5	1.4	1.5	1.4	1.3	1.3	1.3	1.4
20	1.9	1.9	1.7	1.8	1.7	1.5	1.6	1.6	1.7
21	2.1	2.0	1.8	2.0	1.8	1.6	1.5	1.6	1.8
22	2.1	2.0	2.3	2.1	1.8	1.7	1.6	1.7	1.9
23	2.1	2.1	2.8	2.4	1.7	1.6	1.7	1.7	2.0
24	1.8	2.1	3.1	2.3	1.8	1.7	1.9	1.8	2.1
25	1.9	2.0	3.0	2.3	1.4	1.8	2.0	1.7	2.0
26	1.7	2.2	2.8	2.2	1.5	1.8	2.0	1.7	2.0
27	1.8	2.3	3.5	2.5	1.5	1.7	2.0	1.7	2.1
28	2.1	2.2	3.2	2.5	1.7	2.0	2.2	2.0	2.2
29	2.2	2.4	3.7	2.8	2.0	2.0	2.3	2.1	2.4
30	2.2	2.9	3.1	2.7	2.1	2.4	2.9	2.4	2.6

YOTI © 2026 Yoti Ltd

Mean Absolute Error by year

Age	Gender								
	Female				Male				All
	Skin Tone								
	Type 1	Type 2	Type 3	All	Type 1	Type 2	Type 3	All	
	MAE	MAE	MAE	Average MAE	MAE	MAE	MAE	Average MAE	Average MAE
31	2.5	2.9	4.2	3.2	1.9	2.6	2.3	2.3	2.7
32	2.7	2.9	3.9	3.2	2.3	2.9	3.4	2.9	3.0
33	3.0	3.5	4.1	3.5	2.6	3.0	3.0	2.9	3.2
34	2.7	3.2	3.7	3.2	2.4	3.1	3.0	2.8	3.0
35	2.9	3.0	4.5	3.5	2.4	2.8	2.9	2.7	3.1
36	2.4	2.6	4.2	3.1	2.5	2.8	3.0	2.8	2.9
37	3.1	3.5	3.5	3.4	2.4	2.8	2.7	2.6	3.0
38	2.2	2.6	3.1	2.6	2.0	3.1	2.9	2.7	2.7
39	2.6	3.0	2.9	2.8	2.2	2.3	3.6	2.7	2.8
40	2.3	2.4	3.1	2.6	2.4	2.2	2.1	2.2	2.4
41	2.2	2.6	3.1	2.6	2.0	2.2	2.6	2.3	2.5
42	2.4	2.6	2.5	2.5	2.3	2.7	2.9	2.6	2.5
43	2.5	2.8	3.7	3.0	2.4	2.3	2.9	2.5	2.8
44	2.3	2.7	2.8	2.6	2.1	2.4	2.9	2.5	2.5
45	2.3	2.6	3.1	2.7	2.1	2.5	2.6	2.4	2.5
46	2.7	2.9	3.3	2.9	2.5	2.6	2.7	2.6	2.8
47	2.8	2.5	3.3	2.9	2.4	2.6	2.9	2.6	2.8
48	2.6	2.8	2.8	2.7	2.4	2.6	3.0	2.7	2.7
49	2.7	2.8	3.4	3.0	2.6	2.6	3.3	2.8	2.9
50	2.5	2.3	2.7	2.5	2.5	2.5	2.9	2.7	2.6
51	2.8	2.4	4.7	3.3	2.5	2.6	3.4	2.9	3.1
52	2.6	2.0	3.6	2.7	2.6	2.4	4.2	3.1	2.9
53	2.6	2.7	3.4	2.9	2.6	2.8	3.6	3.0	3.0
54	2.6	2.6	3.4	2.9	2.7	3.8	3.6	3.4	3.1
55	2.6	3.6	3.5	3.2	2.5	3.5	3.5	3.2	3.2

YOTI © 2026 Yoti Ltd

Mean Absolute Error by year

Age	Gender								
	Female				Male				All
	Skin Tone								
	Type 1	Type 2	Type 3	All	Type 1	Type 2	Type 3	All	
	MAE	MAE	MAE	Average MAE	MAE	MAE	MAE	Average MAE	Average MAE
56	2.2	3.2	3.2	2.9	2.4	3.4	3.7	3.2	3.0
57	2.8	2.7	3.8	3.1	2.7	2.7	3.4	2.9	3.0
58	2.4	2.6	3.1	2.7	2.5	3.0	3.2	2.9	2.8
59	2.3	4.1	2.9	3.1	2.7	3.6	3.3	3.2	3.1
60	2.3	3.0	3.2	2.8	2.5	2.9	4.2	3.2	3.0
61	2.6	3.6	3.0	3.0	2.7	2.9	3.6	3.1	3.0
62	2.3	3.2	4.2	3.2	2.7	3.2	4.1	3.3	3.3
63	2.6	3.1	5.3	3.6	2.9	3.3	4.8	3.7	3.7
64	2.9	3.1	3.9	3.3	2.8	2.8	4.1	3.2	3.3
65	2.7	3.1	3.1	3.0	2.9	3.0	3.2	3.0	3.0
66	2.8	4.6	6.2	4.5	2.8	3.5	3.3	3.2	3.9
67	3.5	3.4	3.2	3.3	3.7	3.9	4.5	4.0	3.7
68	3.5	3.4	9.6	5.5	3.9	3.6	4.5	4.0	4.8
69	4.1	4.2	–	4.1	3.1	3.3	2.8	3.1	3.6
70	3.1	5.1	0.6	2.9	3.0	3.5	7.7	4.7	3.8
Avg	**2.2**	**2.5**	**3.0**	**2.6**	**2.1**	**2.4**	**2.7**	**2.4**	**2.5**

YOTI © 2026 Yoti Ltd

Training data (data used to build the model)

Since 2015, we have invested in building an experienced R&D team working across a range of AI projects.

We mainly train our facial age estimation models using images from Yoti ID app users. During onboarding, we explain how we use biometrics and provide links to further information. This includes our Privacy Notice, which explains how our R&D team use user data for internal research.[10]

Users can opt out of R&D use of their data at any time through the app settings. When they do, this prevents further data from that user being sent to R&D. It deletes all the data associated with that user held on our R&D server. We also delete this data if a user closes their account. We do this even though GDPR Article 17(3)(d) provides an exemption for certain research activities.[11]

We follow a privacy-by-design approach. This means that we can find the data of a specific user to delete it, but there is no way to use that data to recreate a specific user's identity.

To enhance our coverage of particular demographics, we have previously collected additional age-verified images, with consent, to improve coverage across different age groups and demographics. In Nairobi, Kenya, through the Share2Protect campaign, parents and children were able to support the extension of facial age estimation for children aged 6 to 13.[12] We have also purchased additional images of children, with parental consent and their month and year of birth. We carry out thorough due diligence checks on all our data sources.

In 2021, we took part in the Information Commissioner's Office (ICO) Sandbox to expand facial age estimation to children under 13. This is to help check the ages of people who don't have identity documents.[13] The project ended in May 2022, and the ICO later published its exit report.

Our participation helped the ICO gain insights into age assurance for young people. It also provided evidence that facial age estimation does not process special category data. The ICO has since updated its guidance on special category data as a result of this sandbox.[14]

10. https://www.yoti.com/privacypolicy
11. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) https://eur-lex.europa.eu/eli/reg/2016/679/oj
12. https://www.yoti.com/blog/protecting-kids-safer-internet-day-2021/
13. https://ico-newsroom.prgloo.com/news/ico-supports-projects-to-strengthen-childrens-privacy-rights
14. https://ico.org.uk/media/for-organisations/documents/4020427/yoti-sandbox-exit_report_20220522.pdf

Onboarding and R&D opt-out screens in the Yoti ID app

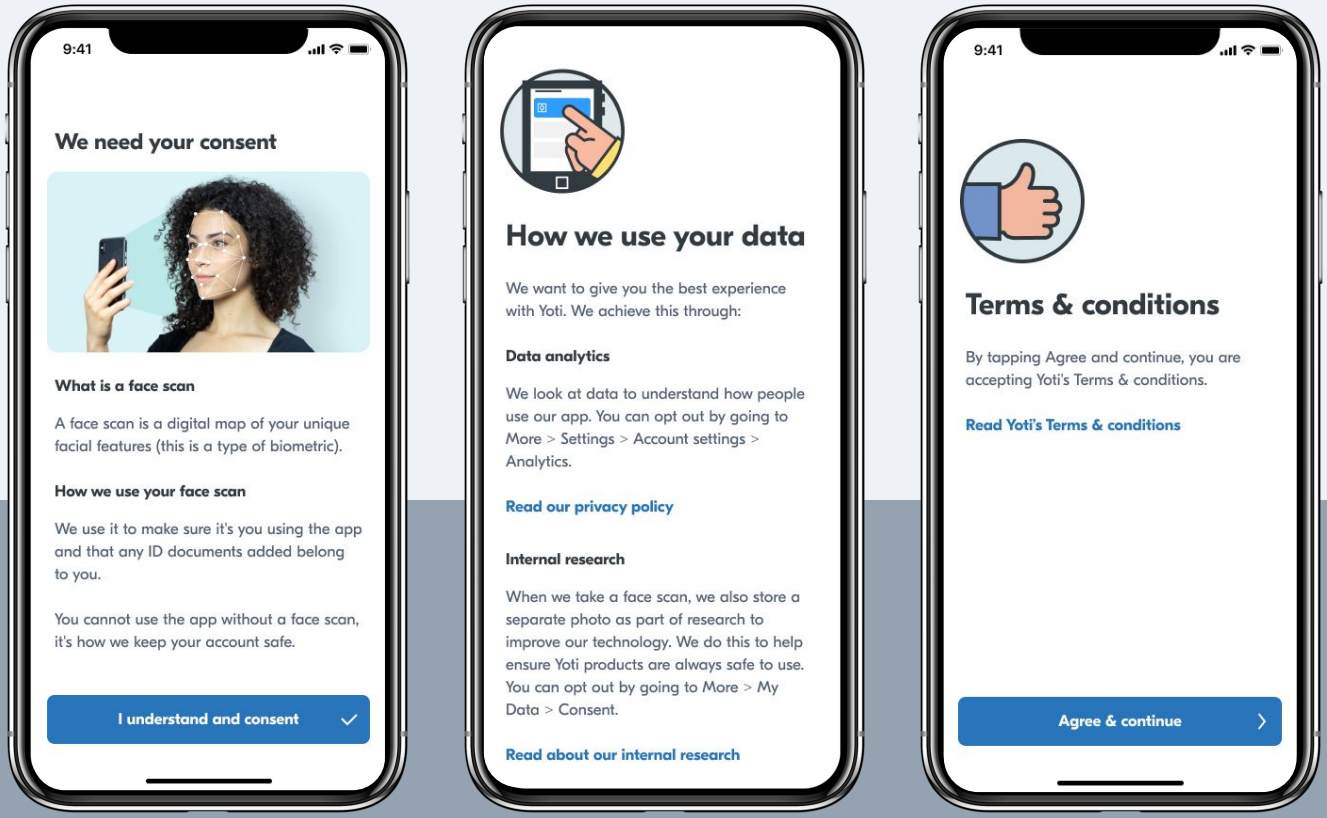

We provide information to users at onboarding about our use of biometrics. This includes links to further information, including our full privacy notice. This explains in details how our R&D team use user data for research. Users can opt out of Yoti using their data for R&D at any time through their app settings.

YOTI © 2026 Yoti Ltd

Data used for testing

We use images from Yoti ID app users around the world to test our models, following the same approach as our training data. We aim to make sure that our test data represents a broad range of people across different ages, genders and skin tones. This helps us understand how the model performs across a wide range of real-world situations.

Accuracy across the entire dataset

In our latest tests of yoti-005, we used hundreds of thousands of test data images with verified ages. The model's Mean Absolute Error (MAE) across all ages is 2.4 years. The MAE is 2.0 years for females and 1.9 years for males. This difference reflects the higher proportion of male images in the training data across most age groups. The range of errors tends towards a normal distribution, with a standard deviation of 1.7 years. Standard deviation measures how much the results vary around the average. The chart below illustrates this distribution.

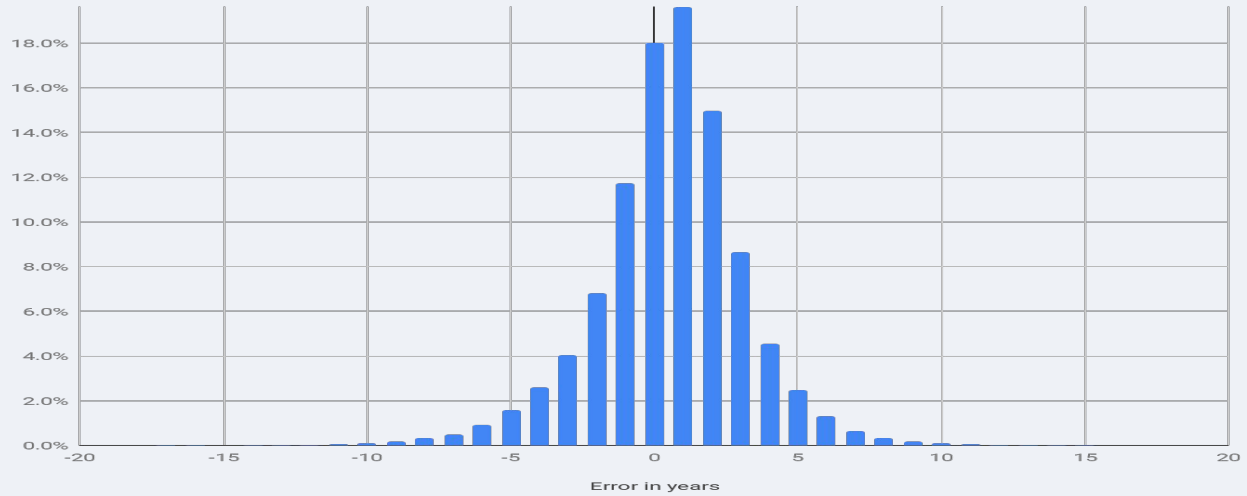

Accuracy by age, gender and skin tone

We tagged hundreds of thousands of images in our test dataset with the subject's verified age, gender and skin tone. We used the gender recorded on each subject's identity document. Our research team assessed skin tone using a three-point scale, from 1 (lightest) to 3 (darkest).

Our team carried out most of the tagging manually, with some automated tagging. We have quality checks in place to help make sure our manual tagging remains consistent and minimises bias.

The average aims to deskew the test data set in order to present equal contributions from the three skin tone groupings and both genders.

YOTI © 2026 Yoti Ltd

Standard deviation of absolute error

Age Band	Gender								
	Female				Male				All
	Skin Tone								
	Type 1	Type 2	Type 3	All	Type 1	Type 2	Type 3	All	
	SD	SD	SD	Average SD	SD	SD	SD	Average SD	Average SD
6-9	0.9	1.0	1.0	1.0	0.9	1.1	1.2	1.0	1.0
10-12	1.0	0.9	1.4	1.1	0.8	0.9	1.1	1.0	1.0
13-15	1.0	1.1	1.2	1.1	0.8	1.1	1.1	1.0	1.1
16-17	0.9	1.0	1.1	1.0	0.8	1.1	1.1	1.0	1.0
18-24	1.3	1.4	1.5	1.4	1.2	1.3	1.3	1.2	1.3
25-29	1.7	1.8	2.4	2.0	1.4	1.5	1.7	1.5	1.8
30-39	2.0	2.5	3.0	2.5	1.7	2.0	2.3	2.0	2.2
40-49	2.0	2.4	2.8	2.4	1.9	2.2	2.2	2.1	2.2
50-59	2.0	2.6	3.0	2.5	2.1	3.0	2.7	2.6	2.6
60-69	2.2	3.0	3.2	2.8	2.3	2.6	3.0	2.6	2.7
All	1.7	2.0	2.2	2.0	1.6	1.9	2.1	1.8	1.9

Mean Absolute Error (MAE) and Absolute Error Standard Deviation

Standard deviation adds context to MAE by showing how consistently a model estimates age.

A higher standard deviation tells us that the errors vary more widely. A lower standard deviation means the errors are closer to the average, showing more consistent performance across the dataset.

YOTI © 2026 Yoti Ltd

Absolute versus percentage errors across age and skin tone

Errors can appear larger in older age groups when measured in years. However, it's worth noting that when we compare errors as a percentage of a person's age, the results can look different.

For example, an error of 2 years represents 13% of a 15-year-old's age, but only 5% of a 50-year-old's age.

People also age differently due to environmental, genetic and health factors. These differences may affect how accurately the model estimates age across different age groups. As a result, older age groups can contribute more to the overall MAE, making it appear higher even when the relative error is smaller.

The chart below illustrates this.Errors can appear larger in older age groups when measured in years. However, it's worth noting that when we compare errors as a percentage of a person's age, the results can look different.

For example, an error of 2 years represents 13% of a 15-year-old's age, but only 5% of a 50-year-old's age.

People also age differently due to environmental, genetic and health factors. These differences may affect how accurately the model estimates age across different age groups. As a result, older age groups can contribute more to the overall MAE, making it appear higher even when the relative error is smaller.

The chart below illustrates this.

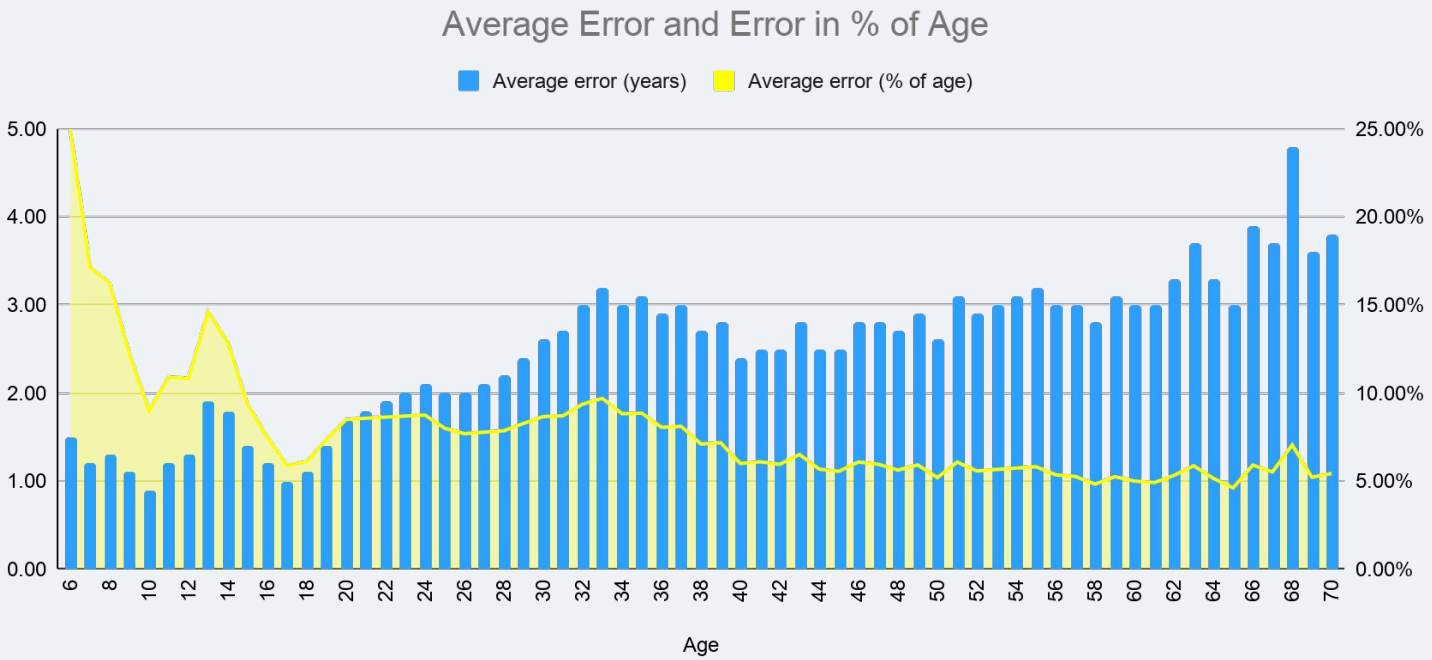

YOTI © 2026 Yoti Ltd

Improvement in model accuracy as the training data set grows and evolves

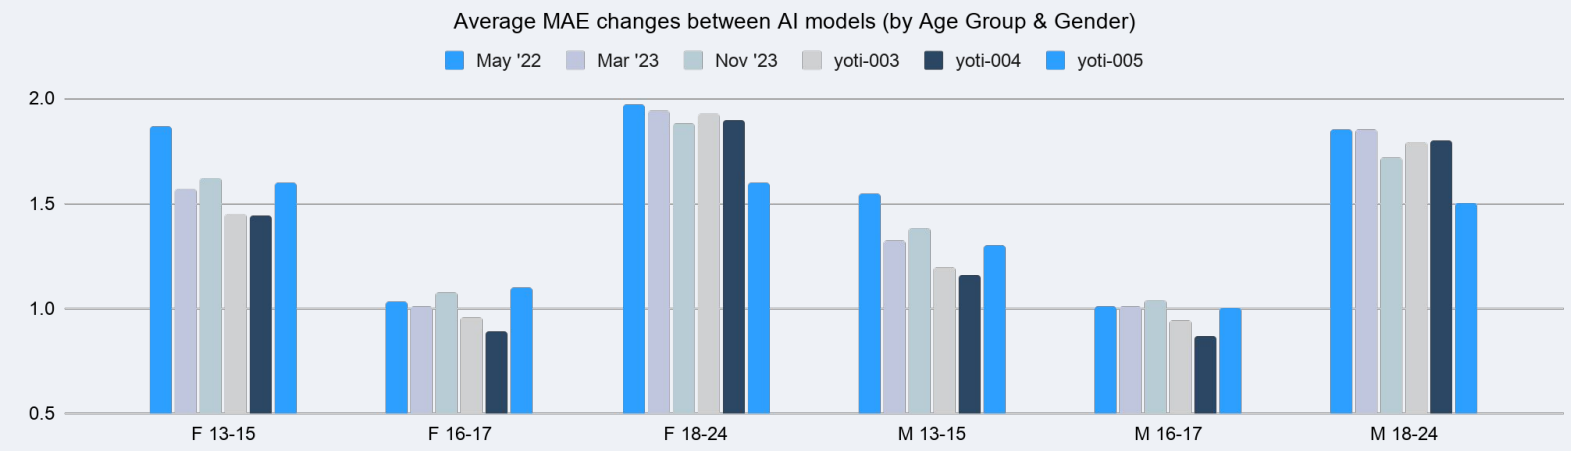

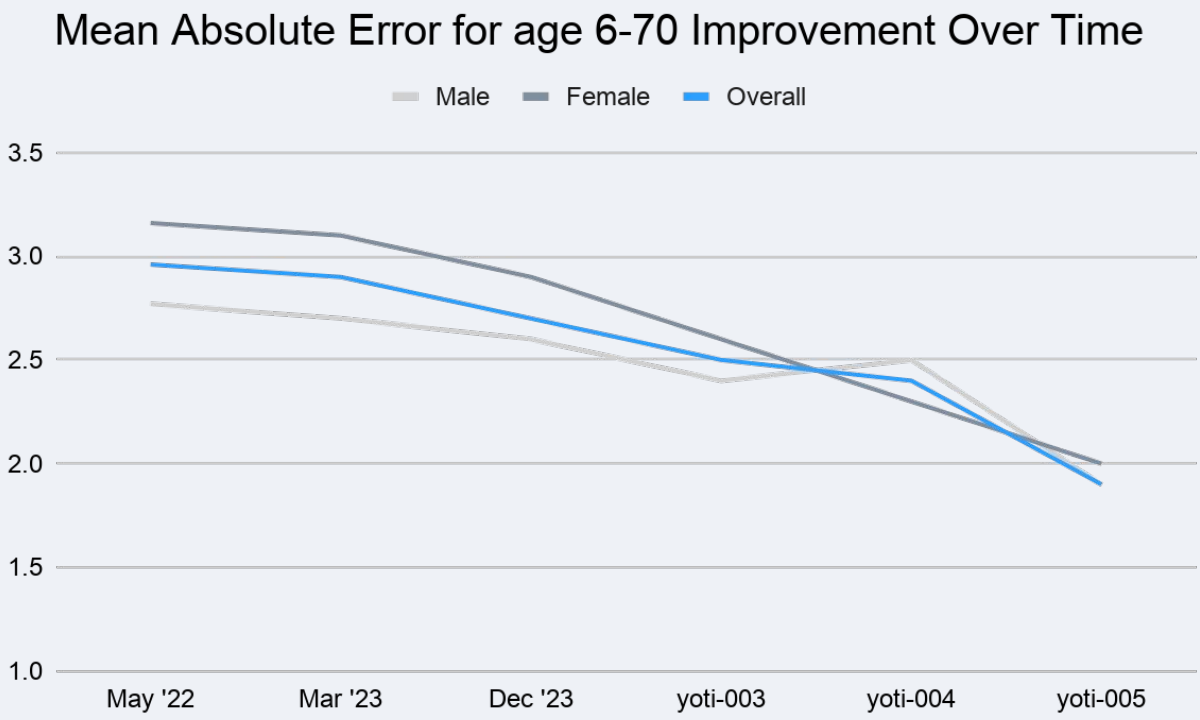

YOTI © 2026 Yoti Ltd

False positive rates at selected thresholds for an age of interest of 18

						Average False Positive Rate (weighted equally for each age)
		14	15	16	17	
Test Sample Size		*3,160*	*7,393*	*10,251*	*10,412*	
Thresholds (years)	20	0.22%	1.10%	2.20%	6.63%	**2.54%**
	21	0.03%	0.39%	1.00%	3.00%	**1.10%**
	22	0.03%	0.11%	0.43%	1.30%	**0.47%**
	23	0.00%	0.07%	0.18%	0.55%	**0.20%**
	24	0.00%	0.07%	0.10%	0.24%	**0.10%**
	25	0.00%	0.04%	0.05%	0.07%	**0.04%**
	26	0.00%	0.01%	0.02%	0.04%	**0.02%**
	27	0.00%	0.01%	0.01%	0.02%	**0.01%**
	28	0.00%	0.00%	0.00%	0.00%	**0.00%**
	29	0.00%	0.00%	0.00%	0.00%	**0.00%**
	30	0.00%	0.00%	0.00%	0.00%	**0.00%**

YOTI © 2026 Yoti Ltd

False positive rates at selected thresholds for an age of interest of 21

		Actual Age					Average False Positive Rate*
		16	**17**	**18**	**19**	**20**	
Test Sample Size		10,258	10,417	9,066	5,236	4,066	
Thresholds (years)	**24**	0.10%	0.24%	0.65%	2.33%	6..83%	**2.03%**
	25	0.05%	0.07%	0.21%	1.03%	2.88%	**0.85%**
	26	0.01%	0.04%	0.08%	0.52%	1.11%	**0.35%**
	27	0.00%	0.02%	0.04%	0.23%	0.47%	**0.15%**
	28	0.00%	0.00%	0.01%	0.06%	0.20%	**0.05%**
	29	0.00%	0.00%	0.00%	0.02%	0.10%	**0.02%**
	30	0.00%	0.00%	0.00%	0.02%	0.05%	**0.02%**
	31	0.00%	0.00%	0.00%	0.02%	0.05%	**0.01%**
	32	0.00%	0.00%	0.00%	0.00%	0.05%	**0.01%**
	33	0.00%	0.00%	0.00%	0.00%	0.02%	**0.00%**
	34	0.00%	0.00%	0.00%	0.00%	0.02%	**0.00%**
	35	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**
	36	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**
	37	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**
	38	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**
	39	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**
	40	0.00%	0.00%	0.00%	0.00%	0.00%	**0.00%**

YOTI © 2026 Yoti Ltd

Live retail trials in UK

In 2022, UK supermarkets including Asda, Morrisons, Tesco and Co-op took part in Home Office trials of our digital age assurance technology at supermarket self-checkouts.. The trials highlighted several benefits:

- Participating supermarkets confirmed they support digital age assurance and said they would welcome changes to legislation in this area.
- The trials recorded no reported cases of underage customers buying age-restricted products when using our technology.
- We gathered informed consent from all customers. They could choose to use the technology or show an identity document to a member of staff.
- The majority of shoppers who used Yoti's digital proof of age solutions liked the technology and said they would use it again, if available.
- Digital age assurance reduced the number of physical age interventions, giving retail staff more time to monitor other activities, including spotting proxy sales.
- Our facial age estimation technology estimate age more accurately than humans, reducing the risk of incorrectly assessing a shopper's age.
- Customers who meet the required age threshold can prove their age without carrying a physical identity document.
- Digital age assurance helps retailers meet their licensing objectives.

Supermarkets can run facial age estimation and liveness checks directly on the self-checkout terminal, without needing an internet connection.

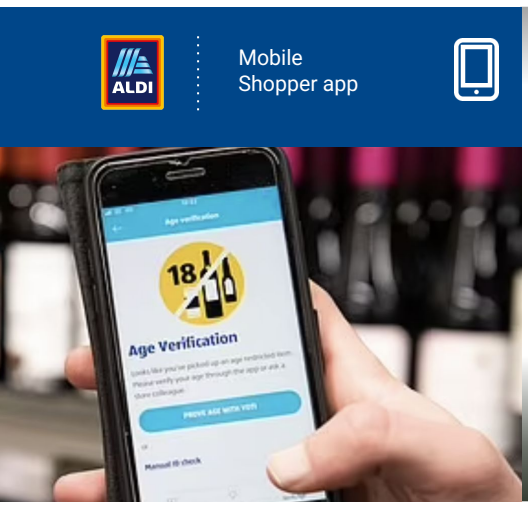

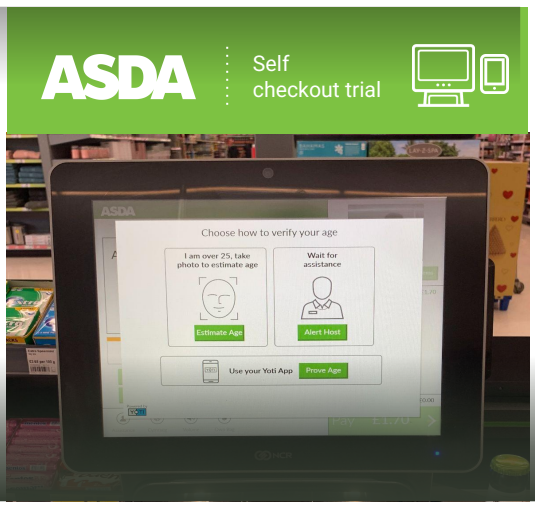

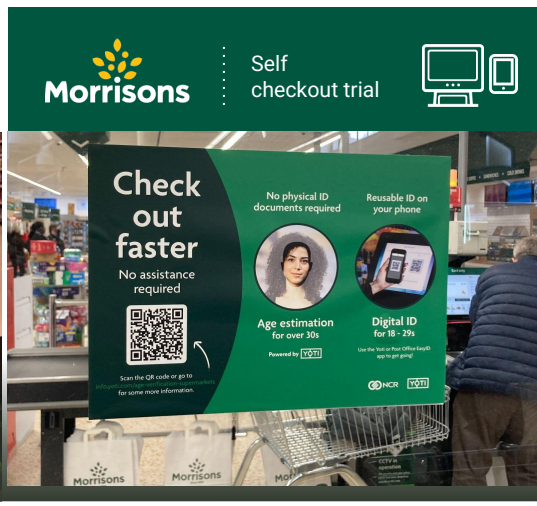

Memberships, associations and accreditations

Reviewed by

To find out more visit **yoti.com**

© 2026 Yoti Ltd