Yoti Facial Age Estimation

White Paper | Full version

July 2025



Executive summary

Yoti's facial age estimation technology can determine a person's age from an image of their face, with no need for a physical document check or human intervention. It is accurate across gender and skin tone.

Yoti's facial age estimation technology is built in accordance with the GDPR principle of 'privacy by design' and data minimisation. Yoti's model has not been trained to recognise faces or match these against other faces in a database. Crucially, this is the difference between facial age estimation and facial recognition. Yoti immediately deletes all images of users as soon as the age is estimated.

Yoti's True Positive Rate[1] (TPR) for 13 to 17 year olds correctly estimated as under the age of 21 is 99.3% and there is no discernible bias across genders or skin tones. The TPRs for female and male 13 to 17 year olds are 99.3% and 99.5% respectively. The TPRs for skin tones 1, 2 and 3 are 99.6%, 99.0% and 98.9% respectively.

The TPR for 6 to 12 year olds correctly estimated as under the age of 13 is 99.0%. The TPRs for female and male 6 to 12 year olds are 99.4% and 98.6% respectively. The TPRs for skin tones 1, 2 and 3 are 99.0%, 98.4% and 99.4% respectively. So there is no material bias in TPR rates in this age group either.

Yoti's facial age estimation is accurate for 6 to 12 year olds, with a mean absolute error (MAE) of 1.3 years, and an MAE of 1.1 years for 13 to 17 year olds. Regulators are focused and most concerned with these two age ranges to ensure that under 13s and under 18s are only able to access age appropriate goods and services.

Yoti takes its ethical responsibilities very seriously when developing its technology. The data used to train the algorithm are obtained, in accordance with GDPR guidelines, during the onboarding process for the Yoti app. We also perform consented data collection exercises and purchase consented data from vetted suppliers where we require training data in efforts to ensure equality of performance across the population. Yoti's latest model continues to show improvements in accuracy.

1. True Positive Rate - the probability that an actual positive will test positive, such as an 17 year old being correctly estimated to be under the age of 21.

Expanding the data set & improving accuracy

Our first white paper, published in December 2018, contained accuracy data by year across the 13-60 age range. Since May 2021, we have added data for the 6-12 age range, and from May 2022 included data for age range 60-70, with performance broken down by year of age, gender and skin tone.

Key takeaways

- TPR for 13 to 17 year olds correctly estimated as under 21 is 99.3%.
- TPR for 6 to 12 year olds correctly estimated as under 13 is 99.0%.
- Mean Absolute Errors (in years) are 2.4 for ages 6 to 70, 1.1 for ages 13 to 17 & 1.3 for ages 6 to 12.
- Users are not individually identifiable.
- Helps organisations to meet Children's Codes or Age Appropriate Design Codes.
- Does not result in the processing of 'special category' data.
- Gender and skin tone bias is minimised.
- NIST ranked number 1 MAE for 13-16 year olds
- Training data is collected in accordance with GDPR.
- Independently tested and certified.
- A secure, privacy respecting solution that protects individuals.
- Yoti liveness and facial age estimation are very hard to 'fool'.
- Over 850 million checks performed worldwide.
- Solution is fast and scales to over 25 million checks per day, or 300 checks per second.

NIST evaluation October 2024

In October 2024, NIST published the results of their evaluation of Yoti's Sep 2024 model, which ranked Yoti in first place when measured for accuracy on 'Visa' images for 13 to 16 year olds and second place when measured for accuracy on 'Mugshot' images for 18-30 year olds.

On 5 July 2025, Yoti submitted its latest Jun 2025 model to NIST for independent evaluation. We expect NIST to publish the results in August 2025. We will update this white paper with the key NIST results.

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**
Skintone	**Tone 1**	**Tone 2**	**Tone 3**	**All**	**Tone 1**	**Tone 2**	**Tone 3**	**All**	
6-12	1.1	1.3	1.5	1.3	1.1	1.2	1.3	1.2	**1.3**
13-17	1.0	1.2	1.5	1.2	0.8	1.0	1.3	1.0	**1.1**
18-24	2.3	2.2	2.5	2.3	2.0	1.8	1.8	1.9	**2.1**
25-70	2.4	2.7	3.2	2.8	2.3	2.6	3.1	2.7	**2.7**
6-70	2.2	2.4	2.8	2.5	2.0	2.3	2.6	2.3	**2.4**

Whilst there are differences in MAEs between skin tones and gender, by setting appropriate thresholds, the true positive rates (TPRs) between skin tones and gender published earlier in this summary are very similar so avoiding discrimination.

With facial age estimation, once you know you're dealing with a child, you can...

- Turn off excessive notifications.
- Minimise the data you collect and do not store it.
- Set geolocation to off but give the child the ability to turn it on if needed.
- Shield their data. It shouldn't be used for purposes not in their interest.
- Provide age-appropriate content.
- Use child-friendly language to explain platforms.
- Be certain the online community is within the same age threshold.
- Always be sure to treat a child like a child.

About 'Mean Absolute Error'
Yoti facial age estimation can make both positive and negative errors when estimating age (that is, it can estimate too high or it can estimate too low). By taking 'absolute' values of each error, this ignores whether the error is positive or negative, simply taking the numerical size of the error. We then take the average (or 'arithmetic mean') of all those absolute error values, to produce an overall 'MAE'. For example, people aged 6-12 have an MAE of 1.3. That means they would be estimated to be within 1.3 years of their age. A table of MAE by year can be found in the appendix on pages 22-24.

© 2025 Yoti Ltd

Contents

What is facial age estimation and what can it do? 6

How accurate is facial age estimation 7

Safety thresholds 8

Buffers and the waterfall approach 9

NIST evaluation 10

NIST evaluation comparison 11

NIST ethnicity approach and methodology 11

ACCS evaluation October '24 12

Growing adoption of facial age estimation 13

Learnings from practical use 14

Live trials in UK & Europe 15

Retail of Alcohol Standards Group report of live trials 16

Supporting Children's Codes 17

Tackling bots, deepfakes and generative AI 18

Legal compliance 19

Yoti's commitment to ethical use of AI technologies 20

Appendix 21

Mean average error by year 22

Data used to build the model ('training data') 25

Data used for testing 27

Accuracy across the entire data set 27

Accuracy by age, gender and skin tone 27

Standard deviation of absolute error 28

Absolute versus percentage errors 29

Improvement in accuracy as the data set grows 30

False positives 32

YOTI © 2025 Yoti Ltd

What is facial age estimation and what can it do?

Yoti facial age estimation is a secure, effective age-checking service that can estimate a person's age using just an image of a face. It provides accurate age checks for businesses providing any age-restricted goods, services or content, both online and in-person. It also helps ensure age checks are more inclusive, given the significant numbers of individuals around the world who do not own a state-issued photo ID document

Yoti's facial age estimation service is designed with user privacy and data minimisation in mind. It does not require individuals to register with Yoti, or provide any documents to prove their identity. It is unable to personally identify an individual. It simply estimates a person's age from analysing an image of their face.

The images are not stored, shared, re-used or sold on. Images are immediately and permanently deleted according to GDPR best practice, and we do not use them for our own learning or training purposes. Our approach is externally reviewed as part of our SOC2 / PAS1296 assessment under control PAS-2.

In a retail setting, facial age estimation can be used at a point-of-sale terminal with a camera, letting a customer choose to prove age at a self-checkout without the need for staff assistance. This is not only quicker and more convenient for shoppers, but can greatly reduce friction and disputes between shoppers and retail staff.

For general online use, it can be embedded into web pages or incorporated into apps, and receive an image of the person's face from the device's camera. This is ideal for controlling access to age-restricted gaming, social media, e-commerce, online dating, gambling and adult content.

Facial age estimation can play an important role in safeguarding children online. As well as preventing minors from accessing adult content, it can prevent predatory adults lying about their age to enter social media spaces designed for children and teenagers. This is illustrated by Yoti's partnership with the Yubo social networking platform. Yubo uses facial age estimation within its app to help identify user profiles where there is suspicion or doubt about the user's age, and flags these cases to its moderation team.

Deployment on premise and on device
Facial age estimation can also be deployed on premise by law enforcement to assess ages of victims and perpetrators in child abuse images. We have also developed a more efficient and lightweight age estimation model that can run on platforms, such as self checkouts, gambling and vending machines, with limited or low computational resources and mobile devices. This model has no reliance on internet connectivity to send an image and receive results from our servers.

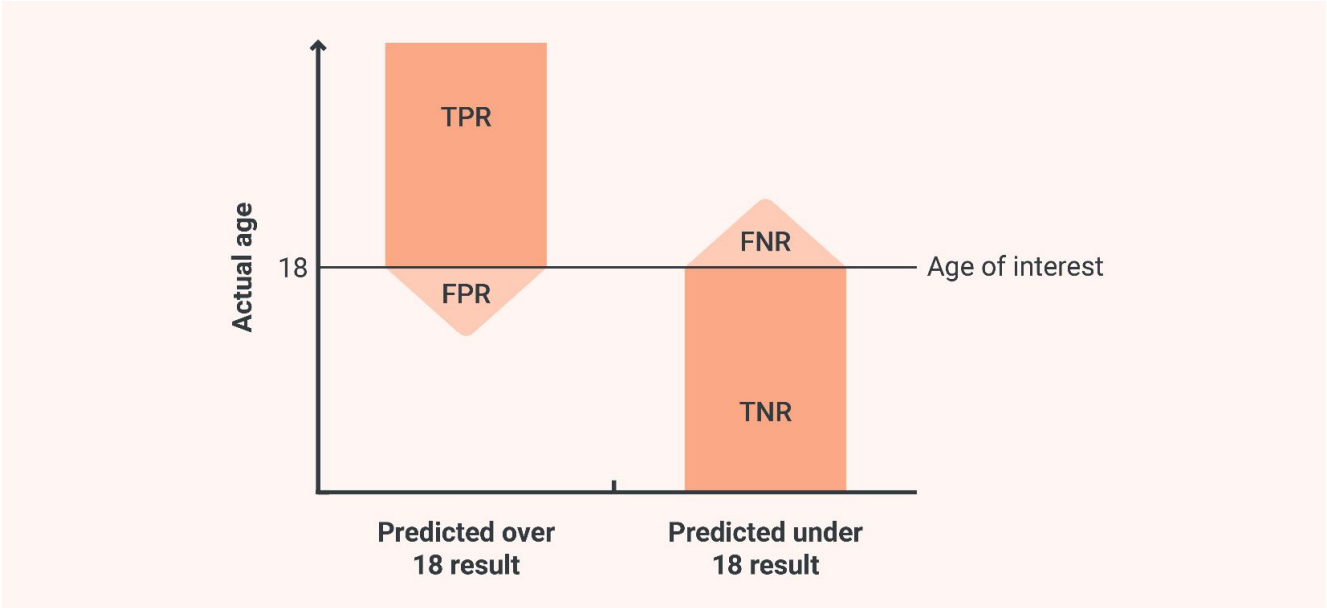
A further potential use is at the entrances to age-restricted premises such as bars, nightclubs and casinos. In this kind of application, facial age estimation offers clear advantages – it does not get fatigued on a long shift,[2] and it cannot show favour to personal friends, or selective bias against individual customers. It is very hard for under 18s to 'fool'. It also reduces the burden on staff to try and estimate customer ages and it reduces abuse to staff.

2. Studies have shown that the objectivity of human judgement of this kind can be significantly affected by hunger and fatigue – for example, see Danziger, Levav, Avnaim-Passo (2011) *Extraneous factors in judicial decisions*, PNAS April 26, 2011 108 (17) 6889-6892; *https://doi.org/10.1073/pnas.1018033108*

How accurate is facial age estimation?

Accuracy is one of the first questions regulators, businesses and users ask. To date we have mainly presented the Mean Absolute Error (MAE) as a proxy for accuracy.

However, the reality is slightly more complicated. Following the COVID-19 pandemic, many people will be more aware of the terms such as 'true positive' and 'false negative' when it comes to vaccinations and testing, and the 'sensitivity' of a test. That is, what level of risk are we prepared to accept, in return for a 'true' result. We can use the same terms when it comes to accuracy of age estimation.



Here is an example of how the terms apply to facial age estimation. The column on the left shows where a model returns a positive result - this person is over the age of 18. The right hand column the model would return a negative result - this person is under 18. Whilst MAE is an easy proxy for accuracy, this does not consider the intricacies within a mean result.

TPR = True Positive Rate - a result that **correctly** estimates the person as **over** 18
FPR = False Positive Rate - a result that **incorrectly** estimates the person as **over** 18
FNR = False Negative Rate- a result that **incorrectly** estimates the person as **under** 18
TNR = True Negative Rate - a result that **correctly** estimates the person as **under** 18

The goal is to have as high a rate in the two 'True' (correct or accurate) categories as possible.

Beyond the 'accuracy' (TPR and TNR) it is also very important to consider the false results. FNR is inconvenient and annoying for customers and organisations, as anyone who has been lucky enough to be asked to show ID when they are over 18 will know, but no law is being breached. Many customers can choose another method to prove their age.

FPR is problematic as it can put operators at risk of breaking the law depending on local regulations. So when we talk about accuracy, this is the key area of concern - incorrectly allowing some underaged individuals access to goods or services when it is illegal for them to do so. This is why we recommend operators use appropriate safety thresholds.

YOTI © 2025 Yoti Ltd

Safety thresholds

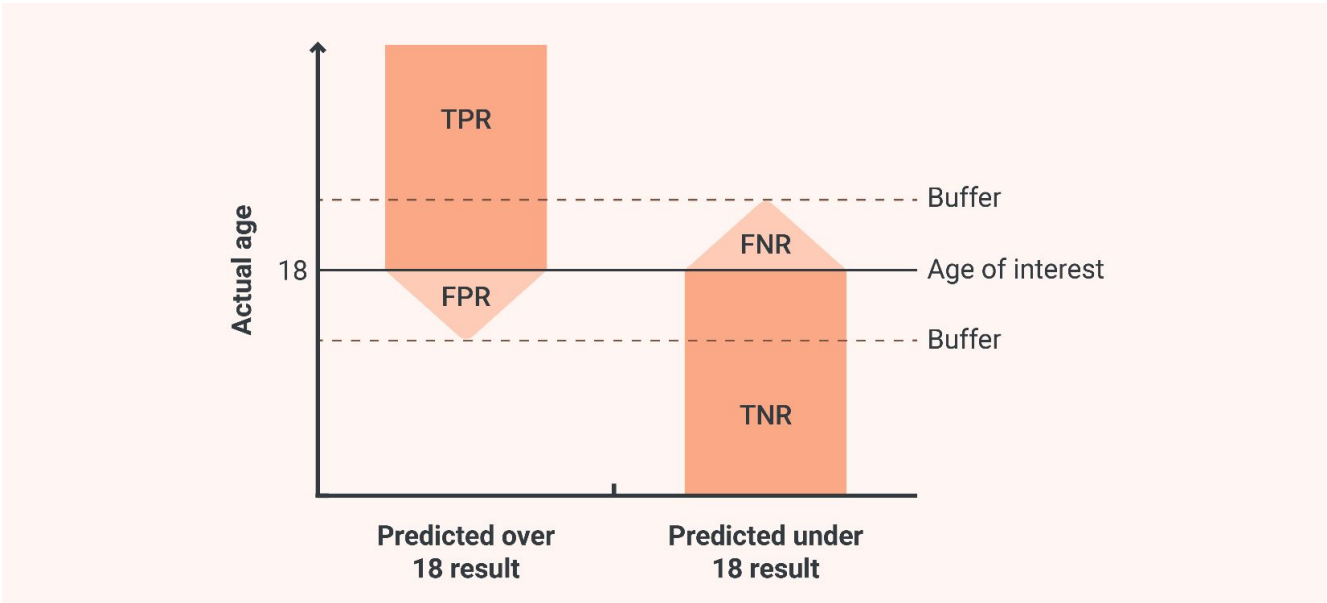
To manage the potential for errors, we recommend using facial age estimation as part of their compliance strategy. An example of this is the British Beer & Pub Association's 'Challenge 21' scheme, which is already widely adopted by publicans and their bar staff in England and Wales.[6] As it is difficult for human staff to be sure whether someone is over 18 just by looking at them, many stores adopt a policy to only require customers to prove their age if they appear to be under 21. Most supermarkets in England use a 'Challenge' / 'Think' 25 policy.

Facial age estimation can be configured to work with legal age thresholds in a similar way. Unlike human staff, facial age estimation's capacity for error is well quantified statistically. This makes it easier to choose a suitable buffer that is outside the technology's margin of error. The system can then be configured to estimate whether customers are above or below that threshold.

For a Challenge 25 scenario, if a customer is estimated to be below 25, they will then be directed into a user flow where they need to present documentary proof of their age. For example, a customer could use their Yoti app that is pre-verified to their passport, driving licence or national ID card, or in a retail setting, revert to an existing ID check by staff.

Since early 2019, we have reviewed the appropriate size of buffer for a number of use cases. We have come to the conclusion that this depends on a number of variables. The primary one is the demographic of users. The under-18 age group is the chief area of concern for regulators globally in terms of age-restricted goods and services. Given the improvements in accuracy of facial age estimation for this demographic, for the 13-25 age band we suggest a buffer of 3–5 years as an appropriate buffer for highly regulated sectors (e.g. adult content, gambling, alcohol, tobacco).

In some countries, more cautious regulators may initially look for a higher buffer. For a jurisdiction with legal age restriction of 18 and a threshold set to 28 (a 10 year buffer) we would have a 0.0% error rate for 13-17 year olds.[7] With a threshold set to 25 years, the current error rate is 0.1%. For a threshold of 21 years, the error rate is 0.7%.



6. See https://beerandpub.com/campaigns/challenge-21/
7. For more information see page 31

Buffers and the waterfall method of age assurance

But what happens for users who fall on the cusp on the age threshold or buffer? The “waterfall” approach is a way of offering alternative methods of age assurance when facial age estimation, usually with buffers, is the first check.

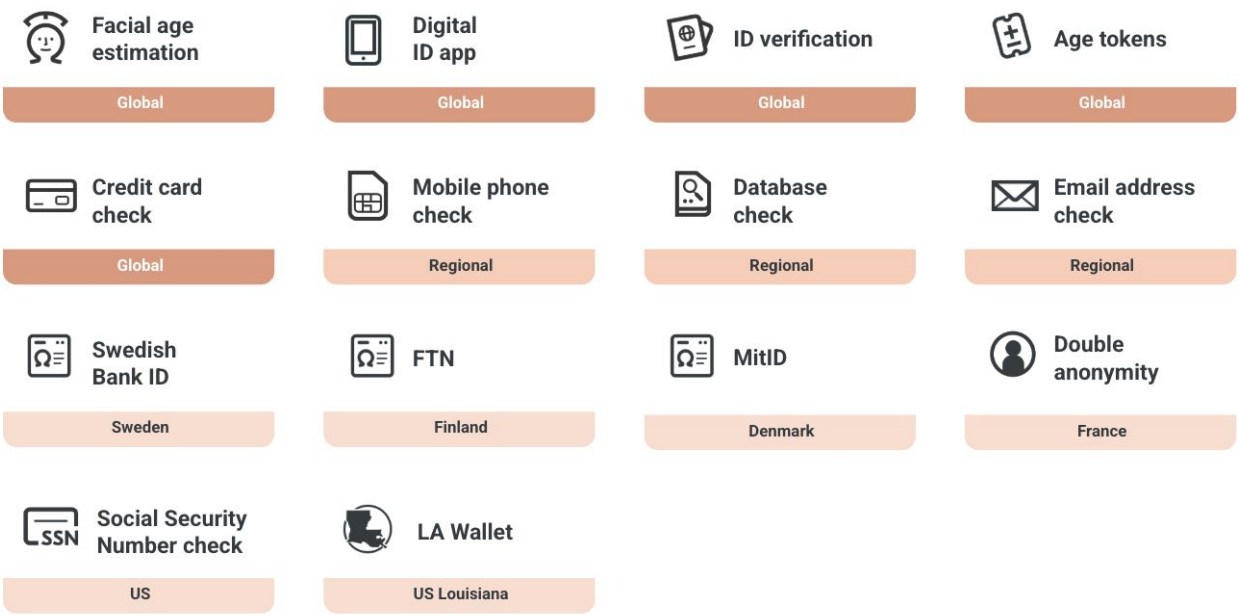
So if someone doesn’t pass the facial age estimation check, for example a 19 year old has been estimated as under 18 or a 17 year old as over 18. They can then be asked to prove their age another way, for example with Digital ID

For a user age 19 in an 18+ use case, for example, organisations are able to offer additional methods of age assurance so the user can still access their services.

This layered strategy enhances inclusivity, reduces user friction, and maintains regulatory compliance – without over-reliance on any single age checking method.

Whilst facial age estimation offers an accessible, data minimising way of accessing platforms, by adding in additional methods, platforms can still enable access to users on the margin who may not meet the platform’s requirements in terms of age estimation.

Yoti’s age check services include patent-protected technology that relates to a process of authorisation by combining detecting the human characteristics of a living person at a terminal which are then used to estimate age failing which there is a verification of an age related identity attribute such as a digital ID or an ID document. For further details of Yoti’s patents, please go to www.yoti.com/patents/.



Age assurance methods offered by Yoti

NIST evaluation

In September 2023, we submitted our facial age estimation model to the US National Institute of Standards and Technology (NIST). NIST assessed vendor Facial Age Estimation models using 4 data test sets at certain image sizes:

Data set	Image size
NIST Application - taken in US immigration offices	300 x 300
NIST Visa - collected in consular offices in Mexico	252 x 300
NIST Mugshot - collected in the US after arrest	480 x 600
NIST Border - collected at US border crossings	240 x 300

NIST note in their report that age estimation accuracy "will depend on the quality of the images" and the type of facial images captured.

Yoti have trained our model primarily on selfies of people looking into a mobile phone or laptop camera. This is because that is how most online organisations capture facial images to be age estimated. We capture facial images at an optimal resolution of 720 x 800 pixels, with the face closely cropped to maximise the facial detail. Furthermore, our face capture module automatically checks to ensure the images captured have acceptable light and facial pose to maximise potential success rates (TPR and TFR).

Our training and testing on mobile phone images with closely cropped faces are key reasons why Yoti's published FPRs (and MAEs) are lower (more accurate) for our Yoti model than the performance data published by NIST using their 4 different test data sets.

Yoti Facial Age Estimation Performance			
	Age Thresholds for 17 year olds		
	NIST Application	**NIST Mugshot**	**Yoti**
Image size	**300x300**	**480x600**	**720x800**
10% FPR	26	21	19
5% FPR	28	23	20
1% FPR	32	28	23
		MAE	
	NIST Application	**NIST Mugshot**	**Yoti**
Image size	**300x300**	**480x600**	**720x800**
18-30 Female	4.7	3.8	2.5
18-30 Male	3.2	2.4	2.0
31-50 Female	3.8	3.9	2.9
31-50 Male	3.2	3.4	2.6
		MAE	
	NIST Application	**NIST Mugshot**	**Yoti**
Image size	**300x300**	**480x600**	**720x800**
6-17	2.0	1.4	1.3
18-24	3.8	2.3	2.2

NIST evaluation comparison

NIST selected FPR objectives of 10%, 5% and 1% in their report as a way to benchmark their evaluation. As can be seen from the table above, NIST publish that Yoti's age estimation model is more accurate on higher image size 'Mugshot' faces than lower image size 'Application' faces. Consequently, the age thresholds required to meet FPRs of 10%, 5% and 1% are lower for Mugshot images than those needed using Application images.

Yoti is fortunate to have a very large set of anonymised facial images, verified to government issued age data, from Yoti app users. By separating out ~120,000 of these images as diverse test data across each year of age, from the many millions of images used to train our model, we have confidence in the accuracy figures we publish in our white papers.

As part of our document authenticity in our identity verification service we compare the age estimation result of the selfie with the real age from their document, which also helps us test the accuracy of the model.

NIST ethnicity approach and methodology

The minimisation and evaluation of bias is of critical importance to regulators, businesses and users. There have been numerous examples of new and AI technology that has been released and found to have bias by ethnicity, skin tone or gender.

Our approach at Yoti has been to tag our data by skin tone on a category of 1 (lightest) to 3 (darkest). We have reported on our bias since December 2018.

NIST's approach uses country metadata from their border image data set as a proxy for ethnicity.

NIST has chosen the following regions:

- **East Africa**: Ethiopia, Kenya, Somalia, Sudan, Tanzania
- **West Africa**: Nigeria, Liberia, Sierra Leone, Benin, Ghana, Mali, Senegal, Togo
- **East Europe**: Poland, Ukraine, Russia, Hungary, Romania, Czechia
- **East Asia**: Korea, China, Japan, Taiwan
- **South East Asia**: Cambodia, Indonesia, Malaysia, Thailand, Vietnam
- **South Asia**: Afghanistan, India, Myanmar, Nepal, Pakistan, Bangladesh

YOTI © 2025 Yoti Ltd

NIST evaluation October 2024

In October 2024, NIST published the results of their second evaluation. NIST placed our September 2024 facial age estimation (yoti-002) model in first place when measured for accuracy on 13-16 year olds, and second place when measured for accuracy on Mugshot images for 18-30 year olds.

Using its Visa test set for 13-16 year olds, NIST measured the average mean absolute error of the September 2024 Yoti model as 1.88.

This also affirms our ongoing work to improve the performance of our model. Within 6 months, our initial model (yoti-001) was evaluated by NIST with an MAE of 2.49 for 13-16 year olds. The model we are reporting on in this white paper (yoti--002 in the table below) has been evaluated by NIST as having an MAE for 13-16 year olds of 1.88.

Algorithm	Child Online Safety (13-16)			
	Age 13-16		Age 8-12	Age 17-22
	MAE	TPR	FPR	FPR
yoti-002	1.876(1)	0.569(1)	0.083(3)	0.01(13)
idemia-001	2.169(2)	0.495(2)	0.047(1)	0.009(12)
yoti-001	2.495(3)	0.477(3)	0.09(4)	0.006(11)
incode-000	2.582(4)	0.368(6)	0.054(2)	0.005(8)
nominder-000	3.03(5)	0.432(4)	0.194(12)	0.005(9)
unissey-001	3.488(6)	0.336(8)	0.216(13)	0.006(10)

ACCS evaluation October 2024

In October 2024, our September 2024 model was been evaluated by the Age Check Certification Scheme (ACCS). ACCS reported our MAE for 18 year olds as just 1.1 years, whilst Yoti's testing showed 1.2 years for that model. NIST reports the Yoti Sep 24 model MAE for 18 year olds as 2.63 years using visa images captured in Mexico consular offices.

It is important to note that ACCS testing is performed using images captured from a mobile device, which have good quality cameras. The images NIST use for testing can vary in terms of quality and source, so it is important to understand the NIST test images are not captured on mobile phones and the ACCS independent testing shows this can make a material difference to accuracy.

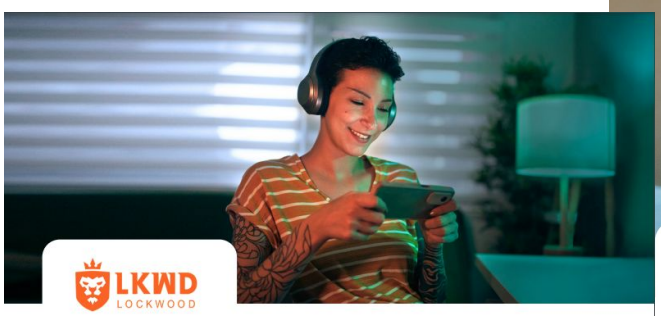
Growing adoption of facial age estimation

Yoti's facial age estimation technology is being used globally by some of the biggest brands across many different use cases.

Trust and safety: The technology is used on social media, e-commerce, streaming and adult platforms to verify that users are of the appropriate age.
Retail: Facial age estimation is a key part of Yoti's age assurance services, which are used globally to provide a secure and anonymous way for customers to prove their age using just a selfie, saving time, resource and potential conflict in retail environments.
Content moderation: On platforms that require age assurance, such as those involving user-generated content, facial age estimation is used to ensure that all individuals involved are of the correct age.
Child safety: Facial age estimation is helping to create age-appropriate experiences online, ensuring children can only access content, goods and services suitable for their age.

A growing number of regulators have approved facial age estimation, including in Germany, the UK, France and the US. This represents a significant global shift in online safety for a challenge once viewed as too difficult to solve.

Meta NCR VOYIX yubo OnlyFans KWS KIDS WEB SERVICES



LKWD LOCKWOOD

Company Lockwood Publishing

Industry Gaming

Solution Facial age estimation

Implementation API and SDK

How Lockwood Publishing increased player protection using Yoti's facial age estimation technology, ensuring a secure and enjoyable experience for players

Challenge



OnlyFans

Company OnlyFans

Industry Social media

Solution Facial age estimation

Implementation API and SDK

OnlyFans is the subscription social platform revolutionising Creator and Fan connections. OnlyFans empowers Creators from all genres to own their potential and gives them the opportunity to monetise the lawful content they

How OnlyFans became the first UK subscription-based platform to protect children and create age-appropriate experiences.

Challenge

As an 18s and over platform it is essential that minors are not able to view, upload or monetise media content on OnlyFans. It was also important to not create undue friction for users who are over 18, or disproportionality impact user privacy by asking them to share excessive amounts of personal data, just

Learnings from practical use

Facial age estimation works quickly, returning an age estimate in around 1 second. 99% of phone users submitting a face image are successfully age estimated. The user needs to present their face to the camera, uncovered (although glasses do not usually present a problem). We recognise that in some areas, internet speed can be challenging which is why we can cater for small image sizes of 50-100KB. We have scaled to handle tens of millions of checks per day, and we are currently able to handle up to 300 checks per second and can scale to do more if needed.

Dim lighting is not helpful; bright ambient light works best. Our research has found that beards and facial disfigurements can have a minor impact, but do not materially affect estimated ages. Following the COVID-19 pandemic, we have been researching how facial age estimation copes when a person is wearing a mask covering the lower half of the face. The results suggest that whilst accuracy is somewhat reduced, acceptable performance can usually still be achieved as long as an appropriate safety buffer is used.

ALL
Countries and territories

4
Scripts

19
Languages

99%
Success globally

Yoti facial age estimation is scalable to handle large volumes, our current performance is as follows but we can easily scale further

300	18K	1.1m	26m	780m
Checks a second	Checks a minute	Checks an hour	Checks a day	Checks a month

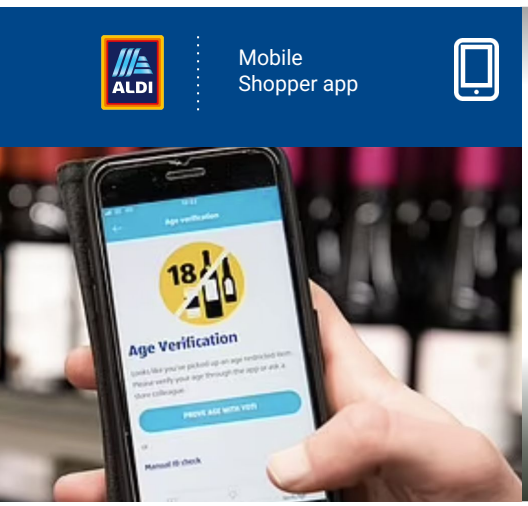
© 2025 Yoti Ltd

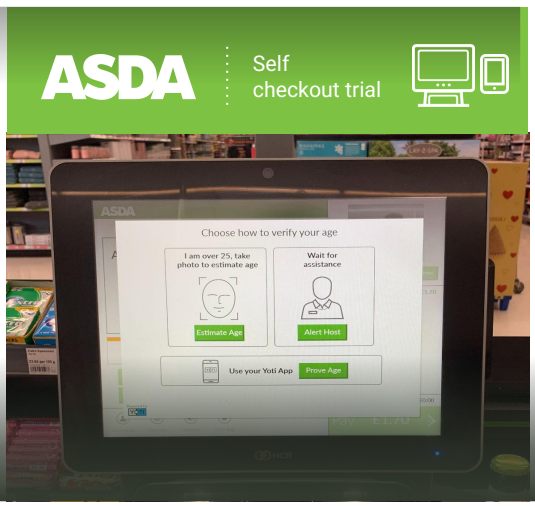
Live retail trials in UK

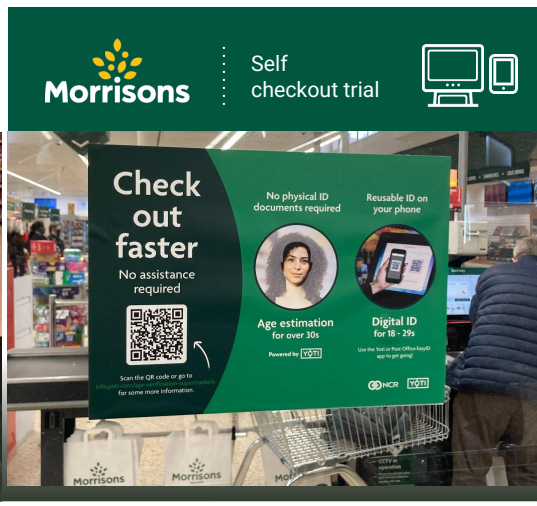
In 2022, UK supermarkets including Asda, Morrisons, Tesco and Co-op trialled our digital age assurance at self-checkouts in a scheme run by the Home Office. Key takeaways from the trials include:

- Participating supermarkets confirmed they support digital age assurance and would welcome legislative change in this area.
- There were no reported sales of underage customers purchasing age-restricted items when using our age assurance technology.
- Informed consent was gathered from all customers, who were given a choice whether to use the technology or present an ID document to a member of staff.
- The majority of shoppers who used Yoti digital proof of age solutions liked the technology and would use it again, once available.
- Digital age assurance technology provided an opportunity to reduce the number of physical age interventions, giving retail staff more time to monitor other activities, including spotting proxy sales.
- Yoti's facial age estimation is more accurate than humans which reduces the risks of incorrectly estimating the age of shoppers.
- Yoti's facial age estimation is more inclusive because anyone who looks over the required age threshold does not need to carry around a physical ID to prove their age.
- Digital age assurance supports the ability for retailers to achieve the Licensing Objectives

Supermarkets can now choose to do all age estimation and liveness checks on the terminal without requiring a connection.







Retail of Alcohol Standards Group report of Home Office trials

The UK Retail of Alcohol Standards Group (RASG) has published a detailed report on the use of Yoti's technology with four of the UK's largest supermarket groups in the Home Office Sandbox trials in 2022.[3] The RASG is an umbrella group of licenced retailers, with the mission to prevent underage drinking and promoting high standards among retailers. The report found the following:

Key findings

- Around 99,800 of shoppers used either facial age estimation or a digital ID app. No underage sales were identified.
- Independent test purchases by two retailers, using 18 to 19 year olds, demonstrated a 100% rejection rate using Yoti.
- All four supermarkets deemed the use of Yoti facial age estimation and digital ID proof of age successful adding that Yoti technology:
    - Was more accurate and consistent than humans in assessing whether a customer needed a Challenge 25 check.
    - Supports the licensing objective of protecting children from harm.
    - Has the potential to support the other licensing objectives.
    - Has the potential to reduce conflict that could be aimed at staff as the technology reduced the number of interactions for physical checks between staff & customers.
- Compliance with the age check approvals process and age assurance policies was higher when trialing digital age checks compared to physical age checks.
- After an initial increase in staff workload supporting customers with the new process, staff workload then reduced freeing up colleagues for other customer-facing activities.

The four supermarkets concluded:

- Digital proof of age and facial age estimation technology should be permitted for alcohol sales. It was more accurate and consistent compared to age verification undertaken by retail staff.
- Yoti's technology improved compliance rates, which if repeated across all stores, would reduce the chance of selling alcohol to minors.
- Lower sales of alcohol to minors will protect children from harm and consequently may reduce anti-social behaviour in the community.
- Yoti's 100% pass rate is higher than the pass rate in conventional age verification checks.
- The use of digital proof of age technology has the potential to reduce conflicts between customers and retail staff.

RASG
Retail of Alcohol Standards Group

3. https://rasg.org.uk/digital-proof-of-age-and-the-use-of-technology-for-alcohol-sales/

YOTI © 2025 Yoti Ltd

Supporting Children's Codes

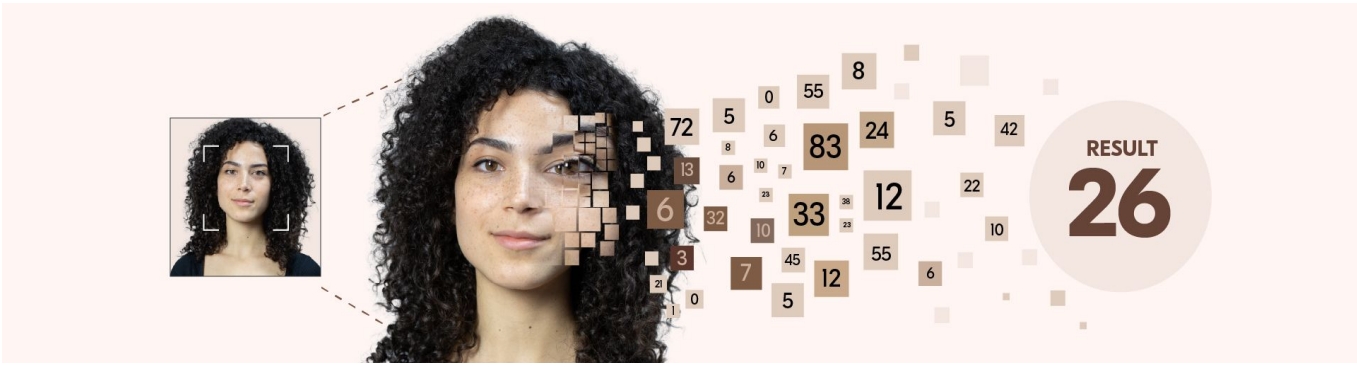
Given the growing importance of age checking online for younger children and teenagers, we have introduced additional training data to allow our algorithm to estimate the age of 6 to 12 year olds.

The UK's Age Appropriate Design Code is driving a global movement to design online interaction 'age appropriately' across the 4 C's - content, conduct, contact or contract.[4] The challenge for designers and platforms is to enable young people to be supported to thrive online whilst also enabling age appropriate interaction, protecting against detrimental content, protecting against grooming and supporting age appropriate content moderation. We can support platforms to recognise child users. Once they know who is a child they can treat them as such, including recommending high privacy settings by default, minimising data collection and having clear privacy information that children can understand. They can also limit features such as private chat functionality, reducing the chance of children speaking to unknown adults.

There are a growing number of countries and states around the world also reviewing legislation for a range of age-restricted goods and services, particularly age assurance for access online. Notably, the Californian Age-Appropriate Design Code Act is currently going through the US courts, but if it comes into effect there will be significant fines for non-compliance. There are also social media, gaming and adult content sites already using Yoti's facial age estimation to successfully prevent children from accessing their websites.

Obtaining consented data to develop our software to accurately estimate the ages of 6 to 12 year olds has been a significant challenge. We have worked hard to ethically obtain parental consent to use anonymous images of children in our training data. This consists of facial images with month and year of birth. For 6-12 year olds, our MAE results are 1.2 years, meaning it could be used effectively for triaging access to 13+ apps.

We will continue to invest more to improve our accuracy to make the internet safer for young people.



Detect face

A face is detected in an image and reduced to pixels. Each pixel is assigned a number that the AI can understand.

Compute numbers

The numbers are computed by a neural network that has been trained to recognise age by looking at millions of images of faces.

Determine age

The AI finds a pattern in the numbers and produces an age.

4. Livingstone, S. and Stoilova, M. "The 4 Cs, Classifying Online Risk to Children." SSOAR, 2021. https://doi.org/10.21241/ssoar.71817.

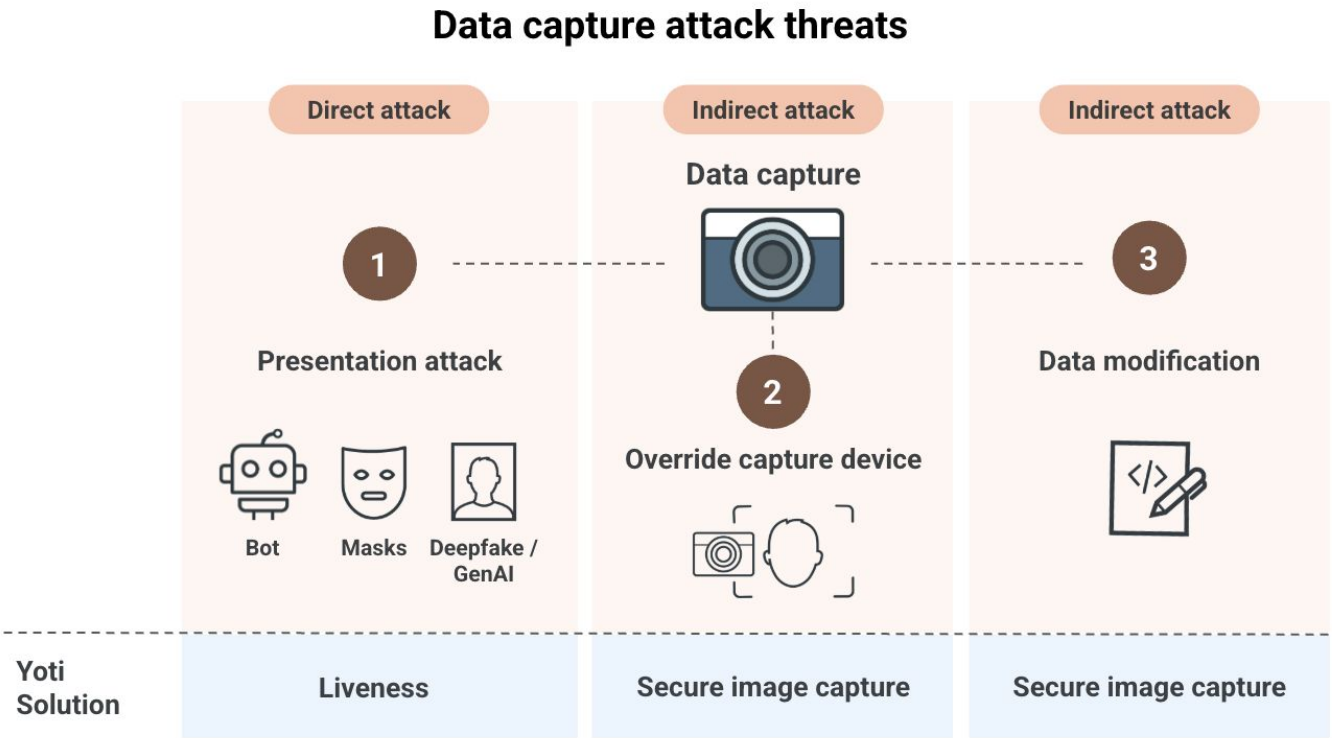
YOTI © 2025 Yoti Ltd

Tackling bots, deepfakes, and Generative AI

As discussed, by using safety buffers according to the age of interest, facial age estimation can be used with a very high level of confidence.

Equally important questions are how secure is the process? Can it be spoofed or are bad actors able to hack into the system to override checks, images or results? There's little point in estimating the age of a face accurately if it's not a real face in front of the camera.

This is why it is important to use a combination of technologies to offer a high level of assurance. There are a number of threat vectors illustrated in the diagram below.



Step 1 as a direct spoofing attack - an attempt to present an image, mask or video, often called a presentation attack. This is an attempt to spoof an age check by appearing older or pretending to be another person. To overcome this we use our NIST 2 certified liveness technology.[5] This ensures that the person undertaking the check is a real person and not someone wearing a mask, or presenting a picture or video of another (older or younger) person to the camera.

Steps 2 & 3 illustrate a newer, more sophisticated but relatively easy way for technically competent individuals to spoof the system. They are called injection attacks. An injection attack involves injecting an image or video designed to pass authentication, rather than the technology using the one captured on the device camera. Using free software and some limited technical ability, a bad actor is able to overwrite the image or video of the camera with pre-prepared images.

Yoti has developed a solution, called SICAP (Secure Image Capture and Processing) for which we have been granted a patent, that makes injection attacks much more difficult for imposters. It is a way of adding security at the point an image is being taken for a liveness check.[6]

5. Yoti MyFace Liveness white paper
6. How Yoti can help combat injection attacks

Legal compliance

There are understandable concerns about the potential unlawful use of personal and biometric data by governments and businesses.

Yoti's facial age estimation technology complies with both EU GDPR, and also our own ethical approach to user data and privacy. KJM (Kommission für Jugendmedienschutz), the German regulator, approved Yoti's facial age estimation for age assurance for online age-restricted content in November 2021. When clients use facial age estimation to verify the ages of their users, Yoti acts as the data processor and our clients are the data controllers. Our clients therefore need a legal basis to use facial age estimation according to their own jurisdiction.

The Yoti Age Portal has a built-in consent option so clients can easily collect consent if that is the lawful basis under which they operate.

Yoti's facial age estimation does not involve the processing of special category data - this has been confirmed by the UK Information Commissioner's Office. This is because the facial age estimation model is unable to allow or confirm the unique identification of a person as it has not been trained to so. Therefore, it is not being used for the purpose of identification which is the key test for special category data.

Put simply, if you put the same face into the model several times, the model would have no idea it is the same face (and no way of working that out) and would give slightly different age estimation results each time. The model is not trained to recognise any particular individual's face, but instead to categorise a presented face into an age.

Definition of special category data in Article 9 of the UK GDPR:

> *Processing of personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data,* ***biometric data for the purpose of uniquely identifying a natural person****, data concerning health or data concerning a natural person's sex life or sexual orientation*

Recital 51 of the UK GDPR further says that:

> *The processing of photographs should not systematically be considered to be processing of special categories of personal data as they are covered by the definition of biometric data only when processed through a specific technical means allowing the* ***unique identification or authentication of a natural person.***

For more information about why Yoti's facial age estimation does not process biometric data, please see our blog here.

Yoti's commitment to ethical use of AI technologies

At Yoti, we take our ethical responsibilities when developing new technology very seriously.

Our Data Protection Officer has completed a formal Privacy and Ethics Impact Assessment for Yoti's age-checking solutions, which is available on request to potential clients. It covers Yoti both as a data controller for our own use of age-checking solutions with our own users, and as a data processor when offering age-checking solutions to corporate customers.

We have set up an internal Ethics & Trust Committee with members from several different areas of our business, to consider ethical issues related to our technology and its uses. We used frameworks such as 'Responsible 100' and 'Consequences Scanning Model' as starting points for the scope of these considerations. Findings of the committee are shared with our senior management teams, Board of Directors and Guardian Council.

External scrutiny

We have obtained an ISAE 3000 assurance report from one of the top four global auditing firms, validating our age-checking services as in accordance with the British Standards Institution's PAS 1296 code of practice.[7]

The German Association for Voluntary Self-Regulation of Digital Media Service Providers (FSM) awarded us its Seal of Approval for our age assurance solutions.[8]

We have hosted seven roundtable sessions to get feedback from a range of industry practitioners on the unintended consequences of our approach.[9]

We also actively engage with organisations representing various minority groups to seek their views and input, including the UK transgender charity, Sparkle and LGBTQ non-profit Mosaic.

- Our accessibility goals
  Follow Web Content Accessibility Guidelines (WCAG) 2.2 Level A and Level AA for our Age Verification Service;
- Work with external companies for independent reviews and user testing of our solutions;
- Engage with people with a range of disabilities, understand their requirements and incorporate them into our scoping and planning.

You can read our accessibility statement here.

Certified B Corporation

FSM

7. PAS 1296: 2018 *Online age checking—Provision and use of online age check services—Code of Practice*. Available from the British Standards Institute shop.bsigroup.com.
8. https://www.fsm.de/de/fsm.de/yoti
9. https://www.yoti.com/blog/insights-from-our-fifth-regulatory-roundtable-exploring-age-assurance-methods/

YOTI © 2025 Yoti Ltd

Appendix

This appendix provides further detail on the current accuracy of facial age estimation technology. Taking confidence from the trends we've seen in past months (illustrated below), we expect these figures to continue to improve as the volume and diversity of our data set increases.

© 2025 Yoti Ltd

Mean Absolute Error by year

(Using test set of 118,403 images)

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	1.1	1.5	2.0	1.5	1.2	1.6	1.7	1.5	1.5
7	1.3	1.1	1.3	1.2	0.9	1.0	1.8	1.2	1.2
8	1.3	1.0	1.4	1.2	1.4	1.4	1.1	1.3	1.3
9	1.1	1.2	1.5	1.3	1.4	1.0	0.9	1.1	1.2
10	1.0	1.2	1.1	1.1	0.8	1.1	0.8	0.9	1.0
11	0.9	1.4	1.8	1.4	1.0	0.9	1.1	1.0	1.2
12	1.3	1.4	1.7	1.5	1.0	1.5	1.2	1.2	1.4
13	1.5	1.9	2.3	1.9	1.1	1.2	1.7	1.3	1.6
14	0.9	1.3	1.8	1.4	0.7	1.1	1.8	1.2	1.3
15	0.7	1.0	1.3	1.0	0.6	0.9	1.3	0.9	1.0
16	0.8	0.9	1.1	0.9	0.6	0.9	1.0	0.9	0.9
17	0.9	0.9	0.8	0.9	0.8	1.0	0.9	0.9	0.9
18	1.4	1.3	0.9	1.2	1.3	1.3	1.0	1.2	1.2
19	1.9	1.8	1.8	1.9	1.7	1.5	1.5	1.5	1.7
20	2.3	2.2	2.1	2.2	2.1	1.7	1.9	1.9	2.0
21	2.7	2.5	2.5	2.5	2.2	2.0	1.8	2.0	2.3
22	2.8	2.5	3.1	2.8	2.3	2.1	2.1	2.1	2.5
23	2.7	2.7	3.2	2.9	2.3	2.0	2.1	2.1	2.5
24	2.4	2.6	3.7	2.9	2.2	2.1	2.4	2.2	2.6
25	2.2	2.3	3.4	2.6	1.7	1.9	2.2	1.9	2.3
26	1.8	2.3	2.9	2.3	1.5	1.8	2.1	1.8	2.1
27	1.9	2.3	3.4	2.5	1.5	1.8	2.0	1.8	2.2
28	2.1	2.3	3.2	2.5	1.5	1.9	2.3	1.9	2.2
29	2.0	2.2	3.4	2.5	1.7	1.9	2.4	2.0	2.3
30	2.1	2.4	2.6	2.4	1.8	2.1	2.5	2.1	2.2

YOTI © 2025 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	2.3	2.5	3.9	2.9	1.6	2.1	2.0	1.9	2.4
32	2.3	2.4	3.7	2.8	1.9	2.4	2.5	2.3	2.5
33	2.5	2.7	3.4	2.9	2.2	2.7	2.5	2.4	2.7
34	2.4	2.8	3.3	2.9	2.1	2.5	2.6	2.4	2.6
35	2.5	2.4	3.9	3.0	2.1	2.4	2.4	2.3	2.6
36	2.0	2.4	4.0	2.8	2.3	2.3	2.5	2.4	2.6
37	2.7	3.3	3.0	3.0	2.1	2.5	2.8	2.5	2.7
38	2.0	2.3	3.0	2.4	1.8	2.7	2.4	2.3	2.4
39	2.3	2.8	2.7	2.6	2.1	2.1	3.2	2.5	2.5
40	2.2	2.3	2.9	2.4	2.2	2.1	2.1	2.1	2.3
41	2.1	2.2	3.1	2.5	2.0	2.1	2.3	2.1	2.3
42	2.3	2.3	2.7	2.5	2.1	2.3	2.6	2.3	2.4
43	2.3	2.4	3.6	2.7	2.3	2.3	2.9	2.5	2.6
44	2.2	2.5	2.9	2.5	2.1	2.3	2.6	2.3	2.4
45	2.3	2.4	3.2	2.6	2.1	2.4	2.5	2.3	2.5
46	2.6	2.9	2.8	2.8	2.5	2.6	2.7	2.6	2.7
47	2.6	2.3	3.2	2.7	2.2	2.5	2.9	2.6	2.6
48	2.6	2.8	2.4	2.6	2.4	2.6	2.7	2.5	2.6
49	2.6	2.5	3.5	2.9	2.6	2.4	2.9	2.6	2.7
50	2.6	2.3	3.3	2.7	2.3	2.6	3.1	2.7	2.7
51	2.7	2.2	4.6	3.1	2.5	2.5	3.4	2.8	3.0
52	2.3	2.1	3.3	2.6	2.6	2.3	4.5	3.1	2.9
53	2.5	2.3	3.6	2.8	2.6	2.8	3.5	3.0	2.9
54	2.6	2.4	3.1	2.7	2.7	3.6	3.3	3.2	2.9
55	2.5	3.3	3.3	3.0	2.5	3.6	3.5	3.2	3.1

YOTI © 2025 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	2.3	3.2	3.1	2.9	2.4	3.3	3.2	3.0	2.9
57	2.6	2.4	3.6	2.9	2.6	2.9	3.3	2.9	2.9
58	2.3	2.6	2.8	2.6	2.2	3.2	2.7	2.7	2.6
59	2.1	3.6	2.9	2.9	2.6	3.3	2.9	2.9	2.9
60	2.1	2.9	2.9	2.6	2.3	2.7	4.3	3.1	2.9
61	2.2	3.3	2.2	2.6	2.5	2.8	4.3	3.2	2.9
62	1.8	3.4	3.0	2.7	2.3	2.7	3.6	2.9	2.8
63	2.3	3.2	1.8	2.4	2.4	3.0	4.3	3.3	2.8
64	2.5	2.8	3.7	3.0	2.5	2.7	3.8	3.0	3.0
65	2.6	2.7	3.1	2.8	2.5	2.5	3.2	2.7	2.8
66	2.7	3.1	4.9	3.5	2.6	3.6	3.0	3.0	3.3
67	3.1	2.7	3.0	2.9	3.0	3.9	3.8	3.6	3.3
68	3.0	3.1	6.2	4.1	3.6	3.5	3.6	3.6	3.8
69	3.9	3.9		3.9	3.2	3.3	3.0	3.2	3.5
70	3.1	4.7	0.5	2.8	2.9	3.2	8.2	4.8	3.8
Avg	**2.2**	**2.4**	**2.8**	**2.5**	**2.0**	**2.3**	**2.6**	**2.3**	**2.4**

YOTI © 2025 Yoti Ltd

Data used to build the model ('training data')

Since early 2015 Yoti has invested significantly in building a leading R&D team who work on a variety of AI initiatives.

The current production model of facial age estimation (July 2025) was built using a training data set taken mainly from users of the Yoti apps. We provide information to users at the onboarding stage about our use of biometrics with links to further information. This includes the Privacy Notice where the use of user data by our R&D team for internal research is extensively detailed.[10]

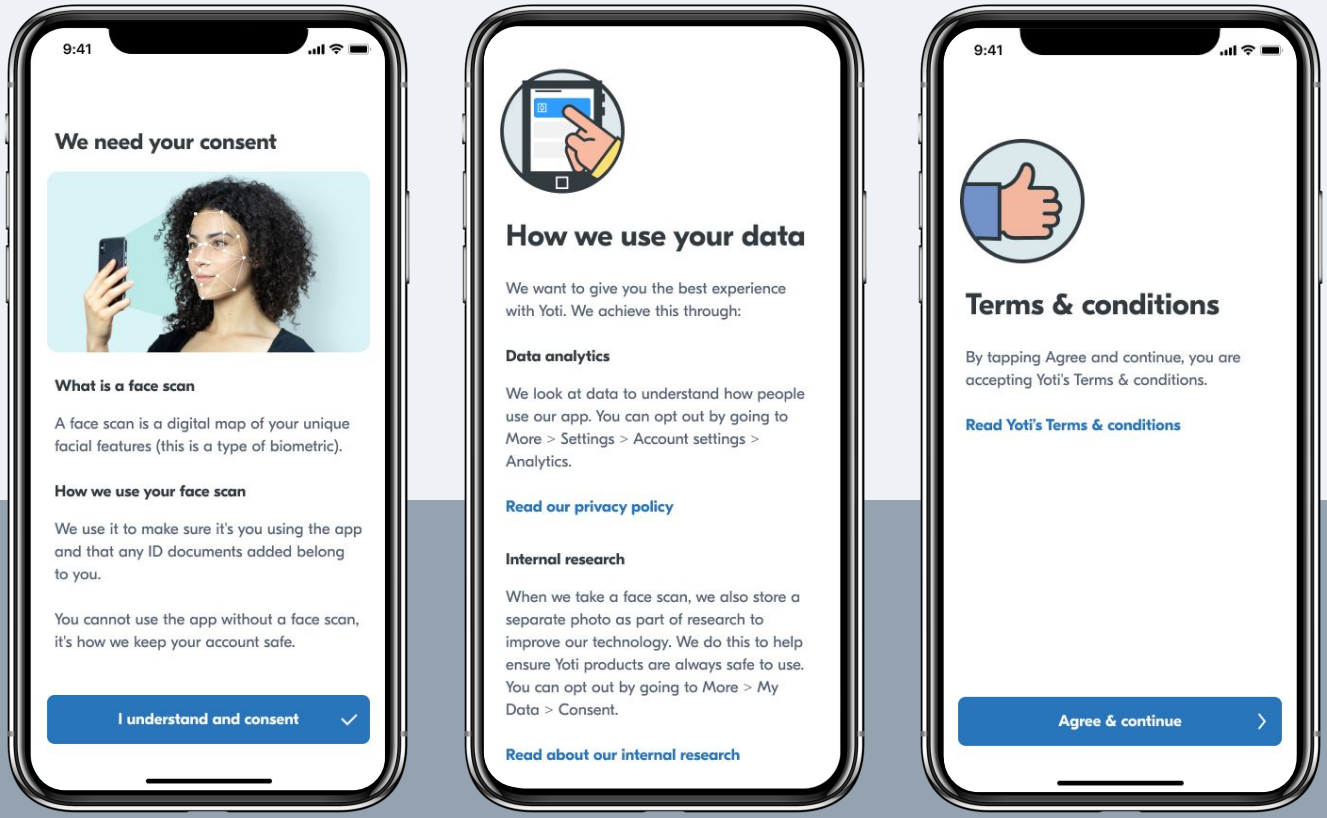
Any user can go to the app settings at any time and opt out of R&D use of their data. This prevents further data from that user being sent to R&D. It deletes all the data associated with that user that is on the R&D server and available for R&D to use. We have chosen to automatically delete the existing data when a user opts out or deletes their account, even though this is not a legal requirement under the research provision in GDPR article 17(3)(d).[11] We employ a privacy-by-design approach. This means that although we can find the data of a specific user to action the data deletion, there is no way to recreate a specific user's identity from that data.

To enhance our coverage of particular demographics, Yoti has previously gathered further age-verified images with consent in Nairobi, Kenya. Through the Share2Protect campaign, parents and children were able to support the extension of facial age estimation to 6-13 year olds.[12] We have also purchased further parent consented child facial images, with month and year of birth, and we undertook thorough due diligence on all our data sources.

In 2021, Yoti was part of the ICO Sandbox to extend our facial age estimation AI programme to those under 13 who don't have ID documents.[13] Ending in May 2022 the ICO has since published their exit report. Our participation helped the ICO gain insights into age assurance for young people . It also helped them to understand that facial age estimation does not process special category data. The ICO have updated their 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

On-boarding and R&D opt-out screens in the Yoti app



We provide information to users at onboarding about our use of biometrics. This includes links to further information, including the full privacy notice, where the use of user data for R&D is extensively detailed. Users can opt out of their data being used for R&D by Yoti at any time, via the settings on the app.

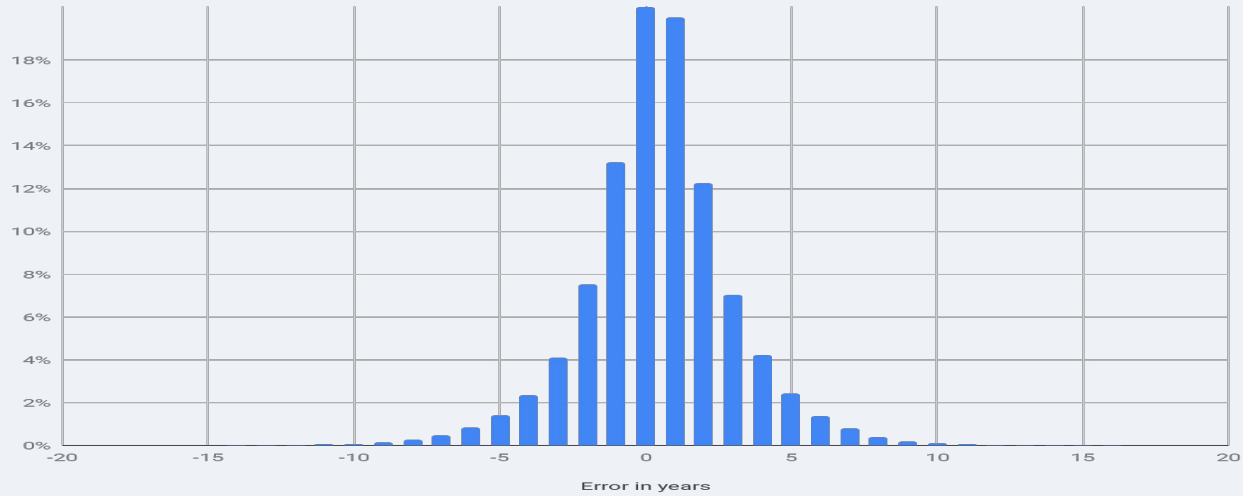
© 2025 Yoti Ltd

Data used for testing

Our testing data is also taken from Yoti users worldwide in the same manner as the training data. We strive to ensure that it represents as broad a demographic as possible, considering age, gender and skin tone. This gives us confidence that the results presented in this white paper are reproducible in a wide variety of real world situations.

Accuracy across the entire dataset

In our most recent testing of the model (July 2025), we used test data that was comprised of hundreds of thousands of images with a verified age. The MAE across all years is now 2.4 years; for females it is 2.0, for males it is 1.9. This reflects a higher number of males in the training data across most years. The range of errors tends towards a normal distribution, with a standard deviation of 1.7. The standard deviation is a measure of the variance of the data around the mean. This is illustrated in the chart below.



Accuracy by age, gender and skin tone

In our testing set hundreds of thousands of facial images with verified age were tagged with the subject's gender and skin tone. The gender was taken from the subject's uploaded identity document, and for skin tone, our research team tagged the images using a a three point numeric scale from 1 to 3, where 1 is the lightest and 3 is the darkest.

The majority of the tagging was performed using a manual process, with some data tagged automatically. We have put quality procedures in place to help ensure our manual tagging is reliable and free from 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.

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	1.0	1.0	1.1	1.0	0.9	1.0	1.1	1.0	1.0
10-12	1.0	1.1	1.6	1.2	0.8	1.1	1.3	1.0	1.1
13-15	0.9	1.0	1.2	1.1	0.6	1.0	1.0	0.9	1.0
16-17	0.9	0.9	1.0	0.9	0.7	1.1	1.0	0.9	0.9
18-24	1.5	1.6	1.7	1.6	1.3	1.5	1.5	1.4	1.5
25-29	1.9	2.1	2.6	2.2	1.5	1.6	1.9	1.7	2.0
30-39	2.0	2.3	3.0	2.4	1.7	2.0	2.1	1.9	2.2
40-49	2.0	2.3	2.7	2.4	1.9	2.1	2.2	2.0	2.2
50-59	2.0	2.6	3.2	2.6	2.1	3.0	2.8	2.6	2.6
60-69	2.1	2.5	2.4	2.3	2.1	2.6	3.0	2.6	2.5
All	1.7	2.0	2.3	2.0	1.6	1.9	2.1	1.9	1.9

Mean Absolute Error and Absolute Error Standard Deviation

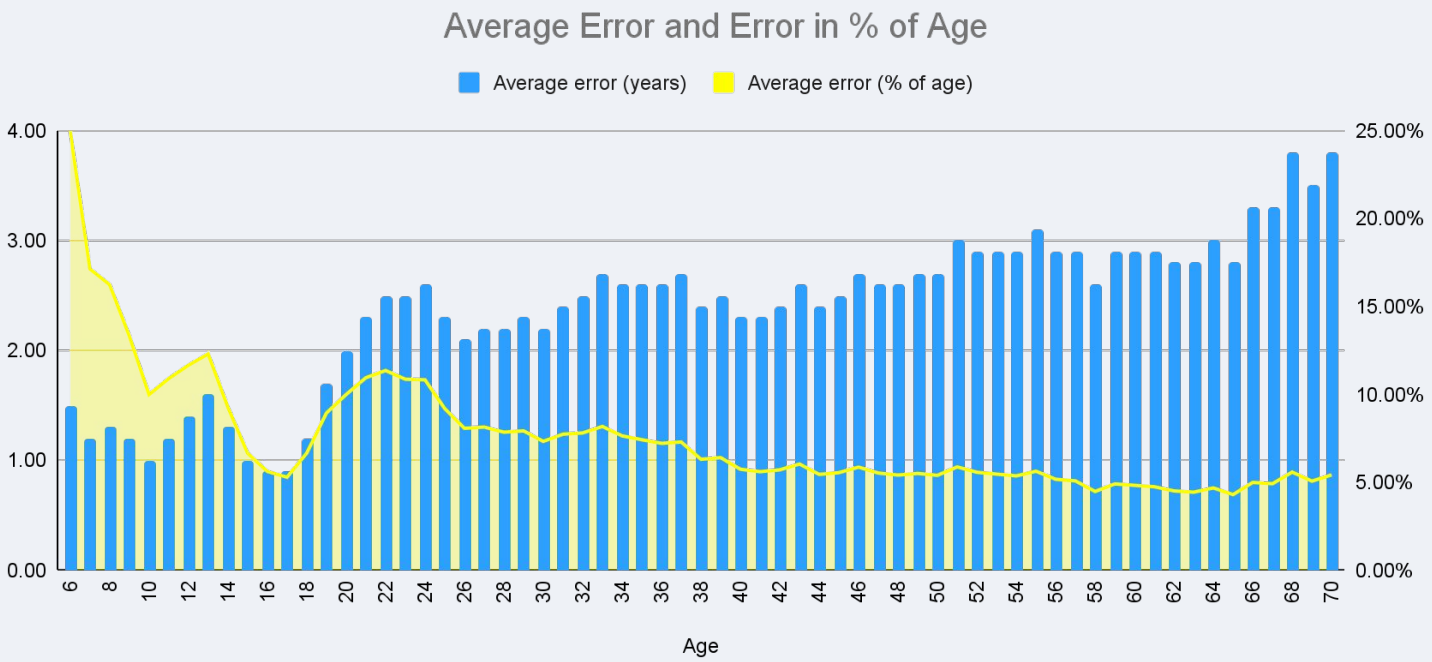
Standard deviation can add qualification to MAE by saying, for example, that we have a low deviation, i.e. the error rate is consistent.

A higher standard deviation tells us that the errors are spread over a bigger range. A lower standard deviation indicates that errors across the data tend to be of a similar range (or more standard).

YOTI © 2025 Yoti Ltd

Absolute versus percentage errors across age and skin tone

It is worth noting that although the magnitude of error may appear larger for older age bands, when this is as a percentage of the individual's age, it often is more accurate in relative terms. For instance, an error of 2 years for a 15 year old is a 13% error, whereas an error of 2 years for a 50 year old is an error of 4%. This is illustrated in the chart below.



Improvement in accuracy as the training data set grows

The differing mean absolute error shown for different groups (age, gender, skin tone) correlates strongly with how well-represented those groups are in the training data set. We are consistently retraining our facial age estimation model on an ever-expanding data set of millions of face images.

The charts below illustrate the significant accuracy improvements that we have observed over time. The size and composition of our test data has itself diversified over this period too, so the comparisons from one model's results to the next are not absolute. However the overall trend is clear and encouraging. Where appropriate, we will endeavour to undertake further targeted fieldwork in this area.

N.B. From September 2021, we have revised our approach to concentrate on achieving a reduction on bias, even where this may have a detrimental effect on accuracy.

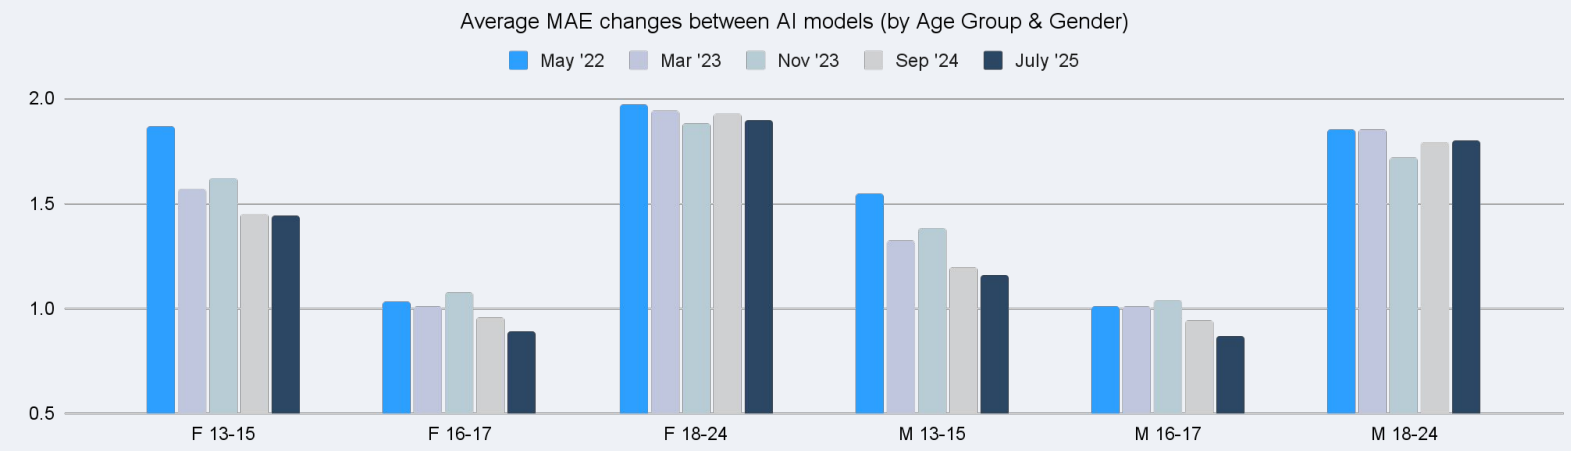
Since our May 2022 update, we have removed some older images from both our training and testing data sets. This is in line with our privacy policy on customer data retention, which states that if a user has been inactive for over 3 years, we delete their data. This has two notable implications:

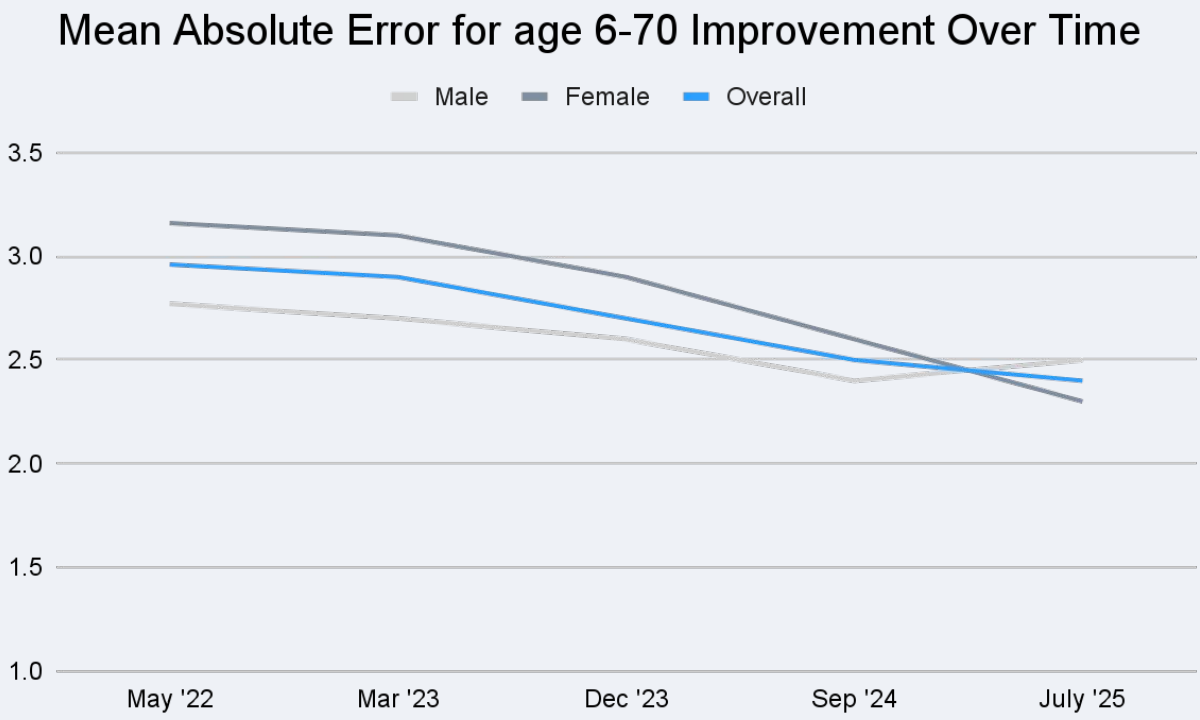
- **Training data** - where deleted data may have a skewed number of images in a certain subcategory, the accuracy in that data range may be affected.
- **Testing data** - changes in this data set mean that results over time are not strictly 100% comparable, as each model is not being tested against exactly the same set of test data.

We do not believe the change in training or testing data will materially affect the accuracy of our model over time. We will also monitor churn of our data sets to ensure that we replace data with the corresponding demographic that may have any significant effect on our accuracy or testing.

YOTI © 2025 Yoti Ltd

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





YOTI © 2025 Yoti Ltd

False positive rates for a selection of thresholds for an age of interest of 18

						Average False Positive Rate (weighted equally for each age)
		14	15	16	17	
Test Sample Size		*3,172*	*7,392*	*10,258*	*10,417*	
Thresholds (years)	20	0.16%	0.50%	1.35%	3.76%	**1.44%**
	21	0.03%	0.31%	0.67%	1.96%	**0.74%**
	22	0.03%	0.18%	0.40%	1.10%	**0.43%**
	23	0.00%	0.11%	0.18%	0.73%	**0.25%**
	24	0.00%	0.07%	0.12%	0.36%	**0.13%**
	25	0.00%	0.05%	0.07%	0.13%	**0.06%**
	26	0.00%	0.04%	0.03%	0.05%	**0.03%**
	27	0.00%	0.04%	0.02%	0.02%	**0.02%**
	28	0.00%	0.04%	0.00%	0.01%	**0.01%**
	29	0.00%	0.01%	0.00%	0.00%	**0.00%**
	30	0.00%	0.01%	0.00%	0.00%	**0.00%**

YOTI © 2025 Yoti Ltd

False positive rates for a selection of 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**	0.12%	0.36%	0.90%	3.23%	8.73%	**2.67%**
	25	0.07%	0.13%	0.38%	1.55%	4.70%	**1.36%**
	26	0.03%	0.05%	0.17%	0.63%	2.02%	**0.58%**
	27	0.02%	0.02%	0.11%	0.27%	0.86%	**0.26%**
	28	0.00%	0.01%	0.06%	0.13%	0.34%	**0.11%**
	29	0.00%	0.00%	0.03%	0.10%	0.20%	**0.07%**
	30	0.00%	0.00%	0.02%	0.10%	0.07%	**0.04%**
	31	0.00%	0.00%	0.02%	0.02%	0.02%	**0.01%**
	32	0.00%	0.00%	0.02%	0.02%	0.02%	**0.01%**
	33	0.00%	0.00%	0.01%	0.00%	0.02%	**0.01%**
	34	0.00%	0.00%	0.01%	0.00%	0.00%	**0.00%**
	35	0.00%	0.00%	0.01%	0.00%	0.00%	**0.00%**
	36	0.00%	0.00%	0.01%	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 © 2025 Yoti Ltd

Memberships, associations and accreditations

Reviewed by

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

© 2025 Yoti Ltd

Discussion paper

Where in the tech stack should age assurance sit and how should it be done?

January 2025



Author
Julie Dawson,
Chief Regulatory and Policy Officer at Yoti



Where in the tech stack should age assurance check(s) occur?

Should they be on device on a software as a service (SAAS) basis, at operating system level or at more than one level in the stack?

How and where should the age assurance be retained?

Age assurance technologies have evolved significantly over the past decade, with approximately 50 organisations developing tools ranging from age verification to age estimation and inference. Some of these tools are now independently audited and comply with global standards. There's also benchmarks like the National Institute for Standards & Technology (NIST) evaluation of AI facial age estimation. These approaches are designed to minimise consumer burden and be simple for organisations to integrate, within a matter of hours.

In this discussion paper, we explore where age assurance checks should sit in the consumer journey and in the tech stack; if it should be a **one-off check** at set up or **both** at set up **and** at the **point of accessing** a service, for added safeguarding. We will explain how tokenisation can lower the cost of repeat checks. Another approach described is a **layered approach** to age assurance combining:

- Age verification during device setup.
- Periodic reauthentication at the device level.
- Age checks at service access points.

However, we need to reflect, with the approach selected, does this really create more burden or less on parents and does this safeguard young people and deliver age appropriate access and experiences?

Yoti's view, weighing the factors discussed in this paper, is that the optimal placement of age assurance within the tech stack and consumer journey depends on

- **Balancing consumer convenience**
- **Privacy**
- **Operational feasibility**

1. https://accscheme.com/registry/
2. https://avpassociation.com/standards-for-age-verification/

Principle of proximity

We advocate that to enhance consumer understanding and acceptance, age checks should occur as close as possible in time to the relevant activity or service which requires age appropriate access. For example, at the very point in time when signing up for a gambling, adult content or dating site platform it is logical that the check is at the direct point of access. This ensures underage people can't access these platforms and content. In the case of a mixed age service such as gaming; this allows a differentiated service to be offered to minors. This principle of context-specific implementation is designed to make it clear to users why they are undertaking the check.

Interoperability and reuse - tokenisation

It stands to reason that once you're clearly over 13, 18 or 21 - you're not getting any younger. So once you've passed an age threshold, you don't want to have to keep repeatedly proving that you have reached that age threshold.

The flip side for regulators and platforms is that they want to be sure that it really is the right person who is accessing a service and not a proxy access, that may have been paid for, to bypass the age check.

To meet this issue and to reduce costs across industries, age verification providers have established interoperable, tokenized approaches. Users can verify their age once and gain continued access to a number of platforms for a certain time period, as approved by regulators, according to the use case and the risk profile. For example, an 18+ token from an adult content site can be reused across other services like dating or gaming platforms, for a certain period.

Binding successful age verification to passkeys saved on device is another way to strengthen repeat authentication.

Parental role and challenges

Many platforms have invested heavily in devising elegant and thoughtful parental controls, no doubt with good intentions. However, despite this, most parents struggle with setup and management of parental controls. This may be due to children pestering them, peer pressure, children outwitting them, or the technical complexity and time required to set up the controls across multiple devices and services.
Industry data shows minimal parental control adoption—for instance, only 1% of U.S. parents use Snapchat and Discord parental control tools.[3] This reality underscores the need for system-level solutions; rather than solely relying on a parent in the loop to set the age of their child within the set up of every device.

---

3. U.S. Senate Judiciary Committee hearing in January https://www.techpolicy.press/transcript-us-senate-judiciary-committee-hearing-on-big-tech-and-the-online-child-sexual-exploitation-crisis/https://www.nbcnews.com/tech/social-media/fewer-1-parents-use-social-media-tools-monitor-childrens-accounts-tech-rcna145592

Proportionality of burden or friction on adults

One of the key considerations in the US has been: is the age assurance journey overly burdensome to an adult? It is worth considering the user experience of unboxing a new phone. Today, when you unbox your new device it is relatively straightforward to set up your new phone, though still some people struggle. And data transfer can take a few minutes to a few hours, depending on the storage on the old and the new phone. However, adding additional friction of undertaking an age check at the moment of unboxing may be unpopular.

In this future world of operating system age checks, you would in addition be asked to go through an age verification, age estimation or age inference check, of your choice. Depending on the options available to an individual, that could take from a few seconds to minutes to a couple of hours.

But this highlights some questions, including:

- How could that be made mandatory across all devices globally?
- Would the device be locked down until an age check had been undertaken?
- Whilst parents may set up devices for young children; teens may unbox their own new devices if they've already had one for several years..
- Given that there are hand-me-down devices in families; would a subsequent re-authentication be required to ensure that the person using the device is actually the same as the one who undertook the initial age check?
- How often should re-authentication be required?
- How would this be relied on with open-source operating systems? (For instance Linux is open source, so it is conceivable that any operating system level age verification can be disabled or modified by the user of the system to generate false results)
- Who is responsible for preventing the age verification information being used by predators to target minors?

There are now calls from certain industries to consider undertaking checks at other levels in the tech stack, including device-level age checks. In the next section we try to outline some of the practicalities as well as the positive and negative, intended and unintended consequences of these suggestions.



System-level age checks: pros and cons

By system-level age checks, we mean Operating System (OS) and App Store Integrations.

Pros

- Streamlined process and reduced burden on smaller organisations.
- Potential to provide a level playing field in terms of regulation.

There are suggestions that age checks at Operating System or App Store level could streamline the process and reduce the burden on certain sectors, perhaps for the long tail of smaller organisations. This has appealed as a concept to some large organisations who want a 'level playing field' in terms of regulation, so it is not just the largest players who are required to undertake age checks. The main proponents of this suggestion are adult content and social media organisations. There is little detail given however, as to how this would happen in practice.

Implementing age assurance at the operating system or app store level raises some practical and ethical questions, which we outline below.

Cons

- Sector exclusion: certain goods and services (e.g., alcohol, nicotine) are not present in app stores.
- Burden on parents to set age levels for every device.
- Device reuse and resale introduce complexities in maintaining accurate age assignments. (For instance hotel rooms, holiday apartments with smart devices - would these be locked by the hotel management and require adults to request family controls be temporarily removed?)
- Device sharing across family members i.e. iPad used by both parents and children.
- Risks of data exploitation and breaches.
- Cost of delivering the service and concerns about corporate dominance.
- Increased complexity for app stores, requiring partnerships with third-party providers.
- Privacy and data risks:
    - Centralised age verification at operating system level may collect sensitive data, increasing the risk of large scale data breaches.
    - Sending age signals from app stores to apps raises tracking and privacy concerns.
    - Anonymous access to apps may be undermined by mandatory age verification.
- Where low cost devices are not supported, in terms of operating systems with age verification enabled capabilities, the device owners may be priced out of content access.
- Legacy devices remain in circulation for many years

- **Sector exclusion:** Adult content is absent from app stores and age-restricted goods, such as alcohol, nicotine and vaping, are purchased online or in-person. How would operating system or app store level checks work across these goods and services?

- **Burden on parents:** Is this really coded language for 'the parent has to set the age of every device in the household'? If only a small percentage of parents are managing to set up parental controls on the most popular social media platforms, can we expect parents to manage resetting the operating system age levels across multiple devices in the home? How is a child supported if their parent does not engage with setting the age at an appropriate level?

- **Device reuse:** How can a platform or a network operator check that age is reassigned when devices are handed down in the family or resold?

- **Device estate**: How to deal with the existing estate of billions of legacy devices, how could they be required to incorporate operating system level age checks?

- **Data exploitation risks:** The data in question - identity or age - also potentially has value for the organisations that access it. Centralised age and identity data might tempt companies or bad actors within companies to turn a blind eye to data minimisation requirements and harvest or monetise personal information despite risking privacy or security violations.

  One of the attractions of independent age verification providers has been that their contracts are predicated on not reselling the age or identity data to third parties, or using that data to mine or build personalised products and services.

- **Data centralisation risk:** Age verification at operating system level may include verification, estimation or inference approaches and may lead to the centralised collection of sensitive personal data, such as birth dates or ID's, which these platforms may then also store, increasing the risk of data breaches, unauthorised data mining and misuse.

- **Managing redress:** How to manage redress, when people are given the incorrect age? Additionally, how do you periodically check that the right person is still using the device? What ongoing re authentication is deemed proportionate by data protection regulators?

- **Cost of delivering the service:** Who should pay for the operating system level checks? Should those offering operating systems be forced to provide them pro bono - is that fair? How should this service be charged for - the initial assessment of age, ongoing customer support, and any reauthentication? The balance between service quality, sustainability and service monetisation must be carefully managed. However, the app stores or operating system organisations may offer a trade - we will provide this age assurance service for free or at this price; however the quid pro quo, is that we wish to mine and reuse this valuable data. Would this be deemed fair by regulators and the public?

- **Market and regulatory considerations:** Competition authorities around the world keep a close eye on platform monopolies. If age assurance is done at operating system level, this could deepen concerns over corporate platform dominance. It is fair to say that much of the innovation, in terms of age assurance, has not stemmed from the large platforms or from the operating systems level organisations. Will continued innovation in age assurance flourish if a small number of operating systems undertake the bulk of checks globally? How will regulators audit the quality of their work? Would operating systems be required to meet the incoming international standards for age assurance? Would dominant global platforms be allowed to charge whatever fees they like for age checks or would pricing have to be regulated? Would age be a Trojan horse for dominant platforms to then dominate the identity market and would their identity fees need to be controlled by competition regulators?
  In sectors such as online gaming, there has been an uneasy relationship between the app stores and gaming sector; where the cost of access / fees charged by app stores has been a bone of contention.
  There has also been concern over the prospect of supra national identity systems, not owned by national governments, but by global corporations. Identity data is closely linked to payments data; so access to that data, would provide a strong competitive advantage to any organisation able to leverage that.

- **Cost and operational burden:** Implementation requires significant financial investment, affecting both app developers and users. It is not just the technical solution that would require investment to develop and maintain. There will also be a requirement for customer service, support for app developers who need to confirm age ranges, and associated legal and compliance costs, transparency audits and requests for information when either data protection or content regulators conduct periodic audits. There is a liability risk; the operating systems would need to consider when undertaking their pricing, that they could face reputational and legal risks if their systems fail, whether by incorrectly verifying ages or being exploited to bypass safeguards.
  Yet, there is no doubt that providing the service of ascertaining age requires investment to provide a secure and robust check; in potential rechecks, in customer support for dispute resolution with consumers, and subsequent audit by regulators. So at whichever level in the stack an age check is undertaken, it will not happen without a sustainable business model.

Shift of accountability

Accountability may shift in other directions for instance to parents if they have to setup or remove restrictions.
Shifting age assurance responsibility to app stores may reduce accountability for developers, who are better positioned to implement tailored systems.
Do relying parties for example hotel or leisure facility management become accountable for children using their facilities to access services such as adult content (pornography)?
Likewise, does any company / work environment become accountable for a child accessing adult content if the child is using company equipment that could or should be locked and age verified?



Device-level age checks

Device-level checks can address specific contexts but face practical and ethical challenges:

- Trust in device security and assurance against data leakage.
- Risks of unauthorised image capture and data sharing.
- Challenges in maintaining privacy and ensuring robust re authentication mechanisms.
- Device reuse poses significant challenges in family and resale contexts.
- Reliance on device manufacturers for ongoing compliance and auditing of practices.

Device-level versus user-level challenges

- Age settings applied at the device level may not suit multi-user households, leading to excessive restrictions or unintended access.
- Individual account-based verification is more flexible but may require more complex implementations.

First of all, it's worth being clear as to what is meant by 'device level'. Device level checks could be construed to be solely consumer devices - handsets, games consoles, VR (virtual reality) headsets etc. Or device level checks could also be considered to be on any hardware device including in store hardware - such as gambling terminals or electronic point of sale cash points in store.

Many age assurance approaches already use encryption, threat monitoring and minimal data retention, ensuring a high level of security. Given this, requiring additional on-device processing for age checks may be redundant. While on-device checks might be suitable for specific contexts, like retail terminals, broader implementations should leverage secure cloud-based verification to avoid unnecessary device-level complexity and maintain privacy safeguards.

In the case of device-level age checks, there are further complexities:

- A relying party would have to trust that the device has not been compromised and to rely on the response from the device. A device based approach allows a device to undertake multiple attacks.
- Parents may input false age information to bypass restrictions under pressure from children.
- Where age controls are tied to devices rather than individual accounts, this provides less flexibility and may fail to address the nuanced needs of multi-user scenarios.

There must also be a level of trust with the on-device provider, relying party and consumer that any images captured by the device are not stored or shared. Without auditing it is not possible to be sure that an on device offering does not deliberately share or unintentionally leak personal information.

Is a one-off check at set up sufficient?

This begs a further question; in the instance that the operating system check has failed as the device is no longer in the hands of an adult, what other protections are in place? Take for instance the scenario of a 3 year old child whose parent does not set up their age as a child, but as an adult from the outset.
Another suggestion is to not just have a one-off check at set-up; but to also have an age check at the point of access

Layered checks both at set up and an additional check at the point of accessing a service

This could lead to a layered set of 'just in time checks' so both at set up and an additional check at the access point of say an age-restricted service, such as live chat in gaming or 18+ gambling, dating, adult content.
Stage one could be to bind age checks to passkeys and integrate mandatory age verification during device setup. And then subsequent periodic reauthentication could be required at device level. This could be supplemented by checks at the point of access of various services.

Potential unintended consequences

- Some apps and services may exclude minors entirely to avoid regulatory burdens, depriving them of safe online experiences.
- Increased friction in app downloads could discourage legitimate users.
- There could be complexities if the operating system level or device level age verification is different to a subsequent point of access age verification (i.e. because the device is being reused or shared) then would other services that rely on the historic app store age or device age be stopped?
- Small developers may struggle to meet certification requirements, stifling innovation and competition.
- Parents could face more rather than less complexity.
- Global platforms may face challenges in adhering to varying local regulations.

By requiring a layered approach, with more age checks than is the current requirement, this may lead some platforms to decide to only provide their services to adults and no longer service the youth market, thereby reducing the offerings to young people.
By requiring age assurance at several levels in the stack, this could also alienate users and discourage adoption. Parents who are already struggling to deal with parental controls, may give up on supporting operating system controls as well as parental controls, if they are perceived to be an additional complexity and burden.
Additionally, whilst larger operating systems and organisations may adapt to meet requirements; smaller organisations may struggle.

These are nuances that have not been discussed widely though.

Conclusion

The optimal placement of age assurance within the tech stack and consumer journey depends on balancing consumer convenience, privacy, and operational feasibility.

Today's approach is to undertake age checks 'at proximity' at the point where a user accesses a service. This is working now and at scale. There are billions of such age checks being undertaken annually; following international standards, which can be audited by independent accredited audit bodies. The more this develops across age restricted sectors, the lower the cost for all the industries involved. These approaches can be tokenized, with each token lasting for a limited time. There can be regular authentication to the current user of the device.

A future, multi-layered approach could be devised, which extends the number of places where age checks are undertaken. This type of approach would need to consider service-specific context, minimise data sharing, and align with global standards for age assurance that offers the most robust path forward. However this will take time to think through. We need to avoid creating an even more complex environment for parents and ensure that users understand the rationale for any change.

The existing age assurance industry, which has decades of experience in delivering independent and auditable approaches, is keen to support regulators to think through the practicalities as well as the positive and negative intended and unintended consequences of the sweeping statement to 'undertake checks at operating system level'.

Age checks are already required by law now, and there are approaches available that meet the requirements. Before, in one penstroke demolishing what exists today, a thorough review of the questions in this paper is recommended.



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

Robin Tombs' Post

Robin Tombs
Co-founder and CEO – Yoti

Yoti has waited a long time to see independent, high profile, well funded technology trials of our facial age estimation (FAE). We're delighted to read the results for Yoti just published in the Australia Age Assurance Technology Trial (AATT) Report.

Yoti released our first FAE model in Dec 2018 & published the performance results (tested against over 100,000 images with ground truth age across 13-60 by year of age, gender & skin tone).

Whilst almost all stakeholders welcomed our transparency, some stakeholders understandably were sceptical of our published internal testing results. They wanted independent validation.

Confidence improved from autumn 2023 when US **National Institute of Standards and Technology (NIST)** started publishing evaluation results for the Yoti Mar 2023 ('Yoti-001') FAE model. NIST currently judges Yoti's Sep 24 ('Yoti-003') FAE model the most accurate for estimating 13-16 year olds.

But today, the Australia AATT report shows Yoti (Sep 24) FAE to be the most accurate FAE model for its 12-23 year olds test set (with an MAE of 1.80 years for 12-23 year olds) & that:

- Yoti's Target of Evaluation is at Technology Readiness Level 9 (the highest level).
- Manual functional testing confirmed interoperability, robustness & privacy aspects of the system.
- Automated lab testing showed high accuracy for Age Gates 13 & 16, with True Positive Rates consistently above 94% from age 13 upwards and Mean Absolute Error (MAE) values under 2 years for ages 13–20.
- There were very few usability issues arising from Mystery Shopper trials. While response time is not as fast as some systems tested, but users reported that the system met or exceeded their expectations in terms of task time. The majority of users had no issues in completing the evaluation task & reported the system to be either easy or very easy to use.
- 80% of users reported being either satisfied or very satisfied with the experience, & the majority had no concerns about privacy.
- 66% reported accurate age estimation results, which is consistent with the overall average accuracy of 70% for the mystery shopper estimation task.
- We note that some users had an expectation of the system estimating their age exactly (to within a few months) & thus reported "no" for the accuracy question without allowing for any margin.

The AATT tested for a lot of important functions & Yoti passed every one - reassuring stakeholders that Yoti FAE does not store PII data, can detect paper, screen & mask attacks & can be effective for 13+, 16+ & 18+ age gating Indigenous (Aboriginal & Torres Strait) Australians

As can be seen below, our Jun 25 FAE model shows further improvement, with our (~125,000 global test set) MAE now at 1.60 years for 12-23 year olds & just 1.04 years for 14-18 year olds.

Australian Age Trial Facial Age Estimation Results			
		Average MAE	
Provider	Vendor	12-23	14-18
provider_N	Yoti	1.80	1.22
provider_E		1.89	1.77
provider_I		2.64	1.57
provider_G		2.83	2.60
provider_A		2.95	2.77
provider_F		3.19	2.88
provider_L		3.91	4.05
provider_K		4.26	4.30
provider_B		5.96	5.30

Yoti Facial Age Estimation Results				
	Australia Trials Results		Yoti Internal Test Results	
	Average MAE		Average MAE	
Yoti Model	12-23	14-18	12-23	14-18
Sep 24 Model	1.80	1.22	1.68	1.15
Jun 25 Model			1.60	1.04

System Testing: Manual Functional Testing

The Age Estimation scenarios are listed in the table below.

Test Scenarios	Results
Valid case for 16+ age estimation	Pass
Age Estimation Correctly Identifies User in 13+ Bracket	Pass
Age Estimation for Aboriginal Individuals	Pass
User Partially Accepts Privacy Terms but Does Not Provide Consent for Age Estimation	Pass
User Partially Accepts consent but Does Not accept terms for Age Estimation	Pass
Age Estimation Fails Due to Camera Permissions Denied	Pass
User Attempts to Use a Photo Instead of Real-Time Face	Pass
Age Estimation Accuracy Across Different Ethnicities and Skin Tones	Pass
Facial Structure Altered Due to Medical Conditions	Pass
User's Distance from Camera Affects Age Estimation	Pass
Exaggerated Facial Expressions Impact Age Estimation	Pass
Age Estimation Correctly Identifies User in 16+ Bracket	Pass
Age Estimation Correctly Identifies User in 18+ Bracket	Pass
User Accesses Age Estimation Feature from Different Australian States	Pass
User Does Not Properly Position Face in Frame	Pass
Age Estimation Affected by Poor Lighting Conditions	Pass
Multiple Faces in Frame during Age Estimation	Pass
Verify System Deletes the Image Post-Age Estimation	Pass
Facial Hair or Beard Impacting Age Estimation	Pass
User wearing mask trying to use face estimation	Pass
User Attempts to Bypass Estimation with Deepfake or AI Generated Image	Pass
User wearing hat during age estimation	Pass
Age estimation correctly works end-to-end using different devices. Laptop, PC, tablet, smartphone.	Pass
Age estimation correctly works end-to-end using different Browser (Chrome, Firefox, Microsoft Edge, Safari)	Pass
Solution does not show excessive and/or PII data in results.	Pass
Age Estimation Correctly Identifies Aboriginal User in 18+ Bracket	Pass
Age Estimation Correctly Identifies Aboriginal User in 16+ Bracket	Pass
Age Estimation Correctly Identifies Aboriginal User in 13+ Bracket	Pass
Age Estimation Correctly Identifies Torres Strait User in 13+ Bracket	Pass
Age Estimation Correctly Identifies Torres Strait User in 16+ Bracket	Pass
Age Estimation Correctly Identifies Torres Strait User in 18+ Bracket	Pass

Solution does not show excessive and/or PII data in results.	Pass