

Yoti Facial Age Estimation

White Paper | Full version

September 2024



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 Rate¹ (TPR) for 13 to 17 year olds correctly estimated as under the age of 21 is 99.2% and there is no discernible bias across genders or skin tones. The TPRs for female and male 13 to 17 year olds are 99.1% and 99.3% respectively. The TPRs for skin tones 1, 2 and 3 are 99.5%, 98.8% and 98.7% respectively.

The TPR for 6 to 12 year olds correctly estimated as under the age of 13 is 99.5%. The TPRs for female and male 6 to 12 year olds are 99.6% and 99.4% respectively. The TPRs for skin tones 1, 2 and 3 are 99.7%, 99.2% 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 1.2 years, and an MAE of 1.3 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 18 year old being correctly estimated to be under the age of 25.

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 99.2%.
- TPR for 6 to 12 year olds correctly estimated as under 13 is 99.5%.
- Mean Absolute Errors (in years) are 2.5 for ages 6 to 70, 1.3 for ages 13 to 17 & 1.2 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 for,....
- 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 650 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.

Mean Absolute Error by age band

YOTI 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	1.4	1.3	1.1	1.2	1.3	1.2	1.2
13-17	1.1	1.4	1.6	1.4	0.9	1.2	1.4	1.2	1.3
18-24	2.4	2.3	2.5	2.4	2.1	1.9	1.9	2.0	2.2
25-70	2.6	2.9	3.6	3.0	2.5	2.8	3.2	2.8	2.9
6-70	2.3	2.6	3.1	2.6	2.2	2.4	2.7	2.4	2.5

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 1.2. That means they would be estimated to be within 1.2 years of their age. A table of MAE by year can be found in the appendix on pages 20-22.

Contents

What is facial age estimation and what can it do?	6	Appendix	19
How accurate is facial age estimation	7	Mean average error by year	20
Safety thresholds	8	Data used to build the model ('training data')	23
NIST evaluation	9	Data used for testing	25
NIST evaluation comparison	10	Accuracy across the entire data set	25
NIST ethnicity approach and methodology	10	Accuracy by age, gender and skin tone	25
Growing adoption of facial age estimation	11	Standard deviation of absolute error	26
Learnings from practical use	12	Absolute versus percentage errors	27
Live trials in UK & Europe	13	Improvement in accuracy as the data set grows	28
Retail of Alcohol Standards Group report of live trials	14	False positives	29
Supporting Children's Codes	15		
Tackling bots, deepfakes and generative AI	16		
Legal compliance	17		
Yoti's commitment to ethical use of AI technologies	18		

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.

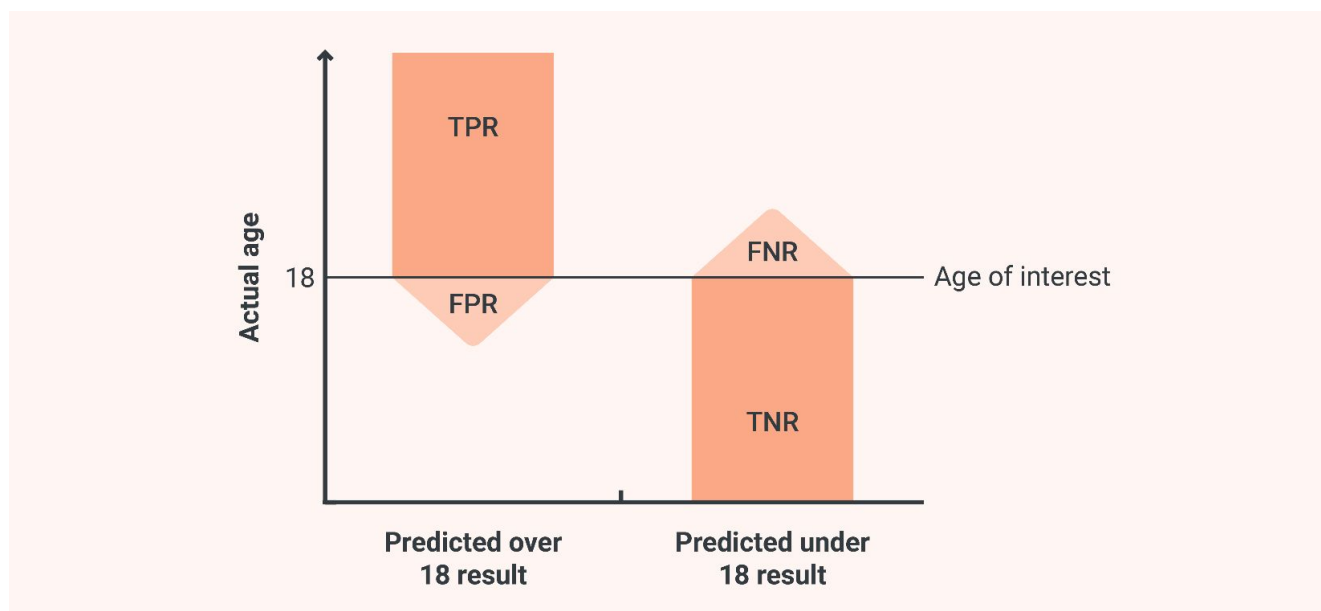
A further potential use is at the entrances to age-restricted premises such as bars, nightclubs and casinos. In this kind of application, facial age estimation offers clear advantages – it does not get fatigued on a long shift,² and it cannot show favour to personal friends, or selective bias against 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.

Safety thresholds

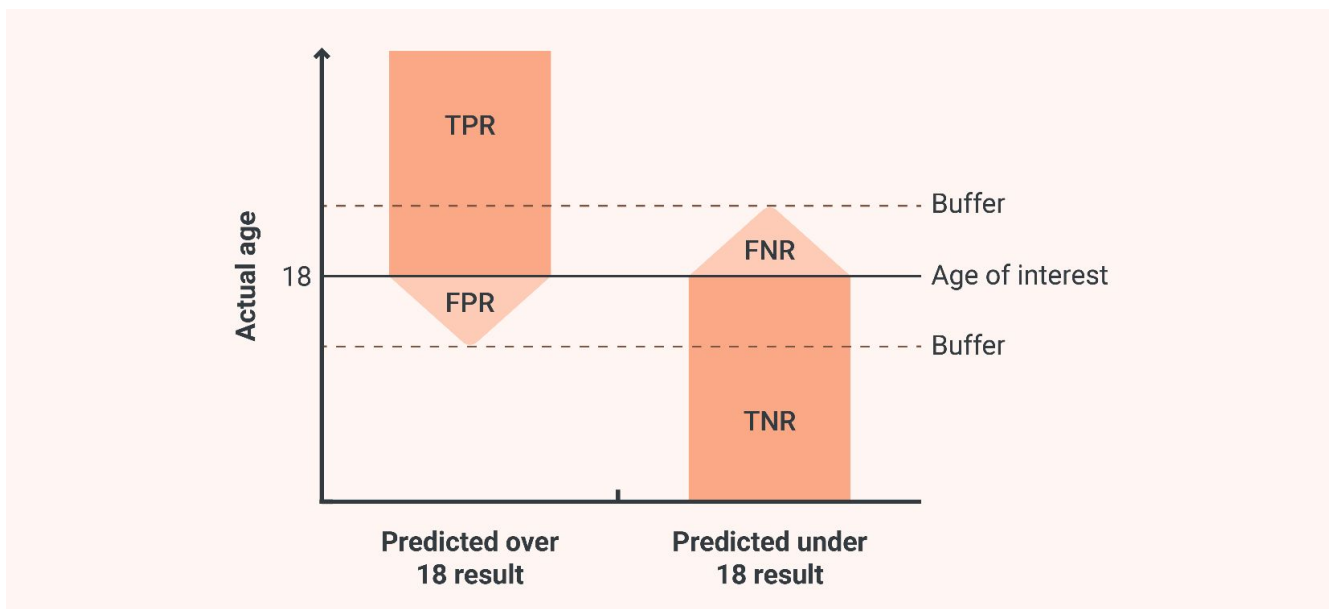
To manage the potential for errors, we recommend using facial age estimation as part of their compliance strategy. An example of this is the British Beer & Pub Association’s ‘Challenge 21’ scheme, which is already widely adopted by publicans and their bar staff in England and Wales.⁶ As it is difficult 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 0.8%.



6. See <https://beerandpub.com/campaigns/challenge-21/>

7. For more information see page 31

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 is reviewing age estimations compared to dates of birth on government ID docs for the above six regions and will publish data on this in the next white paper.

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 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⁾

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 verification 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 verification, 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.



- 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



- Company**
OnlyFans
- Industry**
Social media
- Solution**
Facial age estimation
- Implementation**
API and SDK

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

Challenge


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 create.

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 to prove their age.


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
Checks a second

18K
Checks a minute

1.1m
Checks an hour

26m
Checks a day


780m
Checks a month

Live retail trials in UK & Europe

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

- Participating supermarkets confirmed they support digital age verification and would welcome legislative change in this area.
- There were no reported sales of underage customers purchasing age-restricted items when using our age verification 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 verification 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 verification supports the ability for retailers to achieve the Licensing Objectives

Yoti facial age estimation is being rolled out by Salling in Denmark and we have live pilots underway in Germany with Edeka.



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 verification 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.



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

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>.

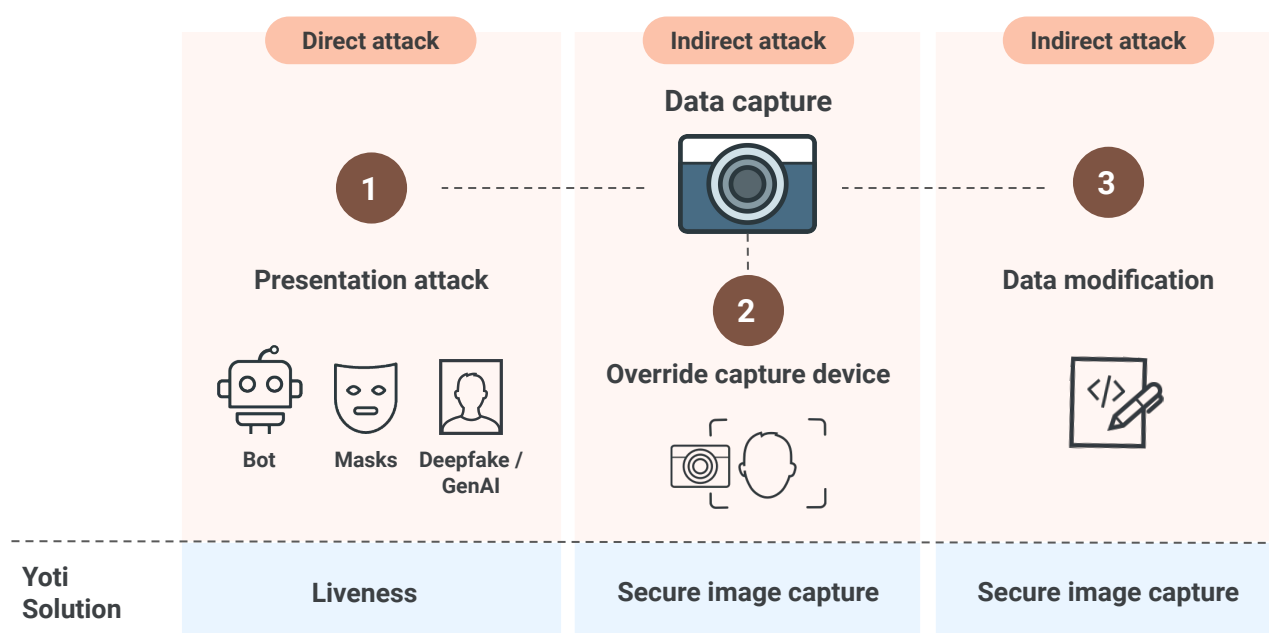
Tackling bots, deepfakes, and Generative AI

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.

Data capture attack threats



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 do. 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 verification solutions.⁸

We have hosted five 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.

Certified



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/>

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.

Mean Absolute Error by year

(Using test set of 118,720 images)

Age	Gender								All
	Female				Male				
	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	
6	1.1	1.6	1.7	1.5	1.0	1.7	1.9	1.5	1.5
7	1.2	1.1	1.3	1.2	0.8	1.1	2.1	1.4	1.3
8	1.3	1.1	1.5	1.3	1.5	1.5	1.2	1.4	1.3
9	1.2	1.2	1.6	1.3	1.4	1.2	1.2	1.3	1.3
10	1.0	1.1	0.9	1.0	1.0	1.1	1.0	1.0	1.0
11	0.8	1.2	1.7	1.2	1.0	0.9	0.9	0.9	1.1
12	1.3	1.4	1.4	1.4	0.9	1.2	1.0	1.0	1.2
13	1.8	2.3	2.4	2.2	1.4	1.5	1.9	1.6	1.9
14	1.2	1.6	2.0	1.6	1.0	1.3	1.9	1.4	1.5
15	0.9	1.2	1.5	1.2	0.7	1.1	1.4	1.1	1.1
16	0.8	1.0	1.1	1.0	0.7	1.0	1.0	0.9	0.9
17	1.0	0.9	1.0	1.0	0.9	1.1	1.0	1.0	1.0
18	1.4	1.3	1.0	1.2	1.3	1.3	1.1	1.3	1.2
19	2.0	1.9	1.7	1.9	1.8	1.6	1.5	1.6	1.7
20	2.4	2.3	2.2	2.3	2.2	1.8	1.9	2.0	2.1
21	2.8	2.6	2.3	2.6	2.3	2.0	2.0	2.1	2.3
22	2.9	2.6	3.0	2.8	2.4	2.1	2.2	2.3	2.5
23	2.8	2.9	3.8	3.2	2.4	2.1	2.2	2.2	2.7
24	2.5	2.5	3.8	3.0	2.3	2.2	2.3	2.3	2.6
25	2.4	2.5	3.7	2.9	1.9	2.1	2.3	2.1	2.5
26	2.0	2.6	3.1	2.6	1.7	1.9	2.3	2.0	2.3
27	1.9	2.5	3.6	2.7	1.7	1.9	2.1	1.9	2.3
28	2.2	2.3	3.2	2.6	1.7	2.0	2.4	2.0	2.3
29	2.2	2.3	3.8	2.8	2.0	2.0	2.4	2.1	2.5
30	2.3	2.8	3.0	2.7	2.0	2.3	2.7	2.3	2.5

Mean Absolute Error by year

Age	Gender								All
	Female				Male				
	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	
31	2.4	2.8	3.8	3.0	1.8	2.4	2.2	2.1	2.5
32	2.5	2.8	4.1	3.1	2.1	2.6	2.6	2.4	2.8
33	2.8	3.0	3.9	3.2	2.5	2.8	2.6	2.7	2.9
34	2.6	3.3	3.3	3.0	2.3	3.0	2.8	2.7	2.9
35	2.8	3.3	4.3	3.4	2.4	2.6	2.5	2.5	3.0
36	2.3	2.6	3.7	2.9	2.6	2.7	2.9	2.7	2.8
37	3.1	3.3	3.0	3.2	2.5	2.9	2.7	2.7	2.9
38	2.2	2.7	3.1	2.6	2.1	3.0	2.8	2.6	2.6
39	2.4	2.8	3.0	2.7	2.4	2.3	3.5	2.7	2.7
40	2.5	2.5	3.1	2.7	2.5	2.3	2.6	2.5	2.6
41	2.4	2.7	2.9	2.7	2.2	2.6	2.6	2.4	2.6
42	2.4	2.6	2.7	2.6	2.4	2.5	2.7	2.6	2.6
43	2.5	2.8	3.8	3.0	2.4	2.4	2.8	2.5	2.8
44	2.3	2.8	3.2	2.8	2.2	2.5	2.9	2.6	2.7
45	2.5	2.8	3.3	2.8	2.1	2.7	2.7	2.5	2.7
46	2.8	2.8	3.5	3.0	2.6	2.7	2.7	2.7	2.8
47	2.9	2.8	3.7	3.1	2.5	2.8	2.8	2.7	2.9
48	2.8	3.2	3.4	3.1	2.5	2.9	3.2	2.9	3.0
49	2.9	2.8	3.7	3.1	2.7	2.7	2.8	2.7	2.9
50	2.7	2.5	3.3	2.8	2.6	2.8	3.0	2.8	2.8
51	3.0	2.4	4.8	3.4	2.7	2.6	3.5	3.0	3.2
52	2.6	2.5	3.2	2.7	2.8	2.6	4.2	3.2	3.0
53	2.7	2.6	3.7	3.0	2.8	3.0	3.7	3.2	3.1
54	2.7	2.6	3.7	3.0	2.7	3.8	3.8	3.4	3.2
55	2.7	3.2	2.8	2.9	2.6	4.0	3.9	3.5	3.2

Mean Absolute Error by year

Age	Gender								All
	Female				Male				
	Skin Tone								Average MAE
	Type 1	Type 2	Type 3	All	Type 1	Type 2	Type 3	All	
	MAE	MAE	MAE	Average MAE	MAE	MAE	MAE	Average MAE	
56	2.5	3.1	3.4	3.0	2.5	3.3	3.7	3.1	3.1
57	2.9	2.7	4.6	3.4	2.8	3.2	3.4	3.1	3.3
58	2.3	2.6	4.1	3.0	2.5	3.3	3.3	3.0	3.0
59	2.4	3.1	2.7	2.7	2.4	3.1	3.3	3.0	2.8
60	2.2	3.3	2.4	2.6	2.5	3.0	4.4	3.3	3.0
61	2.4	3.5	2.9	3.0	2.7	3.0	3.4	3.0	3.0
62	2.1	3.6	2.7	2.8	2.5	3.0	4.0	3.1	3.0
63	2.4	2.5	3.3	2.7	2.6	2.9	4.4	3.3	3.0
64	2.6	3.2	4.5	3.4	2.6	2.9	3.4	3.0	3.2
65	2.4	3.0	4.6	3.3	2.6	2.7	2.9	2.8	3.0
66	2.8	2.5	4.3	3.2	2.8	3.5	3.5	3.3	3.2
67	3.1	3.2	4.2	3.5	3.2	3.7	3.6	3.5	3.5
68	3.2	3.0	5.6	3.9	3.7	3.5	4.0	3.7	3.8
69	3.6	4.5		4.1	3.5	4.1	1.3	3.0	3.5
70	3.3	5.2	4.0	4.2	3.0	4.2	9.3	5.5	4.8
Avg	2.3	2.6	3.1	2.6	2.2	2.4	2.7	2.4	2.5

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 (September 2024) 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 The images overleaf display the current onboarding screen and the screen where users can opt out of their data being used for R&D.

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>

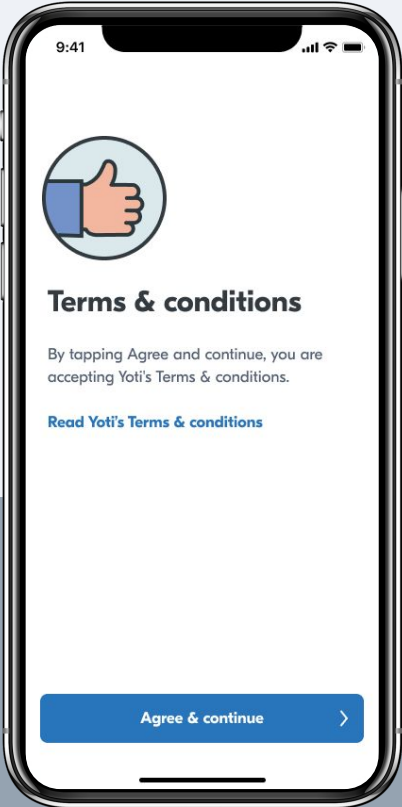
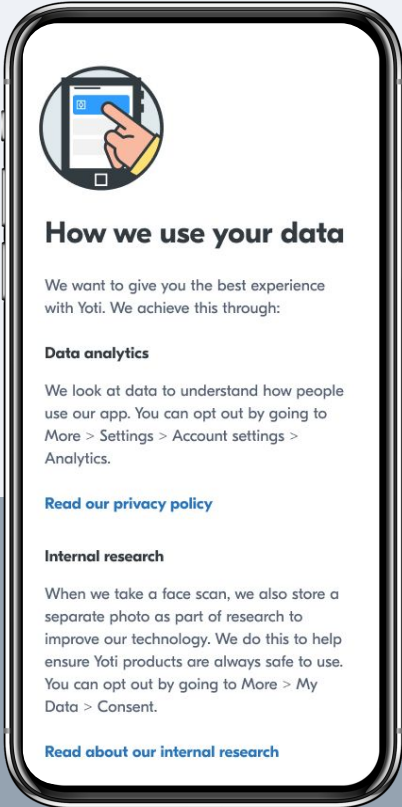
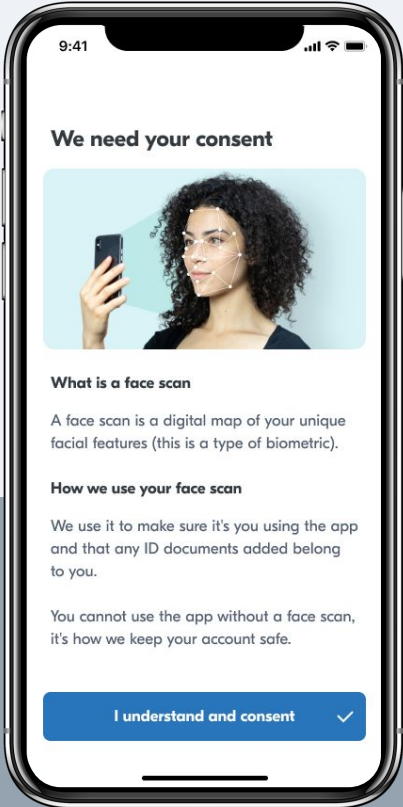
11. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) <https://eur-lex.europa.eu/eli/req/2016/679/oj>

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



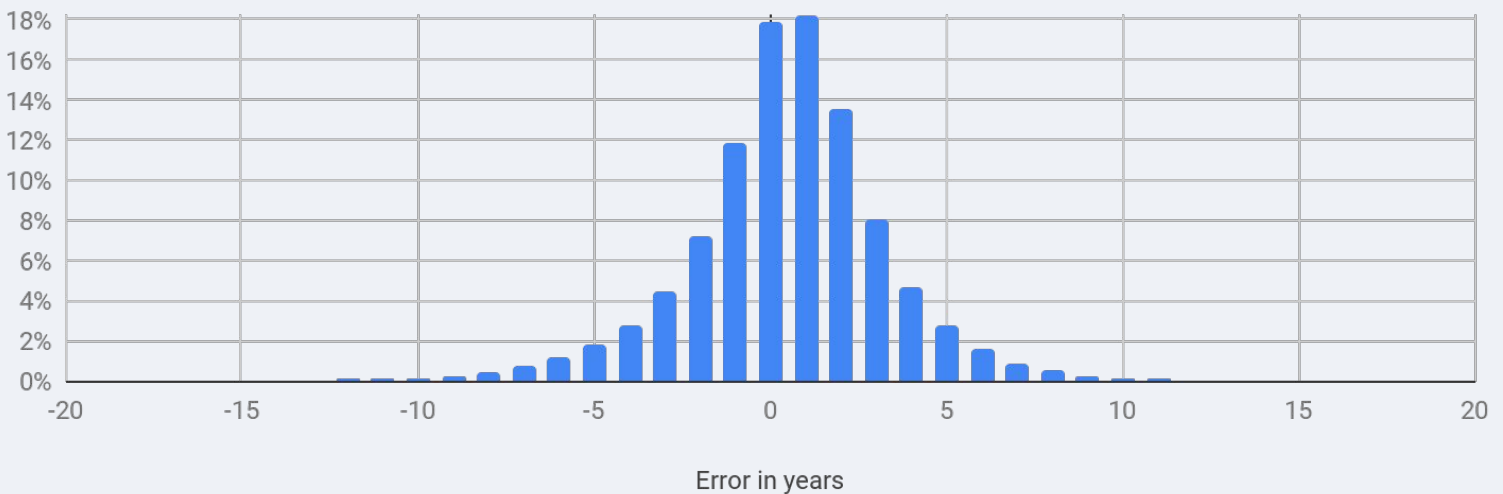
We provide information to users at onboarding about our use of biometrics. This includes links to further information, including the full privacy notice, where the use of user data for R&D is extensively detailed. Users can opt out of their data being used for R&D by Yoti at any time, via the settings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

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 (September 2024), we used test data that was comprised of hundreds of thousands of images with a verified age. The MAE across all years is now 2.5 years; for females it is 2.6, for males it is 2.4. 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 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								All
	Female				Male				
	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	1.0	1.0	1.3	1.1	1.0
10-12	1.0	1.0	1.3	1.1	0.9	0.9	0.8	0.8	1.0
13-15	0.9	1.1	1.2	1.1	0.7	1.0	1.0	0.9	1.0
16-17	0.9	0.9	1.0	0.9	0.7	1.2	1.0	1.0	1.0
18-24	1.5	1.6	1.7	1.6	1.3	1.5	1.5	1.5	1.5
25-29	1.9	2.2	2.8	2.3	1.6	1.7	1.9	1.7	2.0
30-39	2.0	2.4	3.0	2.5	1.7	2.0	2.2	2.0	2.2
40-49	2.1	2.6	2.9	2.5	1.9	2.3	2.2	2.1	2.3
50-59	2.0	2.3	3.5	2.6	2.1	3.0	2.8	2.7	2.6
60-69	2.0	2.5	2.8	2.4	2.2	2.5	2.9	2.5	2.5
6-70	1.5	1.7	2.1	1.8	1.4	1.7	1.7	1.6	1.7

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).

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.

Average Error and Error in % of Age



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.

Improvement in accuracy as the training data set grows and changes

Average MAE changes between AI models (by Age Group & Gender)



Average MAE changes between AI models



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		3,190	7,425	10,265	10,424	
Thresholds (years)	20	0.22%	0.59%	1.51%	4.42%	1.69%
	21	0.03%	0.34%	0.73%	2.43%	0.88%
	22	0.03%	0.23%	0.39%	1.35%	0.50%
	23	0.03%	0.11%	0.23%	0.65%	0.26%
	24	0.03%	0.09%	0.15%	0.36%	0.16%
	25	0.00%	0.03%	0.08%	0.24%	0.09%
	26	0.00%	0.01%	0.05%	0.12%	0.05%
	27	0.00%	0.01%	0.01%	0.01%	0.01%
	28	0.00%	0.01%	0.01%	0.00%	0.01%
	29	0.00%	0.00%	0.01%	0.00%	0.00%
	30	0.00%	0.00%	0.01%	0.00%	0.00%


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	
<i>Test Sample Size</i>		10,265	10,424	9,088	5,246	4,069	
Thresholds (years)	24	0.15%	0.36%	1.03%	3.58%	10.08%	3.04%
	25	0.08%	0.24%	0.47%	1.83%	5.63%	1.65%
	26	0.05%	0.12%	0.17%	0.93%	2.92%	0.84%
	27	0.01%	0.01%	0.07%	0.44%	1.18%	0.34%
	28	0.01%	0.00%	0.04%	0.17%	0.49%	0.14%
	29	0.01%	0.00%	0.03%	0.08%	0.20%	0.06%
	30	0.01%	0.00%	0.02%	0.04%	0.10%	0.03%
	31	0.01%	0.00%	0.01%	0.04%	0.10%	0.03%
	32	0.01%	0.00%	0.01%	0.02%	0.05%	0.02%
	33	0.01%	0.00%	0.01%	0.00%	0.02%	0.01%
	34	0.00%	0.00%	0.01%	0.00%	0.02%	0.01%
	35	0.00%	0.00%	0.01%	0.00%	0.02%	0.01%
	36	0.00%	0.00%	0.01%	0.00%	0.02%	0.01%
	37	0.00%	0.00%	0.01%	0.00%	0.02%	0.01%
	38	0.00%	0.00%	0.01%	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%	

Memberships, associations and accreditations



Reviewed by



To find out more visit yoti.com

© 2024 Yoti Ltd