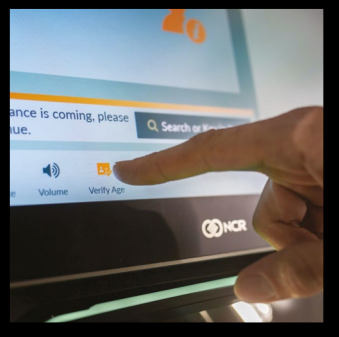


White paper

Yoti Age Scan - Public version

April 2019





Doing things differently

The solutions we develop make it faster, simpler and safer to prove who you are, and who other people are, online and in person.

The way we build those solutions is a little different to other tech companies. With Yoti, you have full control of your own data. We will never mine your data and once your personal details have been verified by Yoti, only you can see them. Even we can't see them.

But we don't stop there. Using our business as a force for good is something everyone at Yoti is passionate about. That's why we're a certified BCorp - internationally awarded for our commitment to rigorous standards of social, environmental performance, accountability and transparency.

Yoti Age Scan is just one innovative use of our digital identity technology. We built it to give everyone a secure and private way of proving how old they are in different everyday scenarios: from age checking in online stores and social platforms, to supermarket self-checkouts, bars and clubs. In this white paper we'll explain everything you need to know about the new way to prove your age.

Contents

What is Yoti Age Scan and what can it do?	04
Data privacy and network security	05
How does it actually work?	05
Tackling the challenge of age determination	06
Human ability to determine age	07
More on how it works	08
Practical use	09
How accurate is Yoti Age Scan?	10
Public acceptance of AI technologies	12
Appendix	13
Data used to build the model	13
Accuracy across the entire dataset	13
Accuracy by age, gender and skin tone	14

What is Yoti Age Scan and what can it do?

Yoti Age Scan is a secure age-checking service that can estimate a person's age by looking at their face. We consider it to have wide application in the provision of any age-restricted goods and services, both online and in person. It is also a means to combat social exclusion for the significant numbers of individuals around the world who do not possess state-issued photo ID.

Yoti Age Scan is designed with user privacy and data minimisation in mind. It does not require users to register with us, nor to provide any documentary evidence of their identity. It neither retains any information about users, nor any images of them. It simply estimates their age.

In a retail setting, Yoti Age Scan can be used at a point-of-sale terminal with a dedicated camera, letting a consumer use a self-checkout without the need for staff assistance. This is not only quicker and less of a nuisance for shoppers, but can greatly reduce friction between them and retail staff.

For general online use, it can be embedded into web pages or incorporated into apps, and receive an image of the user's face from a webcam connected to their computer or the camera in their mobile device, ideal for controlling access to age-restricted gaming, gambling and also adult content (pornography).

We believe Yoti Age Scan can play an important role in safeguarding and child protection online, not only in preventing minors from accessing adult content, but also in preventing predatory adults from accessing social media spaces for children and teenagers. This is illustrated well by Yoti's partnership with Yubo social networking platform. Yubo uses Yoti Age Scan within its app to help identify user profiles where there is suspicion or doubt about the user's age, and flag these to its moderation team.

A further potential use is at the entrances to age-restricted premises such as bars, nightclubs and casinos. In this kind of application, Yoti Age Scan offers clear advantages — it does not get fatigued on a long shift, and it cannot show favour to personal friends or bias against individual customers¹.

Yoti Age Scan 'always forgets a face'

Yoti Age Scan is an emerging technology, and its age estimates are subject to a margin of error. To allow for this, the system is configurable to set whatever threshold a business or regulator requires for instance, requiring those over 18 to be estimated as at least 25 - a buffer of seven years. Where someone is over 18 but appears to be under 25, they can use either the Yoti app, where their account is anchored with a verified ID document, or by a manual photo ID check from a member of staff).


Data privacy and network security

Yoti Age Scan has been designed with data privacy and security as primary considerations.

The user does not have to register to use the service, and does not have to provide any information about themselves. They simply present their face in front of the camera. Their image is not stored locally on the point of sale terminal. It is securely transmitted to the Yoti backend server (currently hosted in the United Kingdom), secured by TLS 1.2 encryption. After the age estimate is performed, the captured facial image is deleted from Yoti's backend servers. The images are not stored, not re-shared, not re-used and not sold on.

How does it actually work?

Yoti Age Scan is based on a computing technique known as a 'neural network', which we have trained to be able to estimate human age using a process of 'machine learning'. This is a form of artificial intelligence (AI), and is increasingly used in a wide variety of applications, from driverless cars to medical diagnosis, from tailoring online advertising to detecting credit card fraud. We discuss machine learning in more detail below, but first some context on the problem we are using it to solve.



1. Studies have shown that the objectivity of human judgement of this kind can be significantly affected by hunger and fatigue—see for instance Danziger, Levav, Avnaim-Pesso (2011) Extraneous factors in judicial decisions, Proc Natl Acad Sci USA, 108 (17): 6889–6892. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3084045/>

Tackling the challenge of age determination

Determining a person's exact age in the absence of documentary evidence of their date of birth is a difficult task. Indeed, the truism that 'age is just a number' could be said to have a sound scientific basis. By 'ageing' in a medical sense, we mean the physiological changes which occur when individuals develop and grow from juvenile to mature forms, and then the types of damage that progressively accumulate within the human body as time passes. The important point is that the rate at which human bodies 'age' in this way is influenced by numerous external factors other than simple passage of time. Factors that affect the ageing process, both in the long and short term, can include: quality of diet and nutrition, exposure to disease, adverse environmental conditions, use of narcotics, physical labour, stress and lack of sleep. Clearly, there are large variations throughout populations as to how different individuals are exposed to these aging factors. The more extensively we look through different countries, ethnicities, and socio-economic groups, the wider these variations in exposure to ageing factors become.

It may be surprising to learn that there are currently no entirely reliable medical or forensic methods to determine human age. Two of the more commonly attempted medical techniques focus on trying to ascertain whether the subject is above or below the legal age of maturity. These are X-ray or Magnetic Resonance Imaging of bone structure in the wrists (the degree to which the cartilage between the carpal bones has ossified) and dental X-rays (examining the maturity of wisdom teeth). However both of these methods have a typical margin of error of at least two or three years, and for individuals with an atypical history to the general population, the error can be significantly worse. Due to this unreliability, their use has proved controversial — for instance, their use by immigration authorities to attempt to differentiate between child and adult refugees who have no documentation.

Other medical techniques examine 'biomarkers' taken from blood or tissue samples. Examples include measuring the degree of DNA methylation present, the length of the 'telomere' portion of chromosomes, or the serum levels of the metabolite C-glycosyl tryptophan. Whilst these biomarker techniques tend to provide good indicators of ageing processes in an individual, they do not correlate reliably with their chronological age from date of birth.

Ultimately, it could be argued that much of the difficulty in trying to measure 'age' (that is, a person's chronological age from their date of birth) arises because 'age' defined this way is a rather arbitrary quantity that does not mean anything definite in physiological terms. Science can accurately measure the extent to which a person's body has aged (that is, how to what extent it has developed, grown, matured and decayed), but cannot always reliably determine how many years it took for their body to arrive at that state.

Human ability to determine age


Notwithstanding the difficulty in devising an accurate forensic test for age, people still possess a reasonably good ability to guess someone's age simply by looking at them. We can all do it, usually coming within a few years of the right answer. How do we manage it? In terms of facial features, what are the tell-tale signs we look for?

The most obvious visual cues include bone structure (bones grow and develop as we pass from child to adulthood), skin tone (wrinkles, elasticity) and hair colour (greyness), male baldness or facial hair after puberty. We could add dozens more cues to this list. However, whatever the detailed nature of the visual cues, the more general point is this: as humans, we simply learn "that's what people of a particular age look like". As we go through life, we encounter other people, we see what they look like and we learn how old they are, with varying degrees of precision (e.g. "a baby", "14", "mid-40s", "79" and so on). We accumulate this information and experience throughout our lives, and our brains can use it to make quick intuitive judgements. The extent of our previous experiences will be an important factor in how good our guesses are. We will be more accurate at guessing the age of someone from our own familiar peer group than from one we've not encountered.

It is worth emphasising that, although we might be able to retrospectively rationalise or refine our guess at someone's age, our initial judgement is more or less intuitive. We are not consciously following some step by step, rule-based method (for instance "add five years if there are wrinkles", or "add ten years for grey hair"). In effect, we don't 'know how we do it' — generally, our brains process the image and form an instinctive judgement, in line with what we've learnt from past experience, faster than any conscious deliberation or systematic evaluation of facial features. It turns out that this 'black box' approach to describing our cognitive process (that is, simply training our brain with data, without worrying too much about how it works) can actually be employed as a successful technique in machine learning too.

More on how it works

The first challenge for Yoti Age Scan is 'face detection'. It has to examine the image it gets from the camera, and work out which bit of it is an actual human face. Only this portion of the image is then fed into the neural network to get an age estimate. This stage also allows for basic error checking: if the system can't find a face in the image (for example, because a customer didn't position themselves properly in front of the camera, or some inappropriate object is put there) then the system can return an error message instead.



We now come to the interesting bit. The facial image is made up of pixels. To the computer, each pixel is just a set of numbers. These numbers are fed into the artificial neural network. This is a network of mathematical processing nodes, arranged in layers, that is roughly analogous to the connections in the human brain. Whilst a typical brain has around 100 billion neurons, the artificial neural network has just hundreds of thousands of nodes. We feed numbers (pixel data) in, and they percolate through the neural net. Each node performs a mathematical function on the pixel data, and passes the result on to nodes in the next layer, until a number finally emerges out the other side. This number is an age estimate.

It's an obvious question to ask 'how is the neural network processing the data? What is it looking for — wrinkles? grey hairs?' and so on. However this is a rather human way of thinking about it, and it's not really a very useful question to ask: to the computer, it is just being fed numbers. It doesn't 'know' what the numbers represent or what they mean. We don't try to tell it that. What we have told it, in the training phase when Yoti Age Scan was being developed, was what the right answers were. In the training phase, we fed it tens of thousands of diverse facial images, for which we knew the subject's age with confidence. The neural network keeps digesting the pixel data from each image, processing the numbers, and trying to get a result which matches the right answer. It keeps repeating the process, adjusting the processing, keeping the variations which bring it closer to the right answer, rejecting the variations which don't help — in other words, it is 'learning'. After repeating the process a huge number of times, it arrives at sets of processing formulae which


work best. To a human, these formulae would be bafflingly long and complex, and next to meaningless (and no, we're not going to print them here...for one thing, they wouldn't fit on the page). However it has effectively created a very complex model of age determination that is far superior to relying on a set of hand-crafted instructions that a human programmer might supply.

The quality of the training data is crucial to any machine learning process. We use many thousands of images from Yoti users (as explained to them at onboarding²). Yoti users want Yoti to make their lives safer and simpler, and they understand that using their data for research and development purposes is how we are able to improve and develop the products and technology to achieve this. We will publish white papers that demonstrate such

applications. For Yoti Age Scan, these research images are tagged with only two attributes taken from a verified ID document that they have uploaded: their gender and their year of birth. Supported documents include passports, driving licences and national ID cards. We believe the size, diversity and verified age accuracy of this training data set gives Yoti Age Scan an advantage over competing solutions.

Practical use

Yoti Age Scan works quickly, returning an age estimate in around 1 to 1½ seconds. The user needs to present their face to the camera, uncovered (although glasses do not usually present a problem). Dim lighting is not helpful; bright ambient light works best.



How accurate is Yoti Age Scan?

We believe that when presented with a clear facial image, Yoti Age Scan's ability to estimate age compares favourably with human abilities.

Research in this area suggests that the root mean square error in human guesses across an age range of 7 to 70 approaches 8 years³. Furthermore, when viewing a succession of faces, a person's judgement tends to be influenced by the preceding faces they have just seen, which is not a problem that affects Yoti Age Scan. Humans tend to systematically underestimate the ages of older people, and overestimate the age of younger people. The latter problem clearly has particular implications for provision of age-restricted goods and services, where we need to check whether teenagers are above or below a required legal age.

The scatter graph overleaf shows how our model performs with our test data set. Each circle represents a single human subject; the graph plots the subject's actual age (x axis) against the age estimate produced by Yoti Age Scan (y axis). If Yoti Age Scan always guessed correctly, the plot would be a straight diagonal line. As can be seen, whilst not perfect, there is a good correlation between actual age and estimated age, and the deviation is reasonably consistent throughout the entire range. Currently, the mean absolute error across the entire data set is 3.22 years. Further detail accuracy, broken down by gender, skin tone and age range, is presented in this paper's appendix. We believe this accuracy will improve still further in years to come, as Yoti Age Scan is trained on an ever greater set of data from Yoti users. We intend to continue comparing Yoti Age Scan's accuracy against that of ordinary human estimators, and against people who believe they

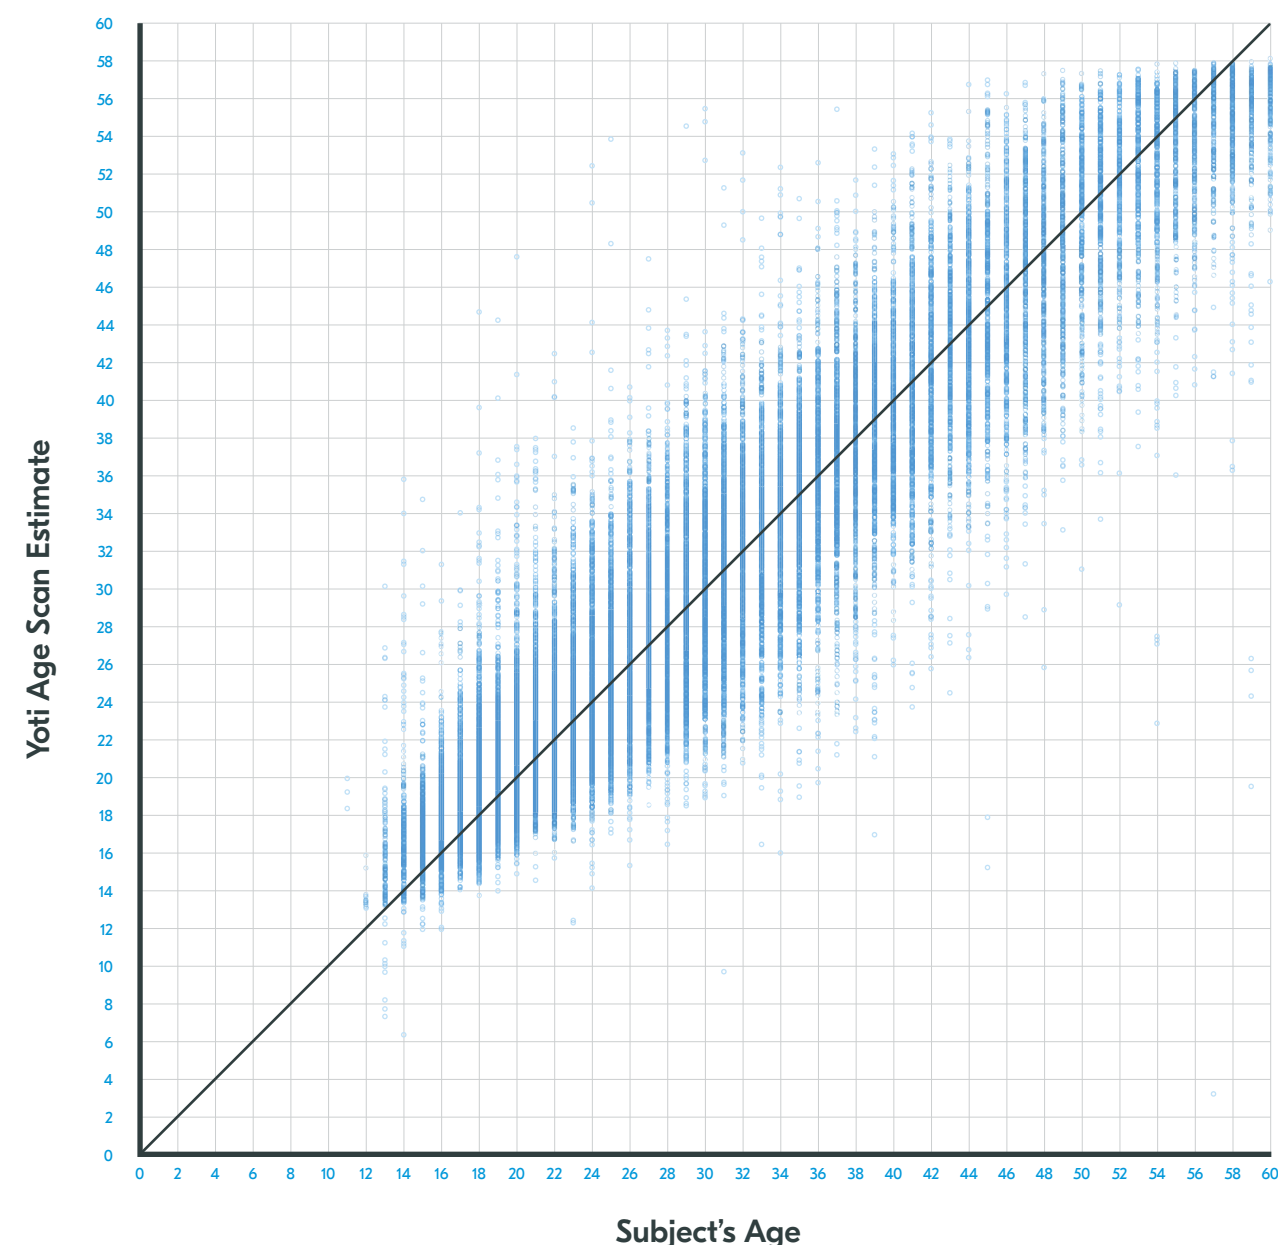
have a special aptitude at estimating age, to demonstrate that Yoti Age Scan is usually a more accurate approach (and cheaper and faster).

For this reason, Yoti recommends using Yoti Age Scan as part of a strategy such as the UK Wine and Spirit Trade Association's 'Challenge 25'. This type of retail strategy works as follows: Certain goods and services can only be sold to customers over a particular age (e.g. 18 years old). However it is difficult for human staff to be sure whether someone is over 18 just by looking at them. Conversely though, it is fairly easy to tell if someone is significantly older than 18, and customers in this age range would find it an unjustifiable inconvenience to have to show ID to prove their age. Therefore, the store's policy is to only require customers to prove their age if they appear to be under 25.

Yoti Age Scan can be configured to work with legal age thresholds in a similar way: choose a suitable buffer that is comfortably outside Yoti Age Scan's margin of error, and configure the system to estimate whether customers are above or below that threshold. At present we would not permit a buffer of less than five years.

For instance, selling alcohol in a jurisdiction with a legal age restriction of 21, Yoti suggests choosing an initial threshold of 30. If Yoti Age Scan estimates that the customer is at least 30 years old, then no further age checking is required. If Yoti Age Scan estimates that the customer is below 30, then they will be directed into a user flow where they need to present documentary proof of their age (for example, using their Yoti app that is anchored to their passport, driving licence or national ID card).

This is great news for the 30 plus population — they will not need to provide ID document evidence of their age and they will be able to happily leave their documents at home. In some countries regulators may initially look for a high buffer of over ten years. With a 15 year buffer we would currently have a nil per cent error rate and with a ten year buffer we would have a 0.86% error rate. Over time regulators will be able to set lower buffers with confidence.




Graph displays results obtained with Yoti age estimation model as of 18-March-2019

3. Clifford CWG, Watson TL, White D. (2018) Two sources of bias explain errors in facial age estimation. R. Soc. open sci. 5:180841. <http://dx.doi.org/10.1098/rsos.180841>

Public acceptance of AI technologies

When discussing the accuracy of Yoti Age Scan, it is worth considering a general point about machine learning and the public's attitude to AI technologies of this kind: namely, how unforgiving humans tend to be in regard to mistakes made by AI.

Whilst we feel it is fair to claim that the accuracy of Yoti Age Scan generally compares favourably with human judgement in the broad majority of cases, there will inevitably be rare occasions where it 'makes mistakes'. Of course, humans make mistakes too. However, sometimes machine learning systems make mistakes that no human would have made. This is illustrated in the Venn diagram below:



■ Errors made by humans ■ Errors made by machines
 ▨ Errors humans react more badly to

As can be seen, typically, humans make errors, just as a well-trained machine learning system does. Furthermore, in most of the cases where the machine system gets it wrong, a human would make the same mistake. However, humans tend to be much more bothered by the small percentage of cases on the right of the diagram — these are cases where the machine learning system makes a mistake, but a human would not have been fooled. It can be argued that this is an irrational reaction, and objectively, the machine learning system is no worse than the human judgement it is replacing (sometimes it may even be better overall). Nevertheless, the general public may often unduly focus their attention on the machine failings, until they become comfortable with the new technology.

Yoti has signed the Safe Face Pledge⁴, which encourages companies using artificial intelligence to ensure that facial recognition technology is not misused.

Appendix

This appendix provides further detail on the current accuracy of Yoti Age Scan's estimates. We expect these figures to continue to improve as the volume and diversity of our dataset increases.

Data used to build the model

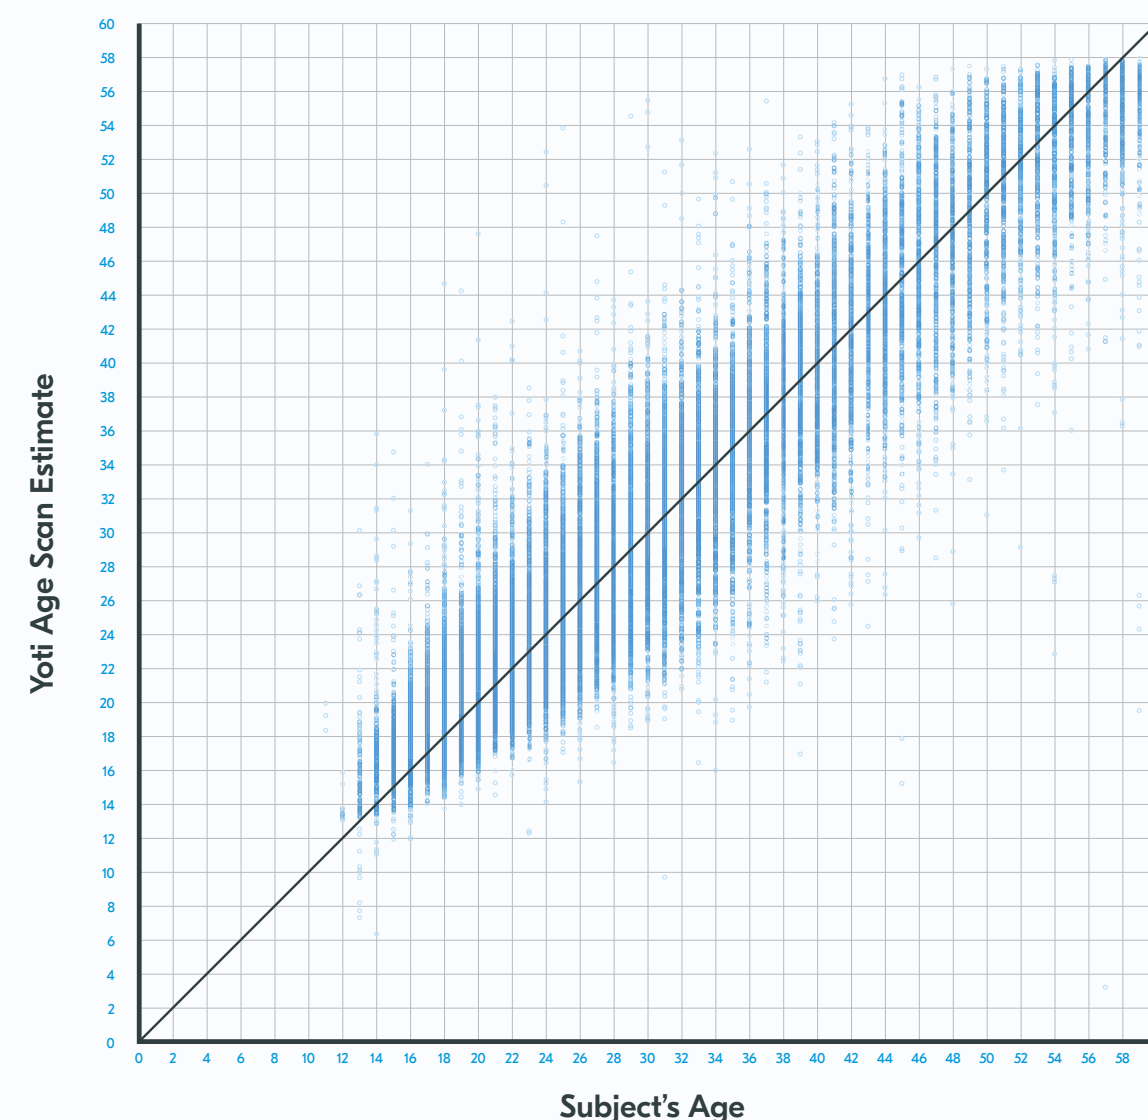
Yoti has invested significantly in building a leading R&D team since early 2015 working on a variety of AI initiatives.

The current production model of Yoti Age Scan (18th March 2019) was built using a training data set taken mainly from Yoti users⁵, and from a public domain source⁶. Further age-verified images were gathered by Yoti in Nairobi, Kenya.

Accuracy across the entire dataset

In our most recent testing of the model (performed 23rd April 2019), we used test data comprising 35,301 facial images of verified age, taken from Yoti users. The mean absolute error in age estimates (across entire data set) was 3.22 years.

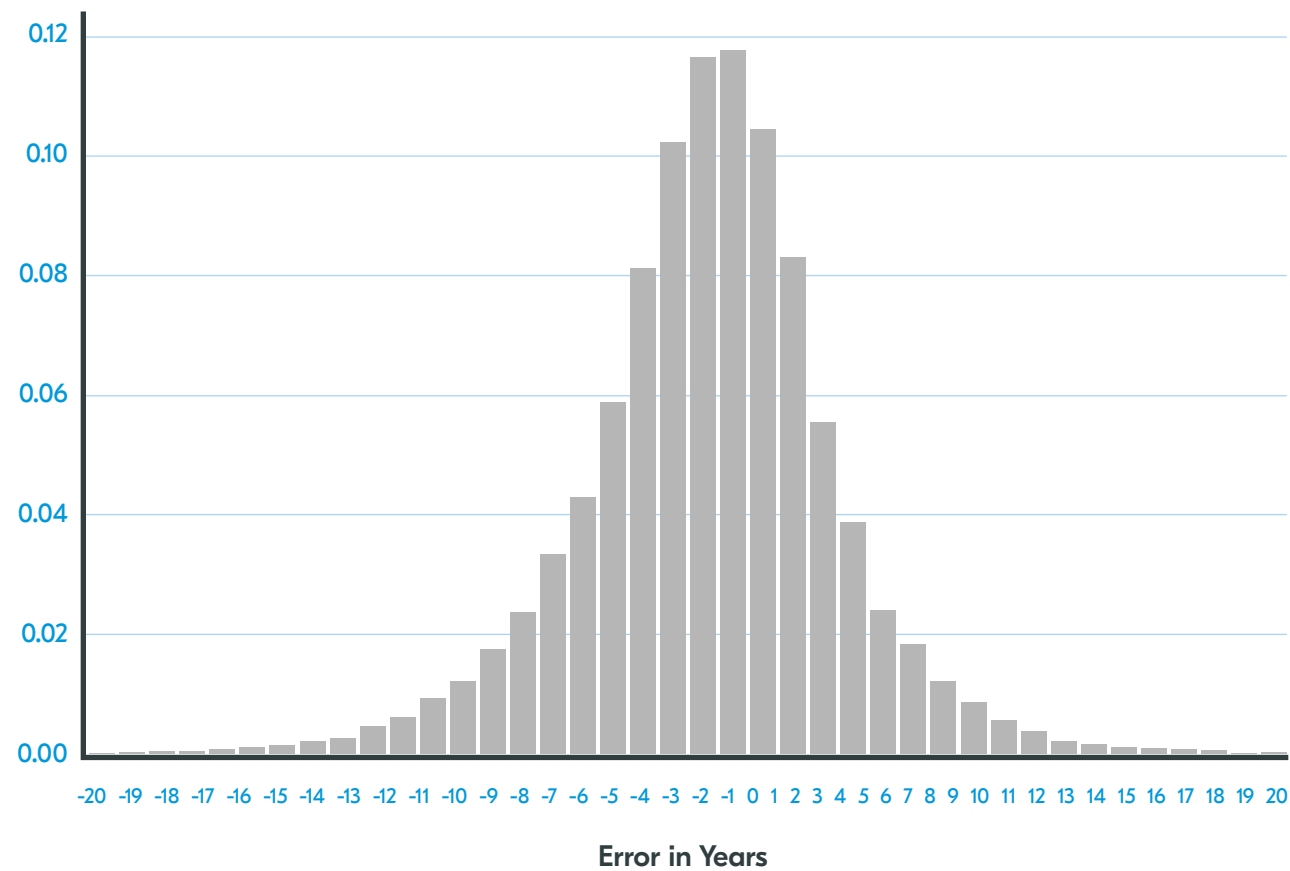
Graph displays results obtained with Yoti age estimation model as of 18-March-2019



5. The Yoti app explains to new users when they are in the process of creating an account that their biometric data (specifically, their photo and year of birth) may be used by our R&D team for internal research. More detail is set out in our Privacy Policy <https://www.yoti.com/privacypolicy/>

6. Images were taken from the Computer Vision Center and University of Barcelona's APPA-REAL Database, <http://chalearnlap.cvc.uab.es/dataset/26/description/>

The range of errors tends towards a normal distribution, with a standard deviation of 4.31. This is illustrated in the chart below.



Accuracy by age, gender and skin tone

We explored how the accuracy (mean absolute error) of Yoti Age Scan varies with age, gender and skin tone. The 35,301 facial images of verified age in our test set were tagged with the subject’s gender and skin tone. Gender was taken from the subject’s uploaded identity document. For skin tone, our research team manually tagged the images according to the widely used Fitzpatrick⁷ scale. Fitzpatrick uses six bands, from Type I (lightest) to Type VI (darkest). For now, we have used three bands (combining Fitzpatrick Types I & II, Types III & IV, and Types V & VI).

In presenting the data, we have grouped it into age bands, focusing on bands which are of particular concern to regulators as regards to the safeguarding of minors and access to age-restricted goods, services, websites and premises.

For each age band, we present the mean absolute error (MAE) in Yoti Age Scan’s age estimates across six categories of subject: female (for three different skin tones), and male (for three different skin tones). There were at least 100 test subjects in every category. For each age band, the table also displays the mean MAE for females (of all skin tones), for males (of all skin tones), and the mean MAE for all six.


Mean Absolute Error (MAE) of Yoti Age Scan for different genders and skin tones, across age bands of interest.

Age Band	Gender								
	Female				Male				All
	Skin Tone (Fitzpatrick Scale)								
	Type I & II	Type III & IV	Type V & VI	All	Type I & II	Type III & IV	Type V & VI	All	
	MAE	MAE	MAE	Average MAE	MAE	MAE	MAE	Average MAE	
13—15	3.37	3.77	4.61	3.92	2.17	3.03	3.67	2.96	3.44
16—17	2.25	2.84	4.46	3.18	1.53	2.17	3.01	2.24	2.71
18—24	2.04	2.20	4.28	2.84	2.13	2.35	2.80	2.43	2.63
25—29	3.60	3.86	4.40	3.95	3.17	3.05	3.45	3.22	3.59
30—39	4.56	5.60	4.86	5.00	3.19	3.39	3.87	3.48	4.24
40—49	4.37	5.85	6.42	5.54	3.50	3.72	4.30	3.84	4.69
50—59	4.07	6.23	6.99	5.76	2.97	3.63	4.10	3.57	4.67

We believe the differing mean absolute errors shown for different categories (age, gender, skin tone) correlates strongly with how well represented those groups are in the training data set. Additionally it seems reasonable to hypothesise that any error will tend to be higher for older people than younger people, because older people will have been exposed to various unpredictable environmental factors for longer. Also, it is worth noting that although the magnitude of error may appear larger for older age bands, that when considered as percentage of the subject’s age, in relative terms it may be more accurate (for instance an error of 2 years for a 15 year old is a 13.3% error, whereas an error of 2 years for a 50 year old is an error of only 4.0%).

Following the initial pilot study, we envisage continual periodic retraining of the model on an ever-expanding data set, as we continue to add further age-verified images taken from Yoti users at onboarding. We are optimistic that we will be able to broaden the diversity of our training data over the coming months. Where appropriate we shall endeavour to undertake further targeted fieldwork in this regard.

7. Fitzpatrick, T, (1988) ‘The Validity and Practicality of Sun-Reactive Skin Types I Through VI’. Archives of Dermatology 1988; 124 (6): 869–871



To find out more visit
yoti.com

