Yoti MyFace ® Liveness

White Paper | Full version

January 2026

Executive Summary

The growing use of biometric and authentication solutions, online and offline, has raised the risk of 'spoofing' attacks, an attempt to trick the system with an artificial representation. Having robust technology to mitigate against spoofing is essential as part of a mix of tools to verify someone. This is true whether that be for verifying age, identity or authenticating a returning customer.

The purpose of liveness is to make sure the person you are verifying is present in front of the device camera in real time. This is also sometimes known as Presentation Attack Detection (PAD). Liveness does not recognise who the person is (that is facial recognition). It is most commonly used in combination with other authentication factors to ensure that authentication or verification isn't being spoofed.

"The purpose of liveness is to make sure the person you are verifying is present in front of the device camera in real time"

In combination with wider AI services, liveness can help provide stronger security for individuals and businesses for verification and authentication. Liveness models can be small enough to exist on a device and therefore work offline - such as on a mobile device, self checkout or terminal.

Without liveness, businesses and individuals are vulnerable to presentation attacks (PAD) and injection attacks (IAD). Such attacks include:

- Paper image
- Video calls
- Screen image
- Video imagery
- Deepfake video
- Injection attacks
- Bot attacks
- Disguises and props

Yoti has found that genuine customers do not object to efficient and effective security layers. In fact most customers appreciate companies taking the time and effort to ensure their online accounts, assets or finances are being protected.

Yoti MyFace is the first, and at time of publication, only liveness model to conform to iBeta Level 3, passing with 100% attack detection. The iBeta Level 3 MyFace model discussed in this paper is currently our development model. We expect it to be in production available to customers before the end of March.

Contents

Why liveness is important

Liveness detection ensures that an identity check is being performed by a real, physically present human being, rather than a static image, recorded video, or synthetic representation. As identity verification has increasingly shifted from in-person checks to remote digital processes, new opportunities have emerged for bad actors to attempt to trick systems.

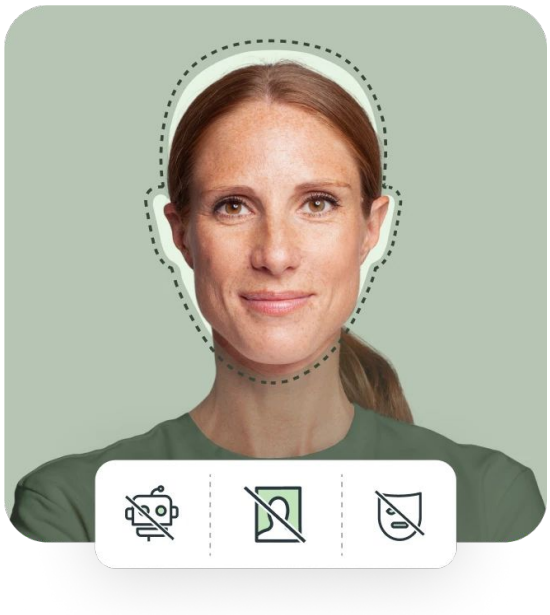

Attackers could be underage children trying to access adult content or sophisticated fraudsters trying to hack bank accounts. Liveness addresses this by providing assurance that biometric data is captured in real time from a genuine individual.

Without liveness detection, biometric systems are vulnerable to spoofing methods such as photos, screen replays, video recordings, masks, or AI-generated deepfakes. This protection is increasingly important as the quality and accessibility of GenAI and synthetic media continue to improve.

Liveness detection also plays a critical role in protecting individuals from fraud. Even if personal data or identity documents are compromised, liveness provides an additional layer of security by requiring the physical presence of the legitimate user at the moment of verification. This reduces the risk of account takeover, unauthorised access and financial fraud.

Beyond security, liveness is fundamental to establishing trust in digital interactions. Sectors such as financial services, employment checks, healthcare, education, online dating, and age-restricted services all rely on remote identity verification. Liveness helps replicate the assurance traditionally provided by face-to-face checks, enabling organisations to verify users confidently without physical meetings.

Finally, liveness supports regulatory and compliance requirements across areas such as KYC, AML, age verification, and right-to-work or right-to-rent checks. It strengthens digital identity processes by demonstrating that biometric data was captured at a specific point in time and not reused, replayed, or fraudulently generated.

In summary, liveness detection confirms real-time human presence, reduces fraud, protects users, and underpins trust in modern digital identity systems.

Where can Liveness be used?

	**Identity Verification** - used as part of the verification process to give a high confidence that the check is effective. **Why important?** - a stolen genuine ID document plus an image of the person in the document could spoof an identity check. **Example use case** - your customer wants to sign up to a new bank account and is required to prove their identity. Liveness ensures the person signing up is a real person. Liveness is used in combination with data extraction, document authenticity and face match for a secure verification process.
	**Age Verification - online** - ensuring the person is not only the right age but also not attempting to spoof the system with a presentation attack. **Why important?** - to strengthen age verification and provide quick, privacy-preserving age checks. **Example use case** - a gaming site might have over and under age restrictions - so that adult chat openly and younger people can enjoy an age appropriate experience. .
	**Age Verification - person** - retail outlet point of sale, self checkouts and kiosks can perform age checks without the requirement for human intervention. Liveness ensure an image or other spoofing attack cannot bypass the check. **Why important?** - to strengthen age verification and provide quick, privacy-preserving age checks, particularly at un-manned transactions. **Example use case** - a shopper uses facial age estimation at a supermarket self-checkout. Liveness prevents someone from attempting to use a picture, mask or other spoofing attack to pass an age check.
	**Digital ID** - to create a reusable Digital ID we require a high level of confidence to determine that the person presenting the document is real and matches the face in the document. **Why important?** - pre-verified, reusable Digital ID provides the highest level of security and confidence for a business to ensure the customer is who they say they are. **Example use case** - A returning customer can easily access their digital ID wallet account, and a business can be confident the customer is who they say they are.

Where can Liveness be used?

	**Authentication** - liveness can be an additional form of authentication for high risk / regulated environments, adding an extra layer which makes it harder for spoofers to scam. **Why important?** - multi-factor authentication is now required by many regulators and an efficient, low friction way to do this is actually desirable for customers. **Example use case** - a genuine customer is accessing their account, or changing important information such as their bank details. A liveness check with a face match can add an authentication layer with low friction. Liveness can prevent bots from signing in to accounts.
	**Verified Video Calls** - checking for liveness at the start, or during, a video call, can help ensure attendees on a call are not fraudsters, bots or deepfakes. **Why important?** - confidential, recruitment, financial services, or education calls are susceptible to fraudsters and can be costly. **Example use case**: a C-suite level call requesting the finance department to transfer funds to a new account could be spoofed by a fraudster using relatively basic tools.
	**eSign** - enhance contract signing even remotely, ensure a real person signing a document, just like an in person signing ceremony. **Why important?** - quickly, easily and securely completing the document signing ceremony can rapidly speed up completion agreements, to a high standard of security. **Example use case**: signing an employment contract, or a rental agreement, can be faithfully achieved remotely, by adding liveness as part of an authentication check.

Types of liveness - active and passive

Active liveness requires the user to present their face on camera and then follow one or more instructions during a check; for example, nodding their head, looking to the left and right, smiling, or repeating random words. Then AI is used to check the instructions were followed, to complete the check.

Active liveness can cause problems for some genuine users. For example the instructions may not be in the user's native language, and adding movement to the check increases the margin for error. Additionally, not all individuals follow instructions carefully.

Most concerningly the rapid development of deepfake AI enables bad actors to mimic the given instructions in real time. In general, the simpler the instruction(s) and the less user time required, the better the user experience and completion rate, but the easier it is for deepfake AI to spoof the instruction(s).

Passive liveness doesn't involve instructio(s) or require any action from the user. It works from a single selfie. Users no longer have to undertake head or hand movements to prove their 'liveness'. This reduces friction for users. It is simpler for people with accessibility needs and so more socially inclusive. This reduces drop off and speeds up the task of verifying genuine customers.

Passive liveness takes 1 to 5 seconds, whilst active liveness takes between 10 and 20 seconds.

A comparison between passive and active liveness

	Passive	**Active**
User feedback	Near Instant feedback	User has to wait for video validation
Time to complete	1-5 seconds on average*	10-20 seconds
Complexity and accessibility	Take a selfie - this can be either using auto face capture or at the click of a button.	User has to record a video and maintain the correct position for the duration. Language, audio input and noisy environments can be a problem for some users.
Permissions	Camera access	Camera (& sometimes microphone) access
Network traffic	Upload a selfie	Upload a video recording

*A blend of existing models and test models

How passive liveness works

From a user perspective, passive liveness means they are just taking a selfie. Behind the scenes there is a considerable amount of technology working to ensure an effective outcome for both users and organisations detecting either a genuine user or an attack.

All the below tech is all processed on average in under one second.

1. **Image capture**: in real-time, AI models ensure the best, right-sized human face image is captured for processing, by analysing the position of the face - avoiding tilted or angled capture so optimising 'head on' to camera.

2. **Image fidelity**: simultaneously, in real time, the image capture assesses the quality of the image, in terms of lighting or blurriness - giving real time feedback if required.

3. **Secure Image Capture (SICAP)**: to ensure the image is genuine, SICAP detects whether the image itself is being captured, live, from the device camera, or has been subjected to an injection attack. See page 9 for further details.

4. **Liveness assessment**: the resulting image is cropped many times to produce the inputs for the relevant neural network models and is then processed using multiple models (over 10 simultaneously).

5. **Data processing**: results from each model are assessed together to produce a response and a confidence level that the image is of a real person.

6. **Results**: a response is returned **between 1-5 seconds**. Relying parties are able to configure their checks to pass only above a given confidence level, depending on their risk profile and regulatory requirements of the territories in which they operate.

Yoti MyFace® proprietary liveness

Yoti MyFace® is a passive liveness software model that uses a single selfie image to detect presentation attacks. MyFace can be configured to also detect injection attacks. Passive liveness doesn't require any action from the user and just works from a selfie, which is processed through a sequence of deep neural networks.

The network has been trained using a variety of models that analyse images in a different way. We have invested considerable time and effort over several years to fine tune these models to optimise how they work together to create world class performance just from capturing a selfie.

As an example, one of the models, for which we have been granted (filed in August 2018), uses 3D and 2D images as training inputs to produce a model that can detect depth in a 2D image. We also have a filed a patent for a technique for detecting accessories used to try to deceive liveness and age estimation.

Independent testing: iBeta levels 1, 2 and 3

iBeta is a NIST accredited US testing laboratory providing conformance to ISO standards for liveness the relevant standard is ISO/IEC 30107-3.

NIST Level 1 involves testing against materials that could be found in a normal home or office. Materials used for testing should not cost more than $30. Masks are excluded. The tester is not an expert. To pass iBeta Level 1, the Liveness service must detect every attack and limit false negatives to less than 15%. (False negatives occur when a system incorrectly rejects a genuine user).

iBeta Level 2 involves testing against more expensive, specialist attacks, such as those using 3D printers, and resin or latex face masks. Materials used for testing should not cost more than $300. The tester has 'moderate' expertise. To pass iBeta Level 2, the service must detect 99% of attacks and limit false negatives to less than 15%.

iBeta Level 3 is the pinnacle of PAD testing. Testers can use hyper realistic masks, and develop custom spoofing attempts. Testers are experts in PAD testing and testing takes place over several weeks. To pass Level 3, a liveness service has to stop practically all attacks. *If an attacker successfully uses a spoof attack to beat the liveness model, they will continue to use the same attack to beat the model. If user successfully beats the model using a single artefact more than 5% of total attacks, the model fails level 3. 5% error rate isn't just relevant over the entirety of 1,000s of attacks, but also by each specific artefact, or attack. As an example, a high quality resin or latex masks can reach $7,000.

You can read more about iBeta Liveness testing here.

Level	Time	Cost	Expertise	Attack Error limit	BPCER Rate
L1	8 hours	$30 per artefact	None	0%	<15%
L2	2-4 days	$300 per artefact	Moderate	1%	<15%
L3	~30 days	No budget limit	Expert	5%*	<10%

Example of a widely available latex mask (c. $750)

Yoti MyFace® iBeta testing timeline - continual improvement

- In **February 2022**, MyFace® achieved iBeta level 1 compliance with 100% attack detection rate (TNR). We worked with iBeta for our level 1 compliance.
- In **February 2023**, MyFace® achieved iBeta Level 2 compliance, again working with iBeta, with 100% attack detection (TNR).
- In **January 2026**, Yoti MyFace® achieved iBeta Level 3 compliance with 100% attack detection (TNR). Yoti MyFace is the the first liveness model (passive or active) to pass iBeta Level 3.

iBeta testing levels, other testing labs and ISO/IEC 30107-3 testing standards

iBeta have been testing liveness models since 2019 and launched their level 3 testing standard in June 2025. It is currently the highest level of conformity for liveness (PAD).

iBeta is a testing lab accredited by NIST under the NVLAP (National Voluntary Laboratory Accreditation Program) to issue conformation for ISO/IEC 30107-3 applications. IBeta are recognised globally for verifying biometric models for security and performance.

As of publication date (27 January 2026),

- 108 vendors have passed iBeta Level 1
- 66 vendors have passed iBeta Level 2
- 1 vendor (Yoti) has passed iBeta Level 3

It's important to note, to achieve Level 3, a model submitted to iBeta must consecutively pass all three levels in a given test. For a Level 3 test, as Yoti had already passed L1 and L2, iBeta perform regression testing across Level 1 and Level 2 criteria, to be able to gain Level 3. This ensures models maintain levels of performance across the overall set of tests.

Each testing lab operates a different scale of conformity so levels for one lab may not be equivalent to levels for another lab.

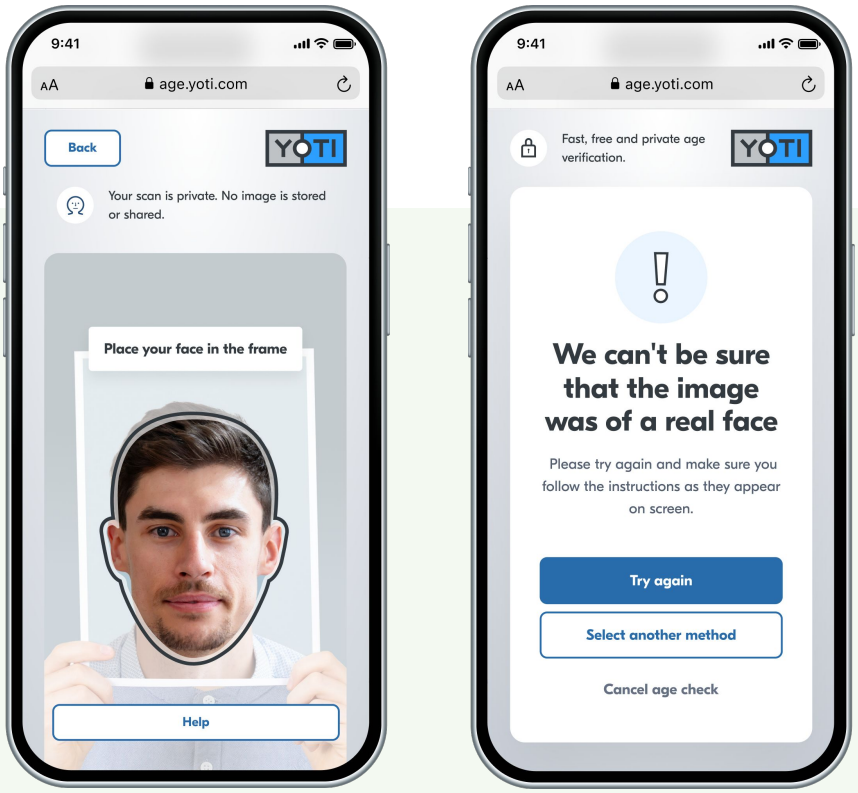

MyFace® performance in live environments

We have been using MyFace® in our facial age estimation solution, alongside other liveness providers. We complete millions of liveness checks globally every day in real world settings which in turn, enables us to gather vast amounts of information on emerging threats.

Typically, we allow users three attempts to complete a liveness check. As can be seen below this helps maximise completion rates for genuine users. This can be for a number of factors including lighting and not capturing a good quality image.

We perform liveness across a variety of environments and applications. Performance is influenced by camera quality, image size, and environmental factors.

For example, during the Home Office trial in supermarkets in the UK, we were able to gain further intelligence on how the use of our technology on self-checkouts can be adversely affected by factors such as sun glare, overhead lighting, camera quality and camera positioning in the self checkout.

This also explains the discrepancy between the mobile only first attempt rates of 94% and lower success rates on laptops or self checkouts. Compared to cameras in most laptop and self-checkouts, mobile phones have much higher quality cameras and most users are used to taking good / clear 'selfies' on their phone. Some users may not follow the active instructions successfully within the Active liveness process.

The vast majority of liveness transactions are completed on phones - our current experience is 97% of checks on mobile and 3% on desktop.

Passive and active liveness success rates improvements over time

		Passive liveness			Active liveness		
		Mar 23	May 25	January 26	Mar 23	May 25	January 26
After one attempt	Mobile	89%	92%	94%	74%	76%	72%
	Desktop	78%	90%	89%	56%	60%	57%
After three attempts	Mobile	95%	97%	98%	83%	87%	92%
	Desktop	87%	94%	94%	67%	68%	85%

Our liveness performance across regions in 2025

For facial age estimation assessment, NIST use regions as a proxy for fairness across ethnicities, with the understanding that this is imperfect. However, it is a practical way to assess performance between skin tones and ethnicities.

The below table shows liveness pass rates by region, on a mobile phone.

Region	MyFace
South East Asia	98.0%
East Asia	96.0%
South Asia	95.0%
East Europe	97.0%
East Africa	93.0%
West Africa	93.0%
Other	99.0%

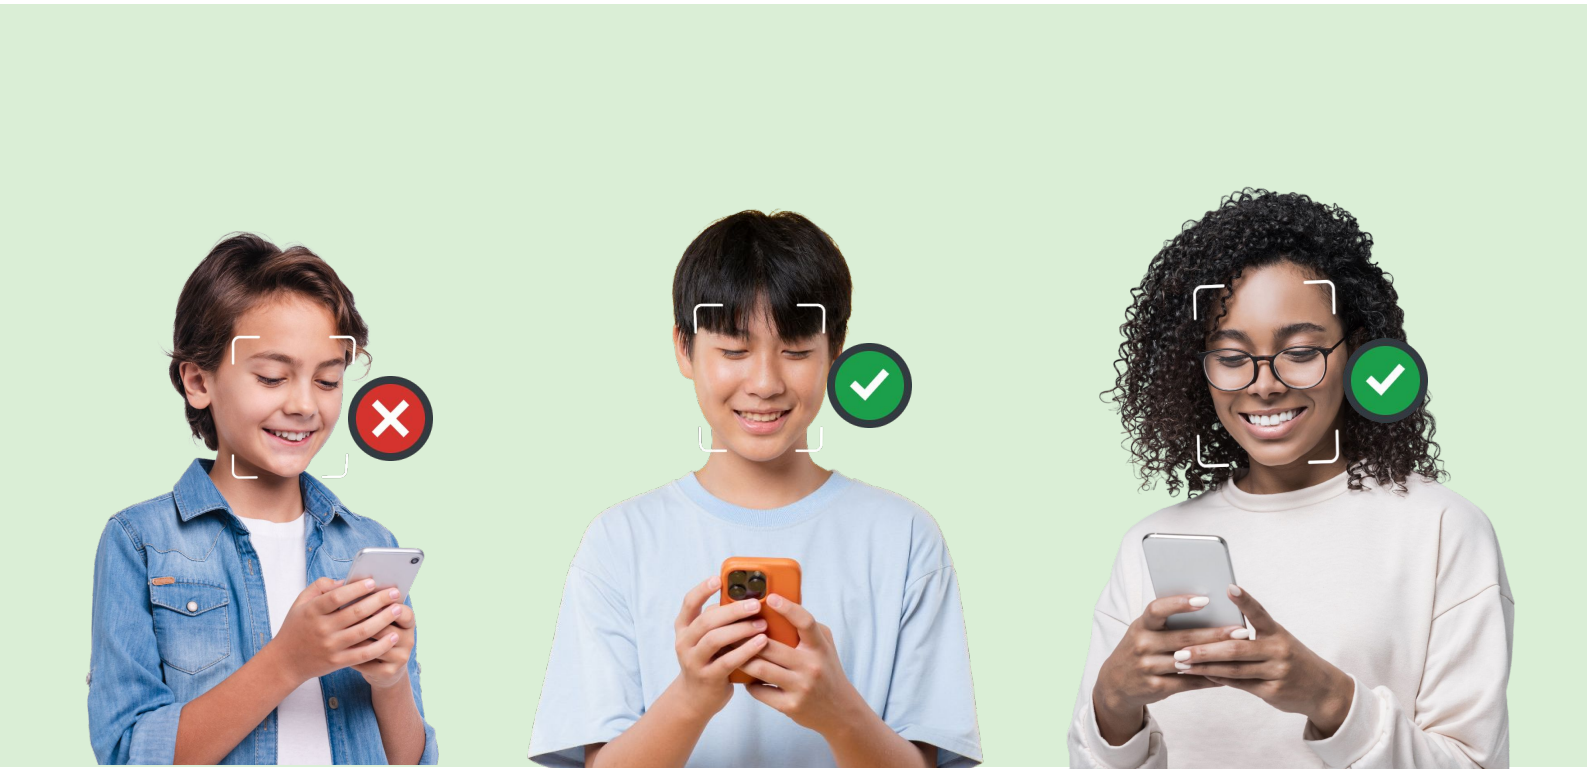

The threat of injection attacks

An important question for any new technology, is how secure is the process? Can it be spoofed? Can bad actors hack into the system to override checks, images or results?

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

Data capture attack threats

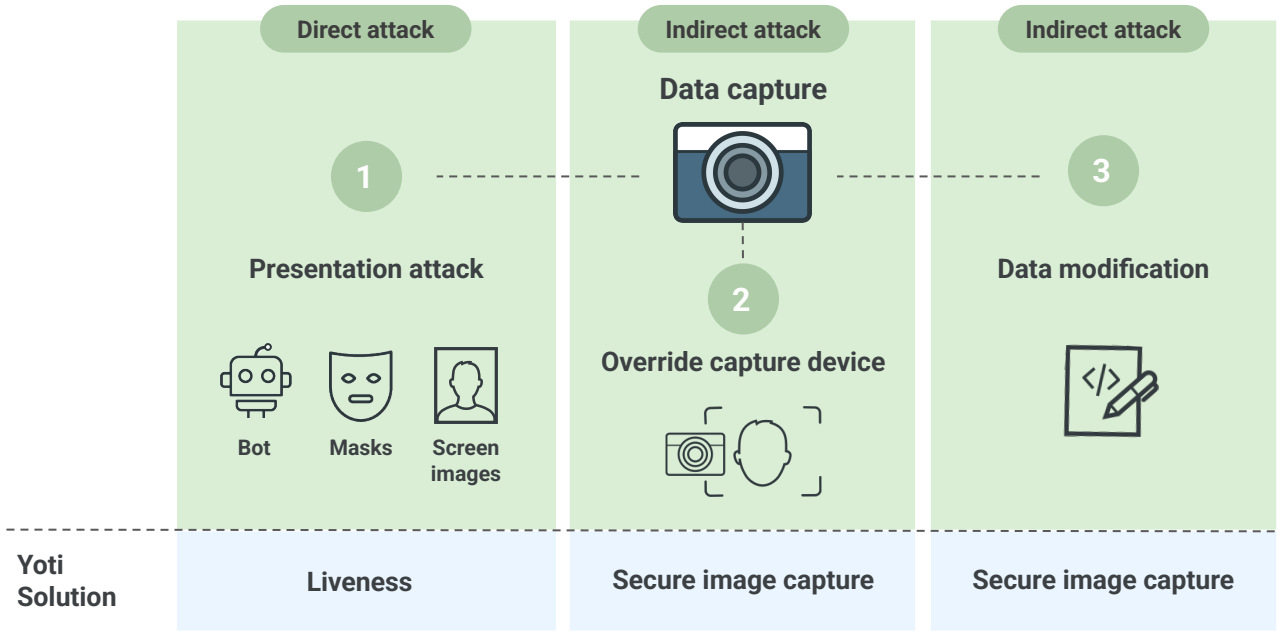

Step 1 as a direct spoofing attack - an attempt to present an image, mask or video, often called a presentation attack. This is an attempt to spoof a check by appearing older or to be another person. To overcome this we use our MyFace® liveness technology. This ensures that the person undertaking the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

Steps 2 & 3 is 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 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.

In 2020, Yoti recognised the upcoming risk of injection attacks and started to develop a solution, called SICAP (**S**ecure **I**mage **CAP**ture). Our patent for SICAP was granted in November 2023, which makes injection attacks significantly more difficult for imposters.

Our experience of rising fraud trends and injection attacks

In 2025, we witnessed an increase in the number of deepfake attacks, or injection attacks during age and identity verification checks. We detected a total of 3.2 million injection attacks in 2025.

In absolute terms, this is a significant rise in the total number of attacks we have detected as our services grew over the year. We now perform over 8 million checks per week across all our services. With the introduction of various regulations globally, companies have been obliged to implement more robust age or identity checks for their users.

We experienced a peak in August with 527,013. This coincided with the implementation of the Online Safety Act in the UK.

The rise in adoption and the improvement in output of AI models enables more individuals to create deepfakes or images of real people to spoof online checks. Message boards and chat rooms openly share tips and tricks of how to spoof checks. The barrier to entry and cost of creating high quality fake images is rapidly diminishing.

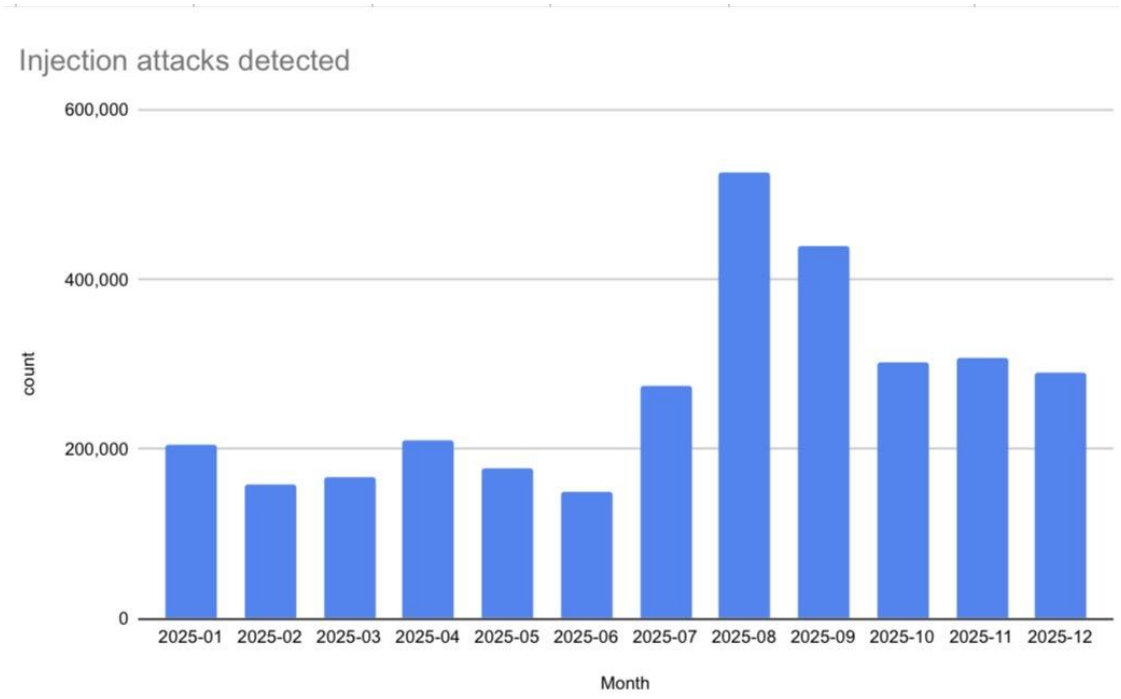

The importance of model efficiency

Using and deploying small, specialised models that each solve a clearly defined problem. gives our customers stronger security, better privacy outcomes, faster performance and greater confidence in how decisions are made.

This layered approach helps avoid reliance on any single model. If one signal is inconclusive, others can still provide strong evidence. The result is greater resilience against spoofing, bots and deepfakes. Examples of these checks include:

- Light refraction analysis - helping to detect masks or printed images
- 3D background inference from 2D imagery - helping identify uploaded or printed images

We process millions of checks every week, operate independent testing programmes and run an ethical hacker bounty scheme. This creates a real-world feedback loop. When new threats appear, we are able to detect them early, and can develop and deploy technology to combat those threats.

As a result of this approach, a check time is extremely fast and we can quickly update models. Processing power required is very low, minimising compute power and therefore cost. Importantly model size can be configured to be small enough to run entirely on-device.

That could be a mobile phone, self-checkout terminal or a vending machine. This means checks can work offline and can mean that personal data doesn't need to leave a device.

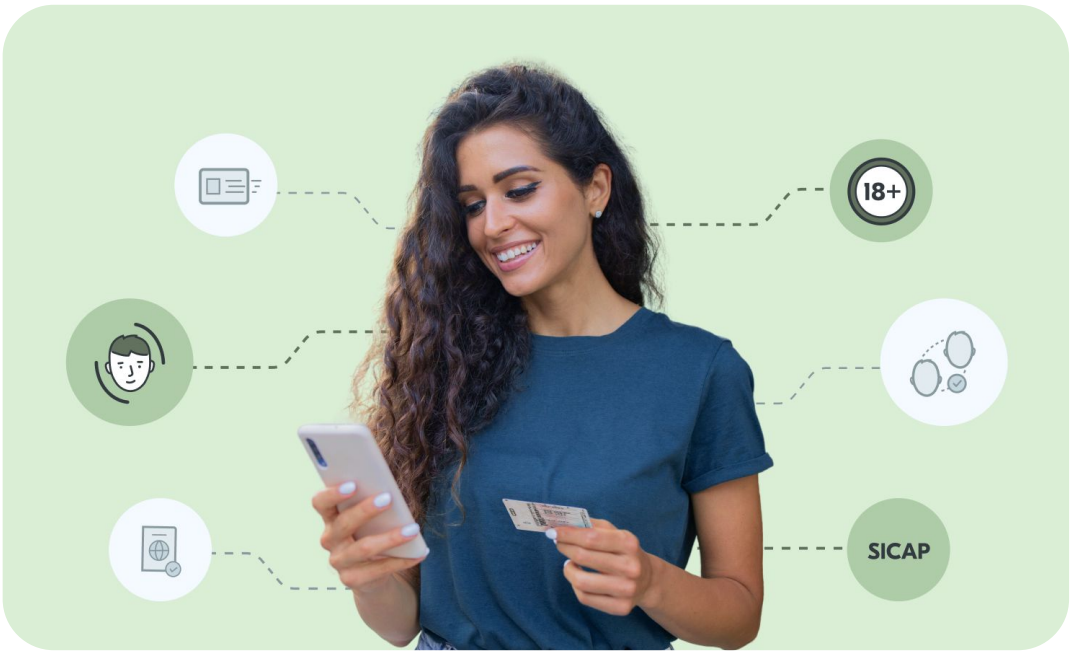

Memberships, associations and accreditations

To learn more about Yoti MyFace® liveness and our AI Solutions please **get in touch**.