iProov

Demystifying Biometric Face Verification

Contents

Introduction

Ostensibly, identity assurance can be seen in a binary way. I.e., do we have one-time passcode (OTP) authentication or do we not? Yet, not all forms of identity assurance are the same.

Traditional identity assurance technologies rely on either possession, like a device, or knowledge, such as a password. Passwords can be stolen/shared, and devices can be compromised, meaning these methods no longer provide the requisite defense against today's threat landscape, nor do they meet user demands for convenient, low-friction digital experiences. Organizations are therefore migrating to the third factor of authentication, inherence, such as face biometrics as proof of identity.

This has many advantages. It can verify the identity of a remote user by tying them to their trusted government-issued ID (passports, national ID scheme). This is known as biometric face verification. Liveness detection is a key component incorporated into biometric face solutions, that detects if a user asserting their identity is a real, 'live' person, and not a presented artifact (such as a photo or mask), or generative AI (deepfakes or other synthetic imagery).

Not being able to determine whether the individual is 'live' and authenticating in real-time, leaves biometric technology vulnerable to spoofs – or 'biometric attacks'. These attacks range from rudimentary masks that change the threat actor's appearance, to highly sophisticated generative AI-created imagery digitally injected into the data stream.

For mission-critical use cases, selecting a technology that is not resilient or adaptable to these evolving threats can lead to severe consequences. Among the most damaging are: monetary and reputational loss, huge fines, and media scrutiny. Identity fraud, particularly synthetic identity fraud, is estimated to account for $2.42 billion in illicit funds obtained in 2023.[1] Meanwhile, global law enforcement deems cyber-enabled financial crimes, such as money laundering, as the greatest threats now and into the future.[2]

The biometric threat landscape is evolving. Threat actors are continually advancing how they attempt to circumvent identity verification technology. As such, biometric face verification as a whole must evolve to adapt to emerging attack vectors.

Biometric solutions are not created equal, they vary significantly in their performance, accessibility, inclusivity, and protection against attack vectors. Selecting the correct biometric solution for the appropriate use case, while providing a positive user experience, is paramount.

This eBook outlines the different types of biometric face verification technology on the market, the various ways they are deployed, their usability, as well as their strengths and weaknesses. It's intended to enable organizations to select the correct solution for their use case.

[1] *The Aite Group*
[2] *INTERPOL*

Liveness

Live-ness

The specific detection of whether a sensor is viewing a live biometric – as opposed to a recording, picture, or another non-living spoof – is commonly known as liveness.

– The Biometric Institute

Liveness Detection in Biometric Face Verification

Face matching

Liveness

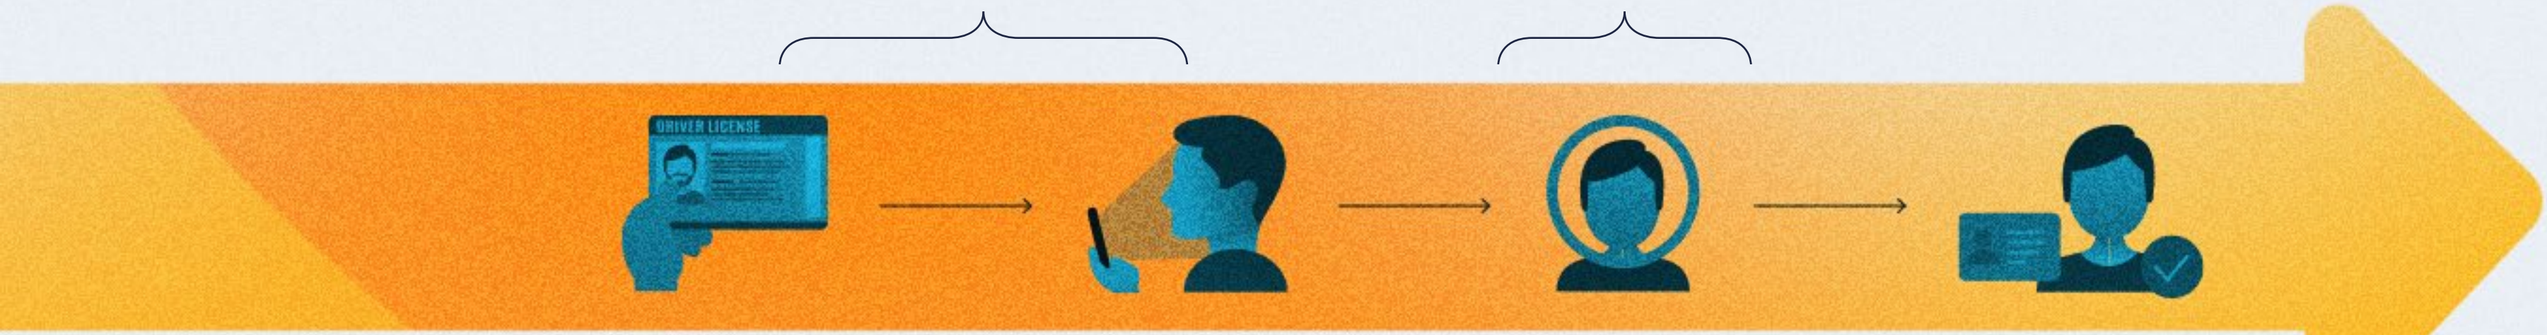

Onboarding:

An individual captures information from their trusted government-issued ID

They take a selfie. Face matching determines if the two images are the same

The system checks whether the image capture of the individual is ‘live’ and not a spoof

Biometric face verification check complete

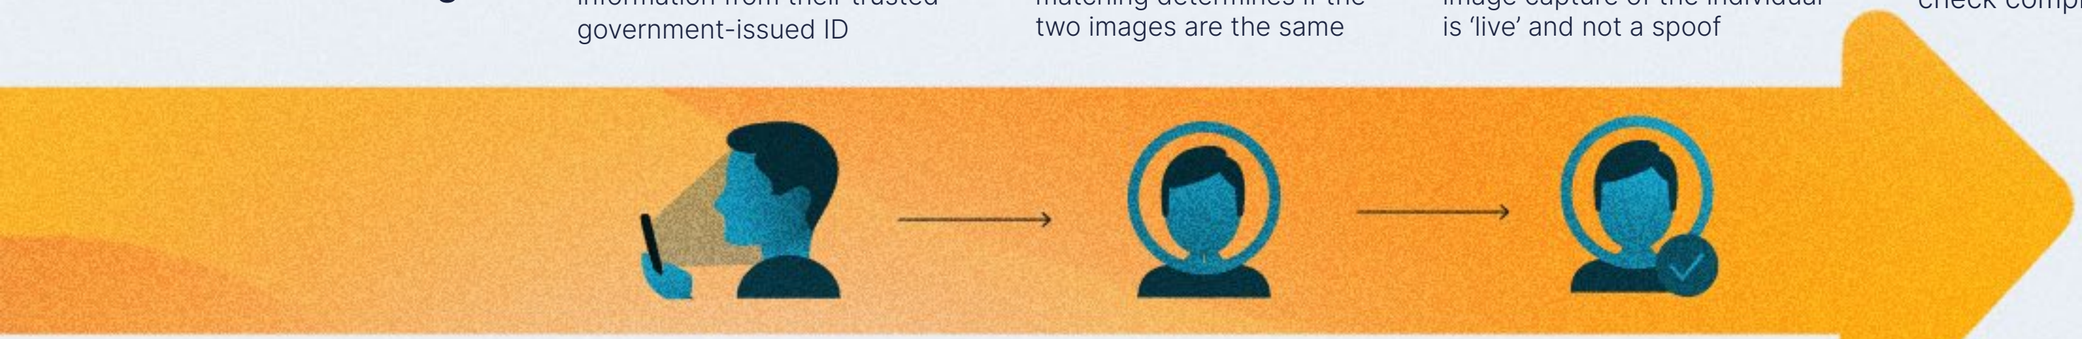

Authentication:

The individual scans their face to prove it’s the same person who onboarded

The system checks whether the image capture of the individual is ‘live’ and not a spoof

Biometric face verification check complete

Face matching

Liveness

The robustness of the solution an organization requires depends on its specific use cases and threat profile. A commercial bank, for example, may need a more secure solution than a gaming application.

Biometric Face Verification vs Face Recognition

Face Verification

Verification is when the user:

- Collaborates **with** the process
- **Does** directly benefit
- **Is** assured their privacy **is** protected

Face Recognition

Recognition is when the user:

- Has **no** knowledge or control over the process
- Does **not** directly benefit
- Has **no** control over their privacy

Concerning Other Biometrics

Non-Face Biometrics

While other biometric modes – such as fingerprint and iris – can incorporate liveness detection, we focus on face biometrics. Other modes require special sensors, impairing their utility to users with devices that do not have these sensors. These modes are mostly used for the re-authentication of returning users, rather than identity verification at onboarding.

Client-Side Biometrics

Likewise, this eBook does not cover on-device biometrics, like Apple's Face ID. While on-device biometrics can contain liveness detection and provide a high level of user convenience, they cannot assure the user is who they claim to be – only that it is the device owner re-authenticating.

Face Matching

Face matching is the ability of biometric systems to ensure two images, such as a selfie and a government-issued ID, match. Face matching is a component of biometrics, but it should not be misinterpreted as liveness detection. Face matching cannot assure the user is 'live' and provides no presentation attack detection (PAD) - presented artifacts (such as a photo or mask), or generative AI mitigation like digital injection attack detection (DIAD). As such, we have omitted it from the eBook.

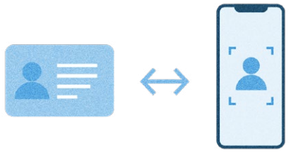

Understanding Biometric Attacks: Presentation vs Digital Injection

Presentation Attacks

The threat actor attempts to circumvent biometric liveness detection technology and gain unauthorized access by presenting an ‘artifact’ to the camera.

Artifacts can include a printed image, silicone mask, replay video, or deepfake presented on screen.

Presentation attacks are limited in scale and have progressed little in sophistication. Many biometric liveness vendors have been accredited for **Presentation Attack Detection (PAD)** by NIST and iBeta.

Digital Injection Attacks

The threat actor attempts to bypass the device camera and inject imagery, such as generative AI-produced synthetic imagery, directly into the data stream.

Synthetic imagery can range from simple images to highly sophisticated deepfakes, such as face swaps, re-enactments, and imagery developed using Generative Adversarial Networks (GANs).

Digital injections are highly scalable, challenging to detect, and have outpaced presentation attacks in terms of sophistication and frequency.[3]

There are no industry-wide benchmark testing or accreditation for **Digital Injection Attack Detection (DIAD).**

An illustration of a **face swap**. Face swaps are advanced synthetic imagery that are often digitally injected. The threat actor merges the traits of one face, such as motion, with the features of another.

[3] *iProov Threat Intelligence Report 2023*

Understanding Bias and Accessibility in Biometric Face Verification

Large-scale organizations, like governments and banks, have a responsibility to ensure their services are accessible to the maximum number of people. No matter how secure the biometrics deployed, if they exclude people based on skin tone, socioeconomic status, digital literacy, or physical or cognitive abilities, their value is greatly diminished.

Any system that learns to distinguish between facial features, skin tone, or texture based on appearance is potentially prone to overfitting to subsections of face types. When this happens, the technology performs inconsistently across some face types than it does for others. To mitigate this bias, systems must ensure the datasets used to train the algorithms are balanced according to age, gender, and skin tone.

However, bias does not only arise from different face types. Causes of bias also include differences in camera types, user behavior, and environmental conditions. These different factors affect certain groups more than others. For example, people in lower-income regions may have lower-quality device cameras. Systems must perform operational testing to ensure they perform equally well for everyone.

Ensuring Accessibility

Biometric face verification solutions can exhibit high accessibility by complying with industry standards.The Web Content Accessibility Guidelines (WCAG 2.2 AA) is the best-practice standard for a range of digital experiences. It ensures accessibility for all, irrespective of age, literacy, language, cognitive ability, disability, or other constraints. Organizations that have a large and diverse user base should seek a vendor that complies with WCAG standards. In certain regions, like the EU and the UK, governments offering digital services must comply with WCAG 2.2 AA and Section 508 (US) as the minimum standard for web accessibility.

Deployment: API vs SDK

Biometric face verification solutions are either deployed via a headless API (Application Programming Interface) endpoint to which imagery can be submitted directly for processing, or as an SDK (Software Development Kit), a more complete solution that handles the collection and pre-processing of imagery before it's sent for processing.

Key differences between APIs and SDKs regarding biometric face verification include user feedback, the imagery captured, and how they are integrated. With an SDK-deployed solution, the technology vendor can provide user guidance, such as informing them to move their face clearly into frame or to correct problematic lighting.

This guidance can ensure that the correct imagery – the data needed for the liveness check to make an assured pass/fail decision – is captured. Failure to do this will mean the verification attempt will fail and the individual must try again. Without the correct imagery, the user will not pass liveness verification. Although an individual can try again, it can lead to frustration and in some cases, abandonment.

With APIs, the organization can build in user guidance and other functionality, but this will likely slow the integration process. SDKs also have some security advantages over APIs. They provide the data needed to enhance the vendor's pass/fail decision, such as whether the user's device has been compromised, such as jailbroken or rooted, and perform other integrity checks against the source of the imagery.

Nevertheless, there are some scenarios where only an API-deployed solution can work. These include embedded systems that do not support an operating system for which the vendor has an SDK.

Plus, an SDK can increase the organization's app size. While the SDK size of some vendors is as little as 2MB, which will unlikely create an issue for most organizations, other vendors' SDKs can reach 10MB, which may be a sticking point for organizations sensitive to app size. SDKs can also make a difference to the app's interface. This may be an issue for organizations that want to build the entirety of UI themselves.

Strengths and Weaknesses

SDK

Strengths

+ User feedback increases pass on first go, **improving completion rates.**
+ Non-image data can be captured, enhancing the **vendor's decision making.**
+ Prebuilt functionality **speeds up integration and time to market.**

Weaknesses

− Consider the organization's app size. Size of integration can **impact the performance,** especially in a region with low network connectivity.
− It can be different from the organization's user interface, so considerations will need to be made around **user experience.**

API

Strengths

+ **Compatibility:** In some cases, only an an API can be integrated.
+ Organization has **100% control over the user interface.**
+ **Lightweight:** An API will not increase the organization's app size as much as an SDK.

Weaknesses

− No real-time user guidance and feedback. This may **reduce completion rates** and impair user convenience.
− Organization must build functionality and user interface on top of the API. This may **slow integration and lengthen time to market.**

Recommendations

To speed up integration and ensure high completion rates, organizations should choose an SDK-deployed solution. However, in certain cases, whereby an SDK-deployed solution simply would not work, organizations may need to choose an API.

Hosting: Cloud-Based vs Server-Side

Cloud-based solutions are becoming ubiquitous. According to Gartner, over 95% of new digital workloads will be deployed and controlled via cloud platforms by 2025 – up from 30% in 2021.[4]

The pervasiveness of cloud-based solutions extends to biometric face verification. Organizations have a choice: to host the solution locally and manage it on their infrastructure or to have it hosted and managed on the vendor's servers.

GDPR has singled out biometric data as a 'special category' of personal data, requiring extra protection.[5] Considering this, on-premise solutions can seem ostensibly appealing as the organization has 100% control over the biometric data.

However, this compromises security. A deployed server-side solution is open to reverse engineering, effectively making defenses more susceptible to being easily spoofed. Furthermore, the organization must accept to make any algorithm and operating system updates – they are not automatically applied. This can slow the process and leave the solution outdated and vulnerable as the threat landscape evolves.

The threat landscape changes rapidly, meaning that even day-old deployments can be vulnerable to the latest attacks. Cloud-based solutions allow for real-time defense to ensure the technology remains resilient to new threats.

Protecting Biometric Data

To protect the privacy of users, many cloud-based vendors pseudonymize the biometric data gathered. Privacy firewalls create a biometric template of the user's face. These templates can be processed by machines but are ineligible to humans.

Some cloud-based vendors also enact data minimization, whereby they process the user's face without storing any accompanying personal identifiable information like names, rendering the data worthless if leaked or stolen.

[4] *Gartner*
[5] *Information Commissioner's Office*

Strengths and Weaknesses

Cloud-based

Strengths

+ **Real-time defense**, which can ensure resilience to the evolving biometric threat landscape.
+ Some vendors can provide **active threat monitoring,** meaning they can gather intelligence on the threat landscape and adapt defenses accordingly.
+ **Flexible commercial models** enable the organization to scale solution as needed.

Weaknesses

− **Privacy:** Organization must ensure that the vendor has robust privacy controls in place, such as a privacy firewall and data minimization.

On-premise

Strengths

+ Organizations has **100% control and responsibility for the biometric data.**

Weaknesses

− **Vulnerability:** Slows deployment of defense updates and algorithm changes.

− **Resource intensive:** Organization must have an in-house team and hardware to manage the infrastructure.

− **Costly:** Organization must pay the cloud server costs.

Recommendations

Due to the threat landscape, organizations should only use on-premise solutions in low-risk scenarios.

Cloud-based solutions ensure technology remains resilient against the rapidly evolving biometric threat landscape.

When choosing a cloud-based solution, ensure that the vendor adheres to strict privacy regulations, such as GDPR and SOC 2 Type II.

Challenge-Response: Active and Passive

Many biometric solutions incorporate a challenge-response mechanism to ensure that the user is a ‘live’ person and not a spoof attempt.

Active systems vary in their ability to thwart attacks. Single-action challenge responses include asking the user to blink or smile. These are predictable and vulnerable to scaled injection attacks, as seen in the iProov Biometric Threat Intelligence Report 2023.[6]

Variable-action active solutions ask the user to do something different each time. This can include turning the head in different directions or reading out loud a sequence of characters. More difficult to reverse-engineer, variable actions provide higher security than single-action solutions.

However, active solutions can impair the inclusivity and accessibility of a solution, as not everyone can perform these actions. Plus, the authentication sequence is exhibited, giving information to the threat actor to reverse engineer.

Passive solutions, conversely, hide the authentication process from the attacker and can improve inclusivity. Passive challenge-response can ensure a high level of assurance without impeding inclusivity and accessibility. They provide the highest assurance that the user is not only ‘live’ but also authenticating in real-time.

Active Authentication

The user is asked to perform actions, such as blinking, smiling, or reading characters aloud.

Passive Authentication

The user does not perform actions. They look at the camera and the authentication process is complete.

Passive Challenge-Reponse

Challenge-response biometrics can either be active or passive. Either the technology does something different each time (passive) or the user is asked to perform something different each time (active).

With passive challenge-response the mechanism is randomized by the technology itself (not by actions performed by the user), making the authentication process unpredictable, impervious to replay attacks and highly challenging to reverse-engineer.

[6] *iProov Biometric Threat Intelligence Report 2023*

Strengths and Weaknesses

Active Authentication

Strengths

+ It's clear that the authentication process is taking place, **reassuring the user.**
+ Vendor can make the variety of actions more complex, providing higher assurance than single-action solutions

Weaknesses

− **Impairs inclusivity.** Asking the user to perform actions risks excluding people with physical or cognitive disabilities.

− **Risks reverse-engineering**. The authentication process is made clear to threat actors. They have the necessary information to potentially reverse-engineer the system.

− Advanced synthetic imagery attacks, such as **face swaps, can perform actions in real-time**, circumventing active systems.

Passive Authentication

Strengths

+ Not soliciting a user response ensures the solution is **accessible to the maximum number of people**, irrespective of physical or cognitive ability.
+ There are no actions for synthetic imagery attacks to perform, making it **more secure against generative AI.**
+ Authentication process is hidden from the attacker, **mitigating the risk of reverse-engineering.**

Weaknesses

− If the experience is too passive, it may not be clear to the user that the authentication process is taking place. **This can be disconcerting.**

Recommendations

For high-risk use cases, such as user onboarding, challenge-response biometrics are essential to establish that a remote user is who they claim to be and therefore defend against the most advanced attacks. Organizations that deliver services to a large and diverse user base should deploy passive solutions to ensure maximum inclusivity and accessibility.

The Biometric Liveness Ecosystem

Liveness Type	Testing Parameters
API Single-Frame	NIST FRVT PAD
API Multi-Frame	NIST FRVT PAD + iBeta
SDK Multi-Frame	iBETA
SDK Multi-Frame	National / Red Team
Passive challenge-response	Testing

Liveness Components	
Imagery Capture:	Single-Frame or Multi-Frame
Deployment:	API or SDK
Attack Detection:	PAD or DIAD
Hosting:	Cloud-Based or On-Premise
Challenge-Response Mechanism:	Active and Passive

Biometric face verification technologies can use a combination of the different components. Whereas a single-frame solution may be deployed via an SDK and hosted on-premise, another may be a multi-frame solution, delivered via an API and cloud-based.

Single-Frame

Single-frame biometric solutions determine whether a user is who they claim to be and is ‘live’, from a single image. Vulnerable to simple attacks, such as doctored images, they provide little assurance against threat actors.

Multi-Frame

Multi-frame biometric solutions determine whether a user is who they claim to be and is ‘live’ from multiple images. Circumventing multi-frame solutions is exponentially more challenging as the attacker needs to forge a video, or use generative AI to create moving synthetic imagery.

Accreditations and Standardized Testing

For an unbiased assessment of a biometric vendor's security, organizations can look for what accreditations they have. Testing and accreditation for biometric security vary.

The **US National Institute of Standards and Technology (NIST)** has developed the **Face Recognition Verification Test (FRVT)** to assess the accuracy and Presentation Attack Detection ability of biometric systems.

Currently, NIST FRVT only tests and accredits API-deployed single-frame solutions, meaning SDK-deployed multi-frame solutions cannot apply.

SDK multi-frame solutions can be tested and accredited for Presentation Attack Detection by **iBeta** to the international standard, ISO 30107-3. iBeta is accredited by NIST NVLAP as an independent test lab.

While Presentation Attack Detection is accredited by NIST FRVT and iBeta, no such testing exists for the detection of digital injection attacks, the greater threat.

Vendors can show proven defenses against current and future biometric attacks to national-level security standards by undergoing testing from security-conscious government agencies.

Summary

As organizations continue to recognize face biometrics with liveness detection as a secure, usable means to verify and authenticate users, the technology will become further ingrained into the digital identity ecosystem.[7]

The impact of this is two-fold.

1. More biometric liveness providers will enter the market, and an array of different solutions will be available.

2. The threat landscape will develop further aided by access to generative AI technologies.

Biometric solutions, therefore, must continually advance and adapt to the evolving threat landscape, developing defenses to thwart new and emerging attack vectors.

It's imperative that organizations understand the different types of biometric technology available and which type is best suited to their use case and risk profile.

Recommendations

- Organizations with the highest online risk exposure, such as financial services and governments, need the highest level of identity assurance. To achieve this, challenge-response biometrics and a cloud-based vendor are essential.
- Choose vendors that have robust processes to mitigate bias in all its forms to ensure maximum equality and inclusivity.
- To ensure maximum accessibility and achieve high performance and completion rates, an SDK-deployed passive authentication solution is needed.
- For low-risk use cases, such as an existing user re-authenticating for a gaming site, a lower identity assurance and lower-performing solution may suffice. This can include an API-deployed single frame, hosted on-premise.
- Organizations with a diverse user base should ensure their chosen vendor delivers accessibility. Choose a vendor that complies with WCAG 2.2 AA and Section 508 standards.

[7] *Goode Intelligence*

iProov is used by leading organizations worldwide to reduce the risk of identity fraud.

Government clients include the U.S. Department of Homeland Security, the UK Home Office, the UK National Health Service, and GovTech Singapore.

Financial services clients include UBS, ING, Rabobank, and Knab.

contact@iproov.com

iproov.com

IPThrReen-USLE01/23