


Whitepaper

Liveness detection

in biometrics


is essential for mobile authentication and onboarding



A W A R E

Table of Contents

- 01** Introduction
- 02** Liveness detection techniques:
UX and opaqueness are key
- 03** Types of presentation attacks:
false match vs. false non-match
- 04** Further advances: machine learning,
multimodality, and browser-based capture
- 05** Certification of liveness detection products
- 06** Liveness detection in biometrics is essential



It's essential to apply robust liveness detection when using facial recognition for unattended mobile applications.

Introduction

Mobile biometrics offer banks and other businesses a way to more effectively verify the identities of their new and existing customers. Mobile onboarding allows a prospective banking customer to open an account without visiting a branch. Biometrics make this identity proofing process more robust and secure. New customers can then use their biometrics in place of passwords for more secure and convenient mobile authentication.

Facial recognition is a particularly useful biometric modality for mobile onboarding and authentication:

- Nearly all mobile devices have built-in cameras that support it.
- The user experience of capturing a “selfie” is exceptionally intuitive and convenient.
- Face recognition is 99.7% accurate and getting more so by the year, according to NIST.

While facial recognition is an ideal biometric modality for mobile applications, it is also particularly vulnerable to “presentation attacks.” A presentation attack is an attempt by a fraudster to intentionally defeat biometric security measures by presenting non-live biometric data.

To do so, a fraudster might use a printed or digital photograph, video, or mask to either impersonate a targeted victim or to assert a false identity. Such an attack is also called a “spoof.”

Facial recognition algorithms can be spoofed with relatively little effort due to the wide availability of facial images throughout the internet.

For this reason, it's essential to apply robust liveness detection when using facial recognition for unattended mobile applications. There are at least two approaches to mitigating the risk of facial presentation attacks:

- Liveness detection algorithms: Analyze facial images to determine whether they are of a live human being or a reproduction.
- Multimodal biometrics: Add a second biometric modality, such as voice.

Without such spoof attack prevention measures, facial recognition-based biometric security is not sufficiently secure.



Passive and hybrid liveness detection approaches have an advantage over active methods of requiring little or no user interaction and therefore present a more frictionless user experience.

Liveness detection techniques: UX and opaqueness are key

Liveness detection algorithms can be categorized as follows:

Active liveness detection. This entails a challenge and response; a user may be prompted to blink, smile, or move their device during a facial recognition capture. Users are fully aware of the liveness detection measures being applied.

Passive liveness detection. This happens in the background and relies on algorithms that can identify and assess those artifacts in an image that indicate its content, including masks, cutouts, skin, texture, borders, and other indicators of a false representation of a user's face. The process is opaque to the user, making it more difficult for a fraudster to learn how to circumvent it.

Hybrid. A hybrid method does not require user interaction but nevertheless is observable by a fraudster, making it potentially more vulnerable to circumvention than a purely passive approach.

Ideally, liveness detection techniques are implemented without degrading the user experience. Passive and hybrid liveness detection approaches have an advantage over active methods of requiring little or no user interaction and therefore present a more frictionless user experience. The opaqueness of truly passive approaches is beneficial in that no clues are given that instruct how to defeat the liveness detection measure.



In a mobile authentication scenario, the goal of the fraudster is always to match the genuine biometric reference sample of the victim.

Types of presentation attacks: false match vs. false non-match

Presentation attacks can also be categorized by the type of false result that the fraudster is aiming to achieve. When a fraudster successfully impersonates his victim in a one-to-one biometric comparison, this can be called a false match. When a fraudster avoids detection in a watchlist search or duplicate search, this can be called a false non-match.

In a mobile authentication scenario, the goal of the fraudster is always to match the genuine biometric reference sample of the victim. This could allow the fraudster to access that victim's account. In mobile onboarding, the fraudster might also try to impersonate a victim using a false match presentation attack.

In doing so, they can falsely use their victim's identity to open a new account. An example is when a biometric match to a driver's license photo is part of the onboarding process.

But in onboarding there is another category of risk in play, which is a false non-match attempt.

In this case, the fraudster attempts to register a new account with a facial image that cannot be used in a biometric watchlist search or duplicate search.

For example, a fraudster might attempt to register an image of the following:

- A random person or celebrity.
- A picture of themselves that is not biometrically searchable (e.g. by obstructing or distorting facial features, applying makeup, wearing dark glasses, etc.).
- A person who does not actually exist.
- A non-human/non-person.

Such false non-matches could allow a fraudster to register one or multiple fraudulent accounts that are difficult to detect. For this reason, it is critical for mobile onboarding to use liveness detection algorithms that can detect these types of spoofs that don't necessarily match to a targeted victim.



Just as with matching algorithms, machine learning is a promising approach for improving liveness detection.

The technique offers a frictionless user experience, while its opaqueness offers better security.

Further advances: machine learning, multimodality, and browser-based capture

Machine learning

Just as with matching algorithms, machine learning is a promising approach for improving liveness detection.

The technique offers a frictionless user experience, while its opaqueness offers better security. It requires the training of algorithms with a variety of spoof data, such as paper images, digital images, digital video, masks, etc.

Multimodality

Facial recognition can also be enhanced with a second modality, such as voice. Face and voice used together improves the performance of both the biometric matching and the liveness detection by an order of magnitude. Applying two modalities, with each leveraging liveness detection, make it difficult for fraudsters to defeat biometric systems.

Analogous to facial recognition, active and passive liveness detection can also be employed for voice authentication:

- Active: The user must speak a randomly generated phrase or number.
- Passive: Algorithms analyze the sample to detect artifacts of recorded or synthetic voice.

Browser-based capture

Mobile onboarding is made even more convenient if a prospective customer does not need to install a mobile app before applying for an account. Instead, they can simply visit a mobile-optimized web page in their mobile browser. For this reason, the ability to capture facial images and perform liveness detection in a browser is increasingly in demand.

Testing involves the fabrication of a few spoof samples and their use to try to defeat the biometric security features of the product.

Certification of liveness detection products

There are international standards (e.g. ISO 30107-3) that define best practices for assessing the performance of commercial liveness detection products. The standards are used by accredited independent laboratories to design and perform tests and issue test reports. They do so for a fee typically paid by the product supplier. Testing involves the fabrication of a few spoof samples and their use to try to defeat the biometric security features of the product.

The test reports offer a valid datapoint in the evaluation of liveness products, but they can also lead to a false sense of security. There are several reasons for this:

- The tests can be performed using different settings than how the technology must be deployed in production. For example, the thresholds used in testing might result in false-positive rates that are too high for full-scale deployment.
- The tests are not sufficiently representative of the diversity of a real-world deployment; they should ideally aim to assess the impact of variations in face types, lighting, devices, and other factors.
- The tests may not be adequately rigorous for important attack modes.

- The tests may not represent performance with the relevant architecture or configuration, e.g. FIDO Certified vs. server-based, single modality vs. multimodal, etc.
- Some tests provide only imposter match rates (IAPMR), which are not relevant for onboarding processes not involving a 1:1 biometric match. For onboarding, assessment of Attack Presentation Classification Error Rate (APCER) is warranted. This measurement indicates success in detecting all spoofs, not just those that biometrically match a target victim¹.

All security mechanisms can be ultimately defeated with enough effort, and liveness detection is no exception. So it is critical for product evaluators to be intimately familiar with whatever vulnerabilities they may have that are relevant to their particular use case, security threats, and customer base.

This information is not available in a certification report. In fact, a perfect score on a test should possibly be seen as an indication that the test is not sufficiently rigorous. A liveness detection product survey should always include internal testing, ideally under conditions that reflect the targeted use case and customer base.

¹ IAPMR stands for Imposter Attack Presentation Match Rate. It is a metric that reflects the susceptibility of a biometric authentication solution to spoofing. If during testing, a spoof can 1) successfully bypass liveness detection mechanisms and also 2) successfully match to the trusted biometric reference sample, then a higher IAPMR will result. It follows that if a spoof does not biometrically match, then it is not considered a successful spoof attempt. APCER stands for Attack Presentation Classification Error Rate. This metric measures the reliability of a liveness detection algorithm in determining whether a spoof sample is truly a spoof, or actually a live image. Biometric matching is not reflected in this metric.

Liveness detection in biometrics is essential

Thanks in large part to machine learning, facial recognition algorithms have become extremely accurate. Liveness detection algorithms are following a similar trajectory.

Commerce is possible only when a person can instill confidence in their claimed identity to their counterpart. Humans have done this for thousands of years through visual recognition of our faces. Biometric facial recognition algorithms allow the task to be performed by computers with extreme speed and accuracy. The technology can be used to establish trusted relationships between businesses and their customers across a digital channel.

But today's technological advances can also be exploited by fraudsters. Modern cameras, displays, 2D and 3D printers, and even computer animation can be used to convincingly simulate a live human face. Liveness is foundational to biometrics; they are secure only when they can confidently demonstrate uniqueness in the present tense.

Thanks in large part to machine learning, facial recognition algorithms have become extremely accurate. Liveness detection algorithms are following a similar trajectory. But as with matching algorithms, liveness detection must work reliably for everyone, regardless of their physical appearance, the device they use, or the lighting environment in which they use it. Solutions that pair biometric matching performance with robust liveness detection will deliver on the promise of biometrics to make mobile commerce more convenient and secure.

Contact Aware for more information about biometric liveness detection and matching algorithms and their Knomi™ solution for mobile biometric onboarding and authentication.



A W A R E

www.aware.com/contact