Biometrics Institute Concepts and Solutions Report

Biometrics: Keeping it Real

March 2026

Silver Jubilee edition celebrating 25 years of Responsible Biometrics

Entrust: Is AI outsmarting you? Deepfakes, Injection attacks, and the next generation of responsible biometrics

1. Introduction: when realism is no longer proof

Biometric systems were long built on a simple but powerful assumption: if a face looks real and is captured live by a camera, it represents a real person, physically present at the time of verification. Advances in generative artificial intelligence have disrupted this assumption. In 2026, deepfakes have reached a level of realism where visual inspection—by humans or traditional computer vision models—is no longer a reliable indicator of authenticity.

Deepfakes are no longer marginal. Industry data now shows that approximately one in five biometric fraud attempts involves a deepfake[1], reflecting both the accessibility of generative tools and the professionalisation of fraud operations. The central question is no longer whether deepfakes can be detected visually, but how systems can reliably prove real-time presence and capture integrity.

2. What deepfakes look like in real identity flows

In operational identity systems, deepfakes typically take three forms:

- **Face swaps**, where synthetic faces are overlaid onto real heads in a live or recorded video stream.
- **Fully synthetic media**, generated entirely by AI models and not corresponding to a real individual.
- **Animated selfies**, where a static image is transformed into a moving video mimicking natural facial motion.

[1] The quantitative data points, figures, and diagrams referenced in this paper are drawn from *Entrust, 2026 Identity Fraud Report: The Changing Face of Fraud*.

All three variants are observed across the identity lifecycle, including onboarding, authentication, and recovery flows. While they differ technically, they share a common objective: impersonation at scale. Increasingly, these attacks succeed not because visual imperfections are obvious—the media is often visually indistinguishable from genuine capture—but because it is delivered through technical pathways that bypass or weaken traditional safeguards.

3. The shift from presentation attacks to injection attacks

Historically, biometric spoofing focused on *presentation attacks*, such as photos, masks, or screens shown to a camera. While these methods of attack persist—including cases where deepfakes are presented as videos of screens—the most consequential evolution has been the rise of *injection attacks*. In practice, they have become the primary vector for submitting deepfakes into biometric systems. Injection attacks bypass the camera entirely by introducing falsified images or video streams directly into the camera feed, using techniques such as virtual cameras, device emulation, browser manipulation, or network-level payload tampering.

Injection-based attacks have increased by roughly 40% year-on-year and are now widely recognised as the primary delivery mechanism for sophisticated deepfakes. These attacks are effective because the software convincingly emulates legitimate devices and camera inputs, undermining controls that rely on environmental cues, metadata, or basic liveness checks. This marks a structural change in the threat model: the attack surface is no longer limited to the biometric itself, but extends to the device, software stack, and transmission channel.

Figure 1 — Suspected injection attack rate in biometric verification flows (2023–2025)

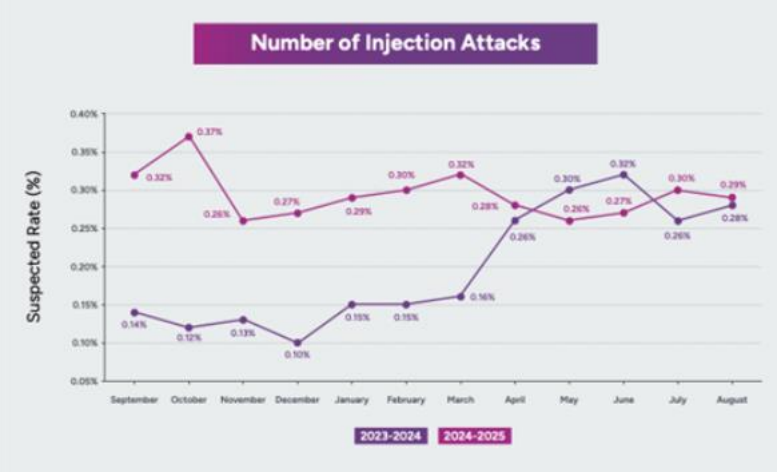

The chart shows a step increase in suspected injection attacks in late 2023–24, followed by sustained stabilisation through 2024–25, suggesting they are now a persistent feature of the biometric threat landscape.

4. The evolving role of liveness detection

Liveness detection remains essential. Passive and active techniques help establish that biometric data was captured from a living person in real time by introducing motion, temporal depth, and randomness.

However, liveness alone is insufficient against modern attacks. Predictable prompts can be learned and replayed, while visually convincing deepfakes can satisfy purely visual-based checks. In controlled deployments, active liveness with enforced randomness can reduce the rate of successful biometric fraud attempts to well below 0.1%, raising the cost and complexity of attacks. As verification systems harden, fraud predictably migrates toward more sophisticated techniques—especially injection-based attacks—making effective defence increasingly reliant on richer, more resilient passive signals to expose anomalies that cannot be spoofed at the source.

Effective defence therefore requires liveness to be combined with capture integrity and non-visual signals, in addition to enforced randomness, to ensure real-time presence rather than replayed or injected content.

5. Measuring effectiveness responsibly

As biometric fraud becomes more sophisticated, organisations, including government agencies and private-sector service providers, must adopt more nuanced performance metrics. Industry guidance converges on three decisive measures for deepfake resilience: False Acceptance Rate (FAR / APCER), False Rejection Rate (FRR), and Missed Fraud Rate (MFR / BPCER).

Optimising a single metric in isolation can distort the balance between FAR, FRR, and MFR, either by excluding legitimate users or by allowing fraud to scale silently. Responsible biometric systems used by government authorities and private-sector organisations alike require balanced optimisation and transparency around trade-offs.

6. Standards and their limits

International standards establish baseline trust. ISO/IEC 30107-3 remains the global benchmark for Presentation Attack Detection, but it was not designed to address injection attacks. Emerging specifications, such as CEN/TS 18099, begin to fill this gap by focusing on capture circumvention and media injection in response to evolving attack vectors.

These frameworks should be treated as minimum foundations rather than complete solutions. Attack techniques evolve faster than certification cycles, making continuous adaptation essential.

This evolution reflects a broader industry shift: deepfakes are now understood not merely as presentation threats, but as system-level risks tied to capture integrity and injection resilience.

7. Best practices for preventing deepfake-enabled fraud

Many effective defenses against deepfake fraud rely on novel capabilities—particularly in the use of passive and integrity-based signals—and on the application of machine learning and automation to detect patterns at scale. However, their impact ultimately depends on disciplined implementation choices that remove unnecessary attack surface. Experience across public- and private-sector identity deployments highlights several priorities:

- prefer secure, native capture over web;
- keep capture components up to date;
- minimise opportunities for injection by disabling uploads and limiting retries;
- harden web-based flows through integrity-aware capture;
- treat fraud prevention as a continuous process through monitoring, automation and adaptation.

8. Conclusion: keeping biometrics “real”

AI succeeds where systems still equate realism with authenticity. Deepfakes exploit weak assumptions about capture, integrity, and trust rather than the limits of human perception. By combining layered controls, metrics, and evolving standards, the biometric ecosystem can deliver security and trust.

Joined in 2016

Organisation: Entrust
Name: Virginia Chiarentin Senior Product Manager;
Samuel Steg, Senior Principal Regulatory Compliance Strategist
Telephone number: +33 6 83 75 0253
Contact details: samuel.steg@entrust.com

RESPONSIBLE BIOMETRICS

2001 BIOMETRICS INSTITUTE 2026

SILVER JUBILEE

A global community promoting the **responsible**, **ethical** and **effective** use of biometrics since 2001

biometricsinstitute.org