Revisiting Biometrics in Cybersecurity: Do AI Methods and Zero-Trust Architectures Drive Innovation?

Siphesihle Sithungu[1] and Christoph Lipps[2]
[1]Academy of Computer Science and Software Engineering, University of Johannesburg, South Africa
[2]Intelligent Networks Research Group, German Research Center for Artificial Intelligence, Kaiserslautern, Germany

siphesihles@uj.ac.za
christoph.Lipps@dfki.de

Abstract: Biometric authentication has long been regarded as a foundational element of identity verification, leveraging unique physiological and behavioral traits to enhance security beyond traditional passwords. While it offers notable advantages such as convenience and resistance to identity theft, concerns are mounting regarding privacy, susceptibility to spoofing, and the irreversibility of compromised biometric identifiers. These weaknesses are becoming increasingly critical as digital infrastructures evolve into distributed, dynamic environments in which static trust models are no longer sufficient. Moreover, several traditional modalities -such as fingerprints, iris scans, and voice recognition- have already been breached. However, Artificial Intelligence (AI) methods are reshaping this landscape by introducing adaptive and context-aware features into biometric systems. Machine Learning (ML) techniques enhance accuracy, enable continuous authentication, and support multimodal fusion, while anomaly-detection mechanisms improve resilience against sophisticated attacks. Generative AI (GenAI) plays a particularly significant role, though it introduces a paradox: it empowers defenders through realistic attack simulations and robustness testing, yet simultaneously equips attackers with tools for producing deepfakes and synthetic identities, thereby expanding the attack surface. In this evolving security landscape, Zero-Trust Architectures (ZTA) have gained prominence as a model that replaces assumptions of inherent trust with continuous verification mechanisms. The use of biometric data within ZTA can enhance the reliability of identity verification; however, it also intensifies several existing issues. Biometric identifiers must be handled and stored in ways that safeguard individual privacy and align with relevant legal requirements, and the incorporation of AI-based assessment methods introduces additional concerns regarding potential bias, transparency, and oversight. Moreover, combining AI-supported biometric systems with Zero-Trust principles raises further questions about scalability, system compatibility, and the broader ethical consequences of more pervasive identity monitoring. This work therefore examines the convergence of biometrics, AI, and Zero-Trust principles from a critical perspective. It highlights the dual role of AI as both a source of innovation and a generator of new threats, while identifying opportunities for adaptive security, real-time threat detection, and improved user experience. By analyzing technical and operational dimensions, the work proposes a roadmap for integrating biometrics into ZTA that balances innovation with accountability and supports trustworthy, resilient cybersecurity frameworks.

Keywords: Biometrics, Zero-trust architecture, Generative AI, Machine learning, Cybersecrutiy

1. Biometric Authentication: From Face Recognition to AI-driven Identity Verification

Current developments in Artificial Intelligence (AI) methods are -again- shifting biometric authentication from the fringes of research to everyday security. Facial recognition unlocks smartphones and payments, fingerprint readers secure workplaces, and voice biometrics speed up customer support. With applications increasingly distributed across clouds and devices, identity has become the new security perimeter, and biometrics provide a human-centric authentication factor improving both usability and security.

This development did not happen overnight, but rather biometric technologies have gone through four phases of development: (i) early research in the 1960s; (ii) large-scale government implementations in the 1990s and 2000s; (iii) consumer adoption via smartphones in the 2010s; and (iv) today's AI-driven boom in the 2020s (Jain, Nandakumar, and Ross, 2016)(Buriro, & Luccio, 2025)(Wayman et al. 2025). Thereby, each phase was triggered by advances and changing security requirements shaping this landscape.

And now, the further developments of AI methods have again changed both, the accuracy and speed of biometric recognition, enabling systems to operate at scale and adapt to different conditions. But these same advances have also opened new attack vectors: deepfakes, high-resolution voice clones, adversarial inputs, and synthetic identities (Farooq et al., 2025). This duality raises a central question: *Do current AI-based methods offer security innovations, or do they merely replace old weaknesses with new ones?*

Historically, biometric systems were designed for closed, static environments with cooperative users and calibrated sensors, as well as risk models balancing false acceptance and false rejection. But, today, authentication takes place in dynamic, distributed contexts: Bring-your-own-device (BYOD) endpoints, variable networks and sensors, evolving user behavior, and continuous software updates. Attackers exploit these

conditions with automated, low-cost tools. In such environments, static thresholds and point-in-time checks are insufficient. What matters is continuous risk assessment, rapid model adaptation, origin-aware detection, and deep defense that assumes compromise, a mindset embodied by Zero Trust. AI methods, on the other hand, are driving this change forward for both defenders and attackers.

On the defense side, deep learning (DL) has improved detection accuracy, enhanced liveness detection, enabled cross-modality fusion, and supported continuous authentication through behavioral biometrics such as keystroke dynamics and gait. On the attacker side, generative models reduce the cost and expertise required to create convincing fakes and adversarial examples. Added to this is the large-scale sharing of biometric data -from photos to voice samples- which highlights a biometric problem: irreversibility. Unlike passwords, biometric features cannot be rotated after a compromise; they can be converted into erasable templates or bound to additional secrets, but the original signal remains visible or is lost forever (Lipps, Herbst, and Schotten, 2021).

This reveals a major gap between research and practice: Many implementations are still based on static statistical models with questionable assumptions about distributions and user cooperation. Calibration drifts as populations diversify and contexts evolve; thresholds tuned into laboratories deteriorate in production; liveness checks that prevent simple spoofing fail against AI-assisted attacks. Furthermore, companies often use biometrics as their sole access control rather than as a risk signal within a broader zero-trust framework, where identity assurance is contextual, continuously updated, and combined with device integrity, network configuration, and behavioral analytics. Thereby, the result is a misalignment as AI-powered biometrics is marketed as an end-to-end solution in environments requiring composable, continuously verifiable signals.

The critical question therefore is not whether AI increases accuracy -it does- but whether it increases security assurance against adaptive attackers and real-world constraints. Any evaluation must go beyond closed benchmarks and focus on operational robustness: resilience to presentation and injection attacks, transferability of adversarial examples, resilience to distribution shifts, privacy-compliant template management, and graceful degradation in the event of partial compromise. In parallel, biometrics must be transformed from immutable identifiers to revocable, policy-bound signals within zero-trust architectures: anchored in device attestations, combined with cryptographic protocols, and continuously monitored with anomaly detection and explainable risk assessment (Uppal, et al., 2024).

Against this backdrop, this work essentially pursues three goals: First, to highlight the research gap between laboratory results and operational requirements for resilience, recoverability, and trustworthy calibration in dynamic, distributed environments. Second, to analyze the dual role of AI -how it enables stronger detection and liveness defense while introducing new classes of attacks and risks in the supply chain. and third, to position biometrics within Zero Trust as a composable signal rather than a single point of truth, outlining architectural patterns that integrate biometrics with device trust, continuous authentication, and privacy-preserving design. Ultimately, innovation in AI-driven biometrics is less about incremental gains in recognition benchmarks and more about end-to-end security engineering: rigorous threat modeling, measurable robustness, revocability in the event of compromise, built-in privacy, and principled integration into Zero Trust workflows. Redefining biometrics from this perspective separates hype from progress and sets a research agenda that strengthens digital trust, not just rationalizes it.

2. The Background: From Traditional Biometrics to Zero-Trust Principles

Biometric authentication has developed from the use of physical traits such as fingerprints and iris patterns to advanced behavioural and continuous authentication methods. Its evolution reflects increasing computational capabilities, usability demands and the shift toward probabilistic, ongoing identity verification within Zero Trust frameworks.

2.1 Definition of Biometric Authentication

Biometric authentication is a process making use of unique physiological or behavioural characteristics of entities (i.e. fingerprints, voice patterns, facial features, gait, iris scans) to determine the legitimacy of access they claim to have in a particular system (Amiri, et al., 2024). In computing, biometrics are considered a more reliable access control mechanism as opposed to passwords (which rely on memory) or tokens (which rely on possession) as authentication is linked directly to the individual (Ryu, et al., 2023). Figure 1 illustrates the three levels into which authentication factors are divided.

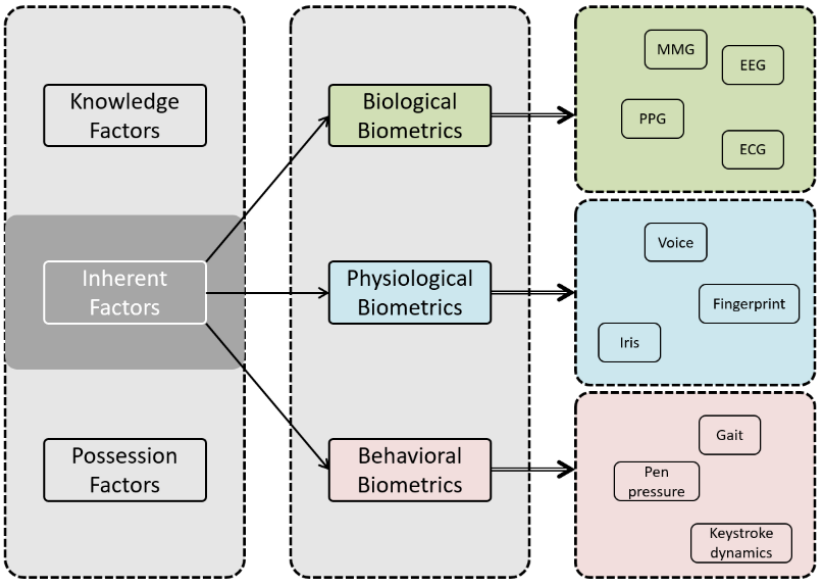

Figure 1: Authentication Factors are basically divided into three levels: Knowledge-, Inherence- and Possession Factors. Inherence Factors are further specified into the sub-groups: biological-, physiological- and behavioural biometrics, which could be further specified (Lipps, et al., 2022)

As indicated in Figure 1, authentication traditionally relies on three categories of factors: knowledge-based, possession-based, and inherence-based factors. Knowledge-based factors depend on information the user knows, such as passwords or PINs, whereas possession-based factors require something the user physically holds, such as a smart card or security token. Inherence-based factors, by contrast, rely on attributes inherent to the individual. Biometrics therefore fall within this third category and can be further differentiated, for instance, into biological, physiological, and behavioural characteristics. Examples of these inherent traits -such as fingerprints, iris patterns, voice characteristics, or gait- are illustrated in the right-most part of the Figure (Lipps, et al., 2022).

Due to advances made in the field of computer science within the last two decades, biometric authentication systems may also apply learning mechanisms to improve pattern recognition. The ability to learn also enables biometric systems to adapt over time and refine their ability to detect biometric features (Amiri, et al., 2024).

2.2 Evolution of Biometrics from Fingerprints and Iris Scans and the Rise of AI-based Methods

Early biometric systems progressed from measuring what a person *touches*, to what can be *seen*, and eventually to how a person *behaves*. This evolution was primarily driven by the desire to capture richer information, reduce user friction, and leverage increasing computational capabilities. The earliest widely adopted biometric modality was fingerprint recognition, due to its high distinctiveness, long-term stability, and the relative simplicity of comparison algorithms. However, fingerprint-based systems also faced notable limitations as they required direct physical contact, raised hygiene concerns, and were vulnerable to physical wear or intentional alteration of the skin. These challenges were less about accuracy and more about usability and practicality in real-world environments. (Herbst et al., 2024).

The next wave of biometrics focused on iris and retina scans, which moved the medium to “inside” the body where information is richer and replication is more difficult, resulting in a significant improvement in accuracy. However, the biggest trade-off was cost (i.e. specialised sensors and controlled lighting) and intrusiveness (user discomfort). These factors made iris detection suitable for high-security environments (i.e. border control) and less practical for most use cases. Furthermore, face recognition became the most ideal option in many cases due to the increasing availability of high-quality cameras and the advantages that come from using faces. Faces are socially natural, do not require physical contact and inexpensive to capture. The associated costs (lighting, ageing, pose and bias) resulting in this kind of detection being statistical rather than absolute.

More recently, biometrics evolved into voice and behaviour where voice biometrics analyse pitch, cadence and spectral features whereas behavioural biometrics analyse a range of behaviours depending on context (i.e. typing or swiping style, gait or how you hold a mobile device). Such biometrics are ongoing and passive, turning authentication from a one-off activity into an ongoing process. The idea of continuous authentication has also

led to another emerging security paradigm (canonically unrelated to biometrics) called Zero Trust, which is described in the following section.

2.3 What is Zero Trust

Zero Trust -often referred to as Zero Trust Architecture (ZTA)- emerged within systems security and is fundamentally concerned with answering a key question: "Given a stream of requests, what is the likelihood that this is still the same user?" Because this question is both probabilistic and time-dependent, it cannot be resolved through a single authentication event conducted prior to granting access. Instead, it necessitates a model of continuous authentication.

Traditional perimeter-based security operated under the assumption of a clear boundary between the "inside" and the "outside." This model collapsed with the rise of cloud computing, mobile devices, and the increasing prevalence of compromised credentials. A one-time login check is no longer adequate to establish -or maintain- trust. As outlined in guidance from the National Institute of Standards and Technology (NIST), Zero Trust shifts the focus toward safeguarding resources and continuously evaluating trust rather than granting it implicitly (Rose, et al., 2020).

A notable connection between ZTA and behavioural biometrics is that, despite their different origins -behavioural biometrics in signal processing and pattern recognition, and Zero Trust in systems security- both disciplines adopt the principle of continuous authentication. Behavioural biometrics poses a parallel question: "Given a stream of actions, what is the likelihood that this is still the same user?" This, too, is a probabilistic and dynamic assessment. Consequently, both approaches converge on the idea that trust is not binary, confidence naturally degrades over time, evidence arrives continuously, and authentication decisions must remain open to revision.

3. Rethinking Biometric Security in the AI era: Fundamental Risks and the Limits of Traditional Trust Models

Despite significant progress in AI-driven biometrics, critical gaps remain between laboratory performance and real-world security requirements. First, current systems face new vulnerabilities introduced by AI itself, including adversarial attacks, deepfakes, and synthetic identities that undermine trust in biometric authentication. Second, the inherent irreversibility of biometric features means that once compromised, they cannot be easily revoked or replaced, creating long-term exposure risks. Third, many deployments still rely on outdated, static trust models that fail in dynamic, distributed environments where continuous verification and adaptive risk assessment are essential. Addressing these challenges is key to ensuring that AI-driven biometrics deliver genuine security innovation rather than shifting old weaknesses into new forms.

3.1 The Main Issues Considering AI in Biometric Authentication

One key aspect, however, is the dual purpose of today's AI in biometric authentication, leading to a significant conflict between improved accuracy and increased vulnerability to attacks. Generative and discriminative techniques designed to improve recognition performance also lower the barriers to presentation and injection attacks -including deepfakes, voice clones, and synthetic fingerprints-, enabling spoofing to scale and identity theft in remote onboarding and electronic Know Your Customer (eKYC) contexts (Verma et al., 2023). The limitations of current mechanisms for detecting liveness and presentation attacks (PAD) against model-aware, AI-powered attackers point to a shift from artifact-based identity spoofing to hostile manipulation of decision boundaries. At the same time, large-scale aggregation and scraping of biometric data poses significant privacy risks, as compromised features are irreversible and common methods for protecting templates are insufficient to ensure revocation semantics comparable to passwords. Taken together, these factors reveal a security gap between benchmark accuracy and operational robustness, necessitating a transition from accuracy-oriented reporting to risk- and attack surface-oriented design, assessment, and governance.

3.2 The Irreversibility of Compromised Biometric Features

Due to these challenges, the irreversibility of biometric features exacerbates systemic risks in an AI-driven threat landscape. Unlike passwords, physiological features cannot be replaced once disclosed, meaning that any compromise becomes a permanent vulnerability rather than a remediable incident. Model inversion and reconstruction attacks can regenerate facial images, fingerprints, or voice patterns from stolen embeddings, and modern AI techniques further exacerbate this problem, allowing attackers to replicate a user's biometric identity at will (Kim et al., 2024). Without meaningful revocation semantics or mature lifecycle controls, compromised

templates remain exploitable indefinitely in systems based on the same modality. At the same time, advances in generative models enable increasingly precise biometric replay and injection attacks that completely bypass sensors and target authentication pipelines directly. These developments underscore the need to reevaluate biometric data as static anchors of trust and instead integrate them into adaptive, risk-oriented authentication frameworks capable of mitigating the lasting effects of compromised features.

3.3 Outdated, Static Trust Models in Dynamic, Distributed Environments

Nonetheless, conventional biometric authentication frameworks continue to rely on centralized trust models no longer aligned with current decentralized, multi-party identity ecosystems. As authentication processes increasingly involve cloud services, third-party providers, federated identity platforms, and remote onboarding channels, static trust anchors cannot capture the fluidity and interdependence of modern digital infrastructures. At the same time, these systems lack continuous verification and contextual risk assessment, granting long-lasting trust based on a single biometric match, an assumption that is incompatible with session hijacking, device switching, and adaptive threats. The problem is exacerbated by AI-driven attackers who are able to iteratively refine spoofing inputs in response to system feedback, rendering static threat models obsolete. Furthermore, older biometric trust anchors conflict with the fundamental principles of Zero Trust, which requires continuous verification, context awareness, and the minimization of implicit trust. These limitations underscore the need to move beyond static notions of identity and transition to architectures that position biometrics as dynamic signals within a continuous, risk-based Zero Trust framework.

4. Transformation by Artificial Intelligence Methods

The rapid maturation of artificial intelligence (AI) has materially advanced the design, operation, and assurance of biometric systems. In particular, AI makes a continuous, risk-based approach -fully aligned with Zero-Trust principles- both feasible and practical at enterprise scale. Rather than treating authentication as a one-off verdict, AI enables adaptive, context-aware decision-making that evolves with each interaction and adjusts to changing user states and environmental conditions. Concretely, machine learning improves recognition accuracy and robustness; multimodal fusion and anomaly detection increase resilience and fault tolerance; and generative AI (GenAI) augments training, simulation, and stress-testing while also creating new forms of attack that must be mitigated. Operationally, this AI-driven transformation changes how signals are captured, features are represented, uncertainty is handled, and policy decisions are enforced. Signals are no longer evaluated in isolation but as part of a streamed evidence pipeline in which risk scores are updated continuously and policy engines (e.g., ZTA policy decision points) can step up, maintain, or revoke access dynamically. Data governance and privacy safeguards must keep pace with these capabilities to ensure lawful, fair, and explainable use of AI in identity systems.

4.1 Adaptive and Context-Aware Biometric Systems via AI Methods

AI allows biometric systems to progress from static templates to incrementally updated models that adapt to users and contexts over time. Instead of verifying identity only once at login, the system can re-evaluate whether access remains appropriate at any given moment and under specific conditions. This is achieved by continuously incorporating contextual signals -such as device posture and health, geolocation patterns, time-of-day regularities, network characteristics, session history, and recent behavioural markers- into a dynamic risk score. In practice, a user who typically authenticates from a specific laptop and office network may be granted low-friction continuation, whereas a session that suddenly shifts device, IP range, and typing cadence may trigger step-up controls (e.g., re-capture of a face or voice sample, or a possession check).

Architecturally, adaptive systems rely on feature drift monitoring and model-update strategies (e.g., constrained online learning or periodic batch refresh) to avoid overfitting to transient anomalies while still accommodating natural change (e.g., ageing, stress). They also benefit from calibrated uncertainty estimates (e.g., temperature scaling for score calibration) so that downstream policy engines can interpret scores consistently. From an implementation standpoint, policy-as-code mechanisms tie these evolving risk estimates to clear actions: maintain, limit, re-authenticate, or terminate.

These architectural properties translate directly into practical deployments across high-variability settings. In remote learning and assessment, for instance, conditions change frequently; contextual indicators -such as device swaps, location shifts, and irregular activity patterns- become integral to maintaining trust over time (e.g., proctoring augmented with continuous, context-aware checks). In large, distributed enterprises, multi-agent arrangements coordinate sensors, models, and policy components across edge and cloud, allowing risk to be reassessed as sessions move between networks or devices (Shukur, 2017). Across both scenarios, trust in the

biometric signal is situational rather than absolute, adapting to surrounding factors to increase security while minimising unnecessary user friction.

4.2 Machine Learning for Accuracy and Continuous Authentication

Machine learning (ML) enhances biometric systems by enabling them to learn continuously rather than only at the enrolment stage. Models adjust representations and decision boundaries as user characteristics and operating conditions evolve, accounting for factors such as ageing, fatigue, stress, medication effects, ergonomic changes, and shifts in behaviour. At the same time, ML-based systems filter noise and extract more stable, higher-level features from real-world data (Bengio, Courville, & Vincent, 2013). In practical deployments, this means that the system can refine user-specific embeddings over time and calibrate individual thresholds so that specificity remains high without sacrificing sensitivity across diverse populations.

A key consequence of these capabilities is that ML naturally supports continuous authentication. Rather than relying on a single decisive event, the system can update confidence scores at interaction time -whether through keystroke windows, cursor-movement segments, micro-gesture batches, or brief voice snippets. This finer temporal granularity allows early detection of deviations from a user's behavioural patterns and makes it possible to combine low-friction, passive checks with active verifications (such as short face recapture) only when the assessed risk increases. The result is a model of risk-proportional friction, which enhances security while minimising unnecessary user interruption.

To operate reliably in this dynamic setting, ML-driven biometric systems must incorporate several design safeguards. They require mechanisms to prevent catastrophic forgetting, ensuring that new behavioural observations do not overwrite stable historical patterns. They also need concept-drift detection to distinguish natural behavioural evolution from adversarial manipulation or account compromise. In addition, systems must integrate fairness and bias monitoring to maintain equitable performance across demographic groups, as well as explainability artefacts -such as feature-attribution summaries or confidence decompositions- that support governance, auditability, and transparent risk decisions.

4.3 Multimodal Fusion and Anomaly Detection to Increase Resilience

Biometric systems based on a single modality often create single points of failure: they perform reliably under normal conditions but become vulnerable when a sensor is spoofed, degraded, or unavailable. Multimodal fusion mitigates this fragility by combining evidence from multiple modalities -such as face, voice, fingerprint, keystroke dynamics, gait, and contextual signals- to produce a more robust and discriminative identity estimate (Ross & Jain, 2004). Fusion can occur at the feature level (early fusion), score level (mid-level fusion), or decision level (late fusion), each offering different trade-offs in terms of transparency, computational cost, and response time. In adversarial settings, multimodal systems significantly increase attacker workload, as spoofing several modalities simultaneously -often under active liveness or presentation-attack-detection (PAD) controls- becomes far more challenging.

In parallel, anomaly detection introduces a complementary layer of protection by focusing not on identity confirmation alone, but on detecting deviations from typical behavioural or interaction patterns. This approach is especially well suited to continuous authentication, where behaviour naturally unfolds as a sequence rather than a single event (Gholami, Alaca, & Zulkernine, 2025). Techniques such as one-class classifiers, autoencoders, isolation forests, and density-based models allow systems to learn a user's characteristic manifold and to flag activity that falls outside expected bounds. Such mechanisms can detect impostors who might bypass a single check, identify legitimate users whose behaviour has been altered due to coercion or malware, and guide risk-adaptive responses that restrict or escalate access when anomaly scores exceed thresholds.

When viewed through a Zero-Trust lens, the integration of multimodal fusion with anomaly detection creates a layered and adaptive defence posture. Evidence becomes redundant (tolerant of sensor failure), diverse (harder to spoof), and temporal (evaluated continuously rather than once). This enriched evidential base enables finer-grained policy decisions—for example, allowing reduced-privilege access under moderate uncertainty, triggering step-up authentication under elevated risk, or isolating a session entirely when anomalous behaviour persists. Taken together, these capabilities shift the system from static recognition towards situational inference, where identity, context, and behavioural consistency are assessed continuously to maintain an appropriate and dynamically calibrated level of trust.

4.4 Generative Artificial Intelligence: Opportunities and Risks

Generative AI (GenAI) introduces strong advantages for biometric security but also creates significant new attack vectors. On the opportunity side, GenAI enhances simulation and red-teaming, enabling synthetic attack artefacts -such as altered illumination, voice perturbations, or spoofing materials (Khan and Khan, 2025)- to be generated for controlled robustness testing. It also supports data augmentation and domain randomisation, enriching training sets with rare poses, sensor noise patterns, or environmental variations, and improves rare-event modelling by synthesising edge cases that would otherwise be difficult to capture. Together, these capabilities help systems train against synthetic attacks and improve resilience without relying solely on large, hard-to-collect real-world datasets.

At the same time, GenAI significantly strengthens adversarial capabilities. Attackers can produce high-fidelity deepfakes, synthetic behavioural traces, and automated content that scales social-engineering efforts and enables synthetic identity fraud. This shortens the cycle between new attack techniques and their real-world deployment, expanding the overall attack surface.

Effective use of GenAI therefore requires robust presentation-attack detection, media-provenance checks, and challenge–response mechanisms that exploit signals difficult to synthesise. Continuous model hardening -through adversarial training and post-deployment monitoring- is essential. Beyond technical measures, sound governance is needed, including dataset lineage tracking, synthetic-data labelling, and clear documentation such as model cards and evaluation reports to ensure that GenAI contributes safely to the security stack.

In conclusion, AI does not improve biometrics by making them stricter. Instead, AI improves biometrics by making them continuous, contextual and probabilistic. Security transitions from being a verdict into being an ongoing assessment.

5. Zero Trust Architecture and the Integration of Biometrics

At this point, the role of biometrics in Zero Trust architectures is redefined from static trust anchors to dynamic, continuously evaluated signals. It outlines the core principles of Zero Trust, identifies optimal integration points for biometric verification, and demonstrates how adaptive trust models counter AI-driven identity threats in distributed environments.

5.1 The Principles of Zero Trust Architectures

Zero Trust -as already mentioned-, focuses on the basic principles of continuous verification, least privilege, and the elimination of implicit trust. Identities, devices, networks, and workloads are treated as untrustworthy by default. Access is granted only when needed and to an appropriate extent, based on dynamic, context-sensitive risks. Verification is event- and time-driven, policies reevaluate signals (identity status, device status, network context, behavioral telemetry) before access and throughout the session. Implicit trust (e.g., “inside the perimeter,” “recently authenticated,” or “biometric match == secure”) is removed; Instead, policy engines -so called policy decision points (PDPs)- create repeatable, auditable decisions that policy enforcement points (PEPs) enforce at every boundary. The architecture emphasizes deep defense (segmentation, strong identity, hardened endpoints, application-level controls) and telemetry fusion to detect anomalies and adjust permissions in real time. Crucially, Zero Trust decouples identity from the duration of permission: authentication is not a permanent pass, but a revocable, continuously assessed status. This mindset enables resilient responses to AI-powered threats, session hijacking, and synthetic identities in distributed environments.

5.2 Integration of Biometric Data into ZTA application

A coherent integration of biometrics into Zero Trust Architectures requires that biometric signals be treated as dynamic, context-dependent inputs within a continuously evaluated trust calculation, rather than as static authentication anchors. Their role only becomes effective when they are embedded in the policy loop -where PDPs coordinate decisions and PEPs enforce them- along with device status, behavioural telemetry, and cryptographically verifiable login data:

1. Step-up at sensitive boundaries (PDP-triggered): Biometric verification should be applied selectively when the PDP detects elevated risk conditions such as privilege escalation, anomalous device configuration, suspicious geolocation, or behavioural anomalies. In these situations, the PDP issues a step-up requirement that mandates new biometric verification with robust liveness and anti-spoofing guarantees. Biometric evidence serves as one of several confirming signals, including device attestations, possession-based evidence, and contextual risk assessments.

2. Continuous, seamless session assurance (PEP-enforced): During active sessions, PEP can maintain identity assurance through passive behavioural please biometrics and seamless indicators of user continuity. By periodically reevaluating policies—either time-based or event-driven—PEP can detect deviations or anomalous activity. Active biometric checks are only triggered when passive signals indicate increasing uncertainty, reducing the duration of session takeovers without creating constant friction points.
3. Account recovery and re-binding of identity: To prevent takeover through compromised recovery channels, re-binding workflows should combine hardware-based possession factors (e.g., FIDO2/WebAuthn) with multimodal biometrics and PAD-verified, trusted capture. This ensures that recovery cannot be undermined by synthetic identities or injected biometric artifacts.
4. High-risk remote onboarding (eKYC): In the onboarding context, biometrics complements cryptographic document verification with liveness detection, trusted device capture, and anti-injection attestations. Templates should be protected by revocable or transformation-based schemas, with clearly defined mode switching policies to transition users to alternative modalities when compromise is suspected.

5.3 Technical and Legal Requirements for the Processing and Storage of Biometric Data

Given the sensitivity and inherent irreversibility of biometric identifiers, their processing and storage must adhere to strict technical and regulatory requirements. From a technical standpoint, biometric pipelines require trusted capture, strong PAD, and protection against injection and replay, ensuring that only authentic signals enter the decision flow. Templates must be secured through cancellable or transform-based protection schemes that reduce the impact of compromise and prevent template inversion. Storage environments should use hardware-rooted isolation, encryption in transit and at rest, and strict segmentation to limit correlation across systems.

Legally, biometric data constitutes high-risk, special-category information, requiring explicit consent, purpose limitation, data minimization, and well-defined retention policies. Privacy-by-design measures -including transparent processing, documented risk assessments, and enforceable user rights-, are essential to ensure that biometric systems remain compliant, auditable, and proportionate. Together, these requirements establish the baseline for safe deployment within adaptive, risk-oriented identity architectures

6. Conclusion and Future Work

This work shows that the most consequential change in biometric authentication is architectural rather than a matter of squeezing out more recognition accuracy. In concert, AI methods and Zero-Trust principles shift authentication from a static, one-off judgement to a continuous, contextual, and probabilistic process. Systems engineered for stable signals, trusted environments, and cooperative users cannot meet the demands of cloud-native, distributed, and adversarial settings shaped by AI-enabled threats. In such environments, single checks and fixed thresholds are no longer adequate; what is required is a pipeline accumulating evidence over time, recalibrates confidence as conditions change, and ties those updates to clear policy actions.

Therefore, AI algorithms provides the technical substrate for this shift. Machine Learning supports temporal modelling of identity and per-user calibration, strengthening robustness as behaviour and operating conditions evolve. Multimodal fusion mitigates single-channel fragility by combining complementary signals, while Anomaly Detection reframes security as situational inference, surfacing deviations that simple matches miss. At the same time, Generative AI lowers the cost of high-fidelity spoofing and synthetic identities, exposing the limits of instance-based trust and expanding the attack surface. The result is an arms race demanding for continuous hardening, calibrated uncertainty, and policy-driven responses rather than static acceptance criteria.

Zero-Trust Architectures supply the operational fabric for these capabilities, wherein biometrics cease to function as immutable anchors and instead act as dynamic, policy-bound signals within a broader risk calculus -combined with device posture, network context, workload sensitivity, and behavioural telemetry. Trust becomes scoped, revisable, and revocable; policy decision points can step up, sustain, or curtail access as evidence evolves across the session.

Accordingly, any progress hinges on new evaluation criteria. Beyond headline accuracy, systems must demonstrate operational robustness to presentation and injection attacks, maintain tolerance to drift, and degrade gracefully when conditions deteriorate. They must also support revocability through cancellable templates, modality pivots, and secure re-binding processes. In parallel, privacy-by-design principles need to be embedded throughout the lifecycle, and biometric components must integrate cleanly into Zero-Trust

enforcement rather than operating as isolated mechanisms. Governance plays an essential role in making these properties transparent and auditable. Clear model cards, calibration reports, provenance tracking, and incident playbooks should enable independent verification of system behaviour and support accountable deployment. Looking ahead, the agenda is clear. Research and practice must standardise robustness benchmarks under generative-AI-driven threats, develop calibrated and explainable risk-scoring mechanisms, scale multimodal fusion with temporal anomaly modelling, and create practical revocation semantics that recognise the permanence of biometric data. Framed this way, AI-driven biometrics can contribute meaningfully to trustworthy digital identity - not as a single point of truth, but as a composable signal within a resilient, risk-based security ecosystem.

Acknowledgement

This work has been supported by the Federal Ministry of Education and Research of the Federal Republic of Germany (Förderkennzeichen 16KIS2239K SUSTAINET_guarDian). The authors alone are responsible for the content of the paper.

Ethics and Data Use Statement: This research used publicly available, anonymized secondary data. No identifiable personal information was accessed. Ethical approval was not required in accordance with institutional policies.

AI Assistance Statement: The authors used generative AI tools solely for linguistic editing and improving readability (e.g., grammar, phrasing, and structure). No AI system contributed to the formulation of the research questions, methodology, data interpretation, theoretical reasoning, or conclusions. All content was critically reviewed and approved by the authors, who take full responsibility for the manuscript.

References

Amirir, Z., Heidari, A., Jafari, N., and Hosseinzadeh, M., „Deep study on autonomous learning techniques for complex pattern recognition in interconnected information systems“, *Computer Science Review*, vol. 54, 2024, ISSN: 1574-0137, DOI: 10.1016/j.cosrev.2024.100666

Bengio, Y., Courville, A., Vincent, P., „Representation Learning: A Review and New Perspectives”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798 – 1828, 2013, DOI: 10.1109/TPAMI.2013.50

Buriro, A., and Luccio, F., “Mobile Biometrics: Innovations, Challenges, and Emerging Trends”, *Advanced Information Networking and Applications. AINA 2025. Lecture Notes on Data Engineering and Communications Technologies*, vol 252. Springer, Cham., DOI: 10.1007/978-3-031-87784-1_16

Farooq, M.U., Khan, A., Uddin, K., and Malik, K.M., “Transferable Adversarial Attacks on Audio Deepfake Detection”, *IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, Tucson, AZ, USA, 2025, DOI: 10.1109/WACVW65960.2025.00178

Gholami, A., Alaca, F., and Zulkernine, M., “An anomaly detection based approach for continuous authentication with smartwatch inertial sensors”, *Computer & Security,* vol. 159.,2025, DOI: 10.1016/j.cose.2025.104656

Herbst, J., Rüb, M., Sanon, S.P., Lipps, C., and Schottem, H.D., “Medical Data in Wireless Body Area Networks: Device Authentication Techniques and Threat Mitigation Strategies Based on a Token-Based Communication Approach”, *network*, vol. 4, no. 2, 2024, DOI: 10.3390/network4020007

Jain, A.K., Nandakumar, K., and Ross, A., “50 years of biometric research: Accomplishments, challenges and opportunities”, *Pattern Recognition Letters,* Elsevier, 2016, DOI: 10.1016/j.patrec.2015.12.013

Kahn, F. A., and Kahn, M.K., „ Generative AI and Deepfake Detection in Biometric Systems“, *Cognitive Computation*, vol. 17, no. 112, pp. 1—21, 2025, DOI: 10.1007/s12559-025-10469-3

Kim, S., Tan,, Y.K., Jeong, B., Mondal, S., Aung, K.M.M., and Seo, J.H., “Scores Tell Everything about Bob: Non-adaptive Face Reconstruction on Face Recognition Systems”, *IEEE Symposium on Security and Privacy (SP)*, San Francisco, CA, USA, 2024, DOI: 10.1109/SP54263.2024.00161

Lipps, C., Herbst, J., and Schotten, H.D. “How to Dance your Passwords: A Biometric MFA-scheme for Identification and Authentication of Individuals in IIoT Environments”, *16th International Conference on Cyber Warfare and Security (ICCWS),* Cookeville, Tennessee, US, 2021.

Lipps, C., and Schotten, H.D., “Physical Layer Security: About Humans, Machines and the Transmission Channel”, *European Conference on Cyber Warfare and Security (ECCWS),* Chester, UK, 2022, DOI: 10.34190/eccws.21.1.403

Lipps, C., Bergkemper, L., Herbst, J., and Schotten, H.D., „ I Know You by Heart: Biometric Authentication based on Electrocardiogram (ECG) signals “, International Conference on Cyber Warfare and Security (ICCWS), Albany, NY, USA, 2022, DOI: 10.34190/iccws.17.1.12

Lipps, C., Reddy, R., Rüb, M., and Schotten, H.D., “Authentication in a Hyperconnected World: Challenges, Opportunities and Approaches”, *International Conference on Cyber Warfare and Security (ICCWS),* Johannesburg, South Africa, 2024, DOI: 10.34190/iccws.19.1.2070

Rose, S., Borchert, O., Mitchell, S., and Connelly, S., “Zero Trust Architecture”, NIST Special Publication 800-207, pp. 2—3, 2020, DOI: 10.6028/NIST.SP.800-207

Ross, A., and Jain, A.K., "Multimodal biometrics: An overview", 12th European Signal Processing Conference, Vienna, Austria, pp. 1221—1224, 2004, Print ISBN:978-320-0001-65-7

Ryu, R., Yeom, S., Herbert, D., and Dermoudy, J., „ The design and evaluation of adaptive biometric authentication systems: Current status, challenges and future direction", *ICT Express*, vol. 9, no. 6, pp. 1183—1197, 2023, DOI: 10.1016/j.icte.2023.04.003

Shukur, F., "Using Multiagents for Context-Aware Adaptive Biometrics", *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pp. 5209—5210, 2017, DOI: 10.24963/ijcai.2017/764

Uppal, S., Banga, V., Neeraj, S., and Singhal, A., "A Comprehensive Study on Mitigating Synthetic Identity Threats Using Deepfake Detection Mechanisms", *14th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*, Noida, India, 2024, DOI: 10.1109/Confluence60223.2024.10463324

Verman, K., Kumar, R., Rao, A.P., and Ranjan, R., "Efficient e-KYC Authentication System: Redefining Customer Verification in Digital Banking", 9th International Conference on Signal Processing and Communication (ICSC), NOIDA, India, 2023, DOI: 10.1109/ICSC60394.2023.10441596

Villegas, W. E., Gutierrez, R., Navarro, A. M., and Mera-Navarrete, A., "Adaptive Authentication and Access Control System in Dynamic Educational Environments Based on AI", *Computer*, vol. 58, pp. 53—63, 2025, DOI: 10.1109/MC.2025.3565149

Wayman, J., Jain, A., Maltoni, D., and Maio, D., "Biometric Systems – Technology, Design and Performance Evaluation", Springer London, 2025, ISBN: 978-1-85233-596-0, DOI: 10.1007/b138151