Partners Universal Innovative Research Publication (PUIRP)
Volume: 04 Issue: 04 | July - August 2026 | ISSN: 3048-586X | www.puirp.com

Big Data, Behavioral Biometrics, and the Era of Continuous Authentication - A Review

Dr.A.Shaji George

Independent Researcher, Chennai, Tamil Nadu, India.

Abstract - The combination of pervasive sensing, machine learning and large-scale data collection of human behavior has brought a paradigm change in digital identity assurance from one shot authentication to a continuous authentication process. This review aims to provide a synthesis of the peer-reviewed literature, systematic surveys, industry telemetry and regulatory developments that span the period from approximately 2015 to 2026 in order to characterize this transition. It explores the big data underpinnings of behavioral biometrics, including the volume, velocity and variety of the interaction signals e.g., keystroke dynamics, touch and swipe gestures, gait, mouse movements, and contextual device usage, and the machine learning architectures that transform these signals into a continuous identity confidence score. The review then examines the documented performance of the paradigm, key technical limitations behavioral drift, variability across devices, adversarial imitation and generative attacks, and lack of standardized benchmarks, and unique privacy and governance issues, as behavioral biometrics can be used to make inferences about health, emotion, and demographics, which are distinct from physiological biometrics. The main claim is that continuous authentication has progressed much further technically than evaluatively and ethically the field boasts impressive equal error rates on curated data, but lacks standardized adversarial benchmarks, longitudinal field evidence, and consent frameworks commensurate with a technology whose very nature is that users can't tell when it's working. The review ends with a research agenda on privacy preserving architectures, disaggregated evaluation and regulatory alignment.

Keywords: Behavioral biometrics, Continuous authentication, Big data, Keystroke dynamics, Machine learning, Zero trust, Privacy, Adaptive authentication.

1. INTRODUCTION

Traditional authentication is episodic in that it is done at a point of entry password, fingerprint, one-time code, and all that comes afterwards within the session is assumed to be true. In this design, a vulnerability is known and built-in. The common thread in all these attacks is that the system's trust in a user's identity is greatest at the time of login, and diminishes, unchecked, after that. With longer sessions, increased value of the sessions, and remote work and mobile banking becoming a reality, point-in-time verification became a first order security issue and not just a theoretical problem.

Behavioral biometrics is a measurement of how a person behaves, not what a person has or what his or her body is. The rhythm of typing, the geometry and pressure of touchscreen swiping, gait from inertial sensors, sequence of applications used, and even the micro-timing of scrolling behaviors are all measurable and can be passively collected and are present throughout a session and not just at the beginning. These signals can be fed into machine-learning algorithms that are combined with the user's

behavioral profile that has been enrolled, resulting in continuous authentication a continuously running, probabilistic evaluation of whether the person at the keyboard is the person who logged in.

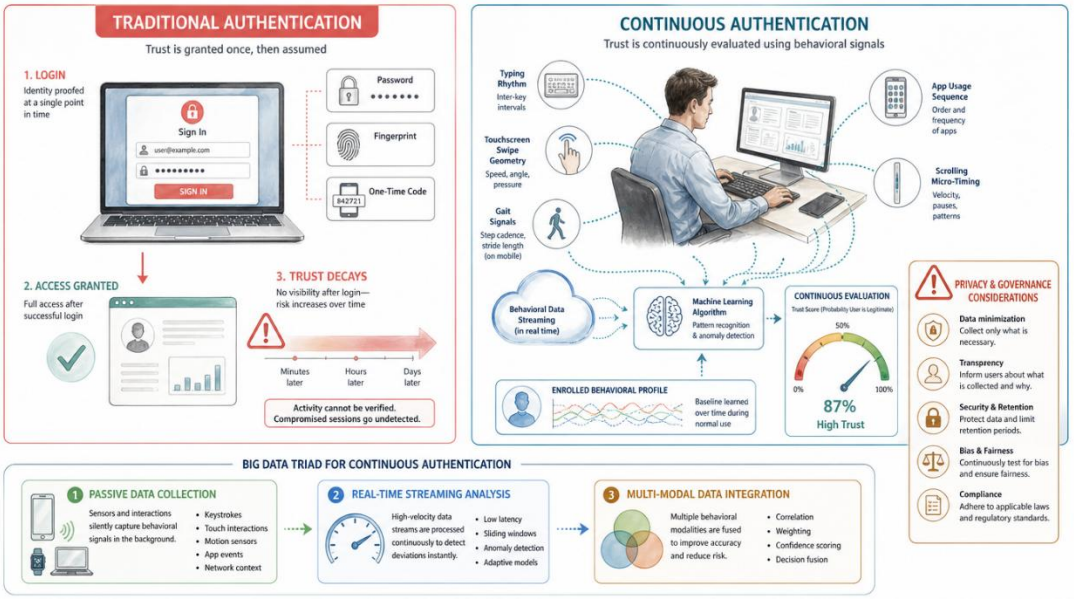

Fig -1: Traditional Authentication Vs Continuous Authentication

This would not be feasible without the use of big data. Behavioral signals are individually very weak a single keystroke interval discriminates almost nothing, and only identifying in the aggregate, across thousands of events, sessions and contexts. Thus, the paradigm is based on three essential features of the big-data triad passive collection of data from commodity sensors, streaming analysis of the data with low latency decisions, and fusion of data from different types of modalities. This reliance is the technology's driving force and its governance issue, as the data that makes continuous authentication possible is necessarily a running record of human activity, and has implications that extend well beyond authentication.

2. SCOPE AND METHOD

This is a narrative review based on four categories of sources:

1. Peer-reviewed surveys and primary studies in behavioral biometrics and continuous authentication 2017-2025.
2. Benchmark and dataset literature mainly in keystroke dynamics, touch dynamics and gait recognition.
3. Industry market and threat reporting with the caveat that vendor-originated figures are driven by commercial interests.
4. And regulatory and standards documents including GDPR treatment of biometric data, the EU AI Act and U.S. state biometric privacy statutes.

The sources cover the period of time from about 2015 to the middle of 2026. Quantitative performance claims are given with their evaluation context because in this field, headline accuracy figures are more

susceptible to the size of the dataset, conditions of text entry fixed versus free and evaluation protocol than in other fields.

3. THE BIG-DATA FOUNDATIONS OF BEHAVIORAL IDENTITY

3.1 From samples to streams

Physiological biometrics rely on samples such as a face image, a fingerprint impression, or an iris scan, which are discrete artifacts, acquired under controlled conditions. Behavioral biometrics are based on streams. The many tens of thousands of touch events, accelerometer readings at 50-100Hz, and continuous application context metadata, can be produced by a single hour of smartphone use. The identity signal is not contained in a single datum but rather in the statistical structure of the stream, such as distributions of key hold durations, inter key latencies, swipe curvature and velocity profiles and gait cycle periodicity. Hence the close match between the field's emergence and the big data era the storage and streaming infrastructure, and the learning algorithms used to derive stable identity from the noisy, high volume streams of behaviors simply did not exist at deployable cost until the 2010s.

3.2 The modality landscape

There are four modality families which have received the most attention in literature. The oldest and most studied is keystroke dynamics which captures typing timing hold time, up down and down-down latencies and, in the case of touchscreens, pressure and touch area it requires no special hardware, and has a large ecosystem of data sets, but performance is highly dependent on whether the text is fixed or free. Touch and swipe events on mobile devices provide information on the geometry, velocity, and pressure of gestures, and are one of the most naturally continuous signals on the mobile device. Motion and gait biometrics involve using phone and wearables mounted inertial sensors to record the way a person moves or walks, which can be used for authentication when the user is not actively interacting with the phone. Location routines, sequences of applications, network environments, typically distinct for weekdays and weekends, provide a fourth, less discriminative, but useful for fusion level of information. The common empirical result in all surveys is that the fusion of multiple modalities is superior to any one modality in terms of accuracy but is more difficult to collect data and more complex to build and maintain the system.

3.3 The economic and architectural context

Continuous behavioral authentication has gone from being a lab concept to a major reality in the market, and it has been primarily propelled by financial services. According to industry estimates, the total spending on behavioral biometrics in the world will be between $3.4 billion and $4.0 billion in 2026, and this amount is expected to increase by approximately 25% per year, with almost half of the spending projected to be in the banking sector and fraud detection as the top use case. The architectural paradigm is the "zero-trust" paradigm, which is a security doctrine that no user or device should be trusted by default and that verification should be ongoing and not based on perimeter. Behavioral biometrics is the only element that can be re-assessed at each moment without any friction on the user end, and is the natural authentication substrate for zero trust, with industry predicting that by the late 2020s most zero trust tooling will include behavioral and AI-driven signals. Another factor is the deepfake crisis in physiological biometrics, while the generation of face and voice images is now a commonplace occurrence, the generation of behavioural images is more difficult to accomplish, as it requires dynamic and interactive generation of images that cannot be synthesized as a static artefact, giving it a complementary status as a more difficult to attack layer.

4. ARCHITECTURE AND PERFORMANCE OF CONTINUOUS AUTHENTICATION

4.1 The canonical pipeline

Passive sensing at the device, feature extraction to convert raw events into timing, geometric, and statistical features, a per user model which is typically a deep architecture, such as CNNs for spatial touch patterns and LSTMs and transformers for temporal sequences, but ensemble methods like random forests are still competitive and much cheaper on-device, a streaming decision layer that outputs a rolling confidence score, and a policy layer that maps confidence to action are all deployed systems that converge on this common pipeline. The key usability insight of the paradigm is that continuous authentication is not binary but rather is graduated this insight is captured in the policy layer. As confidence wanes, friction is increased in a direct proportional manner, whether by limiting sensitive transactions, triggering re-verification or raising the flag to fraud teams, the vast majority of legitimate sessions will have zero additional friction.

4.2 Documented performance

In good circumstances the reported performance is indeed good. The EERs for fixed-text scenarios with ensemble classifiers are reported to be below 1% and mobile keystroke systems using pressure and touch area features have been reported to have EERs as low as 0.45%. On free text, the EERs of deep-learning approaches are in the range of 1–7% depending on the dataset and window size for fusion with gait. Must have three qualifications. First, most results are based on small populations dozens to hundreds of subjects and short collection times and hence extrapolation to the population level has not been demonstrated. Second, fixed-text performance is an overstatement of real-world, free-text performance. Third, and most importantly, the metrics used in the field are not the same as the ones that are operationally relevant. The field's metrics address the question can the model distinguish enrolled users from random others. while the operationally relevant question is, can the model detect a motivated impostor who has studied the target.

4.3 An emerging application frontier

Continuous behavioral authentication, which extends beyond banking to virtual and metaverse workplaces, IoT and wearable ecosystems, and intelligent personal agents that handle authentication decisions across device edge cloud tiers, is detailed in the literature of 2024-2026. The deployment of these models in resource constrained devices, such as wearables and IoT devices, is an active field of research, with the need to make explicit tradeoffs between model accuracy, energy consumption, and latency, driving lightweight models and split device edge architectures.

5. TECHNICAL LIMITATIONS AND FAILURE MODES

5.1 Behavioral drift and template aging

Behavior changes fatigue, injury, mood, medication, aging, new keyboard or just time, unlike a fingerprint. Research on samples taken weeks or months apart has demonstrated measurable degradation in the absence of adaptation, while adaptive schemes that train with the latest data have a security conundrum an adaptive template is one that can be walked, slowly morphing the template into one that represents the attacker's behavior by adding a little contamination at a time. Under adversarial assumptions, template-update policies that are robust remain an open question.

5.2 Cross-device and cross-context variability

Behavioral profiles are entangled with the hardware and context of enrollment. Keystroke timing differs across keyboards; touch dynamics differ across screen sizes and cases, gait differs across footwear and

terrain. Free text keystroke studies document meaningful keyboard dependency effects, and mobile systems often require context specific profiles for example, separate weekday and weekend mobility models, multiplying enrollment burden and failure surface.

5.3 Adversarial and generative attacks

The biggest gamechanger of the mid-2020s is that it has been proven that the intrinsically spoof resistant behavioral biometrics can be spoofed by trained imitators and by generative models. Some documented results are that the false acceptance rate is increased by 5.8–43.66 percentage points when the gait is imitated using a treadmill, and that the equal error rate is increased by 28–84% when compared to the baseline through the use of generative keystroke attacks, which are models that are trained to synthesize the typing statistics of a victim. Such attacks are still less commoditized, more difficult, and are truly a structural defense because they are interactive and closed loop, meaning that behavioral verification is performed. However, there is no standardized adversarial benchmarks and accredited attack resistance testing in the field today that comes from ISO/IEC 30107. This lack has been found in various systematic reviews as the most important methodological gap in the field. The problem is exacerbated by injection attacks that inject synthetic streams of behavioral events into the pipeline, below the sensor layer, through instrumented drivers or emulators, and for which there is no way to validate the origin of the stream.

5.4 Evaluation debt

There is considerable evaluation debt in the field lack of longitudinal field trials, lack of reporting on the demographic makeup of the field, lack of consistency in the protocols used, and lack of homogeneous datasets. It is difficult to compare EERs across different studies without comparison under similar conditions and there is no well-established counterpart to NIST's face recognition evaluation program for behavioral modalities. Until now, vendor performance claims are virtually impossible to audit.

6. PRIVACY, FAIRNESS, AND GOVERNANCE

6.1 The inference problem

Behavioral biometric data is qualitatively different from a face template, being inferentially promiscuous. These keystroke streams have been shown to be successful in identifying gender with more than 95% accuracy and related literature has shown that typing, gait and interaction patterns are related to neurological disorders specifically Parkinsonian motor symptoms, emotional state, fatigue, and age. Continuous authentication, therefore, is if not a health and demographics sensor by design certainly one by data content. This creates a purpose limitation issue which point in time biometrics did not have at a similar scale the data collected for security is constantly renewed, there is a longitudinal dimension, and it can be used for profiling, and the user cannot selectively opt out without losing access to the system.

6.2 The consent problem

The principle of continuous authentication is its principle of consent. If a system is designed to be imperceptible, then it cannot maintain meaningful consent as the user who consented to the collection at the time of enrollment will not be reminded of the collection afterwards. However, systematic reviews indicate that there is limited evidence on the consent and acceptance of such systems, within a framework like GDPR. The legal landscape is also unclear: GDPR considers biometric data used for unique identification as special category data which can only be processed with explicit consent or on one of the other very limited bases, but there is a debate about the difference between behavioral biometrics for identification special category data and behavioral analytics for fraud scoring arguably not special category data and

vendors have commercial interest in claiming that the latter i the case. The definitions of biometric identifiers in state laws like BIPA Illinois are drafted with fingerprints and face geometry in mind and thus raise further uncertainty about whether behavioral profiles are considered biometric identifiers under those laws.

6.3 Fairness and exclusion

There is no well-developed literature on demographic differential evaluation of face recognition that has been as influential in behavioral biometrics as the literature in face recognition has been in the field. The principled concerns are that the distributions of behaviours may plausibly be shifted by motor impairments, tremor, arthritis, use of assistive technology, and atypical interaction styles, and that a system trained on the normative distributions of interaction styles may be rejecting disabled and elderly users or suspecting them of fraud. Continuous systems are opaque, meaning that they affect a user without the user knowing that there has been a biometric judgment, which may lead to blocked transactions, reduced service or account flags that are not disclosed to the user.

6.4 Privacy preserving architectures

There is now a large technical literature that focuses on these risks cancelable behavioral templates and Bloom filter encodings that allow for revocation, on-device processing within trusted execution environments, so that raw behavioral streams never leave the device and zero knowledge protocols that allow proving identity claims without revealing behavioral features, which is already commercialized in FIDO certified products. The same pattern as in physiological biometrics prevails, though privacy preserving methods are published long before being used it is simply easier and cheaper to have a centralized system, and more convenient for the vendors to have a secondary use business model. The biometric breach record lesson is directly applicable, Profiles of behaviors stored in a central repository are a compromise of unlimited potential harm, and like a fingerprint, behaviors are not meaningfully revocable.

7. DISCUSSION OF AN ASSESSMENT OF THE PARADIGM

Evidence is arranged under four conclusions. For one, there is a paradigm shift, and it is warranted. This is a very real problem with point in time authentication, the zero-trust doctrine that continuous verification is a fundamental part of is now mainstream enterprise architecture, and behavioral signals are the only factor that can be re-checked continuously without any user friction. Continuous authentication should not be thought of as a replacement for primary credentials it cannot be used to log in or verify a user's identity before they start a session, so it is essentially a layer on top of passkeys or physiological biometrics, and is more of the connective tissue of layered, risk adaptive identity systems. Second, the maturity of technical is in front of the evaluative maturity. The field can achieve sub-1% equal error rates on curated benchmarks, but cannot yet show standardized adversarial robustness, disaggregation norms, or field performance over time because the benchmarks, disaggregation norms and evaluation infrastructure have not yet been developed. History indicates that it is independent, adversarial, disaggregated evaluation the NIST model that transforms a good idea into a good technology that becomes accountable, and behavioral biometrics has yet to have its NIST moment. Third, it is the governance crux that is dependent on big data. All the unique characteristics of continuous authentication the passivity, persistence and accuracy stems from continuous high volume behavioral collection, and so does all of the unique characteristics of its risk profile. The paradigm's sustainability thus relies on architectures which remove the need to store data on a device on device inference, federated adaptation and zero-knowledge assertion. These are there the only uncertainty is whether or not regulation and procurement will ensure that they are the default, and not the

exception, following the first big behavioral profile violation. Fourth, there must be compensating transparency in order to have invisibility. The disclosure, contestability and audit burden is proportionate to the value proposition, and imperceptibility means that the burden is greater. Where continuous behavioral scoring is used, users should be able to know that it is taking place, what it has resulted in, and challenge negative outcomes requirements that the EU AI Act transparency and human oversight requirements for high risk biometric systems start to put into law, but which existing deployments are far from fulfilling.

8. LIMITATIONS

There are three limitations in this review. Market size and adoption data are based on industry information that has a vested interest in the market's expansion and are therefore not necessarily accurate. The small cohort and short duration of most of the performance literature carries over to the characterization of achievable accuracy. The attack defense balance outlined here is a snapshot as of mid-2026 as the field continues to evolve, especially when it comes to the generative AI behavioral spoofing nexus.

9. CONCLUSION

Behavioral biometrics is made possible by big data, is made viable by continuous authentication and is silently changing the definition of logged in. The evidence presented here is in favor of qualified endorsement. Continuous behavioral authentication is a security architecture that can't be filled by any point in time mechanism and is critical to the security landscape, not just for its market share. However, the paradigm as such has several characteristics invisible collection, longitudinal profiling, inferentially rich data which make it governance risks, rather than algorithm risks. The research agenda that ensues is similarly institutional as well as technical standardized adversarial benchmarks and independent evaluation programs, disaggregated reporting by demographic and disability dimensions as a standard feature, privacy preserving architectures as a standard feature of procurement, instead of an academic option, and consent and transparency mechanisms that are tailored to a technology that is invisible to the user. Continuous authentication is here and the next five years will tell if it is to become continuous accountability.

REFERENCES

[1] George, D. (2026a). Digital dividend data taxation's potential to transform India's economy and redefine fiscal policy in the mobile era. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.19021896

[2] Gali M, Kassenkhan A, Chinibayev Y, Abshukirova A, Serbin V. Behavioral Biometric Continuous Authentication for Mobile Devices with an Intelligent Personal Agent: A Systematic Review. Technologies. 2026; 14(7):451. https://doi.org/10.3390/technologies14070451

[3] George, D., & Dr.T.Baskar. (2026). India's Unified Payments interface architecture, adoption, and the policy challenges of a public digital payment infrastructure. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.21891881

[4] Fares, H., Bakardjieva, T., & Ivanova, A. (2026). The Invisible Guardian: Big Data, Behavioral Biometrics, and the Era of Continuous Authentication. Engineering Proceedings, 150(1), 16. https://doi.org/10.3390/engproc2026150016

[5] Gupta, D. (2026). AI-Powered adaptive authentication and behavioral biometrics: The enterprise guide 2026. [online] Guptadeepak.com. Available at: https://guptadeepak.com/ai-powered-adaptive-authentication-and-behavioral-biometrics-the-enterprise-guide-2026/

[6] George, D. (2026b). Digital Inclusion vs. Fiscal Revenue Assessing India's Per-Gigabyte Data Tax. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.20922074
[7] Sasikumar K and Nagarajan S (2026) LSTM-based ensemble models for keystroke dynamics authentication: integrating explainable AI for transparency. Front. Artif. Intell. 9:1860248. doi: 10.3389/frai.2026.1860248
[8] George, D. (2025c). Sanchar Saathi Digital Security versus Civil Liberty in India 's Smartphone Era. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.17838468
[9] de-Marcos, L., Martínez-Herráiz, J.-J., Junquera-Sánchez, J., Cilleruelo, C., & Pages-Arévalo, C. (2021). Comparing Machine Learning Classifiers for Continuous Authentication on Mobile Devices by Keystroke Dynamics. Electronics, 10(14), 1622. https://doi.org/10.3390/electronics10141622
[10] George, D. (2025b). Digital Watermarking in Cloud Environments for Copyright Protection: A Comprehensive review. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.17726895
[11] Martins, A., Dias, T., Dias, A. et al. Keystroke dynamics for intelligent biometric authentication with machine learning. Discov Appl Sci 7, 992 (2025). https://doi.org/10.1007/s42452-025-07449-5
[12] George, D. (2025a). Beyond the dashboard critical perspectives on digital employee monitoring and the paradox of productivity measurement in contemporary organizations. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.17702042
[13] Lamiche, I., Bin, G., Jing, Y. et al. A continuous smartphone authentication method based on gait patterns and keystroke dynamics. J Ambient Intell Human Comput 10, 4417–4430 (2019). https://doi.org/10.1007/s12652-018-1123-6
[14] George, D., Dr.T.Baskar, & Dr.M.M.Karthikeyan. (2026). The Grammar of Icons: decoding the origins, meanings, and cognitive power of everyday digital and cultural symbols. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.19959167
[15] AbdelRaouf, H., Chelloug, S. A., Muthanna, A., Semary, N., Amin, K., & Ibrahim, M. (2023). Efficient Convolutional Neural Network-Based Keystroke Dynamics for Boosting User Authentication. Sensors, 23(10), 4898. https://doi.org/10.3390/s23104898
[16] George, D., George, A., Dr.T.Baskar, & Pandey, D. (2026). Digital Dependency and Temporal Distortion: A Critical review of smartphone use, cognitive impact, and Behavioral intervention in Post-Pandemic Society. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.19864919
[17] Ganiga, D. V. (2026, July 30). Beyond the password: Why behavioral biometrics is becoming banking's last line of defense. Biometric Update | Biometrics News, Companies and Explainers. https://www.biometricupdate.com/202607/beyond-the-password-why-behavioral-biometrics-is-becoming-bankings-last-line-of-defense
[18] George, D., & George, A. (2025). How artificial intelligence systems function as digital migrants creating more profound societal disruption than human immigration. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.16112307
[19] Gupta, D. (2026, May 25). AI-Powered Adaptive Authentication and Behavioral Biometrics: The Enterprise Guide 2026. Deepak Gupta. https://guptadeepak.com/ai-powered-adaptive-authentication-and-behavioral-biometrics-the-enterprise-guide-2026/
[20] George, D., Dr.T.Baskar, & Siranchuk, D. (2026). Digital Addiction in Children: economic impact, global age restrictions, and protective solutions. Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo.18818134
[21] Martins, A., Dias, T., Dias, A., Vitorino, J., Maia, E., & Praça, I. (2025). Keystroke dynamics for intelligent biometric authentication with machine learning. Discover Applied Sciences, 7(9). https://doi.org/10.1007/s42452-025-07449-5
[22] Shadman, R., Wahab, A. A., Manno, M., Lukaszewski, M., Hou, D., & Hussain, F. (2023). Keystroke Dynamics: Concepts, techniques, and applications. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2303.04605
[23] Volobuyeva, O. (2026). Biometric Security Trends 2025: Fusion models and Behavioral Indicators. International Journal of Innovative Science and Research Technology (IJISRT), 2687. https://doi.org/10.38124/ijisrt/25dec1561