AI-Powered Zero Trust Access Evaluation using Behavioral Fingerprinting

Hritesh Yadav
Independent Researcher
California

Ganapathy Subramanian Ramachandran
Independent Researcher
California

Kshitij Sharma
Independent Researcher
California

ABSTRACT
In today's cybersecurity landscape, the traditional perimeter-based defense model has become obsolete, giving rise to the Zero Trust Architecture (ZTA), where no entity—whether internal or external—is automatically trusted. While ZTA provides a robust security posture, its effectiveness heavily depends on accurate and context-aware access evaluation. Conventional authentication techniques, such as static credentials and multi-factor authentication (MFA), are often insufficient to detect subtle identity compromise or insider threats.

This paper introduces a novel framework that leverages Artificial Intelligence (AI) and behavioral fingerprinting to enable continuous and adaptive access evaluation within a Zero Trust environment. Behavioral fingerprinting, which includes unique user-specific patterns such as keystroke dynamics, mouse movement patterns, application access sequences, and response times, is used to construct a dynamic trust profile for each user. Our system continuously collects telemetry data, extracts behavioral features, and uses supervised and unsupervised learning models to assess risk in real-time. By combining these insights with contextual parameters (such as geolocation, device hygiene, and network indicators), our AI engine computes a Behavioral Trust Score (BTS) to grant, deny, or conditionally allow access.

The results from our prototype demonstrate a significant improvement in detecting anomalous behavior compared to traditional rule-based systems, with a notable reduction in false positives and latency. Our contributions aim to enhance the granularity and responsiveness of Zero Trust security models while maintaining user transparency and compliance.

Keywords
Zero Trust Architecture, Behavioral Fingerprinting, Adaptive Access Control, Behavioral Trust Score, User Behavior Analytics, Insider Threat Detection, Continuous Authentication, Federated Learning, AI in Cybersecurity

1. INTRODUCTION
As enterprises face increasingly sophisticated cyber threats and move toward perimeter-less architectures, the **Zero Trust model** has become a cornerstone of modern cybersecurity strategy. Unlike legacy models that rely on implicit trust based on network location or one-time authentication, Zero Trust operates on the principle of "never trust, always verify"—treating every access request as potentially malicious and requiring continuous validation based on user identity, context, and behavior [1], [2]. This paradigm shift is fueled by the widespread adoption of **hybrid work models**, **Bring Your Own Device (BYOD)** policies, and **cloud-hosted services**, all of which have contributed to making identity the new security perimeter [3]. While modern authentication mechanisms such as **Multi-Factor Authentication (MFA)** provide an added layer of protection, they remain susceptible to credential phishing, session hijacking, and insider threats [4], [5].

Static or one-time validation methods are increasingly inadequate when the threat landscape is dynamic and capable of evolving during an active user session. Consider a scenario where an attacker hijacks a valid authenticated session or where a legitimate insider begins to deviate from normal behavior patterns. In such instances, a Zero Trust framework must extend its capability to **continuously evaluate user trust** based on real-time behavioral signals. To address this challenge, we propose a novel approach that leverages **behavioral fingerprinting**—the non-intrusive capture and analysis of user-specific interaction patterns—as a foundation for adaptive access control [6], [7]. When combined with **AI-driven anomaly detection**, this behavioral insight enables the system to compute real-time **risk profiles** and make intelligent access decisions without interrupting legitimate user activity [8], [9].

This paper makes the following key contributions

(1) We present a detailed system architecture that incorporates real-time behavioral fingerprinting into Zero Trust access models; (2) We propose a multi-model AI framework that computes a Behavioral Trust Score (BTS) by analyzing both historical behavior and real-time interaction data; (3) We demonstrate integration strategies with modern Zero Trust Network Access (ZTNA) and Secure Access Service Edge (SASE) platforms to enhance dynamic policy enforcement; (4) We provide empirical evaluation results that show significant improvements in anomaly detection accuracy while maintaining low system overhead; and (5) We explore ethical and practical considerations in deploying behavior-based security systems at scale, and outline promising directions for future research in this domain.

2. RELATED WORK
A. Zero Trust Access Control Models
Zero Trust Architecture (ZTA) has been formalized through NIST's Special Publication 800-207, which underscores the importance of continuously validating access based on identity, device posture, and context [1]. Traditional access models such as Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) rely on static policies that fail to adapt to dynamic threat scenarios [10]. Enterprise frameworks like Google's BeyondCorp and Microsoft's Zero Trust Maturity Model enhance access decisions by incorporating telemetry signals such as device health, location, and authentication anomalies [3], [4]. However, these approaches primarily assess access at the session initiation phase. Once access is granted, user behavior is not continuously monitored, leaving systems exposed to risks such as lateral movement, session hijacking,

and post-authentication misuse.

B. Behavioral Biometrics and Fingerprinting

Behavioral biometrics use subtle, non-intrusive traits such as keystroke dynamics, mouse movement, and app navigation to continuously assess user identity [6], [7]. These behavior-based identifiers are inherently difficult to spoof, offering an advantage over static credentials or even one-time authentication mechanisms like MFA [5]. Prior studies have demonstrated success in using machine learning classifiers such as Random Forests, Support Vector Machines (SVMs), and Hidden Markov Models (HMMs) for user verification based on typing and interaction patterns [6], [7]. However, most of these systems operate in isolation or depend on periodic re-authentication, which may interrupt usability. Our work builds upon this by fusing multiple behavioral modalities into a real-time behavioral trust score, which is actively integrated into the access control pipeline. This enables dynamic privilege adjustment throughout a session, moving beyond the re-authentication-focused models of prior literature.

C. AI in Adaptive Access Systems

Artificial Intelligence (AI) and Machine Learning (ML) are widely used in cybersecurity for anomaly detection and behavior-based threat prediction [8], [9]. In the identity and access management space, commercial solutions like Azure AD Identity Protection and Okta's Behavioral Detection apply supervised models to assign risk scores to login attempts [4]. However, these tools generally focus on pre-access evaluation and do not monitor session-level behavioral patterns. Academic literature has explored unsupervised learning techniques such as clustering, Principal Component Analysis (PCA), and autoencoders to uncover unknown or evolving threats. Our approach extends these efforts by integrating both supervised and unsupervised models, ensuring precision in recognizing known anomalies while adapting to new behavioral drift patterns [11].

D. Integration with SASE and ZTNA Platforms

Secure Access Service Edge (SASE) and Zero Trust Network Access (ZTNA) platforms are becoming the default for enforcing identity-aware, perimeter-less access across enterprise networks [4]. These platforms typically ingest telemetry from endpoint agents, network signals, and cloud environments, providing a solid foundation for behavior-aware access control. While current ZTNA implementations enforce access based on identity and device posture, they often do not incorporate real-time behavioral signals. Our system enhances these platforms by embedding a **behavioral trust engine** that evaluates live session activity, enabling context-aware policy decisions and improving the system's ability to respond to active threats dynamically.

3. ARCHITECTURAL OVERVIEW

The architecture of our proposed system is purpose-built to support Zero Trust principles by continuously evaluating access through AI-powered behavioral fingerprinting. It achieves this by operating as a modular and scalable framework that is adaptable across on-premises, cloud-native, and hybrid infrastructures. Drawing inspiration from adaptive security platforms like ZTNA and SASE [4], [9], [13], the system enables real-time user session analysis without imposing friction on legitimate users. At the heart of this architecture lies a tightly integrated data and decision pipeline. It starts with the collection of user telemetry—typing cadence, mouse dynamics, interaction sequences, and contextual signals like process activity and IP geolocation. These data streams are locally filtered and securely transmitted using TLS 1.3, preserving both efficiency and privacy.

Once collected, the behavioral data is converted into meaningful signals using techniques such as time-series segmentation, entropy modeling, and Fourier transforms. This step transforms noisy human behavior into structured vectors that can be analyzed in real time. The AI-based risk evaluation engine processes these vectors through a combination of supervised and unsupervised models. Random Forests and SVMs handle known behavioral patterns, while autoencoders and clustering algorithms help identify outliers or zero-day anomalies [6], [8], [11]. Model interpretability is maintained using SHAP values, which allow system analysts to review and audit trust decisions [17]. The output of this pipeline is a Behavioral Trust Score (BTS), a continuously updated metric ranging from 0 to 100. Depending on the score, users may receive full access, conditional access, or complete restriction. These decisions are dynamically enforced through integrations with ZTNA or SASE systems like Prisma Access, Zscaler, and Microsoft Entra [4], [9].

A final, but critical, component is the system's learning and feedback loop. It incorporates input from security analysts and adjusts to long-term behavior shifts—such as job role changes or new tools—using reinforcement and federated learning strategies [19]. This makes the architecture not only resilient but continuously evolving, providing adaptive, real-world protection that aligns with user context and behavior. By harmonizing telemetry, AI risk analysis, trust scoring, and real-time enforcement, the system delivers a Zero Trust implementation that is not just reactive but predictive and self-improving.

Table 1. Table captions should be placed above the table

Graphics	Top	In-between	Bottom
Tables	End	Last	First
Figures	Good	Similar	Very well

4. BEHAVIORAL FINGERPRINTING MODEL

Behavioral fingerprinting enables our system to continuously distinguish users based on how they interact with digital systems. Unlike static biometrics, these behavioral traits are dynamic and context-sensitive, capturing elements like keystroke timing, mouse movement, application usage patterns, and geolocation. These signals are collected passively and securely transmitted for analysis. Using normalization, segmentation, and statistical modeling, the system transforms raw behavior into structured feature vectors. Advanced techniques such as ARIMA and DTW help model temporal behavior [15], [16], while Random Forest and SVM classifiers handle known patterns [6], [8]. Autoencoders and Isolation Forests assist in detecting zero-day or anomalous behavior [11], [16]. Ensemble decision models further refine accuracy while minimizing false positives [13], [20].

Each session's behavior is compared with historical norms using cosine similarity or Mahalanobis distance. Deviations may trigger alerts or retraining, supported by cumulative drift scoring and heatmap visualization. To ensure privacy, the system implements on-device preprocessing and anonymized data transmission while adhering to GDPR and CCPA regulations [19]. Altogether, this model supports real-time identity assurance in Zero Trust environments without interrupting legitimate user activity.

5. ZERO TRUST ACCESS EVALUATION

Traditional access control mechanisms rely on static user attributes and device posture checks, which often fall short in detecting insider threats or session-level anomalies. To overcome these limitations, our framework introduces a real-time, behavior-driven trust model built around the Behavioral Trust Score (BTS). BTS is a dynamic metric ranging from 0 to 100, computed continuously during a user session. It assesses how closely the current behavior aligns with historical norms while accounting for context (e.g., device and network changes) and behavioral drift [6], [9], [13]. The BTS formula weighs three components—behavioral similarity, contextual consistency, and drift adjustment—using tunable coefficients that reflect organizational risk posture. This model is updated every 30–60 seconds and integrated into real-time policy enforcement systems.

Based on the computed BTS, users are categorized into risk tiers that inform access decisions. A high score allows full access, a moderate score may require additional authentication or restricted privileges, and a low score triggers session termination or isolation. These decisions are enforced via existing Policy Enforcement Points (PEPs) such as ZTNA gateways, identity providers, CASBs, and endpoint agents [4]. If a user's behavior shifts mid-session—such as accessing an admin panel without proper authorization or exhibiting typing anomalies—the BTS drops and corresponding enforcement actions are triggered. The real-time flow begins at login, continues with telemetry monitoring, BTS computation, and culminates in policy enforcement without disrupting the user unnecessarily. This system offers adaptive access, continuous risk evaluation, and smooth integration with enterprise security infrastructure [14].

Table 1: Users mapping based on BTS

BTS Range	Risk Level	Action
86–100	Low Risk	Full access granted
60–85	Medium Risk	Step-up authentication/ restricted access
< 60	High Risk	Access denied/session terminated

6. EVALUATION AND RESULTS

To theoretically validate our AI-powered behavioral fingerprinting framework within a Zero Trust architecture, we conducted a comprehensive experimental study using synthetic data designed to emulate enterprise usage patterns. The prototype system included client-side telemetry collectors, a Kubernetes-based processing backend, and a GPU-accelerated inference engine optimized for real-time scoring. The synthetic dataset included a wide array of simulated user behaviors and attack scenarios such as bot activity, credential sharing, session hijacking, and insider misuse. This allowed us to rigorously evaluate the system's anomaly detection capabilities in a controlled and repeatable environment [6], [7].

Our evaluation relied on established metrics including accuracy, precision, recall, false positive rate, and mean time to detect (MTTD). The behavioral fingerprinting model consistently outperformed traditional Role-Based Access Control (RBAC) and Contextual Access Control baselines. Notably, our model achieved a detection accuracy of 94.7% and reduced false positives to 2.8%. It also demonstrated adaptability, effectively reclassifying drifted behavior within two to three days. In a case study involving a simulated session hijack, the model flagged deviations—such as unusual keystroke timing and navigation behavior—within 90 seconds, triggering immediate session termination and forensic alerts [13], [16].

While the system proves highly effective in real-time access evaluation, it has certain limitations. New users experience a cold start period due to limited historical data, and advanced attackers may attempt to mimic behavioral traits. Furthermore, hybrid environments such as virtual desktops may introduce inconsistencies in telemetry. These challenges will be addressed in future work through federated learning [19], adversarial training with generative models, and adaptive retraining protocols.

Table 2: Summary of Theoretical Evaluation Results

Metric	RBAC	Contextual Model	Proposed Model
Accuracy	78.5%	86.4%	94.7%
Precision	74.2%	85.1%	92.3%
Recall	70.6%	82.5%	96.1%
F1 Score	72.3%	83.8%	94.2%
False Positive Rate	9.4%	5.6%	2.8%
Mean Time to Detect	N/A	47.2 sec	18.5 sec

7. DISCUSSION

The empirical findings from our evaluation underscore the promise of AI-powered behavioral fingerprinting in improving Zero Trust access frameworks. This section highlights the most important contributions and real-world trade-offs of our approach.

A. Strengths and Innovations

Our system enables continuous access monitoring by integrating behavioral telemetry into the security lifecycle—detecting session hijacks, credential misuse, and insider threats in real-time [6]. It passively captures high-dimensional signals like typing cadence and mouse movement to build trust profiles that are difficult to spoof, making the system resilient to phishing and social engineering [5], [7]. Unlike frequent prompts used in traditional MFA, our approach allows seamless operation unless risk thresholds are exceeded, which is particularly helpful in high-frequency operational environments [4]. Finally, its modular design ensures compatibility with existing SASE and ZTNA platforms, enabling scalable deployment across enterprise systems [9].

B. Challenges and Trade-offs

Despite its robustness, the system faces operational challenges. New users lacking behavioral history trigger a “cold start” effect, which we address through contextual bootstrapping. There is also theoretical susceptibility to mimicry attacks, prompting future integration of GAN-based anomaly detectors [12]. Model drift due to evolving user behavior remains an ongoing challenge, and the frequency of retraining must be carefully tuned [16]. Additionally, even with anonymization, the collection of behavioral data raises privacy concerns that we address through transparency, opt-in consent, and legal compliance [19].

C. Ethical Implications

Behavioral monitoring brings ethical questions around user autonomy and fairness. Over-monitoring could affect trust and morale, and AI model bias could lead to disproportionate treatment. These risks demand fairness testing and the use of explainable models such as SHAP for auditability [17].

D. Real-World Applicability

The system's design is especially relevant to sectors like government, finance, healthcare, and remote work environments where secure session management is crucial [1], [18]. For enterprises already deploying Zero Trust strategies, our model enhances adaptive policy enforcement without needing a complete overhaul. With minimal user disruption and strong security posture, it shows practical viability across modern enterprise settings.

8. CONCLUSION

As digital infrastructures become increasingly decentralized and cyber threats grow in sophistication, traditional perimeter-based security approaches are no longer sufficient. In response, organizations are embracing the Zero Trust paradigm, which advocates for continuous verification and contextual access decisions rather than relying on implicit trust based on network boundaries or one-time authentication [1], [4]. This paper presented a comprehensive AI-driven framework for Zero Trust access evaluation using behavioral fingerprinting. Our system captures and analyzes a range of behavioral signals—including typing rhythms, cursor movements, application usage flows, and contextual indicators like location and device state—to compute a real-time Behavioral Trust Score (BTS). These scores inform dynamic access decisions, enabling organizations to adapt security posture continuously and automatically.

We demonstrated that our approach achieves strong performance across key metrics, with a detection accuracy of 94.7%, low false positive rates, and a mean threat detection time of under 20 seconds. The system's modular design ensures seamless integration into enterprise environments and supports large-scale deployments with minimal latency. While effective, the system presents limitations such as cold-start behavior for new users and susceptibility to mimicry attacks. Future enhancements, including federated learning, adversarial training using GANs, and dynamic trust graph modeling, can further strengthen its applicability. Overall, this work contributes a scalable and privacy-conscious solution to evolving Zero Trust strategies, embedding intelligent, behavior-aware mechanisms at the core of adaptive access control.

9. REFERENCES

[1] Stafford, V. "Zero trust architecture." *NIST special publication* 800.207 (2020): 800-207.

[2] Ferraiolo, David F., et al. "Proposed NIST standard for role-based access control." *ACM Transactions on Information and System Security (TISSEC)* 4.3 (2001): 224-274.

[3] Google, "BeyondCorp: A New Approach to Enterprise Security," Google White Paper, 2014. [Online]. Available: https://cloud.google.com/beyondcorp

[4] Ward, Rory, and Betsy Beyer. "Beyondcorp: A new approach to enterprise security." *; login:: the magazine of USENIX & SAGE* 39.6 (2014): 6-11.

[5] Das, Sanchari, Andrew Dingman, and L. Jean Camp. "Why Johnny doesn't use two factor a two-phase usability study of the FIDO U2F security key." *Financial Cryptography and Data Security: 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26–March 2, 2018, Revised Selected Papers 22*. Springer Berlin Heidelberg, 2018.

[6] Ahmed, Ahmed A., and Issa Traore. "Biometric recognition based on free-text keystroke dynamics." *IEEE transactions on cybernetics* 44.4 (2013): 458-472.

[7] Kamezaki, Yuto, and Kazutaka Matsuzaki. "User Identification Based on Mouse Behavior--Enhancing Accuracy with Velocity Features and Evaluating Practicality." *IEICE Technical Report; IEICE Tech. Rep.*

[8] Khanan, Akbar, et al. "From bytes to insights: a systematic literature review on unraveling IDS datasets for enhanced cybersecurity understanding." *IEEE Access* (2024).

[9] Gadde, Hemanth. "AI-Driven Anomaly Detection in NoSQL Databases for Enhanced Security." *International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence* 14.1 (2023): 497-522.

[10] Sandhu, Ravi S. "Role-based access control." *Advances in computers*. Vol. 46. Elsevier, 1998. 237-286.

[11] Jazzar, Mahmoud, and Aman Jantan. "A novel soft computing inference engine model for intrusion detection." *IJCSNS International Journal of Computer Science and Network Security* 8.4 (2008): 1-9.

[12] Mahoney, Matthew V., and Philip K. Chan. "An analysis of the 1999 DARPA/Lincoln Laboratory evaluation data for network anomaly detection." *International Workshop on Recent Advances in Intrusion Detection*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003.

[13] Liang, Yunji, et al. "Behavioral biometrics for continuous authentication in the internet-of-things era: An artificial intelligence perspective." *IEEE Internet of Things Journal* 7.9 (2020): 9128-9143.

[14] Nasir, Rida, et al. "Behavioral based insider threat detection using deep learning." *IEEE Access* 9 (2021): 143266-143274.

[15] Bereziński, Przemysław, Bartosz Jasiul, and Marcin Szpyrka. "An entropy-based network anomaly detection method." *Entropy* 17.4 (2015): 2367-2408.

[16] Abuhamad, Mohammed, et al. "Sensor-based continuous authentication of smartphones' users using behavioral biometrics: A contemporary survey." *IEEE Internet of Things Journal* 8.1 (2020): 65-84.

[17] Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." Advances in neural information processing systems 30 (2017).

[18] MacDonald, Neil, Lawrence Orans, and Joe Skorupa. "The Future of Network Security Is in the Cloud." *Gartner. Viitattu* 1 (2019): 2021.

[19] Yang, Qiang, et al. "Federated machine learning: Concept and applications." *ACM Transactions on Intelligent Systems and Technology (TIST)* 10.2 (2019): 1-19.

[20] Killourhy, Kevin S., and Roy A. Maxion. "Comparing anomaly-detection algorithms for keystroke dynamics." *2009 IEEE/IFIP international conference on dependable systems & networks*. IEEE, 2009.