

Adaptive AI Frameworks for Retirement Account Security: Integrating Behavioral Biometrics with Dynamic Graph-Based Fraud Detection

D. Westbrook

University of Toronto, Canada

ARTICLE INFO

Article history:

Submission: January 01, 2026

Accepted: January 16, 2026

Published: February 02, 2026

VOLUME: Vol.11 Issue 02 2026

Keywords:

Behavioral biometrics, retirement account security, graph neural networks, financial fraud detection, dynamic graphs, artificial intelligence

ABSTRACT

The rapid digitization of retirement account management has introduced unprecedented efficiencies alongside complex security vulnerabilities. Among these, fraud targeting defined contribution retirement plans, particularly four zero one k accounts, has emerged as a critical concern due to their high monetary value, long-term nature, and increasing exposure to remote access channels. Traditional authentication mechanisms and rule-based fraud detection systems have proven insufficient against adaptive adversaries who exploit behavioral mimicry, credential stuffing, and low-homophily transaction networks. Recent scholarship has therefore turned toward artificial intelligence-driven behavioral biometrics and graph-based learning paradigms as promising solutions for strengthening retirement account security. This article develops a comprehensive theoretical and methodological synthesis of these approaches, grounded strictly in existing literature, with particular emphasis on behavioral biometrics as articulated in contemporary retirement security research and on graph neural network innovations in financial fraud detection.

The study integrates insights from behavioral modeling, attention-based neural architectures, temporal and dynamic graph learning, and barely supervised fraud detection to propose an advanced conceptual framework for protecting retirement accounts. Behavioral biometrics, encompassing keystroke dynamics, interaction rhythms, and longitudinal user patterns, are examined as a continuous authentication layer capable of detecting subtle deviations indicative of account compromise. These techniques are situated within a broader graph-based fraud intelligence ecosystem that captures relational, temporal, and contextual dependencies across users, devices, and transactions. The analysis critically engages with debates surrounding deep learning efficacy on tabular data, interpretability versus performance trade-offs, and the ethical implications of pervasive behavioral monitoring.

Methodologically, the article elaborates a text-based, publication-ready research design that combines behavioral feature extraction with dynamic graph representation learning, leveraging self-attention and transformer-based mechanisms to address low-label regimes and evolving fraud strategies. Results are discussed descriptively, drawing on comparative interpretations of prior empirical findings to demonstrate how integrated behavioral and graph-based models enhance detection robustness, reduce false positives, and adapt to adversarial drift. The discussion section offers an extensive theoretical interrogation of limitations, counter-arguments, and future research directions, including zero trust architectures, blockchain integration, and regulatory alignment.

By synthesizing behavioral biometrics and graph neural fraud detection within the specific context of retirement account security, this article contributes a deeply elaborated academic perspective that advances both theoretical understanding and applied security design. The work underscores the necessity of multidisciplinary, AI-driven defenses to safeguard long-term financial well-being in an increasingly hostile digital environment.

INTRODUCTION

The transformation of financial services through digital platforms has fundamentally altered how individuals interact with long-term savings instruments, including employer-sponsored defined contribution retirement plans. While this transformation has enhanced accessibility, transparency, and administrative efficiency, it has simultaneously expanded the attack surface for financial fraud, particularly in the context of four zero one k account management systems. These systems are increasingly accessed through web and mobile interfaces, often relying on static credentials and limited multi-factor authentication, thereby creating fertile ground for sophisticated adversaries who exploit both technical vulnerabilities and human behavioral patterns (Valiveti, 2025). The introduction of artificial intelligence into fraud detection and account security has been widely proposed as a necessary evolution, yet the theoretical and practical integration of AI-driven behavioral biometrics with relational fraud intelligence remains an underexplored domain.

Behavioral biometrics refers to the measurement and analysis of patterns in human behavior as a means of identity verification and anomaly detection. Unlike physiological biometrics, behavioral signals are inherently dynamic, context-dependent, and continuous, encompassing interaction rhythms, navigation habits, and temporal usage patterns. In the context of retirement account security, behavioral biometrics offers the promise of continuous authentication without imposing additional cognitive or procedural burdens on legitimate users (Valiveti, 2025). This promise is particularly salient given the demographic diversity of retirement account holders, whose technological proficiency and tolerance for friction vary widely. However, behavioral biometrics alone cannot fully capture the complex relational structures through which fraud propagates, such as coordinated attacks across multiple accounts, devices, or networks.

Parallel to the rise of behavioral biometrics, the field of financial fraud detection has witnessed a shift toward graph-based learning models that explicitly represent relationships among entities. Graph neural networks, attention mechanisms, and temporal graph frameworks have demonstrated superior performance in capturing non-Euclidean structures and dynamic interactions inherent in financial ecosystems (Velickovic et al., 2018; Rossi et al., 2020). These models are particularly valuable in scenarios characterized by low homophily, sparse labels, and adversarial adaptation, all of which are common in retirement account fraud (Wang et al., 2023). Yet, much of the existing literature focuses on transactional fraud in payment systems or credit card networks, leaving a conceptual gap regarding their applicability to long-horizon retirement accounts.

The convergence of behavioral biometrics and graph-based fraud intelligence thus represents a critical frontier in securing retirement assets. This convergence raises fundamental theoretical questions about representation learning, temporal dynamics, and the balance between individual-level behavioral signals and network-level relational patterns. It also invites practical considerations regarding data governance, privacy, and regulatory compliance, especially in jurisdictions with stringent data protection frameworks. Scholars have debated whether deep learning architectures are necessary or even appropriate for tabular and behavioral data, arguing that simpler models may suffice under certain conditions (Shwartz-Ziv and Armon, 2022). Others contend that the complexity and adaptiveness of modern fraud necessitate sophisticated architectures capable of modeling high-dimensional interactions and temporal evolution (Lin et al., 2024).

Against this backdrop, the present article undertakes an extensive theoretical elaboration of AI-driven behavioral biometrics and graph-based fraud detection as applied to retirement account security. Grounded strictly in the provided references, the study seeks to articulate a coherent research narrative that bridges disciplinary silos and addresses the specific vulnerabilities of four zero one k systems. The introduction situates the problem within broader trends in digital finance, reviews relevant theoretical foundations, and identifies a clear literature gap concerning integrated behavioral and relational security models for retirement accounts (Valiveti, 2025). By doing so, it establishes the intellectual foundation for a detailed methodological exposition and interpretive analysis that follow in subsequent sections.

The urgency of this inquiry is underscored by the increasing prevalence of credential stuffing, session replay attacks, and automated bot activity targeting financial platforms (Barkworth et al., 2022; Sadeghpour, 2024). Retirement accounts are particularly attractive targets due to their perceived lower monitoring frequency and

the complexity of withdrawal processes, which can obscure fraudulent activity until substantial losses occur. Traditional perimeter-based security models and post hoc fraud detection mechanisms are ill-suited to this environment, prompting calls for zero trust architectures and continuous risk assessment frameworks (Idialu, 2025). Behavioral biometrics, when embedded within such architectures, offers a means of continuously reassessing trust based on real-time user behavior rather than static credentials alone (Valiveti, 2025).

Furthermore, the relational nature of fraud campaigns necessitates analytical tools that move beyond isolated user profiles to consider network-level patterns. Graph-based models enable the detection of coordinated behavior, shared infrastructure, and temporal sequences that may not be apparent through individual analysis alone (Yu et al., 2024). The integration of attention mechanisms within these models allows for adaptive weighting of relationships, aligning with theoretical insights from cognitive science regarding selective focus and relevance (Vaswani et al., 2017). These theoretical underpinnings provide a rich conceptual palette for reimaging retirement account security as a dynamic, learning-driven process.

Despite these advances, the literature remains fragmented, with behavioral biometrics and graph learning often treated as separate research streams. This fragmentation limits the development of holistic security frameworks capable of addressing both individual-level impersonation and network-level fraud orchestration. The present study addresses this gap by synthesizing these streams into a unified theoretical model, explicitly tailored to the unique characteristics of retirement accounts. In doing so, it responds to calls for more context-sensitive fraud detection research that accounts for domain-specific constraints and threat models (Teimoory, 2025).

The remainder of this article proceeds as follows. The methodology section provides a highly detailed, text-based research design that outlines how behavioral biometric features and graph representations can be jointly leveraged within an AI-driven security architecture. The results section offers a descriptive and interpretive analysis of expected outcomes, grounded in comparative readings of existing empirical studies. The discussion section engages in deep theoretical interpretation, scholarly debate, and critical reflection on limitations and future research directions. The conclusion synthesizes key insights and underscores the broader implications for financial security research and practice. Throughout, the study maintains a rigorous academic tone and adheres strictly to the cited literature, with particular attention to the foundational contribution of behavioral biometrics in retirement account security articulated by Valiveti (2025).

METHODOLOGY

The methodological foundation of this study is conceptual and integrative, designed to synthesize existing empirical and theoretical insights into a coherent research framework for AI-driven retirement account security. Rather than presenting new empirical data, the methodology articulates a detailed design logic that could underpin future implementation and evaluation efforts. This approach aligns with prior conceptual studies in financial fraud detection that emphasize architectural reasoning and methodological transparency as prerequisites for robust empirical validation (Rossi et al., 2020). Central to the methodology is the integration of behavioral biometric analysis with graph-based representation learning, situated within a zero trust security paradigm tailored to four zero one k account systems (Valiveti, 2025).

The first methodological component concerns behavioral biometric feature modeling. Behavioral data in retirement account platforms typically include keystroke timing, mouse movement trajectories, navigation sequences, session duration patterns, and device interaction rhythms. These features are inherently temporal and user-specific, requiring modeling techniques that capture both short-term fluctuations and long-term behavioral baselines (Arora and Kanji, 2019). Drawing on behavioral modeling literature, the methodology assumes that each user exhibits a stable yet evolving behavioral signature that can be learned over time and used for continuous authentication (Valiveti, 2025). Feature extraction is conceptualized as a multi-layer process, beginning with raw interaction logs and progressing toward higher-level behavioral embeddings.

In designing this feature extraction process, the methodology acknowledges debates regarding the necessity of deep learning for tabular and behavioral data. Critics argue that deep architectures may introduce unnecessary complexity and reduce interpretability without commensurate performance gains (Shwartz-Ziv and Armon, 2022). In response, the proposed framework emphasizes hybrid modeling, wherein deep learning components are selectively applied to capture temporal dependencies and non-linear interactions, while simpler statistical descriptors are retained for transparency and baseline comparison. This design choice reflects a pragmatic

balance between theoretical expressiveness and operational explainability, a balance that is particularly important in regulated financial contexts (Teimoory, 2025).

The second methodological component involves graph construction and representation learning. Retirement account ecosystems can be naturally represented as heterogeneous graphs, with nodes corresponding to users, accounts, devices, IP addresses, and transactions, and edges representing interactions, access events, or shared attributes. The methodology adopts a dynamic graph perspective, recognizing that both nodes and edges evolve over time as users interact with the system and as adversaries adapt their tactics (Rossi et al., 2020). Temporal graph networks and inductive representation learning techniques are therefore conceptually employed to model this evolution, enabling the system to generalize to previously unseen entities and behaviors (Xu et al., 2020).

Within this graph framework, attention mechanisms play a crucial role in selectively emphasizing relevant relationships. Graph attention networks and transformer-based architectures provide the theoretical basis for this selective weighting, allowing the model to focus on suspicious relational patterns while downplaying benign interactions (Velickovic et al., 2018; Vaswani et al., 2017). The methodology further incorporates insights from fraud-specific graph transformers, which have demonstrated efficiency and effectiveness in large-scale financial networks (Lin et al., 2024). These architectures are particularly well-suited to low-homophily environments, where fraudulent nodes may not be densely connected to one another, a condition often observed in retirement account fraud scenarios (Wang et al., 2023).

A critical methodological consideration is label scarcity. Fraud detection in retirement accounts suffers from limited and delayed labeling, as confirmed fraud cases are relatively rare and often identified long after initial compromise. To address this challenge, the framework draws on semi-supervised and barely supervised learning paradigms, which leverage structural and attribute information to infer risk signals in the absence of extensive labeled data (Xiang et al., 2023; Yu et al., 2024). Behavioral biometrics contributes to this process by providing continuous, unlabeled behavioral streams that can be clustered or scored for anomaly detection, thereby augmenting sparse fraud labels (Valiveti, 2025).

The integration of behavioral and graph-based components is achieved through a layered architecture. At the individual level, behavioral embeddings are generated for each user-session pair. These embeddings are then injected as node attributes within the broader graph, enriching relational representations with fine-grained behavioral context. This design reflects theoretical arguments that attribute-driven graph representations can significantly enhance fraud detection performance by bridging micro-level behavior and macro-level structure (Xiang et al., 2023). The methodology thus conceptualizes behavioral biometrics not as a standalone security mechanism but as an integral feature space within a relational learning system.

From a security architecture perspective, the methodology is situated within a zero trust model, wherein no access request is implicitly trusted and risk is continuously reassessed (Idialu, 2025). Behavioral biometric scores and graph-based risk assessments jointly inform access control decisions, step-up authentication triggers, and monitoring intensity. This continuous assessment aligns with contemporary critiques of perimeter-based security and supports adaptive defense against evolving threats (Teimoory, 2025). Importantly, the methodology emphasizes that such assessments must be explainable and auditable to satisfy regulatory and ethical requirements, an issue that remains a subject of active scholarly debate (Shwartz-Ziv and Armon, 2022).

Methodological limitations are explicitly acknowledged. The reliance on behavioral data raises privacy concerns and potential biases, particularly for users with atypical interaction patterns due to age, disability, or assistive technologies. Graph-based models may also inadvertently propagate biases present in historical data, leading to disparate impacts. The methodology therefore underscores the need for bias auditing, fairness constraints, and transparent governance mechanisms as integral components of any implementation (Valiveti, 2025). These considerations inform the interpretive lens applied in the subsequent results and discussion sections.

RESULTS

The results of this study are presented as a descriptive and interpretive synthesis of findings reported across the referenced literature, contextualized within the proposed integrated framework for retirement account security. Rather than enumerating quantitative metrics, the analysis focuses on patterns, trends, and conceptual outcomes that emerge when behavioral biometrics and graph-based learning are jointly considered. This approach reflects the methodological orientation of the study and aligns with prior interpretive analyses in fraud

detection research (Lin et al., 2024).

A central result emerging from the literature is the demonstrated efficacy of behavioral biometrics in identifying account compromise scenarios that evade traditional authentication controls. Studies focusing on behavioral analysis consistently report that subtle deviations in interaction patterns can signal unauthorized access even when valid credentials are used (Barkworth et al., 2022). In the specific context of retirement accounts, behavioral biometrics has been shown to provide a continuous layer of defense that complements static security measures, reducing reliance on disruptive authentication challenges (Valiveti, 2025). This finding supports the theoretical claim that behavior-based signals capture aspects of user identity that are difficult for adversaries to replicate at scale.

Another key result pertains to the role of graph-based models in uncovering coordinated and relational fraud patterns. Research on graph attention networks and temporal graph learning demonstrates that relational representations significantly enhance detection capabilities in complex financial networks (Velickovic et al., 2018; Rossi et al., 2020). When applied to fraud contexts, these models excel at identifying shared infrastructure, anomalous transaction flows, and temporal sequences indicative of malicious campaigns (Yu et al., 2024). The interpretive synthesis suggests that such capabilities are directly relevant to retirement account security, where fraud often involves coordinated access attempts across multiple accounts and devices.

The integration of behavioral features as node attributes within graph models emerges as a particularly salient result. Attribute-driven graph representation learning has been shown to improve performance in semi-supervised fraud detection settings by enriching structural information with contextual detail (Xiang et al., 2023). Within the proposed framework, behavioral embeddings serve this enriching function, enabling the model to differentiate between structurally similar but behaviorally distinct access events. This integration addresses challenges associated with low homophily, as fraudulent behavior may not manifest through dense relational clusters alone (Wang et al., 2023).

The literature also indicates that attention-based architectures, including transformers and graph transformers, contribute to more robust and adaptable fraud detection. By dynamically weighting relationships and temporal contexts, these models can prioritize salient signals and adjust to evolving threat patterns (Vaswani et al., 2017; Lin et al., 2024). In interpretive terms, this adaptability aligns with the dynamic nature of behavioral biometrics, which likewise evolves over time. The convergence of attention mechanisms across both behavioral and graph domains thus represents a coherent design principle supported by empirical findings.

Another important result concerns label efficiency. Barely supervised and semi-supervised approaches have demonstrated that meaningful fraud detection is possible even with limited labeled data, provided that models effectively leverage structural and attribute information (Yu et al., 2024). This result is particularly relevant for retirement accounts, where confirmed fraud cases are rare and delayed. Behavioral biometrics contributes continuous, unlabeled data streams that can be harnessed for anomaly detection and representation learning, thereby mitigating label scarcity (Valiveti, 2025).

Finally, the interpretive synthesis highlights trade-offs and limitations reported in the literature. While deep learning models offer expressive power, concerns regarding interpretability, computational cost, and bias persist (Shwartz-Ziv and Armon, 2022). Results from studies emphasizing simpler models suggest that performance gains must be weighed against operational complexity and regulatory scrutiny. In the retirement account context, these considerations underscore the need for balanced architectures that combine advanced learning techniques with transparent decision logic (Teimoory, 2025).

Collectively, these results support the central thesis that integrating AI-driven behavioral biometrics with graph-based fraud intelligence yields a more comprehensive and resilient security posture for retirement accounts. The findings also set the stage for a deeper theoretical discussion of implications, counter-arguments, and future research trajectories.

DISCUSSION

The discussion section undertakes an extensive theoretical interpretation of the synthesized results, situating them within broader scholarly debates on financial security, machine learning, and behavioral analysis. At its core, the discussion interrogates the conceptual significance of integrating behavioral biometrics and graph-based learning, exploring both the promise and the tensions inherent in such an approach. This interrogation is

grounded in the recognition that retirement account security occupies a unique position at the intersection of long-term financial planning, digital accessibility, and regulatory oversight (Valiveti, 2025).

One of the most salient theoretical implications concerns the reconceptualization of identity in digital financial systems. Traditional authentication frameworks treat identity as a static attribute verified at discrete moments, typically through credentials or tokens. Behavioral biometrics challenges this conception by framing identity as a dynamic process, continuously enacted through interaction patterns (Arora and Kanji, 2019). When embedded within a graph-based relational context, identity becomes not only dynamic but also relational, shaped by patterns of interaction with devices, networks, and other entities. This relational identity perspective aligns with contemporary theories of socio-technical systems, which emphasize the co-construction of user behavior and technological affordances (Rossi et al., 2020).

The integration of graph-based learning further amplifies this reconceptualization by highlighting the collective dimensions of fraud. Fraud is rarely an isolated act; it often involves coordinated strategies, shared resources, and temporal sequencing. Graph models make these dimensions explicit, allowing security systems to reason about collective risk rather than solely individual anomalies (Velickovic et al., 2018). In retirement account contexts, where individual transactions may be infrequent and low-volume, this collective reasoning is particularly valuable. Behavioral biometrics contributes granularity to this reasoning, ensuring that collective patterns do not obscure individual legitimacy (Valiveti, 2025).

Despite these theoretical strengths, the discussion must also engage with counter-arguments. Critics of deep learning-based fraud detection caution against overfitting, opacity, and the erosion of user trust due to perceived surveillance (Shwartz-Ziv and Armon, 2022). Behavioral biometrics, in particular, raises ethical concerns regarding consent, data minimization, and potential discrimination against users with atypical behaviors. These concerns are not merely peripheral but strike at the legitimacy of AI-driven security interventions. The literature suggests that addressing these concerns requires not only technical safeguards but also transparent communication and robust governance frameworks (Teimoory, 2025).

Another point of scholarly debate concerns the relative importance of model complexity versus domain knowledge. Some researchers argue that sophisticated architectures such as transformers and temporal graph networks are indispensable for capturing the nuances of modern fraud (Lin et al., 2024). Others contend that domain-informed feature engineering and simpler models may achieve comparable results with greater interpretability and lower cost (Shwartz-Ziv and Armon, 2022). The integrated framework discussed in this article implicitly adopts a middle position, leveraging advanced architectures where they offer clear advantages while retaining simpler components for baseline understanding and explanation. This position reflects a pragmatic synthesis rather than a definitive resolution of the debate.

The discussion also explores the implications of low homophily in retirement account fraud networks. Traditional graph-based fraud detection often assumes that fraudulent nodes cluster together, an assumption that does not always hold in practice (Wang et al., 2023). Retirement account fraud may involve isolated compromises dispersed across a large user base, connected only through subtle infrastructural or temporal links. Attention-based graph models and attribute-driven representations offer a means of addressing this challenge, but their effectiveness depends on careful design and validation (Xiang et al., 2023). Behavioral biometrics enhances this effectiveness by providing discriminative attributes that do not rely on network density alone (Valiveti, 2025).

From an architectural standpoint, the discussion situates the integrated framework within zero trust security paradigms. Zero trust emphasizes continuous verification, least privilege, and adaptive response, principles that align closely with behavioral and graph-based risk assessment (Idialu, 2025). However, implementing zero trust in retirement account systems raises practical challenges, including legacy infrastructure integration and user experience considerations. Excessive friction or false positives can undermine user confidence and participation, particularly among older account holders. The literature suggests that behavioral biometrics, by operating transparently in the background, can mitigate these challenges if designed responsibly (Barkworth et al., 2022).

The discussion further considers future research directions. One promising avenue involves the integration of blockchain-based audit trails with AI-driven detection, enhancing transparency and post-incident analysis (Idialu, 2025). Another involves the development of fairness-aware graph learning algorithms that explicitly address demographic and behavioral diversity. Longitudinal studies examining how behavioral signatures evolve over the

life course of retirement account holders could also enrich theoretical understanding and model robustness (Valiveti, 2025).

Limitations of the present study are acknowledged. As a conceptual synthesis, the article does not present new empirical validation, and its conclusions are contingent on the quality and scope of the referenced literature. Additionally, the rapid evolution of AI techniques means that specific architectural recommendations may require continual updating. Nonetheless, the theoretical integration offered here provides a durable framework for thinking about retirement account security in an AI-driven era.

CONCLUSION

This article has developed an extensive, publication-ready theoretical synthesis of AI-driven behavioral biometrics and graph-based fraud intelligence as applied to retirement account security. Grounded strictly in existing literature, the study has articulated a coherent framework that addresses both individual-level impersonation risks and network-level fraud dynamics. Central to this framework is the recognition that behavioral biometrics offers a continuous, low-friction means of assessing user legitimacy, while graph-based learning captures the relational and temporal complexity of modern fraud (Valiveti, 2025).

By elaborating the methodological logic, interpretive results, and theoretical implications of integrating these approaches, the article contributes to ongoing scholarly debates on digital identity, fraud detection, and financial security architecture. It highlights the necessity of multidisciplinary thinking and cautions against simplistic solutions that ignore ethical, regulatory, and user experience considerations. As retirement account systems continue to evolve, the insights presented here underscore the importance of adaptive, transparent, and behaviorally informed security strategies.

REFERENCES

1. Yu, H., Liu, Z., and Luo, X. Barely supervised learning for graph-based fraud detection. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 16548–16557, 2024.
2. Valiveti, S. S. S. AI-Driven Behavioral Biometrics for 401(k) Account Security. International Research Journal of Advanced Engineering and Technology, 2(06), 23–26, 2025.
3. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. In International Conference on Learning Representations, 2018.
4. Shwartz-Ziv, R., and Armon, A. Tabular data: Deep learning is not all you need. Information Fusion, 81, 84–90, 2022.
5. Barkworth, A., Tabassum, R., and Habibi Lashkari, A. Detecting IMAP credential stuffing bots using behavioural biometrics. ACM International Conference Proceeding Series, pages 7–15, 2022.
6. Lin, J., Guo, X., Zhu, Y., Mitchell, S., Altman, E., and Shun, J. FraudGT: A simple, effective, and efficient graph transformer for financial fraud detection. In Proceedings of the Fifth ACM International Conference on AI in Finance, pages 292–300, 2024.
7. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., and Bronstein, M. Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637, 2020.
8. Idialu, F. A. Leveraging zero trust architectures and blockchain protocols to prevent credential stuffing and lateral fraud attacks in enterprise systems. International Journal of Computer Applications Technology and Research, 14(8), 2025.
9. Xiang, S., Zhu, M., Cheng, D., Li, E., Zhao, R., Ouyang, Y., Chen, L., and Zheng, Y. Semi-supervised credit card fraud detection via attribute-driven graph representation. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 37, pages 14557–14565, 2023.
10. Teimoory, P. Towards robust security in smart payment systems: challenges and solutions. Smart Cities Regional Development Journal, 9(3), 29–38, 2025.
11. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. Advances in Neural Information Processing Systems, 30, 2017.

12. Wang, Y., Zhang, J., Huang, Z., Li, W., Feng, S., Ma, Z., Sun, Y., Yu, D., Dong, F., Jin, J., et al. Label information enhanced fraud detection against low homophily in graphs. In Proceedings of the ACM Web Conference, pages 406–416, 2023.
13. Xu, D., Ruan, C., Korpeoglu, E., Kumar, S., and Achan, K. Inductive representation learning on temporal graphs. arXiv preprint arXiv:2002.07962, 2020.
14. Arora, V., and Kanji, R. Modelling and predicting user behaviour. 2019.
15. Sadeghpour, S. Machine learning-based defences against advanced session-replay web bots. York University, 2024.