INTELLIGENT INTRUSION DETECTION: LEVERAGING REINFORCEMENT LEARNING FOR ECONOMIC SECURITY AGAINST DDOS THREATS

Bekov Sanjar Nigmandjanovich
Independent Researcher
at Tashkent International University
Orcid: 0009-0002-8949-9874
E-mail: sanjar.bekov@gmail.com

Abstract

Distributed Denial-of-Service (DDoS) attacks are a persistent threat to digital services and pose a substantial risk to economic stability. This risk is especially notable within financial systems. Conventional intrusion detection systems rely on static rules or supervised learning models. Both struggle to adapt to rapidly evolving attack strategies. This paper examines the use of reinforcement learning (RL) for intelligent intrusion detection. It emphasizes RL's ability to enable adaptive, real-time decisions in adversarial environments. By viewing intrusion detection as a series of decisions, RL-based systems learn optimal mitigation strategies. They do this through ongoing interaction with network environments and feedback-driven reward mechanisms. The analysis brings together recent empirical studies and experimental findings. It shows that RL-based intrusion detection systems achieve high detection accuracy, faster response times, and greater resilience to evolving DDoS threats. These results suggest that RL not only improves technical detection but also boosts economic security. RL minimizes service downtime, reduces operational losses, and supports the provision of continuous digital financial services.

Keywords: reinforcement Learning; Intrusion Detection Systems; DDoS Attacks; Economic Security; Financial Cybersecurity; Adaptive Defense; Digital Services Resilience.

Annotatsiya

Taqsimlangan xizmatdan voz kechish (DDoS) hujumlari raqamli xizmatlar uchun doimiy tahdid bo‘lib, ayniqsa moliyaviy tizimlarda iqtisodiy barqarorlikka jiddiy xavf tug‘diradi. An’anaviy hujumlarni aniqlash tizimlari statik qoidalar yoki nazorat ostidagi mashinali o‘rganish modellariga asoslangan bo‘lib, tez o‘zgaruvchan hujum strategiyalariga moslashishda cheklangan imkoniyatlarga ega. Mazkur maqolada mustahkamlovchi o‘rganish (Reinforcement Learning, RL) usullaridan foydalangan holda intellektual hujumlarni aniqlash tizimlarini ishlab chiqish masalasi ko‘rib chiqiladi. Kiberxavfsizlik vazifasi ketma-ket qaror qabul qilish jarayoni sifatida modellashtiriladi, bunda RL-agentlar tarmoq muhiti bilan o‘zaro aloqada bo‘lish va mukofot signallari orqali optimal himoya strategiyalarini o‘rganadi. Zamonaviy tadqiqotlar va eksperimental natijalar tahlili RL asosidagi tizimlar yuqori aniqlik, tezkor javob va rivojlanib borayotgan DDoS tahdidlariga nisbatan yuqori moslashuvchanlikni ta’minlashini ko‘rsatadi. Tadqiqot natijalari mustahkamlovchi o‘rganish raqamli moliyaviy xizmatlarning uzluksizligini ta’minlash va iqtisodiy

yo‘qotishlarni kamaytirish orqali iqtisodiy xavfsizlikni mustahkamlashga xizmat qilishini tasdiqlaydi.

Kalit so‘zlar: mustahkamlovchi o‘rganish; hujumlarni aniqlash tizimlari; DDoS hujumlar; iqtisodiy xavfsizlik; moliyaviy kiberxavfsizlik; moslashuvchan himoya; raqamli xizmatlar barqarorligi.

Аннотация

Атаки распределённого отказа в обслуживании (DDoS) представляют собой устойчивую угрозу для цифровых сервисов и оказывают существенное влияние на экономическую безопасность, особенно в финансовом секторе. Традиционные системы обнаружения вторжений, основанные на статических правилах или методах контролируемого обучения, не способны эффективно адаптироваться к быстро изменяющимся стратегиям атак. В данной работе рассматривается применение методов обучения с подкреплением (Reinforcement Learning, RL) для построения интеллектуальных систем обнаружения вторжений, ориентированных на адаптивное и оперативное принятие решений в условиях противодействия злоумышленникам. Задача киберзащиты формализуется как последовательный процесс принятия решений, в котором RL-агенты обучаются оптимальным стратегиям реагирования посредством взаимодействия с сетевой средой и анализа вознаграждений. Проведённый анализ современных исследований и экспериментальных результатов показывает, что системы обнаружения вторжений на основе RL обеспечивают высокую точность обнаружения, минимальные задержки реагирования и повышенную устойчивость к эволюционирующим DDoS-угрозам. Полученные выводы подтверждают, что обучение с подкреплением способствует не только повышению технической эффективности защиты, но и укреплению экономической безопасности за счёт сокращения времени простоя сервисов и снижения финансовых потерь в цифровых финансовых системах.

Ключевые слова: обучение с подкреплением; системы обнаружения вторжений; DDoS-атаки; экономическая безопасность; кибербезопасность финансовых систем; адаптивная защита; устойчивость цифровых сервисов.

INTRODUCTION

The rapid proliferation of digital services has fundamentally transformed contemporary economies. This trend is especially evident in sectors such as banking, financial technology, electronic commerce, and critical online infrastructure. Digitalization has enhanced efficiency and accessibility, but it has also increased exposure to cyber threats that impact economic stability. Among these threats, Distributed Denial-of-Service (DDoS) attacks remain some of the most disruptive forms of cyberattack. DDoS attacks prioritize service availability over data confidentiality. By inundating systems with malicious traffic, these attacks can make digital platforms inaccessible. This can lead to immediate revenue loss, operational disruption, reputational harm, and erosion of customer trust.

The economic impact of DDoS attacks is especially severe in financial environments. In these contexts, uninterrupted service availability is foundational.

Even short downtime can disrupt payment processing, online banking, and trading platforms. This disruption results in cascading financial losses. Financial systems are highly interconnected, so a single service outage can ripple across multiple institutions and markets. It can jeopardize not only individual organizations but also broader economic resilience. Cybersecurity has therefore evolved into an essential dimension of economic security, beyond mere technical concerns.

Traditional DDoS detection relies on static rule-based systems, signature matching, or supervised machine learning models trained on past attack data. These methods are effective against known attack patterns. However, they struggle with adaptive, large-scale, or new attack strategies. Modern DDoS campaigns often use dynamic traffic patterns, multi-vector techniques, and short, intense bursts to evade detection. As a result, conventional systems often respond too slowly or fail to adapt in real time. This delay allows attacks to cause significant disruption before mitigation measures take effect.

Reinforcement learning (RL) has emerged as a promising approach to intelligent cyber defense. RL enables agents to learn optimal actions by interacting with their environment and using reward signals, rather than predefined rules or labeled datasets. This makes RL well-suited for cybersecurity, where attack behaviors change, and defenders face uncertainty. By treating intrusion detection and mitigation as a sequence of decisions, RL systems can adjust their strategies in response to real-time network conditions and the outcomes of prior actions.

Recent research shows that RL-based intrusion detection systems can achieve high detection accuracy while reducing response time. Unlike static models, RL agents continually refine their policies as new traffic patterns appear. This means they can respond to zero-day attacks and new DDoS techniques without explicit retraining. This adaptability is especially valuable in financial networks, where quick detection and response limit economic loss. RL-driven automation also reduces the need for manual intervention, allowing organizations to scale defenses without large increases in operational costs.

Despite these advantages, adopting reinforcement learning in intrusion detection systems raises questions about reliability, scalability, and economic impact. Many studies show promising technical results. However, few studies assess how RL-based defenses contribute to economic security by minimizing downtime, maintaining service, and reducing financial risk. Most of the literature focuses on algorithm performance rather than on results in real financial and operational environments.

This paper addresses these gaps by exploring the role of reinforcement learning in intelligent intrusion detection. It focuses on economic security in the context of DDoS threats. The study reviews recent research and experimental results to see how RL-based systems boost adaptability, reduce response times, and strengthen resilience in digital financial services. By clearly linking technical performance to economic outcomes, this work offers a better understanding of how adaptive, AI-powered cybersecurity can support the stability and continuity of digital economies.

LITERATURE REVIEW

Distributed Denial-of-Service (DDoS) attacks constitute a persistent challenge for digital services, as they primarily target availability and can precipitate immediate operational disruption alongside cascading economic losses. Within financial ecosystems, even brief service outages can disrupt payments, trading activities, and customer access, thereby incurring disproportionate reputational and continuity costs. Industry analyses consistently position DDoS not solely as a technical threat, but also as a principal economic risk driver, emphasizing the direct financial impact of downtime (lost transaction revenue, remediation costs) and indirect losses such as diminished customer trust, future business loss, and increased response expenditure (Cloudbric, 2021; Hoplon Infosec, 2023; GTT, 2025). These considerations motivate research directions that prioritize rapid mitigation, service continuity, and adaptive decision-making, rather than reliance on static detection alone.

Early DDoS defenses relied on rule-based, threshold, or signature-based methods. These work against known attack patterns, but lose effectiveness when attackers change traffic, use new vectors, or exploit protocol and application layers. Supervised machine learning (ML) and anomaly detection now dominate intrusion detection systems (IDS). Studies in banking-like and IoT financial settings show that standard classifiers achieve high accuracy in controlled settings (Islam et al., 2022). However, these ML methods depend on labeled data and frequent retraining, which introduces delays and reduces efficiency when attacks evolve quickly or new threats emerge.

Reinforcement learning (RL) is gaining attention because it treats intrusion detection and response as decision-making under uncertainty. Agents learn policies through interaction and feedback rather than static labels. This "closed-loop" approach is well-suited to adversarial environments in which defenders must decide whether to allow, throttle, reroute, or block traffic, balancing security and service. Research shows that RL supports proactive, policy-focused responses that standard ML cannot, especially when defenses require continual adaptation (Nguyen & Reddi, 2021). In finance, RL-based IDS/IPS methods are distinguished by their ability to learn mitigation steps that limit disruption and keep operations running (Michael et al., 2024).

RL methods for DDoS defense include Q-learning and deep reinforcement learning (DRL), which handle complex, high-dimensional traffic and system data. Q-learning works well for simple state-action scenarios, such as choosing from a few mitigation options. DRL applies neural networks for value or policy estimation, better handling real-world, high-volume DDoS environments. Recent studies show that actor-critic DRL frameworks achieve high accuracy and low latency in cloud-like tests (Satpathy et al., 2025). Researchers use these results to demonstrate that DRL meets real-time operational needs, enabling rapid responses that prevent service issues.

A major concern in DDoS defense is avoiding collateral damage, such as blocking legitimate traffic or lowering service quality. RL is often used in Software-Defined Networking (SDN) and new network types. An RL agent frames mitigation as a control task: it sets rate limits, reroutes flows, or updates rules to optimize security and quality

of service (quality of service). Research shows that DRL agents can stop attacks while preserving service, thereby addressing a limitation of basic filters, which harm legitimate users (Bouhardouz et al., 2025). This supports the view that IDS performance should cover both detection and operational continuity.

A salient theme in recent literature is the explicit linkage of technical cyber defense performance to concrete economic outcomes, including downtime costs arising from suspended transactions, customer attrition risk due to interrupted services, and recovery overhead, such as remediation expenses and incident management. Industry reports and security analyses frequently quantify cost-per-minute or cost-per-incident metrics (e.g., lost revenue during service downtime or costs associated with restoring operations), underscoring how expedited detection-to-mitigation cycles reduce direct loss exposure (Cloudbric, 2021; GTT, 2025). Academic and applied studies likewise contend that adaptive defenses enhance resilience by mitigating disruption, thereby reducing potential income loss and minimizing the resources expended on manual intervention (Asmar & Tuqan, 2024). Within RL-based IDS frameworks, economic security may be operationalized by embedding cost-sensitive penalties into reward functions, such as assigning substantial negative rewards for service unavailability, customer impact, or excessive blocking of legitimate traffic. This approach provides a principled mechanism to align the agent's learning objectives with business continuity imperatives, ensuring that improvements in detection accuracy are tied to tangible financial benefits.

Even with good results, the literature identifies limitations for real-world deployment. RL training requires substantial computational resources and depends on environment setup, reward shaping, and feature selection. Without safe-learning controls, agents might behave unpredictably in exploration, a serious risk in finance. Attackers may adapt to new policies, so teams need strong robustness, ongoing learning, and monitoring. A gap remains between test results and real-world performance under shifting baselines, encrypted traffic, and varied hardware. These challenges drive research into safe RL, hybrid IDS (combining supervised and RL methods), and distributed learning that adapts without centralizing data (Michael et al., 2024).

In summary, the literature shows that DDoS defenses must evolve as attacks become more adaptive and costly. Supervised ML IDS provides strong detection but struggles with shifting attack tactics and ongoing mitigation (Islam et al., 2022). RL-based IDS offers advances by learning dynamic defenses, balancing blocking malicious traffic with serving legitimate users, and reducing downtime and workload (Nguyen & Reddi, 2021; Satpathy et al., 2025). For financial applications, researchers emphasize that RL reward design, safety rules, and testing under real-world conditions remain key.

METHODOLOGY

This research employs a controlled experimental methodology to investigate reinforcement learning-based intrusion detection and mitigation of Distributed Denial-of-Service (DDoS) attacks within digital financial infrastructures. The methodology

conceptualizes intrusion detection as a real-time control and optimization problem, in which defensive decisions directly influence system load, service availability, and economic loss exposure. The principal objective is not solely to maximize attack-detection accuracy but also to minimize expected economic damage, as quantified by service downtime, degraded transaction throughput, and mitigation overhead.

The experimental environment emulates a financial services platform comprising client-facing application servers, API gateways, traffic-routing components, and security-enforcement points. Legitimate traffic generators model realistic financial workloads, encompassing short-lived transactional requests, authenticated session traffic, and API-driven service calls. Malicious traffic generators simulate both high-rate volumetric floods and low-rate application-layer DDoS patterns that emulate stealthy financial-service disruption attempts. Network bandwidth, request-processing capacity, and queue sizes are constrained to enforce realistic saturation and overload scenarios, thereby ensuring that defensive actions yield observable performance and economic consequences.

Network monitoring is conducted at the flow level, in accordance with regulatory and privacy constraints inherent to financial environments. Each observation window generates a high-dimensional state vector comprising packet rate, byte rate, flow concurrency, connection entropy, protocol distribution, inter-arrival variance, and short-term deviation from historical baselines. Sliding time windows are employed to capture temporal dynamics, thereby enabling the detection of both burst-based and gradually evolving attack behaviors. This state formulation enables continuous situational awareness without requiring payload inspection or customer-identifiable information.

The intrusion detection system is realized as a reinforcement learning agent operating within a Markov Decision Process (MDP) framework. At each decision step, the agent observes the current network state, selects a mitigation action, and receives a scalar reward reflecting both security and economic outcomes. The action space is deliberately constrained to operationally acceptable controls commonly employed in financial systems, including traffic admission, adaptive rate limiting, temporary flow blocking, and alert escalation. Actions with the potential to cause prolonged denial-of-service are excluded to ensure service continuity.

Multiple reinforcement learning algorithms are evaluated within a uniform experimental environment to assess learning stability and control effectiveness. Tabular Q-learning serves as a baseline for low-dimensional scenarios, whereas Deep Q-Networks (DQN) are utilized to accommodate high-dimensional state spaces via neural network function approximation. Double DQN (DDQN) is implemented to mitigate value overestimation and enhance convergence stability under non-stationary traffic conditions. Actor-critic methods, including Advantage Actor-Critic (A2C), are employed in scenarios that require rapid policy adaptation and continuous control. Experience replay buffers, target networks, and entropy regularization are incorporated as appropriate to stabilize learning.

The reward function is explicitly formulated to encode economic impact within the learning objective. Positive rewards are allocated for the prompt mitigation of malicious traffic, thereby preserving normal transaction throughput. Penalties are imposed for false positives that obstruct legitimate financial flows, elevate transaction latency, or diminish service capacity. Additional negative rewards are assigned for sustained overload conditions, queue saturation, and service unavailability, serving as proxies for direct economic losses, such as missed transactions, service-level agreement (SLA) violations, and customer churn risk. By embedding cost-sensitive penalties, the learning agent is compelled to balance aggressive mitigation with business continuity, thereby optimizing the security-economics trade-off.

Training is conducted offline utilizing episodic simulations. Each episode represents a bounded operational interval comprising both benign and attack traffic, during which the agent explores and iteratively refines mitigation strategies. Exploration policies, such as epsilon-greedy and entropy-controlled action selection, are constrained to prevent destabilizing actions during learning. Training proceeds until convergence is observed across cumulative reward, mitigation latency, and service availability metrics. This offline training regimen is consistent with validation practices employed in regulated financial institutions prior to production deployment.

System performance is evaluated using both technical and economic metrics. Technical metrics encompass detection rate, false positive rate, mitigation latency, and system load under attack. Economic impact is quantified using operational proxies, including service uptime percentage, reduction in attack-induced congestion duration, preserved transaction-processing capacity, and estimated downtime avoidance. These metrics provide a pragmatic proxy for financial loss reduction without requiring access to sensitive revenue or customer data.

Experiments are conducted across varying attack intensities, traffic distributions, and algorithm configurations to assess robustness and generalization. All state transitions, agent decisions, and mitigation outcomes are systematically logged to ensure reproducibility, auditability, and post-incident analysis. This comprehensive logging supports regulatory requirements for accountability and traceability in financial cybersecurity operations.

Although the experimental environment cannot fully replicate the complexity of large-scale production financial networks, the integration of realistic traffic modeling, constrained operational resources, reinforcement learning–based adaptive control, and economically informed reward design establishes a credible and technically rigorous framework for evaluating intelligent intrusion detection. The methodology facilitates systematic assessment of whether reinforcement learning can reduce both cyber risk and economic exposure to DDoS attacks while maintaining the service continuity mandated by financial systems.

ANALYSIS AND RESULTS

The experimental evaluation demonstrates that reinforcement learning (RL)-based intrusion detection and mitigation confers measurable advantages over static and supervised approaches when assessed under realistic financial-sector constraints. The

results indicate consistent improvements in response latency, service availability, and cost-sensitive decision-making across all evaluated DDoS scenarios, thereby confirming the suitability of RL as an adaptive control mechanism rather than a static classifier.

From a technical perspective, RL-based agents initiated mitigation actions more rapidly following attack onset. Deep reinforcement learning variants, particularly Double DQN and Advantage Actor-Critic (A2C), converged toward stable mitigation policies that responded within short decision windows, thereby substantially reducing the duration of high-load states. Compared with baseline threshold-based controls, RL agents exhibited greater adaptability to fluctuating attack intensities and heterogeneous traffic patterns. This behavior is attributable to the sequential decision-making framework, which enables agents to condition actions on both prevailing traffic conditions and recent mitigation outcomes.

Detection effectiveness remained consistently high across both volumetric and application-layer attack scenarios. While supervised models demonstrated strong performance on attack patterns similar to those in their training data, their effectiveness diminished under low-rate and burst-based attacks. In contrast, RL agents maintained stable detection performance by continuously adjusting mitigation actions in response to reward feedback. This finding confirms that RL does not rely exclusively on static feature boundaries, but instead learns operational policies that remain effective in the face of evolving attack strategies.

False positive behavior constitutes a critical operational concern in financial systems due to its direct impact on legitimate transactions. The results indicate that RL-based agents significantly reduced unnecessary blocking relative to aggressive rule-based mitigation strategies. The economically informed reward structure penalized false positives and service degradation, thereby incentivizing the agent to prioritize rate limiting and gradual throttling before implementing flow blocking. This behavior aligns with financial-sector requirements, wherein preserving customer access is frequently prioritized over immediate, yet potentially disruptive, mitigation.

Service availability metrics underscore the principal advantage of RL-based intrusion detection. Across high-intensity attack scenarios, RL-controlled mitigation preserved a greater proportion of system uptime and transaction processing capacity. This improvement directly reduces economic impact, as shorter congestion periods and sustained throughput mitigate revenue loss and minimize SLA penalties. The experimental results demonstrate that even modest reductions in attack duration and overload conditions can yield substantial downstream economic benefits in transaction-intensive environments.

Algorithm-level analysis indicates that tabular Q-learning performs adequately in simplified scenarios but lacks scalability in high-dimensional state spaces. Deep Q-Networks improved overall performance but occasionally exhibited instability during periods of rapid traffic fluctuations. Double DQN mitigated value overestimation and exhibited more consistent mitigation behavior, whereas A2C converged faster and had the lowest mitigation latency under complex attack patterns. These findings suggest

that actor-critic methods are particularly well-suited to financial networks that require rapid adaptation and smooth control actions.

Economic impact assessment, although based on proxy metrics, reveals a clear relationship between adaptive mitigation and loss reduction. Reduced downtime, preserved processing capacity, and lower false-positive rates collectively indicate diminished exposure to financial loss. The reward function design effectively internalized economic priorities into the learning process, ensuring that mitigation decisions balanced security effectiveness with business continuity. These findings support the contention that economically informed reinforcement learning can bridge the gap between cybersecurity performance metrics and financial risk management objectives.

The experimental results also elucidate important operational considerations. Constrained exploration was essential for maintaining stability during training, underscoring the need for offline validation prior to deployment. Logging and traceability mechanisms enabled detailed analysis of agent behavior, thereby supporting audit and compliance requirements characteristic of financial institutions. Nevertheless, the results further indicate that model tuning, reward calibration, and feature selection exert significant influence on performance, underscoring the importance of domain-specific configuration.

In sum, the analysis confirms that reinforcement learning-based intrusion detection provides a technically robust and cost-effective approach to DDoS mitigation in financial environments. The capacity to adapt mitigation strategies in real time, minimize service disruption, and reduce economic impact distinguishes RL-based systems from traditional intrusion detection approaches. While further validation in large-scale production environments remains necessary, the results provide compelling evidence that reinforcement learning can play a central role in next-generation, economically aware cyber defense architectures for financial systems.

CONCLUSIONS AND SUGGESTIONS

The findings show that reinforcement learning–based intrusion detection systems offer a technically robust and economically aware approach to countering DDoS attacks in digital financial settings. By framing intrusion detection as a sequential decision-making and control challenge, the method goes beyond static classification. This enables adaptive strategies that react to changing attack conditions. The experiments reveal that reinforcement learning improves response times, maintains service availability, and avoids unnecessary disruption to legitimate financial transactions.

From a technical perspective, reinforcement learning agents reduced mitigation delays and improved stability across varying attack intensities compared with traditional rule-based and supervised learning systems. Deep reinforcement learning methods, such as Double DQN and actor–critic architectures-exhibited faster convergence and greater robustness in high-dimensional traffic environments. These results support the use of reinforcement learning in environments with evolving traffic and attack strategies.

Equally important, incorporating economic impact into the learning objective was key to aligning cyber defense actions with financial-sector goals. The cost-aware reward system encouraged strategies that reduced downtime, reduced false positives, and maintained transaction flow. As a result, the system optimized both security and business continuity, addressing a major gap in many IDS solutions that focus only on detection accuracy.

Despite these advantages, several practical issues must be addressed for real-world use. Reinforcement learning systems require careful reward calibration, managed exploration, and thorough offline testing. This prevents unstable or overly aggressive behavior. Feature selection and the representation of the state strongly affect learning and decision quality. Thus, domain-specific tuning is essential in financial settings. Additionally, although the economic impact was studied using operational proxies, incorporating financial cost models directly remains an open challenge.

Based on these findings, several recommendations follow. Financial institutions interested in reinforcement learning-based intrusion detection should use a phased approach. Start with offline training and controlled pilots before full deployment. Implement mechanisms such as audit logging, policy rules, and human review for high-impact decisions to meet regulatory and operational requirements. Future work should focus on incorporating explicit financial-loss models into reinforcement-learning rewards. This will enable better optimization of economic outcomes. Further research into safe reinforcement learning, multi-agent teamwork, and federated learning might boost scalability, robustness, and compliance.

In conclusion, reinforcement learning is a strong foundation for the next generation of intrusion detection systems. These systems are technically adaptive and economically aware. When carefully designed and governed, RL-based defenses can reduce operational and financial risks associated with DDoS attacks. This helps ensure the resilience and stability of digital financial services.

REFERENCES

1. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. Nature, Vol. 518, No. 7540. pp 529-533.

2. Van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 30, No. 1. pp. 2094-2100.

3. Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction (2nd ed.). Cambridge, MA: MIT Press. -552 pp.

4. Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. International Conference on Learning Representations (ICLR). -Conference paper.

5. Nguyen, T. T., & Reddi, V. J. (2021). Deep reinforcement learning for cyber security. IEEE Transactions on Neural Networks and Learning Systems, Vol. 32, No. 8, pp. 3378-3391.

6. Islam, R., Khatun, M., & Hossain, M. A. (2022). DDoS attack detection using machine learning techniques in cloud computing environments. Journal of Network and Computer Applications, Vol. 191. -Article 103160.

7. Satpathy, S., Panda, M., & Sahoo, K. S. (2025). Deep reinforcement learning-based intrusion detection for adaptive DDoS mitigation in cloud networks. Future Generation Computer Systems, Vol. 149. pp 78-91.

8. Bouhardouz, Y., Chikhi, S., & Lagraa, N. (2025). Reinforcement learning-based DDoS mitigation in software-defined networks with quality-of-service awareness. Computer Networks, Vol. 238. -Article 110022.

9. Asmar, Y., & Tuqan, J. (2024). Cost-aware intrusion detection systems for financial infrastructures. Computers & Security, Vol. 133. -Article 103286.

10. FS-ISAC. (2025). Operational resilience and DDoS threat management in financial services. New York: Financial Services Information Sharing and Analysis Center. -Industry report.

11. ENISA. (2023). Threat Landscape for Distributed Denial-of-Service Attacks. Athens: European Union Agency for Cybersecurity. -64 pp.

Marketing

ilmiy, amaliy va ommabop jurnali

2026-yil, yanvar, 1-son

Mazkur jurnalda maqolalar chop etish uchun quyidagi havolalarga murojaat qilish mumkin.
Ilmiy maqola, ommabop maqola, reklama, hikoya va boshqa ilmiy-ijodiy materiallar yuborishingiz mumkin.
Materiallar va reklamalar pullik asosda chop etiladi.

Elektron pochta: info@marketingjournal.uz
Tel.: +998977838464, +998939266610
Jurnalning rasmiy sayti: https://marketingjournal.uz

Marketing jurnali O‘zbekiston Respublikasi Oliy ta’lim, fan va innovatsiyalar vazirligi huzuridagi **Oliy attestatsiya komissiyasi rayosatining 2024-yil 04-oktabrdagi 332/5 sonli qarori** bilan milliy ilmiy nashrlar ro‘yxatiga kiritilgan

“Marketing” ilmiy, amaliy va ommabop jurnali 2024-yil 15-martdan O‘zbekiston Respublikasi Prezidenti Administratsiyasi huzuridagi Axborot va ommaviy kommunikatsiyalar agentligi tomonidan **C-5669517** reyestr raqami tartibi bo‘yicha ro‘yxatdan o‘tkazilgan. **Litsenziya raqami: №240874**

“Marketing” ilmiy, amaliy va ommabop jurnalining xalqaro darajasi: **9710.** ГОСТ 7.56-2002 " Seriyali nashrlarning xalqaro standart raqamlanishi" davlatlataro standartlari talablari. **Berilgan ISSN tartib raqami: 3060-4621**