Baltic Journal of MultiDisciplinary Research - BJMR

Pages: 36-42

Volume-II, Issue-III (2025)

Transferable AI-Causal Models for Proactive Threat Detection Across Heterogeneous Cloud Infrastructures

Author: [1]Arooj Basharat, [2]Anas Raheem

Corresponding Author: aroojbasharat462@gmail.com

Abstract

Cloud infrastructures have rapidly evolved into highly distributed, heterogeneous environments that integrate diverse computing resources, network configurations, and service delivery models. This complexity creates both opportunities and vulnerabilities, especially in the context of cybersecurity. Traditional machine learning methods for threat detection often struggle with domain transfer, adaptability, and explainability, particularly when applied across heterogeneous cloud infrastructures. This paper explores the development and application of transferable AI-causal models that unify causal inference with advanced transfer learning techniques to enable proactive threat detection. By leveraging causal reasoning, these models go beyond correlation-based anomaly detection to uncover causal structures that explain the origins and propagation of cyber threats. Transferability ensures that models trained in one environment can be adapted efficiently to different cloud ecosystems, reducing retraining costs and improving robustness against zero-day attacks. We discuss the architectural design, learning paradigms, and operational deployment strategies for these models, along with their advantages in scalability, transparency, and resilience. Finally, we highlight open challenges and research directions, including causal discovery in dynamic environments, cross-domain explainability, and integration with real-time decision systems.

[1]University of Punjab, Pakistan

[2]Air University, Pakistan

Keywords: Transferable AI, causal models, proactive threat detection, heterogeneous cloud infrastructures, cybersecurity, transfer learning, anomaly detection, cloud security, causal inference, explainable AI

I. Introduction

The adoption of cloud infrastructures has become foundational to digital transformation across industries, enabling agility, scalability, and cost efficiency. However, the heterogeneity of cloud environments—spanning public, private, hybrid, and edge deployments—has introduced significant complexity in security management. Organizations face an ever-expanding attack surface, as cloud platforms integrate a wide variety of virtualized resources, containerized applications, and distributed data pipelines. With the proliferation of sophisticated cyber threats, from ransomware to advanced persistent threats (APTs), conventional security strategies that rely on static rule-based systems or purely statistical anomaly detection fall short in addressing the dynamic and adaptive nature of modern attacks.

Machine learning (ML) has emerged as a promising tool for enhancing cloud security by enabling automated anomaly detection and predictive defense mechanisms. Yet, current ML approaches often encounter limitations when applied across heterogeneous infrastructures. Models trained in one environment may fail to generalize effectively in another due to differences in configurations, workloads, and network topologies. This phenomenon, known as domain shift, severely reduces the efficacy of purely data-driven solutions. Additionally, many ML models function as black boxes, producing predictions without offering causal explanations for why a particular threat is detected, thus limiting trust and interpretability in high-stakes security operations[1].

To overcome these challenges, the integration of causal modeling with transferable AI frameworks offers a transformative solution. Causal models go beyond identifying statistical correlations, aiming instead to uncover the cause-and-effect relationships underlying observed behaviors. By capturing the causal mechanisms of attacks—such as how a vulnerability leads to privilege escalation or how lateral movement propagates across systems—these models provide richer insights into the dynamics of cyber threats. When coupled with transfer learning

techniques, causal models can be adapted from one cloud environment to another, ensuring reusability and reducing the need for exhaustive retraining.

This paper investigates the concept of transferable AI-causal models as a paradigm shift in proactive threat detection for heterogeneous cloud infrastructures. We propose that by embedding causal reasoning within AI architectures, organizations can achieve improved threat attribution, greater interpretability, and heightened resilience against novel attacks. Furthermore, the ability to transfer knowledge across environments addresses the scalability challenge, allowing organizations to deploy consistent security frameworks across multi-cloud ecosystems[2].

The structure of this paper is as follows. First, we discuss the foundational principles and architectures of transferable AI-causal models, highlighting how causal inference and transfer learning synergize in threat detection. Next, we examine their application in heterogeneous cloud infrastructures, focusing on operational benefits, implementation strategies, and challenges. Finally, we conclude by emphasizing the significance of this approach in building proactive, adaptive, and explainable security systems for the future of cloud computing[3].

II. Transferable AI-Causal Models: Principles and Architectures

The foundation of transferable AI-causal models lies in the integration of causal inference with transfer learning to create adaptive, interpretable, and reusable security mechanisms. Unlike conventional deep learning systems that rely solely on correlations, causal models aim to understand the underlying mechanisms that generate data. For example, when monitoring system logs, a correlation-based model may identify that certain access requests co-occur with anomalous traffic, but it cannot explain whether these requests are the cause or the consequence of malicious activity. In contrast, a causal model leverages domain knowledge, graph-based causal structures, or causal discovery algorithms to model the sequence of events and infer the root causes of anomalies[4].

The transferable dimension ensures that these models are not bound to a single domain. Through transfer learning, knowledge acquired in one cloud environment can be fine-tuned or adapted to

another with minimal retraining. For example, a causal model trained to detect insider threats in an AWS-based infrastructure can be repurposed for a Microsoft Azure or hybrid cloud setup by transferring causal dependencies and updating domain-specific parameters. This adaptability addresses the challenges of domain shift and reduces the time, computational cost, and labeled data required for deployment in new contexts[5].

Architecturally, transferable AI-causal models combine three layers: causal representation learning, transferable inference mechanisms, and adaptive threat detection modules. Causal representation learning encodes input data—such as system logs, traffic flows, or authentication events—into representations that capture underlying causal relations rather than surface-level correlations. Transferable inference mechanisms then apply domain adaptation techniques, including adversarial training, multi-domain regularization, and meta-learning, to bridge differences between source and target cloud environments. Finally, adaptive threat detection modules deploy these causal insights for real-time anomaly detection, proactive defense, and automated incident response[6].

Another critical advantage of this architecture is explainability. Because causal models inherently identify cause-and-effect relationships, they can generate interpretable alerts that specify not only that a threat exists but also how it emerged and which pathways it exploited. This interpretability builds trust among security analysts and facilitates compliance with regulations that demand transparency in automated decision-making. Furthermore, causal reasoning can enhance resilience against adversarial attacks, as attackers often exploit correlation-based blind spots, while causal dependencies remain more robust to manipulation[7].

The potential of transferable AI-causal models extends beyond detection to proactive security. By simulating counterfactual scenarios—such as assessing what would happen if a firewall rule were changed or if a suspicious connection were blocked—causal models can support proactive mitigation strategies. This predictive capability enables organizations to act before an attack fully materializes, shifting cybersecurity from reactive incident handling to proactive defense[8].

III. Applications in Heterogeneous Cloud Infrastructures

Heterogeneous cloud infrastructures present unique challenges for security due to their diverse configurations, multi-vendor ecosystems, and distributed nature. Public clouds such as AWS, Azure, and Google Cloud often operate alongside private data centers, edge nodes, and hybrid deployments, resulting in varied resource orchestration, network topologies, and compliance requirements. These differences hinder the deployment of uniform security mechanisms, as a model optimized for one environment may fail to capture the nuances of another[9].

Transferable AI-causal models address this complexity by enabling cross-domain adaptability. For instance, a causal model trained to detect distributed denial-of-service (DDoS) attacks in one cloud environment can be transferred to another with modifications that account for new traffic baselines or different load-balancing mechanisms. Similarly, insider threat detection models can adapt causal structures related to user access patterns across varied identity management systems. This adaptability ensures that organizations maintain consistent and effective threat detection strategies across heterogeneous ecosystems without rebuilding models from scratch[10].

A practical application lies in multi-cloud orchestration security. Organizations increasingly adopt multi-cloud strategies to avoid vendor lock-in and enhance resilience. However, this approach exposes them to inconsistencies in monitoring and incident response. Transferable causal models can unify security analytics across platforms by learning causal dependencies that transcend platform-specific idiosyncrasies. For example, causal chains involving privilege escalation, lateral movement, and data exfiltration remain consistent across environments, even if the technical implementations differ. Thus, transferable models provide a universal layer of defense that adapts to localized contexts[11].

Furthermore, these models facilitate real-time threat prediction by simulating potential attack trajectories in heterogeneous environments. Through counterfactual reasoning, causal models can test hypothetical changes—such as varying access policies or patching strategies—across different platforms, helping security teams anticipate attack vectors before they materialize. This proactive stance aligns with the growing need for resilience in cloud-native architectures, where downtime or breaches have immediate operational and financial consequences[12].

The explainability of causal models also addresses one of the key challenges in heterogeneous infrastructures: compliance and trust. Regulatory frameworks such as GDPR, HIPAA, and industry-specific standards demand transparent and auditable security mechanisms. Causal explanations provide evidence-based narratives that justify security decisions, easing compliance and building confidence among stakeholders. For global organizations operating in multiple jurisdictions, this interpretability is invaluable[13].

Despite these advantages, implementing transferable AI-causal models in heterogeneous cloud infrastructures faces practical challenges. These include the computational overhead of causal discovery in dynamic systems, the scarcity of labeled threat data across environments, and the difficulty of aligning causal graphs with constantly evolving infrastructure configurations[14]. Hybrid approaches that combine causal priors with data-driven learning are emerging as promising solutions to balance scalability with accuracy. Additionally, integrating causal models with existing Security Information and Event Management (SIEM) systems and cloud-native monitoring tools is essential for operational deployment[15].

IV. Conclusion

Transferable AI-causal models represent a paradigm shift in cybersecurity for heterogeneous cloud infrastructures. By uniting causal inference with transfer learning, these models provide adaptability, explainability, and proactive defense capabilities that traditional machine learning approaches lack. Their ability to uncover cause-and-effect relationships ensures not only accurate detection but also meaningful interpretation of threats, while transferability guarantees scalability across diverse cloud ecosystems. Although challenges remain in terms of computational cost, data availability, and real-time integration, the promise of these models lies in their potential to transform cloud security from reactive incident response to proactive, resilient defense. As cloud environments continue to evolve, transferable AI-causal models will be central to ensuring secure, adaptive, and trustworthy digital ecosystems.

References:

[1] Y. Y. Yu, S. Q. Qin, and Q. Y. Wen, "Data Integrity and Availability in Cloud Computing Based on Megastore," *Applied Mechanics and Materials,* vol. 411, pp. 1062-1066, 2013.

[2] M. Y. A. Younis and K. Kifayat, "Secure cloud computing for critical infrastructure: A survey," *Liverpool John Moores University, United Kingdom, Tech. Rep,* 2013.

[3] J. Barach, "Towards Zero Trust Security in SDN: A Multi-Layered Defense Strategy," in *Proceedings of the 26th International Conference on Distributed Computing and Networking*, 2025, pp. 331-339.

[4] K. Yang and X. Jia, "An efficient and secure dynamic auditing protocol for data storage in cloud computing," *Parallel and Distributed Systems, IEEE Transactions on,* vol. 24, no. 9, pp. 1717-1726, 2013.

[5] J. Barach, "Enhancing intrusion detection with CNN attention using NSL-KDD dataset. In 2024 Artificial Intelligence for Business (AIxB)(pp. 15-20)," ed: IEEE, 2024.

[6] B. Wickremasinghe, R. N. Calheiros, and R. Buyya, "Cloudanalyst: A cloudsim-based visual modeller for analysing cloud computing environments and applications," in *Advanced Information Networking and Applications (AINA), 2010 24th IEEE International Conference on*, 2010: IEEE, pp. 446-452.

[7] J. Barach, "Cross-Domain Adversarial Attacks and Robust Defense Mechanisms for Multimodal Neural Networks," in *International Conference on Advanced Network Technologies and Intelligent Computing*, 2024: Springer, pp. 345-362.

[8] C. Wang, Q. Wang, K. Ren, and W. Lou, "Privacy-preserving public auditing for data storage security in cloud computing," in *INFOCOM, 2010 Proceedings IEEE*, 2010: IEEE, pp. 1-9.

[9] J. Barach, "AI-Driven Causal Inference for Cross-Cloud Threat Detection Using Anonymized CloudTrail Logs," in *2025 Conference on Artificial Intelligence x Multimedia (AIxMM)*, 2025: IEEE, pp. 45-50.

[10] H. Takabi, J. B. Joshi, and G.-J. Ahn, "Security and privacy challenges in cloud computing environments," *Security & Privacy, IEEE,* vol. 8, no. 6, pp. 24-31, 2010.

[11] J. Barach, "Federated Learning for Privacy-Preserving Employee Performance Analytics," *IEEE Access,* 2025.

[12] S. Subashini and V. Kavitha, "A survey on security issues in service delivery models of cloud computing," *Journal of Network and Computer Applications,* vol. 34, no. 1, pp. 1-11, 2011.

[13] J. Barach, "Cybersecurity Project Management Failures," *Indexed in,* vol. 38, 2024.

[14] J. Barach, "Integrating AI and HR Strategies in IT Engineering Projects: A Blueprint for Agile Success," *Emerging Engineering and Mathematics,* pp. 1-13, 2025.

[15] P. Kumar and R. Kumar, "Issues and challenges of load balancing techniques in cloud computing: A survey," *ACM Computing Surveys (CSUR),* vol. 51, no. 6, pp. 1-35, 2019.