

Attention-Enhanced Cross-Modal Learning for Detecting Anomalies in System Software

Rachel Stein *

Washington University in St. Louis, USA

* Corresponding Author: rachelstein455@gmail.com

Abstract:

System software anomaly detection has become increasingly critical as modern computing systems grow in complexity and scale. Traditional approaches typically analyze single data modalities in isolation, failing to capture the rich interplay between different system components that often signals anomalous behavior. This paper presents a novel attention-enhanced cross-modal learning framework specifically designed for detecting anomalies in system software environments. Our approach integrates multiple data modalities including system logs, performance metrics, network traffic patterns, and process execution traces through sophisticated convolutional neural network architectures and probabilistic graphical models. The framework employs deep learning-based feature extraction with specialized attention mechanisms that learn to focus on the most discriminative cross-modal relationships, combined with mixture-of-experts approaches that dynamically select the most reliable detection mechanisms based on system context. We introduce adaptive fusion techniques that leverage hierarchical Dirichlet process hidden Markov models for capturing temporal dependencies and switching dynamics in system behavior. Comprehensive evaluation on real-world datasets from enterprise environments demonstrates that our cross-modal approach achieves superior detection performance compared to single-modal baselines, with improvements of 23% in precision and 18% in recall for detecting zero-day attacks and system vulnerabilities. The probabilistic framework provides interpretable explanations for detected anomalies through confidence-based fusion mechanisms that enable security analysts to quickly understand the root causes and take appropriate remediation actions.

Keywords:

Anomaly Detection, Cross-Modal Learning, Convolutional Neural Networks, Hidden Markov Models, System Security, Mixture of Experts

1. Introduction

Modern computing systems generate vast amounts of heterogeneous data from multiple sources including system logs, performance metrics, network communications, and process execution traces[1]. This multi-modal data landscape presents both opportunities and challenges for anomaly detection systems[2]. While the diversity of available data sources provides richer information for identifying malicious activities and system failures, the complexity of analyzing and correlating information across different modalities poses significant technical challenges that traditional single-modal approaches cannot adequately address.

System software anomalies manifest through complex patterns that often span multiple data modalities simultaneously[3]. A sophisticated cyber attack, for example, may involve unusual network traffic patterns, suspicious process executions, abnormal file system access patterns, and characteristic changes in system performance metrics[4]. Traditional anomaly detection systems that analyze each data source independently may miss these multi-faceted attack signatures, leading to decreased detection accuracy and increased false positive rates.

The emergence of deep learning architectures, particularly convolutional neural networks and probabilistic graphical models, has revolutionized how machine learning models process and correlate information across different data types and temporal scales[5]. Convolutional neural networks enable models to extract hierarchical features from raw system data, while hidden Markov models with infinite state spaces provide principled approaches to modeling temporal dependencies and regime changes in system behavior[6].

However, the direct application of standard deep learning architectures to system anomaly detection faces several challenges[7]. First, system data modalities exhibit vastly different statistical properties, temporal characteristics, and semantic meanings that require specialized handling[8]. Second, the imbalanced nature of anomaly detection datasets, where normal behaviors vastly outnumber anomalous events, requires careful model design to avoid bias toward majority classes. Third, the real-time requirements of operational security systems demand efficient architectures that can process high-volume data streams while maintaining detection accuracy[9].

This research addresses these challenges by developing a comprehensive attention-enhanced cross-modal learning framework specifically designed for system software anomaly detection. Our approach recognizes that effective anomaly detection requires understanding both the individual characteristics of different data modalities and the complex relationships between them. The framework integrates domain knowledge about system behavior with the representational power of convolutional neural networks and the temporal modeling capabilities of probabilistic graphical models.

The primary contributions of this work include the development of specialized convolutional architectures for handling heterogeneous system data modalities, probabilistic models that capture temporal dependencies and switching dynamics in system behavior, mixture-of-experts frameworks that dynamically select optimal detection strategies based on system context, and comprehensive evaluation demonstrating superior performance compared to existing approaches. Our framework provides system administrators and security analysts with powerful tools for protecting complex software systems while maintaining the operational efficiency necessary for production environments.

2. Literature Review

The field of system anomaly detection has evolved significantly over the past decade, driven by increasing sophistication of cyber threats and the growing complexity of modern computing environments[10]. Early approaches to anomaly detection focused primarily on single-modal

analysis, examining individual data sources such as network traffic, system logs, or performance metrics in isolation. While these methods achieved reasonable performance for well-defined attack patterns, they struggled with sophisticated threats that manifest across multiple system components[11].

Traditional statistical approaches to anomaly detection, including clustering-based methods, principal component analysis, and support vector machines, provided foundational techniques for identifying outliers in system data[12]. These methods established important principles for anomaly detection but were limited by their inability to capture complex temporal patterns and cross-modal relationships that characterize modern system behaviors.

The introduction of deep learning techniques marked a significant advancement in anomaly detection capabilities[13]. Convolutional neural networks demonstrated exceptional ability to learn hierarchical feature representations from raw system data, enabling automatic feature extraction without manual engineering. Early CNN-based approaches showed promise for analyzing structured system data such as network packet sequences and system call traces, but most implementations remained focused on single-modal analysis[14].

Autoencoder-based approaches demonstrated the ability to learn complex representations of normal system behavior and identify deviations through reconstruction error analysis[15]. These methods provided unsupervised learning capabilities that were particularly valuable for anomaly detection scenarios where labeled data was scarce. However, autoencoder approaches typically struggled with temporal dependencies and cross-modal relationships[16-20].

The development of probabilistic graphical models provided principled approaches to modeling complex dependencies in system data[21]. Hidden Markov Models offered frameworks for capturing temporal patterns and state transitions in system behavior, while hierarchical extensions enabled modeling of multi-scale temporal dependencies. The introduction of non-parametric Bayesian methods, particularly hierarchical Dirichlet processes, addressed the challenge of unknown state space sizes in system modeling[22-26].

Recent advances in mixture-of-experts architectures have shown promise for combining multiple detection mechanisms within unified frameworks. These approaches recognize that different types of anomalies may be best detected by different algorithmic approaches, and provide principled methods for combining diverse detection strategies[27]. The mixture-of-experts paradigm enables dynamic selection of optimal detection mechanisms based on current system context and data characteristics[28].

The challenge of interpretability in anomaly detection has received increasing attention as security systems are deployed in critical infrastructure and enterprise environments[29]. Traditional machine learning approaches often provide limited insights into why specific events are classified as anomalous, making it difficult for security analysts to understand and respond to detected threats. Probabilistic approaches offer promising solutions by providing confidence estimates and uncertainty quantification for detection decisions[30].

Contemporary research has begun to explore the integration of deep learning architectures with probabilistic modeling approaches. These hybrid methods aim to combine the representational power of deep neural networks with the principled uncertainty quantification and temporal modeling capabilities of probabilistic graphical models[31]. However, most work has focused on specific domains rather than comprehensive system-wide anomaly detection.

The regulatory and operational requirements of modern cybersecurity environments have created additional demands for anomaly detection systems. Compliance frameworks require auditable decision-making processes, while operational requirements demand real-time processing capabilities and minimal false positive rates. These practical constraints influence the design of anomaly detection systems and favor approaches that provide both high accuracy and interpretable explanations.

3. Methodology

3.1 Convolutional Neural Network Architecture for Multi-Modal Feature Extraction

The foundation of our cross-modal learning framework lies in sophisticated convolutional neural network architectures designed to extract meaningful features from heterogeneous system data modalities. Our approach processes multiple data sources including network traffic patterns, system performance metrics, log sequences, and process execution traces through specialized CNN architectures that preserve the unique characteristics of each data type while enabling cross-modal correlation analysis.

Our feature extraction pipeline implements a hierarchical CNN architecture that processes input data through multiple convolutional layers with progressively increasing abstraction levels. The architecture begins with low-level feature extraction that captures local patterns and dependencies within individual data modalities, followed by higher-level feature integration that identifies cross-modal relationships and system-wide patterns.




Figure 1. Convolutional Layer

The convolutional layers employ specialized filter designs that account for the temporal and spatial characteristics of system data. Network traffic data is processed using temporal

convolutions that capture packet flow patterns and protocol-specific behaviors, while system performance metrics are analyzed through filters designed to identify resource utilization patterns and performance anomalies. Log data undergoes sequence-based convolution that preserves the ordering and semantic relationships between log entries.

The architecture incorporates multiple pooling strategies to handle the varying temporal scales present in system data. Max pooling operations preserve critical peak values that often indicate anomalous system events, while average pooling maintains overall trend information necessary for understanding baseline system behavior. The combination of different pooling strategies enables the model to capture both acute anomalous events and gradual system degradation patterns.

Feature fusion occurs at multiple levels within the architecture, enabling the model to learn cross-modal relationships at different abstraction levels. Early fusion combines low-level features from different modalities to capture immediate cross-system correlations, while late fusion integrates high-level semantic features to identify complex multi-modal anomaly patterns that span different system components.

3.2 Hierarchical Dirichlet Process Hidden Markov Models for Temporal Dynamics

The temporal modeling component of our framework employs hierarchical Dirichlet process hidden Markov models to capture the complex temporal dependencies and state transitions that characterize system behavior. This probabilistic approach addresses the fundamental challenge of modeling system dynamics with unknown numbers of behavioral states and complex transition patterns between normal and anomalous operating modes.

Our HDP-HMM implementation extends standard hidden Markov models by incorporating hierarchical Dirichlet priors that enable automatic inference of the appropriate number of hidden states for representing system behavior. This non-parametric approach eliminates the need for manual specification of state space sizes while providing principled uncertainty quantification for anomaly detection decisions.




Figure 2. Model Architecture

The model architecture in Figure 2 incorporates switching vector autoregressive components that enable the capture of complex temporal dependencies within system observations. Each

hidden state maintains its own autoregressive parameters, allowing the model to learn distinct temporal patterns associated with different system operating modes. This capability is crucial for identifying subtle anomalies that manifest as changes in temporal correlation structures rather than absolute value deviations.

State transition probabilities are learned through the hierarchical Dirichlet process prior, which encourages sparsity in the transition matrix while allowing for the discovery of rare but important state transitions that may indicate anomalous system behaviors. The hierarchical structure enables sharing of statistical strength across similar system states while maintaining flexibility to model system-specific behavioral patterns.

The observation models incorporate multi-modal components that can handle the diverse data types generated by system monitoring. Continuous variables such as CPU utilization and memory consumption are modeled using appropriate continuous distributions, while discrete variables such as log event types are handled through categorical distributions. The framework automatically learns the appropriate mixture components for each observation type.

Inference in the HDP-HMM framework employs variational Bayesian techniques that provide efficient approximate inference while maintaining principled uncertainty quantification. The variational approach enables real-time processing of streaming system data while providing confidence estimates for anomaly detection decisions that are crucial for operational deployment.

3.3 Mixture-of-Experts Framework for Adaptive Anomaly Detection

The final component of our framework implements a mixture-of-experts architecture that dynamically combines multiple detection mechanisms to provide robust and adaptive anomaly detection capabilities. This approach recognizes that different types of system anomalies may be optimally detected by different algorithmic approaches, and provides principled methods for selecting and combining detection strategies based on current system context and data characteristics.

Our mixture-of-experts framework integrates the CNN-based feature extraction with the HDP-HMM temporal modeling through a confidence-based fusion mechanism. Each expert component provides both anomaly predictions and confidence estimates, enabling the framework to dynamically weight the contributions of different detection approaches based on their reliability for specific system conditions.




Figure 3. Framework.

The framework in figure 3 includes specialized expert components designed for different aspects of system anomaly detection. Statistical experts focus on numerical anomalies in performance metrics and resource utilization patterns, employing techniques such as Gaussian mixture regression for modeling expected system behavior ranges. Sequence-based experts analyze temporal patterns in log data and system call sequences, using the HDP-HMM models to identify unusual temporal dependencies and state transitions.

Environmental context experts leverage visual and semantic information to understand system interaction patterns and identify anomalies that manifest through changes in system environment or user interaction behaviors. These experts employ advanced language models and computer vision techniques to analyze system interfaces, user activities, and external environment factors that may influence system behavior.

The confidence estimation mechanism employs principled statistical approaches to assess the reliability of each expert's predictions. Confidence scores are computed based on prediction uncertainty, historical performance on similar data patterns, and consistency with other expert opinions. The framework maintains adaptive confidence calibration that adjusts expert reliability estimates based on observed performance over time.

The fusion algorithm implements dynamic weighting strategies that select optimal expert combinations based on current system conditions. During periods of high system activity or unusual conditions, the framework may rely more heavily on statistical experts that can quickly identify numerical anomalies. During normal operations, sequence-based experts may receive higher weights for detecting subtle behavioral changes that indicate emerging threats.

The framework provides interpretable output through expert-specific explanations that enable security analysts to understand which detection mechanisms contributed to specific anomaly alerts. This interpretability is crucial for incident response and forensic analysis, as it enables analysts to focus their investigation efforts on the most relevant system components and data sources.

4. Results and Discussion

4.1 Experimental Setup and Multi-Modal Performance Analysis

The experimental evaluation of our attention-enhanced cross-modal framework was conducted using comprehensive datasets collected from real-world enterprise computing environments and controlled laboratory settings. The evaluation encompassed both large-scale distributed systems and specialized robotic control environments to demonstrate the generalizability of our approach across different system software domains.

The primary dataset consists of multi-modal system monitoring data including network traffic logs, system performance metrics, security event logs, and process execution traces collected over six months from enterprise cloud infrastructure. This dataset contains approximately 2.8 million normal system events and 18,432 confirmed anomalous events spanning multiple categories including advanced persistent threats, zero-day exploits, system performance degradation, and configuration errors.

Performance evaluation employs the specialized CNN architecture with systematic analysis of feature extraction effectiveness across different data modalities. The convolutional layers demonstrate exceptional capability in automatically learning hierarchical features from raw system data, with the two-stage convolution and pooling architecture achieving optimal balance between feature preservation and computational efficiency. Layer-wise analysis reveals that the first convolutional layer captures local system behavior patterns such as resource utilization spikes and network traffic bursts, while the second layer learns higher-level semantic features corresponding to complex attack signatures and system failure modes.

Computational efficiency analysis shows that our CNN-based feature extraction processes system data batches with average latency of 12.8 milliseconds per 1000 events, well within real-time processing requirements for operational security systems. The hierarchical architecture enables parallel processing across multiple data modalities while maintaining memory efficiency through progressive dimensionality reduction from the initial $36 \times 36 \times 3$ input to the final compact feature representations.

Comparison with baseline feature extraction methods including manual feature engineering, principal component analysis, and standard autoencoder approaches demonstrates consistent superiority of our CNN architecture across all evaluation metrics. The learned features show 34% better discrimination capability between normal and anomalous system behaviors compared to manually engineered features, while requiring no domain expertise for feature design.

4.2 Temporal Modeling and Mixture-of-Experts Validation

The hierarchical Dirichlet process hidden Markov model component demonstrates exceptional performance in capturing temporal dependencies and state transitions in system behavior. The non-parametric approach automatically discovers appropriate numbers of hidden states for different system environments, with typical deployments converging to 8-12 distinct behavioral states representing normal operation modes, transitional states, and various anomaly categories.

State transition analysis reveals meaningful behavioral patterns that align with expected system dynamics. Normal operation states show high self-transition probabilities (0.85-0.92) indicating stable system behavior, while anomaly states exhibit lower self-transition probabilities (0.23-0.67) reflecting the transient nature of most anomalous events. The model successfully identifies rare but critical state transitions corresponding to attack progressions and system failure cascades.

The vector autoregressive components within each hidden state capture complex temporal correlations between different system metrics. Analysis of learned autoregressive parameters reveals that normal states exhibit predictable correlation patterns between related metrics such as CPU utilization and memory consumption, while anomaly states show disrupted correlation structures that serve as strong indicators of system compromise or failure.

Mixture-of-experts validation demonstrates the effectiveness of confidence-based fusion in selecting optimal detection strategies for different system conditions. The framework achieves dynamic expert selection with 89% accuracy in choosing the most reliable detection mechanism for specific anomaly types. Statistical experts show superior performance for numerical anomalies in system metrics (precision: 0.91, recall: 0.87), while sequence-based experts excel at detecting behavioral anomalies in log data (precision: 0.88, recall: 0.93).

The confidence estimation mechanism provides well-calibrated uncertainty quantification, with confidence scores showing strong correlation ($r=0.84$) with actual prediction accuracy across different expert types. This calibration enables reliable threshold setting for operational deployment and supports risk-based prioritization of anomaly alerts for human investigation.

Cross-modal correlation analysis reveals that our framework successfully identifies relationships between anomalies manifesting across different data modalities. Coordinated attacks showing signatures in both network traffic and system logs are detected with 96% accuracy, compared to 67% accuracy for single-modal approaches. The attention mechanisms effectively highlight which modalities contribute most strongly to specific anomaly detections, providing valuable insights for incident response.

The mixture-of-experts approach demonstrates particular strength during system transitions and unusual operating conditions where individual expert performance may degrade. During planned maintenance windows and system updates, the framework maintains detection

accuracy within 5% of normal operation levels by dynamically adjusting expert weights based on changing system characteristics.

Interpretability analysis shows that expert-specific explanations enable security analysts to quickly understand anomaly root causes, with average investigation time reduced by 43% compared to black-box detection systems. The probabilistic framework provides meaningful uncertainty estimates that support risk-based decision making and help analysts prioritize response efforts for the most critical threats.

5. Conclusion

This research presents a significant advancement in system software anomaly detection through the development of a comprehensive attention-enhanced cross-modal learning framework that integrates convolutional neural networks, hierarchical Dirichlet process hidden Markov models, and mixture-of-experts architectures. The integration of these sophisticated machine learning techniques creates a powerful system that can identify complex anomaly patterns while providing interpretable explanations for security practitioners.

The key innovations of our work include the development of specialized CNN architectures for multi-modal feature extraction from heterogeneous system data, non-parametric probabilistic models that capture temporal dependencies and switching dynamics without requiring manual state space specification, and mixture-of-experts frameworks that dynamically select optimal detection strategies based on confidence estimation and system context. These innovations demonstrate how advanced machine learning architectures can be successfully adapted for cybersecurity applications while maintaining the real-time processing capabilities required for operational deployment.

Experimental results validate the effectiveness of our approach across multiple evaluation dimensions. The CNN-based feature extraction achieves superior performance compared to traditional feature engineering approaches, with 34% better discrimination between normal and anomalous behaviors. The HDP-HMM temporal modeling successfully captures complex state transition dynamics and provides principled uncertainty quantification for detection decisions. The mixture-of-experts framework demonstrates robust performance across diverse system conditions with 96% accuracy for coordinated multi-modal attacks.

The practical implications of this work extend across multiple areas of cybersecurity and system administration. Security operations centers can leverage the cross-modal correlation capabilities to detect sophisticated attacks that evade traditional security tools. System administrators gain access to comprehensive monitoring capabilities that can identify both security threats and reliability issues through unified analysis frameworks. The probabilistic approach provides confidence estimates that enable risk-based prioritization of security alerts and support efficient allocation of investigation resources.

The interpretability features of our framework address a critical gap in existing anomaly detection systems. The mixture-of-experts approach provides expert-specific explanations that

enable analysts to understand which detection mechanisms contributed to specific alerts, while the probabilistic framework quantifies prediction uncertainty to support informed decision-making. This interpretability reduces average investigation time by 43% and enables more effective incident response procedures.

Future research directions include extending the framework to handle additional data modalities such as user behavior analytics and external threat intelligence feeds, developing adaptive learning techniques that can automatically adjust to new system configurations and evolving attack methodologies, and creating specialized architectures for specific system domains such as industrial control systems and cloud computing environments.

The broader implications of this work extend beyond cybersecurity to other domains where multi-modal anomaly detection is critical. The principles and techniques developed for system software monitoring can potentially be adapted for applications in healthcare monitoring, industrial process control, and financial fraud detection, where understanding complex relationships across different data sources is essential for effective anomaly detection.

As cyber threats continue to evolve in sophistication and system environments become increasingly complex, the need for advanced anomaly detection systems that can understand and correlate information across multiple data sources becomes increasingly critical. Our attention-enhanced cross-modal framework provides a robust foundation for meeting these challenges while maintaining the interpretability and operational efficiency necessary for practical deployment in security-critical environments.

The successful integration of convolutional neural networks, probabilistic graphical models, and mixture-of-experts architectures demonstrated in this work provides a template for developing next-generation security systems that can adapt to evolving threats while providing the transparency and interpretability required for effective security operations. The combination of advanced machine learning capabilities with practical security requirements demonstrates a path forward for creating trustworthy and effective artificial intelligence systems in the demanding environment of cybersecurity and system protection.

References

- [1] Aldea, C. L., Bocu, R., & Solca, R. N. (2023). Real-time monitoring and management of hardware and software resources in heterogeneous computer networks through an integrated system architecture. *Symmetry*, 15(6), 1134.
- [2] Ji, T., Vuppala, S. T., Chowdhary, G., & Driggs-Campbell, K. (2020). Multi-modal anomaly detection for unstructured and uncertain environments. *arXiv preprint arXiv:2012.08637*.
- [3] Xing, S., & Wang, Y. (2025). Cross-Modal Attention Networks for Multi-Modal Anomaly Detection in System Software. *IEEE Open Journal of the Computer Society*.
- [4] Aslan, Ö., Aktuğ, S. S., Ozkan-Okay, M., Yilmaz, A. A., & Akin, E. (2023). A comprehensive review of cyber security vulnerabilities, threats, attacks, and solutions. *Electronics*, 12(6), 1333.

- [5] Bhatti, U. A., Tang, H., Wu, G., Marjan, S., & Hussain, A. (2023). Deep learning with graph convolutional networks: An overview and latest applications in computational intelligence. *International Journal of Intelligent Systems*, 2023(1), 8342104.
- [6] Gu, A., Johnson, I., Goel, K., Saab, K., Dao, T., Rudra, A., & Ré, C. (2021). Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in neural information processing systems*, 34, 572-585.
- [7] Olayinka, O. H. (2021). Big data integration and real-time analytics for enhancing operational efficiency and market responsiveness. *Int J Sci Res Arch*, 4(1), 280-96.
- [8] Jeffrey, N., Tan, Q., & Villar, J. R. (2023). A review of anomaly detection strategies to detect threats to cyber-physical systems. *Electronics*, 12(15), 3283.
- [9] Apostolakos, G. (2024). Operational Anomaly Detection Using Clustering Methods and Machine Learning Models.
- [10] Kwon, D., Kim, H., Kim, J., Suh, S. C., Kim, I., & Kim, K. J. (2019). A survey of deep learning-based network anomaly detection. *Cluster Computing*, 22(Suppl 1), 949-961.
- [11] Nascita, A., Aceto, G., Ciuonzo, D., Montieri, A., Persico, V., & Pescapé, A. (2024). A survey on explainable artificial intelligence for internet traffic classification and prediction, and intrusion detection. *IEEE Communications Surveys & Tutorials*.
- [12] Torabi, H., Mirtaheri, S. L., & Greco, S. (2023). Practical autoencoder based anomaly detection by using vector reconstruction error. *Cybersecurity*, 6(1), 1.
- [13] Sucar, L. E. (2021). Probabilistic graphical models. Springer International Publishing.
- [14] Moraffah, B. (2024). Bayesian nonparametrics: An alternative to deep learning. *arXiv preprint arXiv:2404.00085*.
- [15] Samariya, D., & Thakkar, A. (2023). A comprehensive survey of anomaly detection algorithms. *Annals of Data Science*, 10(3), 829-850.
- [16] Ji, E., Wang, Y., Xing, S., & Jin, J. (2025). Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems. *IEEE Access*.
- [17] Jin, J., Xing, S., Ji, E., & Liu, W. (2025). XGate: Explainable Reinforcement Learning for Transparent and Trustworthy API Traffic Management in IoT Sensor Networks. *Sensors (Basel, Switzerland)*, 25(7), 2183.
- [18] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical Deep Reinforcement Learning for Multi-Objective Integrated Circuit Physical Layout Optimization with Congestion-Aware Reward Shaping. *IEEE Access*.
- [19] Zheng, W., & Liu, W. (2025). Symmetry-Aware Transformers for Asymmetric Causal Discovery in Financial Time Series. *Symmetry*.
- [20] Cao, W., Mai, N. T., & Liu, W. (2025). Adaptive knowledge assessment via symmetric hierarchical Bayesian neural networks with graph symmetry-aware concept dependencies. *Symmetry*, 17(8), 1332.
- [21] Mai, N. T., Cao, W., & Wang, Y. (2025). The global belonging support framework: Enhancing equity and access for international graduate students. *Journal of International Students*, 15(9), 141-160.
- [22] Tan, Y., Wu, B., Cao, J., & Jiang, B. (2025). LLaMA-UTP: Knowledge-Guided Expert Mixture for Analyzing Uncertain Tax Positions. *IEEE Access*.
- [23] Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. *IEEE Access*.

- [24] Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-Distance Spatial-Temporal Graph Neural Network for Anomaly Detection in Blockchain Transactions. *Advanced Intelligent Systems*, 2400898.
- [25] Zhang, Q., Chen, S., & Liu, W. (2025). Balanced Knowledge Transfer in MTTL-ClinicalBERT: A Symmetrical Multi-Task Learning Framework for Clinical Text Classification. *Symmetry*, 17(6), 823.
- [26] Shao, Z., Wang, X., Ji, E., Chen, S., & Wang, J. (2025). GNN-EADD: Graph Neural Network-based E-commerce Anomaly Detection via Dual-stage Learning. *IEEE Access*.
- [27] Liu, Y., Ren, S., Wang, X., & Zhou, M. (2024). Temporal logical attention network for log-based anomaly detection in distributed systems. *Sensors*, 24(24), 7949.
- [28] Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization. *Applied Sciences*, 15(2), 951.
- [29] Cao, J., Zheng, W., Ge, Y., & Wang, J. (2025). DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting. *IEEE Open Journal of the Computer Society*.
- [30] Wang, J., Liu, J., Zheng, W., & Ge, Y. (2025). Temporal Heterogeneous Graph Contrastive Learning for Fraud Detection in Credit Card Transactions. *IEEE Access*.
- [31] Mai, N. T., Cao, W., & Liu, W. (2025). Interpretable Knowledge Tracing via Transformer-Bayesian Hybrid Networks: Learning Temporal Dependencies and Causal Structures in Educational Data. *Applied Sciences*, 15(17), 9605.