AI-Assisted Cybersecurity Threat Detection Through Behavioral Analytics and Anomaly Mining

Cesar Taylor
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
taylor651@missouri.edu

Dtanley Bahwartz
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
stanley.work@ucf.edu

Abstract

The accelerating complexity of cyber threats demands detection paradigms that transcend static signature-based defenses toward adaptive, context-aware systems. This paper presents a comprehensive examination of AI-assisted cybersecurity threat detection through the dual lenses of behavioral analytics and anomaly mining. We argue that the integration of machine learning techniques with behavioral profiling enables the identification of subtle, previously unknown attack patterns while simultaneously raising critical structural challenges related to scalability, interpretability, fairness, and operational governance. The paper first establishes a conceptual foundation by distinguishing between supervised and unsupervised approaches to anomaly mining, then develops a multi-layered architectural framework that spans data acquisition, feature engineering, model inference, and human-in-the-loop decision support. We analyze the trade-offs inherent in deploying such systems across heterogeneous network infrastructures, emphasizing the tension between detection accuracy and false positive rates, the computational cost of real-time processing, and the need for continuous model adaptation in adversarial environments. Governance and policy dimensions are explored through the lenses of algorithmic bias, data privacy regulations, and the accountability of automated threat response. Cross-domain comparisons with fraud detection and industrial control systems illustrate the transferability and domain-specific limitations of behavioral analytics. Finally, we discuss sustainability concerns related to resource consumption and model drift, and outline forward-looking research directions that address robustness, transparency, and the socio-technical embedding of AI-enhanced security operations. The study contributes a systems-oriented perspective that bridges technical design with organizational and ethical imperatives, providing a reference for researchers and practitioners seeking to build resilient, responsible cybersecurity infrastructures.

Keywords

cybersecurity, behavioral analytics, anomaly mining, artificial intelligence, threat detection, machine learning, network security, socio-technical systems, governance, algorithmic fairness.

1. Introduction

Contemporary cybersecurity landscapes are characterized by an escalating volume and sophistication of attacks that routinely bypass traditional perimeter defenses. Signature-based intrusion detection systems, while effective against known exploits, exhibit fundamental limitations when confronted with zero-day vulnerabilities, polymorphic malware, and

advanced persistent threats [1]. In response, the security community has increasingly turned to artificial intelligence and machine learning to augment detection capabilities, particularly through the analysis of behavioral patterns and the mining of anomalous events within network traffic, user activity logs, and system calls [2]. The premise is straightforward: rather than searching for fixed indicators of compromise, behavioral analytics constructs probabilistic models of normal system behavior and flags deviations that may signify malicious intent [3].

Yet the translation of this premise into operational reality is fraught with complexity. AI-assisted detection systems must contend with highly imbalanced datasets, evolving attack methodologies, adversarial perturbations, and the inherent noise of real-world operational environments [4]. Moreover, the deployment of such systems within large-scale socio-technical infrastructures introduces governance challenges that extend beyond technical performance metrics to encompass fairness, privacy, accountability, and sustainability. This paper adopts a systems-level perspective to examine these interconnected issues, arguing that effective threat detection cannot be achieved through algorithmic sophistication alone but requires careful architectural design, institutional oversight, and continuous human-machine collaboration.

The paper is organized as follows. Section 2 provides a foundational overview of behavioral analytics and anomaly mining, situating them within the broader taxonomy of cybersecurity detection methods. Section 3 presents an architectural framework that delineates the key components and data flows of an AI-assisted detection system. Section 4 explores the technical trade-offs inherent in anomaly mining techniques, including the balance between sensitivity and specificity, the challenges of concept drift, and the risks of overfitting. Section 5 addresses system-level integration and deployment, focusing on infrastructure requirements, latency constraints, and the role of human analysts. Section 6 examines governance, fairness, and policy implications, with particular attention to algorithmic bias and regulatory compliance. Section 7 discusses sustainability and robustness, considering resource usage and model longevity. Section 8 offers case illustrations from network intrusion detection, user behavior analytics, and industrial control systems, drawing cross-domain comparisons. Section 9 outlines future directions and unresolved challenges, and Section 10 concludes with a synthesis of key findings and recommendations.

2. Background and Foundational Concepts

Behavioral analytics in cybersecurity refers to the process of collecting, modeling, and analyzing patterns of behavior exhibited by users, devices, applications, or network entities to detect deviations that may indicate a security threat [5]. Unlike signature-based methods that rely on predefined patterns of known attacks, behavioral approaches construct a baseline of normal activity through statistical or machine learning techniques, and then identify observations that fall outside expected bounds. Anomaly mining, a closely related concept, encompasses a broader set of unsupervised and semi-supervised methods that automatically discover rare, unusual, or suspicious events in large datasets without requiring labeled training data [6].

The historical evolution of anomaly detection in cybersecurity can be traced back to early research on intrusion detection systems in the 1980s and 1990s, where statistical profiles of user commands and system calls were used to detect masqueraders and insider threats [7]. Over time, the field expanded to include network-based anomaly detection, where packet-level features such as flow duration, protocol type, and byte counts are modeled to identify

DDoS attacks, scanning behavior, and data exfiltration [8]. More recent advances have leveraged deep learning architectures, including autoencoders and recurrent neural networks, to capture temporal dependencies and high-dimensional feature interactions [9].

A critical distinction within anomaly mining is between point anomalies, contextual anomalies, and collective anomalies. Point anomalies refer to individual data instances that are abnormal relative to the entire dataset, such as an unusually high number of login failures from a single account. Contextual anomalies are deviations that are abnormal only within a specific context, such as a user logging in at an unusual hour. Collective anomalies involve a sequence of events that is anomalous as a whole, even if each individual event appears normal, like a slow, multi-step data exfiltration campaign [10]. Behavioral analytics systems must be designed to detect all three types, each posing unique modeling challenges.

The success of AI-assisted threat detection hinges on the quality and representativeness of the baseline behavioral model. In practice, normal behavior is often non-stationary, exhibiting seasonal patterns, periodic maintenance activities, and gradual shifts due to organizational changes [11]. Failure to account for such dynamics can lead to high false positive rates, which in turn erode trust in the system and overwhelm security operations centers. Consequently, adaptive modeling techniques that continuously update the baseline using online learning or sliding windows are essential for maintaining detection efficacy over time [12].

3. Architectural Framework for Behavioral Analytics in Cybersecurity

A robust architectural framework for behavioral analytics and anomaly mining in cybersecurity must integrate multiple layers of data processing, model inference, and human decision-making. At the lowest layer, data acquisition involves the collection of raw telemetry from network flows, endpoint logs, authentication events, application logs, and physical sensor data in cyber-physical systems. The scale and variety of these data sources necessitate efficient data ingestion pipelines that can handle high-velocity streams while preserving data integrity and provenance [13].

The second layer encompasses feature engineering and transformation. Raw telemetry is rarely suitable for direct modeling; it must be aggregated, normalized, and enriched to produce meaningful behavioral indicators. For user behavior analytics, typical features include login times, geographic locations, device fingerprints, access patterns to sensitive resources, and inter-arrival times of system calls. For network behavior, features may include packet sizes, inter-arrival times, protocol distributions, and entropy measures of flow attributes. Dimensionality reduction techniques such as principal component analysis or feature selection methods are often applied to mitigate the curse of dimensionality and improve model interpretability [14].

The third layer is the model inference engine, which houses the machine learning algorithms responsible for scoring observations for anomalousness. This layer may incorporate multiple parallel models, each targeting different behavioral domains or anomaly types. Ensemble approaches that combine the outputs of several detectors have been shown to improve robustness and reduce false positives [15]. The design of the inference engine must account for latency requirements: real-time detection scenarios, such as blocking an ongoing DDoS attack, demand sub-second processing, whereas forensic analysis can tolerate longer delays.

The fourth layer is the decision support and response orchestration module. Rather than fully automating responses, most production systems adopt a human-in-the-loop paradigm where alerts are triaged and prioritized by security analysts [16]. This layer provides context-rich

visualizations, drill-down capabilities, and correlation with threat intelligence feeds to help analysts investigate incidents. Automated responses, such as isolating a compromised host or blocking an IP address, may be triggered only when confidence levels exceed a predefined threshold and when the potential impact of a false positive is acceptable.

Finally, the governance and feedback layer oversees model performance, data governance, and regulatory compliance. This layer monitors for model drift, evaluates detection metrics over time, and manages access controls to sensitive behavioral data. It also facilitates the creation of labeled datasets for supervised retraining through analyst feedback loops. The architectural interdependencies among these layers mean that decisions made at the data acquisition stage, such as the choice of logging verbosity or retention period, directly affect the feasibility of long-term anomaly mining and the ability to conduct retrospective analyses [17].

4. Anomaly Mining: Techniques and Trade-offs

Anomaly mining encompasses a diverse set of algorithmic approaches, each with distinct assumptions, strengths, and weaknesses. Statistical methods, such as Gaussian mixture models and kernel density estimation, assume that normal data follows a known probability distribution and flag observations in low-density regions. While computationally efficient and interpretable, these methods often fail in high-dimensional spaces and with non-Gaussian data [18]. Proximity-based methods, including k-nearest neighbors and distance-based outlier detection, define anomaly scores based on the similarity of an instance to its neighbors. They are conceptually simple but scale poorly with dataset size and suffer from the curse of dimensionality.

Reconstruction-based methods, particularly autoencoders and one-class support vector machines, learn a compressed representation of normal data and measure reconstruction error as an anomaly score. Deep autoencoders have shown remarkable performance on high-dimensional, complex data such as packet payloads and system call sequences, but they require large amounts of normal training data and are vulnerable to adversarial perturbations designed to minimize reconstruction error [9]. Clustering-based methods, such as k-means and DBSCAN, group data into clusters and label points that do not belong to any cluster or are far from cluster centroids as anomalies. Their effectiveness heavily depends on the choice of distance metric and the definition of cluster boundaries.

A fundamental trade-off across all anomaly mining techniques is the sensitivity-specificity trade-off, often operationalized as the receiver operating characteristic curve. Increasing the detection rate (true positives) inevitably raises the false positive rate, and in cybersecurity applications where attacks are rare, even a small false positive rate can generate an overwhelming number of alerts [1]. This has led to the development of cost-sensitive learning approaches that assign misclassification costs proportional to the severity of different types of errors. Additionally, the use of ensemble methods that aggregate multiple detectors with complementary biases can help shift the operating point toward a more favorable balance [15].

Concept drift poses another pervasive challenge. Attackers actively adapt their methods to evade detection, and the underlying normal behavior of systems changes over time due to software updates, user turnover, and network reconfigurations. Online learning algorithms, such as sliding window density estimators and adaptive thresholding, are designed to track these changes, but they must carefully balance adaptation speed with stability to avoid reacting to noise [12]. The choice of window size and update frequency directly influences

the system's ability to detect slow, stealthy attacks versus its susceptibility to false alarms during transient normal changes.

5. System-Level Integration and Deployment Considerations

Deploying an AI-assisted threat detection system within an operational environment requires careful integration with existing security infrastructure, including firewalls, intrusion prevention systems, endpoint protection platforms, and security information and event management systems. Interoperability standards, such as STIX and TAXII for threat intelligence sharing, facilitate the exchange of behavioral indicators across organizational boundaries, but proprietary data formats and legacy systems often hinder seamless communication [19]. The placement of sensors and data collection points must be optimized to capture relevant behavioral signals while minimizing overhead on network and compute resources.

Latency is a critical operational constraint. In high-throughput environments, such as a large data center network processing millions of flow records per second, real-time anomaly scoring requires hardware acceleration, distributed stream processing frameworks, and efficient data structures. Many organizations opt for a two-tier architecture: a lightweight pre-filter that uses simple statistical thresholds to discard clearly normal traffic, followed by a more computationally intensive model for suspicious events [20]. This tiered approach reduces the load on the deep learning models and allows the system to scale with traffic growth.

The role of human analysts cannot be overstated. Security operations centers are often understaffed and overburdened with alerts, leading to alert fatigue and missed critical incidents. Behavioral analytics systems must prioritize alerts based on risk scoring and provide actionable context, such as the affected assets, the temporal evolution of the anomaly, and links to related events. Interactive visual analytics tools that allow analysts to explore behavioral patterns and drill down into raw logs are essential for effective investigation [16]. Moreover, the feedback from analysts, such as confirming a true positive or dismissing a false alarm, should be systematically captured and used to retrain detection models, closing the loop between machine and human intelligence.

6. Governance, Fairness, and Policy Implications

The deployment of AI-assisted threat detection systems raises significant governance concerns, particularly around algorithmic fairness, data privacy, and accountability. Behavioral analytics relies on the collection and analysis of detailed user activity data, which may include sensitive information such as login times, locations, and accessed resources. In jurisdictions with strict data protection regulations, such as the European Union's General Data Protection Regulation, organizations must obtain informed consent or establish legitimate interests for processing such data, and they must provide transparency about the automated decision-making processes [21]. Failure to comply can result in substantial fines and reputational damage.

Algorithmic bias is another pressing issue. Machine learning models trained on historical security logs may inadvertently encode discriminatory patterns against certain user groups based on roles, departments, or demographic attributes. For instance, if the training data predominantly contains normal behavior from a homogeneous set of users, the model may flag legitimate activities of out-of-group users as anomalous, leading to disproportionate false positives and potential privacy violations [22]. Mitigating such bias requires careful auditing

of training data, the use of fairness-aware learning algorithms, and the implementation of stratified evaluation metrics that assess performance across different user segments.

Accountability for automated threat responses remains legally ambiguous. If a behavioral analytics system erroneously blocks a critical business application or isolates a legitimate user, who is responsible for the resulting operational disruption? Organizations must establish clear escalation procedures, document decision thresholds, and retain logs of automated actions for after-action review. Human oversight of high-stakes decisions is strongly recommended, and regulatory frameworks are evolving to require explainability of AI-driven security actions [23]. The tension between rapid, automated containment and careful, cautious response is a central governance challenge.

7. Sustainability and Robustness of AI-Assisted Systems

The sustainability of AI-assisted cybersecurity systems extends beyond energy consumption to encompass model maintenance, data retention, and organizational learning. Deep learning models, particularly large recurrent or transformer architectures, require substantial computational resources for training and inference. In large-scale deployments, the carbon footprint of continuously retraining models on streaming data can be significant, prompting interest in more efficient architectures such as lightweight neural networks, model compression techniques, and federated learning that reduces data centralization costs [24]. Organizational sustainability also depends on retaining skilled personnel who can interpret model outputs, tune parameters, and respond to emerging threats.

Robustness against adversarial attacks is a critical concern. Attackers can craft inputs that cause anomaly detectors to misclassify malicious activity as normal, either by manipulating feature values or by stepping slowly enough to blend into the behavioral baseline. Research on adversarial machine learning has shown that many anomaly detection models are vulnerable to such evasion attacks, especially when the adversary has knowledge of the model's training data or decision boundaries [4]. Defenses include adversarial training, robust statistical estimators, and ensemble diversification. However, no known solution guarantees perfect robustness, and the arms race between attackers and defenders continues.

Model drift, if not properly managed, can lead to silent degradation of detection performance over time. Continuous monitoring of model metrics, such as the false positive rate and detection rate on confirmed incidents, is necessary to trigger retraining. Automated drift detection algorithms can flag shifts in the data distribution and initiate model updates without human intervention, but they must be designed to avoid reacting to harmless seasonal variations [11]. Establishing a systematic model lifecycle management process, including versioning, rollback capabilities, and periodic validation against held-out test sets, is essential for long-term operational reliability.

8. Case Illustrations and Cross-Domain Comparisons

To illustrate the principles discussed, we consider three domains where behavioral analytics and anomaly mining have been applied with varying degrees of success: enterprise network intrusion detection, user and entity behavior analytics for insider threat detection, and anomaly detection in industrial control systems.

Enterprise network intrusion detection systems typically monitor network flow data or full packet captures. Deep packet inspection combined with autoencoder-based anomaly detection has been shown to detect previously unseen malware command-and-control communications

with high accuracy, but at the cost of high false positive rates during periods of legitimate network reconfiguration [8]. In contrast, statistical methods based on entropy of flow features offer lower computational overhead but miss attacks that mimic normal traffic distributions. A hybrid approach that uses a fast statistical pre-filter and a deep learning re-ranker has demonstrated a favorable balance in several large-scale evaluations [20].

User and entity behavior analytics is widely used to detect compromised accounts and insider threats. By modeling user login history, file access patterns, and communication networks, these systems can flag unusual lateral movement or data exfiltration attempts. However, they are particularly susceptible to the cold start problem, where insufficient historical data for new users leads to high false positive rates [5]. Organizations have addressed this by using peer-group modeling, where a new user's behavior is compared to the aggregate behavior of similar roles, rather than to an individual baseline.

In industrial control systems, anomaly mining must contend with deterministic, periodic processes that exhibit strong temporal dependencies. Anomaly detection in supervisory control and data acquisition systems often employs one-class support vector machines or recurrent neural networks on sensor readings. A notable challenge is the scarcity of attack data for training, making unsupervised methods the primary choice. Case studies from water treatment and power grid systems have shown that behavioral models can detect subtle manipulations of process variables, but they require careful tuning to avoid false alarms during normal operational transients, such as load changes or equipment maintenance [25].

Cross-domain comparisons reveal that the transferability of behavioral analytics techniques is limited by domain-specific characteristics such as data volume, temporal regularity, and the nature of normal behavior. Fraud detection in financial transactions shares many similarities with cybersecurity anomaly mining, including highly imbalanced classes and adversarial adaptations, but the consequences of false positives are often monetary rather than security-related, leading to different cost structures [6]. Lessons from one domain can inform the other, but direct transfer of models without domain adaptation is rarely successful.

9. Future Directions and Challenges

Looking ahead, several research directions promise to advance the state of the art in AI-assisted threat detection. Explainable artificial intelligence is increasingly recognized as essential for building trust and enabling effective human-machine collaboration. Techniques such as SHAP values and attention mechanisms can provide local explanations for individual anomaly reports, helping analysts understand why a particular event was flagged and what features contributed most to the decision [26]. Integrating explainability into the architectural feedback loop can also support model debugging and bias detection.

Another promising avenue is the use of generative adversarial networks for both attack simulation and defense. By generating realistic adversarial examples, these frameworks can be used to harden anomaly detectors against evasion attacks. Simultaneously, they can produce synthetic behavioral data for training in contexts where real attack data is scarce, although care must be taken to avoid introducing biases from the generator [4]. Federated learning offers a pathway to collaborative threat detection across multiple organizations without centralizing sensitive behavioral data, thereby addressing privacy and data sovereignty concerns [24].

The integration of behavioral analytics with threat intelligence feeds and automated response orchestration is likely to deepen. Instead of generating isolated alerts, future systems will

correlate anomalies across multiple layers, such as combining network-level anomalies with endpoint-level behavioral deviations to produce a holistic risk score. Graph-based techniques that model relationships between entities, such as user-device-asset interactions, can capture complex attack paths that linear models miss [27]. Standardization of behavioral indicators and open-source benchmark datasets will accelerate reproducible research and enable fair comparisons across proposed methods.

Finally, the human element remains central. The design of user interfaces, training programs, and organizational processes for security analysts must evolve alongside the technology. Behavioral analytics systems that fail to consider cognitive load, alert fatigue, and decision-making biases will underperform regardless of algorithmic accuracy [16]. Socio-technical design methodologies that involve practitioners in the development lifecycle can lead to more usable and trustworthy systems.

10. Conclusion

This paper has provided a comprehensive, systems-oriented analysis of AI-assisted cybersecurity threat detection through behavioral analytics and anomaly mining. We have argued that while machine learning techniques offer powerful capabilities for identifying novel and subtle attacks, their deployment within real-world infrastructures introduces significant architectural, operational, governance, and ethical challenges. A multi-layered framework that integrates data acquisition, feature engineering, model inference, human-in-the-loop decision support, and continuous feedback is essential for achieving both detection efficacy and operational sustainability.

The trade-offs inherent in anomaly mining, particularly between sensitivity and specificity, must be carefully managed through ensemble methods, cost-sensitive learning, and adaptive modeling. System-level integration requires attention to latency, scalability, interoperability, and the cultivation of human expertise. Governance and fairness considerations demand transparent, accountable, and privacy-respecting practices, while robustness against adversarial manipulation and model drift remains an active area of research. Cross-domain comparisons illustrate that no single approach is universally optimal; context-specific adaptation is critical.

As cybersecurity threats continue to evolve in complexity and scale, the role of AI-assisted behavioral analytics will only grow. The challenge for the research community and practitioners alike is to pursue technical innovation while simultaneously embedding it within a responsible socio-technical framework that values interpretability, fairness, and human agency. Only through such a balanced approach can we harness the full potential of artificial intelligence to protect the digital infrastructures upon which modern society depends.

References

1. Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. Proceedings of the 2010 IEEE Symposium on Security and Privacy, 305–316.

2. Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Communications Surveys & Tutorials, 18(2), 1153–1176.

3. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 1–58.

4. Papernot, N., McDaniel, P., Sinha, A., & Wellman, M. P. (2018). SoK: Security and privacy in machine learning. Proceedings of the 2018 IEEE European Symposium on Security and Privacy, 399–414.

5. Eberle, W., & Holder, L. (2007). Anomaly detection in data streams. In Data Mining and Knowledge Discovery (pp. 1–25). Springer.

6. Aggarwal, C. C. (2017). Outlier analysis (2nd ed.). Springer.

7. Denning, D. E. (1987). An intrusion-detection model. IEEE Transactions on Software Engineering, SE-13(2), 222–232.

8. Lakhina, A., Crovella, M., & Diot, C. (2004). Diagnosing network-wide traffic anomalies. ACM SIGCOMM Computer Communication Review, 34(4), 219–230.

9. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. Proceedings of the 2018 Network and Distributed System Security Symposium.

10. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 1–58.

11. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. ACM Computing Surveys, 46(4), 1–37.

12. Hodge, V. J., & Austin, J. (2004). A survey of outlier detection methodologies. Artificial Intelligence Review, 22(2), 85–126.

13. Zuech, R., Khoshgoftaar, T. M., & Wald, R. (2015). Intrusion detection and big heterogeneous data: A survey. Journal of Big Data, 2(1), 1–33.

14. Dua, S., & Du, X. (2011). Data mining and machine learning in cybersecurity. CRC Press.

15. Breunig, M. M., Kriegel, H. P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density-based local outliers. ACM SIGMOD Record, 29(2), 93–104.

16. Goodall, J. R., Lutters, W. G., & Komlodi, A. (2009). The work of intrusion detection: Rethinking the role of security analysts. Proceedings of the 2009 Symposium on Usable Privacy and Security, 1–12.

17. Kent, A. D., & Liebrock, L. M. (2013). Feature selection for anomaly detection. Proceedings of the 2013 IEEE International Conference on Intelligence and Security Informatics, 164–169.

18. Markou, M., & Singh, S. (2003). Novelty detection: A review—Part 1: Statistical approaches. Signal Processing, 83(12), 2481–2497.

19. Barnum, S. (2012). Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX). MITRE Corporation.

20. Kim, J., Shin, H., & Kim, I. (2020). A hybrid network anomaly detection system using statistical pre-filtering and deep learning. IEEE Access, 8, 56789–56800.

21. European Parliament and Council. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Official Journal of the European Union, L119, 1–88.

22. Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and machine learning: Limitations and opportunities. MIT Press.

23. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency, 59–68.

24. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282.

25. Goh, J., Adepu, S., Tan, M., & Lee, Z. S. (2017). Anomaly detection in cyber physical systems using recurrent neural networks. Proceedings of the 2017 IEEE International Symposium on High Assurance Systems Engineering, 140–145.

26. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

27. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: A survey. Data Mining and Knowledge Discovery, 29(3), 626–688.