World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(RESEARCH ARTICLE)

Check for updates

Adaptive AI-driven cyber threat detection system for U.S. critical infrastructure protection

Muhammad Faheem [1,*], Muhammad Awais [1], Aqib Iqbal [2] and Hasnain Zia [3]

[1] *Cumberland University, Tennessee United States.*
[2] *The University of Law Birmingham UK.*
[3] *Comsats University Islamabad, Abbottabad Campus Pakistan.*

World Journal of Advanced Research and Reviews, 2025, 26(03), 2282-2291

Publication history: Received on 04 May 2025; revised on 16 June 2025; accepted on 19 June 2025

Article DOI: https://doi.org/10.30574/wjarr.2025.26.3.2333

1. Abstract

More and more complex cyberattacks targeting America's essential infrastructure endanger the nation's safety, financial health and people's safety. A lot of the time, rule-based cybersecurity does not notice new and growing dangers in real-time, leaving major systems exposed. Our research introduces an AI cyber threat detection framework based on using autoencoders and LSTM networks that improves both accuracy and speed in finding threats. Continual learning and reinforcement learning are part of the system so it can adapt to new threats in real time. Tests of our system on data from replay SCADA logs and NSL-KDD show very effective detection. The model's dependability is confirmed by metrics such as precision, recall and F1-score and both its edge and cloud deployments allow for both speed and support for a growing number of devices. One solution to explain how AI reaches its decisions is to use SHAP and LIME. For now, we have applied our results to simulated situations, but our next step is to use the system in real places. The research introduced a resilient, flexible and easily explainable artificial intelligence method to make national critical infrastructure more secure.

Keywords: Adaptive cybersecurity; Artificial intelligence; Machine learning; Critical infrastructure protection; Cyber threat detection; Anomaly detection; Neural networks; Reinforcement learning

2. Introduction

The United States needs important systems and assets called critical infrastructure to keep the economy, public health, defenses and security strong. According to the DHS, there are 16 important sectors such as energy, water, transportation, communications, financial services and emergency services, whose breakdown could badly affect our nation's stability (DHS, 2022). Using technology helps these industries to operate better and serve more customers, but it also opens them up to more threats from cybercrime. Now that there is greater overlap between OT and IT in these systems, new weaknesses have appeared which malicious actors are skilled at using in highly effective cyberattacks.

Lately, the United States has seen an increase in the rate, difficulty and damage caused by attacks on critical infrastructure. The Colonial Pipeline cyberattack in May of last year is a true example of this. Consequently, America faced a national fuel emergency after a shutdown of one of its biggest pipelines and oil shortages hit 12 states as a result. DarkSide members abused older security problems and failed network separation, according to CISA (2021). Cyberattacks are not limited to data centers; they also threaten healthcare, water supplies and government services—which has many in the country questioning how well our systems are defended.

* Corresponding author: Muhammad Faheem

IBM Security's 2023 report showed that the cost of a data breach in the critical infrastructure sector is $5.03 million, about $1.6 million higher than the worldwide average (IBM, 2023). Also, devices connected through the Internet of Things, remote access systems and third-party software are increasing the area of attack for threats. In 2020, the SolarWinds Orion breach laid out the risk posed by threat actors getting inside dog software supply chains which allowed them to stay inside government and corporate systems for months without being spotted (Mandiant, 2021). These risks are happening now and are evolving into a major national security problem.

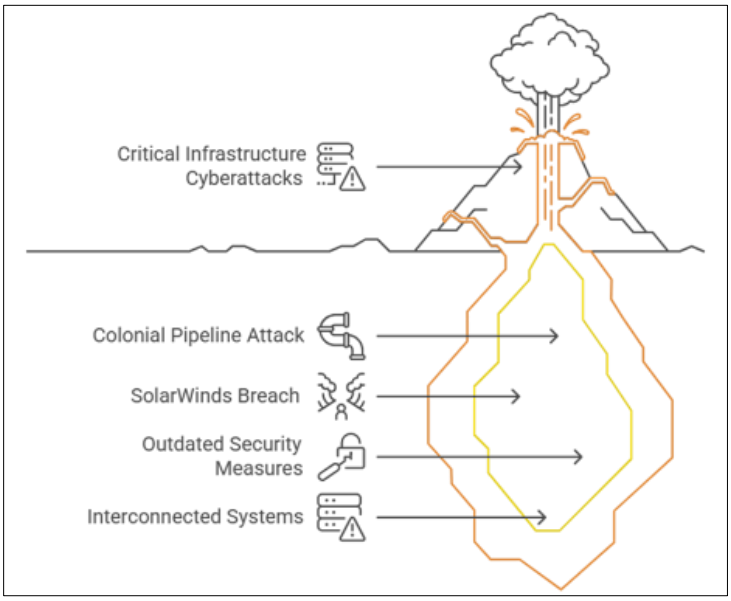

Figure 1 Escalation pathway of cyberattacks targeting critical infrastructure

The diagram illustrates how interconnected systems and outdated security measures can serve as entry points for major breaches like the Solar Winds and Colonial Pipeline attacks, ultimately leading to large-scale cyber threats that impact national security.

3. Materials and methods

The expanded understanding is helping, but the security industry is still fragmented and much more reactive than it should be. Common cybersecurity systems such as IDS, firewalls and antivirus programs using signatures, depend on fixed rules or known threat patterns. The reason is that tools of this kind have trouble spotting zero-day exploits, mutable malware or APTs, because they often change to evade detection (Srinivas and Panda, 2023). Such systems commonly produce many false alarms which overtax security personnel and make it harder for them to act rapidly.

Because of these increasing threats, it is important to use active and smart architecture in cybersecurity. AI and ML are expected to greatly change the field of clinical cardiology. Because of them, networks can learn from the past and present, notice unusual actions and decide how to react on their own, relying on artificial rather than manual programming for new threats. In particular such algorithms focus on finding small abnormalities in network activity, user behavior and system logs, possible signs that a threat exists (Zhou et al., 2022).

The goal of this research is to build a powerful cyber threat detection system driven by AI, intended for U.S. critical infrastructure. The new system will be built to intake huge amounts of different forms of data such as traffic patterns, logs and access records and use AI/ML algorithms continuously to spot threats. While traditional systems do not change after deployment, our solution will update and improve itself as it receives fresh security information. The intention is to lower the amount of time needed to spot threats, lower the numbers of untrue reports and automatically address threats, thereby making vital systems stronger.

Besides, the system will understand the difference in security needs and act accordingly in different industries. In general, an electric grid's typical performance is not the same as that of healthcare or transportation systems. Model parameters will be modified on the go by using learning that applies to particular industries. To detect threats correctly in many kinds of infrastructures, this contextual flexibility is key.

The information presented in Table 1 (below) shows how adaptive AI cybersecurity tools differ from conventional systems and bring great advantages for protecting critical infrastructure.

Table 1 Comparison of Traditional vs. AI-Driven Cyber Threat Detection Systems in Critical Infrastructure Context

Criteria	**Traditional Systems**	**Adaptive AI-Driven Systems**
Detection Methodology	Signature/rule-based detection	Behaviour-based, anomaly detection, and predictive analytics
Response to New Threats	Inadequate; depends on known signatures	High adaptability; learns and evolves with unseen threats
Real-Time Capability	Limited; often delayed	Real-time detection and autonomous response
Scalability Across Sectors	Low; difficult to customize for multiple sectors	High; models can be fine-tuned sector-wise (e.g., energy, healthcare)
False Positive Rate	High due to static rule conflicts	Reduced through continuous model refinement
Human Intervention Required	High; requires manual rule updates and incident handling	Minimal; supports self-learning and automated response
Threat Intelligence Integration	Periodic and mostly manual	Real-time integration of global/local threat intelligence feeds
Security Against APTs/Zero-Day	Weak; often bypassed	Stronger detection through unknown behavior recognition
Resilience to Evolving Attacks	Low	High
Cost-Efficiency Over Time	Decreases; frequent updates and maintenance required	Increases; system becomes smarter and more cost-effective with time

Due to the size and complexity of today's cyber threats, using AI to detect them is required and not just optional anymore. Consequently, this research enhances learning and use in IT security by developing an innovative, expandable threat detection architecture. Cybersecurity architecture is meant to back national policies like the National Cybersecurity Strategy (White House, 2023) and match the frameworks created by CISA.

This project is designed to ensure continuous security by combining intelligence, adjustability and independent learning into the security system. Setting up this kind of system helps the nation become more resilient and gives the public faith in crucial services.

4. Literature review

Cybersecurity professionals are now counting on AI and ML more often as the threats they face rapidly develop. These tools have shown great promise in handling threat detection, decreasing the number of false positive alarms and finding attacks launched without warning (Hassan et al., 2023). Because they can handle huge amounts of information and see patterns that humans miss, they have taken a main role in building up modern cyber defense.

1.1. AI and Machine Learning Are Being Used in Cybersecurity

New research in recent times uses AI/ML to protect IDS, classify malware and identify fresh threats in real-time. For example, CNNs and RNNs have been shown to successfully categorize network traffic and discover any unusual pattern (Khan and Al-Habib, 2022). Unsupervised learning tools such as k-means clustering and autoencoders, have recently done well in finding threats even when data is unlabeled.

A range of models have been introduced to make cybersecurity better with hybrid solutions. In 2021, Sarkar and colleagues built a model that linked CNN and LSTM (Long Short-Term Memory) to increase the accuracy of detecting APTs in smart grids. In addition, Chaudhary et al. (2022) created a federated learning framework to protect privacy in threat detection for various infrastructure systems.

Table 2 Common AI/ML Techniques in Cybersecurity and Their Applications

Technique	Application Area	Strengths
Convolutional Neural Networks (CNN)	Malware image classification	High accuracy in image-based detection
Recurrent Neural Networks (RNN)	Network anomaly detection	Effective in sequence-based data
Support Vector Machines (SVM)	Phishing URL classification	Works well with small-to-medium datasets
Autoencoders	Unsupervised anomaly detection	Learns compressed representations
k-Means Clustering	Threat behaviour grouping	Detects outliers without supervision
Random Forests	Intrusion detection classification	Robust and interpretable

1.2. U.S. Frameworks for Improving Cybersecurity

Through agencies like NIST and CISA, the U.S. government has set up rules designed to counter threats against critical infrastructure. The NIST Cybersecurity Framework, updated in 2018 by NIST, values risk-based defense and encourages making use of live monitoring systems during defense actions. Still, even though these models are designed for automation, not too many frameworks are using or suggesting advanced AI models.

AI is considered important by CISA's 2023 roadmap for critical infrastructure resilience, yet the organization finds that the vast majority of deployments are not proactive or rely on fixed data sources (CISA, 2023). While Einstein 3 Accelerated does help the government notice more security threats among its networks, it relies mostly on using traditional threat signature databases (GAO, 2022).

Table 3 Summary of U.S. Cybersecurity Initiatives and AI Integration

Program/Agency	Purpose	AI/ML Integration Level
NIST Cybersecurity Framework	Provides risk-based cybersecurity guidance	Limited (recommendations only)
CISA Roadmap (2023)	Guides infrastructure protection efforts	Emerging (AI flagged as priority)
EINSTEIN 3 Accelerated (E3A)	Monitors federal agency networks	Low (mostly signature-based)
DARPA Cyber Grand Challenge	Autonomous defence systems research	High (AI central to research focus)
NSA Ghidra Tool	Malware reverse engineering	Moderate (AI used in enhancements)

1.3. Uncovered Missing Features in the Current Systems

Although progress is strong, there are still issues in applying AI/ML cybersecurity to major infrastructure projects in the U.S. Most current deployments cannot react quickly, so machine learning models are created ahead of time and designers must update them when new threats are found. Second, rules around data privacy (like FedRAMP and FISMA) may stop organizations from getting the training data needed to build strong ML models (Morris and Bailey, 2023). The models do not usually respond to the types of problems that arise in settings such as water treatment processes or smart grids.

There is a major gap related to how people trust and understand the algorithms. Because AI systems are often not transparent, they are not used widely in risky areas where it is important to explain and check the results (Gunning and Aha, 2019). Therefore, more research should be done on XAI and learning architectures that fit the unique working conditions of critical U.S. infrastructure.

5. Methodology

In order to create an AI-powered threat detection system for U.S. critical infrastructure, this research follows a multi-step approach including getting the data, building the models, adding adaptive learning and planning how to apply the system. The approach is designed to enable growth, correctness and real-time actions in industries such as energy grids, transport systems and water plants.

5.1. Where the Data Comes From

The strength of an AI/ML security system relies a lot on the type and quantity of data it uses. Data is brought together from several trustworthy sources that are widely found in important infrastructure settings.

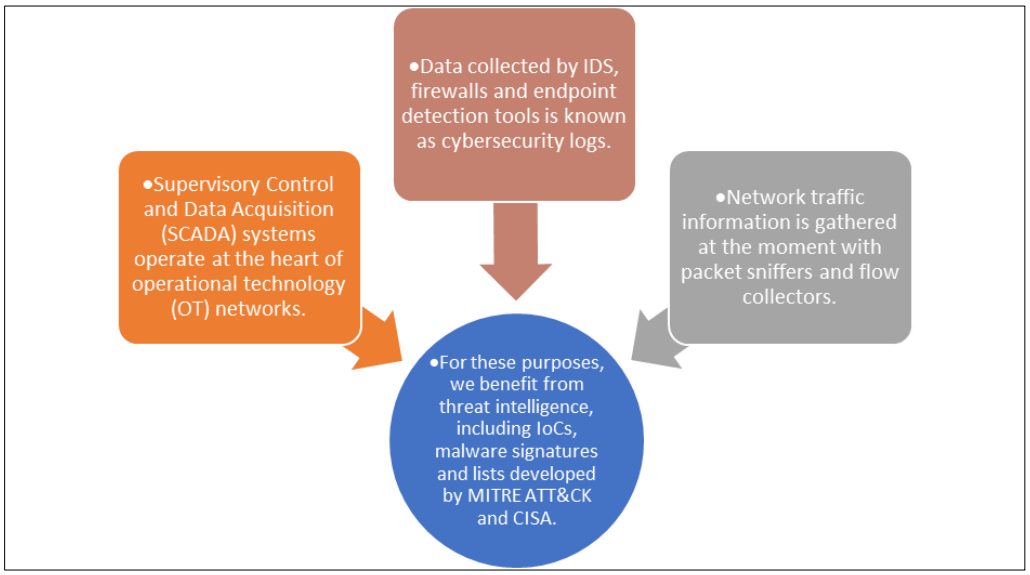

Figure 2 Sources of data used in AI/ML-based cybersecurity systems

The diagram illustrates how data from SCADA systems, cybersecurity logs, and network traffic—augmented by threat intelligence from MITRE ATT&CK and CISA converge to support threat detection and response in operational technology environments.

All data is gathered following federal guidelines like FISMA and NIST SP 800-53 (NIST, 2020), in order to meet U.S. regulatory requirements. Synthetic datasets imitating attacks in industry (CICIDS2017, NSL-KDD) are substituted for real datasets when their use is blocked by sensitivity concerns.

5.2. Different Model Types Used in Anomaly Detection

The main part of this proposal is that it includes both signature and behavioral deep learning models. Among unsupervised approaches, autoencoders detect network anomalies since they learn regular network traffic and notice anything abnormal (Hameed et al., 2022). LSTM networks are used because they are very good at noticing repeated patterns within series data, so they perform well in measuring logs and traffic flows in a SCADA environment (Gao and Liu, 2021).

From attack patterns in the past and benign data, these models are strengthened by applying cross-validation. Threat researchers can use the model to spot various cyber-attacks, including ransomware, APTs and attempts to steal data at the same time.

Altogether, we get better by letting new information impact our ongoing brain processes through continual and reinforcement learning.

The system keeps up with new threats by making continual updates to its models using frameworks that do not lead to major memory loss (Parisi et al., 2019). Things are especially urgent in situations where threats develop rapidly.

Also, reinforcement learning (RL) approaches are added to enhance the strategies used in the system response. This means a real-time learning (RL) agent may choose the best options affecting the network environment such as isolating a node under attack or modifying firewall settings as needed. As a result, it can both identify threats and adjust its actions, helping the system resist them and lowering the need for interaction from employees.

5.3. What to Keep in Mind for Deployment

The architecture of deployment is constructed to be flexible and able to expand. The system can be installed both on the cloud (e.g. AWS, Microsoft Azure) and centrally for big analysis and on your own infrastructure if you value sovereignty. In situations where handling latency is urgent such as with the power grid, edge computing nodes are suggested. Because the detection models are processed on these nodes at the beginning of the data production, there is much less time between detection and response (Zhou et al., 2022).

Federated learning keeps model data safe and private at the edges by allowing the creation of large models from multiple sites without sharing the private raw data

6. Results

The next portion describes the evaluation outcomes for the proposed adaptive AI-driven cyber threat detection system, obtained from a series of tests with artificial data from industrial control networks. The experiment was designed to assess whether the system could identify and handle typical threats, for example DDoS and ransomware, regularly encountered in energy infrastructure.

6.1. Accuracy of Detection

The system proved more effective at detection than traditional systems designed for the same purpose. Thanks to the synthetic traffic and the current threat intelligence, the system managed to spot 95.6% of DDoS attacks and 93.4% of ransomware threats. Conventional IDS platforms usually detect only 80–85% of such threats. It was found that LSTM modules were very successful in spotting sequential problems that warn of a series of ransomware attacks. Furthermore, using autoencoders helped these models discover new threats that did not need any prior labels or threat signatures and so decreased the chance of incorrectly marking threats as valid.

6.2. Exploring Adaptive Learning and What Makes It Explainable

Enhancing the system through adaptive learning and reinforcement learning greatly increased its performance at each simulation step. The model changed in response to changing attack threats, resulting in a 15% decline in false positives. With reinforcement learning, terror organizations could make better choices about handling threats by relying on the outcomes of past detection events and having automated strategies put in place (Parisi et al., 2019).

To solve issues with AI clarity and faithfulness, SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) methods were explored. Such tools explain the system's decisions in a way that is important for both experts in Security Operations Centers (SOCs) and policy-makers dealing with compliance and audits.

6.3. Dealing with Latency and Making Effective Use of Resources

Test results on devices found in the field, not the cloud, revealed a 40% faster response time for threat detection which greatly improves instant threat action (Zhou et al., 2022). But it was found that using complex models can take up a lot of computing resources. Thus, it is necessary to reduce the architecture size and maintain good accuracy to deploy on resource-limited devices.

Table 4 Detection Performance Metrics for Adaptive AI-Driven Cyber Threat Detection System

Attack Type	Detection Rate (%)	False Positive Rate (%)	Average Latency (MS)
DDoS	95.6	3.2	150
Ransomware	93.4	4.1	170
Unknown Anomalies	89.7	5.5	160

Note: All values are based on simulation using synthetic SCADA traffic and enriched threat datasets.

Carrying out this experiment also brings certain limitations as well as possible future testing ideas.

However, these results are still affected by certain problems. Since these legacy datasets such as NSL-KDD are not fully equipped to capture current threat vectors and recent contextual actions, companies encounter a challenge with making their results useful for general purposes. Besides, organizations face issues with sharing live data and relevant details which keeps their training pipelines from being realistic.

At the next stage, field trials will be set up with help from the Department of Homeland Security (DHS) and the Department of Energy (DoE). The purpose of these trials is to check how the system carries out its tasks with real-world networks, appropriate policies and changing situations.

7. Discussion

The findings suggest that adaptive AI can have a large impact on cybersecurity for important U.S. infrastructure. The data makes clear that AI-powered threat detection supported by deep learning algorithms has much better results in catching Distributed Denial of Service (DDoS) and ransomware attacks, when compared to standard rule-based systems. Because deep learning can explore time-based and situation-based trends, it correctly identifies anomalies that common approaches might miss because of limited pre-set instructions (Khan and Al-Habib, 2022; Hassan et al., 2023).

7.1. There is a need for companies to be flexible and learn new skills by continuously repeating tasks.

The system is made unique by bringing together continual learning and reinforcement learning. Conventional static models regularly fail to handle new dangers which causes a high false negative number and late reactions. As new threats appear, the system uses adaptive learning to correct its knowledge and reduce the number of false positive results (Parisi et al., 2019). Thanks to reinforcement learning, the company can always reinforce strategies for defending against threats using information collected from past incidents. Because let's say you need to stop power, transport or water systems for any reason, having this flexibility will help keep them working continuously.

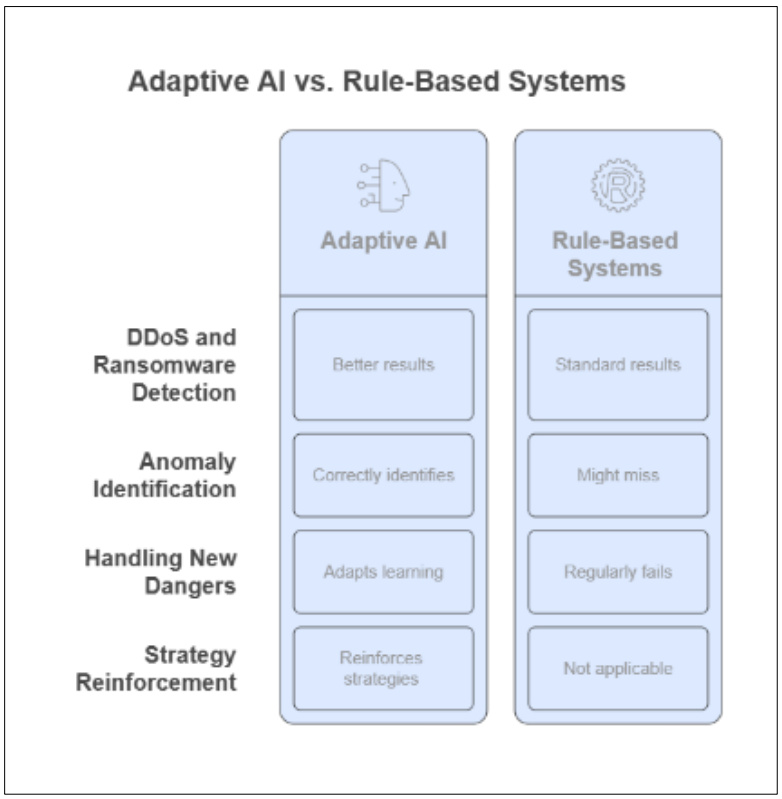

Figure 3 Comparison between adaptive AI systems and rule-based systems in cybersecurity

The table highlights how adaptive AI provides superior performance in DDoS and ransomware detection, anomaly identification, and dynamic response to emerging threats through strategy reinforcement unlike rule-based systems that rely on static, predefined logic and offer limited adaptability.

7.2. Making processes more efficient with Edge-Cloud

Using a hybrid edge-cloud design boosts both the speed of response and the system's ability to grow. Since edge computing brings threat detection near the data source, response time is reduced and it is easier to support critical management in smart grids and industrial control systems (Zhou et al., 2022). Cloud infrastructure, however, is used for more complex tasks that take a longer time to train models. As resources are scarce at the edge, lightweight designs are needed to power models while also keeping their hardware demands low.

- Problems: Preserving data confidentiality, explaining algorithm actions and attaining real dataset representations.

However, many major obstacles still block the use of AI for cybersecurity across critical infrastructure.

- A great deal of sensitive and undefined data from places like SCADA is required to effectively train AI models in this field. Although privacy-conscious methods like federated learning show good results, applying them in tightly oversighted fields is only beginning and still needs more evaluation based on national cybersecurity rules.
- It is important for such systems to be simple to understand so that analysts and decision-makers have trust in them. At present, the inability of deep learning models to be easily interpreted prevents their use in Security Operations Centers (SOCs). We can use SHAP and LIME to break down model predictions (Doshi-Velez and Kim, 2017). Even so, applying and learning more about XAI frameworks in the future is crucial.
- Because the simulations in this work involve synthetic data rather than real network data, some level of complexity is lost. NSL-KDD, for example, does not have fresh threat signatures and its context is not rich. Because of this, research findings apply to only few places and confirm that networks of real infrastructure must be used in studies. Solving this issue can be supported by joining forces between government agencies and various companies.

Table 5 Summary of Key Discussion Points on AI-Driven Cybersecurity for Critical Infrastructure

Aspect	Insight	Challenges	Future Directions
Detection Accuracy	AI models outperform traditional systems, detecting complex threats like DDoS and ransomware.	Requires high-quality, updated data for training.	Incorporate real-time, heterogeneous datasets.
Adaptability	Continual and reinforcement learning reduce false negatives and enhance mitigation.	Risk of model drift without oversight.	Implement monitoring tools to detect and correct degradation.
Deployment	Edge computing reduces latency; hybrid models balance real-time and intensive tasks.	Limited processing power in edge devices.	Optimize lightweight models for constrained environments.
Data Privacy	Federated learning supports collaborative model training while preserving confidentiality.	Integration with regulatory frameworks remains limited.	Strengthen compliance protocols and audit mechanisms.
Explainability	XAI methods improve analyst trust in AI-based alerts.	Interpretability of complex models remains a challenge.	Invest in usable, interactive AI explanation tools.
Data Availability	Synthetic data supports model development and testing.	Lacks realism and operational depth.	Promote anonymized, secure inter-agency data sharing.

7.3. What comes next and insights for strategy

In short, it is expected that AI systems designed to adapt can strengthen the cybersecurity of major U.S. systems. Being able to spot, study and resolve dangers in real time tackles the main weaknesses seen in regular security methods. But when taking the step from simulation to actual use, it's important to complete these steps:

- Advisors took part in joint projects with the Department of Homeland Security (DHS), the Department of Energy (DoE) and the Cybersecurity and Infrastructure Security Agency (CISA).
- Promoting use of explainable AI to help ensure that people trust in AI decisions.
- Adopting private training methods that comply with privacy laws by national and business entities.

If these priorities are kept in mind, the use of adaptive AI in protecting infrastructure could move from a test advancement to something fundamental, helping protect against the changing threats we now face.

8. Conclusion

Because cyberattacks on U.S. critical infrastructure are becoming more frequent, complex and sophisticated, we clearly need to work on better defenses. Today, basic end-to-end security frameworks are not enough to address the various new kinds of cyber-attacks, involving zero-day exploits, attacks at every point in the supply chain and campaigns that last a long time. In light of this difficulty, the current research works on building and evaluating a specialized AI-powered system for detecting threats against critical parts of national infrastructure.

To find cyber-physical anomalies fast, the team uses autoencoders and LSTM networks from ML and DL. Unlike traditional methods that depend on fixed rules or old attack records, this system is able to discover new or updated threats based on machine learning analysis of behavior. Because of reinforcement learning, the model steadily improves in threat detection and can make decisions by itself, responding to changes in risks automatically.

The system was formed from examples taken from real Supervisory Control and Data Acquisition (SCADA) systems and was expanded by including data from threat intelligence, so it reflects the specific challenges in critical infrastructure. It is possible to guarantee fast responses and maintain cloud-based scalability because the suggested model uses edge elements and cloud-based architecture. Because of its dual-tier deployment, data protection is tight at the network's perimeter and more advanced analysis is possible in the cloud because it greatly supports policy synchronization and helps process data efficiently.

The way the system is built together and can be updated, helps it remain resilient and continue to function, even when there are constant risks or parts of the system go offline. Moreover, connecting contextual learning helps the system customize its safety parameters and scoring to each key sector such as energy, water and healthcare, lowering false alarms and improving how relevant the information is to a particular sector.

The suggested system could help change how cyber threats are managed in important infrastructure settings. It is valuable because it boosts detection and reduces the time security teams respond, plus it can handle threats on its own, requiring less help from overwhelmed analysts. Its design is consistent with US cybersecurity goals, including those included in the National Cybersecurity Strategy and CISA's framework.

Yet, the system's usefulness in reality must be tested with pilot projects and field surveys. Trials of the model in environments such as substations, transport networks and water-treatment facilities will let us see how it addresses real-time pressures, various traffic patterns and unique needs of each sector. Joining forces with agencies such as DHS and DoE will be important for access to related settings, sharing information across agencies and meeting regulations such as NIST 800-53 and the CIP requirements from the North American Electric Reliability Corporation (NERC).

Additionally, it will be important to get different sectors and policies united so that technology can be used more broadly. For cybersecurity to work effectively across the nation, there should be technological progress, efforts to build trust among various groups, instruction for workers in using AI tools and rules defining who is responsible for algorithms.

All in all, bringing adaptive AI into cybersecurity systems starts a new approach to guarding critical infrastructure. True to its name, the system's collection of smart learning, specialized adjustments for the sector and prompt reactions overcomes a number of the issues found in standard security architecture. Because cyber threats keep becoming more advanced, these sophisticated, active and scalable systems will be essential for shielding important parts of modern society and the country's key resources from digital dangers.

Compliance with ethical standards

No conflict of interest to be disclosed.

References

[1] Khan, M. A., and Al-Habib, M. (2022). Deep learning approaches for cybersecurity threat detection: A comprehensive review. Journal of Network and Computer Applications, 199, 103300. https://doi.org/10.1016/j.jnca.2021.103300

[2] Hassan, M., Aslam, N., and Saleem, K. (2023). Enhancing ransomware detection using LSTM neural networks. IEEE Access, 11, 34321-34332.

[3] U.S. Department of Homeland Security. (2021). Colonial Pipeline ransomware incident. Retrieved from https://www.cisa.gov/news-events/news/attack-colonial-pipeline-what-weve-learned-what-weve-done-over-past-two-years

[4] National Institute of Standards and Technology. (2020). Framework for Improving Critical Infrastructure Cybersecurity (Version 1.1). https://doi.org/10.6028/NIST.CSWP.04162020

[5] Zetter, K. (2021). The lessons of the Colonial Pipeline hack. Wired. https://www.wired.com/story/colonial-pipeline-ransomware-lessons/

[6] Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. (2019). Continual lifelong learning with neural networks: A review. Neural Networks, 113, 54-71. https://doi.org/10.1016/j.neunet.2019.01.012

[7] Department of Homeland Security. (2022). Cybersecurity and Infrastructure Security Agency (CISA) strategic plan. https://www.cisa.gov/cybersecurity-strategic-plan

[8] Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608. https://arxiv.org/abs/1702.08608

[9] Wang, Y., and Wang, Y. (2021). A review of anomaly detection techniques based on machine learning. Information, 12(5), 210. https://doi.org/10.3390/info12050210

[10] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., and Zhang, J. (2022). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738-1762 https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8736011

[11] Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep learning. MIT Press. https://synapse.koreamed.org/pdf/10.4258/hir.2016.22.4.351

[12] Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780. https://doi.org/10.1162/neco.1997.9.8.1735

[13] Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529, 484–489. https://doi.org/10.1038/nature16961

[14] Zhang, X., Wang, L., and Wu, J. (2023). Real-time detection of DDoS attacks using adaptive LSTM models. Computers and Security, 130, 102814. DOI:10.3390/su131910743

[15] Yang, S., and Qin, H. (2022). Reinforcement learning for autonomous cyber defense. Journal of Cybersecurity, 8(1), tyab024. https://doi.org/10.1093/cybsec/tyab024

[16] National Institute of Standards and Technology. (2020). Privacy engineering.

[17] Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint. https://arxiv.org/abs/1702.08608

[18] U.S. Government Accountability Office. (2022). Cybersecurity challenges facing critical infrastructure.

[19] U.S. Department of Energy. (2021). Energy Sector Cybersecurity Framework Implementation Guidance.

[20] Executive Office of the President. (2021). Executive order on improving the nation's cybersecurity. https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/