

Predictive Response Engine for Dynamic Intelligence and Cyber Threat Mitigation (PRE-DICT): Architecture, Mathematical Formulation, and Empirical Evaluation

S.Karthikeyan¹, J.Pradeep Rajalu², M..Ganeshraj³, S.Abdul Rahuman⁴, V.Prem kumar⁵

¹ Assistant Professor, Department of Artificial Intelligence and Data Science

²³⁴⁵ Under graduate student, Department of Artificial Intelligence and Data Science

¹technokarthi@gmail.com, ²pradeeprajalu5@gmail.com, ³ganeshraj.1418@gmail.com, ⁴abdurahuman.ar2417@gmail.com,
⁵vpremkumar508@gmail.com

S

Abstract— Modern enterprise infrastructures face threats from complex multi-stage cyber attacks carried out by smart and evolving adversaries. Traditional tools like Security Information and Event Management (SIEM) and Security Orchestration, Automation and Response (SOAR) operate in a way. They depend on Indicators of Compromise (IoCs) and fixed response plans. This approach leads to delays in action and a lot of alarms especially during zero-day attacks. This paper presents a system called the Predictive Response Engine for Dynamic Intelligence and Cyber Threat Mitigation (PRE-DICT). It is an automated framework that predicts enemy moves and takes actions in real time. PRE-DICT combines a -source Dynamic Cyber Threat Intelligence (CTI) data pipeline, a Temporal Graph Neural Network (TGNN) for forecasting attack paths with probabilities and a Reinforcement Learning (RL) engine that makes safe risk-controlled decisions. The whole system models threat prediction, as a Constrained Partially Markov Decision Process (CPOMDP) which helps reduce threats while keeping service disruptions to a minimum. In tests using a cyber-range with 1,200 systems PRE-DICT cuts the average time to respond by 78.4% and stops 94.2% of advanced attack chains. It also keeps business operations running.

Keywords—component; formatting; style; styling; insert (key words) *Autonomous cyber defense, cyber threat intelligence, temporal graph neural networks, reinforcement learning, incident response, network security, risk-bounded optimization.*

I. INTRODUCTION

I notice that the modern cyber threat landscape is shifting toward automated distributed attack vectors that rely on evasion techniques and zero-day exploits. In Security Operations Center (SOC) settings defensive operations remain mainly reactive. Telemetry pipelines only fire alerts, after payloads unauthorized lateral movements or anomalous signature matches appear on host or network sensors. When human security analysts finally read telemetry, triage alerts and run response playbooks the adversaries have usually already completed internal discovery, credential acquisition and critical data exfiltration. The modern cyber threat landscape keeps evolving and automated distributed attack vectors keep growing with sophisticated evasion techniques and zero-day exploits. Telemetry pipelines and alerts still serve as the line of

defense while human security analysts work hard to triage before adversaries finish internal discovery, credential acquisition and critical data exfiltration.

, However, the existing tools used by the traditional Cyber Threat Intelligence solutions such as STIX 2.1 and TAXII have made sharing the information on the threats much easier; the process of processing the information in the current systems has its issues. The issue in these systems is that they generate a number of indicators but don't demonstrate how these threats relate to the operations of an organization [3]. The insertion of indicators into the existing rule-based systems produces a huge number of useless alerts with a high level of false positive events. The existing solutions don't reveal how the threat evolves and relates to the organization's infrastructure [4]. Monitoring threat patterns is just insufficient in modern defensive systems that must predict the behavior of the attack.

This study aims to close the gap between identifying threats and dealing with those threats by introducing the PRE-Predictive Response Engine, for Dynamic Intelligence and Cyber Threat Mitigation) framework. The things this research adds are as follows:

• Dynamic Multi-Source Telemetry

Fusion: We create a streaming system that connects outside threat intelligence data with our device monitoring and response tools. We also bring in network flow data and user access information. All of this data is put together into a kind of graph that shows connections and features.

• Temporal Graph-Based Attack

Forecasting: The problem of lateral movement and privilege escalation is formulated as a dynamic prediction of node-edge sequences for heterogeneous IT networks by utilizing a TGNN [5].

• Risk-Bounded Autonomous

Orchestration: Proactive defense is represented as a Constrained Partially Markov Decision Process (CPOMDP). Proactive defense uses relaxed policy gradients to reduce threat while imposing a strict limit, on collateral disruption of services.

• **Empirical Validation and Cyber-Range Benchmarking:** We do tests, on a realistic fake cyber-range that has 1,200 nodes. We run these tests across 150 simulated persistent threat (APT) campaigns

II. RELATED WORK

A. CYBER THREAT INTELLIGENCE AND CORRELATION

CTI systems automatically ingest IoC and TTP data stored in threat intelligence libraries like MISP and OpenCTI [7]. Classical correlation techniques still heavily rely on string-matching or basic heuristic approaches that create substantial analytical and alerting bottlenecks [3]. Current research efforts apply natural language processing and semantic knowledge graphs to CTI enrichment tasks. These solutions remain largely disconnected from the run-time internal telemetry and network topology information [8].

B. GRAPH REPRESENTATION LEARNING IN INTRUSION DETECTION

The Graph Neural Networks have become very popular for finding things in network setups because they can handle data that doesn't fit in a regular space [9]. The Static Graph Convolutional Networks and Graph Attention Networks are good at finding patterns inside parts of a network [10]. The problem is that attacks on computers happen in a sequence and come in groups over time. Static methods can't follow how the network changes over time which shows the need, for Temporal Graph Neural Networks.

C. AUTONOMOUS DECISION-MAKING IN CYBER DEFENSE

Incident response systems have evolved from SOAR playbooks to DRL agents [12][13]. I noticed that DRL agents can always figure out what strategy needs to be followed in tough cyber attacks. DRL agents without boundaries might even select an extremely risky strategy, like shutting down all core switches or blocking all the domains of attackers, just to stop the attackers [14]. Therefore, Safe RL algorithms, which use relaxation techniques or Control Barrier Function, are required [6] [15].

III. SYSTEM ARCHITECTURE AND THREAT MODEL

A. THREAT MODEL

We model a goal-oriented threat actor who carries out multi-stage intrusions that follow the MITRE ATT&CK enterprise matrix [16]. This adversary focuses on operational assets, such as domain controllers, internal financial databases and sensitive IP vaults by moving through four clear stages. First the adversary gains ingress through a perimeter exploit or a spear-phishing attack. Second the adversary conducts reconnaissance and harvests credentials by dumping LSASS data or impersonating tokens. Third the adversary moves laterally across subnets. Fourth the adversary delivers impact either by deploying ransomware or by exfiltrating data. The adversary may also use zero-day exploits or native administrative binaries—Living-off-the-Land (LotL)—to avoid detection, by static signatures [17].

B. ARCHITECTURE OVERVIEW & BLOCK DIAGRAM

PRE-DICT is structured across three asynchronous, decoupled tiers operating over streaming message brokers, as illustrated in the system block diagram below:




Fig. 1. Architectural Block Diagram of the PRE-DICT Framework illustrating data flow across Ingestion, TGNN Forecasting, and Constrained RL Orchestration.

- **Tier 1 (Dynamic Ingestion & Graph Fusion):** Tier 1 brings in CTI feeds (STIX 2.1/TAXII) and internal telemetry (Sysmon EDR traces, NetFlow Active Directory auth logs). Tier 1 then normalises events and puts them together into an attributed property graph $G_t = (V_t, E_t, X_v(t), X_e(t))$.
- **Tier 2 (Temporal Graph Forecasting Engine)** looks at snapshot sequences. Tier 2 (Temporal Graph Forecasting Engine) uses graph attention together with Gated Recurrent Units (GRUs). Tier 2 (Temporal Graph Forecasting Engine) estimates how likely each node and edge will be compromised in the k steps.
- **Tier 3 (Risk-Bounded Autonomous Response Engine):** This system uses a reinforcement learning agent with a policy optimization algorithm. The agent selects actions that are based on risk levels. These actions include things, like micro-segmentation and token revocation. The actions are designed to isolate threats. They do not impact the operation of enterprise applications. The response stays within boundaries. The system learns over time to make decisions. It balances security and business continuity.

IV. MATHEMATICAL AND ALGORITHMIC FORMULATION

A. TEMPORAL GRAPH NEURAL NETWORK FORMULATION

Let G_t be the attributed enterprise graph at time t . For each node i , in V_t that has a feature vector $x_i(t)$ in R^{d_v} spatial graph attention gathers neighborhood embeddings at layer l :

$$h_i^{(l)}(t) = \sigma(W_0 h_i^{(l-1)}(t) + \sum_{j \in N(i)} \alpha_{ij} W_1 h_j^{(l-1)}(t))$$

Here $N(i)$ means the local neighborhood of node i , α_{ij} is the normalized attention coefficient that is computed using a

LeakyReLU attention function and W_0 and W_1 are projection matrices that can be learned.

Temporal transitions are captured across snapshots using a recurrent gating unit:

$$z_i(t) = \text{GRU}(h_i^L(t) z_i(t-1))$$

The output prediction calculates the probability that a node or an edge will be compromised in the k hops using the equation: $P(v_{t+k} \in e_{t+k} | G_{\leq t}) = \text{Softmax}(W_p z_i(t))$.

B. CONSTRAINED POMDP RESPONSE PIPELINE FLOWCHART

The way a system makes decisions, on its own is described using something called a Constrained Partially Markov Decision Process. This is shortened to CPOMDP. It is made up of parts which are listed as S, A, O, T, R, C and γ . The way this system works is shown in the decision flowchart that comes next.:




Fig. 2. Operational Decision Flowchart of PRE-DICT demonstrating risk evaluation and safe policy enforcement via Lagrangian constraints.

The parameterized policy π_θ aims to maximize the expected reward. At the time it keeps the operational disruption cost below a given threshold κ . This is written as:

$$\text{maximize } \theta \text{ of } E_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \gamma^t R(s_t a_t) \right]$$

$$\text{subject to } E_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \gamma^t C(s_t a_t) \right] \leq \kappa$$

To solve this constrained optimization we use relaxation. The Lagrangian function $L(\theta, \lambda)$ combines the reward objective and the cost constraint using a λ . It is defined as:

$$L(\theta, \lambda) = E_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \gamma^t R(s_t a_t) \right] - \lambda \left(E_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \gamma^t C(s_t a_t) \right] - \kappa \right)$$

Here λ is a -negative scalar that controls how strictly the cost constraint is enforced. We update λ using gradient ascent. The update rule is:

$$\lambda \leftarrow \max(0, \lambda + \eta_\lambda (J_C(\pi_\theta) - \kappa))$$

This means we increase λ if the average cost under policy π_θ exceeds the threshold κ and clamp it to zero if it goes negative. The step size, for λ is η_λ .

V. EXPERIMENTAL SETUP AND METHODOLOGY

A. NETWORK TOPOLOGY AND SEGMENTATION

The evaluation environment was set up as a cyber-range simulation. It had N equals 1,200 nodes. These nodes were split into five security zones. Each zone had its specific operational criticality weight. The weight for each zone was, between 0.1 and 1.0

ID	Zone	N	W	Services	Policy
Z1	DMZ	50	.3	NGINX/HAProxy/VPN/SSH	80/443/1194
Z2	Workstations	800	.2	Win/macOS/SMB/DHCP	802.1X
Z3	Apps	250	.7	PG/Redis/Kafka/K8s	API-Z2
Z4	Identity	20	1.0	AD/Kerb/LDAP	RPC/Kerb
Z5	Data/OT	80	.9	SAN/Backup/SCADA	Micro-seg.

TABLE I: NETWORK SEGMENT ALLOCATIONS, NODE DISTRIBUTION, AND CRITICALITY WEIGHTS

C. TOPOLOGICAL DOT DIAGRAM AND THREAT TRAJECTORY

To show how lateral spread happens in stages, across structural network zones and to explain how PRE-DICT places proactive mitigation barriers the topological dot diagram can be seen in Figure 3:

Topological Dot Diagram: Multi-Stage Attack Trajectory vs. PRE-DICT




Fig. 3. Topological Dot Diagram showing simulated adversary traversal across zones Z-1 through Z-3, TGNN predicted trajectory into Z-4, and preemptive barrier placement by PRE-DICT.

D. TELEMETRY SYNTHESIS & EVALUATION METRICS

The simulated telemetry generator ran for thirty days. It produced one point eight times ten to the power of seven events. Background user activity, which was considered normal was created using log- inter-arrival times for sessions. The parameters for this were μ to one point two and σ equal, to zero point four. These values came from the CIC-IDS2017 dataset [18]. The DARPA OpTC dataset [19].

Adversary campaigns were generated by automating the emulation of APT29, APT33, FIN7 and Cobalt Strike tactics. There was a zero-day injection probability set at point one five. This means that in every one hundred attempts about fifteen would include a zero-day exploit.

- Mean Time To React (MTTR): time it takes to carry out a countermeasure after the first sign of an intrusion. The formula is $MTTR = (1/M) \sum_{m=1}^M (t_{\text{mitigate}}^{(m)} + T_{\text{ingress}}^{(m)})$.
- Mitigation Success Ratio (MSR): The percentage of attack campaigns that are stopped before they reach infrastructure at Z-4 or Z-5.
- Service Disruption Cost Index (SDCI): A measure that reflects the cost of service disruptions, over time calculated as $SDCI = (1/T$

V. EMPIRICAL RESULTS AND COMPARATIVE ANALYSIS

PRE-DICT was benchmarked against three baseline paradigms under identical cyber-range conditions across 150 independent attack campaigns:

Defense Framework	Mean Time to Detect (s)	Mean Time to Respond (s)	Mitigation Success (%)	Disruption Cost (SDCI)
Static SIEM/SOAR	412 s	340 s	58.3%	6.8
Anomaly-Based IDS	195 s	210 s	69.1%	8.2
Reactive Human Triage	1,840 s	1,420 s	72.4%	2.1
PRE-DICT (Proposed)	28 s	45 s	94.2%	3.4

TABLE II: COMPARATIVE PERFORMANCE EVALUATION ACROSS DEFENSIVE PARADIGMS (150 ATTACK CAMPAIGNS)

A. LATENCY REDUCTION & CONTAINMENT EFFICACY

As seen in Table II PRE-DICT reduced the duration by 78.4 percent. It went from 340 seconds using SOAR down, to 45 seconds. By doing micro-segmentation and revoking access tokens before lateral movement finished Tier 3 used k-hop trajectories. This happened before the adversary gained any advantage.

B. OPERATIONAL DISRUPTION VS. SECURITY TRADE-OFFS

Anomaly-based ids still had an sdc i of 8.2 even though anomaly-based ids achieved a containment rate of 69.1 percent. this happened because anomaly-based ids performed segregation of the network. in contrast pre-dict was able to keep its sdc i at 3.4. the improvement comes from the fact that the lagrangian multiplier λ penalised any service disruption forcing the agent to take actions.

VI. DISCUSSION AND THREATS TO VALIDITY

- Construct validity (restrictions on synthetic data): we built traffic patterns using a log distribution drawing from known datasets.. real business data tends to be much more chaotic. there are tails, in the data because of shadow it and misconfigurations. the noise level we used—4.5 percent—in our range might not provide enough friction. this could lead to false positives than we expect.
- Internal validity (sensor observability): this test assumes that all zones have data available.. in businesses using older ot equipment, unauthorized iot devices and end-to-end encrypted tunnels can block access, to data. when parts of g_t are missing the k hop prediction accuracy of the tgnn suffers because the model cannot see the picture.

- External validity (living-off-the-land exploits): when attackers use existing tools such as windows management instrumentation (wmi) or powershell during normal working hours they can hide their actions while the network infrastructure is being reconfigured to look legitimate. to stop living-off-the-land attacks it is essential to record process history, for each node.[20][17]

VII. CONCLUSION AND FUTURE DIRECTIONS

In addition, we propose a system called PRE-DICT for protection against cyber attacks. The system leverages real-time threat intelligence and employs Temporal Graph Neural Networks (TGNNs) to model the temporal evolution of threats. It further integrates risk-constrained reinforcement learning to enable decision-making under low risk. Evaluated on a network of 1,200 nodes, PRE-DICT reduced the mean time to respond by 78.4 percent and achieved a mitigation success ratio of 94.2 percent, without degrading system performance or availability. Future work will focus on graph learning in large enterprise environments.

REFERENCES

- [1] M. Husák, J. Komárková, E. Bou-Harb, and P. Čeleda, "Survey of attack projection, prediction, and forecasting in cyber security," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1760-1783, 2019.
- [2] A. Al-Sallami and M. T. Raza, "Collaborative automated incident response in enterprise security operations centers," *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4102-4115, 2021.
- [3] A. Dimitriadis, C. Prassas, J. L. Flores, B. Kulvatunyou, N. Ivezic, D. A. Gritzalis, and I. K. Mavridis, "Contextualized filtering for shared cyber threat information," *Sensors*, vol. 21, no. 14, p. 4890, 2021.
- [4] C. Sauerwein, C. Sillaber, M. Musmann, and R. Brey, "Threat intelligence sharing platforms: An exploratory study of software vendors and research prototypes," *Computers & Security*, vol. 103, p. 102174, 2021.
- [5] Y. Zhou, H. Peng, J. Li, Y. He, and L. Wang, "Dynamic graph neural networks for multi-stage cyber attack path forecasting," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 2894-2907, 2022.
- [6] C. Tessler, D. J. Mankowitz, and S. Mannor, "Reward constrained policy optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2019, pp. 1-14.
- [7] T. D. Wagner, K. Mahbub, E. Palomar, and A. E. Abdallah, "Cyber threat intelligence sharing: Survey and research directions," *Computers & Security*, vol. 87, p. 101589, 2019.
- [8] X. Sun, J. Dai, P. Liu, A. Singhal, and J. Yen, "Using cyber threat intelligence for automated cybersecurity knowledge graph construction," *IEEE Transactions on Dependable and Secure Computing*, vol. 20, no. 3, pp. 2314-2329, 2023.
- [9] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4-24, 2021.
- [10] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," in *Proc. International Conference on Learning Representations (ICLR)*, 2018, pp. 1-12.
- [11] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, "Temporal graph networks for deep learning on dynamic graphs," in *ICML Workshop on Graph Representation Learning*, 2020, pp. 1-12.
- [12] T. T. Nguyen and V. J. Reddi, "Deep reinforcement learning for cyber security: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1394-1426, 2021.
- [13] M. Standen, M. Lucas, D. Wardman, J. Bowman, and D. Kim, "Cyborg: A gym for the development of autonomous cyber agents," in *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2021, pp. 1677-1679.
- [14] A. Ridley, "Machine learning for autonomous cyber defense," *The Next Wave*, vol. 22, no. 1, pp. 7-14, 2018.
- [15] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. International Conference on Machine Learning (ICML)*, 2017, pp. 22-31.
- [16] B. E. Strom et al., "MITRE ATT&CK: Design and philosophy," MITRE Technical Report, MTR180183, 2018.
- [17] W. U. Hassan, S. Guo, D. Li, Z. Chen, K. Jee, Z. Li, and A. Bates, "NoDoze: Combatting threat alert fatigue with automated provenance triage," in *Proc. Network and Distributed System Security Symposium (NDSS)*, 2019, pp. 1-15.
- [18] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proc. 4th International Conference on Information Systems Security and Privacy (ICISSP)*, 2018, pp. 108-116.
- [19] DARPA, "Transparent computing (TC) & Operationally Transparent Cyber (OpTC) data repository," Defense Advanced Research Projects Agency, Tech. Rep., 2020.
- [20] M. Ring, S. Wunderlich, D. Scheuring, D. Landes, and A. Hotho, "A survey of network-based intrusion detection data sets," *Computers & Security*, vol. 86, pp. 147-167, 2019.