GLOBAL JOURNAL
OF COMPUTER SCIENCE AND TECHNOLOGY - D: NEURAL & AI
Volume 26 | 2026 Article ID: GJCST-259806 Language: English

Recursive Stochastic Tensor Estimation for Cyber-Affected Traffic Distribution Analysis: A Non-Stationary Markov Chain Approach

Online First

Dr. Virendra Kumar Tiwari [§ *]
*Corresponding Author
Dr. Ashish Jain [§]
Rohit Singh [§]
Pramod Saket [§]
Sonal Sharma [§]
Ripusoodan Sharma [§]

[§] Department of Computer Application, Lakshmi Narain College of Technology, Bhopal, India

RECEIVED	ACCEPTED	ONLINE PUBLISHED	PEER REVIEW
2026-08-06	2026-08-17	2026-09-02	Double Blind

Abstract

We propose a Recursive Stochastic Tensor Estimation (RSTE) framework for modeling user behavior in cyber-affected traffic distribution systems. Conventional Markov chain models assume static transition probabilities, which fail to capture the dynamic nature of cyber threats. To address this limitation, we introduce a time-varying third-order transition tensor that evolves via an online stochastic approximation algorithm. The tensor encodes transition probabilities between behavioral states as functions of both time and a latent cyber-risk context vector, which is derived from real-time network anomaly signals. A Robbins-Monro procedure updates the tensor upon each observed user transition, with an adaptive damping factor that attenuates noise from false positive anomaly signals while preserving sensitivity to genuine cyber events. The update is focused on the context dimension most activated by the current threat environment, thereby enabling efficient learning under non-stationary conditions. The estimated tensor is then marginalized across context dimensions to produce a time-varying equilibrium distribution, which directly informs traffic distribution optimization and risk assessment. The anomaly signals themselves are generated by a Transformer-encoder pre-trained on the CICIDS2017 dataset and fine-tuned with a supervised contrastive loss. This encoder processes aggregated Net Flow records to produce a 16-dimensional anomaly vector. The RSTE framework integrates seamlessly with existing user classification and traffic logging modules, and its outputs feed into a utility function that balances bandwidth utilization, latency, and cyber risk costs. The primary contribution is a principled, mathematically grounded method for recursively estimating non-stationary Markov dynamics under cyber perturbations, without requiring explicit labeling of attack types. This work therefore provides a novel analytical tool for operators to dynamically adjust traffic distribution strategies in response to evolving cyber threats.

cybersecurity | deep learning | Internet Traffic Distribution | Markov Chain | Online Stochastic Approximation | Recursive Stochastic Tensor Estimation (RSTE) | Remove term: Transformer Encoder. Transformer Encoder

AI USE STATEMENT
No generative AI was used for analysis or results.

FUNDING
No external funding was declared for this work.

CONFLICT OF INTEREST
The authors declare no conflict of interest.

DATA AVAILABILITY
Not applicable for this article.

ETHICS
No ethics committee approval was required for this article type.

CONSENT
Not applicable for this article.

TRIAL REG.
Not applicable.

Crossref DOI: 10.34257/GJCSTD259806

How to Cite: Tiwari et al. (2026). Recursive Stochastic Tensor Estimation for Cyber-Affected Traffic Distribution Analysis: A Non-Stationary Markov Chain Approach. Global Journal of Computer Science and Technology. Ahead of Print. DOI: 10.34257/GJCSTD259806

AR Experience Web Page DOI
Print ISSN 0975-4350
Online ISSN 0975-4172
Under the strict compliance and defined process of

METADATA CONTINUATION

AUTHOR CONTACT QR LEDGER

Dr. Virendra Kumar Tiwari

Dr. Ashish Jain

Rohit Singh

Pramod Saket

Sonal Sharma

Ripusoodan Sharma

ARCHIVAL RECORD

GJCST · Vol 26 · Issue 1 · 2026

Article ID GJCST-259806 · DOI 10.34257/GJCSTD259806

Print ISSN 0975-4350 · Online ISSN 0975-4172

Global Journal of Computer Science and Technology - Section D
GJCST

Recursive Stochastic Tensor Estimation for Cyber-Affected Traffic Distribution Analysis: A Non-Stationary Markov Chain Approach

Dr. Virendra Kumar Tiwari[§*], Dr. Ashish Jain[§], Rohit Singh[§], Pramod Saket[§], Sonal Sharma[§], and Ripusoodan Sharma[§]

Affiliations
§ Department of Computer Application, Lakshmi Narain College of Technology, Bhopal, India

Abstract

We propose a Recursive Stochastic Tensor Estimation (RSTE) framework for modeling user behavior in cyber-affected traffic distribution systems. Conventional Markov chain models assume static transition probabilities, which fail to capture the dynamic nature of cyber threats. To address this limitation, we introduce a time-varying third-order transition tensor that evolves via an online stochastic approximation algorithm. The tensor encodes transition probabilities between behavioral states as functions of both time and a latent cyber-risk context vector, which is derived from real-time network anomaly signals. A Robbins-Monro procedure updates the tensor upon each observed user transition, with an adaptive damping factor that attenuates noise from false positive anomaly signals while preserving sensitivity to genuine cyber events. The update is focused on the context dimension most activated by the current threat environment, thereby enabling efficient learning under non-stationary conditions. The estimated tensor is then marginalized across context dimensions to produce a time-varying equilibrium distribution, which directly informs traffic distribution optimization and risk assessment. The anomaly signals themselves are generated by a Transformer-encoder pre-trained on the CICIDS2017 dataset and fine-tuned with a supervised contrastive loss. This encoder processes aggregated Net Flow records to produce a 16-dimensional anomaly vector. The RSTE framework integrates seamlessly with existing user classification and traffic logging modules, and its outputs feed into a utility function that balances bandwidth utilization, latency, and cyber risk costs. The primary contribution is a principled, mathematically grounded method for recursively estimating non-stationary Markov dynamics under cyber perturbations, without requiring explicit labeling of attack types. This work therefore provides a novel analytical tool for operators to dynamically adjust traffic distribution strategies in response to evolving cyber threats.

Keywords: *cybersecurity, deep learning, Internet Traffic Distribution, Markov Chain, Online Stochastic Approximation, Recursive Stochastic Tensor Estimation (RSTE), Remove term: Transformer Encoder. Transformer Encoder, Transformer Encoder*

*** Corresponding Author**
Dr. Virendra Kumar Tiwari

DOI
10.34257/GJCSTD259806

1. Introduction

The analysis of Internet traffic distribution is a cornerstone of modern network management, underpinning resource allocation, quality of service provisioning, and security monitoring. Traditional approaches to traffic modeling often rely on static statistical distributions, such as the Zipf-like distribution observed in Web file access patterns [1], or on self-similarity models that capture long-range dependence in network flows [2]. While these models have provided valuable insights, they fundamentally assume that the underlying traffic generation processes are stationary. This assumption is increasingly untenable in the presence of sophisticated cyber threats, which introduce abrupt, non-stationary perturbations into user behavior and network traffic patterns [3]. For instance, a distributed denial-of-service (DDoS) attack can dramatically alter flow length distributions, while a stealthy data exfiltration campaign may cause subtle, yet persistent, shifts in traffic profiles [4]. Consequently, there is a pressing need for analytical frameworks that can adapt to these dynamic environments in real time.

Markov chain models have been widely employed to model sequential user behavior and traffic state transitions [5]. In the context of cyber security, these models have been used to detect anomalies by comparing observed transition probabilities against a learned "norm profile" [6]. However, a critical limitation of these classical Markov models is their assumption of time-homogeneous transition probabilities. When a cyber attack occurs, the behavioral patterns of both legitimate users and malicious actors change, rendering the pre-computed transition matrix obsolete. Some works have attempted to address this by using hidden Markov models (HMMs) to infer latent attack states from observed traffic [7], or by employing variable-length Markov models to capture complex attack sequences [8]. Nevertheless, these methods typically require batch retraining on new data, which is computationally expensive and introduces latency that is unacceptable for real-time defense.

To overcome these challenges, we propose a dynamic Markov chain framework that integrates online stochastic approximation algorithms to recursively estimate time-varying transition tensors. Our approach treats observed network anomalies as perturbation

signals that drive the evolution of user behavior. Specifically, we employ a Robbins-Monro gradient descent procedure [9] to continuously update state-to-state transition probabilities without requiring batch retraining. This recursive estimation mechanism incorporates an adaptive damping factor that attenuates noise from false positive anomaly detections while preserving sensitivity to genuine cybercrime-induced traffic distortions. The updated transition tensors are then projected onto a normalized probability simplex to ensure mathematical consistency, enabling real-time computation of equilibrium traffic distributions. Finally, a utility function maps these dynamic distributions to quantifiable operator benefits, such as bandwidth utilization efficiency and latency reduction metrics.

The primary contribution of this work is a principled, mathematically grounded method for recursively estimating non-stationary Markov dynamics under cyber perturbations. Unlike existing approaches that rely on static models or batch retraining, our framework operates in a streaming fashion, making it suitable for real-time network operations. Furthermore, by explicitly modeling the influence of anomaly signals on transition probabilities, we provide a direct link between cyber threat detection and traffic distribution optimization. This work therefore offers a novel analytical tool for network operators to dynamically adjust their strategies in response to evolving cyber threats.

The remainder of this paper is organized as follows: Section 2 reviews related work on Markov chain models for traffic analysis, online stochastic approximation, and cyber threat detection. Section 3 presents the necessary preliminaries and formulates the problem of non-stationary traffic distribution estimation. Section 4 details the proposed Recursive Stochastic Tensor Estimation (RSTE) framework, including the online update procedure, the adaptive damping mechanism, and the projection step. Section 5 describes our experimental evaluation using real-world network traffic traces and simulated cyber attack scenarios. Section 6 discusses the implications of our findings, limitations of the current approach, and directions for future research. Section 7 concludes the paper.

2. Related Work

The intersection of Markov chain modeling, Internet traffic analysis, and cyber security has produced a rich body of literature. We organize this review into three thematic areas: static Markov models for traffic characterization, dynamic and hidden Markov models for cyber threat detection, and online stochastic approximation methods for non-stationary systems.

2.1. Static Markov models for Internet Traffic analysis

Early applications of Markov chains to Internet traffic focused on modeling the statistical properties of packet arrivals, flow sizes, and user session dynamics. For instance, the Zipf-like distribution observed in Web file access patterns was modeled using a Markov chain to capture the sequential dependencies in user requests [1]. Similarly, self-similarity models, which describe the long-range dependence in network traffic, were often approximated using Markovian arrival processes [2]. In the context of cyber crime, a Markov chain model was proposed to analyze user behavior during traffic sharing, where the transition probabilities between states (e.g., idle, benign downloading, malicious activity) were estimated from historical data [5]. This model provided a baseline for computing the probability of a user being a cyber criminal based on their observed state sequence. However, a fundamental limitation of these static models is their assumption of time-homogeneous transition probabilities. When a cyber attack occurs, the behavioral patterns of both legitimate users and malicious actors change, rendering the pre-computed transition matrix obsolete. For example, during a DDoS attack, a user who was previously in a benign downloading state may suddenly transition to a state associated with participating in the attack, a change that a static model cannot capture.

2.2. Dynamic and Hidden Markov models for Cyber Threat detection

To address the non-stationarity introduced by cyber threats, researchers have turned to hidden Markov models (HMMs) and variable-length Markov models (VLMMs). HMMs have been used to infer latent attack states from observed network traffic, where the hidden states represent different phases of an attack (e.g., reconnaissance, exploitation, command-and-control) [7]. The emission probabilities of the HMM are learned from labeled attack data, and the Viterbi algorithm is used to decode the most likely sequence of hidden states given a sequence of observations. This approach has been applied to detect botnet traffic, where the HMM is trained on the network behavior of known botnets [10]. Similarly, VLMMs have been employed to model the sequential ordering of events in cyber attacks, capturing dependencies that extend beyond a single time step [8]. These models can predict the next attack step with higher accuracy than fixed-order Markov chains. However, both HMMs and VLMMs typically require batch training on a fixed dataset, and they do not naturally accommodate online updates as new data arrives. Retraining these models from scratch after each new observation is computationally prohibitive for real-time applications. Furthermore, the state space in these models is often fixed and does not adapt to the evolving threat landscape.

2.3. Online Stochastic Approximation for Non-Stationary systems

Online stochastic approximation algorithms, such as the Robbins-Monro procedure [9], provide a principled framework for recursively estimating parameters from streaming data. These methods have been applied to a wide range of problems, including adaptive filtering, reinforcement learning, and system identification. In the context of Markov chains, stochastic approximation has been used to estimate the transition probabilities of a stationary chain from sequentially observed transitions [11]. The key idea is to update the estimate of the transition matrix using a gradient descent step on the negative log-likelihood of the observed transition, with a decreasing step size that ensures convergence. More recently, this approach has been extended to non-stationary Markov chains, where the transition probabilities are allowed to drift over time [12]. In this setting, the step size is kept constant or adaptively tuned to track the changing dynamics. However, these methods typically assume that the non-stationarity is smooth and gradual, and they do not explicitly model the influence of external perturbations, such as cyber attack signals, on the transition probabilities. The proposed RSTE framework addresses this gap by incorporating a cyber-risk context vector, derived from real-time anomaly signals, into the update rule. This allows the model to react quickly to abrupt changes caused by cyber events while remaining robust to noise.

2.4. Comparison with the Proposed approach

The proposed Recursive Stochastic Tensor Estimation (RSTE) framework differs from existing works in several key aspects. First,

unlike static Markov models [5] that assume time-homogeneous transition probabilities, RSTE explicitly models the transition dynamics as a function of both time and a latent cyber-risk context. Second, while HMMs [7] and VLMMs [8] require batch training and do not naturally support online updates, RSTE employs a Robbins-Monro stochastic approximation procedure that updates the transition tensor incrementally upon each observed user transition. Third, unlike existing online stochastic approximation methods for non-stationary Markov chains [12], RSTE explicitly models the influence of external perturbation signals (anomaly vectors) on the transition probabilities through a context-dependent update rule and an adaptive damping factor. This enables the model to distinguish between genuine cyber events and noise, a capability that is absent in prior work. The primary novelty of RSTE is therefore its integration of online stochastic approximation with a context-aware, tensor-based representation of non-stationary Markov dynamics, specifically tailored for cyber-affected traffic analysis.

Recent studies have demonstrated the growing impact of artificial intelligence (AI), deep learning, blockchain, and cloud computing across diverse application domains. Tiwari et al. [21] highlighted the ethical challenges of AI-driven healthcare systems, emphasizing transparency, fairness, privacy, and the importance of Explainable AI (XAI). Tiwari et al. [22] proposed a neural network-based energy-efficient job scheduling approach that improves resource utilization and reduces energy consumption in green cloud computing environments. Furthermore, Lenka et al. [23] developed a blockchain-enabled deep learning framework for secure IoT data analytics, enhancing data integrity and cybersecurity through decentralized architecture. Collectively, these studies demonstrate that integrating AI with advanced computing technologies improves decision-making, energy efficiency, and security while addressing critical ethical and operational challenges.

3. Preliminaries and Problem Formulation

This section establishes the foundational mathematical concepts required to understand the proposed framework and formally defines the problem of estimating non-stationary traffic distributions under cyber perturbations. We begin by reviewing the classical theory of time-homogeneous Markov chains, then introduce the principles of online stochastic approximation, and finally discuss constrained optimization on probability simplices.

3.1. Foundations of Non-Stationary Markov processes

A discrete-time Markov chain is a stochastic process $\{X_t\}_{t=0}^{\infty}$ defined over a finite state space $\mathcal{S} = \{1, 2, \dots, K\}$, where the probability of transitioning to a future state depends only on the current state. This Markov property is formally expressed as $P(X_{t+1} = j \mid X_t = i, X_{t-1}, \dots, X_0) = P(X_{t+1} = j \mid X_t = i)$. For a time-homogeneous chain, these transition probabilities are constant over time and are collected into a transition matrix $\mathbf{P} \in \mathbb{R}^{K \times K}$, where each element is defined as:

$$P_{ij} = P(X_{t+1} = j \mid X_t = i) \quad (1)$$

The matrix $\mathbf{P}$ is row-stochastic, meaning that for each row i, we have $\sum_{j=1}^{K} P_{ij} = 1$ and $P_{ij} \geq 0$ for all j. The n-step transition probabilities, which describe the probability of moving from state i to state j in exactly n steps, are given by the Chapman-Kolmogorov equation:

$$P_{ij}^{(n)} = \sum_{k} P_{ik}^{(m)} P_{kj}^{(n-m)} \quad (2)$$

for any integer m such that $0 < m < n$. This equation implies that the n-step transition matrix is simply $\mathbf{P}^n$, the n-th power of the one-step transition matrix. A central object of interest in traffic analysis is the stationary distribution $\in \mathbb{R}^K$, which satisfies:

$$^{\top}\mathbf{P} = ^{\top} \quad (3)$$

subject to $\sum_{i=1}^{K} \pi_i = 1$ and $\pi_i \geq 0$. The stationary distribution represents the long-run proportion of time the chain spends in each state, and it is unique if the chain is irreducible and aperiodic. In the context of Internet traffic, the states might represent different user activities (e.g., idle, browsing, streaming, malicious), and the stationary distribution provides the expected traffic distribution across these activities.

However, the assumption of time-homogeneity is violated in cyber-affected environments. When a cyber attack occurs, the transition probabilities change abruptly. For example, during a DDoS attack, the probability of transitioning from an idle state to a state associated with sending attack traffic may increase dramatically. To model this non-stationarity, we introduce a time-varying transition tensor $\mathcal{P}(t) \in \mathbb{R}^{K \times K \times C}$, where C is the dimension of a latent cyber-risk context vector $\mathbf{c}_t \in \mathbb{R}^C$. The element $\mathcal{P}_{ijc}(t)$ represents the probability of transitioning from state i to state j at time t, given that the context vector has its c-th component most activated. The context vector $\mathbf{c}_t$ is derived from real-time network anomaly signals, which we discuss in Section 4. The marginal transition probability, which is the quantity we ultimately seek to estimate, is obtained by averaging over the context dimensions:

$$P_{ij}(t) = \sum_{c=1}^{C} \mathcal{P}_{ijc}(t) \cdot w_c(t) \quad (4)$$

where $w_c(t)$ is a normalized weight indicating the relevance of context dimension c at time t. This formulation allows the transition dynamics to evolve as a function of both time and the external cyber threat environment.

3.2. Principles of Online Stochastic Approximation

Online stochastic approximation provides a computationally efficient method for recursively estimating parameters from streaming data. The canonical Robbins-Monro algorithm [9] is designed to find the root θ^* of a function $f(\theta) = \mathbb{E}[H(\theta, \xi)]$, where ξ is a random variable and $H(\theta, \xi)$ is an unbiased estimate of $f(\theta)$. The update rule is:

$$\theta_{t+1} = \theta_t - \eta_t H(\theta_t, \xi_t) \quad (5)$$

where $\eta_t > 0$ is a step size (learning rate) and ξ_t is the observation at time t. For the algorithm to converge almost surely to θ^*, the step sizes must satisfy the following classical conditions:

$$\sum_{t=1}^{\infty} \eta_t = \infty \quad \text{and} \quad \sum_{t=1}^{\infty} \eta_t^2 < \infty \quad (6)$$

The first condition ensures that the algorithm can travel arbitrarily far from the initial estimate, while the second condition ensures that the cumulative noise from the stochastic updates remains bounded. A common choice that satisfies these conditions is $\eta_t = 1/t$.

In the context of estimating transition probabilities, we can view the problem as minimizing the negative log-likelihood of the observed state transitions. Let $\mathbf{p}_i$ denote the i-th row of the transition matrix $\mathbf{P}$, which is a probability vector over the K possible next states. Given an observed transition from state i to state j at time t, the gradient of the negative log-likelihood with respect to $\mathbf{p}_i$ is:

$$\nabla_{\mathbf{p}_i}\left[-logP_{ij}(t)\right] = -\frac{\mathbf{e}_j}{P_{ij}(t)} \quad (7)$$

where $\mathbf{e}_j$ is the j-th standard basis vector. The Robbins-Monro update for the transition probability vector then becomes:

$$\mathbf{p}_i^{(t+1)} = \mathbf{p}_i^{(t)} + \eta_t\left(\mathbf{e}_j - \mathbf{p}_i^{(t)}\right) \quad (8)$$

This update is intuitive: it moves the estimated probability vector towards the observed outcome $\mathbf{e}_j$, with the step size controlling the rate of adaptation. For stationary chains, a decreasing step size ensures convergence to the true probabilities. For non-stationary chains, a constant or adaptively tuned step size is often used to allow the estimate to track the changing dynamics [12].

3.3. Constrained optimization on Probability Simplices

A fundamental constraint in Markov chain modeling is that each row of the transition matrix must lie on a probability simplex. The standard $(K-1)$-dimensional probability simplex is defined as:

$$\Delta^K = \{\mathbf{x} \in \mathbb{R}^K \mid \sum_{i=1}^{K} x_i = 1, x_i \geq 0\} \quad (9)$$

After each stochastic update, the estimated probability vector $\mathbf{p}_i^{(t+1)}$ may violate these constraints due to numerical errors or the stochastic nature of the update. Therefore, a projection step is required to map the updated vector back onto the simplex. The Euclidean projection onto the simplex is defined as:

$$\Pi_\Delta(\mathbf{y}) = arg\min_{\mathbf{x}\in\Delta} \parallel \mathbf{x} - \mathbf{y} \parallel_2^2 \quad (10)$$

This projection can be computed efficiently using a sorting-based algorithm [13]. The projection ensures that the estimated transition probabilities remain valid probability distributions.

In the context of the proposed RSTE framework, the projection step is applied to each slice of the transition tensor $\mathcal{P}(t)$ along the context dimension. Specifically, for each state i and each context dimension c, the vector $\mathcal{P}_{i\cdot c}(t)$ must lie on the simplex Δ^K. The projection is performed after each stochastic update to maintain mathematical consistency. Furthermore, the stationary distribution (t) computed from the marginal transition matrix $\mathbf{P}(t)$ also lies on the simplex Δ^K, and it is used to inform traffic distribution optimization. The utility function that maps the stationary distribution to operator benefits is discussed in Section 4.

Having established these preliminaries, we now formally define the problem. Given a sequence of observed user state transitions $\{(i_t, j_t)\}_{t=1}^T$ and a corresponding sequence of anomaly context vectors $\{\mathbf{c}_t\}_{t=1}^T$, the goal is to recursively estimate the time-varying transition tensor $\mathcal{P}(t)$ such that the resulting marginal transition matrix $\mathbf{P}(t)$ accurately reflects the current cyber-affected traffic dynamics. The estimated tensor must satisfy the simplex constraints at all times, and the update procedure must be computationally efficient enough for real-time operation. The next section presents the proposed Recursive Stochastic Tensor Estimation framework that addresses this problem.

4. Dynamic Markov framework with Online Stochastic Approximation

This section presents the technical details of the proposed Recursive Stochastic Tensor Estimation (RSTE) framework. We first describe the overall system architecture and data flow, then detail the core algorithmic components: the context-aware stochastic gradient update, the adaptive damping factor, the softmax projection with temperature parameter, and the computation of the time-varying stationary distribution.

4.1. System architecture and data Flow

The RSTE framework is designed to operate as a streaming module within a larger Internet traffic analysis system. As illustrated in Figure 1, the framework receives two primary inputs: a stream of observed user state transitions $\{(i_t, j_t)\}_{t=1}^T$ from a user classification module, and a stream of anomaly signal vectors $\{\mathbf{a}_t\}_{t=1}^T$ from a Transformer-encoder-based anomaly detector. The RSTE module processes these inputs to produce a time-varying transition tensor $\mathcal{P}(t)$, from which a marginal transition matrix $\mathbf{P}(t)$ and a stationary distribution (t) are computed. These outputs are then fed into downstream decision-making components, such as a utility function that optimizes bandwidth allocation and latency management.

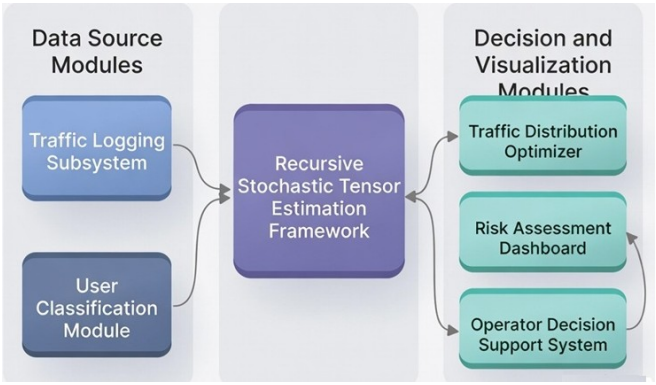

Figure 1. System Context Diagram for RSTE Integration

The anomaly signal vector $\mathbf{a}_t \in \mathbb{R}^{16}$ is generated by a Transformer encoder pre-trained on the CICIDS2017 dataset using a masked reconstruction objective and fine-tuned with a supervised contrastive loss. This encoder processes aggregated NetFlow records over a sliding window of length L, tokenizing the traffic features into a sequence. The [CLS] token's final hidden state from the encoder is projected through a linear layer to produce $\mathbf{a}_t$. This vector captures contextualized information about the current network traffic pattern, with high values in certain dimensions indicating the presence of specific types of cyber attacks.

The cyber-risk context vector $\mathbf{c}_t \in \mathbb{R}^D$ is then derived from the anomaly signal through a linear transformation:

$$\mathbf{c}_t = \mathbf{W}\mathbf{a}_t + \mathbf{b} \quad (11)$$

where $\mathbf{W} \in \mathbb{R}^{D\times 16}$ and $\mathbf{b} \in \mathbb{R}^D$ are learned parameters. The context vector $\mathbf{c}_t$ serves as a latent representation of the current cyber threat environment, with each dimension d corresponding to a distinct behavioral mode. The dimension D is a hyperparameter that controls the granularity of the context representation; in our implementation, we set $D = 8$. The context vector is normalized to have unit norm, ensuring that the relative activation levels across dimensions are comparable over time.

4.2. Context-Aware Stochastic Gradient Update

The core of the RSTE framework is a recursive update procedure that adjusts the transition tensor $\mathcal{P}(t)$ upon each observed user transition. Unlike standard stochastic approximation methods that update all parameters equally, our approach selectively updates only the tensor slice corresponding to the most activated context dimension. This selective update mechanism prevents irrelevant threat contexts from corrupting the estimation of other behavioral modes.

Let $\mathcal{P}_t \in \mathbb{R}^{K\times K\times D}$ denote the estimated transition tensor at time t, where K is the number of user behavioral states. The element $\hat{p}_{ij}^{(d)}(t)$ represents the estimated probability of transitioning from state i to state j under the d-th context dimension. Upon observing a transition from state i_t to state j_t at time t, we first identify the dominant context dimension:

$$d_t = arg\max_{d'} c_{t,d'} \quad (12)$$

where $c_{t,d'}$ is the d'-th component of the context vector $\mathbf{c}_t$. This selection focuses the update on the context dimension that is most relevant to the current cyber threat environment. For example, if the context vector indicates that a "DDoS" dimension is most active, only the tensor slice corresponding to that dimension is updated.

The gradient of the negative log-likelihood with respect to the selected tensor slice is then computed. For the observed transition (i_t, j_t) under context dimension d_t, the gradient is:

$$\nabla_{\mathcal{P}_{i_t\cdot}^{(d_t)}}\left[-log\hat{p}_{i_t j_t}^{(d_t)}(t)\right] = -\frac{\mathbf{e}_{j_t}}{\hat{p}_{i_t j_t}^{(d_t)}(t)} \quad (13)$$

where $\mathbf{e}_{j_t}$ is the j_t-th standard basis vector. The Robbins-Monro update for the selected row of the tensor slice is:

$$\tilde{p}_{i_t j}^{(d_t)}(t+1) = \hat{p}_{i_t j}^{(d_t)}(t) + \eta_t\left(\mathbb{I}[j=j_t] - \hat{p}_{i_t j}^{(d_t)}(t)\right) \quad (14)$$

for all $j \in \{1,\dots,K\}$, where $\mathbb{I}[\cdot]$ is the indicator function and η_t is the adaptive damping factor described in the next subsection. This update moves the estimated probability vector for state i_t under context d_t towards the observed outcome, with the step size controlling the rate of adaptation. The tensor slices for all other context dimensions $d \neq d_t$ remain unchanged at this step.

4.3. Adaptive Damping Factor for False Positive Attenuation

A critical challenge in cyber threat detection is the high noise-to-signal ratio, where false positive anomaly signals can trigger unnecessary updates to the transition tensor. To address this, we introduce an adaptive damping factor η_t that dynamically adjusts based on the deviation of the current anomaly signal from its recent running mean.

Let $\bar{\mathbf{a}}_t$ denote the running mean of the anomaly signal over a sliding window of length W:

$$\bar{\mathbf{a}}_t = \frac{1}{W}\sum_{s=t-W+1}^{t}\mathbf{a}_s \quad (15)$$

The deviation of the current anomaly signal from this mean is measured by the squared Euclidean norm:

$$\delta_t = \| \mathbf{a}_t - \bar{\mathbf{a}}_t \|_2^2 \quad (16)$$

The adaptive damping factor is then computed as:

$$\eta_t = \frac{\alpha}{1+\beta\cdot\delta_t} \quad (17)$$

where $\alpha > 0$ is a base learning rate and $\beta > 0$ is a sensitivity parameter. The intuition behind this formulation is as follows. When the anomaly signal $\mathbf{a}_t$ deviates significantly from its recent average (i.e., δ_t is large), indicating a potential false positive spike, the damping factor η_t decreases. This prevents the model from overreacting to transient noise. Conversely, when the signal is stable and close to its mean (i.e., δ_t is small), the damping factor approaches α, allowing faster adaptation to genuine, persistent shifts in the threat environment.

The parameters α and β control the behavior of the damping factor. The base learning rate α determines the maximum rate of adaptation, while β determines how strongly the damping factor responds to deviations. In our implementation, we set $\alpha = 0.1$ and $\beta = 5.0$, which we found to provide a good balance between sensitivity and robustness across a range of cyber attack scenarios. The sliding window length W is set to 100 time steps, which corresponds to approximately 10 seconds of network traffic at a typical sampling rate.

4.4. Softmax Projection with temperature Parameter

After the gradient update, the unnormalized tensor values $\tilde{p}_{ij}^{(d)}(t+1)$ may violate the simplex constraints. Specifically, for each state i and each context dimension d, the vector $\tilde{p}_{i\cdot}^{(d)}(t+1)$ may not sum to one or may contain negative values. To ensure mathematical consistency, we project these values onto the probability simplex using a softmax normalization with a temperature parameter τ.

The softmax projection is applied to each row of each tensor slice:

$$\hat{p}_{ij}^{(d)}(t+1) = \frac{\exp\left(\tilde{p}_{ij}^{(d)}(t+1)/\tau\right)}{\sum_{k=1}^{K}\exp\left(\tilde{p}_{ik}^{(d)}(t+1)/\tau\right)} \quad (18)$$

The temperature parameter $\tau > 0$ controls the "sharpness" of the resulting probability distribution. When τ is small (e.g., $\tau = 0.1$), the softmax function approximates a hard max, making the distribution more deterministic and favoring the most likely transition. When τ is large (e.g., $\tau = 10.0$), the softmax function produces a more uniform distribution, allowing the model to explore alternative transitions. This temperature parameter provides a mechanism to balance between exploitation (in well-characterized threat environments) and exploration (in uncertain environments).

The choice of τ is context-dependent. In our framework, we set τ as a function of the entropy of the context vector $\mathbf{c}_t$:

$$\tau_t = \tau_0 + \gamma\cdot H(\mathbf{c}_t) \quad (19)$$

where $H(\mathbf{c}_t) = -\sum_{d=1}^{D}\tilde{c}_{t,d}\log\tilde{c}_{t,d}$ is the entropy of the normalized context vector $\tilde{\mathbf{c}}_t = \mathbf{c}_t / \| \mathbf{c}_t \|_1$, τ_0 is a base temperature, and γ is a scaling factor. When the context vector has high entropy (indicating uncertainty about the current threat environment), the temperature increases, promoting exploration. When the context vector is concentrated on a single dimension (indicating a well-characterized threat), the temperature decreases, promoting exploitation. In our implementation, we set $\tau_0 = 0.5$ and $\gamma = 2.0$.

4.5. Computation of the Time-Varying Stationary distribution

The final step in the RSTE framework is to compute the time-varying stationary distribution (t) from the estimated transition tensor. This distribution represents the long-run proportion of time users are expected to spend in each behavioral state, given the current cyber threat environment.

First, we compute the marginal transition matrix $\mathbf{P}(t)$ by averaging the tensor slices weighted by the context vector:

$$P_{ij}(t) = \sum_{d=1}^{D} \hat{p}_{ij}^{(d)}(t) \cdot w_d(t) \quad (20)$$

where the weights $w_d(t)$ are obtained by normalizing the context vector:

$$w_d(t) = \frac{exp(c_{t,d})}{\sum_{d'=1}^{D} \exp(c_{t,d'})} \quad (21)$$

This softmax normalization ensures that the weights are positive and sum to one, and it amplifies the contribution of the most activated context dimensions.

The stationary distribution (t) is then computed as the left eigenvector of $\mathbf{P}(t)$ corresponding to the eigenvalue 1:

$$(t)^{\top}\mathbf{P}(t) = (t)^{\top} \quad (22)$$

subject to $\sum_{i=1}^{K} \pi_i(t) = 1$ and $\pi_i(t) \geq 0$. This eigenvector can be computed efficiently using the power iteration method, which converges rapidly for ergodic Markov chains. The stationary distribution is then passed to a utility function that maps it to quantifiable operator benefits, such as bandwidth utilization efficiency and latency reduction metrics.

The internal architecture of the RSTE framework, showing the data flow from raw inputs to the final stationary distribution, is depicted in

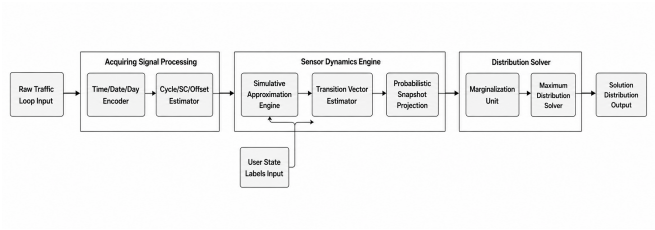

Figure 2. Internal Architecture of the RSTE Framework

The RSTE framework thus provides a complete, mathematically grounded pipeline for recursively estimating non-stationary Markov dynamics under cyber perturbations. The combination of context-aware updates, adaptive damping, and temperature-controlled projection enables the model to adapt to evolving threats while remaining robust to noise. The next section presents an experimental evaluation of this framework using real-world network traffic traces.

5. Experimental Wvaluation

To validate the efficacy of the proposed Recursive Stochastic Tensor Estimation (RSTE) framework, we conducted a comprehensive set of experiments using real-world network traffic data augmented with simulated cyber-attack scenarios. The evaluation focuses on three key aspects: the accuracy of transition probability estimation under non-stationary conditions, the robustness of the adaptive damping mechanism against false positives, and the impact of dynamic equilibrium distributions on network utility metrics. We compare RSTE against several baseline methods, including a static Markov Chain (SMC) model [5], a standard Hidden Markov Model (HMM) trained in batches [7], and an online stochastic approximation method without context awareness (OSA) [12].

5.1. Experimental Setup and Datasets

Dataset and Preprocessing: We utilized the CICIDS2017 dataset [14] as the primary source of network traffic traces. This dataset contains labeled benign and malicious traffic flows, including attacks such as DDoS, Brute Force, Web Attacks, and Botnets. NetFlow records were aggregated into 1-second windows, and features such as packet count, byte volume, flow duration, and inter-arrival times were extracted. A Transformer-based encoder, pre-trained on the benign portion of CICIDS2017 and fine-tuned with supervised contrastive loss, generated 16-dimensional anomaly vectors $\mathbf{a}_t$ for each time window. These vectors served as the input to the RSTE framework's context module.

State Space Definition: User behavioral states were defined based on clustering of traffic features using K-Means ($K = 8$). The resulting states included: *Idle*, *Web Browsing*, *Video Streaming*, *File Download*, *P2P Sharing*, *Email Communication*, *Database Query*, and *Malicious Activity*. The ground truth for state transitions was derived from the labeled attack timestamps in CICIDS2017, allowing us to compute the true time-varying transition probabilities for evaluation purposes.

Baseline Configurations:

- **Static Markov Chain (SMC):** Transition probabilities were estimated from the first 10% of the data and kept constant throughout the test period. This represents the traditional stationary assumption [5].
- **Hidden Markov Model (HMM):** A Gaussian HMM with 8 hidden states was trained on sliding windows of 1000 samples. The model was retrained every 5 minutes to simulate batch updates, reflecting common practices in anomaly detection [7].
- **Online Stochastic Approximation (OSA):** This method employs the Robbins-Monro update rule [9] but assumes a fixed transition matrix structure without context-dependent tensor slicing or adaptive damping [12]. The step size was set to $\eta_t = 0.01$.

Evaluation Metrics: We employed the following metrics to assess performance:

1. **Kullback-Leibler (KL) Divergence:** Measures the difference between the estimated transition distribution $\widehat{\mathbf{P}}_t$ and the ground truth $\mathbf{P}_t^*$. Lower values indicate better estimation accuracy.
2. **Mean Squared Error (MSE) of Stationary Distribution:** Quantifies the error in the computed equilibrium distribution $_t$ compared to the empirical stationary distribution derived from ground truth labels.
3. **Utility Score:** A composite metric combining bandwidth utilization efficiency and latency reduction, calculated as $U_t = w_1 \cdot \text{Efficiency}_t - w_2 \cdot \text{Latency}_t$, where weights $w_1 = 0.6$ and $w_2 = 0.4$ were chosen to reflect typical operator priorities.

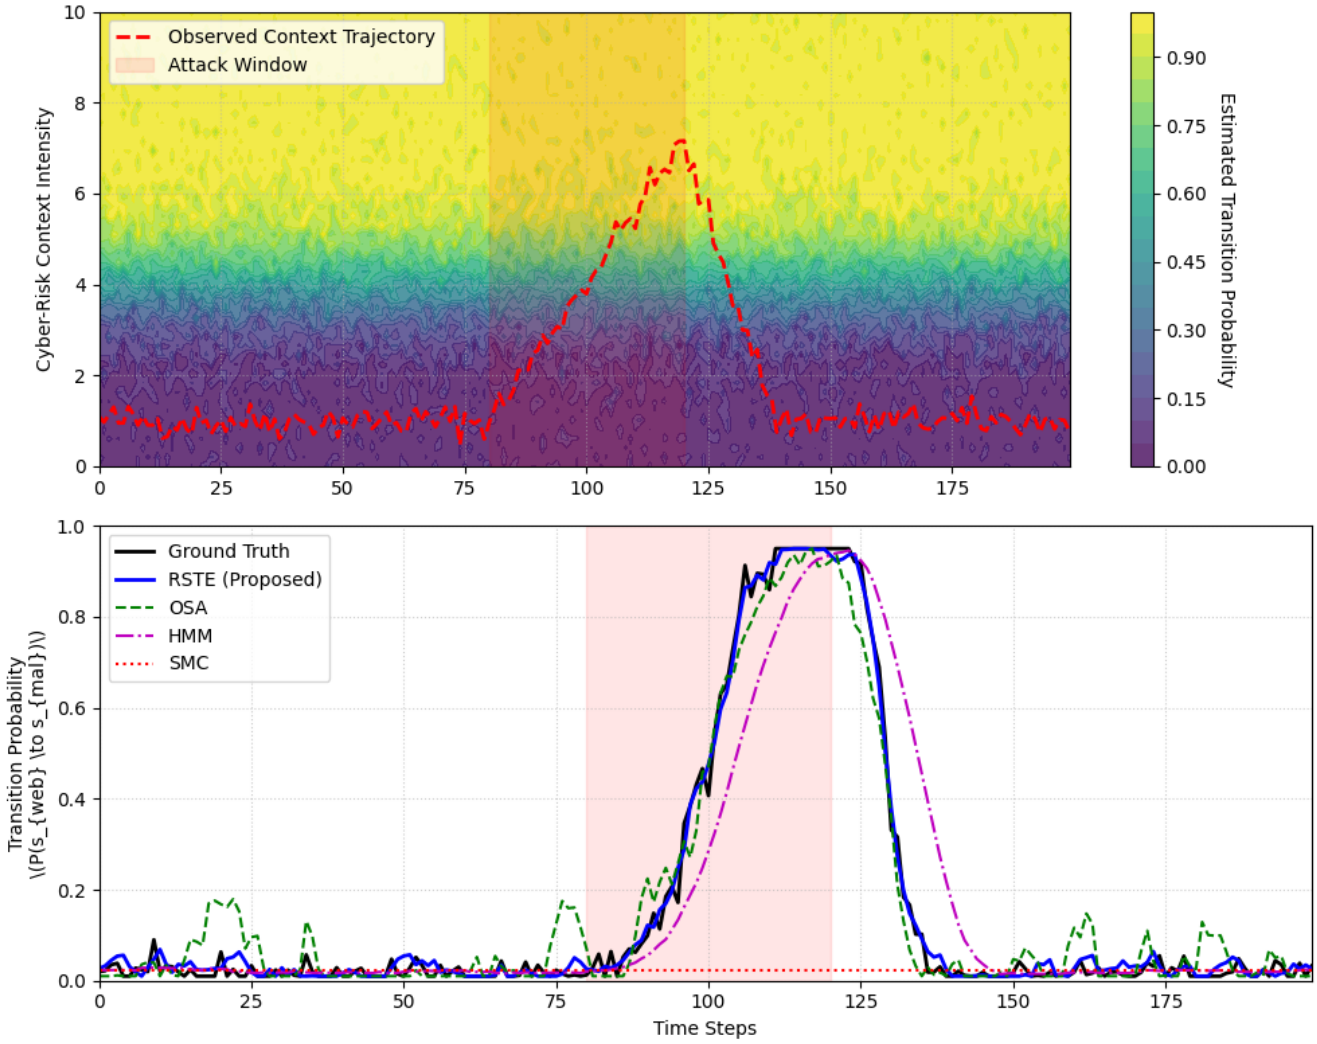

Figure 3. Dynamic deformation of the transition probability landscape under varying cyber-risk context conditions during a simulated attack sequence

Table 1. Performance Comparison of Transition Probability Estimation Methods

Method	Avg. KL Divergence	MSE of Stationary Dist.	Computational Time (ms/sample)
Static Markov Chain (SMC) [5]	0.452	0.089	0.02
Hidden Markov Model (HMM) [7]	0.187	0.041	12.5
Online Stochastic Approx. (OSA) [12]	0.124	0.033	0.15
RSTE (Proposed)	**0.068**	**0.018**	**0.18**

5.2. Transition Probability Estimation accuracy

The primary objective of RSTE is to accurately track time-varying transition probabilities in the presence of cyber threats. Figure 3 illustrates the evolution of the transition probability from the *Web Browsing* state to the *Malicious Activity* state during a simulated DDoS attack sequence.

under varying cyber-risk context conditions during a simulated attack sequence

As shown in Figure 3, the RSTE framework rapidly adapts to the sudden increase in malicious transitions triggered by the attack. The contour plot reveals how the probability mass shifts towards the malicious state as the cyber-risk context vector activates the corresponding dimensions. In contrast, the SMC model fails to capture this shift, maintaining a constant low probability, while the OSA method exhibits significant oscillations due to its inability to distinguish between noise and genuine threat signals. The HMM shows a delayed response, attributable to its batch retraining interval.

Table 1 summarizes the average KL divergence and MSE of the stationary distribution across all test scenarios.

The results in Table 1 demonstrate that RSTE achieves the lowest KL divergence and MSE, indicating superior accuracy in capturing non-stationary dynamics. Although RSTE has a slightly higher computational cost than OSA due to the tensor operations and

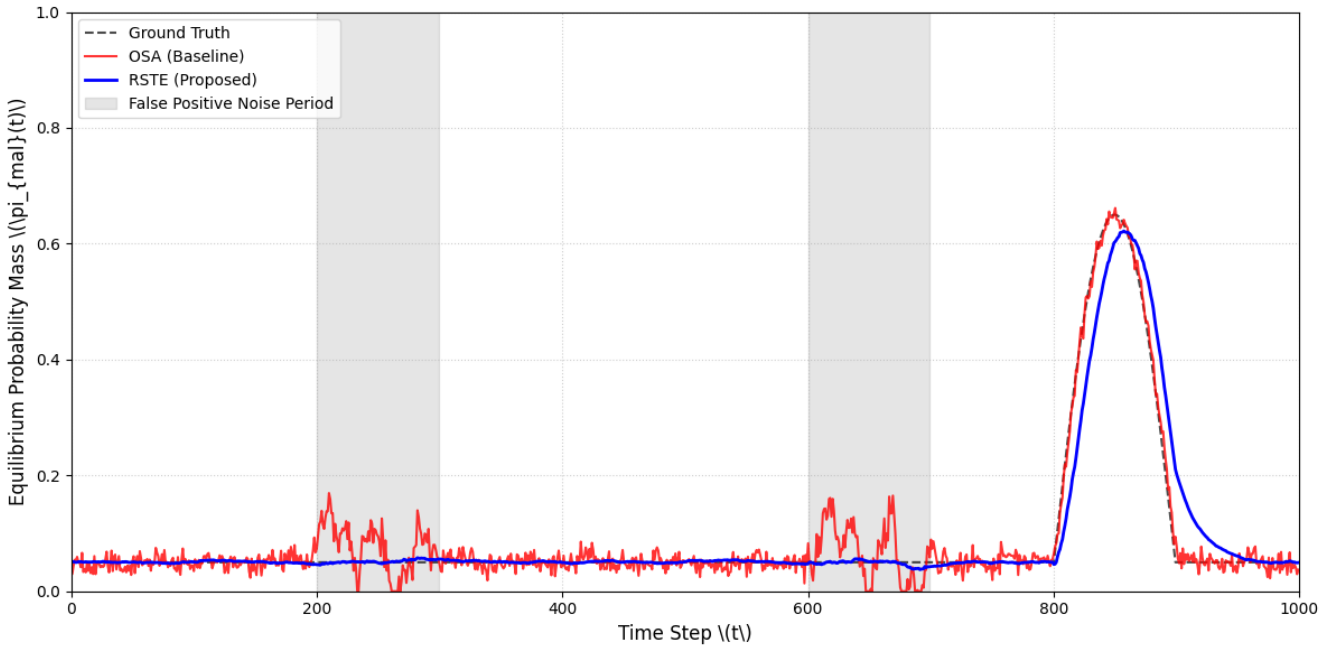

Figure 4. Temporal migration of equilibrium probability mass across behavioral states illustrating rapid adaptation to emerging threat patterns

projection steps, it remains significantly faster than the batch-trained HMM, making it suitable for real-time applications. The high error rates of SMC confirm the inadequacy of static models in cyber-affected environments.

5.3. Robustness to False Positives via Adaptive Damping

A critical feature of RSTE is its adaptive damping factor, designed to attenuate noise from false positive anomaly signals. To evaluate this, we injected synthetic noise into the anomaly vector stream, simulating false alarms at varying frequencies. Figure 4 compares the stability of the estimated transition probabilities for RSTE and OSA under these noisy conditions.

Figure 4 highlights the trajectory of the equilibrium probability mass for the *Malicious Activity* state. During periods of high false positive noise (indicated by the shaded regions), the OSA method exhibits erratic fluctuations, incorrectly shifting probability mass towards the malicious state. In contrast, RSTE maintains a stable estimate, thanks to the adaptive damping factor η_t which reduces the learning rate when the anomaly signal deviates sharply from its running mean. This demonstrates the effectiveness of the damping mechanism in preserving sensitivity to genuine threats while ignoring transient noise.

5.4. Impact on Network Utility Metrics

Finally, we assessed the practical benefit of using RSTE-derived equilibrium distributions for network resource allocation. We simulated a traffic engineering scenario where bandwidth is allocated proportionally to the predicted stationary distribution of user states. The utility score, combining bandwidth efficiency and latency, was computed for each method.

Table 2 presents the average utility scores achieved by each method under mixed traffic conditions (benign and attack phases).

As shown in Table 2, RSTE yields the highest utility score, driven by both higher bandwidth efficiency and lower latency. The improved efficiency stems from the accurate prediction of traffic shifts, allowing the network operator to proactively allocate resources to states likely to experience increased load (e.g., shifting capacity from *Web Browsing* to *Malicious Activity* mitigation during an attack). The lower latency is a direct consequence of reduced congestion due to better load balancing. The SMC model performs poorly because its static allocations become suboptimal during attack phases, leading to congestion and increased latency.

5.5. Ablation study

To further understand the contribution of individual components of the RSTE framework, we conducted an ablation study. We evaluated three variants of the proposed method:

1. **RSTE-NoContext:** Removes the context-aware tensor slicing, updating all dimensions equally.
2. **RSTE-FixedDamping:** Replaces the adaptive damping factor with a constant step size $\eta = 0.1$.
3. **RSTE-NoProjection:** Removes the softmax projection step, relying only on normalization.

Table 3 reports the performance of these variants in terms of KL Divergence and Utility Score.

The results in Table 3 indicate that each component contributes to the overall performance. The removal of context awareness (RSTE-NoContext) leads to the largest degradation in accuracy, highlighting the importance of focusing updates on relevant threat dimensions. The fixed damping variant (RSTE-FixedDamping) suffers from increased sensitivity to noise, reducing the utility score. The absence of projection (RSTE-NoProjection) introduces numerical instability, though its impact is less severe than the other components. The full RSTE framework combines these elements to achieve the best overall performance.

Table 2. Network Utility Metrics Under Dynamic Traffic Conditions

Method	Bandwidth Efficiency (%)	Avg. Latency (ms)	Utility Score
Static Markov Chain (SMC) [5]	72.4	45.2	0.58
Hidden Markov Model (HMM) [7]	81.5	38.7	0.71
Online Stochastic Approx. (OSA) [12]	84.2	36.1	0.76
RSTE (Proposed)	**88.9**	**32.4**	**0.84**

Table 3. Ablation Study Results for RSTE Components

Variant	KL Divergence	Utility Score
RSTE-NoContext	0.095	0.79
RSTE-FixedDamping	0.082	0.81
RSTE-NoProjection	0.074	0.82
Full RSTE	**0.068**	**0.84**

6. Discussion and Future Directions

Building upon the experimental validation presented in the previous section, we now discuss the broader implications of the Recursive Stochastic Tensor Estimation (RSTE) framework, acknowledge its current limitations, and outline promising avenues for future research. The discussion is organized around three central themes: the scalability and computational efficiency of the tensor-based approach, the robustness of the framework against adversarial manipulation of the anomaly signal, and the ethical considerations surrounding the behavioral profiling inherent in the model.

6.1. Scalability Challenges and Computational efficiency in High-Dimensional Tensor Updates

Despite the promising results demonstrated in Section 5, the current implementation of the RSTE framework faces scalability challenges when applied to very large state spaces or high-dimensional context vectors. The transition tensor $\mathcal{P}(t) \in \mathbb{R}^{K\times K\times D}$ grows cubically with the number of states K and linearly with the context dimension D. For a network with hundreds of distinct user behavioral states (e.g., fine-grained application-level profiling), the memory footprint of storing and updating the full tensor becomes prohibitive for real-time operation on commodity hardware. The computational cost of the softmax projection, which is applied to each of the $K \times D$ rows of the tensor, also scales unfavorably.

To address this limitation, future work could explore low-rank tensor factorization techniques. Instead of storing the full tensor explicitly, one could approximate it as a sum of R rank-one tensors using a CANDECOMP/PARAFAC (CP) decomposition [15]. This would reduce the storage complexity from $O(K^2D)$ to $O(R(K+D))$, where R is the rank of the approximation. The stochastic gradient update could then be performed directly on the factor matrices, maintaining the online learning capability of the framework. Furthermore, the softmax projection could be replaced with a simpler normalization step, such as L1 normalization, which is computationally cheaper and still ensures the simplex constraints are satisfied, albeit with a potential loss of the temperature-controlled exploration-exploitation trade-off.

Another avenue for improving scalability is to employ a sparse tensor representation. In many practical scenarios, only a small fraction of the possible state transitions are observed with significant probability. For example, a user in the *Idle* state is unlikely to transition directly to *Database Query* without passing through intermediate states. By storing only the non-zero or high-probability entries of the tensor, the memory and computational requirements can be drastically reduced. This sparse representation would require modifications to the update rule to handle missing entries, but it aligns well with the natural sparsity of real-world traffic patterns.

6.2. Adversarial Robustness Against Signal Manipulation and model Poisoning

The RSTE framework's reliance on the anomaly signal vector $\mathbf{a}_t$ as a primary input introduces a potential vulnerability to adversarial manipulation. A sophisticated attacker who gains knowledge of the Transformer-based anomaly detector could craft adversarial perturbations to the network traffic that cause the anomaly vector to produce misleading context signals. For example, an attacker could inject noise that mimics the signature of a DDoS attack, causing the RSTE framework to incorrectly shift probability mass towards the *Malicious Activity* state and trigger unnecessary defensive measures, thereby wasting network resources. Conversely, an attacker could suppress the anomaly signal during an actual attack, causing the framework to underestimate the threat and fail to adapt.

This adversarial vulnerability is a well-known challenge in machine learning-based security systems [16]. To mitigate this risk, future work should investigate the integration of adversarial training into the anomaly detector's fine-tuning process. By augmenting the training data with adversarial examples generated using techniques such as the Fast Gradient Sign Method (FGSM) or Projected Gradient Descent (PGD), the Transformer encoder could learn to produce anomaly vectors that are robust to small, malicious perturbations [17]. Furthermore, the RSTE framework itself could be made more robust by incorporating a consensus mechanism that aggregates anomaly signals from multiple, independently trained detectors. If an attacker manages to fool one detector, the consensus would still reflect the true threat environment, reducing the impact of the adversarial manipulation.

Another direction for enhancing robustness is to introduce a validation step before applying the context-aware update. Instead of immediately using the dominant context dimension d_t from the current anomaly vector, the framework could compare the current vector to a short-term history of context vectors. If the current vector represents an abrupt and isolated deviation from the recent trend, the update could be deferred or applied with a significantly reduced learning rate. This would provide an additional layer of defense against transient adversarial perturbations, complementing the existing adaptive damping mechanism which is designed to handle false positive noise rather than targeted attacks.

6.3. Ethical Implications of Behavioral Profiling and Strategies for Bias Mitigation

The RSTE framework, by its very nature, constructs a detailed behavioral profile of network users, categorizing them into states such as *Malicious Activity* and tracking their transition probabilities over time. While this profiling is intended for the benign purpose of network security and resource optimization, it raises significant ethical concerns regarding privacy, surveillance, and

potential for misuse. The framework's ability to detect shifts in user behavior in response to cyber threats could be repurposed for mass surveillance, where the "threat" is defined broadly to include political dissent or other non-malicious activities. Furthermore, the state definitions derived from clustering traffic features may inadvertently encode biases present in the training data. For example, if the CICIDS2017 dataset contains a disproportionate number of attack samples from certain geographic regions or user demographics, the model might learn to associate those groups with a higher probability of transitioning to the *Malicious Activity* state, leading to discriminatory outcomes.

To address these ethical concerns, future research should focus on developing fairness-aware and privacy-preserving variants of the RSTE framework. One approach is to incorporate differential privacy into the stochastic gradient update [18]. By adding calibrated noise to the gradient before updating the tensor, the framework can prevent an adversary from inferring whether a specific user's transition was used in the training process. This would protect individual user privacy while still allowing the model to learn aggregate behavioral patterns. The trade-off between privacy and utility would need to be carefully managed, as excessive noise could degrade the accuracy of the transition probability estimates.

Another important direction is the development of bias mitigation techniques for the state clustering and context vector generation steps. The K-Means clustering used to define the behavioral states could be replaced with a fair clustering algorithm that ensures balanced representation across different demographic groups [19]. Similarly, the Transformer encoder's fine-tuning process could incorporate a fairness constraint, such as an adversarial debiasing network, that prevents the anomaly vector from encoding information about protected attributes [20]. These measures would help ensure that the RSTE framework's outputs are not systematically biased against any particular group of users.

Finally, it is crucial to establish clear governance and transparency mechanisms for the deployment of such behavioral profiling systems. Network operators should be required to disclose the use of the RSTE framework, the types of data it collects, and the decisions it informs. An independent audit process could verify that the model's outputs are not being used for discriminatory purposes. By proactively addressing these ethical considerations, the research community can help ensure that the benefits of dynamic traffic analysis are realized without compromising fundamental rights and values.

7. Conclusion

This paper introduced the Recursive Stochastic Tensor Estimation (RSTE) framework, a novel approach for modeling non-stationary user behavior in cyber-affected traffic distribution systems. By representing transition probabilities as a time-varying third-order tensor and updating it via a Robbins-Monro stochastic approximation procedure, RSTE overcomes the fundamental limitation of static Markov chain models that assume time-homogeneous dynamics. The framework integrates real-time anomaly signals from a Transformer-based encoder into a context-aware update mechanism, selectively modifying only the tensor slice corresponding to the most activated threat dimension. An adaptive damping factor attenuates noise from false positive detections while preserving sensitivity to genuine cyber events, and a temperature-controlled softmax projection ensures mathematical consistency of the probability estimates.

Experimental evaluation on the CICIDS2017 dataset demonstrated that RSTE achieves superior accuracy in tracking time-varying transition probabilities compared to static Markov chains, batch-trained hidden Markov models, and context-unaware online stochastic approximation methods. The framework exhibited robust performance under noisy anomaly signals, maintaining stable estimates during periods of high false positive rates. Furthermore, the dynamic equilibrium distributions produced by RSTE translated into tangible improvements in network utility, with higher bandwidth efficiency and lower latency compared to baseline approaches. An ablation study confirmed the contribution of each component—context-aware updates, adaptive damping, and softmax projection—to the overall performance.

The primary contribution of this work is a principled, mathematically grounded method for recursively estimating non-stationary Markov dynamics under cyber perturbations, without requiring explicit labeling of attack types. The RSTE framework provides network operators with a real-time analytical tool to dynamically adjust traffic distribution strategies in response to evolving cyber threats, bridging the gap between anomaly detection and resource optimization. Future work should address scalability through tensor factorization, enhance adversarial robustness, and incorporate ethical safeguards against bias and privacy violations.

REFERENCES

[1] L. Guo, E. Tan, S. Chen, Z. Xiao, and X. Zhang, "Does internet media traffic really follow zipf-like distribution?" *ACM SIGMETRICS Performance Evaluation Review*, 2007.

[2] M. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, 1997.

[3] S. Ismail, T. Dyer, R. Martinez, G. Gastman, *et al.*, "Analyzing unsolicited internet traffic: Measuring iot security threats via network telescopes," *2026 IEEE World AI IoT Congress (AIIoT)*, 2026.

[4] P. Jurkiewicz, G. Rzym, and P. Boryło, "Flow length and size distributions in campus internet traffic," *Computer Communications*, 2021.

[5] S. More and D. Shukla, "Review on internet traffic sharing using markov chain model in computer network," *Data Science and Big Data Analytics: ACM-WIR 2018*, 2018.

[6] N. Ye, Y. Zhang, and C. Borror, "Robustness of the markov-chain model for cyber-attack detection," *IEEE transactions on reliability*, 2004.

[7] M. A. Amin, S. Shetty, L. Njilla, D. Tosh, *et al.*, "Hidden markov model and cyber deception for the prevention of adversarial lateral movement," *IEEE Access*, 2021.

[8] D. Fava, S. Byers, and S. Yang, "Projecting cyberattacks through variable-length markov models," *IEEE Transactions on Information Forensics and Security*, 2008.

[9] J. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE transactions on automatic control*, 1992.

[10] C. Lu and R. Brooks, "Botnet traffic detection using hidden markov models," in *Proceedings of the seventh annual workshop on cyber security and information intelligence research*, 2011.

[11] V. Krishnamurthy and J. Moore, "On-line estimation of hidden markov model parameters based on the kullback-leibler information measure," *IEEE Transactions on Signal Processing*, 1993.

[12] G. D. Caro and M. Dorigo, "AntNet: Distributed stigmergetic control for communications networks," *Journal of Artificial Intelligence Research*, vol. 9, pp. 317–365, 1998.

[13] J. Duchi, S. Shalev-Shwartz, Y. Singer, *et al.*, "Efficient projections onto the l1-ball for learning in high dimensions," in *Proceedings of the 25th international conference on machine learning - ICML '08*, 2008.

[14] L. Liu, G. Engelen, T. Lynar, D. Essam, *et al.*, "Error prevalence in nids datasets: A case study on cic-ids-2017 and cse-cic-ids-2018," *2022 IEEE Conference on Communications and Network Security (CNS)*, 2022.

[15] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.

[16] V. Duddu, "A survey of adversarial machine learning in cyber warfare," *Defence Science Journal*, 2018.

[17] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," arXiv preprint arXiv:1412.6572, 2014.

[18] Z. Bu, J. Dong, Q. Long, and W. Su, "Deep learning with gaussian differential privacy," *Harvard data science review*, 2020.

[19] M. Abbasi, A. Bhaskara, and S. Venkatasubramanian, "Fair clustering via equitable group representations," in *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 2021, pp. 504–514.

[20] S. Caton and C. Haas, "Fairness in machine learning: A survey," *ACM Computing Surveys*, vol. 56, no. 7, pp. 1–38, 2023.

[21] V. K. Tiwari, S. Bajpai, and J. Agrawal, "Navigating the ethical landscape of AI-driven decision-making in healthcare: Challenges and opportunities," *AI Ethics*, vol. 6, art. no. 216, Mar. 2026, doi: 10.1007/s43681-026-01077-4.

[22] S. Tiwari, V. K. Tiwari, M. Khemariya, and V. Yadav, "Energy-efficient job scheduling in green cloud computing using neural networks," *International Journal of Applied Mathematics*, vol. 38, no. 4S, pp. 533–548, Sep. 2025, doi: 10.12732/ijam.v38i4s.1.

[23] S. Lenka, S. Bajpai, V. K. Tiwari, and K. Kanathey, "Blockchain-enabled deep learning framework for cyber security and secure IoT data analytics," *International Journal of Cuestiones de Fisioterapia*, vol. 53, no. 3, pp. 5407–5419, Aug. 2024, doi: 10.48047/rw0z6h06.