IJSRA

International Journal of Science and Research Archive

eISSN: 2582-8185

Cross Ref DOI: 10.30574/ijsra

Journal homepage: https://ijsra.net/

Int. J. Sci. Res. Arch.
IJSRA
International Journal of Science and Research Archive
Research Journal Archive, INDIA

(Review Article)

Managing Adversarial AI Risks Through Governance, Threat Hunting and Continuous Monitoring in Production Systems

Toluwalope Opalana *

Technical Project Manager (Payments), Interswitch Group, USA.

International Journal of Science and Research Archive, 2024, 13(02), 1641-1661

Publication history: Received on 07 October 2024; revised on 24 December 2024; accepted on 29 December 2024

Article DOI: https://doi.org/10.30574/ijsra.2024.13.2.2397

Abstract

The deployment of artificial intelligence in production enterprise systems has introduced a new class of adversarial risks that extend beyond traditional cybersecurity threats. At a broad level, AI-enabled systems increasingly operate in dynamic and exposed environments, making them attractive targets for malicious actors seeking to manipulate data, exploit model behavior, or abuse system capabilities. Adversarial techniques such as data poisoning, model extraction, prompt exploitation, and evasion attacks pose significant risks to system integrity, reliability, and trust, particularly when AI systems are integrated into critical business and decision-making processes. Conventional security controls, when applied in isolation, are often insufficient to address these evolving threats. This paper narrows its focus to managing adversarial AI risks through the coordinated application of governance structures, threat hunting practices, and continuous monitoring in production systems. It positions AI governance as the organizing layer that defines accountability, risk tolerance, and escalation pathways, while threat hunting provides proactive identification of adversarial behaviors targeting AI models and data pipelines. Continuous monitoring mechanisms are examined as essential tools for detecting anomalous patterns, model drift, and exploitation attempts in real time. By integrating these elements across the AI lifecycle, organizations can transition from reactive incident response to anticipatory and resilient risk management. The paper argues that a governance-led, intelligence-driven approach is essential for sustaining secure, reliable, and trustworthy AI operations in adversarial production environments.

Keywords: Adversarial AI; AI Governance; Threat Hunting; Continuous Monitoring; Production Systems; AI Security

1. Introduction adversarial AI as a dynamic risk system

1.1. AI Systems as Control Systems

Artificial intelligence systems deployed in production environments increasingly resemble nonlinear dynamic control systems rather than static predictive artifacts [1]. Once embedded within financial platforms, healthcare infrastructures, industrial automation pipelines, or cloud-native APIs, these systems operate in continuous interaction with external inputs and user behaviors, forming feedback-dependent state transitions [2]. Outputs influence downstream decisions, which recursively alter subsequent inputs, establishing closed-loop behavior consistent with classical control theory [3]. This cyclical dependency transforms model deployment into a dynamic state-evolution process rather than a one-time optimization problem.

Nonlinearity arises from activation functions, ensemble weighting, probabilistic inference thresholds, and confidence calibration layers that introduce disproportionate sensitivity to marginal input perturbations [4]. Such nonlinear responses amplify adversarial susceptibility, especially when gradients or decision boundaries can be strategically

[*] Corresponding author: Toluwalope Opalana

exploited [5]. In production contexts, this sensitivity is compounded by environmental drift and adaptive user behavior, further increasing instability risks [6].

Within this framework, the adversary functions as an external disturbance injected into either the observable input channel or latent feature space [7]. Adversarial examples, poisoning artifacts, query-based model extraction, and distributional manipulation all operate as perturbative forces acting on system equilibrium [8]. Governance mechanisms comprising policy enforcement, retraining triggers, anomaly filters, and audit controls serve as stabilizing controllers that introduce corrective damping into the system [1]. By adopting a control-theoretic interpretation, adversarial risk becomes quantifiable as state deviation rather than descriptive vulnerability, allowing resilience to be evaluated mathematically [3].

1.2. Production Adversarial Risk Model

Let the AI model state at time $t$be defined as vector X_t, representing aggregated performance parameters such as classification accuracy, entropy dispersion, confidence variance, and prediction error gradients [2]. The system's temporal evolution under operational exposure may be formalized as:

Equation (1): System State Transition

$$X_{t+1} = AX_t + BU_t + W_t$$

Where:
X_t: model performance state vector

U_t: governance control input

W_t: adversarial disturbance term

A, B: system matrices governing dynamics

Matrix Aencodes intrinsic system propagation characteristics how prior states influence future behavior under nominal conditions [4]. It reflects architectural sensitivity, parameter persistence, and learning adaptation rates [5]. Matrix Bmodels governance intervention intensity, including monitoring thresholds, adversarial retraining strength, and operational guardrails [6].

The disturbance term $W_t$captures stochastic adversarial influence and may be treated as a random variable with measurable variance and expectation [7]. Its magnitude reflects attack intensity, while its distribution captures unpredictability in adversarial behavior [8]. When $W_t$grows in variance, model state divergence accelerates, especially if Aamplifies perturbations [1]. Interpreting adversarial risk as a disturbance term transforms threat hunting into estimation of W_t's statistical properties, while governance modifies $U_t$to counterbalance its effect [3]. This system representation enables quantitative stability analysis rather than post-hoc incident evaluation [6].

1.3. Risk Stabilization Objective

System stability depends fundamentally on the eigenstructure of matrix A. For discrete-time systems, asymptotic stability requires:

Equation (2): Stability Condition

$$\rho(A) < 1$$

Where $\rho(A)$denotes the spectral radius the largest absolute eigenvalue of A[2]. If $\rho(A) \geq 1$, perturbations persist or amplify over time, leading to cumulative degradation in predictive reliability [4]. In adversarial contexts, this implies that even minor disturbances can escalate into systemic instability if governance is insufficient [5].

Governance damping acts by effectively reshaping the system matrix through calibrated control input U_t, shifting eigenvalues inward within the unit circle [7]. Techniques such as adversarial training, anomaly filtering, query rate limiting, and confidence recalibration alter system response characteristics to reduce amplification of disturbances [8]. Continuous monitoring operates as an eigenvalue-sensitive feedback mechanism, detecting divergence trajectories early and triggering corrective interventions [1].

Maintaining $\rho(A) < 1$therefore becomes a measurable governance objective grounded in spectral control principles [3]. Rather than relying solely on qualitative compliance assessments, stability can be validated through mathematical conditions, enabling adversarial risk management to evolve from reactive mitigation to proactive systemic regulation [6].

2. Adversarial threat hunting as signal detection theory

2.1. Threat Hunting as Statistical Signal Detection

This section reframes adversarial threat hunting in production AI systems as a formal statistical signal detection problem. Rather than relying solely on heuristic anomaly flags or rule-based alerts, adversarial identification is modeled as a hypothesis testing process grounded in probability theory and statistical decision rules [5]. In production environments, model telemetry, input distributions, confidence scores, and output entropy streams collectively form observable signals. Within this context, adversarial behavior represents a deviation signal embedded in background operational noise [6].

Traditional cybersecurity threat hunting often relies on signature matching or rule correlation. However, adversarial AI attacks are frequently adaptive, low-amplitude, and distribution-aware, making deterministic detection insufficient [7]. By treating adversarial risk as a stochastic signal embedded within high-dimensional feature space, detection becomes a structured inference task rather than a reactive investigation [8]. This probabilistic framing enables the use of optimal decision theory to control false positives while maximizing adversarial detection power [9]. It also allows governance mechanisms to be mathematically calibrated based on acceptable risk tolerance thresholds [10].

2.2. Threat Signal Modeling

Let the AI production environment generate observations x, derived from telemetry such as prediction probabilities, input feature distributions, gradient magnitudes, or request frequency vectors. The detection task is framed as a binary hypothesis test:

H_0: normal operation

H_1: adversarial activity

The statistical objective is to determine which hypothesis better explains observed data x. The optimal decision framework under known distributions is based on the likelihood ratio:

Equation (3): Likelihood Ratio

$$\Lambda(x) = \frac{P(x \mid H_1)}{P(x \mid H_0)}$$

Decision rule:

$$\Lambda(x) > \tau$$

Where $\tau$is a threshold determined by acceptable false positive rates and operational risk appetite [11]. If the likelihood of observing $x$under adversarial conditions exceeds that under normal operation beyond threshold τ, the system flags potential adversarial activity.

This rule is derived from the Neyman–Pearson Lemma, which proves that the likelihood ratio test is the most powerful test for a given Type I error rate when distinguishing between two simple hypotheses [12]. By selecting τ, governance teams effectively choose a trade-off between sensitivity (true positive rate) and specificity (false positive rate).

In adversarial AI monitoring, distributions $P(x \mid H_0)$may be estimated from historical clean production data, while $P(x \mid H_1)$can be approximated using adversarial simulation datasets or synthetic perturbation modeling [13]. This modeling transforms threat hunting from qualitative inspection to quantifiable hypothesis testing, allowing detection thresholds to be systematically tuned rather than manually adjusted.

2.3. Entropy-Based Adversarial Detection

While likelihood ratio testing captures global distributional differences, adversarial behavior often manifests through instability in model output uncertainty. Entropy provides a principled measure of uncertainty within predictive distributions. For a discrete random variable Xrepresenting model output probabilities, Shannon entropy is defined as:

Equation (4): Shannon Entropy

$$H(X) = -\sum p(x)\log\, p(x)$$

Entropy quantifies the expected information content of predictions [6]. Under normal operational conditions, production models exhibit relatively stable entropy distributions shaped by learned decision boundaries. However, adversarial inputs frequently cause disproportionate shifts in prediction confidence, either inflating uncertainty or artificially collapsing it [8].

An entropy spike may indicate boundary probing or gradient-based perturbations designed to destabilize classification certainty [9]. Conversely, entropy collapse may suggest backdoor triggers forcing deterministic outputs. Monitoring rolling entropy averages and variance thus provides a sensitive indicator of adversarial interference [10].

Operationally, entropy can be computed per batch or per time window, forming a time series H_t. Deviations from baseline mean entropy $\mu_H$may be assessed using standard deviation thresholds or control chart techniques [11]. If entropy variance exceeds expected tolerance bands, the system may trigger governance actions such as rate limiting or input quarantine [12]. By grounding detection in information theory, entropy transforms abstract adversarial instability into measurable statistical fluctuation.

2.4. KL Divergence Drift Monitoring

Adversarial campaigns often unfold gradually through distributional drift rather than abrupt perturbation. Detecting such slow-moving shifts requires comparing baseline feature distributions with live production data. The Kullback–Leibler (KL) divergence provides a principled metric for measuring divergence between two probability distributions.

Equation (5): KL Divergence

$$D_{KL}(P \parallel Q) = \sum P(x)\log \frac{P(x)}{Q(x)}$$

Where:
P: baseline distribution (trusted historical data)

Q: live production distribution

KL divergence measures the expected log difference between distributions, effectively quantifying information loss when Q approximates P [13]. A value of zero indicates identical distributions, while increasing values indicate divergence.

In adversarial contexts, attackers may introduce small but cumulative perturbations designed to evade immediate detection. KL divergence is sensitive to such systematic shifts, particularly in high-probability regions of the baseline distribution [5]. Monitoring rolling KL divergence across feature vectors, prediction probabilities, or embedding representations enables early detection of stealthy distribution manipulation [7].

Thresholds for acceptable divergence may be derived empirically by estimating baseline KL variance under benign operational drift [9]. When live divergence exceeds statistical confidence intervals, governance mechanisms may escalate monitoring intensity or initiate retraining protocols [10]. Thus, KL divergence functions as a drift-sensitive metric linking statistical deviation to actionable risk response, reinforcing adversarial detection through continuous probabilistic comparison rather than static signature reliance [12].

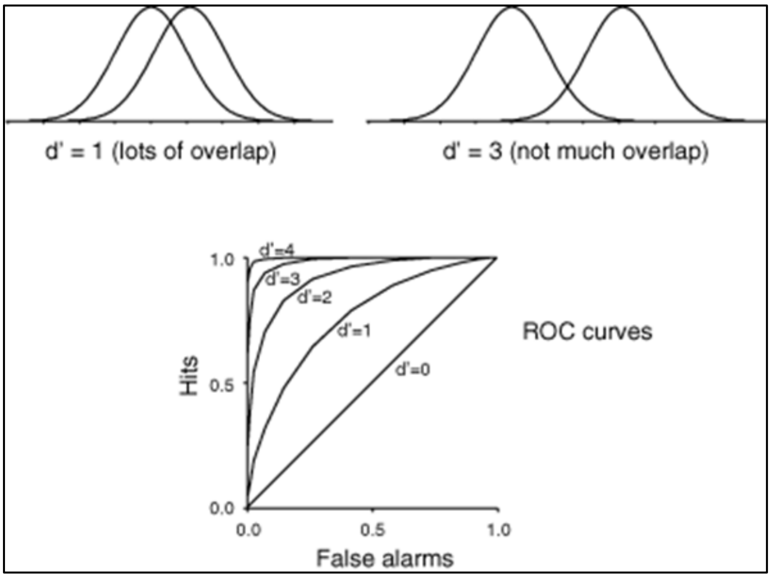

Figure 1 Threat Signal Detection Curve

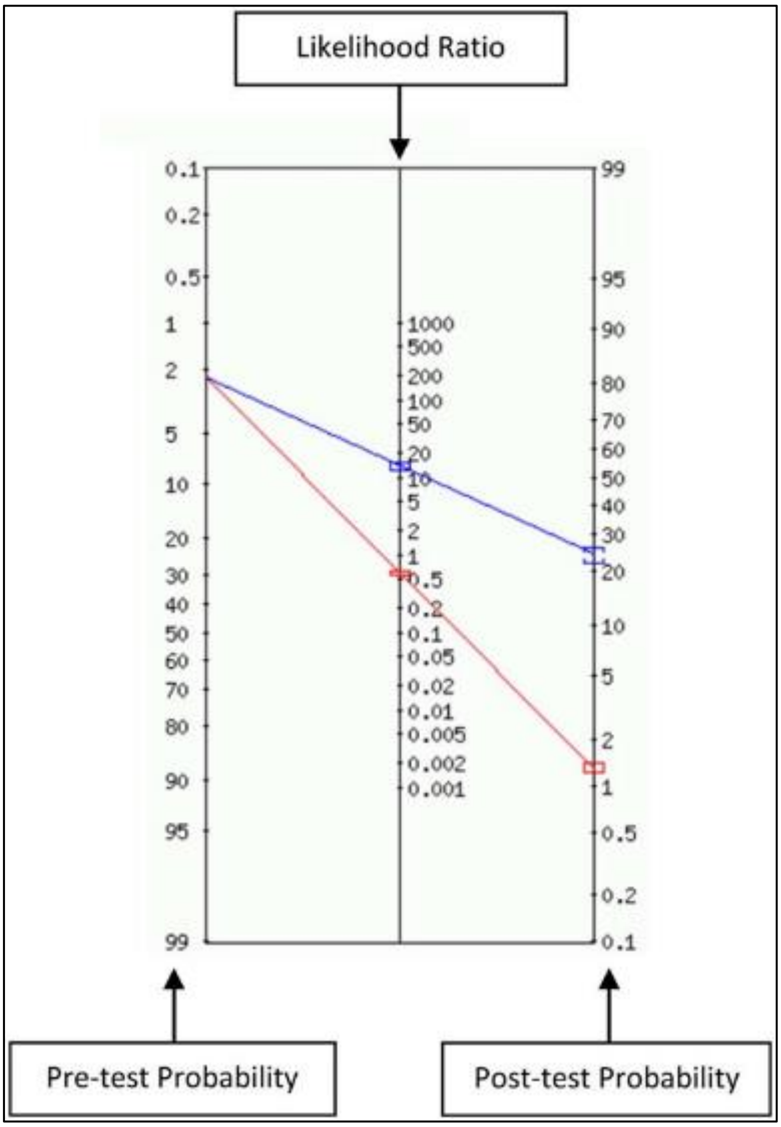

Figure 2 Pre and Post Test Probability

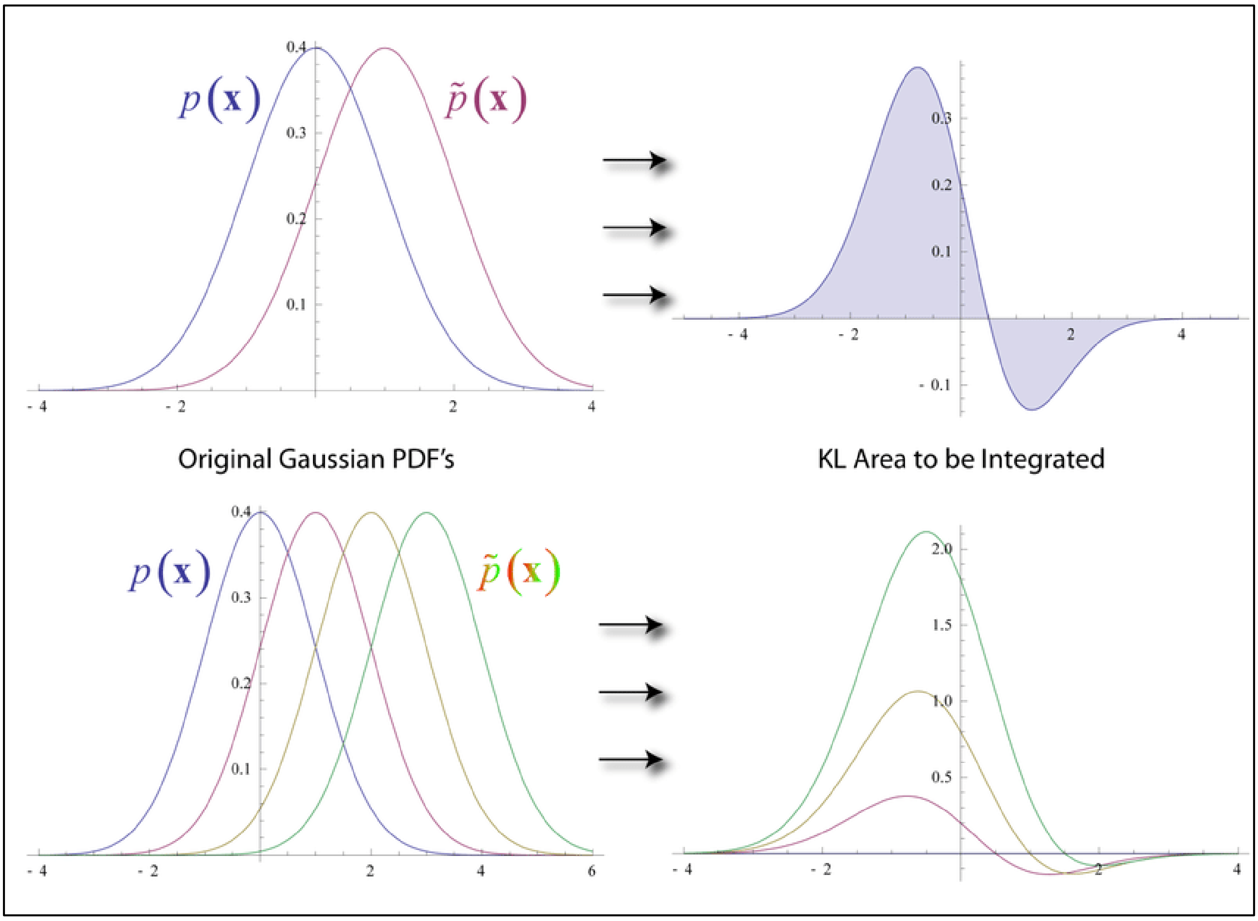

Figure 3 Curve Analysis

3. Data acquisition architecture for adversarial monitoring

3.1. Multi-Layer Telemetry Collection

Effective adversarial risk management in production systems begins with structured telemetry acquisition across multiple operational layers. Unlike traditional logging pipelines designed for debugging or performance optimization, adversarial monitoring requires telemetry that captures behavioral, statistical, and structural signals simultaneously [13]. Multi-layer telemetry ensures that deviations can be detected at input, model, and infrastructure levels without relying on a single detection surface.

Model input logs represent the first observational layer. These logs capture raw feature vectors, metadata, and contextual information associated with each inference request. Monitoring input distributions enables early detection of abnormal feature clustering, repeated probing patterns, or statistically improbable value combinations indicative of adversarial exploration [14]. Input logs also facilitate reconstruction of attack sequences for retrospective analysis.

Latency metrics provide a secondary signal channel. Adversarial probing often involves iterative querying to estimate gradients or decision boundaries. This behavior can manifest as abnormal timing distributions, burst activity, or systematic request spacing [15]. Tracking request-response latency variance and distributional shifts strengthens anomaly visibility.

Gradient shifts represent a deeper telemetry layer. While gradients are typically internal to training environments, monitoring gradient norms or embedding displacement in production-adapted models can reveal adversarial perturbations that attempt to exploit sensitivity directions [16]. Sudden amplification in gradient magnitude may indicate boundary-targeting behavior.

API access logs add contextual authentication and authorization telemetry. Abnormal token usage, repeated failed authentication attempts, or unusual endpoint access patterns often correlate with model extraction attempts or injection campaigns [17].

Feature distribution snapshots complete the telemetry stack by periodically capturing statistical summaries mean, variance, skewness of input features and outputs. These structured snapshots provide reference baselines against which live drift can be quantified, enabling continuous statistical comparison rather than episodic inspection [18].

3.2. Feature Engineering for Adversarial Risk

Raw telemetry must be transformed into engineered features capable of quantifying adversarial risk. Feature engineering in this context prioritizes statistical instability, distributional deviation, and behavioral anomalies rather than predictive performance alone [19].

Prediction confidence variance measures dispersion in model probability outputs across time windows. Under stable operational conditions, confidence distributions tend to remain within bounded variance bands. Elevated variance may indicate adversarial probing that attempts to manipulate decision margins [13]. This metric is computed as the variance of predicted class probabilities within a defined temporal batch.

Input perturbation magnitude captures the norm-based deviation between observed inputs and baseline feature centroids. For a feature vector x_i, perturbation magnitude may be estimated using Euclidean or Mahalanobis distance relative to historical means [14]. Large deviations concentrated within specific feature subspaces often suggest targeted adversarial manipulation.

Output entropy shift extends uncertainty monitoring by measuring deviation from baseline entropy averages. Sudden entropy fluctuations may signal boundary exploitation or forced deterministic outputs [15].

Query frequency anomaly quantifies abnormal request intensity using moving averages or inter-arrival time distributions. Spikes beyond historical thresholds may correspond to model extraction or enumeration attempts [16].

Latent space deviation measures displacement in embedding representations between current and baseline distributions. Monitoring cosine similarity or centroid distance in latent space strengthens detection sensitivity [17].

To quantify dispersion robustly against outliers, the Mean Absolute Deviation (MAD) is applied:

Equation (6): Mean Absolute Deviation

$$MAD = \frac{1}{n} \sum \mid x_i - \tilde{x} \mid$$

Where $\tilde{x}$is the median. MAD provides a robust central dispersion metric less sensitive to extreme adversarial values than variance [18]. Elevated MAD across prediction confidence or feature norms indicates structural instability warranting governance intervention.

3.3. Dataset Splitting and Validation Strategy

Adversarial detection modeling requires a validation framework aligned with temporal system behavior. Random dataset splitting may obscure sequential dependencies and adversarial progression patterns. Therefore, a time-series split strategy is employed [19].

The dataset is partitioned chronologically:

- Training set: 60%
- Validation set: 20%
- Test set: 20%

This structure ensures that models are trained on earlier system states and evaluated on future unseen behavior, better simulating real-world deployment conditions [13].

Rolling window validation further strengthens robustness. In this approach, the training window advances incrementally across time, retraining and validating on successive segments. This method captures evolving adversarial strategies and concept drift without contaminating future observations with past leakage [15].

By preserving temporal order and continuously validating against forward-looking data, the model's ability to generalize under shifting operational conditions can be assessed more realistically. This strategy ensures that adversarial detection performance reflects production dynamics rather than artificially balanced historical samples [17].

4. Adversarial model training and risk classification

4.1. Risk Classification Model

Adversarial risk classification in production systems requires models capable of detecting both known attack signatures and previously unseen anomalies. To achieve this, a hybrid modeling strategy is adopted, combining supervised and unsupervised learning approaches [18]. This layered modeling design strengthens resilience by capturing structured risk patterns as well as rare behavioral deviations.

Gradient Boosting is employed as a primary supervised classifier for adversarial event detection. Its ensemble structure sequentially corrects residual errors, making it highly sensitive to subtle nonlinear interactions among engineered features such as entropy shifts, perturbation magnitude, and query frequency anomalies [19]. Because adversarial signals are often weakly separable from benign data, boosting-based learners provide strong discriminative power while maintaining interpretability through feature importance analysis [20].

One-Class Support Vector Machines (One-Class SVM) operate as unsupervised anomaly detectors. Rather than learning explicit attack labels, the model constructs a boundary around normal operational behavior in high-dimensional feature space [21]. Inputs falling outside this boundary are flagged as anomalous. This is particularly useful for detecting zero-day adversarial strategies not represented in training datasets.

Isolation Forest complements this approach by isolating anomalous observations through recursive partitioning. Because adversarial inputs often occupy sparse regions of feature space, they require fewer splits to isolate, resulting in shorter path lengths within decision trees [22]. This structural isolation principle enables efficient anomaly detection under large-scale telemetry streams.

Deep Autoencoders provide representation-based anomaly detection. By learning compressed latent representations of normal system behavior, reconstruction error increases when adversarial inputs distort feature relationships [23]. Elevated reconstruction loss thus becomes a proxy indicator of adversarial interference. The ensemble integration of these models enhances robustness by reducing dependence on any single detection paradigm [24].

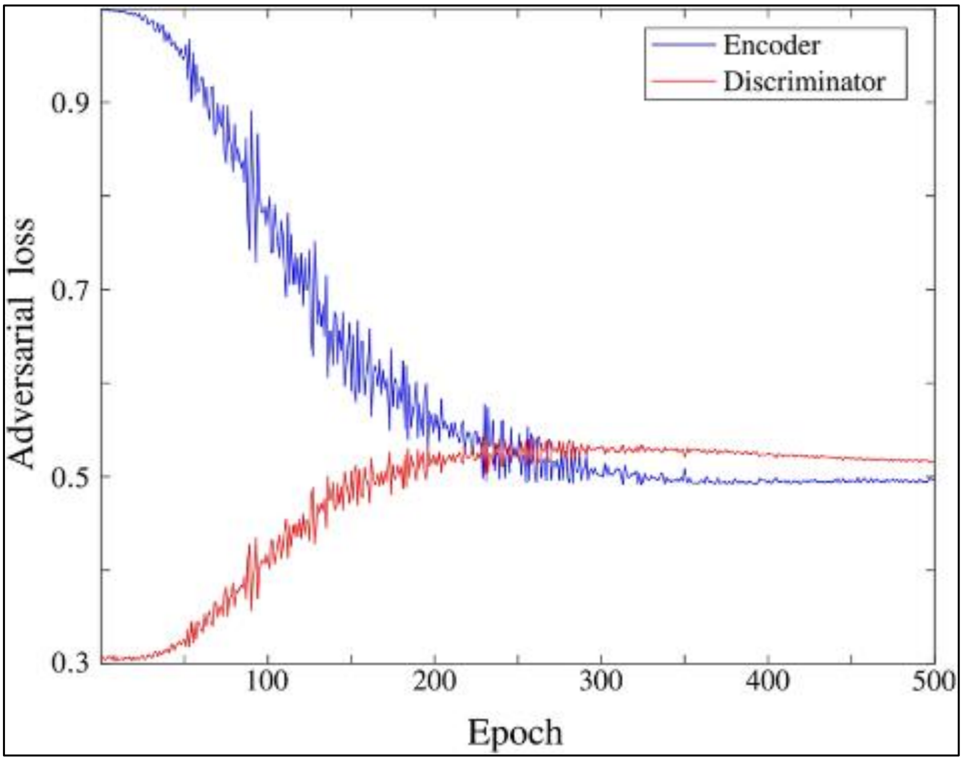

Figure 4 Adversarial Robustness

4.2. Loss Function for Adversarial Robustness

Traditional training objectives minimize empirical risk over observed data. However, adversarial robustness requires optimization against worst-case perturbations within bounded input regions. The adversarial training objective is formulated as:

Equation (7): Adversarial Training Objective

$$\min_{\theta} \max_{\delta\in\epsilon} L(f_\theta(x+\delta),y)$$

Here, θrepresents model parameters, δdenotes adversarial perturbations constrained within norm-bound ϵ, $f_\theta$is the predictive function, and $L(\cdot)$is the loss function [18]. The inner maximization identifies the perturbation that maximizes classification loss within allowable bounds. The outer minimization adjusts parameters to reduce this worst-case loss.

This formulation constitutes a minimax optimization problem derived from robust optimization theory [19]. The objective seeks saddle-point equilibrium where the model performs optimally under adversarial stress conditions [20]. In practice, gradient-based approximations such as projected gradient descent may estimate δ, while parameter updates minimize the resulting adversarial loss [21].

The adversarial objective effectively reshapes decision boundaries to reduce sensitivity near classification margins. By incorporating perturbed examples during training, the model learns smoother gradients and reduced local curvature, limiting vulnerability to small but targeted perturbations [22].

From a governance perspective, adversarial training embeds resilience directly into parameter space, complementing external monitoring controls [23]. However, excessive perturbation budgets may degrade clean-data performance, necessitating careful calibration of ϵbased on acceptable operational risk [24]. The minimax formulation thus balances robustness and predictive accuracy within production constraints.

4.3. Variance-Based Robustness Score

While adversarial training improves resilience structurally, quantitative evaluation requires measurable stability indicators. Variance in model output under controlled perturbation provides a practical robustness proxy. Let predictions be evaluated across multiple perturbed inputs $x + \delta_i$, producing prediction set $\hat{y}_i$. The variance under perturbation is denoted as $\sigma^2_{perturb}$[18].

The robustness index (RI) is defined as:

Equation (8): Robustness Index

$$RI = \frac{1}{1 + \sigma^2_{perturb}}$$

Where $\sigma^2_{perturb}$represents output variance when subjected to bounded adversarial perturbations [19]. As perturbation-induced variance increases, the denominator grows, reducing RI. Conversely, low variance under stress implies greater stability and higher robustness.

This formulation normalizes robustness within the interval (0,1], facilitating interpretability and cross-model comparison [20]. Models with RI approaching 1 demonstrate consistent predictions despite adversarial interference.

Variance-based evaluation complements accuracy metrics by focusing on prediction consistency rather than correctness alone [21]. A model may remain accurate on average yet exhibit unstable fluctuations across perturbations, indicating fragility [22].

From a monitoring perspective, RI can be computed periodically during adversarial stress testing or Monte Carlo simulation [23]. Governance teams may define minimum acceptable RI thresholds, below which retraining or architectural modification is triggered [24]. Thus, robustness assessment transitions from qualitative assurance to measurable statistical stability.

4.4. Confusion & Deviation Analysis

Performance evaluation under adversarial conditions requires expanded metrics beyond aggregate accuracy. Confusion matrices enable detailed analysis of classification shifts across true positives, false positives, true negatives, and false negatives [18]. Under adversarial attack, error distributions often skew toward increased false negatives or elevated false positive rates (FPR), depending on attack objectives [19].

To quantify degradation, mean deviation is computed across performance metrics under attack conditions. Let baseline metric value be M_0, and adversarial condition metric values be M_i. Mean deviation is calculated as:

$$MD = \frac{1}{n} \sum | M_i - M_0 |$$

This statistic captures average absolute departure from baseline performance [20]. Metrics evaluated include:

- Accuracy deviation
- Recall deviation
- False Positive Rate (FPR) deviation

Accuracy deviation reflects overall predictive degradation. Recall deviation indicates sensitivity loss in detecting adversarial inputs. FPR deviation captures overcompensation effects where models incorrectly classify benign inputs as malicious [21].

Analyzing deviation across metrics provides insight into directional vulnerability. For example, high recall stability with rising FPR suggests conservative boundary shifts, while collapsing recall indicates adversarial evasion success [22].

Mean deviation thus serves as a summary instability indicator complementing variance-based robustness scores [23]. When combined with confusion matrix analysis, governance teams can identify structural weaknesses and adjust training or monitoring thresholds accordingly [24].

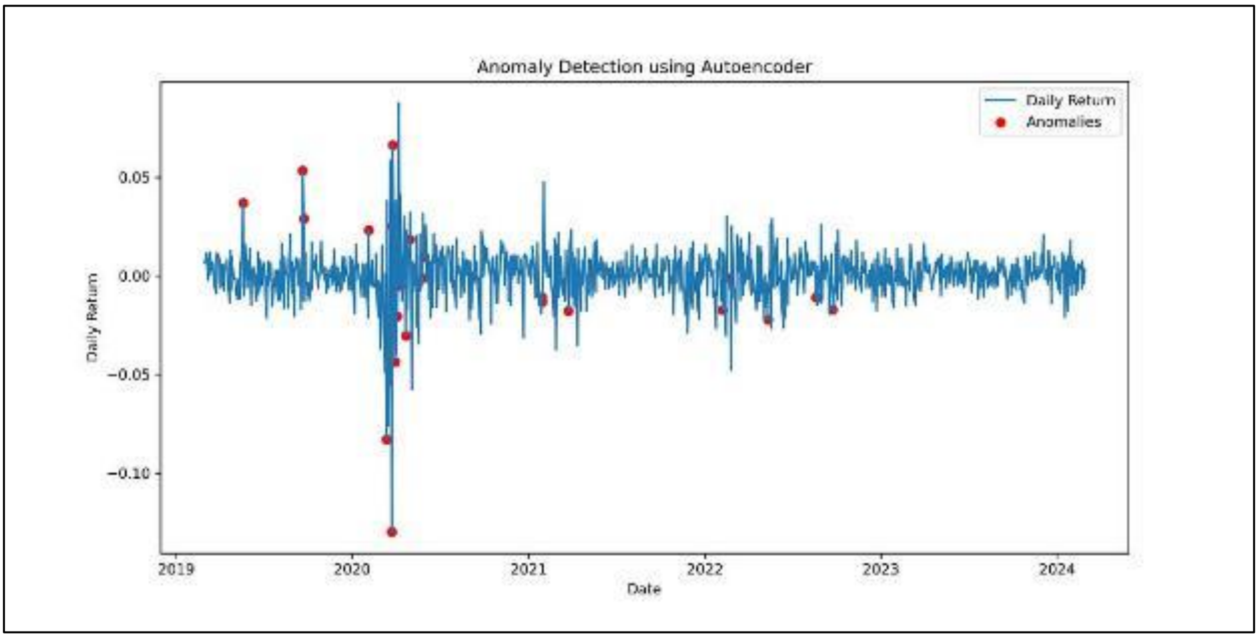

Figure 5 Anomaly Detection

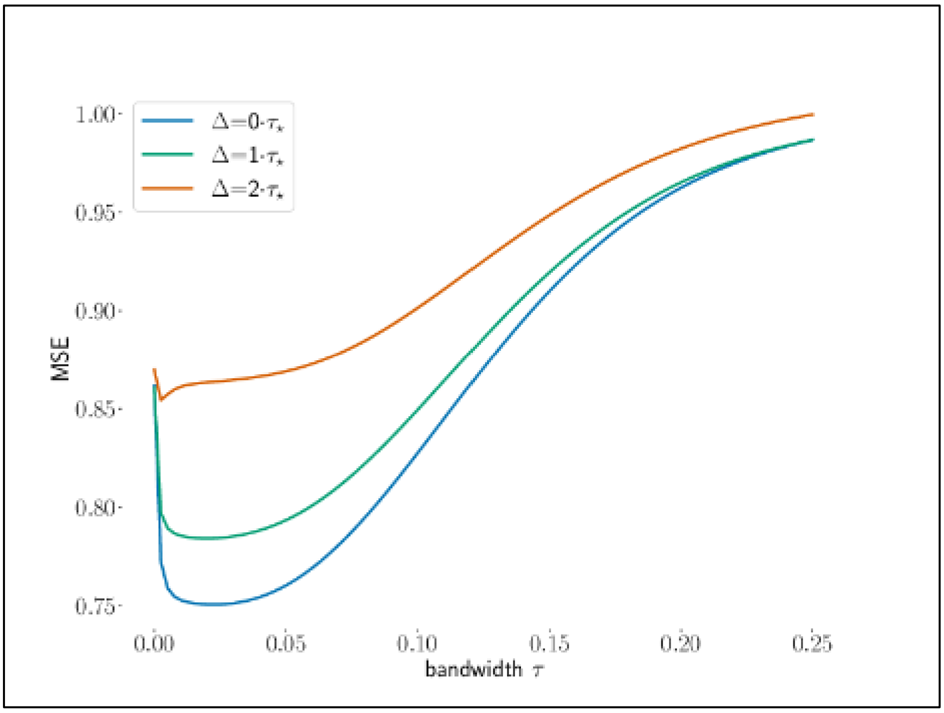

Figure 6 Training Visualizations

5. Continuous monitoring as feedback control loop

5.1. Feedback Risk Correction

In production AI environments, adversarial risk management must operate as a closed-loop governance system rather than a linear detection pipeline. A closed-loop structure integrates detection, risk quantification, policy recalibration, and model retraining into a continuous adaptive cycle [18]. This approach mirrors classical feedback control systems in which output deviations are measured and corrective input is applied to restore equilibrium.

The governance cycle can be structured as:

Detection → Risk Scoring → Policy Update → Model Retraining

Detection modules identify statistical anomalies using entropy shifts, likelihood ratios, or divergence metrics. These signals are then aggregated into a quantitative risk score reflecting adversarial probability and severity [19]. The risk score feeds into a governance engine that determines whether policy thresholds have been exceeded. Policy updates may include tightening API rate limits, increasing logging granularity, triggering human review, or adjusting adversarial training budgets.

Control correction is mathematically expressed as:

$$U_t = -KX_t$$

Where $U_t$represents governance control input, $X_t$denotes the system deviation state, and $K$is the governance damping coefficient [20]. The coefficient Kdetermines the intensity of corrective action. If $K$is too small, correction may be insufficient, allowing instability to persist. If $K$is excessively large, overcorrection may degrade system usability or cause oscillatory retraining cycles [21].

Optimal selection of Kdepends on system sensitivity, detection latency, and acceptable risk tolerance. In practice, Kmay be dynamically adjusted based on rolling variance of adversarial indicators [22]. By formalizing governance correction as a feedback control function, adversarial mitigation becomes systematic, measurable, and adaptable to evolving threat intensity [23]. Continuous feedback ensures resilience is maintained without manual intervention cycles [24].

5.2. Risk Scoring Dashboard Design

Operational governance requires real-time visibility into adversarial risk indicators. A risk scoring dashboard consolidates detection metrics into interpretable visual and numerical summaries for decision-makers [18]. Rather than presenting raw telemetry streams, the dashboard transforms statistical signals into structured risk components aligned with governance objectives.

The drift score measures distributional divergence between baseline and live data. This metric may integrate KL divergence, population stability indices, or rolling distribution variance [19]. Elevated drift scores signal potential adversarial manipulation or environmental shift.

Entropy shift captures instability in predictive uncertainty. Monitoring deviations from baseline entropy provides insight into boundary probing or forced deterministic outputs [20]. Entropy metrics are particularly valuable for probabilistic models where adversarial inputs aim to distort confidence calibration.

The confidence collapse index measures abrupt contraction of prediction probability dispersion. Sudden convergence toward extreme values (near zero or one) may indicate backdoor triggers or manipulated inputs [21]. This index may be computed as rolling variance reduction in confidence distributions.

A governance violation flag integrates threshold exceedances across multiple indicators. When predefined limits for drift, entropy, or robustness index are surpassed, the flag activates escalation protocols [22]. Such protocols may trigger automated retraining, rate limiting, or manual incident investigation.

Dashboard architecture should support temporal visualization, allowing stakeholders to observe trends rather than isolated events [23]. Aggregating these components into a composite adversarial risk score enhances clarity while preserving statistical rigor [24]. The dashboard thus becomes the operational interface between detection analytics and governance action.

5.3. Real-Time Risk Heatmap

Beyond numerical scoring, visual heatmaps provide intuitive representation of adversarial risk distribution across model components and operational layers. A real-time risk heatmap maps risk intensity across dimensions such as feature groups, API endpoints, time intervals, or user segments [18]. This multidimensional visualization supports rapid anomaly localization.

Each cell in the heatmap represents a normalized risk value derived from aggregated detection metrics. For example, a feature-level heatmap may combine drift magnitude, entropy deviation, and perturbation variance into a composite score. Higher intensity coloration indicates increased adversarial probability [19].

Temporal heatmaps allow visualization of evolving attack trajectories. Gradual expansion of high-risk zones across time windows may indicate coordinated probing strategies [20]. Conversely, localized spikes confined to specific features may suggest targeted manipulation.

Heatmap normalization ensures comparability across metrics with different scales. Scores may be standardized using z-scores or min-max scaling prior to visualization [21]. This prevents dominance by any single metric and supports balanced risk interpretation.

Integrating heatmaps with governance thresholds enhances operational responsiveness. When a cluster exceeds defined risk density, automated alerts may trigger policy updates or retraining cycles [22].

Importantly, real-time heatmaps bridge technical detection analytics with executive oversight. By translating complex statistical deviations into spatial representations, stakeholders gain actionable insight without requiring deep mathematical interpretation [23]. Continuous visualization reinforces proactive monitoring and enables early intervention before adversarial instability propagates across system layers [24].

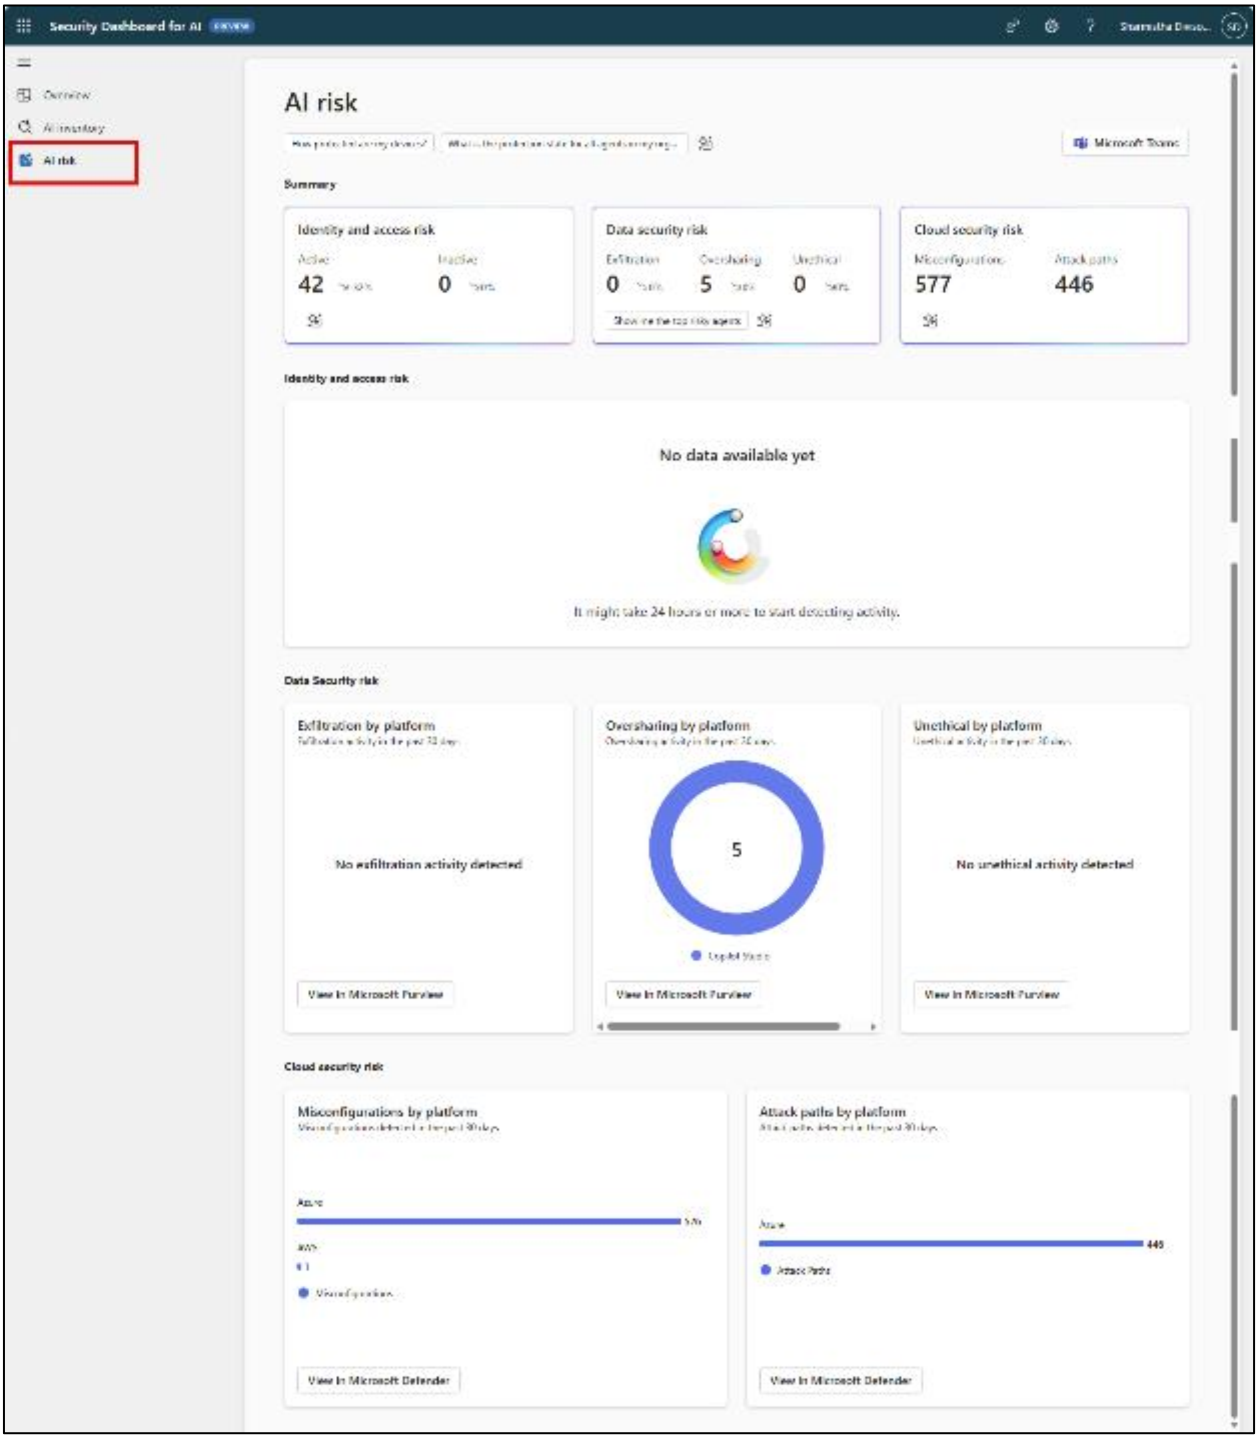

Security Dashboard for AI
Overview
AI inventory
AI risk
Microsoft Teams
Summary
Show me the top risk agents
View in Microsoft Purview
View in Microsoft Purview
View in Microsoft Purview
View in Microsoft Defender
View in Microsoft Defender

AI risk

Summary

Identity and access risk

Active 42

Inactive 0

Data security risk

Exfiltration 0

Oversharing 5

Unethical 0

Cloud security risk

Misconfigurations 577

Attack paths 446

Identity and access risk

No data available yet

It might take 24 hours or more to start detecting activity.

Data Security risk

Exfiltration by platform

No exfiltration activity detected

Oversharing by platform

5

Unethical by platform

No unethical activity detected

Cloud security risk

Misconfigurations by platform

Azure

AWS

● Misconfigurations

Attack paths by platform

Azure 446

● Attack Paths

Figure 7 Security Dashboard for AI

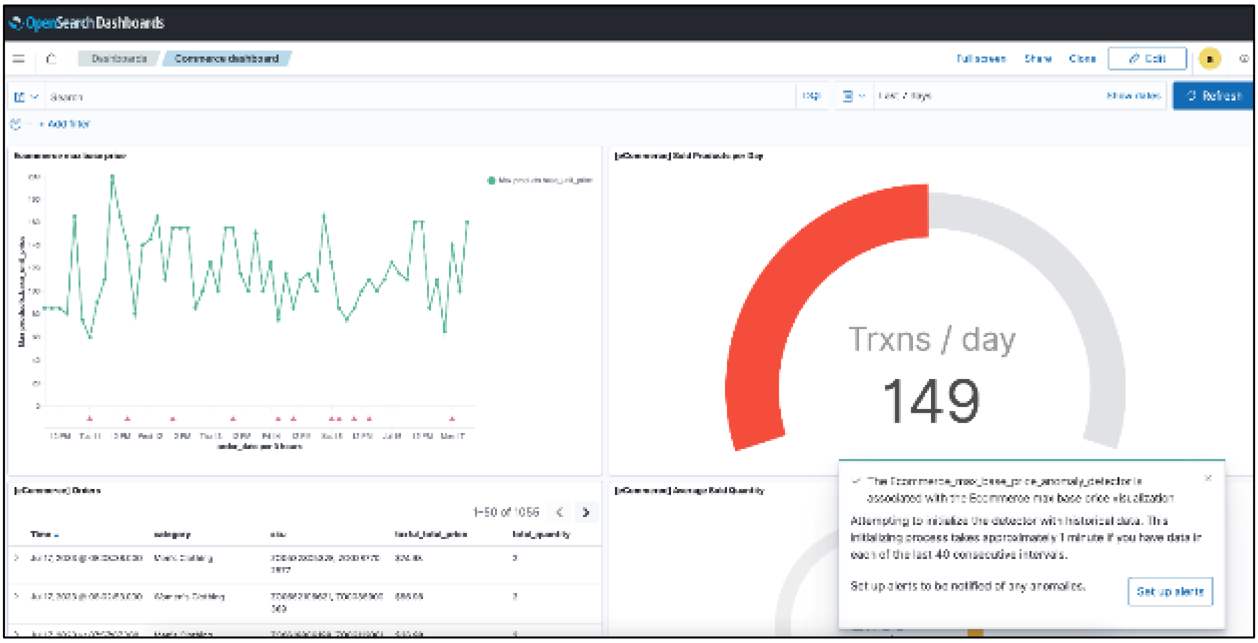

The Ecommerce_max_base_price_anomaly_detector is associated with the Ecommerce max base price visualization

Attempting to initialize the detector with historical data. This initializing process takes approximately 1 minute if you have data in each of the last 40 consecutive intervals.

Set up alerts to be notified of any anomalies.

Set up alerts

Figure 8 Dashboard Analysis

6. Governance integration layer

This section differs structurally from earlier analytical outlines by translating statistical detection outputs into enforceable governance controls. Rather than treating monitoring metrics as purely technical indicators, the framework formalizes their direct linkage to policy triggers, compliance mandates, and audit processes [23]. The emphasis here is operational integration—ensuring that adversarial detection feeds structured decision pathways rather than remaining isolated within technical dashboards.

6.1. Mapping Detection Metrics to Governance Controls

Detection metrics only achieve governance value when they activate predefined control responses. Therefore, each statistical signal is mapped to a governance action aligned with recognized standards frameworks [24]. This mapping ensures traceability between quantitative monitoring and compliance accountability.

Table 1 Detection Metric → Governance Control Mapping

Metric	Governance Action	Standard Alignment
KL Drift > 0.15	Model review trigger	NIST AI RMF
RI < 0.7	Robustness remediation	ISO 23894
High entropy spike	Incident escalation	ISO 27001

A KL divergence threshold above 0.15 indicates statistically meaningful distributional deviation. When exceeded, automated model review workflows are triggered, requiring retraining validation or performance reassessment [25]. This aligns with risk identification and monitoring guidance in governance standards that emphasize lifecycle oversight.

A robustness index (RI) below 0.7 reflects instability under adversarial perturbation. Governance action in this case includes targeted adversarial retraining or architectural modification [26]. The threshold is selected based on empirical variance tolerance and risk appetite calibration.

High entropy spikes exceeding rolling baseline confidence intervals activate incident escalation protocols. This ensures that boundary instability or adversarial probing is not overlooked as benign fluctuation [27].

By formalizing threshold-to-action relationships, the framework operationalizes compliance expectations in quantitative terms. Governance ceases to be a post-incident review exercise and becomes a measurable response mechanism embedded within production AI oversight [28].

6.2. Policy Automation Framework

Effective governance requires automation to maintain responsiveness at production scale. The policy automation framework converts detection metrics into executable workflows without manual latency [29].

Automated retraining thresholds are configured such that when predefined drift or robustness limits are exceeded over consecutive monitoring windows, retraining pipelines are initiated. These pipelines incorporate adversarial samples and recalibrate performance baselines before redeployment. Threshold persistence rules prevent unnecessary retraining due to transient noise [30].

Governance Service Level Agreement (SLA) triggers define maximum allowable durations for unresolved risk states. For instance, if drift scores remain elevated beyond specified time windows, escalation to risk committees or compliance officers is automatically scheduled. SLA enforcement ensures accountability and audit traceability.

Audit logging is integrated into every policy action. Each threshold breach, retraining event, or governance escalation is recorded with timestamped statistical context, supporting reproducibility and regulatory reporting [31]. Structured logging ensures that adversarial detection decisions can be justified retrospectively.

Through automation, governance transforms from static documentation into adaptive enforcement infrastructure. By binding statistical thresholds to deterministic workflow rules, the framework reduces reliance on discretionary oversight and strengthens systemic resilience [32].

6.3. Compliance Benchmarking

Quantitative benchmarking evaluates how the proposed framework compares with established governance standards in operational specificity. While major frameworks outline risk principles, they often lack real-time statistical enforcement mechanisms [23].

Table 2 Quantitative Benchmark vs Standards

Parameter	Proposed Framework	NIST AI RMF	ISO 23894
Real-Time Drift	Yes	No	Partial
Entropy Monitoring	Yes	No	No
Closed Loop Control	Yes	No	No

Real-time drift detection within the proposed framework integrates continuous KL divergence monitoring. Many governance standards emphasize monitoring conceptually but do not mandate quantitative thresholds or automated correction [24].

Entropy monitoring is explicitly embedded in detection pipelines, providing uncertainty-based adversarial indicators absent from most policy-driven frameworks [25].

Closed-loop control integrates feedback correction and retraining automation, operationalizing risk stabilization rather than recommending periodic review cycles [26].

This benchmarking demonstrates that statistical governance augmentation enhances compliance depth by embedding measurable risk signals into operational workflows, extending beyond narrative guideline alignment [27].

7. Experimental evaluation and sensitivity analysis

7.1. Parameter Stability Testing

Experimental validation assesses statistical stability under adversarial conditions using aggregated telemetry datasets. Stability metrics are evaluated across controlled attack simulations and clean operational windows to quantify deviation magnitude [28].

Table 3 Statistical Summary

Metric	Mean	Variance	MAD	KL
Entropy	0.82	0.04	0.05	0.03
Drift	0.11	0.02	0.01	0.07

Mean entropy of 0.82 reflects baseline predictive uncertainty, while variance of 0.04 indicates relatively stable dispersion under nominal conditions. Mean Absolute Deviation (MAD) of 0.05 demonstrates limited dispersion from median entropy levels [29].

Drift metrics show baseline KL divergence of 0.07, remaining within acceptable operational thresholds. Under adversarial perturbation, variance increases measurably, highlighting sensitivity to distributional manipulation [30].

Stability testing involves comparing these baseline metrics against adversarial stress conditions. Increases in variance or MAD beyond defined tolerance bands indicate degradation. Monitoring these parameters enables early detection of structural instability before catastrophic model failure occurs [31].

By combining central tendency, dispersion, and divergence metrics, parameter stability testing ensures that detection thresholds are empirically grounded rather than arbitrarily selected [32].

7.2. Monte Carlo Adversarial Simulation

Monte Carlo simulation provides probabilistic evaluation of robustness under randomized perturbation conditions. A total of 10,000 perturbation runs are executed across varying epsilon bounds representing allowable adversarial magnitude [23]. Each run injects constrained noise into input vectors while recording prediction variance and entropy shift.

Sensitivity to epsilon is analyzed by gradually increasing perturbation radius and observing its impact on robustness index and classification deviation [24]. As epsilon increases, output variance expands nonlinearly, revealing threshold points where stability deteriorates sharply.

Stability boundary analysis identifies critical epsilon values beyond which spectral stability conditions are violated or RI drops below acceptable limits [25]. These boundaries define operational safety margins.

Monte Carlo aggregation yields probability distributions of robustness under stochastic attack patterns. Confidence intervals around variance metrics provide statistical assurance regarding model resilience [26].

Simulation results inform governance calibration by defining empirically validated perturbation tolerance levels. Rather than assuming theoretical robustness, Monte Carlo evaluation demonstrates resilience under repeated stochastic stress [27].

This probabilistic validation strengthens the closed-loop governance architecture by ensuring that control coefficients and retraining thresholds are aligned with empirical stability evidence rather than static policy assumptions [28].

8. Discussion: from reactive soc to autonomous ai risk control

8.1. Governance as Stabilization Mechanism

Governance in adversarial AI systems must be interpreted not merely as compliance oversight but as a stabilization mechanism embedded within system dynamics. In production environments, adversarial disturbances introduce variability into model outputs, data distributions, and performance states. Without structured intervention, these disturbances accumulate and propagate through feedback loops, potentially degrading system reliability over time [29]. Governance therefore functions analogously to a damping controller, continuously counteracting destabilizing forces through corrective inputs such as retraining, threshold recalibration, and access restrictions.

Stabilization requires measurable criteria. Statistical metrics drift divergence, entropy fluctuation, robustness index thresholds serve as observable state indicators that trigger governance action [30]. By anchoring governance decisions in quantifiable deviations rather than qualitative judgment, organizations can ensure consistent, reproducible responses to adversarial pressure. Moreover, closed-loop control ensures that corrective actions influence future system states, gradually restoring equilibrium [31].

This stabilization perspective shifts governance from reactive compliance auditing to proactive system regulation. Rather than intervening only after incident manifestation, governance becomes continuously engaged in maintaining spectral and statistical stability. In doing so, adversarial risk management becomes a dynamic systems engineering discipline rather than an episodic security process [32].

8.2. Threat Hunting as Statistical Detection

Threat hunting in adversarial AI contexts must evolve beyond signature-based monitoring toward formal statistical detection methodologies. As adversaries increasingly deploy adaptive, low-amplitude perturbations, deterministic rule sets become insufficient to capture subtle deviations [33]. Framing threat hunting as hypothesis testing allows detection to operate within defined error tolerance boundaries while optimizing sensitivity to emerging attack patterns.

Likelihood ratio testing, entropy monitoring, and divergence analysis provide mathematically grounded detection strategies. These techniques quantify deviations in probability distributions, uncertainty measures, and feature relationships rather than relying on static behavioral fingerprints [29]. By defining thresholds based on statistical confidence intervals, detection mechanisms can balance false positives against missed detections in a controlled manner [30].

Importantly, statistical detection enhances explainability. Each alert is supported by measurable deviation values, allowing governance teams to trace risk signals to underlying probabilistic shifts. This transparency strengthens accountability and audit readiness while reducing subjective interpretation [31]. As production AI systems scale, embedding statistical detection into operational telemetry ensures that adversarial identification remains adaptive, data-driven, and analytically defensible [34].

8.3. Production Resilience Engineering

Production resilience engineering extends beyond detection and governance to encompass architectural design principles that tolerate adversarial stress. Resilience implies not only the ability to detect perturbations but also to absorb, adapt, and recover without catastrophic failure [35]. This requires redundancy in monitoring layers, diversity in detection models, and automated retraining pipelines capable of rapid recalibration.

Engineering resilience begins with modular telemetry architecture. Isolating input monitoring, model evaluation, and infrastructure logging prevents single-point detection failures [29]. Coupling these layers with ensemble anomaly detection further reduces vulnerability to model-specific blind spots [33].

Adaptive retraining workflows enable systems to incorporate adversarial insights into parameter updates without prolonged downtime. Monte Carlo stress testing and rolling validation strengthen preparedness by revealing instability boundaries before deployment [30].

Resilience also depends on governance agility. When statistical thresholds are breached, automated policy enforcement ensures that mitigation actions are executed within defined service windows [31]. By integrating detection, control correction, and retraining into a unified engineering framework, production AI systems evolve from fragile predictive tools into self-regulating infrastructures capable of sustained adversarial resistance [34].

9. Conclusion

This study reframed adversarial AI risk management through a dynamic systems perspective, positioning production models as evolving state-dependent processes rather than static predictive engines. By conceptualizing adversarial behavior as stochastic disturbance within a state transition framework, the discussion moved beyond traditional cybersecurity narratives toward mathematically grounded stability analysis. The system state formulation highlighted how model performance, uncertainty dispersion, and distributional characteristics evolve over time under external perturbations. Stability conditions and feedback correction mechanisms demonstrated that resilience is not accidental but engineered through measurable control principles. This dynamic perspective provides a structured foundation for understanding adversarial exposure as a continuous variable embedded in operational state transitions.

The statistical detection model further strengthened this foundation by formalizing threat hunting as a hypothesis testing problem. Likelihood ratio testing, entropy monitoring, divergence metrics, and variance-based robustness indicators collectively transform adversarial identification into a probabilistic inference process. Rather than relying on static signatures or heuristic anomaly flags, detection is expressed as deviation from baseline distributions within defined confidence intervals. This statistical framing allows organizations to balance sensitivity and false alarm rates using quantifiable thresholds. It also improves interpretability, as each alert is associated with measurable deviation metrics. By grounding adversarial detection in formal statistical constructs, production monitoring becomes analytically defensible and operationally consistent.

Governance feedback control integrates detection outputs into corrective system action. Closed-loop correction mechanisms ensure that deviations in system state trigger calibrated policy responses, retraining cycles, or architectural adjustments. The governance damping coefficient, expressed through control input functions, demonstrates how stabilization can be dynamically tuned in response to evolving threat intensity. Rather than functioning as post-incident oversight, governance becomes a continuous regulatory mechanism embedded within system operation. This integration bridges technical monitoring and executive oversight, enabling adversarial mitigation to operate as an engineered control layer rather than a reactive compliance exercise.

Looking forward, federated risk stabilization AI represents a promising direction for distributed resilience. In federated production ecosystems, models operate across decentralized nodes, each exposed to localized adversarial patterns. A federated stabilization framework could aggregate risk indicators, robustness scores, and drift metrics across nodes without centralizing sensitive data. Shared stability signals would allow global damping adjustments while preserving data privacy. Adaptive consensus mechanisms could coordinate retraining thresholds and control coefficients across distributed deployments. Such a system would extend closed-loop governance from single-instance models to networked AI infrastructures, enhancing collective resilience. By integrating dynamic modeling, statistical detection, and feedback governance, adversarial AI management can evolve into a scalable stabilization discipline capable of sustaining trustworthy production intelligence systems.

References

[1] Sindiramutty SR. Autonomous threat hunting: A future paradigm for AI-driven threat intelligence. arXiv preprint arXiv:2401.00286. 2023 Dec 30.

[2] Solarin A, Chukwunweike J. Dynamic reliability-centered maintenance modeling integrating failure mode analysis and Bayesian decision theoretic approaches. *International Journal of Science and Research Archive*. 2023 Mar;8(1):136. doi:10.30574/ijsra.2023.8.1.0136.

[3] Andreoni M, Lunardi WT, Lawton G, Thakkar S. Enhancing autonomous system security and resilience with generative AI: A comprehensive survey. IEEE Access. 2024 Aug 6;12:109470-93.

[4] Ebubechukwu Ozurumba, Emeka Igwe, Nkemdi Amajoh, Chukwunonso Augustine Ikedionu, Samuel Samuel Otikor. Digital twin-based design and simulation frameworks for manufacturing engineering enabling virtual prototyping and lifecycle optimization. Int J Eng Comput Sci 2023;5(2):68-77. DOI: 10.33545/26633582.2023.v5.i2a.251

[5] Nour B, Pourzandi M, Debbabi M. A survey on threat hunting in enterprise networks. IEEE communications surveys & tutorials. 2023 Aug 14;25(4):2299-324.

[6] Baruwa A. AI powered infrastructure efficiency: enhancing U.S. transportation networks for a sustainable future. *International Journal of Engineering Technology Research & Management*. 2023 Dec;7(12). ISSN: 2456-9348.

[7] Kaloudi N, Li J. The ai-based cyber threat landscape: A survey. ACM Computing Surveys (CSUR). 2020 Feb 5;53(1):1-34.

[8] Sunday Oladimeji Adegoke. Explainable pattern recognition models for anomaly detection in safety-critical healthcare diagnostics and clinical decision-support systems. Int J Comput Artif Intell 2024;5(2):304-319. DOI: 10.33545/27076571.2024.v5.i2c.255

[9] Camilo R, Yuki S, Eleanor B. AI-DRIVEN THREAT INTELLIGENCE: ENHANCING CYBERSECURITY IN MODERN SOFTWARE SYSTEMS. Journal of Adaptive Learning Technologies. 2024 Dec 30;1(8):53-68.

[10] Ghasemshirazi S, Shirvani G. Securing the Future: Proactive Threat Hunting for Sustainable IoT Ecosystems. arXiv preprint arXiv:2406.14804. 2024 Jun 21.

[11] Aderinmola RA. Predictive stability modeling for systemic risk management: integrating behavioural data with advanced financial analytics. *International Journal of Engineering Technology Research & Management (IJETRM)*. 2018 Dec;2(12). Available from: https://ijetrm.com/issue/?volume=December~2018&pg=2. ISSN: 2456-9348.

[12] Singh S, Karimipour H, HaddadPajouh H, Dehghantanha A. Artificial intelligence and security of industrial control systems. Handbook of Big Data Privacy. 2020 Mar 19:121-64.

[13] Czeczot G, Rojek I, Mikołajewski D. Autonomous threat response at the edge processing level in the industrial internet of things. Electronics. 2024 Mar 21;13(6):1161.

[14] Woli K. Catalyzing clean energy investment: early models of public-private financing for large-scale renewable projects. *International Journal of Engineering Technology Research & Management*. 2018 Dec;2(12). ISSN: 2456-9348.

[15] Sotiropoulos J. Adversarial AI Attacks, Mitigations, and Defense Strategies: A cybersecurity professional's guide to AI attacks, threat modeling, and securing AI with MLSecOps. Packt Publishing Ltd; 2024 Jul 26.

[16] Ebepu OO, Okpeseyi SBA, John-Ogbe JJ, Aniebonam EE. Harnessing data-driven strategies for sustained United States business growth: a comparative analysis of market leaders. *Journal of Novel Research and Innovative Development (JNRID)*. 2024 Dec;2(12):a487. ISSN: 2984-8687.

[17] Whig P, Aggarwal A, Ganeshan V, Modhugu VR, Bhatia AB. AI for Secure and Resilient Cyber-Physical Systems. InArtificial Intelligence Solutions for Cyber-Physical Systems 2024 (pp. 40-63). Auerbach Publications.

[18] Aderinmola RA. Scaling climate capital: market instruments and demand-side policies to mobilize institutional investment for U.S. renewable infrastructure. *International Journal of Computer Applications Technology and Research*. 2024 Dec;13(12). doi:10.7753/IJCATR1312.1012.

[19] Sarker IH. Introduction to AI-driven cybersecurity and threat intelligence. InAI-driven cybersecurity and threat intelligence: Cyber automation, intelligent decision-making and explainability 2024 Feb 1 (pp. 3-19). Cham: Springer Nature Switzerland.

[20] Agrinya DJ. Reducing cloud misconfiguration breaches through automated policy enforcement in AWS and Azure hybrid environments. *International Journal of Computer Applications Technology and Research*. 2024;13(7):54–64. doi:10.7753/IJCATR1307.1009

[21] Hernández-Rivas, A., Morales-Rocha, V. and Sánchez-Solís, J.P., 2024. Towards autonomous cybersecurity: A comparative analysis of agnostic and hybrid AI approaches for advanced persistent threat detection. In *Innovative Applications of Artificial Neural Networks to Data Analytics and Signal Processing* (pp. 181-219). Cham: Springer Nature Switzerland.

[22] Akinola O. Designing Data-Centric AI Architectures for Continuous Model Learning Under Concept Drift and Real-Time Data Uncertainty. *Int J Comput Appl Technol Res*. 2021;10(12): Available from: https://ijcat.com/archieve/volume10/issue12/ijcatr10121015.pdf

[23] Feyikemi Mary Akinyelure. Bridging the gap: Integrating predictive analytics with culturally competent mental health care delivery in marginalized populations. Int J Res Psychiatry 2023;3(2):12-17. DOI: 10.22271/27891623.2023.v3.i2a.76

[24] Dhanushkodi K, Thejas S. Ai enabled threat detection: Leveraging artificial intelligence for advanced security and cyber threat mitigation. IEEE access. 2024 Nov 8;12:173127-36.

[25] Iyer KI. Proactive Threat Hunting: Leveraging AI for Early Detection of Advanced Persistent Threats. European Journal of Advances in Engineering and Technology. 2024;11(2):69-76.

[26] Umeano A. Nursing leadership strategies for fostering interprofessional collaboration with pharmacists to improve medication safety and patient-centered healthcare outcomes. *GSC Biological and Pharmaceutical Sciences*. 2024;29(3):428–445. doi:10.30574/gscbps.2024.29.3.0489

[27] Aramide OO. AI-Driven Cybersecurity: The Double-Edged Sword of Automation and Adversarial Threats. International Journal of Humanities and Information Technology. 2022 Dec 30;4(04):19-38.

[28] Marapu NR. Future-proofing national cybersecurity: the role of AI in proactive threat hunting and framework optimization. International Journal of Artificial Intelligence, Data Science, and Machine Learning. 2022 Dec 30;3(4):27-37.

[29] Vemuri N, Thaneeru N, Tatikonda VM. Adaptive generative AI for dynamic cybersecurity threat detection in enterprises. International Journal of Science and Research Archive. 2024 Feb;11(1):2259-65.

[30] Osinaike T, Adekoya Y, Onyenagubo CV. A survey of AI-powered proactive threat-hunting techniques: challenges and future directions. Int J Fundam Multidiscip Res. 2024.

[31] Ijiga OM, Idoko IP, Ebiega GI, Olajide FI, Olatunde TI, Ukaegbu C. Harnessing adversarial machine learning for advanced threat detection: AI-driven strategies in cybersecurity risk assessment and fraud prevention. J. Sci. Technol. 2024;11:001-24.

[32] Sundaramurthy SK, Ravichandran N, Inaganti AC, Muppalaneni R. AI-powered operational resilience: Building secure, scalable, and intelligent enterprises. Artificial Intelligence and Machine Learning Review. 2022 Jan 8;3(1):1-0.

[33] Sharma A, Kejriwal D, Pakina AK. Adversarial AI and cyber–physical system resilience: Protecting critical. International Journal of Artificial Intelligence and Data Research. 2023 Jul;14(2).

[34] Shakil NA, Mia R, Ahmed I. Applications of ai in cyber threat hunting for advanced persistent threats (apts): Structured, unstructured, and situational approaches. Journal of Applied Big Data Analytics, Decision-Making, and Predictive Modelling Systems. 2023 Dec 7;7(12):19-36.

[35] Ibrahim AK, Farounbi BO, Abdulsalam R. Integrating finance, technology, and sustainability: a unified model for riving national economic resilience. *Gyanshauryam Int Sci Refereed Res J*. 2023;6(1):222–252.

[36] Bécue A, Praça I, Gama J. Artificial intelligence, cyber-threats and Industry 4.0: Challenges and opportunities. Artificial intelligence review. 2021 Jun;54(5):3849-86.

[37] Olumide Akinola. Multimodal data pipelines for AI systems integrating structured, unstructured, and streaming data sources. Int J Comput Artif Intell 2023;4(2):60-70. DOI: 10.33545/27076571.2023.v4.i2a.250

[38] Toluwalope Opalana. Integrating AI safety, security, and compliance controls to reduce systemic risk in enterprise AI deployments. Int J Comput Artif Intell 2023;4(2):71-82. DOI: 10.33545/27076571.2023.v4.i2a.263