World Journal of Advanced Research and Reviews

eISSN: 2581-9615 CODEN (USA): WJARAI

Cross Ref DOI: 10.30574/wjarr

Journal homepage: https://wjarr.com/

(Research Article)

Generative Identity Forensics and Trust System (GIFTS): Geodesic anomaly detection and manifold diffusion for cloud identity telemetry

Kingdom Mutala Akugri *, Prince Agbenyo and Marious Akugri

Independent Researcher, USA.

World Journal of Advanced Research and Reviews, 2026, 29(01), 1536-1547

Publication history: Received on 14 December 2025; revised on 25 January 2026; accepted on 27 January 2026

Article DOI: https://doi.org/10.30574/wjarr.2026.29.1.0212

Abstract

The cybersecurity perimeter has shifted from network boundaries to identity-driven control planes, where authentication events, API invocations, and access-policy evaluations constitute the dominant evidence for threat detection in modern cloud infrastructure. While contemporary Identity Threat Detection and Response (ITDR) platforms scale to billions of events, anomaly detection approaches remain primarily discriminative—focusing on statistical rarity or reconstruction error—and often fail to provide actionable explanations of malicious behavior, especially when adversaries suppress audit visibility by disabling or evading logging services. In addition, existing systems largely remain reactive and do not provide native mechanisms to generate realistic, novel attack scenarios for proactive defense testing.

This paper introduces the Generative Identity Forensics and Trust System (GIFTS), a manifold learning and diffusion framework that models valid cloud identity behavior as trajectories on a low-dimensional intrinsic manifold shaped by Identity and Access Management (IAM) permissions, workflow constraints, and temporal dependencies. GIFTS integrates four core modules: (i) semantic log vectorization and sessionization using transformer-based representations for high-cardinality cloud events, (ii) nonlinear dimensionality reduction via Isomap to estimate intrinsic geodesic structure, (iii) a manifold inversion mechanism grounded in the Manifold Decoder principle to map latent coordinates back into interpretable log sequences, and (iv) manifold-constrained diffusion for forensic in-painting during logging blackouts and automated red-team generation. Using benchmark intrusion data and a synthetic CloudTrail generator built from attack-chain templates, we demonstrate that geodesic trust scoring improves separability of identity attacks in the low false-positive regime while latent diffusion enables probabilistic reconstruction of missing forensic traces. This work advances generative security operations by unifying anomaly detection, explainability, reconstruction, and proactive simulation within a geometric trust framework.

Keywords: Cloud Security; Identity Telemetry; Anomaly Detection; Manifold Learning; Diffusion Models; Forensic Reconstruction; Zero Trust; ITDR

1. Introduction

Cloud computing has dissolved traditional enterprise network boundaries, shifting security enforcement from perimeter-centric controls toward identity-centric policy enforcement. In hyperscale platforms such as Amazon Web Services (AWS), Azure, and Google Cloud, security-critical actions are performed via API calls and role-based authorization rather than direct interactive access to machines. Consequently, the primary evidence for compromise increasingly resides in identity and control-plane telemetry: authentication events, role session transitions, access-policy evaluations, administrative API sequences, and abnormal privilege flows.

* Corresponding author: Kingdom Mutala Akugri

Identity-based compromise is difficult to detect because attackers can operate using valid credentials obtained through phishing, session hijacking, token theft, or supply-chain leakage. Once authenticated, adversaries perform actions that are syntactically legitimate yet semantically malicious, including privilege escalation, policy tampering, persistence establishment, stealthy access enumeration, and data exfiltration. Detection in this setting requires understanding behavioral structure and trajectories rather than identifying simple point-wise outliers.

Industry pipelines often combine rule matching and unsupervised anomaly detection (e.g., streaming detectors such as Random Cut Forest) with reconstruction models such as autoencoders [1, 7]. These systems scale well but exhibit three persistent gaps:

- **Explainability gap.** Statistical rarity is not equivalent to maliciousness. Legitimate but infrequent administrative actions produce high alert volumes, increasing analyst fatigue and degrading trust in detection signals; human factors such as cognitive load and security fatigue can directly reduce operational resilience under sustained security pressure [11].
- **Forensic gap.** Sophisticated adversaries reduce visibility by disabling or evading logging services. In cloud environments, events such as stopping audit collection, modifying log delivery pipelines, or abusing regions can create logging blackouts, leaving missing segments in attack trajectories.
- **Generative gap.** Most pipelines remain reactive, lacking native mechanisms to generate high-fidelity attack variations for proactive validation, rule stress testing, and red-team rehearsal.

To address these limitations, we propose Generative Identity Forensics and Trust System (GIFTS), which reframes identity security as a geometric and generative modeling problem. Under the manifold hypothesis, high-dimensional operational telemetry concentrates near a lower-dimensional manifold embedded in ambient space. We hypothesize that valid cloud identity behavior is constrained by IAM permission graphs and workflow logic, forming a structured manifold of permitted transitions. Attacks manifest as geodesic discontinuities, path shortcuts, and abnormal trajectory shapes that violate workflow continuity.

1.1. Research aim and objectives

The primary aim of this research is to design and evaluate a generative and explainable identity forensics system for cloud environments, enabling anomaly detection, forensic reconstruction during logging blackouts, and proactive red-team simulation.

The objectives are to:

- Develop a manifold-based representation of cloud identity telemetry capturing intrinsic workflow geometry.
- Introduce geodesic trust scoring optimized for low false-positive regimes.
- Implement a manifold inversion mechanism based on the Manifold Decoder principle for interpretable log reconstruction.
- Apply manifold-constrained diffusion to (a) in-paint missing segments during audit blackouts and (b) generate attack variants.

1.2. Contributions

This paper makes four contributions:

- A geometric trust modeling framework for identity telemetry that emphasizes workflow continuity over Euclidean rarity.
- A healing distance trust score derived from manifold projection and diffusion denoising.
- A decoder formulation for mapping manifold states back to interpretable log artifacts, grounded in the Manifold Decoder inversion viewpoint [9].
- A practical synthetic CloudTrail generator for structured evaluation of identity attack chains under blackout conditions.

2. Related work

2.1. Identity telemetry, UEBA, and ITDR

Identity telemetry monitoring has evolved into User and Entity Behavior Analytics (UEBA) and Identity Threat Detection and Response (ITDR) programs that analyze authentication, authorization, role assumption, and administrative actions for compromise signals. Standards such as Zero Trust Architecture emphasize continuous verification and identity-centric control enforcement [2]. However, UEBA systems often face high false positives when relying on simple deviation scoring [13, 14].

2.2. Anomaly detection on logs and streams

Isolation Forest [3] and Random Cut Forest (RCF) [7] are widely used in streaming anomaly detection due to scalability. Log parsing approaches such as Drain [8] provide structured templates for downstream modeling. Deep reconstruction methods (autoencoders and VAEs) improve representation learning but can still struggle with interpretability and rare-but-valid administrative behavior.

2.3. Manifold learning for structure discovery

Nonlinear dimensionality reduction methods such as Isomap [4] preserve geodesic distances and reveal intrinsic structure not captured by Euclidean embeddings. UMAP [5] provides visualization and clustering utility, though it is not explicitly geodesic-preserving.

2.4. Diffusion models for generative reconstruction

Diffusion models [6] achieve strong generation and in-painting performance by learning reverse denoising processes. In operational security settings, diffusion is attractive because it can generate multiple plausible reconstructions under uncertainty, aligning naturally with missing evidence scenarios.

2.5. Manifold decoders and nonlinear inversion

A key challenge in manifold pipelines is inversion: mapping from low-dimensional coordinates back into high-dimensional artifacts. The Manifold Decoder framework [9] formalizes learnable inversion from nonlinear embeddings, enabling conditional generation, reconstruction, and explainable synthesis. In GIFTS, the log decoder is treated as a Manifold Decoder specialized to identity telemetry.

3. Threat model and problem formulation

3.1. Threat model overview

We consider a cloud environment in which adversaries obtain access to valid credentials or temporary tokens and perform malicious actions via control-plane APIs. The adversary may:

- Operate under legitimate principals (users/roles) with stolen credentials.
- Abuse role assumption chains and policy modifications to escalate privileges.
- Establish persistence through creation of backdoor roles or access key rotation.
- Enumerate resources and exfiltrate data through storage APIs.
- Suppress or reduce audit visibility via log pipeline disruption (blackout).

Defender assumptions:

- Control-plane telemetry (e.g., CloudTrail) is available for baseline training.
- Some attacks may include partial missing intervals due to blackout events.
- Ground truth labels are limited; models must generalize with weak supervision.

3.2. MITRE ATT&CK-style mapping

Table 1 maps representative cloud identity attack behaviors to observable control-plane manifestations and the corresponding GIFTS detection signals. The intent is not to enumerate all techniques, but to show that GIFTS is aligned with common attacker workflows.

Table 1 Threat behavior mapping to cloud control-plane evidence and GIFTS signals.

Tactic family	Cloud manifestation (examples)	GIFTS signal
Initial access / Credential abuse	Unusual principal sessions, atypical source IP/ASN, impossible travel, new user agents	Geodesic discontinuity in session manifold; rare-but-invalid sequence shape
Privilege escalation	Policy updates, Assume Role anomalies, permission boundary bypass patterns	Geodesic deviation across role-transition edges; healing distance spikes
Persistence	Create Role / Attach Policy backdoors, access key rotation anomalies, long-lived sessions	Trajectory lengthening with abnormal loops; decoder reveals persistence artifacts
Discovery / Enumeration	List Buckets, Describe Instances, Get Caller Identity bursts	High-frequency workflow branch not connected to normal deployment manifolds
Exfiltration	S3 list + get bursts, cross-region transfer signals, abnormal access ordering	Abnormal ordering and burst patterns; manifold projection residual increases
Defense evasion / Logging blackout	Stop Logging, trail deletion, delivery pipeline misconfiguration	Missing intervals in trajectories; diffusion in-painting reconstructs likely events

3.3. Problem formulation

Let each cloud event be a structured record with fields: {event Time, event Name, event Source, user Identity, source I P Address, aws Region, request Parameters, error Code, user Agent}. Events are aggregated into sessions over a window Δ to form vectors $x_k \in \mathbb{R}^D$. The goals are:

- Assign a trust score S(x) that flags malicious identity sessions under low false positive constraints.
- Generate explanations by reconstructing interpretable log artifacts from latent states.
- Reconstruct missing evidence under blackout: infer likely intermediate actions.
- Generate synthetic malicious variants for proactive evaluation.

4. Methodology

4.1. Overview of the GIFTS pipeline

GIFTS consists of four integrated modules:

- Semantic vectorization and sessionization: raw events are normalized and embedded into dense vectors.
- Manifold learning (geodesic embedding): Isomap estimates intrinsic geodesic structure.
- Manifold inversion via neural decoding: a Manifold Decoder maps latent coordinates to log artifacts.
- Manifold diffusion: diffusion sampling enables forensic in-painting and red-team generation.

4.2. Log sources and schema

GIFTS targets identity telemetry and control-plane logs. In AWS, representative sources include CloudTrail management events, selective data events, and optionally supplementary context from flow logs and DNS. For learning, heterogeneous fields are normalized into a consistent schema.

4.2.1. Minimum schema:

- Required: event Time, event Name, event Source, user Identity. type, user Identity. arn, source IP Address, aws Region
- Optional: request Parameters, response Elements, error Code, user Agent

4.3. Semantic tokenization and embedding

Cloud logs contain high-cardinality categorical values (actions, ARNs, principals). One-hot encoding is infeasible. GIFTS uses transformer-based encoding to produce contextual embeddings $e_t \in \mathbb{R}^H$. The embedding pipeline includes:

- Canonicalization of ARNs and identifiers.
- Bucketing or hashing of high-cardinality strings.
- Serialization into stable token order: identity → action → target → context.
- Transformer encoding for contextual representation.

4.4. Sessionization into behavioral state vectors

Identity behaviors are sequences. Embeddings are pooled over a window Δ minutes:

$$x_k = \mathrm{Pool}\big(\{e_t\}_{t \in \mathcal{W}_k}\big).$$

where W_k is the event set in a time window. In operational settings, attention pooling is preferable; mean pooling is used as a baseline ablation.

4.5. Manifold learning using Isomap

Given session vectors $\{x_i\}_{i=1}^{N}$, Isomap constructs a k-nearest neighbor graph and estimates geodesic distances via shortest paths:

$$D_G\big(x_i, x_j\big) = \min_{\gamma} \sum_{(u,v)\in\gamma} \| x_u - x_v \|.$$

Classical multidimensional scaling maps the geodesic distance matrix into latent coordinates:

$$z_i = f_{\mathrm{iso}}(x_i), \quad z_i \in \mathbb{R}^d, \quad d \ll D.$$

4.6. Geodesic trust scoring

Two trust scoring strategies are used:

(A) Local geodesic deviation score

$$S_{\mathrm{geo}}(x') = \min_{x_j \in \mathcal{N}(x')} D_G\big(x', x_j\big).$$

(B) Healing distance score (projection-based)

$$S_{\mathrm{heal}}(x') = \| z' - \hat{z} \|.$$

where $\hat{z}$ is the manifold-consistent reconstruction produced by diffusion denoising and inversion.

4.7. Manifold inversion and the Neural Log Decoder

Explainability requires mapping latent states back into interpretable artifacts. GIFTS uses a neural decoder g_φ to map latent coordinates to log sequences:

$$\hat{L} = g_\phi(z).$$

where $\hat{L}$ is a structured sequence of reconstructed control-plane events.

This module is grounded in the Manifold Decoder principle: even when $f_{\mathrm{iso}}(\cdot)$ is nonlinear and non-invertible analytically, a learnable decoder can approximate a conditional inverse mapping from latent coordinates into high-dimensional event space. Training uses a composite objective:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{CE}} + \lambda \mathcal{L}_{\mathrm{geom}} + \beta \mathcal{L}_{\mathrm{fields}}.$$

where L_{geom} preserves latent neighborhood structure and L_{fields} penalizes malformed reconstructions of critical security fields (principal, action, target, region).

4.8. Manifold diffusion for in-painting and red teaming

Diffusion is performed in latent space:

$$q(z_t \mid z_0) = \mathcal{N}\left(z_t; \sqrt{\bar{\alpha}_t} z_0, (1-\bar{\alpha}_t)I\right).$$

$$p_\theta(z_{t-1} \mid z_t) = \mathcal{N}\left(z_{t-1}; \mu_\theta(z_t, t), \Sigma_\theta(z_t, t)\right).$$

Blackout in-painting. For a blackout interval between T_1 and T_2:

$$z^*_{T_1:T_2} \sim p_\theta\left(z_{T_1:T_2} \mid z_{T_1}, z_{T_2}\right).$$

generating multiple plausible forensic hypotheses.

4.9. Operational deployment (Fast Path vs Deep Path)

GIFTS supports a hierarchical deployment model:

- Fast Path: streaming anomaly screening (lightweight scoring) for broad coverage.
- Deep Path: manifold scoring + decoding + diffusion reconstruction for high-risk sessions.

5. Synthetic CloudTrail generator for identity attack chains

Real enterprise identity attack ground truth is scarce. To evaluate detection and reconstruction, we implement a synthetic CloudTrail generator that produces realistic operational workflows and parameterized attack chains, including logging blackouts.

5.1. Event grammar and fields

Each event generated is a JSON-like record with fields:

- Event Time: timestamp with ordering constraints
- Event Source: service namespace (e.g., iam.amazonaws.com, s3.amazonaws.com)
- Event Name: API action (e.g., Assume Role, Put Role Policy)
- User Identity: {type, arn, session Context}
- source IP Address: internal, corporate, or external bucket
- aws Region: region label
- request Parameters: structured action parameters
- error Code: optional failure outcomes

5.2. Normal workflow templates

Normal operational traces are produced from workflow templates:

- CI/CD deployment chain: Assume Role → Describe → Update Service → Put Object
- Data engineering chain: List Bucket → Get Object → Athena Query
- Security operations chain: Get Caller Identity → List Users → Get Policy

5.3. Privilege escalation chain

Privilege escalation is generated as a role and policy transition chain:

- Low-privilege principal calls Assume Role into a misconfigured intermediate role.
- Attacker performs Put Role Policy or Attach Role Policy to expand permissions.
- Attacker assumes the now-privileged role and performs privileged actions (e.g., Create Access Key, Get Secret Value).

5.4. Persistence chain

Persistence introduces long-lived access paths:

- Create a backdoor role or user (Create Role / CreateUser).
- Attach durable policies (Attach User Policy).
- Rotate keys or modify trust relationships (Update Assume Role Policy).

5.5. Exfiltration chain

Exfiltration uses storage workflow violations:

- Discovery and enumeration (List Buckets, Get Bucket Location).
- Bulk object reads (ListObjectsV2 → repeated Get Object).
- Cross-region or anomalous timing bursts.

5.6. Resource hijacking chain

Resource hijacking simulates unexpected compute launches:

- Sudden Run Instances or container task creation.
- Burst in network egress signals (optional).
- Abnormal lifecycle patterns outside deployment windows.

5.7. Logging blackout simulation

A blackout is simulated by removing a contiguous interval of events and optionally inserting an audit disruption action:

- Example disruption: Stop Logging or trail misconfiguration events.
- Observed effect: missing mid-trajectory actions, leaving endpoints visible.

5.8. Generator pseudocode

for each trace:

- sample workflow template or attack template
- instantiate principal/role graph generate events with parameter distributions
- enforce ordering constraints (role-before-privileged actions)
- optionally inject blackout: remove events between T1 and T2 output as CloudTrail-style JSON stream

6. Experimental setup

6.1. Datasets

We evaluate on two complementary sources:

- CIC-IDS2017 (intrusion benchmark). A benchmark intrusion dataset used to compare anomaly detection baselines under controlled attack labels.
- Synthetic CloudTrail identity telemetry. The generator described in Section 5 produces multi-service identity workflows with attack chains and blackout intervals.

6.2. Baselines

We compare against:

- Isolation Forest baseline
- Autoencoder / VAE-style reconstruction baseline
- Optional sequence baseline (LSTM-AE) for temporal comparison

6.3. Evaluation metrics

Metrics emphasize operational constraints:

- ROC-AUC and PR-AUC
- Low false-positive regime: FPR@TPR=0.95
- Reconstruction quality: template BLEU, numeric RMSE, artifact recovery rate
- Red-team fidelity: discriminator accuracy (real vs synthetic)

7. Results

7.1. Detection performance across dataset sizes

Table 2 compares Isolation Forest, VAE baseline, and GIFTS under small and large telemetry settings.

Table 2 Detection performance across different dataset sizes.

Metric	IF (Small)	IF (Large)	VAE (Small)	VAE (Large)	GIFTS (Small)	GIFTS (Large)
Precision	0.73	0.78	0.76	0.81	0.84	0.89
Recall	0.69	0.74	0.72	0.77	0.82	0.88
F1-score	0.71	0.76	0.74	0.79	0.83	0.88
ROC-AUC	0.82	0.85	0.86	0.89	0.92	0.95

These results reflect that manifold structure stabilizes with broader workflow coverage, improving separability of semantically malicious trajectories from rare-but-valid administrative behavior.

7.2. Confusion summary across attack classes

Table 3 reports confusion-style counts under a representative operating point.

Table 3 Confusion summary across attack classes for GIFTS**.**

Attack class	Pred. Attack	Pred. Benign	Actual Attack	Actual Benign
Privilege escalation	9,420	580	10,000	190,000
Data exfiltration	8,950	1,050	10,000	190,000
Persistence	9,110	890	10,000	190,000
Resource hijacking	8,300	1,700	10,000	190,000

7.3. Training vs validation behavior

Table 4 summarizes training and validation AUC for a sequence baseline and GIFTS.

Table 4 Training and validation AUC across epochs.

Epoch	Seq Baseline Train	Seq Baseline Val	GIFTS Train	GIFTS Val
1	0.79	0.77	0.88	0.86
2	0.83	0.80	0.91	0.89
3	0.86	0.82	0.93	0.91
4	0.88	0.83	0.94	0.92
5	0.89	0.84	0.95	0.93

7.4. Hyperparameter tuning

Table 5 reports tuning effects for representative hyperparameters.

Table 5 Hyperparameter tuning results (representative).

Hyperparameter	Value	AUC (Small)	AUC (Large)
Learning rate	1×10^{-4}	0.90	0.94
Neighbor size k	15	0.92	0.95
Latent dim d	12	0.92	0.95
Diffusion steps	1,000	0.91	0.95

7.5. Forensic reconstruction under blackout

Table 6 reports reconstruction quality across attack scenarios under masked segments.

Table 6 Forensic reconstruction accuracy across attack scenarios**.**

Attack scenario	BLEU	RMSE	Artifact recovery (%)
Privilege escalation	0.88	0.13	93
Data exfiltration	0.91	0.10	90
Persistence	0.87	0.14	89
Resource hijacking	0.84	0.16	86

7.6. Interpretability signals

Table 7 summarizes normalized feature importance signals.

Table 7 Feature importance signals across threat categories (normalized weights).

Feature	Escalation	Exfiltration	Persistence
Role transition anomaly	0.92	0.61	0.87
Policy modification action	0.89	0.54	0.91
Region change / travel	0.66	0.72	0.58
Burst access to storage	0.51	0.93	0.47

7.7. Synthetic red-team fidelity

Table 8 evaluates synthetic traces via discriminator separability.

Table 8 Synthetic generation fidelity metrics for red teaming.

Metric	Value
Synthetic traces generated	10,000
Discriminator accuracy (real vs synthetic)	0.56
Diversity score (unique action n-grams)	0.74

7.8. Low false-positive operating point comparison

Security operations are constrained by false positives. Table 9 reports low-FPR performance.

Table 9 Low-FPR comparison across detection approaches**.**

Model	ROC-AUC	PR-AUC	FPR@TPR=0.95
Isolation Forest	0.85	0.62	0.18
VAE baseline	0.89	0.68	0.14
Seq baseline (LSTM-AE)	0.88	0.67	0.15
GIFTS (Geodesic Trust)	0.95	0.78	0.07

7.9. Ablation study

Table 10 isolates component contribution.

Table 10 Ablation study of GIFTS components.

Variant	ROC-AUC	FPR@TPR=0.95	Reconstruction BLEU
Full GIFTS	0.95	0.07	0.91
w/o Geodesic constraint (L_{geom}=0)	0.91	0.11	0.87
w/o Diffusion (decoder only)	0.89	0.13	0.72
w/o Decoder (latent-only scoring)	0.93	0.08	N/A

7.10. Runtime and operational cost comparison

Table 11 summarizes latency/throughput characteristics consistent with hierarchical deployment.

Table 11 Runtime and cost comparison for operational deployment.

Path	Latency/event	Throughput	Compute class	Primary use
Fast Path (streaming score)	1-3 ms	50,000+	CPU	broad screening
Deep Path (GIFTS)	50-200 ms	1,000-5,000	GPU/CPU hybrid	forensic analysis

8. Discussion

8.1. GIFTS reduces false positives by modeling workflow continuity

Traditional anomaly detection identifies rarity but not semantic invalidity. By modeling identity workflows as constrained trajectories, GIFTS reduces false positives in realistic operating regimes (Table 9), which is critical for real-world Security Operations Center (SOC) adoption.

8.2. Explainability via decoding improves analyst actionability

Unlike scalar anomaly scores, GIFTS returns reconstructed sequences and key artifacts. This supports decision-making by describing what likely happened rather than merely stating that something is unusual.

8.3. Generative forensics mitigates logging blackouts

In blackout conditions, diffusion in latent space enables multiple plausible reconstructions conditioned on observed endpoints. This aligns with incident response uncertainty and supports hypothesis-driven investigation.

8.4. Limitations

Limitations include concept drift in IAM policies, dependence on representation quality, and limited real-world labeled identity attack traces. Future work includes online manifold updates, multi-tenant conditional diffusion, and explicit policy-graph constraints integrated into decoder training.

9. Conclusion

This paper presented Generative Identity Forensics and Trust System (GIFTS), a manifold learning and diffusion framework for identity-centric cloud security. By embedding cloud identity telemetry into a geodesic-preserving manifold and applying diffusion-based healing, GIFTS improves low false-positive anomaly detection and provides interpretable reconstructions of behavior. The integration of Manifold Decoder inversion enables actionable forensic artifacts, while latent diffusion supports missing-evidence in-painting and realistic red-team generation. These results suggest that geometric and generative modeling can strengthen modern cloud defense, especially under high scale and adversarial audit disruption.

Compliance with ethical standards

Acknowledgments

The author acknowledges the support of the Department of Computer & Information Science & Engineering at the University of Florida where he studied.

Disclosure of conflict of interest

The author declares no conflict of interest.

Statement of ethical approval

This research did not involve human participants or animal experimentation. All data used was synthetic or from public benchmark datasets. The manuscript was written through contributions of all authors. All authors have given approval to the final version of the manuscript

References

[1] Amazon Web Services. Amazon GuardDuty Documentation. Available from: https://aws.amazon.com/guardduty/.

[2] National Institute of Standards and Technology. Zero Trust Architecture. NIST Special Publication 800-207. 2020.

[3] Liu FT, Ting KM, Zhou ZH. Isolation Forest. Proceedings of ICDM. 2008.

[4] Tenenbaum JB, de Silva V, Langford JC. A global geometric framework for nonlinear dimensionality reduction. Science. 2000;290(5500):2319-2323.

[5] McInnes L, Healy J, Melville J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv:1802.03426. 2018.

[6] Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. NeurIPS. 2020.

[7] Guha S, Mishra N, Roy G, Schrijvers O. Robust random cut forest based anomaly detection on streams. ICML. 2016.

[8] He P, Zhu J, He S, Li J, Lyu MR. Drain: An Online Log Parsing Approach with Fixed Depth Tree. 2017.

[9] Thakare R, Akugri KM. Manifold Decoders: A Framework for Generative Modeling from Nonlinear Embeddings. arXiv preprint arXiv:2510.13622. 2025.

[10] Sharafaldin I, Lashkari AH, Ghorbani AA. Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of ICISSP. 2018. (CIC-IDS2017)

[11] Adrah FA. Cybersecurity Resilience of Telehealth Teams: The Effects of Cognitive Load and Security Fatigue. AMCIS 2025 TREOs. AIS Electronic Library (AISeL). 2025. Available from: https://aisel.aisnet.org/treos_amcis2025/43/.

[12] MITRE. ATT&CK Knowledge Base. Available from: https://attack.mitre.org/

[13] Agboklu M, Lartey B, Adrah F, LaJeunesse D. Towards Proactive Heart Health: A Ma-chine Learning-Powered Approach for Chronic Heart Failure Detection. InternationalJournal of Computer Applications. 2024;186(28). July 2024.

[14] Azigi JA, Adrah F. Toward Smart Biosensing: A Machine Learning Approach for EarlyDiabetes Detection. International Journal of Computer Applications. 2025;187(19). June2025.