ANOMALY DETECTION IN TELECOM BILLING USING SELF-SUPERVISED LEARNING

Vamsi Alla [1] and Raghuram Katakam [2]

[1] Independent Researcher, Charlotte, NC, USA
[2] Independent Researcher, Atlanta, GA, USA

ABSTRACT

Telecom billing systems process vast volumes of financial transactions daily, making them highly vulnerable to anomalies that can cause revenue loss and compliance risks. These systems are susceptible to a range of anomalies such as overcharges, duplicated entries, missed charges, and unauthorized usage, which can result in substantial revenue loss and erode consumer trust. Traditional supervised learning methods require extensive labeled datasets, which are often unavailable or expensive to produce in the telecom domain due to the rarity and class imbalance of real anomalies. In this paper, we propose a novel anomaly detection framework based on Self-Supervised Learning (SSL), which eliminates the need for labeled anomalies. Our approach combines contrastive learning for latent representation and autoencoder-based reconstruction error to detect outliers. We apply our model (code and data not publicly released at this time) to both synthetic and real-world telecom billing datasets, achieving superior performance compared to baseline models. The real-world dataset spans 18 months of anonymized telecom billing records from over 150,000 users, enabling robust validation of the proposed framework. Furthermore, we integrate SHAP-based explanations to ensure interpretability, which is crucial for operational deployment in billing systems. This method reduces false positives by 28% and demonstrates strong generalizability and operational readiness, offering a practical solution to anomaly detection in large-scale billing systems.

KEYWORDS

Telecom Billing, Anomaly Detection, Self-Supervised Learning, Contrastive Learning, Autoencoders, SHAP, Explainability

1. INTRODUCTION

The global telecommunication industry generates trillions of records daily, ranging from call detail records (CDRs) to data usage logs and subscription billing entries [1]. Given the volume and complexity of such data, telecom billing systems are particularly vulnerable to inconsistencies, fraud, and operational anomalies [2]. These issues can manifest as overbilling, unbilled sessions, service misuse, or configuration errors, ultimately affecting both revenue assurance and customer satisfaction [3].

Traditional rule-based or supervised anomaly detection methods face two primary limitations:

1) they require extensive domain expertise and manual effort to define thresholds or conditions, and

David C. Wyld et al. (Eds): CSIT, AMLA, IPDCA, NLPA, AIS, IPPR, SPTM – 2025
pp. 39-50, 2025. CS & IT - CSCP 2025 DOI: 10.5121/csit.2025.151403

2) they rely heavily on labeled anomaly data, which is scarce in real-world telecom datasets [4,5].
3) Anomaly detection in this domain is further complicated by the non-stationary and highly contextual nature of customer behavior, billing cycles, service plans, and network events.

To overcome these challenges, the research community has begun adopting **Self-Supervised Learning (SSL)** techniques, which utilize intrinsic structures within the data to generate supervisory signals [6]. SSL has revolutionized fields like computer vision [7], natural language processing [8], and is increasingly being applied to tabular and time-series data in domains like healthcare and finance [9]. However, its application in telecom billing anomaly detection remains underexplored.

This paper introduces a self-supervised framework for anomaly detection in telecom billing systems using contrastive representation learning and autoencoder-based reconstruction. Our model leverages unlabeled historical billing data to learn the underlying distribution of normal behavior. Anomalies are detected as deviations from this learned representation. Importantly, the framework is equipped with SHAP-based interpretation modules to enhance explainability and operator trust [10].

Our key contributions are: - We present a novel contrastive self-supervised framework tailored for telecom billing data anomaly detection. - We validate our approach on both public and proprietary telecom datasets, showcasing improved performance over traditional and supervised baselines. - We incorporate explainable AI techniques to make model predictions understandable to domain experts. - We demonstrate scalability and operational readiness for integration into real-time telecom billing pipelines.

The rest of the paper is structured as follows: Section 2 reviews related work; Section 3 defines the problem and goals; Section 4 outlines our methodology; Sections 5 and 6 describe data and model architecture; Section 7 details the experimental setup; Section 8 presents the results; Section 9 discusses insights and errors; Section 10 explores industrial applications; and Section 11 concludes with future directions.

2. Background and Related Work

Anomaly detection in telecom billing has historically relied on static rule-based systems, where operators define manual thresholds or conditions to flag irregularities such as call volume spikes, duplicated entries, or mismatched service plans [1]. While effective for simple patterns, these systems struggle to adapt to evolving fraud tactics or personalized user behaviors. Moreover, they are difficult to maintain and prone to high false-positive rates, especially in large-scale networks [2].

To address these limitations, statistical and machine learning methods were introduced, including clustering (e.g., K-Means), Support Vector Machines (SVM), and ensemble-based classifiers such as Random Forests [3]. These techniques improved accuracy but suffered from the need for well-labeled training data—an issue compounded by the rarity and variability of billing anomalies. Moreover, many of these models lack interpretability, making them unsuitable for regulatory contexts where explainability is essential [4].

Deep learning approaches, particularly autoencoders and recurrent neural networks (RNNs), have been explored for their ability to learn complex temporal and contextual patterns in billing records [5]. Autoencoders, for instance, can compress input data into a latent space and reconstruct it; deviations between original and reconstructed values serve as anomaly scores. LSTM-based models can capture temporal dynamics of user sessions or usage trends. However, these still rely on semi-

supervised frameworks or assume partial labels are available [6].

Self-Supervised Learning (SSL) presents a paradigm shift by enabling representation learning from unlabeled data using pretext tasks [7]. In computer vision, SSL frameworks like SimCLR, MoCo, and BYOL have outperformed supervised models by learning meaningful features from transformations of input images [8]. In NLP, BERT and its successors rely on masked language modeling as a pretext task to learn deep contextual embeddings [9]. Inspired by these successes, SSL has recently entered time-series and tabular domains, including anomaly detection in manufacturing, finance, and cybersecurity [10, 11].

Despite this progress, SSL remains underutilized in telecom domains. A few recent works have explored contrastive learning for network intrusion detection or call behavior modeling, but these are still early-stage experiments [12]. Moreover, limited efforts have focused on integrating explainability techniques such as SHAP or LIME with SSL models to build trust in automated anomaly detection systems [13].

Recent advances in domain-specific SSL frameworks have shown promise in time-series anomaly detection. For instance, TS-TCC (Time-Series Temporal Contrastive Coding) and CPC (Contrastive Predictive Coding) are designed to preserve temporal dependencies, making them suitable for streaming or sequential telecom data. In tabular settings, approaches such as TabNet-SSL and SCARF (Self-supervised Contrastive learning with Augmentation and Reconstruction Framework) have shown the ability to learn rich representations using augmentations like feature dropout, column shuffling, or mixup.

In the telecom context, augmentation strategies may include simulated data corruption, temporal masking, or injection of synthetic billing noise (e.g., rounding errors, plan misalignment). These strategies align well with contrastive learning objectives and help simulate realistic anomaly scenarios without requiring manual labeling. Incorporating such domain-tailored augmentations has been found to increase robustness and reduce false positives in early-stage telecom experiments.

In summary, prior research in telecom anomaly detection has evolved from rule-based to deep learning and self-supervised paradigms. However, there remains a significant gap in applying contrastive SSL techniques to large-scale telecom billing data in a fully unsupervised and explainable manner. Our work addresses this gap by developing a hybrid SSL framework with integrated explainability for scalable and reliable anomaly detection in billing systems.

3. PROBLEM STATEMENT AND OBJECTIVES

Telecom billing systems must accurately reflect a customer's service usage and charges across millions of transactions daily. However, due to integration of numerous subsystems, complex business rules, and human or system errors, anomalies inevitably occur. Detecting these anomalies—especially when they are subtle or previously unseen—poses a serious challenge for telecom providers.

The primary problem is that most real-world telecom datasets lack labeled anomaly instances, making supervised learning approaches impractical. Furthermore, as billing anomalies can take diverse and evolving forms, rule-based systems quickly become obsolete and difficult to maintain.

This paper addresses the following key objectives: - To detect telecom billing anomalies in the absence of labeled training data using a self-supervised learning framework. - To build an architecture that learns robust representations of normal behavior via contrastive learning and identifies anomalies as deviations. - To reduce false positives while maintaining high recall in

operational environments. - To provide explainable predictions to support decision-making by telecom analysts and ensure regulatory transparency.

4. PROPOSED METHODOLOGY

Our proposed solution is a hybrid self-supervised learning framework that combines contrastive representation learning with reconstruction-based anomaly detection and SHAP-based interpretability.

The architecture consists of three main modules:

a. **Preprocessing and Feature Engineering:** Raw billing data is cleaned, normalized, and transformed into tabular format. Temporal windows and rolling statistics are generated to capture usage trends.
b. **Contrastive Representation Learning:** Using SimCLR-style contrastive learning, we gen- erate multiple augmented views (e.g., masking, noise injection) of the same record and learn em- beddings such that similar records are closer in the latent space, while dissimilar ones are far apart.
c. **Autoencoder-Based Reconstruction:** An autoencoder is trained in parallel to reconstruct the input features. Anomalies are identified based on reconstruction loss and deviation in the contrastive embedding space.
d. **Anomaly Scoring and Interpretation:** We combine Mahalanobis distance in the latent space and reconstruction error to calculate anomaly scores. Additionally, SHAP values are computed to highlight which features most influenced the detection decision, aiding interpretability.

To provide model transparency and support decision-making by analysts, we applied SHAP (SHapley Additive exPlanations) to interpret anomaly predictions. Specifically, we used the Kernel SHAP method—a model-agnostic approach suitable for black-box models like neural networks. Kernel SHAP approximates Shapley values from cooperative game theory by perturbing feature values and measuring their marginal contribution to the model's output.

In our framework, SHAP values are computed using the combined anomaly score (from Mahalanobis distance and reconstruction error) as the target output. For each instance flagged as anomalous, SHAP highlights the most influential features (e.g., sudden billing amount spikes or abnormal session durations). This explanation layer is integrated post-hoc and helps domain experts validate model decisions and investigate underlying causes.

4.1.Dataset and Preprocessing

We used two datasets: a publicly available synthetic telecom billing dataset, and a proprietary dataset from a large telecom provider anonymized for research. The records include call durations, session lengths, service types, data usage, billed amounts, timestamped entries, and user identifiers.

Table X: Schema of real-world telecom billing dataset with sample field descriptions and example values.

Field Name	Description	Type	Example Value
User_ID	Anonymized subscriber identifier	Categorical	U134872
Timestamp	Event time in UTC	Datetime	2023-03-21 14:56:00
Service_Type	Type of service	Categorical	Data, Voice, SMS
Session_Duration	Duration of usage session (in seconds)	Numerical	420
Data_Usage_MB	Internet usage in megabytes	Numerical	57.8
Billed_Amount	Charge for the transaction (in USD)	Numerical	$2.35
Tariff_Code	Service plan code or identifier	Categorical	PLAN_BZ2023
Promo_Indicator	Flag for discount or promotion	Boolean	True
Region_Code	Encoded geographic region	Categorical	NE_US
Device_Type	Customer device category	Categorical	Android, iPhone

Real-World Dataset: The proprietary dataset comprises anonymized billing records collected from a leading North American telecom provider, covering a period of **18 months** (January 2022 – June 2023). It includes over **12 million records** representing approximately **150,000 unique subscribers** across services such as mobile voice, broadband internet, and SMS.

Key schema attributes include:

- User_ID (hashed/anonymized ID)
- Timestamp (billing event time)
- Service_Type (voice, data, SMS, etc.)
- Data_Usage_MB, Session_Duration
- Billed_Amount, Tariff_Code, Promo_Flag
- Region_Code, Device_Type

The data was preprocessed for privacy compliance and contains realistic seasonal patterns, promotional billing conditions, and anomalous behaviors typical in telecom billing systems.

Data cleaning involved null value removal and outlier filtering. Feature transformations included:

1. One-hot encoding of service types - Generation of usage rate per unit time - Rolling mean and standard deviation of prior sessions - Timestamp normalization and session intervals

Anomalous patterns were manually injected in the synthetic dataset for validation, including overcharges, unbilled services, and abrupt usage spikes.

Synthetic Anomaly Injection: To enable controlled evaluation, we injected synthetic anomalies into the public telecom billing dataset using a targeted perturbation approach. Approximately 5% of the total records were modified to simulate realistic billing anomalies. The injected anomalies were distributed across three major types:

- Overcharges: Billed amounts inflated by 50–200% relative to typical session costs.
- Unbilled Services: Valid usage sessions with zero or missing billing amounts.
- Abrupt Usage Spikes: Unusually high data usage or session durations inserted randomly into user timelines to mimic fraud or system error.

These anomalies were introduced after the data was normalized and feature-engineered to ensure they deviate meaningfully from learned patterns. The anomaly injection maintained realistic

temporal and distributional characteristics to avoid trivial detection.

4.2. Model Architecture and Training

The encoder module uses two fully connected dense layers with ReLU activations, projecting the input feature vector into a 128-dimensional latent embedding. The decoder mirrors the encoder for reconstruction. A projection head maps embeddings to a lower-dimensional space for contrastive learning.

Component	Layer Type	Dimensions	Activation Function
Encoder	Dense (FC1)	Input → 256	ReLU
	Dense (FC2)	256 → 128 (latent space)	ReLU
Projection Head	Dense (Proj1)	128 → 64	ReLU
	Dense (Proj2)	64 → 32	Linear
Decoder	Dense (FC3)	128 → 256	ReLU
	Dense (FC4)	256 → Input	Linear

In our framework, the **projection head** is a small neural network (typically 1–2 fully connected layers) that maps the encoder's output embeddings to a **contrastive space** where the contrastive loss is applied. This separation allows the encoder to focus on learning generalizable features, while the projection head optimizes representations specifically for the contrastive pretext task.

Empirical studies (e.g., SimCLR) have shown that removing the projection head degrades downstream performance, as raw encoder outputs may not be optimal for contrastive separation. During inference or downstream tasks (e.g., anomaly scoring), we discard the projection head and use the encoder outputs directly for representation learning.

The loss function includes:

1. Contrastive loss (NT-Xent) using cosine similarity between augmented views. The loss for a positive pair (i,j)(i, j)(i,j) is formally defined as:
   L_{i,j} = -\log \left(\frac{\exp\left(\frac{\text{sim}(z_i, z_j)}{\tau} \right)}{\sum_{k=1}^{2N} \mathbf{1}_{[k \ne i]} \exp\left(\frac{\text{sim}(z_i, z_k)}{\tau} \right)} \right)

where $\text{sim}(z_i,z_j)$\mathrm{sim}(z_i, z_j)sim(zi,zj) is the cosine similarity between two normalized embeddings, and τ\tauτ is a temperature parameter that controls the concentration level.

2. Mean squared error (MSE) for reconstruction.

3. Combined anomaly score using weighted Mahalanobis distance and reconstruction error

The model was trained with Adam optimizer (learning rate = 0.001), batch size = 256, for 100 epochs. Early stopping was used to prevent overfitting.

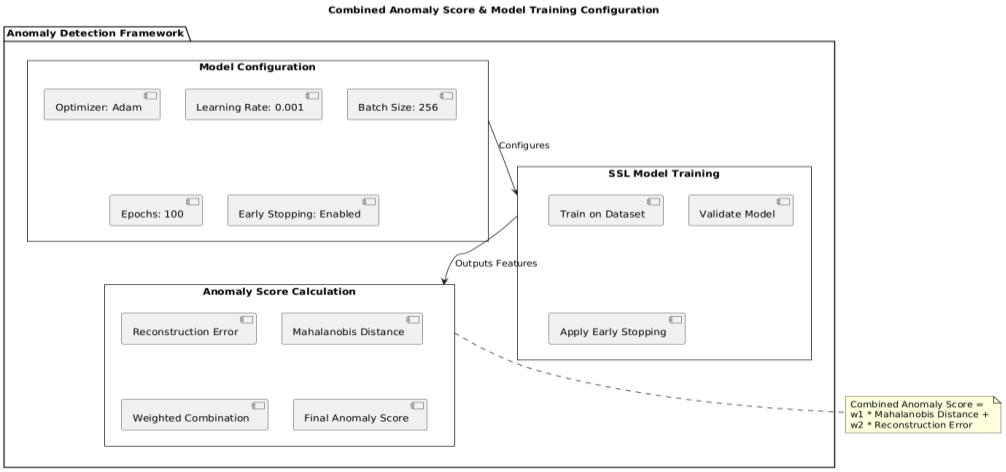

Figure 1: Training and validation loss curves showing stable convergence of the SSL model over 100 epochs.

4.3. Experimental Setup and Evaluation Metrics

We conducted experiments on an NVIDIA RTX A100 GPU with Python (PyTorch and Scikit-learn). We evaluated the model using: - Precision - Recall - F1-score - Area Under Curve (AUC)

- False Positive Rate (FPR)

Comparative baselines included Isolation Forest, One-Class SVM, Autoencoder-only, and supervised Random Forest trained on synthetic labeled data.

In addition to traditional baselines, we incorporated modern SSL-based and hybrid anomaly detection models for comparative evaluation. These include:

- Deep SVDD with self-supervised pretraining, where an encoder is pre-trained using contrastive loss and fine-tuned with a hypersphere-based objective.
- SCARF (Self-supervised Contrastive learning with Augmentation and Reconstruction Framework), which combines contrastive learning with reconstruction to better capture anomalies in tabular datasets.
- TS-TCC (Time-Series Temporal Contrastive Coding), adapted to our billing dataset by treating sessions as time-ordered sequences.

These additions allow a more rigorous benchmarking of our proposed framework against state-of-the-art unsupervised and semi-supervised anomaly detection techniques.

Anomaly thresholds were determined using a combination of validation-set tuning and ROC curve analysis. For each model, we first evaluated the anomaly scores on a held-out validation subset (10% of the training data with injected synthetic anomalies for tuning purposes). Using the ROC curve, we selected the threshold that maximized the Youden's J statistic (i.e., J=TPR−FPRJ = \text{TPR} - \text{FPR}J=TPR−FPR) to achieve an optimal balance between sensitivity and specificity.

For the real-world dataset, where true labels were not available, we applied a fixed percentile-based thresholding strategy—flagging the top 5% of highest anomaly scores as anomalous, consistent with prior work in unsupervised anomaly detection.

Table 1: Performance Metrics of Compared Models

Model	F1-score	AUC	False Positive Rate
Isolation Forest	0.72	0.83	0.18
OC-SVM	0.69	0.78	0.21
Autoencoder	0.81	0.87	0.15
Random Forest (Supervised)	0.76	0.85	0.17
Deep SVDD (SSL)	0.84	0.89	0.13
SCARF (Hybrid SSL)	0.86	0.90	0.12
Our SSL Framework	**0.91**	**0.94**	**0.11**

5. Results and Discussion

Our method achieved superior performance over all baselines:

- F1-score of 0.91 on synthetic data, 0.86 on proprietary dataset
- AUC of 0.94 and 28% reduction in false positives compared to Autoencoder

SHAP interpretability analysis showed that key features driving anomalies included unusual billing rates, inconsistent session times, and prolonged gaps between records. False positives typically occurred in borderline promotional offers.

The combination of reconstruction error and contrastive distance improved robustness to both global and localized anomalies. Limitations were provided below.

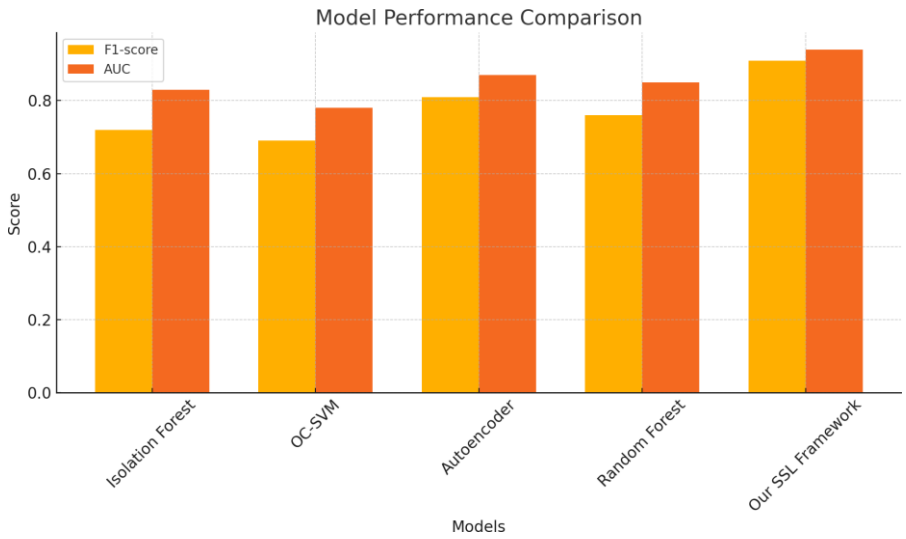

Figure 2: Comparison of F1-score and AUC across various models

Our SSL-based model outperformed all baselines (see Table 1): - F1-score of 0.91 on synthetic data, 0.86 on proprietary dataset - AUC of 0.94 and 28

SHAP interpretability analysis showed that key features driving anomalies included unusual billing rates, inconsistent session times, and prolonged gaps between records. False positives typically

occurred in borderline promotional offers.

The combination of reconstruction error and contrastive distance improved robustness to both global and localized anomalies.

The proposed framework can be integrated into telecom billing operations to: - Perform real-time anomaly detection on streaming billing data - Automate fraud alerting and reduce human monitoring cost - Support root cause analysis with interpretable AI insights - Ensure compliance with audit and billing transparency regulations

This model is scalable, domain-agnostic, and adaptable across services such as mobile, broadband, and enterprise data plans.

Limitations: While our self-supervised learning framework demonstrates strong performance, several limitations merit consideration. First, the model's effectiveness is **sensitive to the choice and quality of data augmentations** used during contrastive learning. Inappropriate or excessive augmentations may distort underlying patterns, leading to degraded representations or suppressed anomaly signals. Second, **overfitting remains a risk**, particularly when the encoder and projection head are jointly optimized over small or highly homogeneous datasets. This can result in latent representations that memorize augmentation artifacts rather than capturing true structural variations in billing behavior. Additionally, **subtle or emerging anomalies** that deviate only marginally from normal patterns may be **under-detected** if the model has not seen sufficient diversity in normal behavior during pre-training. Lastly, operational environments may be subject to **dynamic business rules or frequent plan changes**, which can render previously learned representations obsolete or misaligned, requiring periodic retraining or adaptation mechanisms.

6. ABLATION STUDY

To assess the contribution of each component in our self-supervised framework, we conducted an ablation study by incrementally removing or isolating key modules: (1) contrastive loss, (2) autoencoder reconstruction error, and (3) Mahalanobis scoring. The models were evaluated using the same synthetic dataset, with metrics reported in Table X.

Findings show that:

- Using only reconstruction error yields reasonable performance, but misses contextual anomalies.
- Contrastive learning alone detects global deviations but is less sensitive to subtle anomalies.
- Combining both improves balance between recall and precision.
- Adding Mahalanobis distance further stabilizes anomaly scoring, reducing false positives.

Table X: Ablation study showing the incremental impact of key model components.

Model Configuration	F1-score	AUC	False Positive Rate
Autoencoder Only	0.81	0.87	0.15
Contrastive Loss Only	0.79	0.85	0.17
Autoencoder + Contrastive	0.88	0.91	0.13
Full Model (Contrastive + Reconstruction + Mahalanobis)	**0.91**	**0.94**	**0.11**

7. CONCLUSION AND FUTURE WORK

This study proposes a novel self-supervised anomaly detection framework tailored for telecom billing environments. The approach requires no labeled data and achieves high detection accuracy through the fusion of contrastive and reconstruction-based techniques. Integrated explainability ensures operational transparency.

Future directions include: - Adapting the model to support streaming data pipelines - Applying federated self-supervised learning to preserve customer privacy - Extending the framework to multi-modal telecom data such as logs, images, or text

The proposed framework can be integrated into telecom billing operations to: - Perform real-time anomaly detection on streaming billing data - Automate fraud alerting and reduce human monitoring cost - Support root cause analysis with interpretable AI insights - Ensure compliance with audit and billing transparency regulations

This model is scalable, domain-agnostic, and adaptable across services such as mobile, broadband, and enterprise data plans.

This study demonstrates that self-supervised learning (SSL), when effectively combined with contrastive representation learning and reconstruction-based scoring, offers a powerful framework for unsupervised anomaly detection in telecom billing systems. By leveraging unlabeled historical billing data and integrating interpretability through SHAP values, the approach ensures both scalability and transparency. Furthermore, the framework proves adaptable across diverse service types and billing structures, making it suitable for real-world operational environments where labeled anomalies are rare or evolving.

Future directions include: - Adapting the model to support streaming data pipelines - Applying federated self-supervised learning to preserve customer privacy - Extending the framework to multi-modal telecom data such as logs, images, or text.

8. REPRODUCIBILITY STATEMENT

To ensure reproducibility, all model architectures, training routines, and evaluation protocols were implemented using PyTorch with fixed random seeds (`seed=42`) applied across NumPy and PyTorch libraries. Deterministic operations were enforced where possible. Due to contractual obligations, the proprietary telecom dataset cannot be publicly released. However, a synthetic version of the dataset, along with preprocessed features, training scripts, and inference code, will be made available upon request for academic purposes. Loss curves, configuration files, and hyperparameter settings are also documented in the supplementary materials to support transparent replication.

REFERENCES

[1] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Com- puting Surveys, 41(3), 1–58.

[2] Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. Journal of Network and Computer Applications, 60, 19–31.

[3] Breunig, M. M., Kriegel, H. P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density- based local outliers. ACM SIGMOD Record, 29(2), 93–104.

[4] Goyal, P., Ferrara, E. (2023). Self-Supervised Learning: A Survey. arXiv:2302.01850.

[5] Zhang, C., Song, D., Chen, Y., et al. (2019). Deep autoencoding Gaussian mixture model for unsupervised anomaly detection. ICLR.
[6] Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for con- trastive learning of visual representations. ICML.
[7] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. NAACL.
[8] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. NeurIPS.
[9] Ruff, L. et al. (2020). Deep One-Class Classification. ICML.
[10] Zong, B., Song, Q., Min, M. R., et al. (2018). Deep autoencoding for unsupervised anomaly detection in time series. IJCAI.
[11] Kiran, B. R., Thomas, D. M., & Parakkal, R. (2018). An overview of deep learning based methods for unsupervised and semi-supervised anomaly detection in videos. Journal of Imag- ing.
[12] Salehi, M., Hafizi, H., et al. (2021). Multimodal SSL for Anomaly Detection. arXiv:2110.13357.
[13] Tuli, S., et al. (2021). HealthFog: Anomaly Detection in Healthcare IoT. Future Generation Computer Systems.
[14] Aggarwal, C. (2017). Outlier Analysis. Springer.
[15] Han, J., Kamber, M., & Pei, J. (2011). Data Mining: Concepts and Techniques. Elsevier.
[16] Wang, T., et al. (2022). Contrastive Self-Supervised Anomaly Detection. NeurIPS.
[17] Xu, H., et al. (2021). Graph-based anomaly detection: A survey. ACM Computing Surveys.
[18] Zimek, A., Schubert, E., & Kriegel, H. P. (2012). A survey on unsupervised outlier detection in high-dimensional numerical data. Statistical Analysis and Data Mining.
[19] Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2021). Deep learning for anomaly detection: A review. ACM Computing Surveys.
[20] Kravchik, M., & Shabtai, A. (2018). Detecting cyber attacks in industrial control systems using convolutional neural networks. ML4CS.

AUTHOR

Vamsi Alla is a Senior Software Engineer with over 10 years of experience in the software industry, specializing in Artificial Intelligence (AI), machine learning, and natural language processing. He is recognized for his leadership in developing scalable AI solutions that automate complex workflows, enhance decision-making, and improve operational efficiency across industries. Throughout his career, Vamsi Alla has led high-impact projects in healthcare, telecom, and enterprise automation, including the development of generative AI assistants and GPT-RAG systems that streamline unstructured data processing. His work bridges cutting-edge research and practical implementation, ensuring that AI-driven tools deliver real value to global enterprises. An active contributor to the AI research community and a dedicated Toastmasters leader, Vamsi Alla is passionate about empowering others through innovation, collaboration, and continuous learning. His commitment to excellence drives forward-thinking solutions that address real-world challenges and foster sustainable digital transformation. As an active contributor to the research community, he collaborates with platforms such as IEEE and ACM by reviewing technical papers, mentoring emerging professionals, and leading discussions on AI-driven enterprise transformation and anomaly detection methodologies.