

Selective Noise Injection and Feature Scoring for Unsupervised Request Anomaly Detection

Yu Cheng

Fordham University, New York, USA

ycheng77@fordham.edu

Abstract: This paper addresses the problem of anomaly detection in high-dimensional and complex request scenarios. An unsupervised anomaly request detection method based on diffusion models is proposed. The method integrates generative modeling with a feature discrimination mechanism. It builds a forward noise injection and reverse reconstruction process to effectively capture the distribution characteristics of normal request data. Specifically, a structure-aware module, Selective Noise Injection for Reconstruction (SNIR), is introduced. It selectively injects noise during the forward diffusion phase to preserve key feature dimensions and improve reconstruction quality. On this basis, a Feature-aware Discriminative Scoring (FDS) mechanism is designed. It embeds the semantic features of original and reconstructed requests and computes a combined score using Euclidean distance and cosine similarity. This enables fine-grained discrimination of abnormal requests. The method does not rely on labeled data. It uses only normal samples for both modeling and detection. It offers strong generalization and practical applicability. Experimental results on multiple benchmark datasets show that the proposed method significantly outperforms existing representative methods in terms of AUC, Precision, and F1. It effectively handles diverse attack patterns and changes in data distribution. It also maintains stable detection performance in high-noise environments.

Keywords: Diffusion modeling; anomaly detection; reconstruction mechanism; feature embedding

1. Introduction

With the rapid development of network and information technology, backend systems have become the core hub of the entire application ecosystem. They handle an increasing volume of data requests and service calls[1,2]. As the number of users and the complexity of requests grow, backend systems face more severe challenges in security and stability. Abnormal requests, as potential threats in system operation, may originate from malicious attacks, configuration errors, or system failures. These can cause service interruptions, data leakage, resource abuse, and a decline in user experience. Therefore, designing an efficient and accurate abnormal request detection system is crucial for ensuring system reliability and improving service quality[3].

Traditional anomaly detection methods mainly rely on rule-based statistical features or machine learning techniques such as clustering and classifiers. These approaches have practical value in early-stage applications. However, they often depend on manually crafted features and struggle with complex and dynamically evolving request patterns. In modern distributed systems, API request behavior is typically high-dimensional, non-linear, and weakly structured[4]. Traditional models often fail to capture deep semantic information and contextual dependencies in data. As a result, their detection performance is limited, especially when facing novel attacks or stealthy anomalies. There is an urgent need for more intelligent and expressive modeling methods in anomaly detection[5].

In recent years, generative models, especially deep generative models, have offered new perspectives for anomaly detection. Compared to discriminative models, generative models can learn the latent distribution of normal data. They identify anomalies through reconstruction errors or generation probabilities[6]. This approach is well-suited for unsupervised or semi-supervised settings, as it does not require large amounts of abnormal data. Current mainstream generative methods include Variational Autoencoders and Generative Adversarial Networks. While they have shown promising results, challenges remain in modeling complex temporal data, sparse structures, and high-dimensional semantics. In particular, their reconstruction accuracy and controllability still need improvement[7].

Diffusion models have emerged as a new type of generative model. They show strong modeling capabilities and stable reconstruction quality in fields such as image generation, speech synthesis, and natural language processing. The core idea is to gradually add noise and then reverse the process to denoise, approximating the target distribution from a random one. This progressive mechanism provides better controllability and robustness[8]. In anomaly detection, this makes diffusion models highly promising. For complex request sequences with temporal or contextual dependencies, diffusion models can learn and reconstruct normal behavior step by step. This allows accurate identification of deviations from normal patterns.

A discriminative reconstruction strategy based on diffusion models improves anomaly detection accuracy. It also offers strong interpretability and scalability. In practice, this strategy

can work with modules like log analysis, system monitoring, and access control. It enables early detection, localization, and response to abnormal requests. The general modeling ability of diffusion models also allows them to handle various request data types, including structured, semi-structured, and unstructured data. This supports building a unified and efficient security detection framework. Therefore, applying diffusion models to backend anomaly detection has significant theoretical value and broad engineering potential. It also contributes to the development of intelligent systems with high security and availability.

2. Related work

2.1 Diffusion Model

Diffusion models were originally proposed for data generation tasks. The core idea is to gradually inject noise into data to transform it into a form close to a Gaussian distribution[9,10]. Then, a reverse process is learned to recover the original data step by step from pure noise. This two-stage process of forward perturbation and reverse reconstruction allows the model to capture fine-grained structures and deep distribution information[11]. Compared to traditional generative models, diffusion models do not rely on explicit probability density estimation and do not require complex adversarial training. They offer more stable training and higher-quality generation. This mechanism has shown significant advantages in image generation, speech synthesis, and natural language modeling. It has become one of the mainstream frameworks in generative modeling[12].

In anomaly detection tasks, the main motivation for introducing diffusion models lies in their powerful data reconstruction ability. Diffusion models learn the latent distribution of normal samples[13]. As a result, they can reconstruct normal data accurately. When encountering unseen anomalies that deviate from the training distribution, the reconstruction performance drops significantly. The difference in reconstruction error becomes a key indicator for identifying anomalies[14]. Compared with earlier reconstruction methods using autoencoders or variational autoencoders, diffusion models offer stronger modeling capacity and more stable reconstruction quality. This is especially effective when handling high-dimensional and complex data. Diffusion models can better recover detailed data features, improving sensitivity to fine-grained anomalies and reducing false negatives and false positives[15].

In addition, diffusion models have a natural advantage in unsupervised learning. This makes them suitable for anomaly detection scenarios where labeled data is scarce. The training process relies only on normal data and does not require abnormal samples. This fits common real-world one-class learning settings. With the development of optimization strategies such as downsampling acceleration, conditional modeling, and attention mechanisms, the training efficiency and inference speed of diffusion models have improved significantly. They are gradually becoming applicable to online or near real-time scenarios. These features make diffusion models promising not only in theoretical research but also in practical systems. They provide a solid technical foundation for

building more efficient and intelligent anomaly detection systems[16].

2.2 Reconstruction-discriminative anomaly detection algorithm

Reconstruction-based anomaly detection algorithms have been widely used in unsupervised and weakly supervised settings in recent years. The core idea is to learn reconstruction patterns from normal data and use reconstruction error as the criterion for anomaly detection[17]. These methods typically adopt an encoder-decoder neural network structure, such as autoencoders, variational autoencoders, or generative adversarial networks[18]. They map input data into a low-dimensional latent space and then reconstruct it back to the original space. Since the model is trained only on normal data, it can accurately reconstruct normal samples. However, when faced with abnormal data that deviates from the training distribution, the model usually fails to produce high-quality reconstructions[19]. This leads to larger reconstruction errors, which serve as an effective metric for distinguishing normal from abnormal samples[20,21].

Compared with traditional discriminative methods such as classifiers or clustering models, reconstruction-based detection has advantages like independence from labels and strong adaptability. It is especially suitable for scenarios with rare or unknown anomaly distributions[22]. With the advancement of deep learning, many enhanced models have been proposed to improve the discriminative ability of this approach. For example, attention mechanisms, residual connections, and multi-scale feature fusion have been introduced to improve reconstruction quality and robustness[23]. Time series modeling techniques such as LSTM and Transformer have also been integrated to handle anomalies with temporal dependencies. In addition, some studies have combined reconstruction error with distributional distance in the feature space to build multi-level anomaly scoring systems. This improves the detection of boundary and stealth anomalies[24].

Although these methods have shown strong performance in various applications, they also have limitations. In some situations, the model may produce accurate reconstructions for certain abnormal samples, which leads to false negatives and weakens the reliability of detection. This occurs because the model focuses primarily on reconstructing input data without explicitly learning to differentiate between normal and abnormal patterns. Moreover, relying solely on reconstruction error as the detection criterion can be insufficient, as it may fail to reflect the subtle structural or semantic irregularities present in complex data distributions[25].

To address these challenges, recent research has increasingly explored the integration of reconstruction mechanisms with discriminative modeling. This combined approach typically introduces a discriminator or an auxiliary classification module that operates on the reconstructed outputs. The goal is to enhance the model's sensitivity to anomalies by leveraging both reconstruction consistency and discriminative signals. This hybrid strategy enables a more comprehensive understanding of data behavior, capturing not only how well a sample is reconstructed but also how its feature representations align with the learned patterns of normality.

3. Method

This study proposes an anomaly request detection method based on diffusion models. It combines a generative reconstruction mechanism with a discriminative enhancement strategy to improve anomaly recognition in complex request environments. Compared with existing methods, the proposed approach includes two main innovations. First, a Selective Noise Injection for Reconstruction (SNIR) strategy is introduced. During the forward diffusion process, noise is selectively injected to retain critical structural information.

This enhances the model's ability to reconstruct normal requests more effectively. Second, a Feature-aware Discriminative Scoring (FDS) mechanism is designed. It integrates original request features with reconstruction errors to perform multidimensional evaluation. This improves the model's sensitivity to boundary anomalies and potential threats. These two innovations complement each other. They enable the model to maintain high-quality generation while improving discriminative performance. The method is suitable for diverse and highly dynamic API request scenarios. The model architecture is shown in Figure 1.




Figure 1. Overall model architecture diagram

3.1 Selective Noise Injection for Reconstruction

In the basic framework of diffusion models, the forward diffusion process gradually adds Gaussian noise to input data, effectively transforming the original data distribution into a standard Gaussian distribution over multiple steps. Traditional diffusion models follow a fixed noise schedule, which uniformly applies noise across all feature dimensions regardless of their individual roles or importance. While this approach is mathematically straightforward, it overlooks the heterogeneous nature of real-world request data, where different features often carry varying structural significance. Some dimensions contain critical information essential for accurate reconstruction, and indiscriminate noise injection into these features can result in the degradation of meaningful structural patterns. To mitigate this issue, this study introduces a Selective Noise Injection for Reconstruction (SNIR) mechanism, which incorporates a structure-aware noise scheduling strategy during the forward diffusion phase. By adaptively adjusting the noise applied to each feature based on

its structural relevance, SNIR enables the model to preserve essential characteristics of the input data while still benefiting from the generative capacity of the diffusion framework. Its module architecture is shown in Figure 2.




Figure 2. SNIR module architecture

We assume that the original request sample is $x_0 \in R^d$, and the sample at the t -th step of the diffusion process is represented by x_t , which is defined by the following formula:

$$x_t = \sqrt{\alpha_t} \cdot x_0 + \sqrt{1 - \alpha_t} \cdot z, \quad z \sim N(0, I)$$

Where $\alpha_t \in (0,1)$ is the noise scheduling parameter corresponding to time step t . In SNIR, we introduce the dimensional attention weight vector $w \in [0,1]^d$ and use it to adjust the amplitude of noise injection so that the key dimensions are less perturbed:

$$x_t = \sqrt{\alpha_t} \cdot x_0 + \sqrt{1 - \alpha_t} \cdot (w \otimes z)$$

The symbol $\otimes$ represents the element-wise multiplication operation.

To further enhance the selectivity of noise injection, we regularize the attention weight vector to keep it dynamically learned under certain distribution constraints:

$$\sum_{i=1}^d w_i = d, \quad w_i \geq 0$$

At the same time, during the training process, an auxiliary loss term is designed to guide the selectivity of noise injection:

$$L_{snir} = \lambda \cdot KL(w \| 1)$$

Where KL represents the Kullback-Leibler divergence, which is used to constrain the difference between the current weight distribution and the uniform distribution, and λ is the adjustment coefficient.

Through the above mechanism, SNIR significantly enhances the ability to protect the input structure while maintaining the ability to generate the diffusion model. In actual request data modeling, this strategy can effectively alleviate the problem of important structure destruction, making the subsequent reverse diffusion process more stable and expressive, and laying a solid foundation for high-quality reconstruction results. This method not only improves the reconstruction accuracy of the model, but also provides a more detailed and controllable basis for anomaly detection.

3.2 Feature-aware Discriminative Scoring

To further improve the discriminative performance of anomaly request detection, this study introduces the Feature-aware Discriminative Scoring (FDS) mechanism. It performs multidimensional anomaly measurement by capturing deep semantic differences between the original input features and the reconstructed outputs. After the reconstruction process by the diffusion model, FDS applies parallel feature extractors to encode both the original request and its reconstructed version. This allows the model to detect structural shifts in the high-dimensional feature space that are caused by anomalies. Unlike traditional methods that rely only on reconstruction error, FDS focuses on semantic deviations in the request. This enhances the model's ability to identify subtle and weak anomalies. Its module architecture is shown in Figure 3.




Figure 3. FDS module architecture

Suppose the original request is x_0 , and the reconstruction result generated by back diffusion is $\bar{x}_0$. Both are embedded through the feature encoding network $f(\cdot)$ with shared parameters:

$$h_x = f(x_0), \quad h_{\hat{x}} = f(\bar{x}_0)$$

The FDS mechanism compares the similarity between the original features and the reconstructed features and uses this to measure the degree of anomaly. The basic measurement method is the Euclidean distance:

$$D_{euc} = \|h_x - h_{\hat{x}}\|_2$$

In addition, we introduce cosine similarity as a supplementary metric to enhance the ability to capture directional deviations:

$$D_{cos} = 1 - \frac{h_x^T h_{\hat{x}}}{\|h_x\|_2 \cdot \|h_{\hat{x}}\|_2}$$

On this basis, in order to avoid the limitations of a single metric, we design a weighted combination of anomaly score functions:

$$A(x_0) = \alpha \cdot D_{euc} + (1 - \alpha) \cdot D_{cos}, \quad \alpha \in [0,1]$$

Where α is the weight coefficient, which can adjust the contribution ratio of different similarity indicators. The final anomaly score not only takes into account the reconstruction error itself, but also integrates the differences in the latent semantic representation layer, thereby showing better discrimination ability in complex and difficult-to-perceive request anomalies.

In addition, to improve the feature extractor's ability to align normal patterns, the FDS mechanism introduces an embedding alignment regularization term during the training process to encourage the embedding space of normal data to be as consistent as possible:

$$L_{align} = \|h_x - h_{\hat{x}}\|_1$$

This regularization term promotes the stability of feature alignment and is embedded as an auxiliary constraint term in

the overall training objective, synergistically improving the accuracy of anomaly detection from both structural and semantic dimensions. The FDS mechanism builds an anomaly measurement system that takes into account both spatial distance and directional consistency through a deep fusion comparison of the reconstructed representation and the original representation, providing strong support for anomaly detection in high-dimensional complex data scenarios.

4. Experimental Results

4.1 Dataset

The primary dataset used in this study is UNSW-NB15. This dataset is a comprehensive network traffic dataset designed for cybersecurity research. It is widely used for anomaly detection and intrusion recognition tasks. The dataset contains a mix of real and synthetic traffic. It includes multiple attack types as well as normal traffic samples. Due to its high complexity and representativeness, it is well-suited for evaluating model performance on high-dimensional and unstructured network requests.

The UNSW-NB15 dataset is composed of several subsets. Each record contains a large number of features, including basic TCP/IP protocol fields, content features, time-based statistical features, and behavioral features. The total number of dimensions reaches 49. It includes various attack categories such as analysis, backdoor, DoS, probing, and unauthorized access control. This provides good class diversity and structural complexity. The dataset effectively simulates real-world abnormal request scenarios.

After preprocessing, the dataset has been widely applied in both unsupervised and supervised learning tasks. This study selects UNSW-NB15 mainly due to its broad acceptance in anomaly detection research. It also contains a large number of unlabeled anomalous requests. These characteristics align well with the needs of this study, which is based on diffusion model reconstruction and discriminative mechanisms. The dataset provides strong support for testing model performance in simulated real request scenarios.

4.2 Experimental setup

In the experimental setup, the model is deployed in an environment built on the PyTorch framework. An NVIDIA GPU is used to accelerate both training and inference. The diffusion model uses a multi-step noise scheduling strategy. The forward process includes 1000 steps, and the reverse reconstruction is performed step by step through a neural network trained for this purpose. The feature extraction module and the discriminative mechanism share parameters to ensure semantic consistency. The model is trained using the Adam optimizer, with an initial learning rate of 0.001 and a batch size of 128. Anomaly scores are generated by the FDS module, based on similarity between features before and after reconstruction.

The training and testing sets follow the original dataset ratio. This ensures that the distribution of attack and normal samples in the test set remains natural. No artificial anomalies are introduced. The training process does not use any anomaly samples. It models only normal requests, simulating real-world unsupervised anomaly detection scenarios.

All evaluation metrics, including AUC, F1, and Precision, are calculated on the same test set. The results are compared with those of classical reconstruction-based methods. Table 1 lists the specific parameters used in the experimental setup.

Table 1: Specific parameter

Parameter name	Setting Value
Framework	Pytorch 2.0
GPU	NVIDIA RTX A6000
Optimizer	Adam
Learning Rate	0.001
BatchSize	128
Diffusion Steps	1000
Feature Extractor	3-layer MLP
Training Epochs	200

4.3 Experimental Results

1) Comparative experimental results

This paper first gives the comparative experimental results, as shown in Table 2.

Table2: Comparative experimental results

Method	AUC	Precision	F1-Score
DAGMM[26]	0.872	0.801	0.784
DevNet[27]	0.893	0.819	0.812
Gods[28]	0.867	0.778	0.769
Anomaly Transformer[29]	0.912	0.841	0.834
Ours	0.941	0.869	0.879

As shown in the comparative results in Table 2, the proposed method outperforms several mainstream anomaly detection approaches across multiple key evaluation metrics. It shows particularly strong performance in AUC and F1-Score. As a measure of the model's overall discriminative ability, the AUC of this method reaches 0.941, significantly higher than that of other models. This indicates that the designed discriminative mechanism can more effectively distinguish between normal and abnormal requests. The improvement is attributed to the diffusion model's ability to capture fine-grained structure within the request distribution.

Further analysis of the trends in Precision and F1-Score reveals a common trade-off between precision and recall in traditional generative or discriminative models when handling complex anomalies. For example, although DAGMM and Gods have certain generative and discriminative capabilities, they often misclassify requests with subtle anomaly patterns. In contrast, the Selective Noise Injection (SNIR) mechanism introduced in this study selectively preserves key feature dimensions. This enhances the model's reconstruction of normal requests and reduces false positives. As a result, the Precision reaches 0.869.

At the same time, the increase in F1-Score shows that the proposed method improves recall while maintaining high precision. This improvement is closely related to the Feature-aware Discriminative Scoring (FDS) mechanism. By measuring semantic distances between original inputs and their reconstructed embeddings, FDS provides a multi-perspective evaluation of request abnormality. It strengthens the model's ability to detect anomalies in complex and ambiguous scenarios. This allows improved detection performance without relying on labeled data.

Overall, the experimental results demonstrate that combining diffusion models with structure-aware feature fusion offers a new approach to anomaly detection. Compared with traditional methods that rely only on reconstruction error or classifier output, the proposed method builds a more refined and stable anomaly detection framework. It leverages both the robustness of the generative process and the semantic consistency of the feature space. This makes it especially suitable for high-dimensional, unstructured, and dynamically evolving network request scenarios.

2) Ablation Experiment Results

This paper also presents the results of an ablation experiment designed to evaluate the individual contributions of key components within the proposed model. The purpose of this experiment is to analyze the effectiveness of each module by isolating and comparing different model configurations. Through this controlled setup, it becomes possible to understand how specific mechanisms impact the overall performance of the detection framework. The detailed experimental results and comparisons are provided in Table 3.

Table 3: Ablation Experiment Results

Method	AUC	Precision	F1-Score
BaseLine	0.901	0.821	0.813
+SNIR	0.918	0.838	0.829
+FDS	0.924	0.846	0.841
Ours	0.941	0.869	0.879

As shown in the ablation results in Table 3, the BaseLine model demonstrates a certain level of anomaly detection capability even without any enhancement modules. It achieves an AUC of 0.901 and an F1-Score of 0.813. This indicates that the diffusion model itself has a degree of expressive power in modeling the latent structure of request data. However, the model still shows recognition errors, especially when dealing with abnormal requests that have complex structural details or subtle semantic deviations. Its reconstruction and discrimination accuracy remain limited in such cases.

After introducing the Selective Noise Injection for Reconstruction (SNIR) module, the model performance improves significantly. The AUC rises to 0.918 and the F1-Score increases to 0.829. SNIR injects structure-aware noise into the input features during the diffusion process. This preserves key dimensional information, enabling more accurate reconstruction of normal requests. As a result, the reconstruction error range for normal samples is effectively reduced. This strategy enhances the model's expressive ability

during reconstruction, which indirectly improves its capacity to distinguish abnormal behaviors.

When the Feature-aware Discriminative Scoring (FDS) module is added to the BaseLine model alone, the AUC increases to 0.924 and the F1-Score reaches 0.841. This performance exceeds that of using SNIR alone. It indicates that FDS provides a more powerful semantic distance evaluation in anomaly discrimination. By assessing the differences in feature space representations between original and reconstructed requests, FDS complements the limitations of reconstruction error. It improves the model's sensitivity to detecting anomalies with subtle structural shifts.

Finally, when both SNIR and FDS are applied together, the full model achieves the best performance. The F1-Score increases to 0.879 and the Precision reaches 0.869. This result validates the effectiveness of combining both mechanisms in the model. SNIR ensures structural fidelity in the generative process, while FDS guides the model to better assess abnormality from a discriminative perspective. This dual-path strategy, combining generation and discrimination, reflects the study's focus on interpretability and robustness in anomaly request detection.

3) Hyperparameter sensitivity experiments

Furthermore, this paper gives the experimental results of hyperparameter sensitivity. First, the experimental results of learning rate are given, as shown in Table 4.

Table 4: Hyperparameter sensitivity experiment results (learning rate)

Learning Rate	AUC	Precision	F1-Score
0.004	0.911	0.833	0.826
0.003	0.926	0.847	0.839
0.002	0.934	0.858	0.867
0.001	0.941	0.869	0.879

As shown in the hyperparameter sensitivity results in Table 3, the learning rate has a significant impact on the training stability and final performance of the diffusion model in anomaly request detection. When the learning rate is set to a relatively high value of 0.004, the model achieves an AUC of 0.911 and an F1-Score of 0.826. This indicates that although the model retains some discriminative ability, a high learning rate may cause it to exit the optimal region too early during training, leading to insufficient or unstable convergence.

As the learning rate decreases to 0.003 and 0.002, the overall model performance improves steadily. The F1-Score reaches 0.839 and 0.867, respectively. This suggests that a moderate reduction in learning rate helps to optimize both the diffusion process and the discriminative module more precisely. At this stage, the Selective Noise Injection (SNIR) mechanism shows better effectiveness in preserving structural information. The Feature-aware Discriminative Scoring (FDS) module also achieves more consistent feature alignment, which enhances anomaly recognition accuracy.

When the learning rate is further reduced to 0.001, the model reaches its best performance. The AUC increases to

0.941 and the F1-Score rises to 0.879. This result indicates that with a lower learning rate, the model can more stably optimize the objectives of both the diffusion and discriminative stages. The combined effect of SNIR and FDS is more fully realized, resulting in a more powerful anomaly request detector. This is especially important in high-dimensional request spaces, where fine-grained feature relationships depend on stable parameter updates.

In summary, the learning rate plays a key role in the performance of the proposed method. A suitable learning rate not only improves reconstruction quality but also enhances semantic consistency in the feature space. This leads to better robustness when handling subtle deviations and structurally ambiguous anomaly requests. These findings further validate the adaptability and effectiveness of the proposed mechanisms within a finely optimized generative-discriminative framework.

Furthermore, the experimental results of different optimizers are given, as shown in Table 5.

Table 5: Hyperparameter sensitivity experiment results (Optimizer)

Optimizer	AUC	Precision	F1-Score
RMSProp	0.912	0.834	0.827
AdaGrad	0.897	0.812	0.804
SGD	0.881	0.798	0.791
Adam	0.941	0.869	0.879

As shown in the optimizer sensitivity results in Table 5, the choice of optimizer has a significant impact on model performance in this study. This is especially evident in the joint training of the diffusion model and the feature discrimination module. Differences in gradient update stability and convergence efficiency under different optimization strategies directly affect anomaly detection performance. Among all tested optimizers, SGD performs the worst, with an AUC of 0.881 and an F1-Score of 0.791. This suggests that SGD struggles to achieve high-quality convergence when training complex network structures with multiple components. It may result in insufficient reconstruction or poor feature alignment.

AdaGrad shows strong adaptive adjustment in the early training stage. However, it suffers from rapid learning rate decay during longer training, which limits further optimization. Its F1-Score reaches 0.804, still below the expected level. RMSProp alleviates the decay problem seen in AdaGrad and improves the AUC to 0.912. However, it still fails to maintain consistent performance across the multi-stage generative-discriminative structure. This indicates limitations in its gradient adjustment capability when applied to diffusion models.

In contrast, the Adam optimizer achieves the best performance in this task. It shows superior stability and adaptability compared to the other methods. The F1-Score reaches 0.879, and the Precision improves to 0.869. These results suggest that Adam's first- and second-order moment estimation mechanisms effectively support the joint optimization of the Selective Noise Injection (SNIR) and Feature-aware Discriminative Scoring (FDS) modules. It ensures smooth progression of the diffusion process and better alignment in the semantic feature space.

Overall, the analysis shows that the optimizer affects not only training efficiency but also the coordination between the generative and discriminative modules. For a high-capacity diffusion architecture like the one proposed in this study, using an adaptive optimizer such as Adam can better unlock the model's potential. It improves both the accuracy and robustness of anomaly request detection. These results also highlight the specific training requirements of diffusion-based generative models.

4) Robustness evaluation of the model under various attack types

This paper further gives the experimental results of the model under various attack types, and the experimental results are shown in Figure 4.




Figure 4. Experimental results of the model under various attack types

As shown in the results of Figure 4, the proposed method demonstrates strong robustness across different attack types, with consistently high F1-Scores. In particular, the model

performs exceptionally well on typical attack types such as DoS, Reconnaissance, and Fuzzers, achieving F1-Scores of 0.91, 0.89, and 0.88, respectively. These results indicate that

the method can effectively detect abnormal requests with high traffic intensity and sudden pattern shifts. This performance is attributed to the SNIR module, which preserves key feature dimensions during reconstruction, increasing the model's sensitivity to such structured attacks.

For more stealthy and sparse attacks such as Worms, Analysis, and Shellcode, the F1-Score shows a slight decrease but remains within the range of 0.80 to 0.83. This demonstrates that the proposed Feature-aware Discriminative Scoring (FDS) mechanism can still capture subtle semantic deviations in the feature space. This embedding-based modeling approach compensates for the limitations of traditional reconstruction error in detecting minor anomalies and enhances the model's discriminative ability against hidden attacks.

Moreover, the F1-Score distribution curve shows that performance fluctuations across different attack types remain small. This reflects the model's strong consistency and generalization ability. Such stability is particularly important in real-world deployment, where attack traffic is highly diverse





Figure 5. Performance sensitivity experiment under the change of unlabeled data ratio

As shown in the results of Figure 5, the model's performance varies under different proportions of unlabeled data. This reflects the sensitivity of anomaly detection tasks to supervision in the training samples. When the proportion of unlabeled data is 10 percent or 90 percent, both F1-Score and AUC are at lower levels. This indicates that too few or too many unlabeled samples hinder the effective learning of both the diffusion model and the discriminative module. In particular, when 90 percent of the data is unlabeled, the supervision signal becomes too sparse. The model fails to form a stable decision boundary.

The model achieves the best performance when the proportion of unlabeled data is 50 percent. The F1-Score rises to 0.869 and the AUC reaches 0.941. At this level, the model benefits from a sufficient amount of unlabeled data to learn the underlying structure, while the remaining labeled data provides enough supervision to train the discriminative layer. This demonstrates the proposed method's adaptability in weakly supervised environments. Especially within the Feature-aware Discriminative Scoring (FDS) module, anomaly recognition is

and unpredictable. A model that is heavily biased toward certain attack types may result in serious security gaps. The proposed method maintains reliable detection performance under various attack conditions, indicating its practical value in highly dynamic request environments.

In summary, this experiment further validates the broad adaptability of the proposed diffusion-discrimination fusion mechanism across diverse anomaly patterns. Through the combined guidance of structure-aware noise control and semantic feature alignment, the model improves the expressiveness of reconstruction and enhances the consistency of semantic features during discrimination. This enables stable and efficient recognition of a wide range of attack behaviors.

5) Performance sensitivity experiment under the change of unlabeled data ratio

This paper also presents a performance sensitivity experiment under the change of the proportion of unlabeled data, and the experimental results are shown in Figure 5.

achieved through semantic embedding rather than relying on label-intensive classification.

Additionally, the two subplots show that when the proportion of unlabeled data increases from 30 percent to 50 percent, model performance improves rapidly. However, as it increases to 70 percent, performance begins to decline slightly, although it remains relatively high. This suggests that while the SNIR module enhances the model's ability to retain structural information, the reconstruction process still benefits from some level of label guidance. Without enough supervision, the discriminative module cannot adjust effectively, leading to a drop in overall detection performance.

These experimental results further validate the advantages of the proposed generative-discriminative fusion mechanism in low-label settings. Through sensitivity analysis of the unlabeled data ratio, it is observed that the model can achieve an optimal balance between structural modeling and semantic discrimination under certain weak supervision conditions. This reflects the effectiveness of the coordinated design between the diffusion model and the discriminative enhancement

mechanism. It also supports the method's applicability to real-world scenarios where labeled data is limited.

5. Conclusion

This paper presents an unsupervised anomaly request detection method that integrates diffusion-based generation and discriminative enhancement. The approach leverages the high-quality reconstruction and structural representation capabilities of diffusion models to effectively capture the latent distribution of normal requests. A structure-aware mechanism, Selective Noise Injection for Reconstruction (SNIR), is introduced to preserve critical feature dimensions during the forward diffusion process. This enhances the model's ability to represent request details. Compared with traditional reconstruction-based methods, this strategy maintains generation quality while better preserving the structure of the original data. It provides a solid foundation for distinguishing abnormal requests.

In addition, the paper proposes a Feature-aware Discriminative Scoring (FDS) module. After reconstruction, the module compares embedded features of original and reconstructed requests to measure semantic deviations. This enables a more robust anomaly scoring system. The method addresses the sensitivity of traditional reconstruction error to noise and improves detection performance for complex boundary cases and stealthy attacks. Experimental results show that the proposed method outperforms existing approaches across key metrics. It offers stronger generalization and more stable discrimination, particularly under diverse attack types and weak supervision conditions.

This study not only introduces a novel integration of diffusion generation and discriminative mechanisms at the methodological level but also demonstrates strong scalability in practical implementation. The method does not rely on labeled training data, making it suitable for large-scale real-world systems for online request monitoring and intelligent anomaly detection. It has broad application potential in fields such as cybersecurity, API management, and microservice monitoring. With its modular design, the system can be flexibly integrated into existing log collection and detection frameworks. It also supports deployment in various architectures, including streaming detection systems, enhancing its practical feasibility.

Future work can be explored in several directions to further improve the flexibility and applicability of the proposed method. One promising direction is the integration of multimodal input data, which may include parameter content, request paths, and behavioral context information. By combining different types of input features, the model could achieve a deeper and more comprehensive understanding of complex attack behaviors. This multimodal fusion would allow the detection framework to capture a wider range of semantic and structural patterns, enabling it to handle more diverse and sophisticated request scenarios.

Another valuable direction involves enhancing the model architecture through the use of advanced neural network components. For example, graph neural networks could be employed to model the inherent relationships between different requests, while attention mechanisms could help the model

focus on the most relevant parts of the input. Additionally, investigating lightweight diffusion approximations could lead to improved inference speed and better resource efficiency. These advancements would support the integration of the method into real-time industrial applications. Overall, the proposed framework contributes to the development of anomaly detection modeling and offers a new approach for intelligent security protection in high-dimensional and complex environments.

References

- [1] Wang S, Balarezo J F, Kandeepan S, et al. Machine learning in network anomaly detection: A survey[J]. *IEEE Access*, 2021, 9: 152379-152396.
- [2] Abdallah A M, Alkaabi A S R O, Alameri G B N D, et al. Cloud network anomaly detection using machine and deep learning techniques—Recent research advancements[J]. *IEEE Access*, 2024, 12: 56749-56773.
- [3] Lim W, Yong K S C, Lau B T, et al. Future of generative adversarial networks (GAN) for anomaly detection in network security: A review[J]. *Computers & Security*, 2024, 139: 103733.
- [4] Kim H, Lee B S, Shin W Y, et al. Graph anomaly detection with graph neural networks: Current status and challenges[J]. *IEEE Access*, 2022, 10: 111820-111829.
- [5] Hooshmand M K, Hosahalli D. Network anomaly detection using deep learning techniques[J]. *CAAI Transactions on Intelligence Technology*, 2022, 7(2): 228-243.
- [6] Xu W, Jang-Jaccard J, Singh A, et al. Improving performance of autoencoder-based network anomaly detection on nsl-kdd dataset[J]. *IEEE Access*, 2021, 9: 140136-140146.
- [7] Ding K, Zhou Q, Tong H, et al. Few-shot network anomaly detection via cross-network meta-learning[C]//*Proceedings of the web conference* 2021. 2021: 2448-2456.
- [8] Jain M, Kaur G, Saxena V. A K-Means clustering and SVM based hybrid concept drift detection technique for network anomaly detection[J]. *Expert Systems with Applications*, 2022, 193: 116510.
- [9] Yang L, Zhang Z, Song Y, et al. Diffusion models: A comprehensive survey of methods and applications[J]. *ACM Computing Surveys*, 2023, 56(4): 1-39.
- [10] Chen S, Sun P, Song Y, et al. Diffusiondet: Diffusion model for object detection[C]//*Proceedings of the IEEE/CVF international conference on computer vision*. 2023: 19830-19843.
- [11] Cao H, Tan C, Gao Z, et al. A survey on generative diffusion models[J]. *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [12] Lin L, Li Z, Li R, et al. Diffusion models for time-series applications: a survey[J]. *Frontiers of Information Technology & Electronic Engineering*, 2024, 25(1): 19-41.
- [13] Fuest M, Ma P, Gui M, et al. Diffusion models and representation learning: A survey[J]. *arXiv preprint arXiv:2407.00783*, 2024.
- [14] Livernoche, V., Jain, V., Hezaveh, Y., & Ravanbakhsh, S. (2023). On diffusion modeling for anomaly detection. *arXiv preprint arXiv:2305.18593*.
- [15] Bercea C I, Neumayr M, Rueckert D, et al. Mask, stitch, and re-sample: Enhancing robustness and generalizability in anomaly detection through automatic diffusion models[J]. *arXiv preprint arXiv:2305.19643*, 2023.
- [16] Han J, Feng S, Zhou M, et al. Diffusion Model in Normal Gathering Latent Space for Time Series Anomaly Detection[C]//*Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Cham: Springer Nature Switzerland, 2024: 284-300.
- [17] Zuo H, Zhu A, Zhu Y, et al. Unsupervised diffusion based anomaly detection for time series[J]. *Applied Intelligence*, 2024, 54(19): 8968-8981.
- [18] Mousakhan, A., Brox, T., & Tayyub, J. (2024, September). Anomaly detection with conditioned denoising diffusion models. In *DAGM German Conference on Pattern Recognition* (pp. 181-195). Cham: Springer Nature Switzerland.

- [19] Flaborea A, Collorone L, Di Melendugno G M D A, et al. Multimodal motion conditioned diffusion model for skeleton-based video anomaly detection[C]//Proceedings of the IEEE/CVF international conference on computer vision. 2023: 10318-10329.
- [20] Hu T, Zhang J, Yi R, et al. Anomalydiffusion: Few-shot anomaly image generation with diffusion model[C]//Proceedings of the AAAI conference on artificial intelligence. 2024, 38(8): 8526-8534.
- [21] Cheng K, Pan Y, Liu Y, et al. Denoising diffusion-augmented hybrid video anomaly detection via reconstructing noised frames[C]//Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence. 2024: 695-703.
- [22] Tong A, Wolf G, Krishnaswamy S. Fixing bias in reconstruction-based anomaly detection with lipschitz discriminators[J]. Journal of Signal Processing Systems, 2022, 94(2): 229-243.
- [23] Liu W, Chang H, Ma B, et al. Diversity-measurable anomaly detection[C]//Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2023: 12147-12156.
- [24] Marimont, S. N., Baugh, M., Siomos, V., Tzelepis, C., Kainz, B., & Tarroni, G. (2024, May). Disyre: Diffusion-inspired synthetic restoration for unsupervised anomaly detection. In 2024 IEEE International Symposium on Biomedical Imaging (ISBI) (pp. 1-5). IEEE.
- [25] Huang D, Shen L, Yu Z, et al. Efficient time series anomaly detection by multiresolution self-supervised discriminative network[J]. Neurocomputing, 2022, 491: 261-272.
- [26] Zong B, Song Q, Min M R, et al. Deep autoencoding gaussian mixture model for unsupervised anomaly detection[C]//International conference on learning representations. 2018.
- [27] Zheng H, Sun D, Han X, et al. Research on Network Security Intrusion Detection Based on Devnet[C]//Proceedings of the 3rd International Conference on Signal Processing, Computer Networks and Communications. 2024: 251-256.
- [28] Wang J, Cherian A. Gods: Generalized one-class discriminative subspaces for anomaly detection[C]//Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019: 8201-8211.
- [29] Xu J, Wu H, Wang J, et al. Anomaly transformer: Time series anomaly detection with association discrepancy[J]. arXiv preprint arXiv:2110.02642, 2021.