Proceedings of the 4th International Conference on Signal Processing and Machine Learning
DOI: 10.54254/2755-2721/53/20241281

An improved BiGAN model for anomaly detection in finance

Guanda Chen

Warwick Mathematics Institute, University of Warwick, Coventry, CV4 7AL, United Kingdom

Guanda.Chen@warwick.ac.uk

Abstract. Financial systems play a pivotal role in shaping contemporary society, and the detection of financial anomalies holds immense significance in mitigating the adverse repercussions of market uncertainties on the global economy. In this context, this study presents an innovative LSTM-GANs model, specifically crafted to enhance the detection of anomalies in financial stock markets. The model introduces an "Anomaly Score" as a pivotal metric, which is computed through a combination of factors such as Reconstruction Loss, Latent Space Distance, and Discriminator Score. This composite score provides a quantitative assessment of the anomaly level within the financial data. By applying a predefined threshold to this Anomaly Score, the model efficiently identifies and flags anomalies. In a world where financial markets are increasingly complex and prone to unexpected events, the ability to detect and respond to anomalies swiftly is paramount. This novel LSTM-GANs model offers a promising approach to bolster the accuracy and effectiveness of financial anomaly detection, thereby contributing to the stability and resilience of global financial systems.

Keywords: GAN, Anomaly detection, Quantitative finance.

1. Introduction

Financial systems are one of the most crucial factors in the development of modern society. They integrate capital in a rational manner, enabling various industries to operate and thus promoting societal development. Therefore, the significance of maintaining a stable financial system cannot be overstated in the pursuit of social development. Thereinto, anomalies refer to unexpected changes or deviations from expected patterns in datasets [1, 2]. The identification of such anomalies is essential for items, events, or observations that do not match the expected patterns in a dataset. It serves as an important tool for maintaining a stable financial system. In finance, anomaly detection can be used not only for identifying fraudulent activities but also in algorithmic trading to optimize strategies by recognizing abnormal fluctuations. The challenge of financial anomaly detection has been long-standing and holds great potential for further development.

Generative Adversarial Networks were introduced by Ian Goodfellow in 2014 [3]. Consisting of the generator and the discriminator, the generator aims to generate fake data as close as possible to a real data distribution from random noise. The discriminator's task is to judge whether the input data is from a real dataset or generated by the generator. They train in an adversarial manner, with the generator attempting to generate realistic data to "fool" the discriminator, and the discriminator striving to

DOI: 10.54254/2755-2721/53/20241281

distinguish between real and fake data. Since its inception, GANs have demonstrated remarkable results in image generation and data augmentation and have wide-ranging applications.

Financial systems are inherently complex, yet years of research on financial time-series data have revealed some statistical regularities known as "stylized facts," such as linear unpredictability, fat-tailed distribution, volatility clustering, and gain/loss asymmetry. GANs, including QuantGAN, MasGAN, and DCGAN, have shown promising results in generating and processing financial time-series data [4]. They can accurately replicate these stylized facts, indicating great potential for GANs in finance.

Subsequent research has shown that GANs also hold immense potential in anomaly detection [5]. Models such as Anogan, bigan, egbad, and Ganormality have been progressively developed for this purpose, showing favorable results. Additionally, because GANs can capture high-level features and intrinsic structures of the data, these models have the capability achieving higher accuracy rates than traditional rule-based or statistical anomaly detection methods [6].

Based on the above context, this study aims to use GANs for anomaly detection in financial prices, specifically, a lstm-bigan model is built, it has an additional encoder that maps real data to latent representation. And a anomaly score is defined to represent the anomaly level.

2. Method

2.1. Dataset preparation

2.1.1. Dataset introduction

The dataset of the study source from the S&P 500 in Yahoo Finance. Originally, the dataset is five-dimensional, featuring the following attributes: 'Open', 'High', 'Low', 'Close', 'Adj Close', and 'Volume'. These attributes represent typical, publicly available information about stock performance. The S&P 500 is a stock market index that tracks the performance of the 500 largest companies in the United States, making it a reasonable choice for the data in our experiment. Selected features

2.1.2. Dataset preprocessing

The minmax normalisation is first applied to keep the data in the same range. Important features for time-series are created additionally including daily return percentage, 5-day closing price average, 5-day return percentage. To ensure the model can learn the temporal dependencies from the data, the 30-days windows is created with sliding stride 1 day. Following the above data preprocessing, the data comprises 9 features. The data is then splitting into training set, validating set and testing set with the ratio of 0.64, 0.16, 0.20 respectively.

2.2. Architecture

The model in this study is an extension of BiGan model used for anomaly detection in time series data. The traditional BiGan model consist of a Generator, a Discriminator and an Encoder. The additional component encoder learns to map the original data x to a latent representation z during the Gan training process.

Although the original BiGAN architecture has demonstrated effectiveness across diverse datasets [7], for instance, MNIST and KDD99 dataset are both behaving well, it is hard to capture the temporal dependencies present in time series data. To address this limitation, this study introduced Long Short-Term Memory (LSTM) layers [8, 9] that is a valuable tool for capturing the information from the time series data in all three components of the BiGan model.

2.2.1. Generator

The generator shown in Figure 1 is used to generate synthetic data, in this case, the goal is to generate a S&P stock price and trick the discriminator. It takes a random noise from a standard normal distribution as input, passing through a fully connected layer, activation layer of Leaky ReLU, drop out layer,

DOI: 10.54254/2755-2721/53/20241281

reshape layer and followed by a LSTM layer, then it passes to the output layer to generate time series data.

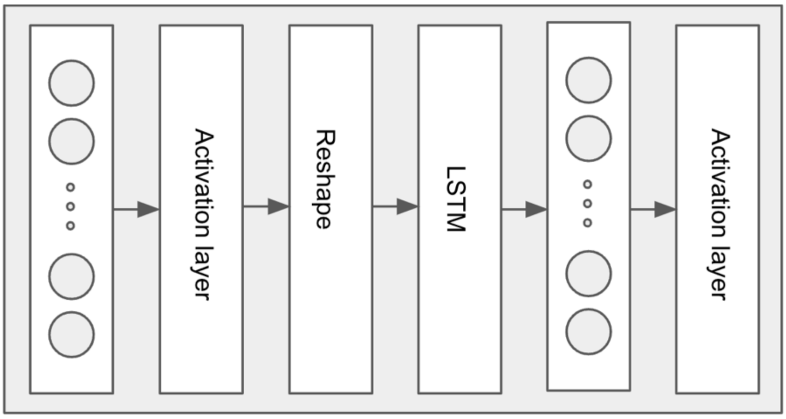

Figure 1. LSTM-BiGAN generator (Photo/Picture credit: Original)

2.2.2. Discriminator

The Discriminator shown in Figure 2 aims to distinguish between real and synthetic data. It is different from the traditional discriminator, in the design, it takes the real or synthetic data and their corresponding latent information as input. The structure first consists of a LSTM layer receiving the real or synthetic data and a fully-connected layer dealing with the latent input followed by the activation layer leaky ReLU and a drop out layer for regularization. Finally, the sigmoid activation function gives the probability of real or fake.

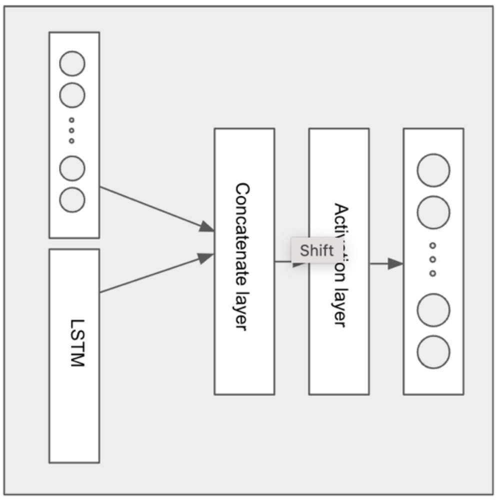

Figure 2. LSTM-BiGAN discriminator (Photo/Picture credit: Original)

2.2.3. Encoder

The Encoder shown in Figure 3 is responsible for transforming real data into the latent space, enhancing the model's capacity to comprehend the underlying structure of the data. It consists of an LSTM layer followed by a LeakyReLU activation layer and a dropout layer for regularization. The output is a latent variable z with the same dimension as the input to the Generator.

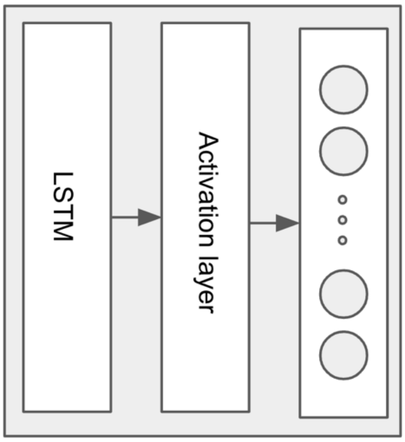

Figure 3. LSTM-BiGAN Encoder (Photo/Picture credit: Original)

The combination of the three models is the proposed LSTM-BiGAN model. Ideally, the model will have a good understanding of the data after training. The following part introduces the methodology of anomaly detection for the time series data.

2.3. Anomaly detection

Assuming the model possesses a solid grasp of the data, an anomaly score is defined in the study. Reconstruction loss, Latent Space Distance, and Discriminator Score serve as effective indicators of anomaly levels, as they take into account the quality of data reconstruction, the consistency of latent representation, and the model's assessment of data normality. And this study therefore used the three metrics to define the anomaly score. The following are the detailed method.

$$\text{Anomaly Score} = w_1 \times \text{Reconstruction Error} + w_2 \times \text{Latent Distance} + \text{w\_3} \times \text{Discriminator score} \quad (1)$$

2.3.1. Reconstruction loss

The reconstruction loss quantifies the dissimilarity between the original data and the reconstructed data. Normal data and anomalous data behave differently in the latent representation. So, the reconstruction loss of the anomalous data is usually higher. The encoder maps the original data X to $Z = E(X)$. The latent representation is then used to reconstruct the data $\hat{X}$. where $\hat{X} = G(Z)$. The Mean squared errors is used as the metric for the reconstruction loss.

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}\left(X_i - \hat{X_\iota}\right)^2 \quad (2)$$

2.3.2. Latent space distance

The latent space distance measures the difference of the real data and generated data from the inherent perspective. As above, this study encoded the real data to its latent representation $E(X)$. And then this study used MSE as the metric to measure latent space loss. Since the generator takes value from the normal distribution, this study therefore assumed the random normal data Z as the latent representation of the generated data

$$\text{Latent Distance} = \frac{1}{N}\sum_{i=1}^{N}(E(X_i) - Z_i)^2 \quad (3)$$

2.3.3. Discriminator score

The discriminator will receive the real data X, and calculate the probability of the data being real using $D(X, E(X))$. The probability output: Discriminator Score $= D(X, E(X))$will record the anomalous level from the perspective of discriminator.

DOI: 10.54254/2755-2721/53/20241281

3. Results and discussion

The BiGAN model was trained and evaluated on the S&P 500 dataset. The anomaly scores are computed for each data point, with a mean anomaly score of 0.7621 and a standard deviation of 0.0979. A threshold for anomaly detection was set at 1.0556, calculated as the mean plus three times the standard deviation of the anomaly scores.

The model identified a total of 106 anomalies in the entire dataset, accounting for approximately 1.25% of the data. The distribution of these anomaly scores is visualized in Figure 4 below, which shows a Gaussian-like distribution centered around the mean.

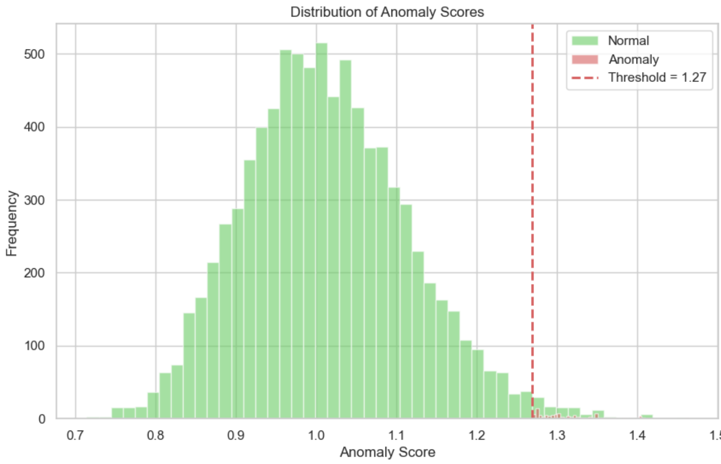

Figure 4. Distribution of Anomaly Score (Photo/Picture credit: Original)

The anomalies detected were plotted over the time series data of the S&P 500 'Close' prices shown in Figure 5. The visual inspection suggests that the model was able to flag significant spikes or drops in the stock prices, which could be indicative of market volatility or specific financial events. However, due to the lack of labeled data, a quantitative evaluation of the model's performance could not be performed.

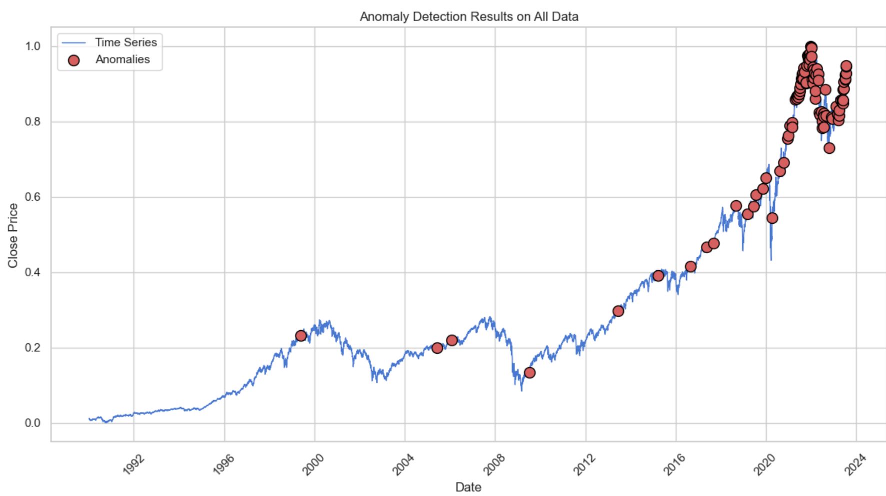

Figure 5. Visualizing Anomaly on S&P 500 (Photo/Picture credit: Original)

Applying BiGAN to anomaly detection in time series data is an emerging research area that presents several challenges. One issue is mode collapse [10], where the discriminator tends to outperform the generator, as evidenced by the discriminator's high model accuracy in our experiments. Another significant challenge is the scarcity of labeled data for anomaly detection. Defining what constitutes an

DOI: 10.54254/2755-2721/53/20241281

anomaly in time series data is not straightforward. In our study, this study defined anomalies as values exceeding the mean plus three times the standard deviation of the anomaly scores, a criterion that lacks theoretical backing. This makes it difficult to quantitatively evaluate the model's performance. Additionally, like most neural networks, the components of the BiGAN model—namely the generator, discriminator, and encoder—are not easily interpretable.

4. Conclusion

The research introduces a GAN model designed for detecting anomalies in the finance domain. More precisely, a LSTM-BiGAN model is developed by extending the conventional BiGAN model. The Anomaly score is determined through a combination of factors, including the reconstruction loss, Latent Space Distance, and Discriminator Score.

In forthcoming work, there are plans to comprehensively assess the model's performance using various evaluation methods. To obtain quantitative results such as accuracy scores and F1 scores, it will be essential to evaluate the model using labeled data. This approach will provide a more thorough understanding of the model's effectiveness in anomaly detection within the finance sector.

References

[1] Eskin E 2000 Anomaly detection over noisy data using learned probability distributions

[2] Blázquez-García A et al 2021 A review on outlier/anomaly detection in time series data. ACM Computing Surveys (CSUR) 54(3) 1-33

[3] Goodfellow I 2014 Generative adversarial networks Advances in neural information processing systems

[4] Eckerli F 2021 Generative Adversarial Networks in finance: an overview arXiv preprint arXiv:2106.06364.

[5] Mattia F D 2019 A Survey on GANs for Anomaly Detection arXiv preprint arXiv:1906.11632.

[6] Bashar M A 2023 ALGAN: Time Series Anomaly Detection with Adjusted-LSTM GAN. arXiv preprint arXiv:2308.06663

[7] Zenati H Foo C S Lecouat B Manek G and Chandrasekhar V R 2018 EFFICIENT GAN-BASED ANOMALY DETECTION. ICLR.

[8] Yu Y et al 2019 A review of recurrent neural networks: LSTM cells and network architectures. Neural computation 31(7) 1235-1270

[9] Staudemeyer R C and Morris E R 2019 Understanding LSTM--a tutorial into long short-term memory recurrent neural networks arXiv preprint arXiv:1909.09586

[10] Srivastava A et al 2017 Veegan: Reducing mode collapse in gans using implicit variational learning. Advances in neural information processing systems 30