INTERNATIONAL RESEARCH JOURNAL OF ADVANCED ENGINEERING AND TECHNOLOGY (IRJAET)

eISSN: 3087-4068

Volume. 02, Issue. 09, pp. 01-07, September 2025"

A Comprehensive Review and Empirical Assessment of Data Augmentation Techniques in Time-Series Classification

Dr. Elena M. Petrovic
Department of Informatics, ETH Zurich, Switzerland

Dr. Rajan V. Subramaniam
Department of Computer Science and Engineering, Indian Institute of Technology (IIT) Bombay, India

Article received: 05/07/2025, Article Revised: 06/08/2025, Article Accepted: 01/09/2025
DOI: https://doi.org/10.55640/irjaet-v02i09-01

ABSTRACT

Time-series data is ubiquitous across various domains, from healthcare and finance to industrial monitoring and human activity recognition. The accurate classification of such data is crucial for informed decision-making and automated systems. However, a common challenge in developing robust time-series classification models, especially deep learning-based ones, is the scarcity of sufficiently large and diverse labeled datasets. Data augmentation has emerged as a powerful technique to address this limitation by synthetically expanding the training data, thereby enhancing model generalization and reducing overfitting. While data augmentation has been extensively studied in domains like image processing and natural language processing, its application and effectiveness in time-series classification present unique challenges and opportunities. This article provides a comprehensive survey of existing data augmentation techniques specifically tailored for time-series classification. Furthermore, it synthesizes empirical findings from a wide range of studies, discussing the efficacy of different augmentation strategies across various datasets and model architectures. We categorize augmentation methods, analyze their underlying principles, and highlight their impact on classification performance. Finally, we identify current limitations and propose future research directions to foster the development of more effective and universally applicable time-series data augmentation methodologies.

KEYWORDS

Time-Series Classification, Data Augmentation, Empirical Evaluation, Deep Learning, Synthetic Data Generation, Machine Learning, Temporal Data.

INTRODUCTION

Time-series data, characterized by sequential measurements over time, is a fundamental data type in numerous scientific and engineering disciplines. Its prevalence spans diverse applications, including but not limited to, medical diagnostics (e.g., Electrocardiogram (ECG) signals) [17, 36, 47, 53, 58], human activity recognition from wearable sensors [6, 7, 34, 41], financial market prediction [16], industrial fault detection [31], and communication signal analysis [26, 28, 33]. The goal of time-series classification (TSC) is to assign a predefined label to an entire time series or a segment thereof. With the advent of deep learning, particularly Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) [3, 10], significant advancements have been made in TSC. These models, however, are inherently data-hungry and typically require large volumes of labeled data to achieve optimal performance and avoid overfitting [5].

A persistent challenge in many real-world TSC scenarios is the limited availability of high-quality, labeled time-series data. Data collection can be expensive, time-consuming, or inherently difficult due to privacy concerns, rare events, or sensor limitations. This data scarcity often leads to models that generalize poorly to unseen data, suffering from high variance and low robustness. Data augmentation, a technique involving the creation of new, synthetic training examples from existing ones, offers a promising solution to mitigate these issues. By artificially expanding the dataset, augmentation helps expose the model to a wider variety of data patterns, thereby improving its ability to learn robust features and generalize effectively [1].

While data augmentation has been a cornerstone in the success of deep learning in computer vision [1, 2, 5, 20] and has seen increasing adoption in natural language processing [21], its application to time-series data is more nuanced. Unlike images, where transformations like rotation and flipping maintain semantic meaning, time-series data possesses inherent temporal dependencies and structural characteristics that must be preserved during augmentation to ensure the generated samples remain realistic and relevant to the original class [8]. Arbitrary transformations can easily distort the underlying patterns, leading to corrupted or misleading synthetic data that can degrade model performance rather than enhance it.

Several reviews have touched upon aspects of time-series classification and deep learning [3, 9, 10], and some have broadly surveyed data augmentation techniques [8, 19]. However, a dedicated and extensive empirical study combined with a comprehensive survey specifically focusing on the impact and efficacy of various data augmentation strategies for time-series classification using neural networks is still needed. This paper aims to fill this gap by providing an in-depth analysis of the current landscape of time-series data augmentation techniques, synthesizing empirical evidence of their performance, and offering a structured overview for researchers and practitioners.

The remainder of this article is structured as follows: Section 2 details the methodology employed for this survey and the empirical assessment. Section 3 presents the identified data augmentation techniques and summarizes their empirical performance across various studies. Section 4 discusses the implications of these findings, addresses current limitations, and outlines future research directions. Finally, Section 5 concludes the article.

METHODS

This section outlines the systematic approach undertaken to conduct this comprehensive review and synthesize empirical findings on data augmentation for time-series classification. The methodology is designed to ensure a broad coverage of relevant literature and a structured analysis of the identified techniques and their reported performance.

2.1 Literature Search Strategy

Our literature search was primarily conducted using a combination of keyword-based searches across major scientific databases (e.g., IEEE Xplore, ACM Digital Library, arXiv, Google Scholar) and a snowballing approach [22]. Initial keywords included "time series data augmentation," "time series classification," "deep learning time series," "synthetic time series," and combinations thereof. The snowballing technique involved examining the reference lists of highly relevant papers and reviews to identify additional pertinent studies. This iterative process ensured a comprehensive collection of research articles, including recent preprints from arXiv. The search encompassed publications from 2012 onwards, coinciding with the rise of deep learning in various domains [5].

2.2 Inclusion and Exclusion Criteria

To maintain focus and relevance, specific inclusion and exclusion criteria were applied:

- Inclusion Criteria:
  - o Papers explicitly discussing data augmentation techniques for time-series data.
  - o Studies that empirically evaluate the impact of data augmentation on time-series classification tasks.
  - o Research involving deep learning models (e.g., CNNs, LSTMs) for TSC.
  - o Articles published in peer-reviewed conferences, journals, or reputable preprint archives.
- Exclusion Criteria:
  - o Papers focusing solely on data augmentation for other time-series tasks (e.g., forecasting, anomaly detection) without classification.
  - o Studies that do not involve deep learning models.
  - o Review papers that do not present novel augmentation techniques or empirical results (though these were used for snowballing).
  - o Papers not available in English.

2.3 Categorization of Augmentation Techniques

Based on the identified literature, a taxonomy of time-series data augmentation techniques was developed. This categorization helps in systematically understanding the diverse approaches and their underlying mechanisms. The techniques are broadly classified into:

- Transformation-based methods: These involve applying various transformations directly to the raw time-series data.

• Generative model-based methods: These leverage generative models to create entirely new synthetic time series.

• Feature-space augmentation: Techniques that operate on extracted features rather than raw data.

• Hybrid approaches: Combinations of the above methods.

Each identified technique was analyzed for its operational principle, common variations, and reported benefits and drawbacks.

2.4 Empirical Findings Synthesis

Given that this is a survey paper and not a new empirical study, this section describes how the empirical findings from the reviewed literature were synthesized. For each relevant paper, information regarding the augmentation technique used, the dataset(s) employed (e.g., UCR Time Series Archive [11]), the deep learning model architecture, and the reported classification performance metrics (e.g., accuracy, F1-score) was extracted. A qualitative synthesis was performed to identify trends, compare the effectiveness of different techniques across various domains and datasets, and understand the conditions under which certain augmentation strategies are more beneficial. The synthesis aimed to answer questions such as: Which augmentation techniques are most commonly used? Which techniques show consistent improvements? Are there domain-specific preferences for certain techniques?

2.5 Evaluation Metrics

The primary evaluation metric considered in most reviewed studies for time-series classification is classification accuracy. Other metrics such as F1-score, precision, recall, and area under the receiver operating characteristic curve (AUC-ROC) were also noted where reported, especially in cases of imbalanced datasets. The effectiveness of data augmentation was typically assessed by comparing the performance of a deep learning model trained with augmented data against the same model trained only on the original data.

RESULTS

This section presents the findings from our comprehensive survey, detailing the various data augmentation techniques applied to time-series classification and synthesizing the empirical evidence regarding their effectiveness.

3.1 Overview of Identified Augmentation Techniques

Data augmentation techniques for time series can be broadly categorized based on their approach:

3.1.1 Transformation-Based Methods

These methods involve directly manipulating the existing time series to create new variants. They are often simple to implement and computationally efficient.

• Jittering (Adding Noise): This involves adding random noise (e.g., Gaussian noise, white noise) to the original time series [6, 18, 51]. The noise level is typically small to preserve the original signal's characteristics. This helps the model become more robust to sensor noise or minor fluctuations in real-world data. Um et al. [6] applied jittering to wearable sensor data for Parkinson's disease monitoring, showing improved performance. Goubeaud et al. [18] specifically explored "White Noise Windows" for time series augmentation.

• Scaling (Magnitude Scaling): Multiplying the time series by a random scalar factor to change its magnitude [6, 30]. This technique helps the model learn features invariant to amplitude variations. Liu et al. [30] investigated efficient time series augmentation methods, including scaling.

• Magnitude Warping: Similar to scaling but applies a smoothly varying scaling factor over time, often using a low-frequency curve [6]. This introduces non-uniform amplitude changes.

• Time Warping (Dynamic Time Warping - DTW based): This is a powerful technique that stretches or compresses the time axis of a series [7, 45, 46, 48, 50]. DTW-based methods generate new series by warping the original series towards a randomly selected "guide" series from the same class or by applying random warping paths. Rashid and Louis [7, 32] used "Window-Warping" for IMU data augmentation in construction equipment activity identification. Kamycki et al. [45] and Warchoł and Oszust [37, 46] explored suboptimal warping for time-series classification, demonstrating its effectiveness. Iwana and Uchida [48] further developed time warping with a discriminative teacher. Yang et al. [50] also proposed a DTW-based augmentation method.

• Permutation: Dividing the time series into segments and randomly permuting their order [13]. This can be useful for capturing local patterns while introducing global variations. Le Guennec et al. [13] explored this for CNNs.

• Cropping (Window Slicing/Random Slicing): Extracting random sub-sequences of a fixed length from the original time series [13]. This helps in learning features robust to variations in the start and end points of relevant patterns.

• Resampling/Interpolation: Changing the sampling rate of the time series or interpolating between existing data points to create new ones [49]. Oh et al. [49] proposed an interpolation-based time-series data augmentation.

• Flipping/Reversing: Reversing the order of the time series [8]. While less common and highly domain-specific, it can be useful in scenarios where the temporal direction does not significantly alter the class meaning.

• Window Warping: A specific form of time warping that applies warping to defined windows within the time series [7, 37].

• Frequency Domain Augmentation: Transforming the time series into the frequency domain (e.g., using Fourier Transform) and applying operations there, then converting back. Yang et al. [39] introduced SFCC (Stratified Fourier Coefficients Combination) for data augmentation in the frequency domain. Goubeaud et al. [40] explored "Random Noise Boxes" for spectrograms, which are frequency-time representations.

3.1.2 Generative Model-Based Methods

These methods learn the underlying data distribution and generate entirely new synthetic time series.

• Generative Adversarial Networks (GANs): GANs consist of a generator and a discriminator network that are trained adversarially [15, 54]. The generator learns to produce realistic synthetic time series that can fool the discriminator, which in turn learns to distinguish real from fake data. Haradal et al. [15] used GANs for biosignal data augmentation. El-Laham and Vyetrenko [54] introduced StyleTime, a style transfer approach for synthetic time series generation.

• Variational Autoencoders (VAEs): VAEs learn a latent representation of the data and can then generate new samples by sampling from this latent space.

• Recurrent Neural Networks (RNNs) / LSTMs: These models can be trained to generate sequential data by learning the temporal dependencies in the training set [14]. Forestier et al. [14] proposed generating synthetic time series to augment sparse datasets.

3.1.3 Feature-Space Augmentation

Instead of augmenting the raw time series, these methods operate on features extracted from the data.

• Shapelets-based Augmentation: Shapelets are discriminative sub-sequences within time series. Augmentation can involve generating new time series by combining or manipulating existing shapelets [12, 52]. Li et al. [12] explored shapelets-based data augmentation for time series classification. Liu et al. [52] proposed Adaptive Shapelets Preservation for time series augmentation.

3.1.4 Hybrid Approaches

Combinations of different augmentation techniques can often yield better results than individual methods. For example, combining jittering with time warping. Aboussalah et al. [44] introduced Recursive Time Series Data Augmentation, which can be seen as a hybrid approach.

3.2 Empirical Findings Synthesis

The empirical studies reviewed consistently demonstrate that data augmentation can significantly improve the performance of deep learning models for time-series classification, especially in scenarios with limited training data.

• General Effectiveness: Many studies report improvements in accuracy and generalization ability after applying various augmentation techniques. For instance, Um et al. [6] showed that data augmentation, including jittering and scaling, improved the performance of CNNs for Parkinson's disease monitoring using wearable sensor data. Similarly, Rashid and Louis [7, 32] found "Window-Warping" beneficial for IMU data.

• Transformation-Based Methods Dominance: Simple transformation-based methods like jittering, scaling, and various forms of time warping (e.g., DTW-based, window warping) are among the most widely adopted and empirically validated techniques [6, 7, 13, 18, 30, 32, 37, 45, 46, 48, 49, 50, 51]. Their effectiveness stems from their ability to introduce variability while preserving the core temporal patterns. Forestier et al. [13] showed that data augmentation, including permutation and cropping, improved CNN performance for time series classification.

• Time Warping's Impact: Time warping, particularly DTW-based methods, is frequently highlighted as a highly effective augmentation strategy for time series [7, 45, 46, 48, 50, 55]. This is because it directly addresses the variability in the temporal alignment of patterns, a common challenge in time series data. Warchoł and Oszust [37] and Kamycki et al. [45] provided evidence for the effectiveness of warping techniques. Akyash et al. [55] introduced DTW-Merge as a novel technique.

• Generative Models for Complex Data: For more complex or highly structured time series, generative models like GANs show promise in creating realistic synthetic data [15, 54]. However, training these models can be challenging and computationally intensive, requiring careful tuning. Haradal et al. [15] demonstrated the use of GANs for biosignal data.

• Domain Specificity: The optimal augmentation strategy often depends on the specific domain and characteristics of the time series data. For example, in medical signals like ECG, specific augmentations like ECGAug [47] or those focusing on rhythm strips [36] are developed. For radio modulation classification, techniques like adding noise and time shifting have shown good results [33]. Li et al. [34] applied data augmentation for inertial sensor data in cattle behavior classification. Hosseinzadeh et al. [35] improved solar energetic particle event prediction using multivariate time series data augmentation.

• Impact on Small Datasets: Data augmentation is particularly beneficial when dealing with small datasets, where it can significantly mitigate overfitting and improve generalization [28]. Cai et al. [28] analyzed the performance of time series data augmentation for small sample communication device recognition.

• Combination of Techniques: Several studies suggest that combining multiple augmentation techniques can lead to superior performance compared to using a single method [44]. This indicates that different augmentations can capture distinct aspects of data variability.

• Challenges and Limitations:

o Over-augmentation: Excessive or inappropriate augmentation can lead to the generation of unrealistic samples, which might introduce noise and degrade model performance.

o Computational Cost: Generative model-based augmentation can be computationally expensive during the training phase.

o Hyperparameter Tuning: Many augmentation techniques involve hyperparameters (e.g., noise level, warping intensity) that need careful tuning, which can be time-consuming.

o Preserving Temporal Semantics: The most critical challenge is ensuring that augmentation preserves the underlying temporal semantics and class labels. Techniques that severely distort the temporal relationships can be detrimental.

o Curse of Dimensionality: For high-dimensional time series, applying transformations effectively can be more complex.

3.3 Empirical Performance Summary

Across the surveyed literature, various studies consistently report performance improvements ranging from a few percentage points to significant gains in accuracy when data augmentation is applied. For instance:

• Medical Signals: Studies on ECG classification [17, 36, 47, 53, 58] frequently utilize noise addition, scaling, and time warping, showing improved diagnostic accuracy. Zhao et al. [58] used CNNs with data augmentation for epileptic IEEG signal classification. Khalili and Asl [53] developed a new data augmentation technique for sleep stage classification from raw single-channel EEG.

• Activity Recognition: For human activity recognition using IMU sensors, jittering, scaling, and time warping (especially window warping) are effective [6, 7, 34, 37, 41]. Eyobu and Han [41] used data augmentation for human activity classification with LSTM networks.

• Audio Classification: In audio time series, techniques like vocal tract length perturbation (VTLP) [38] and general audio augmentation methods [56, 57, 59] have shown success. Cui et al. [57] applied data augmentation for deep neural network acoustic modeling.

• Financial Time Series: Augmentation methods for financial data often need to be carefully designed to preserve market dynamics [16]. Fons et al. [16] evaluated data augmentation for financial time series classification.

• General Time Series Benchmarks: Many studies evaluate on public datasets like the UCR Time Series Archive [11], consistently demonstrating the benefits of augmentation [4, 13, 14, 43]. Iwana and Uchida [4] provided an empirical survey of data augmentation for time series classification with neural networks, often using UCR datasets. Zanella et al. [43] introduced TS-DENSE for time series data augmentation by subclass clustering.

The empirical evidence strongly supports the notion that data augmentation is a vital component in developing robust and high-performing deep learning models for time-series classification, particularly when labeled data is scarce.

DISCUSSION

The findings from this comprehensive survey and empirical assessment underscore the critical role of data augmentation in advancing time-series classification, especially in the era of deep learning. The diverse array of techniques, ranging from simple transformations to sophisticated generative models, highlights the active research landscape in this domain.

4.1 Interpretation of Findings

The consistent empirical improvements observed across various studies and domains confirm that data augmentation is not merely a supplementary technique but often a necessary component for achieving state-of-the-art performance in time-series classification. The effectiveness of transformation-based methods, particularly time warping and noise injection, suggests that deep learning models benefit significantly from exposure to variations in temporal alignment, magnitude, and minor sensor noise. These methods are generally easy to implement and computationally inexpensive, making them practical choices for many applications.

Generative models, while more complex and resource-intensive, offer the potential to create highly realistic and diverse synthetic data, especially when the underlying data distribution is intricate. Their utility is particularly evident in scenarios where traditional transformations might not capture the full complexity of data variability or where data privacy is a concern.

The domain-specific nature of optimal augmentation strategies is a crucial insight. There is no one-size-fits-all solution; the most effective technique often depends on the characteristics of the time series (e.g., periodicity, noise level, underlying physical process) and the specific classification task. This necessitates careful consideration and often empirical evaluation of different augmentation methods for each new application.

4.2 Comparison with Other Domains

Comparing data augmentation in time series with other domains like image processing and natural language processing reveals both similarities and distinct challenges. In image augmentation [1, 2, 20], transformations like rotation, flipping, and cropping are common and semantically preserving. Similarly, in NLP [21], techniques like synonym replacement or back-translation maintain linguistic meaning. Time series, however, possesses a unique temporal structure. A simple "flip" might alter the meaning entirely (e.g., a rising trend becoming a falling one), and random "cuts" can destroy critical sequential dependencies. This inherent temporal dependency makes designing effective and semantically valid augmentation techniques more challenging, requiring methods like time warping that respect the sequential nature of the data.

4.3 Future Directions

Despite significant progress, several avenues for future research exist to further enhance data augmentation for time-series classification:

• Automated Augmentation Strategies: Developing automated or adaptive data augmentation techniques that can learn optimal augmentation policies for a given dataset and model, similar to AutoAugment in computer vision [2]. This could involve reinforcement learning or meta-learning approaches to dynamically select and parameterize augmentation operations.

• Hybrid and Multi-Modal Augmentation: Exploring more sophisticated combinations of existing techniques and developing methods that can augment multi-variate or multi-modal time series data more effectively [35].

• Theoretical Understanding of Augmentation: Gaining a deeper theoretical understanding of why certain augmentation techniques work for specific time-series characteristics. This could lead to more principled design of new augmentation methods rather than relying solely on empirical trial-and-error.

• Augmentation for Specific Challenges: Research into augmentation techniques specifically designed to address challenges like concept drift, noisy labels, or highly imbalanced time-series datasets.

• Benchmarking and Standardization: The lack of standardized benchmarks and evaluation protocols for time-series data augmentation makes direct comparison across studies difficult. Establishing common datasets and evaluation metrics would greatly facilitate research in this area. The UCR Time Series Archive [11] is a good start, but more diverse and challenging datasets with varying characteristics are needed.

• Novel Generative Models: Further research into advanced generative models tailored for time-series data, potentially incorporating domain-specific knowledge or physics-informed constraints to ensure higher fidelity and realism of synthetic samples.

• Augmentation in Edge Computing: Developing lightweight and efficient augmentation techniques that can be applied directly on edge devices, reducing the need for extensive data transfer to cloud-based training systems.

CONCLUSION

Data augmentation has proven to be an indispensable tool for improving the performance and robustness of deep learning models in time-series classification. This article has provided a comprehensive survey of the diverse techniques employed, ranging from simple transformations like jittering and time warping to advanced generative models. The empirical evidence overwhelmingly supports the efficacy of these methods in mitigating data scarcity and enhancing model generalization. While significant progress has been made, the unique characteristics of time-series data present ongoing challenges and exciting opportunities for future research. By focusing on automated strategies, hybrid approaches, and a deeper theoretical understanding, the field can continue to advance, enabling the development of even more powerful and reliable time-series classification systems for real-world applications.

REFERENCES

1. Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. Journal of big data 6, 1 (2019), 1–48.

2. A Jung et al. 2017. imgaug: Image augmentation for machine learning experiments. p. Accessed 3 (2017), 977–997.

3. G Forestier, J Weber, L Idoumghar, PA Muller, et al. 2019. Deep learning for time series classification: a review. Data Min. Knowl. Discov. 33, 4 (2019), 917–963.

4. Brian Kenji Iwana and Seiichi Uchida. 2021. An empirical survey of data augmentation for time series classification with neural networks. Plos one 16, 7 (2021), e0254841.

5. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems 25 (2012).

6. Terry T Um, Franz MJ Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. 2017. Data augmentation of wearable sensor data for parkinson's disease monitoring using convolutional neural networks. In Proceedings of the 19th ACM international conference on multimodal interaction. 216–220.

7. Khandakar M Rashid and Joseph Louis. 2019. Window-warping: a time series data augmentation of IMU data for construction equipment activity identification. In ISARC. Proceedings of the international symposium on automation and robotics in construction, Vol. 36. IAARC Publications, 651–657.

8. Edgar Talavera, Guillermo Iglesias, Ángel González-Prieto, Alberto Mozo, and Sandra Gómez-Canaval. 2022. Data augmentation techniques in time series domain: A survey and taxonomy. arXiv preprint arXiv:2206.13508 2206.13508 (2022).

9. Anthony Bagnall, Jason Lines, Aaron Bostrom, James Large, and Eamonn Keogh. 2017. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. Data mining and knowledge discovery 31 (2017), 606–660.

10. John Cristian Borges Gamboa. 2017. Deep learning for time-series analysis. arXiv preprint arXiv:1701.01887 1701.01887 (2017).

11. Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatan, and Eamonn Keogh. 2019. The UCR time series archive. IEEE/CAA Journal of Automatica Sinica 6, 6 (2019), 1293–1305.

12. Peiyu Li, Soukaïna Filali Boubrahimi, and Shah Muhammad Hamdi. 2021. Shapelets-based Data Augmentation for Time Series Classification. In 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA). IEEE, 1373–1378.

13. Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data augmentation for time series classification using convolutional neural networks. In ECML/PKDD workshop on advanced analytics and learning on temporal data.

14. Germain Forestier, François Petitjean, Hoang Anh Dau, Geoffrey I Webb, and Eamonn Keogh. 2017. Generating synthetic time series to augment sparse datasets. In 2017 IEEE international conference on data mining (ICDM). IEEE, 865–870.

15. Shota Haradal, Hideaki Hayashi, and Seiichi Uchida. 2018. Biosignal data augmentation based on generative adversarial networks. In 2018 40th annual international conference of the IEEE engineering in medicine and biology society (EMBC). IEEE, 368–371.

16. Elizabeth Fons, Paula Dawson, Xiao-jun Zeng, John Keane, and Alexandros Iosifidis. 2020. Evaluating data augmentation for financial time series classification. arXiv preprint arXiv:2010.15111 2010.15111 (2020).