

A Survey on Deep Learning based Time Series Analysis with Frequency Transformation

KUN YI, Beijing Institute of Technology, China

QI ZHANG, Tongji University, China

Longbing CAO, Macquarie University, Australia

SHOUJIN WANG, GUODONG LONG, University of Technology Sydney, Australia

LIANG HU, Tongji University, China

HUI HE, ZHENDONG NIU, Beijing Institute of Technology, China

WEI FAN, University of Central Florida, USA

HUI XIONG, Hong Kong University of Science and Technology, China

Recently, frequency transformation (FT) has been increasingly incorporated into deep learning models to significantly enhance state-of-the-art accuracy and efficiency in time series analysis. The advantages of FT, such as high efficiency and a global view, have been rapidly explored and exploited in various time series tasks and applications, demonstrating the promising potential of FT as a new deep learning paradigm for time series analysis. Despite the growing attention and the proliferation of research in this emerging field, there is currently a lack of a systematic review and in-depth analysis of deep learning-based time series models with FT. It is also unclear why FT can enhance time series analysis and what its limitations in the field are. To address these gaps, we present a comprehensive review that systematically investigates and summarizes the recent research advancements in deep learning-based time series analysis with FT. Specifically, we explore the primary approaches used in current models that incorporate FT, the types of neural networks that leverage FT, and the representative FT-equipped models in deep time series analysis. We propose a novel taxonomy to categorize the existing methods in this field, providing a structured overview of the diverse approaches employed in incorporating FT into deep learning models for time series analysis. Finally, we highlight the advantages and limitations of FT for time series modeling and identify potential future research directions that can further contribute to the community of time series analysis.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Information systems** → **Information systems application**.

Additional Key Words and Phrases: time series, neural networks, frequency transformation

ACM Reference Format:

Kun Yi, Qi Zhang, Longbing Cao, Shoujin Wang, Guodong Long, Liang Hu, Hui He, Zhendong Niu, Wei Fan, and Hui Xiong. 2023. A Survey on Deep Learning based Time Series Analysis with Frequency Transformation. *J. ACM* 37, 4, Article 111 (August 2023), 15 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Authors' addresses: Kun Yi, yikun@bit.edu.cn, Beijing Institute of Technology, Beijing, China; Qi Zhang, zhangqi_cs@tongji.edu.cn, Tongji University, Shanghai, China; Longbing Cao, longbing.cao@mq.edu.au, Macquarie University, Sydney, Australia; Shoujin Wang, Guodong Long, {shoujin.wang,guodong.long}@uts.edu.au, University of Technology Sydney, Sydney, Australia; Liang Hu, milkrain@gmail.com, Tongji University, Shanghai, China; Hui He, Zhendong Niu, {hehui617,niu}@bit.edu.cn, Beijing Institute of Technology, Beijing, China; Wei Fan, weifan@knights.ucf.edu, University of Central Florida, Orlando, USA; Hui Xiong, xionghui@ust.hk, Hong Kong University of Science and Technology, Guangzhou, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

0004-5411/2023/8-ART111 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

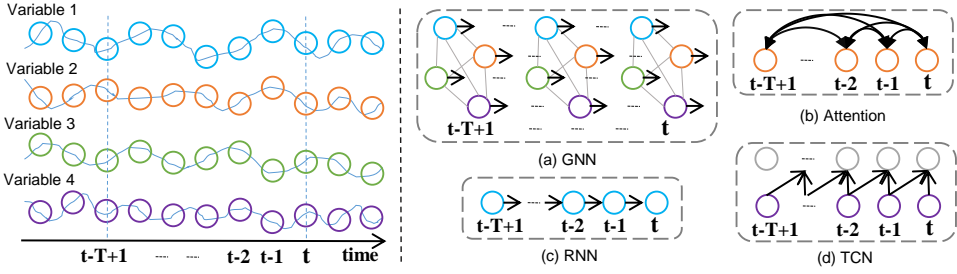


Fig. 1. Illustration of various working mechanisms applied to time series data. We take an example of four variables and T timestamps, as shown in the left portion of the figure. (a) GNN constructs a graph connecting variables for each timestamp. (b) Self-attention builds temporal connections for each variable. (c) RNN creates a recursive cycle for capturing temporal transitions. (d) TCN consists of a stack of causal convolutional layers over timestamps.

1 INTRODUCTION

Time series data is amongst the most ubiquitous data types, and has penetrated nearly every corner of our daily life [13], e.g., user-item interaction series in e-commerce and stock price series over time in finance. In recent years, time series analysis has attracted rapidly increasing attention from academia and industry, particularly in areas such as time series forecasting [5], anomaly detection [14], and classification [20]. Time series analysis has played a critical role in a wide variety of real-world applications to address significant challenges around us long-lastingly, such as traffic monitoring [3], financial analysis [21], and COVID-19 prediction [8]. However, time series analysis is extremely challenging due to the intricate inter-series correlations and intra-series dependencies.

Previous time series models based on deep learning have been devoted to modeling complex intra- and inter-series dependencies in the time domain to enhance downstream tasks. Representative sequential models such as recurrent neural networks (RNNs) [26, 30], temporal convolutional networks (TCNs) [4], and attention networks [56] are utilized to capture intra-series dependencies, while convolutional networks such as convolutional neural networks (CNNs) [31] and graph neural networks (GNNs) [9] are preferred to attend to inter-series correlations. Although achieving good results, those networks have inherent drawbacks of time-domain modeling, limiting their capabilities in capturing critical patterns for time series analysis. For example, GNNs are constructed based on variable-wise connections as illustrated in Fig. 1(a), and the sequential models (i.e., Transformer, RNN, and TCN) are based on timestamp-wise connections as shown in Fig. 1(b), (c), and (d), respectively. These modelings consider point-wise (e.g., variable/timestamp-wise) connections and fail to attend to whole or sub time series. Therefore, they are usually incapable of modeling common but complex global patterns, such as periodic patterns of seasonality, in time series [53, 60]. These inherent drawbacks inspire researchers to address the intricate inter-series correlations and intra-series dependencies of time series from a different perspective.

Recently, deep learning methods leveraging frequency transformation (FT) [42], e.g., Discrete Fourier Transform (DFT) [52], Discrete Cosine Transform (DCT) [1], and Discrete Wavelet Transform (DWT) [44], have gained a surge of interest within the machine learning community [11, 24, 57, 67]. These neural models incorporating frequency transformation have demonstrated an efficient learning paradigm in time series analysis and achieved state-of-the-art performance in terms of both efficiency and effectiveness [56, 62, 69]. This can be attributed to the distinctive advantages of FT (see Section 6.1) that the frequency spectrums generated by FT contain abundant

vital patterns, e.g., seasonal trends, and provide a global view of the characteristics of time series. In addition, FT facilitates obtaining multi-scale representations and multi-frequency components of time series for capturing informative representations and patterns. This motivates us to systematically summarize and analyze the advantages of FT to instruct researchers in this area and to deliver a comprehensive survey on the emerging area, i.e., deep learning based time series analysis with FT, thereby enlightening the time series community. While the literature includes various studies that discuss time series analysis from different perspectives [5, 7, 10, 17, 19, 43], there remains a lack of comprehensive summaries on the topic of time series analysis with FT. To the best of our knowledge, there is a notable absence of such a review covering the latest research progress of existing neural time series models based on FT. Moreover, the reasons why FT can enhance the time series analysis have not yet been summarized, and its limitations have not been thoroughly analyzed. These gaps have hindered the theoretical development and practical applications of time series analysis with FT.

In this paper, we aim to fill the aforementioned gaps by reviewing existing deep learning methods for time series analysis with FT. Specifically, our primary objective is to provide answers to four crucial perspectives: i) the strategies employed by current neural time series models in incorporating neural networks with FT; ii) the specific types of neural networks utilized in conjunction with FT; iii) the representative FT-equipped neural models commonly employed in time series applications; and iv) an exploration of the reasons behind FT to enhance neural models as well as an analysis of its limitations in the context of time series analysis. By addressing these questions, we provide valuable insights into the realm of neural time series analysis with FT. To our knowledge, this paper is the first work to comprehensively and systematically review neural time series analysis with FT and to propose a new taxonomy for this emerging area, as depicted in Fig. 2.

The subsequent sections of this paper are structured as follows: Section 2 initially presents the fundamental concepts of time series analysis and frequency transformation. Following that, Section 3 summarizes existing FT-equipped models in terms of approaches incorporating FT to enhance the accuracy or efficiency of time series analysis. In Section 4, we delve into the practical implementations of these models and examine the types of neural works utilized in conjunction with FT. Subsequently, Section 5 categorizes representative frequency-based methods based on common time series tasks, including forecasting, anomaly detection, and classification. In Section 6, we discuss the advantages and limitations of the frequency domain. Finally, we enlighten new avenues of future directions for time series analysis in Section 7.

2 PRELIMINARIES

2.1 Time Series Analysis

In this section, we provide a brief introduction to the three fundamental tasks of time series analysis before diving into neural time series analysis with Fourier Transform (FT).

2.1.1 Forecasting. Time series forecasting is the task of extrapolating time series into the future [5]. For a given time series $\mathbf{X} = [X_1, X_2, \dots, X_T] \in \mathbb{R}^{N \times T}$ with N series and T timestamps, where $X_t \in \mathbb{R}^N$ denotes the multi-variate values of N distinct series at timestamp t . We consider a time series lookback window of length- L at timestamp t , namely $\mathbf{X}_t = [X_{t-L+1}, X_{t-L+2}, \dots, X_t] \in \mathbb{R}^{N \times L}$; also, we consider a horizon window of length- τ at timestamp t as the prediction target, denoted as $\mathbf{Y}_t = [X_{t+1}, X_{t+2}, \dots, X_{t+\tau}] \in \mathbb{R}^{N \times \tau}$. Then the time series forecasting task is to use historical observations \mathbf{X}_t to predict future values $\hat{\mathbf{Y}}_t$ and the typical forecasting model f_θ parameterized by θ is to produce forecasting results by $\hat{\mathbf{Y}}_t = f_\theta(\mathbf{X}_t)$.

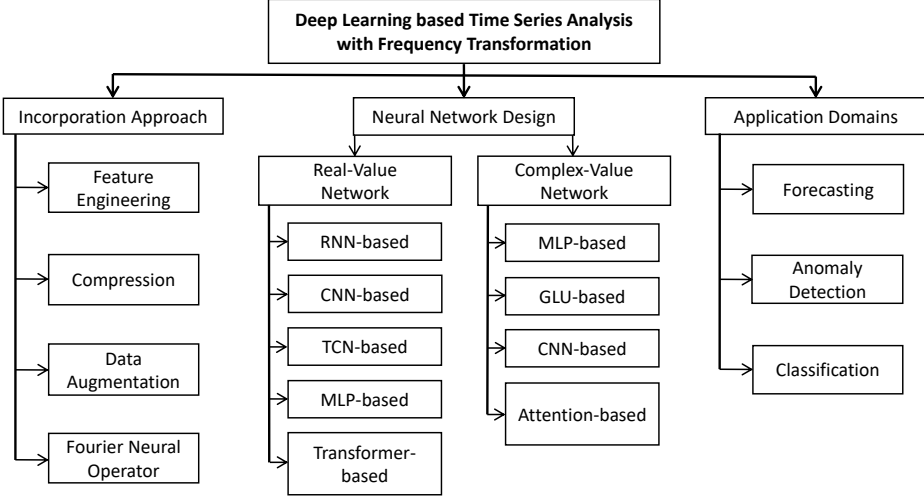


Fig. 2. A taxonomy of deep learning based time series analysis with frequency transformation.

2.1.2 Classification. Time series classification seeks to assign labels to each series of a dataset [19]. Generally, for a time series dataset $D = \{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$ where $X_i \in \mathbb{R}^T$ is a time series with T timestamps and Y_i is its corresponding one-hot vector label. For the dataset D containing K classes, Y_i is a vector of length K where each element $j \in [1, K]$ is equal to 1 if the class of X_i is j and 0 otherwise. Then the classification task is to train a classifier f_θ parameterized by θ on the dataset D to map from the space of possible inputs to a probability distribution over the class variable labels, formulated as $Y_i = f_\theta(X_i)$.

2.1.3 Anomaly Detection. Time series anomaly detection seeks to find abnormal subsequences in a series [10]. The goal is to develop algorithms or models that can effectively distinguish between normal and anomalous behavior, thereby providing early detection and alerting for unusual events or behaviors in the time series data. Given a time series $X = [x_1, x_2, \dots, x_T]$ with T timestamps where x_i represent the data point at time index i , the anomaly detection is to identify a subset of data points $X_s \subseteq X$ that represents the anomalous or abnormal instances.

2.2 Frequency Transformation

In this section, we briefly introduce commonly used frequency transformations that convert time-domain data into the frequency domain, including Discrete Fourier Transform (DFT), Discrete Cosine Transform (DCT), and Discrete Wavelet Transform (DWT). Additionally, we describe the convolution theorem, which is a fundamental property in the frequency domain.

2.2.1 Discrete Fourier Transform. Discrete Fourier Transform (DFT) [52] plays an important role in the area of digital signal processing. Given a sequence $x[n]$ with the length of N , DFT converts $x[n]$ into the frequency domain by:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-j(2\pi/N)kn}, \text{ s.t., } k = 0, 1, \dots, N-1 \quad (1)$$

where j is the imaginary unit and $X[k]$ represents the spectrum of $x[n]$ at the frequency $\omega_k = 2\pi k/N$. The spectrum $X \in \mathbb{C}^k$ consists of real parts $\text{Re} = \sum_{n=0}^{N-1} x[n] \cos(2\pi/N)kn \in \mathbb{R}^k$ and

imaginary parts $\text{Im} = -\sum_{n=0}^{N-1} x[n] \sin(2\pi/N)kn \in \mathbb{R}^k$ as:

$$\mathcal{X} = \text{Re} + j \text{Im} \quad (2)$$

The amplitude part A and phase part θ of \mathcal{X} is defined as:

$$A = \sqrt{\text{Re}^2 + \text{Im}^2} \quad (3)$$

$$\theta = \arctan\left(\frac{\text{Im}}{\text{Re}}\right) \quad (4)$$

2.2.2 Discrete Cosine Transform. Discrete Cosine Transform (DCT) [1] has emerged as the de-facto image transformation in most visual systems. The most common 1-D DCT $C(k)$ of a data sequence $x[n]$ is defined as

$$C(k) = \alpha(k) \sum_{n=0}^{N-1} x[n] \cos\left[\frac{\pi(2n+1)k}{2N}\right] \quad (5)$$

where $k = 0, 1, \dots, N-1$, and $\alpha(k)$ is defined as

$$\alpha(k) = \begin{cases} \sqrt{\frac{1}{N}}, & \text{for } k = 0 \\ \sqrt{\frac{2}{N}}, & \text{for } k \neq 0 \end{cases} \quad (6)$$

DCT only retains the real parts of DFT and is roughly equivalent to DFT that has twice its length. It often performs on real data with even symmetry or in some variants where the input or output data are shifted by half a sample.

2.2.3 Discrete Wavelet Transform. Discrete Wavelet Transform (DWT) [44] has been shown to be an appropriate tool for time-frequency analysis. It decomposes a given signal into a number of sets in which each set is a time series of coefficients describing the time evolution of the signal in the corresponding frequency band.

For a signal $x(t)$, the wavelet transform WT can be expressed as $\text{WT}(a, b) = \int_{-\infty}^{\infty} x(t) \Psi_{a,b}(t) dt = \langle x(t), \Psi_{a,b}(t) \rangle$ where Ψ is the wavelet basis function. The basis generation can be defined by $\Psi_{a,b}(t) = \frac{1}{\sqrt{a}} \Psi\left(\frac{t-b}{a}\right)$ where a and b are the scaling and translation factors respectively. DWT discretizes the scale factor a and the translation factor b as $a = a_0^m, b = ka_0^m b_0, m, k \in \mathbb{Z}$. Typically, a_0 is set to 2, and b_0 is set to 1. Accordingly, the DWT can be defined as:

$$\text{DWT}(a, b) = a_0^{-m/2} \int_{-\infty}^{\infty} x(t) \Psi(a_0^{-m}t - kb_0)(t) dt \quad (7)$$

In contrast to DFT and DCT, a wavelet transform has the ability to identify the locations containing observed frequency content, while the DFT and DCT can only extract pure frequencies from the signal. Hence, DWT can perform time-frequency analysis. In addition, DWT can obtain different resolution representations [35] by changing the scaling and translation factors. In Table 1, we compare the three frequency analysis methods, including their pros and cons of them.

2.2.4 Convolution Theorem. The convolution theorem [45] states the Fourier transform of a circular convolution of two signals equals the point-wise product of their Fourier transforms. Given a signal $x[n]$ and a filter $h[n]$, the convolution theorem can be defined as follows:

$$\mathcal{F}(x[n] * h[n]) = \mathcal{F}(x) \mathcal{F}(h) \quad (8)$$

where $x[n] * h[n] = \sum_{m=0}^{N-1} h[m]x[(n-m)_N]$, $(n-m)_N$ denotes $(n-m)$ modulo N , and $\mathcal{F}(x)$ and $\mathcal{F}(h)$ denote discrete Fourier transform of $x[n]$ and $h[n]$, respectively.

Table 1. Comparison of DFT, DCT, and DWT for time series analysis.

FT	Basis Function	Value Type	Time-Frequency	Pros	Cons
DFT	Sine+Cosine	Complex	No	Shift-invariant	Leakage effect Lack of time localization
DCT	Cosine	Real	No	Computationally efficient	No phase information Lack of time localization
DWT	Wavelet	Real	Yes	Multi-resolution analysis Localization in time and frequency	Computation complexity

According to the convolution theorem, the point-wise product of frequency spectrums of two sequences is equivalent to their circular convolution in the time domain, where the product with a larger receptive field of the whole sequences better captures the overall characteristics (e.g., periodicity) and requires less computation cost [2].

3 INCORPORATION APPROACH

In this section, we present a systematic summary and discussion of the research categorization and progress regarding incorporating the frequency transformation to enhance time series analysis.

3.1 Feature Engineering

Previous works employ frequency transformation (DFT, DCT, and DWT) as feature engineering tools to obtain frequency domain patterns. Basically, they utilize frequency transformation to capture three primary types of information: periodic patterns, multi-scale patterns, and global dependencies.

Periodicity. Compared to the time domain, the frequency domain can provide vital information for time series, such as periodic information. Prior models take advantage of frequency domain information for periodic analysis and use it as an important complement to the time domain information. [60] proposes a frequency-domain block to capture dynamic and complicated periodic patterns of time series data, and integrates deep learning networks with frequency patterns. [62] utilizes a frequency domain analysis branch to detect complex pattern anomalies, e.g., periodic anomalies. [53] learns the trend representations in the time domain, whereas the seasonal representations are learned by a Fourier layer in the frequency domain. [46] is a frequency domain-based neural network model that is built on top of the baseline model to enhance its performance. [54] utilized DFT to design a frequency attention mechanism to replace the self-attention mechanism to identify seasonal patterns. [55] leverages a novel concatenated Fourier features module to efficiently learn high-frequency patterns in time series.

Multi-Scale. One big challenge for time series analysis is that there are intricate entangled temporal dynamics among time series data. To address this challenge, some methods try to solve it in terms of the frequency domain. They disentangle temporal patterns by decomposing time series data into different frequency components. [25] separates the memory states of RNN into different frequency states such that they can explicitly learn the dependencies of both the low- and high-frequency patterns. [63] explicitly decomposes trading patterns into various frequency components and each component models a particular frequency of latent trading pattern underlying the fluctuation of stock price. Recently, wavelet-based models have shown competitive performances since wavelet transform can retain both time and frequency information and obtain multi-resolution representations. [48] proposes a wavelet-based neural network structure for building frequency-aware deep learning models for time series analysis. [50] applies maximal overlap discrete wavelet

transform to decouple time series into multiple levels of wavelet coefficients and then detect single periodicity at each level. [49] devises a novel data-dependent wavelet attention mechanism for dynamic frequency analysis of non-stationary time series analysis. [58] proposes an end-to-end graph enhanced Wavelet learning framework for long sequence forecasting which utilizes DWT to represent MTS in the wavelet domain.

Global Dependencies. Existing time domain methods construct their models based on point-wise connections (see Fig. 1), which prevent them from capturing series-level patterns, such as overall characteristics of time series. By leveraging the global view property of the frequency domain, some works utilize frequency information to attend to series-level patterns. [69] combines Fourier analysis with the Transformer which helps the Transformer better capture the global properties of time series. [62] integrates the frequency domain analysis branch with the time domain analysis branch and detects seasonality anomalies in the frequency domain. Besides, some works introduce frequency domain analysis to improve neural networks in order to address their inherent drawbacks. Vanilla convolutions in modern deep networks are known to operate locally, which causes low efficacy in connecting two distant locations in the network. To mitigate the locality limitation of convolutions, [12] converts data into the frequency domain and proposes spectral residual learning for achieving a fully global receptive field, and [11] harnesses the Fourier spectral theory and designs an operation unit to leverage frequency information for enlarging the receptive field of vanilla convolutions.

3.2 Compression

Previous works utilize frequency transformation to obtain sparse representations and remove redundant information in the frequency domain. Moreover, since noise signals usually appear as high frequencies, it is easy to filter out them in the frequency domain. For example, in [68], authors view time series forecasting from the sequence compression perspective and apply Fourier analysis to keep the part of the representation related to low-frequency Fourier components to remove the impact of noises. [41] proposes spectral pooling that performs dimensionality reduction by truncating the representation in the frequency domain because energy is heavily concentrated in the lower frequencies. [57] proposes a learning-based frequency selection method to identify the trivial frequency components while removing redundant information.

3.3 Data Augmentation

Recently, a few studies investigate data augmentation from a frequency domain perspective for time series [51]. Since the frequency domain contains some vital information for time series analysis, such as periodic patterns, existing methods incorporate frequency domain features with time domain features for data augmentations with the aim of enhancing time series representations. For example, CoST [53] incorporates a novel frequency domain contrastive loss which encourages discriminative seasonal representations and sidesteps the issue of determining the period of seasonal patterns present in the time series data. BTSF [59] fuses the temporal and spectral features to enhance the discriminativity and expressiveness of the representations. TS-TFC [33] proposes a temporal-frequency co-training model for time-series semi-supervised learning, utilizing the complementary information from two distinct views for unlabeled data learning.

More recently, different from CoST and BTSF that apply DFT after augmenting samples in the time domain, one new approach named TF-C [64] introduces frequency domain augmentations that directly perturb the frequency spectrum. It develops frequency-based contrastive augmentation to leverage rich spectral information and directly perturbs the frequency spectrum to leverage frequency-invariance for contrastive learning. Compared to performing data augmentations directly

in the frequency domain (e.g., TF-C), applying the FFT after augmenting samples in the time domain (e.g., CoST and BTSF) may lead to information loss.

3.4 Fourier Neural Operator Learning

According to the convolution theorem, differentiation is equivalent to multiplication in the Fourier domain [32]. This efficiency property makes DFT frequently used to solve differential equations.

Recently, Fourier Neural Operators (FNOs) [32], which is currently the most promising one of the neural operators [29], have been proposed as an effective framework to solve partial differential equations (PDEs). More recently, FNO has been introduced in time series forecasting. [69] proposes Fourier-enhanced blocks and Wavelet-enhanced blocks to capture important structures in time series through frequency domain mapping. [61] reformulates the graph convolution operator in the frequency domain and efficiently computes graph convolutions over a supra-graph which represents non-static correlations between any two variables at any two timestamps.

4 NEURAL NETWORK DESIGN

In this section, we delve deeper into existing related models that utilize specific types of neural networks to leverage frequency information. Considering that frequency transformation outputs can be either complex values or real values (as shown in Table 1), and each value type requires distinct handling methods, we discuss the models from the perspectives of these two value types.

4.1 Complex-Value Data

The DFT output values are complex and can be represented in two ways. One representation is through the real and imaginary parts (as shown in Equation (2)), while the other representation is through the amplitude and phase parts (as depicted in Equations (3) and (4)). While it is possible to simplify the calculation by retaining only one part, such as discarding the imaginary components [23], this approach may result in information loss.

In fact, there are mainly two approaches for performing neural networks on complex values. One approach is to treat each part of the complex value as a feature and then feed them to neural networks, respectively. Afterward, the output of corresponding networks is combined as a complex type (e.g., like Equation (2)), then the inverse DFT is executed and transmitted to the time domain. For example, StemGNN [6] conducts GLU [15] on real and imaginary parts, respectively, which concatenates them as a complex value and applies IDFT. ATFN [60] utilizes two linear layers to process the amplitude part and phase part, respectively, and then combine them as a whole. The other one is to conduct complex multiplication in the frequency domain directly. For example, FEDformer [69] randomly samples a few frequencies and conducts complex multiplication with a parameterized kernel incorporated with attention architecture.

4.2 Real-Value Data

The output value type of DCT and DWT is real, hence commonly used network structures can be directly applied to them, such as RNN and CNN. Besides, although the output value type of DFT is complex, some work discards one part, such as phase part [63], and thus their network design also belongs to a real value network. However, except for capturing frequency patterns, in contrast to other network designs, one main purpose of network design for frequency-based models is the frequency component selection to decide which component is discriminative or critical.

For example, [57] converts the input to the frequency domain by DCT and groups the same frequency into one channel, and then proposes a learning-based dynamic channel selection method to identify the trivial frequency components. [38] proposes to generalize global average pooling to more frequency components of DCT and designs three kinds of frequency components selection

Table 2. Summary of representative FT-equipped neural models in time series analysis.

Models	Frequency Transformation	Incorporation Approach	Value Type	Neural Network	Application Domains	Leveraged Advantages
SFM [63]	DFT	Feature engineering	Real-value	RNN	Forecasting	Decomposition
StemGNN [6]	DFT	Feature engineering	Complex-value	GLU	Forecasting	Decomposition
Autoformer [56]	DFT	Feature engineering	Complex-value	Attention Network	Forecasting	Global view Efficiency
AFTN [60]	DFT	Feature engineering	Complex-value	MLP	Forecasting	Decomposition
DEPTS [18]	DCT	Feature engineering	Real-value	MLP	Forecasting	Decomposition
FEDformer [69]	DFT	Feature engineering Operator learning	Complex-value	Attention Network	Forecasting	Global view Efficiency
CoST [53]	DFT	Data augmentation	Complex-value	MLP	Forecasting	Decomposition
FiLM [68]	DFT	Compression	Complex-value	MLP	Forecasting	Sparse Representation
EV-FGN [61]	DFT	Operator learning	Complex-value	MLP	Forecasting	Efficiency
FreDo [46]	DFT	Feature engineering	Complex-value	MLP	Forecasting	Decomposition
WAVEFORM [58]	DWT	Feature engineering	Real-value	GCN	Forecasting	Decomposition
SR-CNN [39]	DFT	Feature engineering	Real-value	CNN	Anomaly Detection	Decomposition
RobustTAD [22]	DFT	Data augmentation	Complex-value	CNN	Anomaly Detection	Decomposition
TFAD [62]	DWT	Feature engineering	Real-value	TCN	Anomaly Detection	Decomposition
RCF [48]	DWT	Feature engineering	Real-value	CNN	Classification	Decomposition
WD [28]	DWT	Feature engineering	Real-value	CNN	Classification	Decomposition
BTSF [59]	DFT	Data augmentation	Real-value	CNN	Classification Forecasting	Decomposition
TF-C [64]	DFT	Data augmentation	Real-value	Transformer	Classification	Decomposition

criteria. RobustPeriod [50] applies DWT to decouple time series into multiple levels of wavelet coefficients and then proposes a method to robustly calculate unbiased wavelet variance at each level and rank periodic possibilities.

5 APPLICATIONS

In this section, we review the representative FT-equipped neural time series models. We categorize them into three main applications, including forecasting, anomaly detection, and classification. In Table 2, we further compare them from six dimensions.

5.1 Time Series Forecasting

Time series forecasting is essential in various domains, such as decision making and financial analysis. Recently, some methods leverage frequency information to improve the accuracy or efficiency of time series forecasting. SFM [63] decomposes the hidden states of memory cells into multiple frequency components and models multi-frequency trading patterns. StemGNN [6] learns spectral representations which are easier to recognize after DFT. Autoformer [56] leverages FFT to calculate auto-correlation efficiently. DEPTS [18] conducts DCT to extract periodic features and then applies multi-layer perceptrons on these features for periodicity dependencies in time series. FEDformer [69] captures the global view of time series in the frequency domain. CoST [53] learns

the seasonal representations in the frequency domain. FiLM [68] utilizes Fourier analysis to keep low-frequency Fourier components.

5.2 Time Series Anomaly Detection

In recent years, frequency-based models have been introduced in anomaly detection. SR [39] extracts the spectral residual in the frequency domain for detecting the anomaly. RobustTAD [22] explores the data augmentation methods in the frequency domain to further increase labeled data. PFT [37] proposes a partial Fourier transform for anomaly detection with an order of magnitude of speedup without sacrificing accuracy. TFAD [62] takes advantage of frequency domain analysis for seasonality anomaly.

5.3 Time Series Classification

Time series classification is an important and challenging problem in time series analysis. Recently, a few models have considered frequency domain information to perform this task. RCF [48] extracts distinguishing features from the DWT decomposed results. WD [28] uses wavelet functions with adjustable scale parameters to learn the spectral decomposition directly from the signal. BTSF [59] fuses time and spectral information to enhance the discriminativity and expressiveness of the representations. TF-C [64] develops frequency-based contrastive augmentation to leverage rich spectral information and explore time-frequency consistency in time series.

6 SUMMARY OF FREQUENCY TRANSFORMATION

In this section, to investigate why FT can enhance the neural models and what are its limitations for time series analysis, we summarize the advantages and limitations of frequency transformation.

6.1 Advantages

Decomposition. Frequency transformation can decompose the original time series into different frequency components that embody vital information of time series, such as periodic patterns of seasonality. In particular, DWT can decompose a time series into a group of sub-series with frequencies ranked from high to low and obtains multi-scale representations. By decomposing time series in the time domain into different components in the frequency domain, it is naturally helpful to figure out and obtain beneficial information for time series analysis.

Global View. According to Equations (1), (5), and (7), a frequency spectrum is calculated through the summation of all signals over time. Accordingly, each spectrum element in the frequency domain attends to all timestamps in the time domain, illustrating that a spectrum has a global view of the whole sequence of time series. In addition, according to the convolution theorem (see Equation (8)), the point-wise product of frequency spectrums also captures the global characteristics of the whole sequence, inspiring to parameterize global learnable filters in the frequency domain.

Sparse Representation. Frequency transformation enables the provision of sparse representations for sequences. Taking DFT as an example, a substantial number of coefficients are close to zero, indicating that we can employ a reduced number of coefficients to represent the entire sequence. In other words, the corresponding representations in the frequency domain have a property of *energy compaction*. For example, the important features of signals captured by a subset of DWT coefficients are typically much smaller than the original. Specifically, using DWT, it ends up with the same number of coefficients as the original signal where many of the coefficients may be close to zero. As a result, we can effectively represent the original signal using only a small number of non-zero coefficients.

Efficiency. As mentioned earlier, frequency transformation often leads to sparse representations, where a substantial number of coefficients are close to zero. Exploiting this sparsity allows for efficient computations by discarding or compressing the negligible coefficients, resulting in reduced memory requirements and faster processing. Moreover, according to the convolution theorem, convolution in the time domain corresponds to Hadamard's point-wise product in the frequency domain, which allows for convolution to be calculated more efficiently in the frequency domain. Therefore, considering the equivalence of the convolution theorem, convolution calculated in the frequency domain involves significantly fewer computational operations.

6.2 Limitations

Loss of temporal information. Frequency transformation techniques, including DFT and DCT, primarily emphasize capturing the frequency characteristics of a time series. While these techniques offer valuable insights into the frequency domain, they may overlook or inadequately represent temporal information. Certain temporal patterns or dynamics inherent in the time series may not be fully captured in the frequency domain, thereby limiting the comprehensive analysis and understanding of the temporal aspects [23].

Dependence on pre-defined parameters. Frequency transformation techniques often require setting parameters, such as window size, sampling rate, or frequency bands. Selecting appropriate parameter values can be challenging, and suboptimal choices may lead to inaccurate frequency representations or missed important frequency components [28, 36]. Accordingly, parameter tuning and optimization are necessary to ensure the effectiveness of frequency transformation in time series analysis.

7 DISCUSSION FOR FUTURE OPPORTUNITIES

In this section, we explore the prospects for future research in neural time series analysis with frequency transformation. We begin by outlining the current limitations of frequency transformation and propose innovative directions to overcome these challenges. Subsequently, we delve into open research issues and emerging trends in the field of time series analysis that can be addressed through the utilization of frequency transformations.

7.1 From the Perspective of Frequency Transformation

7.1.1 Leveraging New Orthogonal Transform Technology. Recent studies have shown the efficiency and effectiveness of orthogonal transform which serves as a plug-in operation in neural networks, including frequency analysis and polynomial family. Some new orthogonal transform technologies have been introduced in neural networks and achieved good results. For example, FiLM [68] exploits the Legendre projection, which is one type of orthogonal polynomials, to update the representation of time series. [37] proposes Partial Fourier Transform (PFT) to reduce complexity from $O(N \log N)$ to $O(N + M \log M)$ where $M \ll N$. The Fractional Fourier transform (FrFT) has been proven to be desirable for noise removal and can enhance the discrimination between anomalies and background [47]. In [65], authors utilize FrFT to enhance efficient feature fusion and comprehensive feature extraction. [66] leverages FrFT to enable flexible extraction of global contexts and sequential spectral information. In the future, it would be a promising direction to incorporate more new orthogonal transform technologies for deep learning in time series analysis, such as orthogonal polynomials, DCT, and FrFT.

7.1.2 Integrating Frequency Transformation with Deep Learning. The basis functions used in frequency transformation, such as sine, cosine, and wavelet functions, are fixed across different

domains. As a result, the frequency features extracted through these basis functions are domain-invariant. In other words, the features are insensitive to unexpected noise or to changing conditions.

To mitigate the limitation, few previous works combine frequency transformation with the learning ability of neural networks. mWDN [48] proposes a wavelet-based neural network structure, in which all parameters can be fine-tuned to fit training data of different learning tasks. [28] proposes a method to efficiently optimize the parameters of the spectral decomposition based on the wavelet transform in a neural network framework. [36] mimics the fast DWT cascade architecture utilizing the deep learning framework. These methods have shown promising performances, and in the future, the combination of frequency transformation with deep learning deserves further investigation.

7.1.3 Jointly Learning in the Time and Frequency Domain. The frequency domain only uses periodic components, and thus cannot accurately model the non-periodic aspects of a signal, such as a linear trend [23]. Moreover, according to the uncertainty principle [62], designing a model with a single structure that can capture the time and frequency patterns simultaneously is difficult.

As a result, in the future, an interesting direction is to take advantage of corresponding characteristics of learning in the time and frequency domain to improve the accuracy and efficiency of time series analysis. Few works have tried to learn representations in the time and frequency domain, respectively. For example, CoST [53] learns the trend representations in the time domain and the seasonal representations in the frequency domain. However, it only performs data augmentations in the time domain and learns time and frequency representations separately. More time-frequency representation learning methods are required in the future.

7.2 From the Perspective of Time Series Analysis

7.2.1 Applying Frequency Transformation to Enhance Time Series Applications. Applying frequency transformation techniques to a wider range of time series applications has the potential to unlock valuable insights and enhance decision-making in various domains. No matter in detecting anomalies in physiological signals, uncovering market cycles in financial data, or identifying patterns in environmental parameters, frequency transformation enables a deeper understanding of complex temporal patterns and trends. By harnessing the power of frequency analysis, researchers and practitioners can uncover hidden relationships, improve forecasting accuracy, optimize resource management, and advance knowledge in diverse fields, ultimately driving innovation and enabling data-driven decision-making in a wide range of time series applications.

7.2.2 Scalability. Scalability [27] is a key consideration in time series analysis. When coupled with frequency transformation techniques, it offers the potential for efficient and scalable analysis of large-scale time series data. Frequency transformation allows for the extraction of frequency components, reducing the dimensionality of the data and enabling more efficient processing. This reduction in dimensionality can significantly improve the scalability of time series analysis algorithms, as it reduces computational complexity and memory requirements. Scalable time series analysis with frequency transformation can pave the way for analyzing and extracting insights from big data time series applications in domains such as the Internet of Things (IoT), financial markets, or sensor networks.

7.2.3 Interpretability and Explainability. Interpretability and explainability [34, 40] are crucial aspects of time series analysis especially in practical applications. Intuitively, frequency transformation capably transforms time series into a more intuitive and interpretable representation in the frequency domain, offering valuable insights into the underlying patterns and behaviors. The frequency components obtained from the transformation can be analyzed to understand the dominant frequencies, periodicities, or significant events embodied in the data. This not only enhances

the interpretability of the analysis but also enables the explanation of observed phenomena and anomalies in terms of frequency patterns. The interpretability and explainability of time series analysis with frequency transformation offer valuable advantages, enabling analysts and domain experts to attain deeper insights and establish trust in the analysis outcomes.

7.2.4 Privacy-preserving. Leveraging frequency transformation offers a powerful approach to data privacy-preserving [16] in time series analysis. By applying frequency transformation, time series data can be transformed into frequency domain representations without revealing the underlying raw data. This transformation allows for the extraction of frequency components and patterns while maintaining the confidentiality of the original information. Privacy-preserving with frequency transformation techniques can ensure individual privacy and data confidentiality, and enable collaborative analysis, data sharing, and research collaborations while mitigating privacy risks. This approach is particularly valuable in domains where data sensitivity is critical, such as healthcare, finance, or personal monitoring, allowing for the utilization of frequency analysis while protecting the privacy of individuals or organizations involved.

8 CONCLUSION

In this paper, we provide a comprehensive survey on deep learning based time series analysis with frequency transformation. We organize the reviewed methods from the perspectives of incorporation approaches, neural network design, and application domains, and we summarize the advantages and limitations of frequency transformation for time series analysis. To the best of our knowledge, this paper is the first work to comprehensively and systematically review neural time series analysis with frequency transformation, which would greatly benefit the time series community. Additionally, we offer a curated collection of sources, accessible at https://github.com/BIT-Yi/time_series_frequency, to further assist the research community.

REFERENCES

- [1] Nasir Ahmed, T. Natarajan, and Kamisetty R Rao. 1974. Discrete cosine transform. *IEEE transactions on Computers* 100, 1 (1974), 90–93.
- [2] Ahmed M. Alaa, Alex James Chan, and Mihaela van der Schaar. 2021. Generative Time-series Modeling with Fourier Flows. In *ICLR*. OpenReview.net.
- [3] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2020. Adaptive Graph Convolutional Recurrent Network for Traffic Forecasting. In *NeurIPS*.
- [4] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *CoRR* abs/1803.01271 (2018).
- [5] Konstantinos Benidis, Syama Sundar Rangapuram, Valentin Flunkert, Yuyang Wang, Danielle Maddix, Caner Turkmen, Jan Gasthaus, Michael Bohlke-Schneider, David Salinas, Lorenzo Stella, François-Xavier Aubet, Laurent Callot, and Tim Januschowski. 2022. Deep Learning for Time Series Forecasting: Tutorial and Literature Survey. *Comput. Surveys* 55 (2022).
- [6] Defu Cao, Yujing Wang, Juanyong Duan, Ce Zhang, Xia Zhu, Congrui Huang, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, and Qi Zhang. 2020. Spectral Temporal Graph Neural Network for Multivariate Time-series Forecasting. In *NeurIPS*.
- [7] Irene Y. Chen, Rahul G. Krishnan, and David A. Sontag. 2022. Clustering Interval-Censored Time-Series for Disease Phenotyping. In *AAAI*. AAAI Press, 6211–6221.
- [8] Yuzhou Chen, Ignacio Segovia-Dominguez, Baris Coskunuzer, and Yulia Gel. 2022. TAMP-S2GCNets: Coupling Time-Aware Multipersistence Knowledge Representation with Spatio-Supra Graph Convolutional Networks for Time-Series Forecasting. In *ICLR*.
- [9] Zekai Chen, Dingshuo Chen, Xiao Zhang, Zixuan Yuan, and Xiuzhen Cheng. 2022. Learning Graph Structures With Transformer for Multivariate Time-Series Anomaly Detection in IoT. *IEEE Internet Things J.* 9, 12 (2022), 9179–9189.
- [10] Zhipeng Chen, Zhang Peng, Xueqiang Zou, and Haoqi Sun. 2021. Deep Learning Based Anomaly Detection for Multi-dimensional Time Series: A Survey. In *CNCERT (Communications in Computer and Information Science, Vol. 1506)*. Springer, 71–92.

- [11] Lu Chi, Borui Jiang, and Yadong Mu. 2020. Fast Fourier Convolution. In *NeurIPS*.
- [12] Lu Chi, Guiyu Tian, Yadong Mu, Lingxi Xie, and Qi Tian. 2019. Fast Non-Local Neural Networks with Spectral Residual Learning. In *ACM Multimedia*. ACM, 2142–2151.
- [13] Fatoumata Dama and Christine Sinoquet. 2021. Analysis and modeling to forecast in time series: a systematic review. *CoRR* abs/2104.00164 (2021).
- [14] Zahra Zamanzadeh Darban, Geoffrey I. Webb, Shirui Pan, Charu C. Aggarwal, and Mahsa Salehi. 2022. Deep Learning for Time Series Anomaly Detection: A Survey. *CoRR* abs/2211.05244 (2022).
- [15] Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language Modeling with Gated Convolutional Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 70)*. PMLR, 933–941.
- [16] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. 2016. Calibrating Noise to Sensitivity in Private Data Analysis. *J. Priv. Confidentiality* 7, 3 (2016), 17–51.
- [17] Amin Fakhrazari and Hamid Vakilzadian. 2017. A survey on time series data mining. In *EIT. IEEE*, 476–481.
- [18] Wei Fan, Shun Zheng, Xiaohan Yi, Wei Cao, Yanjie Fu, Jiang Bian, and Tie-Yan Liu. 2022. DEPTS: Deep Expansion Learning for Periodic Time Series Forecasting. In *ICLR*. OpenReview.net.
- [19] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: a review. *Data Min. Knowl. Discov.* 33, 4 (2019), 917–963.
- [20] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery* 33 (2019).
- [21] Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. Temporal Relational Ranking for Stock Prediction. *ACM Trans. Inf. Syst.* 37, 2 (2019), 27:1–27:30.
- [22] Jingkun Gao, Xiaomin Song, Qingsong Wen, Pichao Wang, Liang Sun, and Huan Xu. 2020. RobustTAD: Robust Time Series Anomaly Detection via Decomposition and Convolutional Neural Networks. *CoRR* abs/2002.09545 (2020).
- [23] Luke B. Godfrey and Michael S. Gashler. 2018. Neural Decomposition of Time-Series Data for Effective Generalization. *IEEE Trans. Neural Networks Learn. Syst.* 29, 7 (2018), 2973–2985.
- [24] John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. 2022. Adaptive Fourier Neural Operators: Efficient Token Mixers for Transformers. In *ICLR*.
- [25] Hao Hu and Guo-Jun Qi. 2017. State-Frequency Memory Recurrent Neural Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 70)*. PMLR, 1568–1577.
- [26] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Söderström. 2018. Detecting Spacecraft Anomalies Using LSTMs and Nonparametric Dynamic Thresholding. In *KDD*. ACM, 387–395.
- [27] Eamonn J. Keogh and Shrutli Kasetty. 2003. On the Need for Time Series Data Mining Benchmarks: A Survey and Empirical Demonstration. *Data Min. Knowl. Discov.* 7, 4 (2003), 349–371.
- [28] Haidar Khan and Bülent Yener. 2018. Learning filter widths of spectral decompositions with wavelets. In *NeurIPS*. 4606–4617.
- [29] Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. 2021. Neural Operator: Learning Maps Between Function Spaces. *CoRR* abs/2108.08481 (2021).
- [30] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling Long- and Short-Term Temporal Patterns with Deep Neural Networks. In *SIGIR*. 95–104.
- [31] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *ICLR (Poster)*.
- [32] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. 2021. Fourier Neural Operator for Parametric Partial Differential Equations. In *ICLR*.
- [33] Zhen Liu, Qianli Ma, Peitian Ma, and Linghao Wang. 2023. Temporal-Frequency Co-training for Time Series Semi-supervised Learning. In *AAAI*. AAAI Press, 8923–8931.
- [34] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *NIPS*. 4765–4774.
- [35] Stéphane Mallat. 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. *IEEE Trans. Pattern Anal. Mach. Intell.* 11, 7 (1989), 674–693.
- [36] Gabriel Michau, Gaetan Frusque, and Olga Fink. 2022. Fully learnable deep wavelet transform for unsupervised monitoring of high-frequency time series. *Proceedings of the National Academy of Sciences* 119, 8 (2022).
- [37] Yong-chan Park, Jun-Gi Jang, and U Kang. 2021. Fast and Accurate Partial Fourier Transform for Time Series Data. In *KDD*. ACM, 1309–1318.
- [38] Zequn Qin, Pengyi Zhang, Fei Wu, and Xi Li. 2021. FcaNet: Frequency Channel Attention Networks. In *ICCV*. IEEE, 763–772.
- [39] Hansheng Ren, Bixiong Xu, Yujing Wang, Chao Yi, Congrui Huang, Xiaoyu Kou, Tony Xing, Mao Yang, Jie Tong, and Qi Zhang. 2019. Time-Series Anomaly Detection Service at Microsoft. In *KDD*. ACM, 3009–3017.
- [40] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *KDD*. ACM, 1135–1144.

- [41] Oren Rippel, Jasper Snoek, and Ryan P. Adams. 2015. Spectral Representations for Convolutional Neural Networks. In *NIPS*. 2449–2457.
- [42] Richard A Roberts and Clifford T Mullis. 1987. *Digital signal processing*. Addison-Wesley Longman Publishing Co., Inc.
- [43] Patrick Schäfer, Arik Ermshaus, and Ulf Leser. 2021. ClaSP - Time Series Segmentation. In *CIKM*. ACM, 1578–1587.
- [44] Mark J Shensa et al. 1992. The discrete wavelet transform: wedding the a trous and Mallat algorithms. *IEEE Transactions on signal processing* 40, 10 (1992), 2464–2482.
- [45] S. S. Soliman and MD Srinath. 1990. Continuous and discrete signals and systems. *Prentice Hall*, (1990).
- [46] Fan-Keng Sun and Duane S. Boning. 2022. FreDo: Frequency Domain-based Long-Term Time Series Forecasting. *CoRR* abs/2205.12301 (2022).
- [47] Ran Tao, Xudong Zhao, Wei Li, Heng-Chao Li, and Qian Du. 2019. Hyperspectral Anomaly Detection by Fractional Fourier Entropy. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.* 12, 12 (2019), 4920–4929.
- [48] Jingyuan Wang, Ze Wang, Jianfeng Li, and Junjie Wu. 2018. Multilevel Wavelet Decomposition Network for Interpretable Time Series Analysis. In *KDD*. ACM, 2437–2446.
- [49] Jingyuan Wang, Chen Yang, Xiaohan Jiang, and Junjie Wu. 2023. WHEN: A Wavelet-DTW Hybrid Attention Network for Heterogeneous Time Series Analysis. In *KDD*. ACM, 2361–2373.
- [50] Qingsong Wen, Kai He, Liang Sun, Yingying Zhang, Min Ke, and Huan Xu. 2021. RobustPeriod: Robust Time-Frequency Mining for Multiple Periodicity Detection. In *SIGMOD Conference*. ACM, 2328–2337.
- [51] Qingsong Wen, Liang Sun, Fan Yang, Xiaomin Song, Jingkun Gao, Xue Wang, and Huan Xu. 2021. Time Series Data Augmentation for Deep Learning: A Survey. In *IJCAI*. ijcai.org, 4653–4660.
- [52] Shmuel Winograd. 1976. On computing the discrete Fourier transform. *Proceedings of the National Academy of Sciences* 73, 4 (1976), 1005–1006.
- [53] Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven C. H. Hoi. 2022. CoST: Contrastive Learning of Disentangled Seasonal-Trend Representations for Time Series Forecasting. In *ICLR*. OpenReview.net.
- [54] Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven C. H. Hoi. 2022. ETSformer: Exponential Smoothing Transformers for Time-series Forecasting. *CoRR* abs/2202.01381 (2022).
- [55] Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven C. H. Hoi. 2023. Learning Deep Time-index Models for Time Series Forecasting. In *ICML (Proceedings of Machine Learning Research, Vol. 202)*. PMLR, 37217–37237.
- [56] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. In *NeurIPS*. 22419–22430.
- [57] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. 2020. Learning in the Frequency Domain. In *CVPR*. 1737–1746.
- [58] Fuhao Yang, Xin Li, Min Wang, Hongyu Zang, Wei Pang, and Mingzhong Wang. 2023. WaveForM: Graph Enhanced Wavelet Learning for Long Sequence Forecasting of Multivariate Time Series. In *AAAI*. AAAI Press, 10754–10761.
- [59] Ling Yang and Shenda Hong. 2022. Unsupervised Time-Series Representation Learning with Iterative Bilinear Temporal-Spectral Fusion. In *ICML (Proceedings of Machine Learning Research, Vol. 162)*. PMLR, 25038–25054.
- [60] Zhangjing Yang, Weiwei Yan, Xiaolin Huang, and Lin Mei. 2022. Adaptive Temporal-Frequency Network for Time-Series Forecasting. *IEEE Trans. Knowl. Data Eng.* 34, 4 (2022), 1576–1587.
- [61] Kun Yi, Qi Zhang, Liang Hu, Hui He, Ning An, Longbing Cao, and Zhendong Niu. 2022. Edge-Varying Fourier Graph Networks for Multivariate Time Series Forecasting. *CoRR* abs/2210.03093 (2022).
- [62] Chaoli Zhang, Tian Zhou, Qingsong Wen, and Liang Sun. 2022. TFAD: A Decomposition Time Series Anomaly Detection Architecture with Time-Frequency Analysis. In *CIKM*. ACM, 2497–2507.
- [63] Liheng Zhang, Charu C. Aggarwal, and Guo-Jun Qi. 2017. Stock Price Prediction via Discovering Multi-Frequency Trading Patterns. In *KDD*. ACM, 2141–2149.
- [64] Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. 2022. Self-Supervised Contrastive Pre-Training For Time Series via Time-Frequency Consistency. In *NeurIPS*.
- [65] Xudong Zhao, Ran Tao, Wei Li, Wilfried Philips, and Wenzhi Liao. 2022. Fractional Gabor Convolutional Network for Multisource Remote Sensing Data Classification. *IEEE Trans. Geosci. Remote. Sens.* 60 (2022), 1–18.
- [66] Xudong Zhao, Mengmeng Zhang, Ran Tao, Wei Li, Wenzhi Liao, Lianfang Tian, and Wilfried Philips. 2022. Fractional Fourier Image Transformer for Multimodal Remote Sensing Data Classification. *IEEE Trans. Neural Networks Learn. Syst.* (2022), 1–13.
- [67] Kun Zhou, Hui Yu, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Filter-enhanced MLP is All You Need for Sequential Recommendation. In *WWW*. 2388–2399.
- [68] Tian Zhou, Ziqing Ma, Xue Wang, Qingsong Wen, Liang Sun, Tao Yao, Wotao Yin, and Rong Jin. 2022. FiLM: Frequency improved Legendre Memory Model for Long-term Time Series Forecasting. In *NeurIPS*.
- [69] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *ICML*.