From Pixels to Predictions: Spectrogram and Vision Transformer for Better Time Series Forecasting

Zhen Zeng J. P. Morgan AI Research New York, NY, USA zhen.zeng@jpmchase.com Rachneet Kaur* J. P. Morgan AI Research New York, NY, USA rachneet.kaur@jpmorgan.com

Tucker Balch J. P. Morgan AI Research New York, NY, USA tucker.balch@jpmchase.com Suchetha Siddagangappa* J. P. Morgan AI Research New York, NY, USA suchetha.siddagangappa@jpmchase.com

ABSTRACT

Time series forecasting plays a crucial role in decision-making across various domains, but it presents significant challenges. Recent studies have explored image-driven approaches using computer vision models to address these challenges, often employing lineplots as the visual representation of time series data. In this paper, we propose a novel approach that uses time-frequency spectrograms as the visual representation of time series data. We introduce the use of a vision transformer for multimodal learning, showcasing the advantages of our approach across diverse datasets from different domains. To evaluate its effectiveness, we compare our method against statistical baselines (EMA and ARIMA), a state-of-the-art deep learning-based approach (DeepAR), other visual representations of time series data (lineplot images), and an ablation study on using only the time series as input. Our experiments demonstrate the benefits of utilizing spectrograms as a visual representation for time series data, along with the advantages of employing a vision transformer for simultaneous learning in both the time and frequency domains.

KEYWORDS

time series forecasting, image representations, spectrogram, neural networks, attention, transformer

ACM Reference Format:

Zhen Zeng, Rachneet Kaur, Suchetha Siddagangappa, Tucker Balch, and Manuela Veloso. 2023. From Pixels to Predictions: Spectrogram and Vision Transformer for Better Time Series Forecasting. In 4th ACM International Conference on AI in Finance (ICAIF '23), November 27–29, 2023, Brooklyn, NY, USA. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3604237.3626905

ICAIF '23, November 27-29, 2023, Brooklyn, NY, USA

Manuela Veloso J. P. Morgan AI Research New York, NY, USA manuela.veloso@jpmchase.com



Figure 1: Visual representation of time series in the form of a time-frequency spectrogram augmented with intensities of time series at the top

1 INTRODUCTION

Time series forecasting poses significant challenges due to the inherent noise in the data. Traditional statistical methods often utilize linear regression, exponential smoothing [12, 18, 36], and autoregression models [21]. In recent years, deep learning has witnessed substantial progress, focusing on ensemble models and sequenceto-sequence modeling, such as recurrent neural networks (RNNs), long short-term memory (LSTM) [17], and more recently, transformers [34]. These advanced deep learning techniques offer enhanced capabilities in capturing complex non-linear dependencies within time series data.

In recent years, a new perspective has been introduced for forecasting time series where numeric data is converted to images to leverage successful computer vision algorithms for forecasting. These studies are motivated by the fact that traders' decision making is augmented by the visual representation of financial time series images such as charts, graphs, and plots. In addition, this approach can help to capture additional patterns and dependencies that are beneficial for understanding and forecasting time series data. However, these approaches typically use lineplots as the visual representation of time series, which misses some crucial information required for understanding time series. For example, frequency information is not apparent when directly looking at the time series in its raw form.

One could use frequency spectrum of a time series as the visual representation. But this merely displays the various frequency components present in the data. High-frequency components are

^{*}Both authors contributed equally to this research.

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 the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or 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} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0240-2/23/11...\$15.00 https://doi.org/10.1145/3604237.3626905

often associated with noise, while low-frequency components correspond to signals. However, it does not provide insights into the temporal dependencies between occurrences of different frequency components. The time-frequency spectrogram on the other hand visually represents how the frequency spectrum changes over time, enabling the learning of temporal dependencies between different frequency patterns.

The main focus of this research is to use time-frequency spectrogram as visual representation of time series data and leveraging vision transformers to achieve simultaneous learning in both the time and frequency domains, enabling accurate forecasting. The paper presents the following key contributions:

- We demonstrate the superiority of spectrograms as a visual representation of time series data, leading to enhanced predictions across various domains and diverse datasets.
- We leverage transformer's ability to learn cross-modality information across time and frequency domain to enhance the forecasting process.

Through comprehensive experiments conducted on diverse datasets from different domains, including synthetic, temperature, and financial time series, the paper demonstrates the advantages and effectiveness of the proposed approach in time series forecasting.

2 RELATED WORKS

2.1 Time series forecasting

Time series forecasting is a crucial task involving the statistical analysis of historical data to predict future values. In the literature, traditional forecasting techniques have relied on statistical tools like exponential smoothing [12, 18, 36] and autoregressive integrated moving average (ARIMA) [21], primarily applied to numerical time series data for one-step-ahead predictions. Moreover, several traditional multi-horizon forecasting methods [22, 33] have been developed to generate simultaneous predictions for multiple future time steps.

Machine learning (ML) approaches have shown promise in improving forecasting performance by effectively addressing highdimensional and non-linear feature interactions in a model-free manner. These ML methods encompass tree-based algorithms, ensemble methods, neural networks, autoregression, and recurrent neural networks (RNNs) [16]. In recent works, deep learning (DL) methods have gained traction in time series forecasting, demonstrating impressive results on numerical time series data [4, 13, 26, 27, 32, 41]. DL techniques automate the feature extraction process, eliminating the need for domain expertise and offering promising outcomes in forecasting tasks.

In the domain of natural language processing (NLP), RNNs, long short-term memory (LSTM) [17], and gated recurrent units (GRUs) [6] have been widely used for tasks like machine translation and language modeling. Interestingly, these models have been adapted for time series forecasting, where they have outperformed traditional statistical methods. Furthermore, the DeepAR algorithm, utilizing the RNN backbone, has emerged as a potent approach in time series forecasting. DeepAR has been more recently proposed and has shown superior performance compared to traditional forecasting techniques, particularly for datasets with multiple related time series [28]. This highlights the potential of RNN-based methods in handling complex time series data and effectively capturing temporal dependencies.

Transformers [34], on the other hand, have revolutionized NLP applications and have also shown promise in time series forecasting. Unlike traditional RNN and LSTM networks, transformers leverage multi-headed self-attention to process all inputs simultaneously. This parallel computation capability not only reduces training time but also enables transformers to handle longer sequences without suffering from long-term memory dependency issues, often encountered in RNN-based models. These characteristics make transformers well-suited for handling time series forecasting tasks effectively [37, 39, 42]. By effectively capturing complex temporal dependencies within time series data, transformers have demonstrated their ability to outperform traditional methods and are gaining popularity as a powerful tool in the field of time series forecasting. Given the promising merits of transformers in the context of time series forecasting, we will utilize transformers as the main model architecture in this work.

2.2 Visual time series forecasting

In recent years, a novel approach has gained traction in time series forecasting, which involves converting numeric data into images to leverage successful computer vision algorithms. This approach is motivated by the belief that visual representations of financial time series, such as charts and graphs, can enhance traders' decisionmaking. Prior research has explored representing time series data as images of line plots [30, 31], recurrent plots [20], and candlestick plots [7] for both forecasting and classification tasks. Furthermore, recent advancements in video prediction from the computer vision domain have been adapted to improve forecasting accuracy [38]. These studies have bridged the gap between computer vision and time series forecasting, offering innovative ways to capture and analyze temporal patterns in numerical data.

Given the inherent unpredictability in financial time series, researchers have turned to decomposition techniques to reveal more detailed information at varying frequency levels [5] for betterinformed portfolio management [2, 29] and risk modeling [3]. Recent works have explored time series forecasting or classification with frequency representation combined with deep learning [20, 23, 40, 43]. In our work, we used the wavelet transform as a time-frequency representation of the time series. The advantage of the wavelet transform over other frequency representations such as the Fourier transform lies in its ability to extract both local spectral and temporal information, resulting in a more comprehensive representation of the underlying data patterns. Researchers have already explored the utilization of time-frequency spectrograms for time series classification in the finance domain [10]. In our study, we adopt the time-frequency spectrograms for a harder task on time series forecasting. Specifically, we leverage a multimodal image representation consisting of a time-frequency spectrogram augmented with intensities of numerical time series data, capitalizing on the advantages of this combined approach.

Moreover, transformers have been extended into the computer vision domain by integrating self-attention with convolutional neural network (CNN) architectures [9]. Dosovitskiy et al. demonstrated that a pure transformer applied directly to sequences of image patches can achieve outstanding performance in image classification tasks. Recognizing this potential, our study adopts a vision transformer for simultaneous learning in both the time and frequency domains, capitalizing on the spectrogram representation of the time series data. The vision transformer's ability to handle cross-modality makes it well-suited for this task.

3 DATA

This paper analyzes three datasets out of which one is synthetic (Section 3.1) and the other two are real datasets (Section 3.2 and 3.3). This helps us to comprehensively evaluate spectrogram-based forecasting techniques as they encompass varying levels of periodicity and complexity, providing a comprehensive basis for the examination.

3.1 Synthetic Data

We created a synthetic dataset consisting of multiple periods by sampling data from harmonic functions. The initial dataset is artificially generated and aims to have complexity while still exhibiting a prominent and repeated signal. We synthesized the time series s_t using a linearly additive two-timescale harmonic generating function:

$$s_t = (A_1 + B_1 t) \sin(2\pi t/T_1 + \phi_1) + (A_2 + B_2 t) \sin(2\pi t/T_2 + \phi_2)$$
(1)

In the equation, the variable *t* ranges from t = 1 to t = T, where T represents the total length of the time series. The multiplicative amplitudes A_1 and A_2 are randomly selected from a Gaussian distribution $\mathcal{N}(1, 0.5)$. The amplitudes of the linear trends, B_1 and B_2 , are sampled from a uniform distribution $\mathcal{U}(-1/T, 1/T)$. The driving time scales, T_1 and T_2 , are relatively short and long compared to the total length T. Therefore, T_1 follows a normal distribution $\mathcal{N}(T/5, T/10)$, while T_2 follows a normal distribution $\mathcal{N}(T, T/2)$. Finally, the phase shifts ϕ_1 and ϕ_2 are chosen from a uniform distribution $\mathcal{U}(0, 2\pi)$. Overall, we generated 150K time series of length 100 that were used for our experiments. Each time series differ concerning the possible combination of tuning parameters. The first panel in Figure 2 displays a sampled instance of the synthetic data. It is evident that the synthetic time series comprises two distinct time scales: brief oscillations superimposed on much longer wave trains.

3.2 Temperature data

This dataset contains various temperature observations gathered by the Australian Bureau of Meteorology for 422 weather stations across Australia, between 02/05/2015 and 26/04/2017. The dataset was obtained from Monash repo [14]. The dataset contains several attributes as equal lenght time series (725 days). For our experiments, we used the daily mean temperature in Celsius for each station recorded over a 24-hour period. For every station, we sampled 30 time series of length 60 for every station. Overall We had 7320 time series. The second panel in Figure 2 displays a sampled instance of the temperature data.



Figure 2: Illustrations of the inputs of the three datasets: a) synthetic, b) temperature, and c) financial stock prices. The top panels show the raw time series represented as lineplots and the bottom panels depict the augmented time-frequency spectrogram. Each input time series consists of 80 steps for the synthetic and financial datasets, while the temperature dataset has 50 steps. For the financial and temperature data, each time step represents a 1-day time interval.

3.3 Financial data

We obtained the daily stock prices of S&P 500 constituent stocks through the Yahoo! finance database. The dataset comprises daily Adjusted Close values for the stocks that constitute the S&P-500 index, spanning from the year 2000 onwards. In total, our dataset includes around 58K individual time series, with each time series segment consisting of 100 days of data. The third panel in Figure 2 displays a sampled instance of the financial data.

4 METHOD

Our goal is time series forecasting using a vision transformer and a multimodal image of time-frequency spectrogram augmented with intensities of <u>num</u>erical time series, which we refer to as ViT-num-spec. This approach leverages the strengths of both the visual representation from the spectrogram and the numerical information to improve forecasting performance. The approach involves converting the numeric time series into an image using a time-frequency spectrogram (Section 4.2), and then employing a vision transformer encoder appended with a multi-layer perceptron (MLP) head to forecast the future (Section 4.3), as illustrated in Figure 3.

4.1 Preprocessing

The temperature dataset contained missing values, which were handled by performing forward filling. This means that any missing values were replaced with the most recent available value in the dataset, moving forward in time. For each of the three datasets, the data \mathbf{x} was scaled to ensure that the scaled \mathbf{x} fell within the range of [0, 1]. This scaled \mathbf{x} was then utilized to incorporate the original time series information into the spectrogram and to facilitate the learning of the target, as explained in more detail later.



Figure 3: Overview of the proposed approach.

4.2 Time-Frequency Spectrogram

To generate a time-frequency spectrogram from a given time series, two common methods are Short-Time Fourier transforms (STFT) [15] and wavelet transforms [8]. STFT involves a sliding window-based version of the Fourier transform, extracting frequency components within a fixed time window. However, it can only uncover a limited set of frequency components within this fixed window and may not be suitable for non-periodic and transient signals in stock price time series. On the other hand, wavelet transform is well-suited for analyzing time series with transient signals, as it can uncover varying frequency components [11, 25]. Unlike STFT, wavelet transform is not limited by a fixed time window, making it a more suitable choice for our study. Therefore, we choose to use wavelet transforms in our approach for generating time-frequency spectrograms from time series data.

A wavelet is a wave-like oscillation (zero-mean signal) that is localized in both time and frequency space. For our approach, we utilized the morlet wavelet (see Equation (2)), which has been demonstrated to be effective in time series classification [10].

$$\psi(x) = \sqrt{\frac{1}{s}} \pi^{-\frac{1}{4}} e^{\frac{1}{2}\frac{x}{s}} e^{jw\frac{x}{s}},$$
(2)

where *s* represents the scale of the wavelet. The morlet wavelet is essentially a sine wave multiplied by a Gaussian envelope centered at zero. The scale *s* determines the width and frequency of the wavelet. As *s* increases, the frequency of the wavelet becomes wider and lower. The wavelet transform is performed by convoluting wavelets at different scales with the time series. The magnitude of the resulting multiplication coefficients represents the signal strength of different wavelets. This is saved as a heatmap (or spectrogram), as depicted in Figure 2. The rows of the spectrogram correspond to varying wavelet scales, i.e., varying frequencies. The columns correspond to time. Higher frequency components are at the top, and lower frequency components are at the bottom as the *s* increases from the top to the bottom. Note that we augmented the spectrogram by adding an image stripe row at the top, and the resulting full image was used as the input image to our model. Throughout the paper, we will continue to refer to the augmented image as the spectrogram. The reason for this augmentation is as follows: the standard spectrogram only visualizes the strength of varying signals, but it does not retain the sign of the signal, which is available in the time domain. To preserve the sign information, we convert the standardized time series to an image row with intensities represented as integers in the range of [0, 255]. Examples of time-frequency spectrograms are provided in Figure 2.

4.3 Vision Transformer for Forecasting

We utilize a vision transformer encoder with an appended MLP head for time series forecasting, as illustrated in Figure 3. To effectively process the input image, we divide it into non-overlapping patches of equal size. This division results in the horizontal time axis being split into image-patch-sized time intervals. These patches are then converted to tokens through linear projection, and standard 1D position embedding is added. The objective of the forecasting model is to learn temporal dependencies between time and frequency patterns across the horizontal time axis. The encoder learns to encode these patches into latent feature vectors. In our experiments, we use an input image size of 128×128 , and image patches of size 16 \times 16. The top row of the price time series is scaled to 16 \times 128, and the spectrogram is scaled to 112×128 , ensuring that they are divisible by the patch size. The forecast network is constructed by appending a lightweight MLP head to the vision transformer encoder, as depicted in Figure 3. Further details on the network architecture and training can be found in Section 5.2.

5 EXPERIMENTS

In our experimental evaluation, we compared the performance of seven methods: 1) ViT-num-spec - the proposed method utilizing

From Pixels to Predictions: Spectrogram and Vision Transformer for Better Time Series Forecasting

the vision transformer and a multimodal image of time-frequency spectrogram augmented with intensities of numerical time series, 2) ViT-lineplot - utilizing the vision transformer on a lineplot visualization of the time series, 3) ViT-num - an ablation study on using only the intensities of time series as input, 4) DeepAR - a state-of-the-art LSTM-based time series forecasting method, as well as statistical baselines such as 5) EMA, and 6) ARIMA and 7) naive numeric forecasting baseline. The evaluation was conducted using two widely used metrics, namely SMAPE and MASE, to assess performance.

Along with traditional numerical time series forecasting methods, we showcased the advantages of using a visual spectrogram as a representation for time series data compared to other visual representations and an ablation study using only the intensities of time series as input. Additionally, we highlighted the benefits of employing a vision transformer for simultaneous learning in both the time and frequency domains. The experiments were conducted on a Linux machine equipped with 8 NVIDIA T4 Tensor Core GPUs, each with 16GB memory. PyTorch v1.0.0 DL platform in Python 3.6 was used for all models. To ensure reproducibility, we set a fixed random seed of 42.

5.1 Experimental setup

5.1.1 Synthetic Data. Our training dataset comprised 80K samples, while the validation dataset consisted of 20K samples. For the testing phase, we had a dataset of 50K samples available. Each time series consists of 100 time steps, and the task involved predicting the subsequent 20 time steps using the initial 80 time steps as input for each time series.

5.1.2 Temperature Data. For training purposes, data from 2015 and 2016 was utilized, while data from 2017 was reserved for testing. The construction of the training, testing, and validation datasets involved randomly sampling 10 time series of lenght 60 per station. Therefore, each of the train, test, and val datasets consisted of 4220 time series of length 60. The forecasting task entailed using the initial 50 time steps as the input to predict the subsequent 10 time steps for each individual time series. We maintained the length of the past timeseries as 50 to forecast the next 10 steps due to the small size of the daily temperature dataset, and to minimize the overlap in the generated time series. This means that we utilized the temperature data from the past 50 days to predict the temperature for the upcoming 10 days.

5.1.3 Financial Data. For the training data, we randomly sampled information from the period between 2000 and 2014. As for the test data, we sampled information from the years 2016 to 2019. The training set consists of a total of 46,875 time series with a 0.8 train and validation split, and we sampled 15,625 time series as the test set. The forecasting task involved using the first 80 time steps to predict the subsequent 20 time steps for each time series.

For each of the three datasets, the performance evaluation involved averaging metrics over the 20 predicted time steps for the synthetic and financial datasets, and over 10 predicted time steps for the temperature dataset. The overall performance on the test set is represented by the mean and standard deviation of the evaluation metrics calculated across the entire testing set for each dataset.

5.2 Models

5.2.1 ViT-num-spec. ViT-num-spec is our proposed method that utilizes the vision transformer on a multimodal image of time-frequency spectrogram augmented with intensities of numerical time series as the model input. The latent embedding vector has a dimension of 128. We trained the model using a batch size of 128 and the AdamW optimizer with a weight decay of 0.05. The training process was performed for a maximum of 200 epochs, and we implemented an early stopping mechanism with a patience of 10. The base learning rate and learning rate scheduler hyperparameters were tuned for each dataset.

5.2.2 ViT-lineplot. ViT-lineplot is a variant of our proposed method that utilizes the vision transformer on a lineplot visualization of the time series, without including the frequency information. It shares the same architecture with ViT-num-spec, including the loss function. Similar to ViT-num-spec, the base learning rate and learning rate scheduler hyperparameters were tuned for each dataset in the case of ViT-lineplot as well.

5.2.3 ViT-num. ViT-num is an ablation study of our proposed method that uses only the intensities of time series as input, specifically the top price row of a spectrogram image. Similar to ViT-lineplot, ViT-num shares the same architecture with ViT-num-spec, and the learning rate and scheduler were tuned for this variant as well.

5.2.4 DeepAR. DeepAR is an effective supervised learning algorithm used for forecasting univariate time series [28]. It leverages autoregressive recurrent neural networks, specifically built on LSTM architecture, to generate predictions. DeepAR generates probabilistic forecasts by training a single model simultaneously on a vast array of interconnected time series. In our implementation, we utilized DeepAR to forecast the value distribution for the next 20 steps in the harmonic and financial datasets, while using an input time series length of 80. For the temperature dataset, we used an input length of 50 to predict the next 10 steps. We utilized the normal distribution loss and employed the AdamW optimizer with a batch size of 128. The model underwent training for a maximum of 300 epochs, incorporating an early stopping mechanism with a patience of 15. The initial learning rate was set at 1e-3 and adjusted by a learning rate scheduler with a decay factor of 0.1 and a patience of 5. To prevent overfitting, a drop rate of 0.1 was applied. DeepAR generated 200 samples of the prediction target for each time series, and the final prediction was derived by calculating the average across these predicted samples.

5.2.5 ARIMA. ARIMA (AutoRegressive Integrated Moving Average) is a forecasting model that combines autoregressive (AR), differencing (I), and moving average (MA) components to capture dependencies and patterns in time series data, enabling predictions based on past observations [35]. By incorporating lagged values and previous errors, ARIMA models effectively capture trends, seasonality, and other temporal patterns.

To streamline the modeling process, we utilized the auto-ARIMA procedure [1]. This automated approach employs algorithms to systematically search for the optimal ARIMA parameter configuration, utilizing the Akaike Information Criterion. This automated parameter selection saves time and eliminates the need for manual tuning. Additionally, we implemented the stepwise algorithm described in [19] to further enhance the efficiency of the model selection process.

5.2.6 *EMA*. Exponential Moving Average (EMA) is a smoothing technique commonly employed in data analysis and forecasting. EMA distinguishes itself from the simple moving average by giving more significance to recent observations through the use of varying weights assigned to data points. This weighting scheme grants EMA the ability to swiftly detect and highlight short-term trends in the data. By iteratively updating the average, EMA combines the previous EMA value with the current observation, utilizing a smoothing factor known as the smoothing constant. This constant provides control over the level of responsiveness and smoothness exhibited by the EMA.

5.2.7 Naive. We developed a simplistic baseline model that relied on the assumption that the last observed value would serve as the prediction for all future time steps.

5.3 Evaluation metrics

To evaluate the performance of our models, the forecasted prices were initially transformed back to the original scale. All of our models produced continuous predictions, and we assessed their performance using three metrics: SMAPE (Symmetric Mean Absolute Percentage Error), MASE (Mean Absolute Scaled Error) and sign accuracy.

1) **SMAPE:** The symmetric mean absolute percentage error (SMAPE) is defined as,

$$SMAPE = \frac{1}{T} \sum_{t=1}^{T} \frac{|x_t - \hat{x}_t|}{(|x_t| + |\hat{x}_t|)/2},$$
(3)

where *T* denotes the length of the predicted time series (in our case, *T* = 1, 2, 3, 4, 5); x_t and \hat{x}_t is the observed ground truth and it's corresponding forecast at time step *t*. SMAPE values lie within [0, 2], with smaller values imply more accurate forecast.

2) **MASE:** The mean absolute scaled error (MASE) is defined as,

$$MASE = \frac{\frac{1}{T} \sum_{i=t+1}^{t+T} |x_i - \hat{x}_i|}{\frac{1}{t} \sum_{i=1}^{t} |x_i - x_{i-1}|}$$
(4)

It is the mean absolute error of the forecast divided by the mean absolute error of naive one-step forecast on the insample data. Note that MASE becomes unstable when the denominator approaches 0. Similar to SMAPE, lower MASE values indicate better forecasts.

3) Sign Accuracy: This represents the average classification accuracy across all test samples. A higher accuracy indicates more precise forecasts. We define three classes to classify forecast movements: class 1 for "going up", class 2 for "remaining flat," and class 3 for "going down". To determine the class, we compare the difference between the predicted forecasts and the last observed value in the input time series.

5.4 Results

We present a summary of our quantitative benchmark results in Table 1, which includes the mean and standard deviation of the evaluation metrics calculated across the entire testing set. Additionally, we provide qualitative examples comparing different methods across all three datasets in Figure 4.

Overall, our proposed approach, ViT-num-spec, demonstrates its superiority over all state-of-the-art statistical and deep learning baselines, firmly establishing its competitiveness in the field of numerical time series forecasting. Notably, it outperforms ViTlineplot and ViT-num, thus highlighting the advantages of utilizing a multimodal image of time-frequency spectrogram augmented with intensities of numerical time series as the visual representation for time series data. The successful fusion of the multimodal image representation and the vision transformer's simultaneous learning in both the time and frequency domains plays a crucial role in significantly enhancing the forecasting performance. This distinct capability sets ViT-num-spec apart from traditional time series forecasting methods, providing it with a clear advantage in effectively handling complex temporal dependencies. In the following sections, we provide a dataset-wise breakdown of the results to offer a comprehensive understanding of ViT-num-spec's performance across different datasets.

5.4.1 Synthetic data. Among all the evaluated methods, all the ViT variants outperform all baseline models. This shows that a vision transformer model can successfully capture the patterns present in the input signal. Among the Vit methods, ViT-num-spec performing the best, followed by ViT-num, and then ViT-lineplot. Moreover, ViT-num-spec's superiority over ViT-num demonstrates the benefits of using a multimodal setup that includes both spectrogram and intensities. Lastly, ViT-num-spec's outperformance over ViT-lineplot underscores the superiority of using time-frequency spectrograms instead of simple lineplot visualizations for representing time series data. All these findings emphasizes the advantage of utilizing multimodal representations for time series forecasting, where the combination of spectrogram and intensities contributes to improved performance in capturing complex temporal patterns.

5.4.2 Temperature data. Among all the evaluated methods, ViTnum-spec emerges as the top performer in terms of SMAPE and MASE, However, DeepAR achieves the highest accuracy in sign predictions. Note that the sign accuracy of ViT-num-spec is better than other ViT variants. While sign accuracy is crucial for making correct directional predictions, SMAPE and MASE are considered more crucial for achieving precise temperature forecasts. The performance of ViT-num-spec is still better than other ViT variants. This trend is consistent with what was observed in the synthetic datasets. This further highlights the significance of incorporating a vision transformer in a multimodal setup that includes both spectrogram and intensities.

5.4.3 *Financial data.* Among the three datasets, financial time series present the most significant forecasting challenges. Prior research [24] indicates that financial data exhibits properties resembling randomness at short scales and lacks evident periodic patterns at larger scales. When considering SMAPE and MASE, all methods demonstrate similar performance, except for ViT-lineplot,

From Pixels to Predictions: Spectrogram and Vision Transformer for Better Time Series Forecasting

ICAIF '23, November 27-29, 2023, Brooklyn, NY, USA

	Baselines				ViT		
SMAPE ↓	Naive	EMA	ARIMA	DeepAR	ViT-num	ViT-lineplot	ViT-num-spec
Synthetic	1.234 ± 0.442	1.242 ± 0.447	0.549 ± 0.379	1.073 ± 0.403	0.376 ± 0.294	0.531 ± 0.346	$\textbf{0.304} \pm \textbf{0.258}$
Temperature	0.134 ± 0.110	0.132 ± 0.109	0.136 ± 0.109	0.128 ± 0.102	0.132 ± 0.104	0.135 ± 0.107	$\textbf{0.125} \pm \textbf{0.100}$
Financial	0.036 ± 0.028	0.036 ± 0.028	0.041 ± 0.035	0.038 ± 0.030	0.036 ± 0.028	0.068 ± 0.054	0.036 ± 0.028
MASE ↓							
Synthetic	5.272 ± 3.892	5.265 ± 3.899	2.194 ± 2.679	4.225 ± 4.092	0.921 ± 1.078	1.518 ± 1.704	$\textbf{0.757} \pm \textbf{0.974}$
Temperature	1.497 ± 0.885	1.460 ± 0.862	1.499 ± 0.818	1.446 ± 0.792	1.513 ± 0.843	1.526 ± 0.904	$\textbf{1.413} \pm \textbf{0.793}$
Financial	3.486 ± 2.789	3.487 ± 2.788	3.884 ± 3.076	3.659 ± 2.905	3.464 ± 2.796	6.383 ± 4.344	$\textbf{3.461} \pm \textbf{2.801}$
Sign accuracy ↑							
Synthetic	0.0%	44.3%	88.0%	75.1%	94.4%	90.9%	95.7%
Temperature	0.2%	59.5%	60.4%	65.1%	61.3%	60.7%	64.2%
Financial	0.3% / 0.0%	49.7% / 50.7%	50.6% / 50.9	51.5% / 53.0%	54.3% / 57.4%	50.3% / 50.5%	54.4% / 58.4%

Table 1: Summary of the evaluation metrics: For the financial dataset, we report both the original/thresholded sign accuracies (threshold is 20% of the standard deviation of past time series data). The results are reported as the mean \pm standard deviation of the evaluation metrics calculated across the entire testing set.

which performs the worst among the evaluated methods. Upon analyzing sign accuracy, it becomes evident that ViT-num-spec and ViT-num outperform the other methods. Thus, lineplots are not a good visual representation of financial data. For practical decision-making in financial scenarios, buying or selling actions are usually executed only when the predicted change exceeds a certain significant threshold. To address this, we apply a thresholding technique where the threshold value is set to 20% of the standard deviation of the past time series data. By using this threshold, we evaluate the sign accuracy of the predictions, which provides a more realistic representation of the model's performance in making actionable decisions for financial forecasting. Indeed, when considering the threshold accuracy, ViT-num-spec emerges as the top performer, closely followed by ViT-num. This reinforces the advantage of using spectrograms over lineplots for representing time series data, as demonstrated by the superiority of both ViTnum-spec and ViT-num over ViT-lineplot. Furthermore, the ground truth statistics reveal that the distribution of positive, negative, and neutral signs in the dataset is 55.7%, 44.0%, and 0.3% respectively. When considering the dominant trend as the forecasted sign, we achieve an accuracy of 55.7%, which is lower than ViT-num-spec's performance after applying thresholding.

In summary, the results validate the effectiveness of our proposed approach, which combines the multimodal image representation and the vision transformer's simultaneous learning in both the time and frequency domains. This fusion plays a crucial role in significantly enhancing the forecasting performance, demonstrating the capability of achieving state-of-the-art results in time series forecasting.

6 CONCLUSION

In conclusion, this paper introduces a novel approach for time series forecasting, utilizing time-frequency spectrograms as visual representation of raw time sereis and incorporating a vision transformer for multimodal learning. We compare our method against benchmarked methods, including statistical baselines and state-ofthe-art deep learning-based approaches. We also conduct experiments using lineplot as the visual representation of time series and an ablation study using only the intensities of time series to demonstrate the superiority of our proposed approach in forecasting across diverse datasets, including synthetic, temperature, and financial stock price data.

ACKNOWLEDGMENTS

This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates ("J.P. Morgan") and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

REFERENCES

- 2023. Auto ARIMA. https://www.rdocumentation.org/packages/forecast/ versions/8.21/topics/auto.arima. Accessed: 2023-06-20.
- [2] Alvaro Arroyo, Bruno Scalzo, Ljubiša Stanković, and Danilo P Mandic. 2022. Dynamic Portfolio Cuts: A Spectral Approach to Graph-Theoretic Diversification. In ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 5468–5472.
- [3] Federico M. Bandi, Shomesh E. Chaudhuri, Andrew W. Lo, and Andrea Tamoni. 2021. Spectral factor models. *Journal of Financial Economics* 142, 1 (2021), 214–238. https://doi.org/10.1016/j.jfineco.2021.04.024
- [4] Wei Bao, Jun Yue, and Yulei Rao. 2017. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PloS one* 12, 7 (2017), e0180944.
- [5] Shomesh Chaudhuri and Andrew W Lo. 2016. Spectral portfolio theory. Available at SSRN 2788999 (2016).
- [6] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014).
- [7] Naftali Cohen, Tucker Balch, and Manuela Veloso. 2020. Trading via image classification. In Proceedings of the First ACM International Conference on AI in Finance. 1–6.
- [8] Ingrid Daubechies. 1990. The wavelet transform, time-frequency localization and signal analysis. *IEEE transactions on information theory* 36, 5 (1990), 961–1005.



Figure 4: Qualitative examples for predictions for the three datasets: a) synthetic, b) temperature, and c) financial stock prices.

- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020).
- [10] Bairui Du and Paolo Barucca. 2020. Image Processing Tools for Financial Time Series Classification. arXiv preprint arXiv:2008.06042 (2020).
- [11] Mordechai Frisch and Hagit Messer. 1992. The use of the wavelet transform in the detection of an unknown transient signal. *IEEE Transactions on Information Theory* 38, 2 (1992), 892–897.
- [12] Everette S Gardner Jr and ED McKenzie. 1985. Forecasting trends in time series. Management Science 31, 10 (1985), 1237–1246.
- [13] Andre Gensler, Janosch Henze, Bernhard Sick, and Nils Raabe. 2016. Deep Learning for solar power forecasting—An approach using AutoEncoder and LSTM Neural Networks. In 2016 IEEE international conference on systems, man, and cybernetics (SMC). IEEE, 002858–002865.
- [14] Rakshitha Godahewa, Christoph Bergmeir, Geoffrey I. Webb, Rob J. Hyndman, and Pablo Montero-Manso. 2021. Monash Time Series Forecasting Archive. In Neural Information Processing Systems Track on Datasets and Benchmarks.
- [15] Daniel Griffin and Jae Lim. 1984. Signal estimation from modified short-time Fourier transform. IEEE Transactions on acoustics, speech, and signal processing 32, 2 (1984), 236–243.
- [16] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2001. The elements of statistical learning. Springer series in statistics. New York, NY, USA (2001).

From Pixels to Predictions: Spectrogram and Vision Transformer for Better Time Series Forecasting

- [17] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. Neural computation 9, 8 (1997), 1735–1780.
- [18] Charles C Holt. 2004. Forecasting seasonals and trends by exponentially weighted moving averages. *International journal of forecasting* 20, 1 (2004), 5–10.
- [19] Rob J Hyndman and Yeasmin Khandakar. 2008. Automatic time series forecasting: the forecast package for R. Journal of statistical software 27 (2008), 1–22.
- [20] Xixi Li, Yanfei Kang, and Feng Li. 2020. Forecasting with time series imaging. Expert Systems with Applications 160 (2020), 113680.
- [21] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2020. The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting* 36, 1 (2020), 54–74.
- [22] Massimiliano Marcellino, James H Stock, and Mark W Watson. 2006. A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of econometrics* 135, 1-2 (2006), 499–526.
- [23] Fernando Moreno-Pino, Pablo M Olmos, and Antonio Artés-Rodríguez. 2023. Deep autoregressive models with spectral attention. *Pattern Recognition* 133 (2023), 109014.
- [24] Lasse Heje Pedersen. 2019. Efficiently inefficient: how smart money invests and market prices are determined. Princeton University Press.
- [25] Sharif Md Raihan, Yi Wen, and Bing Zeng. 2005. Wavelet: A new tool for business cycle analysis. Federal Reserve Bank of St. Louis Working Paper Series 2005-050 (2005).
- [26] Pablo Romeu, Francisco Zamora-Martínez, Paloma Botella-Rocamora, and Juan Pardo. 2015. Stacked denoising auto-encoders for short-term time series forecasting. In Artificial Neural Networks. Springer, 463–486.
- [27] Alaa Sagheer and Mostafa Kotb. 2019. Time series forecasting of petroleum production using deep LSTM recurrent networks. *Neurocomputing* 323 (2019), 203–213.
- [28] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.
- [29] Bruno Scalzo, Alvaro Arroyo, Ljubiša Stanković, and Danilo P Mandic. 2021. Nonstationary portfolios: Diversification in the spectral domain. In ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 5155–5159.
- [30] Artemios-Anargyros Semenoglou, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2023. Image-based time series forecasting: A deep convolutional neural network approach. *Neural Networks* 157 (2023), 39–53.

- [31] Srijan Sood, Zhen Zeng, Naftali Cohen, Tucker Balch, and Manuela Veloso. 2021. Visual time series forecasting: an image-driven approach. In Proceedings of the Second ACM International Conference on AI in Finance. 1–9.
- [32] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. Advances in neural information processing systems 27 (2014).
- [33] Souhaib Ben Taieb, Antti Sorjamaa, and Gianluca Bontempi. 2010. Multipleoutput modeling for multi-step-ahead time series forecasting. *Neurocomputing* 73, 10-12 (2010), 1950–1957.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- [35] Daniel S Wilks. 2011. Statistical methods in the atmospheric sciences. Vol. 100. Academic press.
- [36] Peter R Winters. 1960. Forecasting sales by exponentially weighted moving averages. Management science 6, 3 (1960), 324-342.
- [37] Neo Wu, Bradley Green, Xue Ben, and Shawn O'Banion. 2020. Deep transformer models for time series forecasting: The influenza prevalence case. arXiv preprint arXiv:2001.08317 (2020).
- [38] Zhen Zeng, Tucker Balch, and Manuela Veloso. 2021. Deep video prediction for time series forecasting. In Proceedings of the Second ACM International Conference on AI in Finance. 1–7.
- [39] Zhen Zeng, Rachneet Kaur, Suchetha Siddagangappa, Saba Rahimi, Tucker Balch, and Manuela Veloso. 2023. Financial Time Series Forecasting using CNN and Transformer. arXiv preprint arXiv:2304.04912 (2023).
- [40] Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. 2022. Self-supervised contrastive pre-training for time series via time-frequency consistency. Advances in Neural Information Processing Systems 35 (2022), 3988–4003.
- [41] Zihao Zhang, Stefan Zohren, and Stephen Roberts. 2019. Deeplob: Deep convolutional neural networks for limit order books. *IEEE Transactions on Signal Processing* 67, 11 (2019), 3001–3012.
- [42] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In Proceedings of the AAAI conference on artificial intelligence, Vol. 35. 11106–11115.
- [43] 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 *International Conference on Machine Learning*. PMLR, 27268– 27286.