TI-MAE: SELF-SUPERVISED MASKED TIME SERIES AUTOENCODERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Multivariate Time Series forecasting has been an increasingly popular topic in various applications and scenarios. Recently, contrastive learning and Transformer-based models have achieved good performance in many long-term series forecasting tasks. However, there are still several issues in existing methods. First, the training paradigm of contrastive learning and downstream prediction tasks are inconsistent, leading to inaccurate prediction results. Second, existing Transformer-based models which resort to similar patterns in historical time series data for predicting future values generally induce severe distribution shift problems, and do not fully leverage the sequence information compared to self-supervised methods. To address these issues, we propose a novel framework named Ti-MAE, in which the input time series are assumed to follow an integrate distribution. In detail, Ti-MAE randomly masks out embedded time series data and learns an autoencoder to reconstruct them at the point-level. Ti-MAE adopts mask modeling (rather than contrastive learning) as the auxiliary task and bridges the connection between existing representation learning and generative Transformer-based methods, reducing the difference between upstream and downstream forecasting tasks while maintaining the utilization of original time series data. Experiments on several public real-world datasets demonstrate that our framework of masked autoencoding could learn strong representations directly from the raw data, yielding better performance in time series forecasting and classification tasks. The code will be made public after this paper is accepted.

1 INTRODUCTION

Time series modeling has an urgent need in many fields, such as time series classification (Dau et al., 2019), demand forecasting (Carbonneau et al., 2008), and anomaly detection (Laptev et al., 2017). Recently, long sequence time series forecasting (LSTF), which aims to predict the change of values in a long future period, has aroused significant interests of researchers. In the previous work, most of the self-supervised representation learning methods on time series aim to learn transformation-invariant features via contrastive learning to be applied on downstream tasks. Although these methods perform well on classification tasks, there is still a gap between their performance and other supervised models on forecasting tasks. Apart from the inevitable distortion to time series caused by augmentation strategies they have borrowed from vision or language, the inconsistency between upstream contrastive learning approaches and downstream forecasting tasks should be also a major cause of this problem. Besides, as the latest contrastive learning frameworks (Yue et al., 2022; Woo et al., 2022a) reported, Transformer (Vaswani et al., 2017) performs worse than CNN-based backbones, which is also not consistent with our experience. We have to reveal the differences and relationships between existing contrastive learning and supervised methods on time series.

As an alternative of contrastive learning, denoising autoencoders (Vincent et al., 2008) are also used to be an auxiliary task to learn intermediate representation from the data. Due to the ability of Transformer to capture long-range dependencies, many of existing methods (Zhou et al., 2021; Wu et al., 2021; Woo et al., 2022b) focused on reducing the time complexity and memory usage caused by vanilla attention mechanism such as sparse attention or correlation to process longer time series. These transformer-based models all follow the same training paradigm, which learns similar patterns from input historical time series segments and predict future time series values from

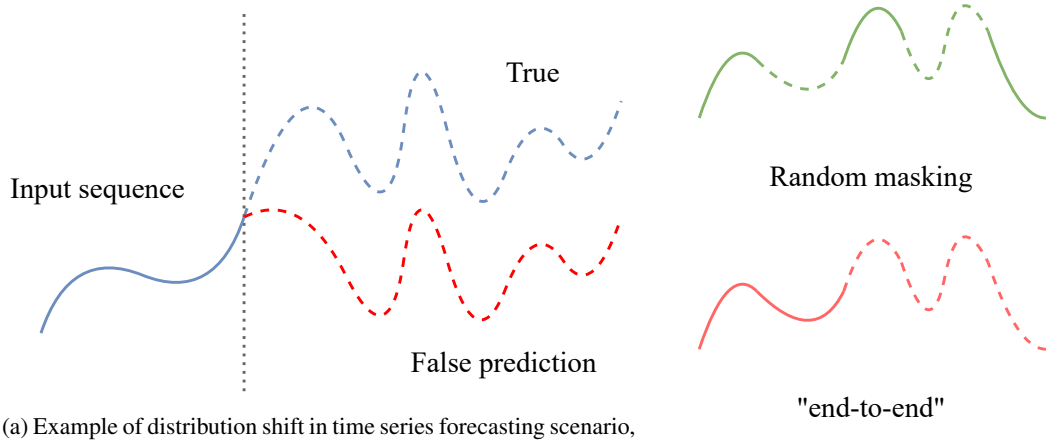

(a) Example of distribution shift in time series forecasting scenario, where the input time series and the prediction should have almost the same shape or seasonality information but different mean value caused by the change of trend.

(b) Two types of masking strategies. Dashed lines indicate the region to be masked out.

Figure 1: Different masking strategies in generative Transformer-based models on time series.

captured patterns. These so-called generative Transformer-based models are actually a special kind of denoising autoencoders, where we only mask the future values and reconstruct them.

However, this fixed masking strategy is usually accompanied by two severe problems. For one thing, fixed masking strategy will limit the learning ability of the model, which captures only the information of the visible sequence and some mapping relationship between the historical and the future segments. Similar problems have been reported in vision tasks (Zhang et al., 2017) . For another, fixed masking strategy will induce severe distribution shift problems as 1a shows, especially when the prediction horizon is longer than input sequence. In reality, most of the time series data collected from real scenarios are non-stationary, whose mean or variance changes over time. Similar problems were also observed in previous studies (Qiu et al., 2018; Oreshkin et al., 2020; Wu et al., 2021; Woo et al., 2022a). Most of them have tried to disentangle the input time series into a trend part and a seasonality part in order to enhance the capture of periodic features and to make the model robust to outlier noises. Specifically, they utilize moving average implemented by one average pooling layer with a fixed size sliding window to gain trend information of input time series. Then they capture seasonality features from periodic sequences, which are obtained by simply subtracting trend items from the original signal. To further clarify the mechanism of this disentanglement, we intuitively propose an easy but comprehensible description of disentangled time series as

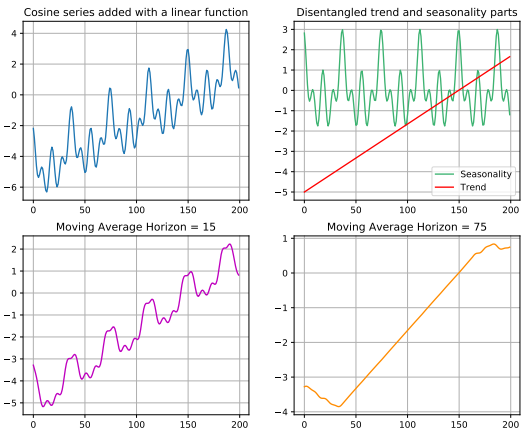

Figure 2: Example of disentanglement. **Top Left**: Simulated input cosine series added with a linear trend. **Top Right**: The true trend and seasonality parts of the input. **Bottom Left**: Disentangled trend part though average pooling with the sliding window size of 15. **Bottom Right**: Disentangled trend part though average pooling with the sliding window size of 75.

$$\boldsymbol{y}(t) = \textit{Trend}(t) + \textit{Seasonality}(t) + \text{Noises}. \tag{1}$$

For better illustration, we simply use polynomial series $\sum_n t^n$ and Fourier cosine series $\sum_n \cos^n t$ to respectively describe trend parts and seasonality parts of the original time series in Eq.(1). Apparently, the seasonality part is stationary when we set a proper observation horizon (not less than the maximum period of seasonality parts), while the moments of the trend part change continuously over

time. Figure 2 illustrates that the size of sliding window in average pooling layer plays a vital role in the quality of disentangled trend part. Natural time series data generally have more complex periodic patterns, which means we have to employ longer sliding windows or other hierarchical disposals. In addition, when moving average is used to capture the trend parts, both ends of a sequence need to be padded for alignment, which causes inevitable data distortion at the head and tail. These observed phenomenons suggest there are still some unresolved issues in the current disentanglement.

To address these issues, this paper proposes a novel Transformer-based framework named Ti-MAE as shown in Figure 3. Ti-MAE randomly masks out parts of embedded time series data and learns an autoencoder to reconstruct them at the point-level in the training stage. Figure 1b shows the difference between random masking and fixed continuous masking in end-to-end models, where we adequately leverage all the input sequence with different combination of visible tokens. Random masking takes the overall distribution of inputs into consideration, which can therefore alleviate the distribution shift problem. Moreover, with the power of pre-training or representation learning embodied in the encoder-decoder structure, Ti-MAE provides a universal scheme for both forecasting and classification. The contributions of our work are summarized as follows:

- We provide a novel perspective to bridge the connection between existing contrastive learning and generative Transformer-based models on time series and point out the inconsistency and deficiencies of them on downstream tasks.
- We propose Ti-MAE, a masked time series autoencoders which can learn strong representations with less inductive bias or hierarchical trick. Masking time series modeling in training stage adequately leverages the input data and successfully alleviates the distribution shift problem. Due to the flexible setting of masking ratio, Ti-MAE can adapt to complex scenarios which require the trained model to make forecasting simultaneously for multiple time windows with various sizes without re-training.
- Ti-MAE has achieved excellent performance for both forecasting and classification tasks on several public real-world time series datasets, demonstrating the power of pre-training or representation learning of Ti-MAE in the time series domain.

2 Related Work

2.1 Transformer-based time series model

Due to the ability of Transformer to capture long-range dependencies, Transformer-based model has been widely used in language and vision tasks. Song et al. (2018); Ma et al. (2019); LI et al. (2019) tried to directly apply vanilla Transformer to time series data but failed in long sequence time series forecasting tasks as self-attention operation scales quadratically with the input sequence length. Child et al. (2019); Zhou et al. (2021); Liu et al. (2022) noticed the long tail distribution in self-attention feature map so that they utilized sparse attention mechanism to reduce time complexity and memory usage of vanilla Transformer for processing longer sequences. Unfortunately, applying too long input sequence in training stage will degrade the forecasting accuracy of the model (Wu et al., 2021), which is in contrast to the ability that Transformer-based model can capture long-range dependencies. Some of the latest works like ETSformer (Woo et al., 2022b) and FEDformer (Zhou et al., 2022) also rely heavily on disentanglement and extra introduced domain knowledge.

2.2 Time Series Representation Learning

Self-supervised representation learning has achieved good performance in time series domain, especially using contrastive learning to learn a good intermediate representation. Lei et al. (2019); Franceschi et al. (2019) used loss function of metric learning to preserve pairwise similarities in the time domain. CPC (van den Oord et al., 2018) first proposed contrastive predictive coding and InfoNCE, which treats the data from the same sequence as positive pairs while the different noise data from the mini-batch as negative pairs. Different data augmentations on time series data were proposed to capture transformation-invariant features at semantic level (Eldele et al., 2021; Yue et al., 2022). CoST (Woo et al., 2022a) introduced extra inductive biases in frequency domain through DFT and separately processed disentangled trend and seasonality parts of the original time series

data to encourage discriminative seasonal and trend representations. Almost all of these methods rely on heavily data augmentation or other domain knowledge like hierarchy and disentanglement.

2.3 Masked Data Modeling

Masked language modeling is a widely adapted method for pre-training in NLP. BERT (Devlin et al., 2019) holds out a portion of the input sequence and predicts the missing content in training stage, which can generate good representations to various downstream tasks. Masked image encoding methods are also used for learning image representations. Pathak et al. (2016) recovered a small portion of missing regions using convolution. Motivated by the huge successes in NLP, recent methods (Bao et al., 2021; Dosovitskiy et al., 2021) are resort to Transformers to predict unknown pixels. MAE (He et al., 2021) proposed to mask a high portion of patches and retain a small set of visible patches received by encoder in pre-training on image data. Zerveas et al. (2021) directly masked a small portion of time series to learn representations. Pang et al. (2022); Tong et al. (2022); Feichtenhofer et al. (2022) have shown that MAE-style methods are effective to learn good representations. Specially designed for enhancing the performance of spatial-temporal graph neural networks, TSformer (Shao et al., 2022) concurrently tried to use MAE to generate intermediate representations for spatial-temporal data. In addition, ExtraMAE (fang Zha, 2022) used RNN as the backbone with the masking scheme for fast time series generation. Different from these methods, Ti-MAE inherits the advantages of Transformer in modeling long time dependencies, and develops a universal time series generation method for both forecasting and classification, with outstanding performance in comparison with state-of-the-art transformer time series models and contrastive learning methods.

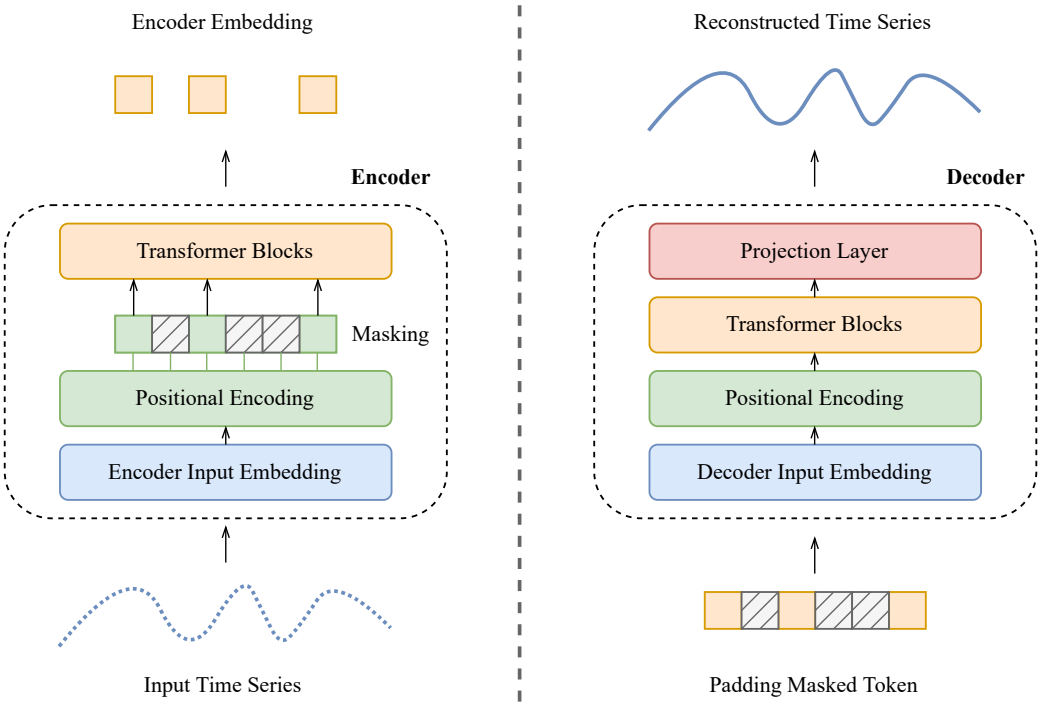

Figure 3: **Ti-MAE structure overview**. **Left**: The encoder receives raw time series inputs. After embedding inputs into tokens on timestamp, we randomly mask a large subset of tokens. Then we feed all the visible tokens into Transformer blocks to capture dependencies. **Right**: The lighter decoder processes encoded tokens padded with masked tokens and reconstructs the original time series at the point-level.

3 METHODOLOGY

3.1 PROBLEM DEFINITION

Let $\mathcal{X} = (\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_T) \in \mathbb{R}^{T \times m}$ be a multivariate time series instance with length of T , where m is the dimension of each signal. Given a historical multivariate time series segment $\mathcal{X}_h \in \mathbb{R}^{h \times m}$ with length of h, forecasting tasks aim to predict the next k steps values of $\mathcal{X}_f \in \mathbb{R}^{k \times m}$. For classification tasks, we should match the categorical ground truth from a set of labels $\mathcal{C}$ and each time series instance $\mathcal{X}$.

3.2 MODEL ARCHITECTURE

The overall architecture of Ti-MAE is shown in Figure 3. Similar as all autoencoders, our framework has an encoder that maps the observed time series signal $\mathcal{X} \in \mathbb{R}^{T \times m}$ to a latent representation $\mathcal{H} \in \mathbb{R}^{T \times n}$, and a decoder that reconstructs the original sequence from the embedding of the encoder on timestamp. Motivated by the great success of other MAE-style approaches (He et al., 2021; Feichtenhofer et al., 2022; Hou et al., 2022), we also adopt an asymmetric design that the encoder only operates visible tokens after applying masking on input embedding, and a lighter decoder processes encoded tokens padded with masked tokens and reconstructs the original time series at the point-level. More details of each component are introduced as follows.

Input embedding. Unlike other time series modeling methods, we have not adopted any multi-scale or complex convolution scheme like dilated convolution. Given a time series segment, we directly use one 1-D convolutional layer to extract local temporal features on timestamp across channels. Fixed sinusoidal positional embeddings are added to maintain the position information. Be different from other temporal data embedding approaches, we do not add any handcrafting task-specific or date-specific embeddings so as to introduce as little inductive bias as possible.

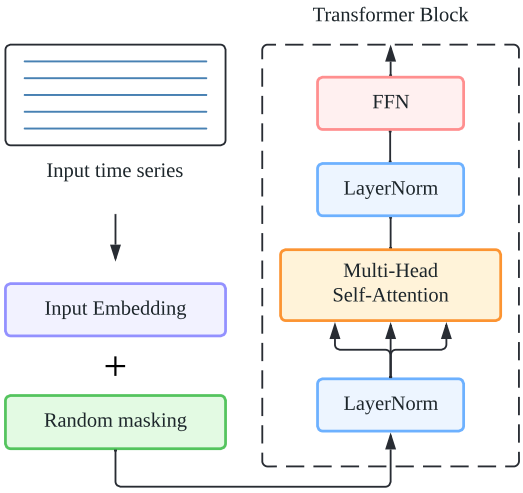

Figure 4: Ti-MAE encoder overview. **Left**: Encoder input embedding. **Right**: Details of one Transformer block used in both Ti-MAE encoder and decoder, where we utilize pre-norm instead of post-norm scheme.

Masking. After tokenizing original temporal data into tokens on timestamp, we randomly sample a subset of tokens without replacement which follows the uniform distribution and mask the remaining parts. It is hypothesized and summarized in (He et al., 2021; Feichtenhofer et al., 2022) that the masking ratio is related to the information density and redundancy of the data, which has an immense impact on the performance of the autoencoders. Generally speaking, natural language has higher information density due to its highly discrete word distribution, while images are of heavy spatial redundancy. Specifically, single pixel in one image has lower semantic information so that we can reconstruct a missing region from neighboring pixels by interpolation with little understanding of contents. Thus, data with lower information density should be applied a higher masking ratio to largely eliminate redundancy and prevent the model from focusing only on low-level semantic information. As a benchmark model often used in natural language, BERT (Devlin et al., 2019) uses a masking ratio of 15% while MAE uses a ratio of 75% for images (He et al., 2021) and 90% for videos (Feichtenhofer et al., 2022). Similar as images, time series data also have local continuity so that we should determine a high masking ratio in training stage. The optimal masking ratio of multivariate time series we observe is also around 75%.

Ti-MAE Encoder. Our encoder is a set of vanilla Transformer blocks with input embedding but utilizes pre-norm instead of post-norm in each block, which is shown as Figure 4. Like other MAE-style methods, Ti-MAE's encoder is applied only on visible tokens after embedding and random

masking. This design significantly reduces time complexity and memory usage compared to full encoding.

Ti-MAE Decoder. Our decoder also contains a set of vanilla Transformer blocks applied on the union of the encoded visible tokens and learnable randomly initialized mask tokens. Following (He et al., 2021), the decoder is designed to be smaller than the encoder. Notably, we add positional embeddings to all tokens after padding to supplement the location information of the missing parts. The last layer of the decoder is a linear projection layer which reconstructs the input by predicting all the values at the point-level. The training loss function is the mean squared error (MSE) between the original time series data and the prediction over masking regions.

The encoder and decoder of Ti-MAE are both agnostic to the sequential data with as less domain knowledge as possible. There is no date-specific embedding, hierarchy or disentanglement in contrast to other architectures (Zhou et al., 2021; Wu et al., 2021; Yue et al., 2022; Woo et al., 2022a). Compared to masked autoencoders used in vision tasks, a lot of parameter settings of Ti-MAE have been adjusted to better fit the time series data. We keep the point-level modeling rather than patch embedding for the consistency between masked modeling and downstream forecasting tasks. Unlike (Shao et al., 2022), we directly generate future values from the decoder as prediction, maintaining the consistency of training and inference stages.

4 Experiments

4.1 Experimental Setup

Datasets. We conduct extensively experiments on several public real-world datasets, covering time series forecasting and classification applications. (1) ETT (Electricity Transformer Temperature) (Zhou et al., 2021) consists of the data collected from electricity transformers, recording six power load features and oil temperature. (2) Weather[1] contains 21 meteorological indicators like humidity, pressure in 2020 year from nearly 1600 locations in the U.S.. (3) Exchange (Lai et al., 2018) is a collection of exchange rates among eight different countries from 1990 to 2016. (4) ILI[2] records the weekly influenza-like illness (ILI) patients data from Centers for Disease Control and Prevention of the United States between 2002 and 2021, describing the ratio of patients observed with ILI and the total number of patients. (5) The UCR archive (Dau et al., 2019) has 128 different datasets covering multiple domains like object outlines, traffic and body posture. We follow the same protocol and split all forecasting datasets into training, validation and test set by the ratio of 6:2:2 for the ETT dataset and 7:1:2 for other datasets. For classification, each dataset of UCR archive has been already divided into training and test set where the size of test set is greatly larger than training set in order to be accord with the actual scenarios.

Baselines. We select two types of baselines, end-to-end and representation learning methods which have public official codes. For time series forecasting tasks, we select four latest state-of-the-art representation learning models: CoST (Woo et al., 2022a), TS2Vec (Yue et al., 2022), TNC (Tonekaboni et al., 2021) and MoCo (Chen et al., 2021) applied on time series and four end-to-end models: Autoformer (Wu et al., 2021), Informer (Zhou et al., 2021), LogTrans (LI et al., 2019) and TCN (Bai et al., 2018). For time series classification tasks, we include more competitive unsupervised representation learning methods: TS2Vec, T-Loss (Franceschi et al., 2019), TS-TCC (Eldele et al., 2021), TST (Zerveas et al., 2021), TNC (Tonekaboni et al., 2021) and DTW (Chen et al., 2013).

Implementation Details. The encoder and decoder of Ti-MAE both use 2 layers of vanilla Transformer blocks with 4 heads self-attention. The number of hidden states dimension is set to 64, which is significantly lower than other existing methods (e.g., 320, 512). Ti-MAE is trained with MSE loss, using the Adam optimizer (Kingma & Ba, 2015) with an initial learning rate of $1e-3$. We apply a batch size of 64 and sampling time of 30 in each iteration. We use mean squared error (MSE) $\frac{1}{n}\sum_{i=1}^{n}(\boldsymbol{y}-\hat{\boldsymbol{y}})^2$ and mean absolute error (MAE) $\frac{1}{n}\sum_{i=1}^{n}|\boldsymbol{y}-\hat{\boldsymbol{y}}|$ as evaluation metrics on forecasting tasks, and average accuracy with critical difference (CD) on classification tasks. All the models are implemented in PyTorch (Paszke et al., 2019) and trained/tested on a single Nvidia V100 32GB GPU.

[1] https://www.ncei.noaa.gov/data/local-climatological-data/

[2] https://gis.cdc.gov/grasp/fluview/fluportaldashboard.html

Table 1: Multivariate time series forecasting results compared to representation learning methods.

Method		Ti-MAE		CoST		TS2Vec		TNC		MoCo	
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh	12	**0.2629**	**0.3462**	0.3374	0.4001	0.5817	0.5217	0.6056	0.5389	0.6419	0.5518
	24	**0.3520**	**0.3924**	0.3832	0.4301	0.5897	0.5312	0.6331	0.5616	0.6491	0.5630
	48	**0.3977**	**0.4173**	0.4342	0.4665	0.6242	0.5545	0.6934	0.6001	0.6804	0.5867
	96	**0.4266**	**0.4301**	0.5229	0.5201	0.6812	0.5699	0.7538	0.6391	0.7618	0.6323
	128	**0.4493**	**0.4436**	0.5742	0.5529	0.7190	0.5919	0.7949	0.6622	0.8062	0.6577
	168	**0.5091**	**0.4594**	0.6326	0.5838	0.7621	0.6387	0.8360	0.6849	0.8201	0.6742
Weather	12	**0.0932**	**0.1460**	0.1652	0.2630	0.1481	0.2367	0.1819	0.2508	0.1642	0.2521
	24	**0.1226**	**0.1809**	0.2719	0.3525	0.3011	0.3551	0.3118	0.3731	0.3112	0.3651
	48	**0.1633**	**0.2280**	0.3662	0.3672	0.3741	0.4178	0.3803	0.4117	0.3717	0.4163
	96	**0.2123**	**0.2735**	0.4119	0.4266	0.4289	0.4507	0.4176	0.4174	0.4077	0.4419
	128	**0.2197**	**0.2805**	0.4302	0.4686	0.4663	0.4839	0.4569	0.4824	0.4582	0.4693
	168	**0.2460**	**0.3049**	0.4636	0.4914	0.4909	0.5061	0.4789	0.4950	0.4820	0.4992
Exchange	24	**0.0697**	**0.1889**	0.1365	0.2721	0.0873	0.2245	0.0834	0.2084	0.1058	0.2553
	48	**0.1255**	**0.2448**	0.2532	0.3783	0.1666	0.3047	0.1648	0.2928	0.2018	0.3588
	96	**0.1912**	**0.2972**	0.5408	0.5645	0.4686	0.5098	0.3756	0.4510	0.4162	0.5002
	128	**0.2208**	**0.3242**	0.6786	0.6334	0.6540	0.6036	0.5483	0.5441	0.5950	0.6050
	168	**0.2151**	**0.3316**	0.8859	0.7338	0.9683	0.7348	0.7701	0.6470	0.8079	0.6997
	196	**0.2123**	**0.3291**	0.9720	0.7703	1.1692	0.8084	0.9495	0.7204	0.9534	0.7591
ILI	24	**2.7474**	1.0740	2.8332	**1.0656**	3.5111	1.1882	3.3729	1.2011	2.9399	1.1014
	36	**2.7124**	**1.0348**	3.1439	1.1197	3.7813	1.2588	4.0722	1.3292	3.4974	1.2212
	48	**2.6138**	**1.0399**	3.4153	1.1725	4.1892	1.3319	4.1239	1.3239	3.7872	1.2713
	60	**2.2889**	**0.8940**	3.7917	1.2553	4.2588	1.3352	3.9937	1.3063	3.8137	1.2695
	72	**2.0820**	**0.8372**	4.0823	1.3232	4.1868	1.3431	4.0423	1.3294	3.8818	1.3023
	96	**2.4419**	**1.0287**	4.2442	1.3755	4.3677	1.3756	4.2162	1.3594	4.2148	1.3530

4.2 Time Series Forecasting

To simulate different forecasting scenarios, we evaluate models under different future horizons, covering short-term and long-term forecasting cases. Tables 1 and 2 summarize the multivariate time series forecasting evaluation results of four datasets. Ti-MAE consistently improves the performance in across all datasets of different prediction horizons. Specifically, Ti-MAE achieves a MAE decrease of 15.7% in ETT, 30.3% in Weather, 19.0% in Exchange and 14.4% in ILI. Ti-MAE shows more compatible performance compared to other self-supervised and supervised methods. Notably, we have trained **only one** Ti-MAE model for different prediction horizons while all the end-to-end supervised models should be trained separately for different prediction length settings, which proves the huge scalability and efficiency of our proposed framework. To further explore the impact of main properties of Ti-MAE, we conduct extensive ablation experiments on Weather under input sequence length of 200 and prediction horizon of 100 setting for evaluation. Table 3 demonstrates all the results of ablation study.

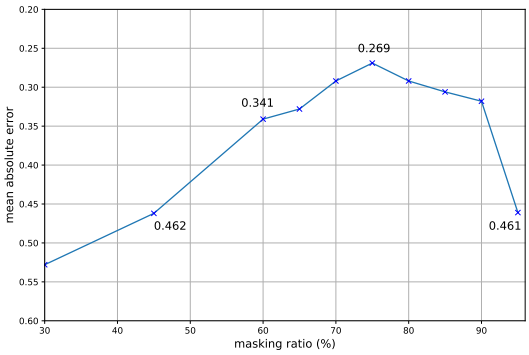

Figure 5: The optimal masking ratio is around 75%. Lower or higher masking ratio will degrade the performance of prediction.

Masking ratio. Figure 5 and Table 3a show the influence of the masking ratio. The optimal ratios are around 75%, which is in contrast to BERT (Devlin et al., 2019) and video-MAE (Feichtenhofer et al., 2022) but similar to MAE for images (He et al., 2021). The high masking ratio induces the model to process fewer tokens and learn high-level semantic information. We can see that lower masking ratios perform worse even if the encoder could see more tokens because the model trained

Table 2: Multivariate time series forecasting results compared to end-to-end methods.

Method		Ti-MAE		Autoformer		Informer		LogTrans		TCN	
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh	12	**0.2629**	**0.3462**	0.5016	0.5204	0.4299	0.4644	0.5191	0.5128	0.4633	0.4969
	24	**0.3520**	**0.3924**	0.5063	0.5309	0.4880	0.4963	0.6861	0.6039	0.5831	0.5467
	48	**0.3977**	**0.4173**	0.5703	0.5563	0.6625	0.5774	0.7658	0.7572	0.6698	0.6063
	96	**0.4266**	**0.4301**	0.6052	0.5663	0.9584	0.7157	0.9632	0.7113	0.7213	0.6794
	128	**0.4493**	**0.4436**	0.6043	0.5726	0.9504	0.7197	0.9991	0.7945	0.7974	0.7137
	168	**0.5091**	**0.4594**	0.7382	0.6199	1.1043	0.7867	1.0017	0.8464	0.8109	0.6804
Weather	12	**0.0932**	**0.1460**	0.2042	0.2960	0.2351	0.3128	0.3354	0.3767	0.3187	0.3762
	24	**0.1226**	**0.1809**	0.2200	0.3141	0.1244	0.2022	0.3054	0.3566	0.3210	0.3671
	48	**0.1633**	**0.2280**	0.2691	0.3542	0.2352	0.3129	0.3632	0.3951	0.3862	0.4228
	96	**0.2123**	**0.2735**	0.2891	0.3673	0.2808	0.3586	0.3758	0.3889	0.3989	0.4341
	128	**0.2197**	**0.2805**	0.2758	0.3469	0.3055	0.3723	0.3956	0.4415	0.4374	0.4856
	168	**0.2460**	**0.3049**	0.2861	0.3506	0.3473	0.4003	0.4227	0.5013	0.4911	0.5018
Exchange	24	**0.0697**	**0.1889**	0.0894	0.2239	0.4963	0.5623	0.5121	0.6039	1.9397	0.8975
	48	**0.1255**	**0.2448**	0.1474	0.2881	1.0477	0.8169	1.0395	0.8122	2.9541	1.3356
	96	**0.1912**	**0.2972**	0.2883	0.3957	1.1038	0.8215	1.0687	0.8312	2.9913	1.3432
	128	**0.2208**	**0.3242**	0.3102	0.4107	1.1978	0.8535	1.1264	0.8436	3.0010	1.3945
	168	**0.2151**	**0.3316**	0.3066	0.4108	1.1564	0.8444	1.0959	0.8377	3.0138	1.4136
	196	**0.2199**	**0.3393**	0.2990	0.4021	1.1679	0.8545	1.1410	0.8615	3.0287	1.4248
ILI	24	**2.7474**	**1.0740**	3.3292	1.2088	4.2526	1.3551	4.1128	1.3142	6.6241	1.8302
	36	**2.7124**	**1.0348**	3.4076	1.1688	4.7647	1.4433	4.2377	1.3967	6.8578	1.8779
	48	**2.6138**	**1.0399**	3.2077	1.1125	4.8189	1.4553	4.8543	1.4653	6.9683	1.8920
	60	**2.2889**	**0.8940**	3.3373	1.1659	4.7974	1.4669	4.8785	1.4766	7.1269	1.9182
	72	**2.0820**	**0.8372**	3.1079	1.1237	4.1188	1.3718	4.7576	1.4645	7.0043	1.9094
	96	**2.4419**	**1.0287**	3.0530	1.1260	4.5218	1.4401	4.7263	1.4572	7.1425	1.9322

with lower masking ratio may simply recover the values by interpolation or extrapolation, focusing on low level semantic features locally.

Sampling Time Tables 3b and 3c study the influence of sampling time in each iteration and data augmentation on Ti-MAE training stage. Ti-MAE works well with proper sampling time in each iteration and even no extra data augmentation, which is different from other existing representation learning methods on time series, especially contrastive learning models which rely on heavily data augmentation. Ti-MAE can directly learn adequate information from masked data. Additionally, introducing extra data augmentation will degrade the performance due to inevitable distortions of the original data, which is different from the result of MAE for images or videos. Random masking in each iteration generates a large number of different views without any distortion so that model can make use of visible tokens to capture more useful features.

Input sequence length. In Table 3d we compare different length of input time series in training stage. Surprisingly, although lengthening the input length of the pre-training stage can improve the performance within limits, too long input sequence may degrade the results of our model because there is a certain conflict between the complex periodic pattern in the long sequence and the short-term prediction task in the downstream.

Decoder Design. Tables 3e and 3f show the influence of the decoder width and depth. A shallow design of the decoder is sufficient for reconstruction tasks. It is because that time series data are not that complicated and thus need lower decoding dimension to reduce redundancy. Such a lightweight decoder can efficiently reduce computational complexity and memory usage.

4.3 Time Series Classification

In the previous section, we have improved the performance of our framework on forecasting tasks by reducing the consistency between upstream tasks and downstream tasks compared to contrastive learning methods. Thus, we should evaluate learning ability of instance-level representation on classification tasks. The results on 128 UCR archive are summarized in Table 4. Compared to other representation learning methods, Ti-MAE achieves more compatible performance. More details

Table 3: Ablation experiments on Weather. The entries marked in **bold** are the same which specify the default settings. Lower MSE and MAE represent better performance. This table format follows (Feichtenhofer et al., 2022).

(a) Masking ratio

Ratio	MSE	MAE
0.45	0.3082	0.3557
0.60	0.2650	0.3414
0.75	**0.2103**	**0.2696**
0.90	0.2483	0.3176

(b) Sampling time

#Sampling	MSE	MAE
20	0.2151	0.2892
25	0.2447	0.3095
30	**0.2103**	**0.2696**
35	0.2171	0.2769

(c) Data augmentation

Case	MSE	MAE
None	**0.2103**	**0.2696**
Scaling	0.2383	0.3022
Shifting	0.2399	0.3388
Jittering	0.2508	0.3316

(d) Input sequence length

Length	MSE	MAE
200	0.2413	0.2844
300	**0.2103**	**0.2696**
400	0.2877	0.3501
500	0.2328	0.2985

(e) Decoder width

#Blocks	MSE	MAE
1	0.2375	0.2954
2	**0.2103**	**0.2696**
3	0.2125	0.3112
4	0.2616	0.3057

(f) Decoder depth

Dim	MSE	MAE
16	0.4374	0.4805
32	**0.2103**	**0.2696**
64	0.2380	0.2813
128	0.3172	0.3758

and full results of each dataset in UCR are listed in the appendix. Following (Yue et al., 2022), Critical Difference diagram (Demsar, 2006) for Nemenyi tests conducted on all datasets is shown as Figure 6, where classifiers that are connected by a bold line do not have a significant difference. This proves that Ti-MAE could learn good instance-level representations directly from the raw time series data without any hierarchical tricks or data augmentation.

Table 4: 128 UCR Archive Classification

method	Ti-MAE	TS2Vec	T-Loss	TS-TCC	TST	TNC	DTW
Avg.Acc.	**0.8231**	0.8201	0.7875	0.7396	0.6385	0.7431	0.7278
Avg.Rank	**2.054**	3.016	4.016	4.445	4.883	3.875	5.711

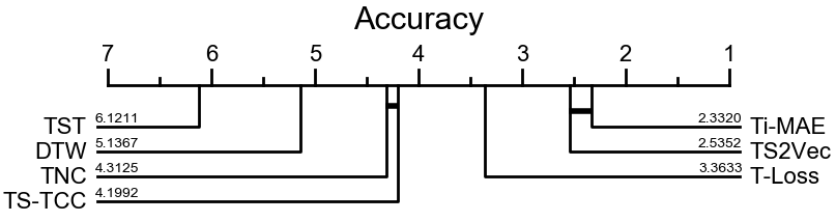

Figure 6: Critical Difference (CD) diagram on UCR classification with a 95% confidence level.

5 Conclusion

This paper proposes a novel self-supervised framework named Ti-MAE for time series representation learning, which randomly masks out tokenized time series and learns an autoencoder to reconstruct them at the point-level. Ti-MAE bridges the connection between contrastive representation learning and generative Transformer-based methods and greatly improves the performance on forecasting tasks due to reducing the inconsistency of upstream and downstream tasks compared to contrastive learning methods. Compared with the fixed continuous masking strategy used in existing Transformer-based models, Ti-MAE adequately leverages all the input sequence and alleviates the distribution shift problem. The flexible setting of masking ratio makes Ti-MAE more adaptive to various prediction scenarios with different time steps. The experiments on real-world datasets and ablation study demonstrate the effectiveness and scalability of our framework. Future work will extend our work for different reconstruction targets according to their requirements.

REFERENCES

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *ArXiv*, abs/1803.01271, 2018.

Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *ArXiv*, abs/2106.08254, 2021.

Réal André Carbonneau, Kevin Laframboise, and Rustam M. Vahidov. Application of machine learning techniques for supply chain demand forecasting. *Eur. J. Oper. Res.*, 184:1140–1154, 2008.

Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9620–9629, 2021.

Yanping Chen, Bing Hu, Eamonn J. Keogh, and Gustavo E. A. P. A. Batista. Dtw-d: time series semi-supervised learning from a single example. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013.

Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *ArXiv*, abs/1904.10509, 2019.

Hoang Anh Dau, A. Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ratanamahatana, and Eamonn J. Keogh. The ucr time series archive. *IEEE/CAA Journal of Automatica Sinica*, 6:1293–1305, 2019.

Janez Demsar. Statistical comparisons of classifiers over multiple data sets. *J. Mach. Learn. Res.*, 7:1–30, 2006.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805, 2019.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2021.

Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, C. Kwoh, Xiaoli Li, and Cuntai Guan. Time-series representation learning via temporal and contextual contrasting. *ArXiv*, abs/2106.14112, 2021.

Meng fang Zha. Time series generation with masked autoencoder. *ArXiv*, abs/2201.07006, 2022.

Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *ArXiv*, abs/2205.09113, 2022.

Jean-Yves Franceschi, Aymeric Dieuleveut, and Martin Jaggi. Unsupervised scalable representation learning for multivariate time series. In *NeurIPS*, 2019.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Doll'ar, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *ArXiv*, abs/2111.06377, 2021.

Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, Chun-Wei Wang, and Jie Tang. Graphmae: Self-supervised masked graph autoencoders. *ArXiv*, abs/2205.10803, 2022.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.

Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018.

Nikolay Pavlovich Laptev, Jason Yosinski, Li Erran Li, and Slawek Smyl. Time-series extreme event forecasting with neural networks at uber. 2017.

Qi Lei, Jinfeng Yi, Roman Vaculín, Lingfei Wu, and Inderjit S. Dhillon. Similarity preserving representation learning for time series clustering. In *IJCAI*, 2019.

SHIYANG LI, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *ArXiv*, abs/1907.00235, 2019.

Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X. Liu, and Schahram Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=0EXmFzUn5I.

Jiawei Ma, Zheng Shou, Alireza Zareian, Hassan Mansour, Anthony Vetro, and Shih-Fu Chang. Cdsa: Cross-dimensional self-attention for multivariate, geo-tagged time series imputation. *ArXiv*, abs/1905.09904, 2019.

Boris N. Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. N-beats: Neural basis expansion analysis for interpretable time series forecasting. *ArXiv*, abs/1905.10437, 2020.

Yatian Pang, Wenxiao Wang, Francis E. H. Tay, W. Liu, Yonghong Tian, and Liuliang Yuan. Masked autoencoders for point cloud self-supervised learning. *ArXiv*, abs/2203.06604, 2022.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.

Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2536–2544, 2016.

Jinwen Qiu, Sreenivasa Rao Jammalamadaka, and Ning Ning. Multivariate bayesian structural time series model. *J. Mach. Learn. Res.*, 19:68:1–68:33, 2018.

Zezhi Shao, Zhao Zhang, Fei Wang, and Yongjun Xu. Pre-training enhanced spatial-temporal graph neural network for multivariate time series forecasting. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, pp. 1567–1577, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393850. doi: 10.1145/3534678.3539396. URL https://doi.org/10.1145/3534678.3539396.

Huan-Zhi Song, Deepta Rajan, Jayaraman J. Thiagarajan, and Andreas Spanias. Attend and diagnose: Clinical time series analysis using attention models. In *AAAI*, 2018.

Sana Tonekaboni, Danny Eytan, and Anna Goldenberg. Unsupervised representation learning for time series with temporal neighborhood coding. *ArXiv*, abs/2106.00750, 2021.

Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *ArXiv*, abs/2203.12602, 2022.

Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *ArXiv*, abs/1807.03748, 2018.

Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *ArXiv*, abs/1706.03762, 2017.

Pascal Vincent, H. Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML '08*, 2008.

Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven Hoi. CoST: Contrastive learning of disentangled seasonal-trend representations for time series forecasting. In *International Conference on Learning Representations*, 2022a. URL https://openreview.net/forum?id=PilZY3omXV2.

Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven C. H. Hoi. Etsformer: Exponential smoothing transformers for time-series forecasting. *ArXiv*, abs/2202.01381, 2022b.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *NeurIPS*, 2021.

Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yu Tong, and Bixiong Xu. Ts2vec: Towards universal representation of time series. In *AAAI*, 2022.

George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. A transformer-based framework for multivariate time series representation learning. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.

Richard Zhang, Phillip Isola, and Alexei A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 645–654, 2017.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wan Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *AAAI*, 2021.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *ICML*, 2022.