

ST-ReP: Learning Predictive Representations Efficiently for Spatial-Temporal Forecasting

Qi Zheng, Zihao Yao, Yaying Zhang*

The Key Laboratory of Embedded System and Service Computing, Ministry of Education,
Tongji University, Shanghai 200092, China
{zhengqi97, yaozihao, yaying.zhang}@tongji.edu.cn

Abstract

Spatial-temporal forecasting is crucial and widely applicable in various domains such as traffic, energy, and climate. Benefiting from the abundance of unlabeled spatial-temporal data, self-supervised methods are increasingly adapted to learn spatial-temporal representations. However, it encounters three key challenges: 1) the difficulty in selecting reliable negative pairs due to the homogeneity of variables, hindering contrastive learning methods; 2) overlooking spatial correlations across variables over time; 3) limitations of efficiency and scalability in existing self-supervised learning methods. To tackle these, we propose a lightweight representation-learning model ST-ReP, integrating current value reconstruction and future value prediction into the pre-training framework for spatial-temporal forecasting. And we design a new spatial-temporal encoder to model fine-grained relationships. Moreover, multi-time scale analysis is incorporated into the self-supervised loss to enhance predictive capability. Experimental results across diverse domains demonstrate that the proposed model surpasses pre-training-based baselines, showcasing its ability to learn compact and semantically enriched representations while exhibiting superior scalability.

Code — <https://github.com/zhuoshu/ST-ReP>

Extended version — <https://arxiv.org/abs/2412.14537>

Introduction

Spatial-temporal (ST) forecasting, a crucial task within multivariate time series (MTS) analysis, plays an essential role in understanding the operational dynamics of urban and geospatial systems. With the rapid expansion of data monitoring technologies, there is an unprecedented increase in unlabeled spatial-temporal series (STS) data, prompting growing interest in self-supervised learning approaches. These approaches exploit large-scale data to generate pseudo-labels for training models to effectively represent raw data. Although self-supervised learning has shown promise in MTS modeling through contrastive and reconstruction-based methods, its application to STS data is challenged by variable homogeneity and the inherent complexity of spatial-temporal attributes:

*Corresponding author.
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

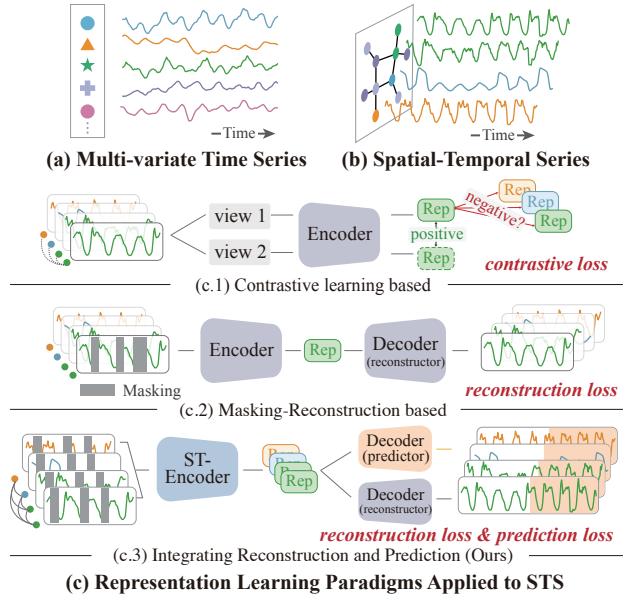


Figure 1: Top: An illustration of typical multivariate series and spatial-temporal series. Different shapes denote types of variables. STS usually has homogeneous ones. Bottom: Comparison of three examples of representation learning paradigms. Our proposed model belongs to the third one.

1) Crafting suitable positive and negative pairs for contrastive methods in the context of STS is challenging due to the variable homogeneity. As depicted in Figure 1(a)(b), in MTS, variables often represent different attributes (such as oil temperature, external load, and location features of a power transformer). When comparing sequences from distinct attributes, it is possible to learn features that exhibit clear differences. In contrast, STS variables, typically from similar sources like traffic network's multiple road sensors, possess homogeneity when measuring the same data type, making variable distinction challenging. When contrastive methods treat representations of different variables within the same batch as negative samples, as shown in Figure 1(c.1), there is a significant risk of generating false negatives. This occurs because these variables, despite being considered negatives, may actually be closely related or exhibit

similar patterns in the spatial-temporal context, thereby hindering the accurate extraction of valuable information.

2) Existing self-supervised methods often overlook spatial correlations between variables. As shown in Figure 1(c.2), many reconstruction-based approaches aim to enhance model understanding and representation of MTS by reconstructing input data. These methods generally focus on individual variables and extract features predominantly from the temporal dimension, with limited consideration of inter-variable relationships. However, an important characteristic of STS data is the spatial correlation between variables, which has been frequently studied in recent years. Ignoring this correlation can lead to insufficient spatial-temporal extraction.

3) The efficiency and scalability of self-supervised methods demand augmentation. As illustrated in Figure 1(c.3), unlike MTS models that treat different variable sequences as distinct samples, STS models primarily consider all variable sequences within a given time period as a single sample to account for their potential spatial correlations. However, this approach can increase the computational burden during training, particularly as the number of nodes grows. In most high-complexity algorithms, the computational cost escalates significantly. Thus, efficient and effective self-supervised models for practical spatial-temporal forecasting needs remain an area for further research.

To address these challenges, we propose a reconstruction-based self-supervised model *ST-ReP* to learn representations for Spatial-Temporal forecasting. It integrates Reconstruction with Prediction, explicitly considering spatial correlations between variables during the encoding process, as illustrated in Figure 1(c.3). Specifically, we extend the reconstruction objective to encompass not only the current series but also unseen future values, thereby encouraging the model to learn representations that have predictive capabilities. Additionally, we propose a new ST encoder, to model the spatial-temporal relationships between variables with compression and decompression. It first reduces the temporal dimension to a compact latent space, then employs a linear encoder to compute spatial correlations between variables, and finally restores the original temporal length. This approach aims to reduce redundancy in extracting spatial-temporal features while also lowering computational complexity.

Moreover, we introduce a loss function to capture the multi-scale temporal information in the self-supervised training process, which links the reconstructed values and the predicted ones. This method constrains the model to understand temporal correlations from multiple perspectives without increasing the computational cost during encoding. With these designs on pre-training framework and the encoder mechanism, our proposed model ST-ReP is capable of learning semantically dense and predictive ST representations efficiently. Our contribution are summarized as follows:

- To capture unique spatial interactions inherent in spatial-temporal data, we propose a reconstruction-based representation learning method called ST-ReP. By combining the reconstruction of current series with the prediction

of future values, ST-ReP is encouraged to learn spatial-temporal correlations and generate predictive representations.

- We propose a spatial-temporal encoding scheme called *Compression-Extraction-Decompression (C-E-D)*. It involves *Compression-Decompression* in the temporal dimension to transform features into a compact latent space, with spatial feature *Extraction* in between. Additionally, we utilize a linear encoder in the spatial dimension to capture correlations between variables, thereby reducing potential redundancy.
- Experiments are conducted on six spatial-temporal datasets from various domains. The results demonstrate that our model achieves superior downstream prediction accuracy compared to advanced self-supervised learning baselines. Furthermore, our model showcases lower footprints and learns representations with lower dimensions, highlighting its good scalability.

Related Work

In this section, we will review self-supervised representation learning methods and their adaptation in STS data. Due to strong generalization and adaptability, the self-supervised learning method has garnered significant attention in MTS modeling (Zhang et al. 2024). Effective self-supervised learning relies on a suitable encoder, clear pseudo-labeling principles, and robust loss constraints, and typically falls into two main categories: contrastive methods and generative methods.

Contrastive-based methods Contrastive methods perform pseudo-classification tasks in the representation space. They primarily rely on data augmentation and the design of positive and negative pairs to make the learned representations as distinguishable as possible. Numerous studies for MTS adopt this paradigm (Yue et al. 2022; Woo et al. 2022; Fraikin, Bennetot, and Allassonniere 2024; Liu and Chen 2024). They often select representations of other variables within the same batch as negative examples. However, due to the homogeneity of variables in STS data, this selection has inherent limitations and may provide misleading information. Then ST-SSL (Ji et al. 2023) introduces self-supervised learning into spatial-temporal forecasting tasks. It augments ST graphs for capturing spatial-temporal heterogeneity under the contrastive setting. In contrast, our proposed ST-ReP does not adopt a contrastive strategy. It avoids multi-view augmentations that increase the encoding burden and sidesteps the potential inductive biases arising from unclear selection of positive and negative pairs.

Reconstruction-based methods Generative methods primarily include reconstruction-based approaches. These methods perform regression tasks in the original input space, typically using an auto-encoder framework to restore original data from masked visible inputs. This approach forces the model to understand underlying patterns in the data and enhances its robustness. Many methods adopt similar strategies for representation learning (Zerveas et al. 2021; Nie et al. 2023; Dong et al. 2023). These methods primarily

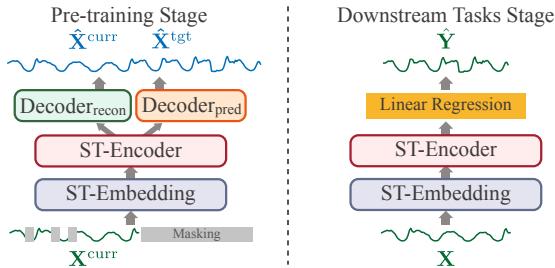


Figure 2: The overall workflow of ST-ReP.

model individual time series without considering spatial correlations between variables. STEP (Shao et al. 2022) and GPT-ST (Li et al. 2023) are two recent studies that introduce self-supervised learning into STS forecasting. However, their focus is primarily on using self-supervised learning as an auxiliary tool to enhance the performance of complex downstream ST models. They do not specifically evaluate the spatial-temporal representations learned by the pre-training models.

Moreover, most methods use highly complex encoding structures or high-dimensional representations under multi-view augmentation, leading to high computational overhead, which becomes a bottleneck when dealing with large-scale datasets with more than thousands of variables. By comparison, our ST-ReP introduces completely unseen future values as partly reconstruction targets to enhance the predictive performance of the representations. It captures correlations between variables while maintaining high encoding efficiency and scalability.

Preliminary

A set of STS data with N variables can be formulated as a three-dimensional tensor $\mathbf{X} \in \mathbb{R}^{N \times T \times C}$ in the previous T time intervals, and C is the number of features of each variable. In this work, rather than focusing on end-to-end prediction, we aim to learn ST representations that can extract information beneficial for downstream prediction tasks. Formally, we need to learn a function f that maps the raw data to a d -dimensional representation space: $\{r_1, \dots, r_i, \dots, r_N\} = f(\{X_1, \dots, X_i, \dots, X_N\})$. Here, $r_i \in \mathbb{R}^{T \times d}$ represents the representation of variable i , with $r_{i,t_j} \in \mathbb{R}^d$ being the representation vector at time step t_j . Once representations of past data are obtained, they can be used for various downstream tasks (such as forecasting and imputation). In this paper, we focus on ST forecasting, which involves using data from the past T time steps to predict data for the future T' time steps.

The primary goal of self-supervised learning is to benefit from pre-training on abundant data, enabling the model to understand the data and convert raw inputs into valuable representations. To assess the quality of these learned representations, we employ simple and lightweight downstream models (e.g., linear regression) instead of relying on complex models (e.g., deep learning methods), thereby demonstrating the predictive performance in the most basic scenarios. The overall workflow of our method is shown in Figure 2. This approach also showcases the adaptability of self-

supervised learning in resource-constrained environments.

Methodology

The core idea of our proposed ST-ReP is to use current ST series data $\mathbf{X}^{\text{curr}} \in \mathbb{R}^{N \times T \times C}$ by masking and reconstructing it, while simultaneously predicting future series values $\mathbf{X}^{\text{tgt}} \in \mathbb{R}^{N \times F \times C}$. F is the predicted length of the targeted value in the pretext task. Figure 3 provides an overview of our model’s architecture, which follows a self-supervised pretraining framework. Similar to many reconstruction-based methods (Shao et al. 2022), we initially applies random masking to the current values \mathbf{X}^{curr} along the temporal dimension (masking is only used during training), obtaining the masked input $\tilde{\mathbf{X}}^{\text{curr}}$. Then we utilize the embedding module and the encoder to generate data representations. Differently, ST-ReP prioritizes a lightweight encoder structure to minimize redundancy during feature extraction. Following encoding, ST-ReP employs a reconstructor ($\text{Decoder}_{\text{recon}}$) to restore the representations to their original values, and an independent predictor ($\text{Decoder}_{\text{pred}}$) to project them onto future values. That means \mathbf{X}^{tgt} is only used to participate in the loss calculation and is totally unseen for encoder, which can be viewed as a complete mask for \mathbf{X}^{tgt} . Lastly, the model integrates the reconstructed current values and the predicted future values, performing multi-scale temporal comparison over long sequences, thereby compelling the model to learn temporal relationships among series elements. In the following subsections, we elaborate three key components in ST-ReP: the ST-Embedding module, the linear ST encoder with the $C-E-D$ structure, and the reconstruction and prediction integrated loss setting.

ST-Embedding Module

To capture both the temporal progression and inherent spatial characteristics of the data, we integrate these aspects using embedding techniques. Firstly, a projector maps the feature dimensions of the raw data into a hidden space:

$$\tilde{\mathbf{E}}^{\text{hidden}} = \Phi_1(\tilde{\mathbf{X}}^{\text{curr}}), \quad (1)$$

where Φ_1 is the projector, it contains two linear layers with a ReLU activation function in between. Following a recent study’s proposed cross-time normalization strategy (Zheng and Zhang 2023), we concatenate data at the most recent time step with the raw data as input to the projector. Next, we use learnable, randomly initialized mask tokens to fill in the hidden data to the original length and obtain $\tilde{\mathbf{E}}^{\text{hidden}} \in \mathbb{R}^{N \times T \times d}$. We add three commonly-used embedding information: E^{tod} , E^{dow} , and E^{spt} . $E^{\text{tod}} \in \mathbb{R}^{T \times d}$, represents the time of day corresponding to each time interval (with an upper limit of T), $E^{\text{dow}} \in \mathbb{R}^{T \times d}$ represents the day of the week (with an upper limit of 7 days), and both are queried from learnable parameters of length T and 7, respectively. Additionally, E^{spt} represents the variable’s index information. These embeddings are then added through broadcasting. Note that even in the masked parts, the temporal and spatial embeddings remain, reflecting the real-world scenario of data missing. After incorporating static temporal

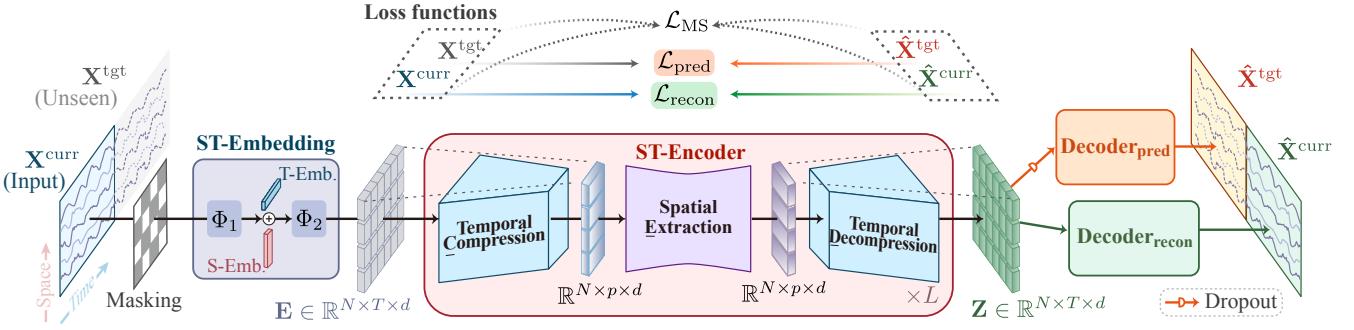


Figure 3: The overall pre-training framework of our proposed ST-ReP. The ST-Encoder utilizes a Compression-Extraction-Decompression structure to enhance modeling efficiency. ST-ReP integrates reconstruction and prediction, employing three types of loss functions to supervise representation learning: reconstruction loss $\mathcal{L}_{\text{recon}}$, prediction loss $\mathcal{L}_{\text{pred}}$, and multi-scale loss \mathcal{L}_{MS} .

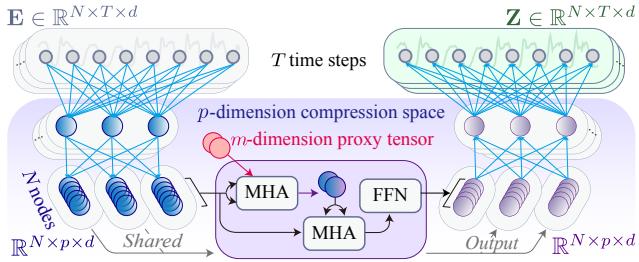


Figure 4: An example of the ST-Encoder. MLPs are used as the compressor and decompressor, with a linear spatial extractor based on proxy tensor.

and spatial identification information, we use another projector to enhance the semantic density of the series embedding:

$$\mathbf{E} = \Phi_2(\tilde{\mathbf{E}}^{\text{hidden}} + \mathbf{E}^{\text{tod}} + \mathbf{E}^{\text{dow}} + \mathbf{E}^{\text{spt}}). \quad (2)$$

Here, a one-layer equal-width convolution along with the time dimension is employed as projector Φ_2 . This results in an embedding $\mathbf{E} \in \mathbb{R}^{N \times T \times d}$ that combines spatial-temporal context.

Compression-Extraction-Decompression Structure for ST Encoding

The encoding is one of pivotal parts in ST representation learning, and various ST mining models can serve as the backbone for it. In this work, we focus on capturing cross-time spatial relationships among variables while emphasizing the use of linear complexity models for ST modeling. This approach aims to enhance encoding efficiency and scalability, enabling the model to handle STS with a large number of variables. Achieving this balance between reduced complexity and effective modeling is a challenging task, which has been increasingly studied in related research (Zhou et al. 2021; Liu et al. 2022; Zhang and Yan 2023; Choromanski et al. 2021; Wu et al. 2023b).

To meet the objective, we propose a sandwich structure called *Compression-Extraction-Decompression (C-E-D)*, as

shown in Figure 3. This structure compresses hidden features along the temporal dimension into a low-dimensional latent space, performs low-complexity spatial relationship extraction, and then restores the ST features to their original temporal length. Figure 4 presents an example that we used in this paper. We elaborate the details of each part as follows.

Temporal Compression and Decompression In this paper, we employ two fully connected layers with a non-linear activation function GELU (Hendrycks and Gimpel 2016) in between, serving as the compressor and decompressor along the temporal dimension.

$$\mathbf{E}^{\text{comp}} = \text{Comp}_{\text{tp}}(\mathbf{E}) \in \mathbb{R}^{N \times p \times d}, \quad (3)$$

$$\mathbf{H} = \text{Extraction}_{\text{spt}}(\mathbf{E}^{\text{comp}}) \in \mathbb{R}^{N \times p \times d}, \quad (4)$$

$$\mathbf{E}^{\text{decomp}} = \text{Decomp}_{\text{tp}}(\mathbf{H}) \in \mathbb{R}^{N \times T \times d}. \quad (5)$$

The compressor maps the input embedding to a p -dimensional latent space, p is a small constant value. The goal here is not to extract temporal correlations within the sequence but to summarize the current time series into a few key hidden states containing high-level semantics. This approach avoids potential redundancy and helps to reduce model overhead. Moreover, the decompression operation following spatial feature extraction (will be described in the next paragraph) is essential. It restores the low-dimensional hidden states to their original temporal length, providing an additional perspective to constrain the temporal semantics of the representation. This can be viewed as a form of temporal autoencoding. Additionally, this setup allows for the stacking of multiple layers, enhancing the model's expressive capacity.

Spatial Extraction We use the spatial node representations at a single time step as input tokens to compute attention among them, with the same encoding model shared across multiple time steps. Additionally, we employ a linear transformer structure to extract spatial correlations between variables, an area that has garnered increasing research interest (Zhou et al. 2021; Liu et al. 2022; Zhang and Yan 2023;

Choromanski et al. 2021; Wu et al. 2023b). Specifically, we reduce the number of input tokens for attention calculation, similar to the router mechanism in (Zhang and Yan 2023) and the bottleneck setup in (Guo et al. 2023). Our approach can be formulated as follows:

$$\begin{aligned} H^p &= \text{MHA}_1(P_w, E^{comp}, E^{comp}) \in \mathbb{R}^{m \times d}, \\ H' &= \text{MHA}_2(E^{comp}, H^p, H^p) + E^{comp} \in \mathbb{R}^{N \times d}, \quad (6) \\ H &= \text{FFN}(H') + H' \in \mathbb{R}^{N \times d}. \end{aligned}$$

For brevity, the subscript denoting the time dimension is omitted here, with unbolded symbols representing the hidden features at a single time step. Here, MHA (Multi-Head Attention) and FFN (Feed-Forward Network) are two fundamental components of the Transformer structure. The former takes three inputs as query, key, and value. $P_w \in \mathbb{R}^{m \times d}$ is learnable parameters of constant size m . It acts as a proxy tensor in the attention computation, reflecting high-level semantics or summary information of the token sequence. This allows for an alternative interaction with the input tokens. Specifically, in MHA_1 , P_w serves as the query, while E^{comp} serves as the key and value. The result H^p is then used in MHA_2 as the key and value to restore the data features to the original token count. This setup significantly reduces computational overhead while maintaining expressive power, showing a linear computational complexity.

When capturing spatial correlations, each token already encapsulates the desired high-level temporal semantics. Consequently, the attention computation at this stage can effectively represent the extent of spatial interactions across time. After concatenating the spatial modeling results across p virtual time steps, we obtain the final results of spatial extraction $\mathbf{H} \in \mathbb{R}^{N \times p \times d}$ for the succeeding decompression. Equation 3 to Equation 5 constitute one layer of the ST Encoder. Without loss of generality, the ST Encoder can be stacked in L layers. We add the output of the final layer to the initial input \mathbf{E} to obtain the final ST representation $\mathbf{Z} \in \mathbb{R}^{N \times T \times d}$.

Reconstruction and Prediction Integrated Loss Constraint

After obtaining the representations, it is crucial to ensure they contain valuable semantic information. A key idea of ST-ReP is to combine reconstruction and prediction from generative methods. We use the learned representations \mathbf{Z} from the masked input $\tilde{\mathbf{X}}^{\text{curr}}$ to both reconstruct the current series \mathbf{X}^{curr} and predict future values \mathbf{X}^{tgt} :

$$\hat{\mathbf{X}}^{\text{curr}} = \text{Decoder}_{\text{recon}}(\mathbf{Z}) \in \mathbb{R}^{N \times T \times C}, \quad (7)$$

$$\hat{\mathbf{X}}^{\text{tgt}} = \text{Decoder}_{\text{pred}}(\text{Drop}(\mathbf{Z})) \in \mathbb{R}^{N \times F \times C}. \quad (8)$$

Specifically, we use two linear layers with a GELU activation function in between as the $\text{Decoder}_{\text{recon}}$, with a residual connection before and after the first linear layer. For the $\text{Decoder}_{\text{pred}}$, we use two linear layers. Drop represents the dropout operation, which is used to mitigate overfitting. These layers map the representations to the target dimensions in both the time and feature dimensions. Both the re-

Dataset	#Nodes	Time Steps	Time Interval	CV*
PEMS04	307	16992	5min	58.82
PEMS08	170	17856	5min	46.75
CA	8600	35040	15min	60.10
SDWPF	134	35280	10min	121.97
Humidity	2048	8784	1hour	17.19
Temperature	2048	8784	1hour	2.19

* CV denotes the coefficient of variation.

Table 1: Summary of datasets.

construction and prediction tasks involve calculating the regression loss against the ground truth values:

$$\mathcal{L}_{\text{recon}} = g(\hat{\mathbf{X}}^{\text{curr}}, \mathbf{X}^{\text{curr}}), \mathcal{L}_{\text{pred}} = g(\hat{\mathbf{X}}^{\text{tgt}}, \mathbf{X}^{\text{tgt}}). \quad (9)$$

where g denotes the distance metric function, we use Huber loss (Huber 1992) in this paper.

Furthermore, to encourage the model to acquire representations that retain as much temporal information as possible, we concatenate the predicted future values $\hat{\mathbf{X}}^{\text{tgt}}$ with the reconstructed current values $\hat{\mathbf{X}}^{\text{curr}}$, denoted as $\hat{\mathbf{X}}^{\text{full}}$. Subsequently, we compare this concatenated sequence $\hat{\mathbf{X}}^{\text{full}}$, with a length of $T + F$, to its corresponding ground truth values \mathbf{X}^{full} across multiple time scales. Specifically, we utilize a set of kernels Ω with different size to conduct average pooling on both sequences separately, calculate the information loss at each scale, and then aggregate the results:

$$\mathcal{L}_{\text{MS}} = \sum_{k \in \Omega} g(\text{AvgPool}_k(\hat{\mathbf{X}}^{\text{full}}), \text{AvgPool}_k(\mathbf{X}^{\text{full}})). \quad (10)$$

This allows the model to capture more temporal granularity information and embed it within the learned ST representation. In contrast to many methods (Wu et al. 2023a; Challu et al. 2023; Chen et al. 2023; Shabani et al. 2023; Wang et al. 2024) that cascade to extract multi-scale information during encoding, we introduce multi-scale analysis solely into the loss computation, thus aiding in reducing model overhead.

Finally, the total loss is a linear combination of these three components:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{recon}} + \beta \mathcal{L}_{\text{pred}} + \gamma \mathcal{L}_{\text{MS}}, \quad (11)$$

where α , β , and $\gamma = 1 - \alpha - \beta$ are weights of three parts.

Experiments

We conduct experiments on datasets from diverse domains to evaluate the performance of our proposed ST-ReP.

Experiment Setup

Datasets Six datasets from three different domains are used for the experiments. (1) Transportation domain: PEMS04, PEMS08 (Song et al. 2020a,b), and CA (Liu et al. 2023). (2) Climate domain: Temperature, Humidity (Rasp et al. 2020). (3) Energy domain: SDWPF (Zhou et al. 2022). Table 1 summarises all datasets. Additionally, we measure the average coefficient of variation (CV) (Brown 1998) for

methods		PEMS04		PEMS08		SDWPF		Temperature		Humidity		CA	
		MSE	MAE										
E2E	HL	0.095	0.200	0.070	0.173	0.416	0.378	0.013	0.062	0.350	0.384	0.226	0.300
	Ridge Reg.	0.077	0.181	<u>0.058</u>	<u>0.157</u>	0.174	0.285	0.007	0.050	0.253	0.349	0.171	0.271
	DLinear	0.081	0.188	0.062	0.164	0.170	0.277	0.013	0.067	0.261	0.361	0.171	0.271
	iTransformer	0.061	<u>0.155</u>	0.044	0.132	0.160	0.225	<u>0.006</u>	<u>0.046</u>	0.237	0.325	<u>0.131</u>	0.222
MTS	TS2Vec	0.088	0.206	0.072	0.194	0.199	0.314	0.036	0.146	0.294	0.395	OOM	OOM
	CoST	0.075	0.191	0.052	0.156	0.192	0.302	0.014	0.092	0.228	0.345	OOM	OOM
	PatchTST	0.069	0.172	0.053	0.150	0.168	0.272	0.007	0.050	0.247	0.346	OOM	OOM
	T-Rep	0.058	0.156	<u>0.042</u>	<u>0.131</u>	0.179	0.279	0.005	0.043	0.217	<u>0.326</u>	OOM	OOM
	TimesURL	0.090	0.206	0.063	0.175	0.195	0.300	OOM	OOM	OOM	OOM	OOM	OOM
STS	STEP*	<u>0.054</u>	0.159	0.047	0.156	0.171	0.279	OOM	OOM	OOM	OOM	OOM	OOM
	ST-SSL*	0.083	0.200	0.059	0.174	0.177	0.283	0.025	0.116	0.254	0.371	OOM	OOM
	GPT-ST*	0.066	0.168	0.052	0.154	0.170	0.275	0.012	0.068	0.239	0.343	OOM	OOM
	ST-ReP	0.044	<u>0.134</u>	<u>0.033</u>	<u>0.120</u>	<u>0.167</u>	<u>0.267</u>	0.005	<u>0.046</u>	0.226	0.335	0.071	<u>0.182</u>

* For self-supervised-based STS forecasting methods, we only use their pre-trained encoders to generate spatial-temporal representations, which are then evaluated for predictive performance using the same downstream model.

Table 2: Comparison of the downstream spatial-temporal forecasting accuracy on diverse datasets. The best results are shown in bold, while the second-best results are underlined. Italicized numbers represent the best results among the self-supervised representation learning methods. OOM indicates that the method encountered an out-of-memory issue during either the pre-training or the downstream task. *E2E* denotes end-to-end methods.

each dataset, which is the ratio of the standard deviation to the mean of each variable, reflecting the variability of the data across all time points. Datasets with high CV values are typically more challenging to predict, whereas the performance gains on datasets with extremely low CV values tend to be minimal or limited.

Baselines and Implementation We choose twelve methods as baselines for comparison with ST-ReP: (1) Two naive methods: **History Last (HL)** (Guo et al. 2019), which uses the last observation as all prediction. And the **Ridge Regression** (McDonald 2009) model providing simple yet effective baseline. (2) Two end-to-end baselines: **DLinear** (Zeng et al. 2023), **iTransformer** (Liu et al. 2024). (3) Five advanced self-supervised representation learning models for MTS: **TS2Vec** (Yue et al. 2022), **CoST** (Woo et al. 2022), **T-Rep** (Fraikin, Bennetot, and Allassonniere 2024), **TimesURL** (Liu and Chen 2024) and **PatchTST** (Nie et al. 2023). (4) Three self-supervised based ST forecasting model: **ST-SSL** (Ji et al. 2023), **STEP** (Shao et al. 2022) and **GPT-ST** (Li et al. 2023). All experiments are conducted on a Linux server with one Intel(R) Xeon(R) Gold 5220 CPU @ 2.20 GHz and one 32GB NVIDIA Tesla V100-SXM2 GPU card. More details about the implementation of baselines and the proposed ST-ReP are in the Appendix.

Evaluation Setup Following the setup used in previous work (Yue et al. 2022; Liu and Chen 2024), we generate representations from T historical observations using the pre-trained model and predict the final T' targets based on the representation $r_{:,t}$ of the most recent time step t . We first train the representation learning model using the training set, and then use the learned representations and corresponding targets with various prediction length to train linear ridge regression models. The regularization weight for the ridge regression model is selected based on performance on the

validation set. Finally, we evaluate the model on the test set using two common-used metrics, MSE and MAE, to measure the predictive performance. We use normalized metrics to handle the large differences in raw value scales across domains and facilitate unified presentation of results. The normalization is applied to all variables collectively, rather than individually.

Note, ridge regression does not require GPU devices or backpropagation algorithms. For training samples already converted into spatial-temporal representations, we perform sampling and use only a small fraction of the representation samples (0.93% to 5.5%) as training data for downstream tasks (keep consistent with previous methods). This approach significantly reduces computational overhead, accelerates processing, and simulates resource-constrained scenarios. By contrast, for end-to-end deep-learning baselines, a specific model needs to be trained for each prediction horizon, using the entire training samples.

Prediction Performance Comparison

Table 2 presents the results for a prediction horizon of 12. Results for additional prediction horizons can be found in the Appendix. Our model outperforms all self-supervised representation learning methods in most cases and even exceeds the performance of end-to-end models on certain datasets. Specifically, on two PEMS datasets from the transportation domain, ST-ReP achieves an average improvement of 19.97%/11.25% in terms of MSE/MAE compared to the second best self-supervised baselines. For SDWPF dataset in energy domain, the improvement is 0.60%/1.87%. On two climate datasets, ST-ReP ranks second among self-supervised methods, with only T-Rep showing marginally better performance. This is likely due to the stronger influence of trends compared to spatial dependencies in these datasets, which are more effectively captured by the tempo-

methods	PEMS04 (N=307)		Humidity (N=2048)		CA (N=8600)	
	Fp	T_{trn}	Fp	T_{trn}	Fp	T_{trn}
ST-SSL	2.87	28.6	9.76	87.4	OOM	
GPT-ST	4.79	34.2	23.58	50.7	OOM	
STReP	1.57	24.6	6.12	25.5	23.58	458.5

Table 3: Comparison of computational efficiency in the pre-training stage. The batch size is fixed to 32. Fp: GPU memory footprints (GB), T_{trn} : the training time (s/epoch).

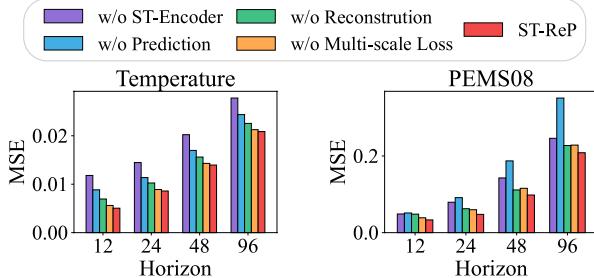


Figure 5: The results of ablation study.

ral module in T-Rep. These results demonstrate the superior predictive performance of ST-ReP and its robust adaptability across different domain datasets with various CV values, validating its ability to learn representations that are beneficial for downstream prediction tasks.

Efficiency Analysis

We compare the computational efficiency of ST-ReP with other STS forecasting models, as they all include sequences of all variables within a single batch. The results are presented in Table 3. The batch size is fixed to 32. STEP is excluded from this comparison because its pre-training requires additional input of historical data that is hundreds of times longer, resulting in significantly higher overhead with the same batch size, making it incomparable.

The results indicate that our model has the smallest memory footprints and the shortest training time on Humidity and PEMS04 datasets. This efficiency is attributed to the linear structure of our ST-Encoder. Moreover, the representations learned by our ST-ReP are significantly lower in dimensionality. Compared to other models, which often learn representations with similar or higher dimensions (e.g., d is 64 in both ST-SSL and ST-ReP, and 320 in four contrastive-based methods), the representations learned by ST-ReP achieves better predictive performance. This highlights that our ST representations have higher semantic density, reducing potential redundancy during the encoding process.

Moreover, on the CA dataset with the largest number of spatial nodes, all self-supervised learning baseline methods encounter OOM issues during both the pretraining and downstream prediction stages. This is due to their large footprint or feature dimensions. In contrast, ST-ReP is able to run on the CA dataset without reducing the model size, demonstrating its good scalability.

Ablation Study

To validate the roles of key components in the model, we conduct a set of ablation experiments, setting up the following four variant models: (1) **w/o ST-Encoder**: The ST-Encoder part of ST-ReP is removed, meaning that the ST representation generated by the model directly degenerates to \mathbf{E} from Equation 2. This model only incorporates static ST contextual information to generate future values. (2) **w/o Prediction and w/o Reconstruction**: The $Decoder_{pred}$ and $Decoder_{recon}$, along with their respective loss constraints, are removed separately to assess the roles of each decoder. (3) **w/o Multi-scale Loss**: The loss \mathcal{L}_{MS} for multiple time scales supervision is removed to evaluate the benefits brought by multi-scale temporal analysis.

Figure 5 presents the results of ablation study. Overall, each component of ST-ReP contributed to improving its predictive performance across datasets from different domains, indicating the benefits of the ST-Encoder and the three-part loss function. Additionally, we find that even without the ST-Encoder, the model’s predictions remained competitive with the best baseline, suggesting that the static context embedding and the designed loss function provide strong supervision. The performance degradation in the experiments where the decoder components were removed demonstrates the importance of these components in supervising the learned representations. Moreover, the impact of removing components varies across datasets. For example, the predictor played a more significant role than the ST encoder on the PEMS08 dataset. This variability helps us understand the unique characteristics and inductive biases of different datasets.

Conclusion

In this paper, we proposed a spatial-temporal representation learning methods ST-ReP for spatial-temporal forecasting, integrating the reconstruction for current value and prediction for unseen future value in a uniform pre-training framework. These two components, combined with multi-scale temporal analysis, constitute a ternary loss that guides the model in learning STS representations with robust predictive performance from various perspectives. Furthermore, we propose a lightweight STS encoding mechanism, build a Compression-Extraction-Decompression architecture with linear complexity. Such an architecture significantly reduces model overhead without compromising performance.

Empirical results demonstrate that the model learns representations that are both semantically dense and predictive, outperforming the best baseline methods in terms of accuracy and scalability. This highlights ST-ReP’s potential to generate compact spatial-temporal representations that can be flexibly and efficiently applied in resource-constrained downstream scenarios. We believe that this capability is an important area for STS community to explore, complementing the emphasis on accuracy improvements achieved with large-scale data and complex downstream models. Our future work will explore fusion strategies between various spatial-temporal mining models and self-supervised learning methods to enhance representation robustness and generalization across diverse downstream tasks.

Acknowledgments

This work was partly supported by the National Key Research and Development Program of China under Grant 2022YFB4501704, the National Natural Science Foundation of China under Grant 72342026, and Fundamental Research Funds for the Central Universities under Grant 2024-6-ZD-02.

References

Brown, C. E. 1998. *Coefficient of Variation*, 155–157. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-80328-4.

Challu, C.; Olivares, K. G.; Oreshkin, B. N.; Ramirez, F. G.; Canseco, M. M.; and Dubrawski, A. 2023. NHITS: Neural Hierarchical Interpolation for Time Series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6): 6989–6997. Number: 6.

Chen, L.; Chen, D.; Shang, Z.; Wu, B.; Zheng, C.; Wen, B.; and Zhang, W. 2023. Multi-Scale Adaptive Graph Neural Network for Multivariate Time Series Forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 35(10): 10748–10761. Conference Name: IEEE Transactions on Knowledge and Data Engineering.

Choromanski, K. M.; Likhoshesterov, V.; Dohan, D.; Song, X.; Gane, A.; Sarlos, T.; Hawkins, P.; Davis, J. Q.; Mohiuddin, A.; Kaiser, L.; Belanger, D. B.; Colwell, L. J.; and Weller, A. 2021. Rethinking Attention with Performers. In *International Conference on Learning Representations*.

Dong, J.; Wu, H.; Zhang, H.; Zhang, L.; Wang, J.; and Long, M. 2023. SimMTM: A Simple Pre-Training Framework for Masked Time-Series Modeling. *Advances in Neural Information Processing Systems*, 36: 29996–30025.

Fraikin, A. F.; Bennetot, A.; and Allassonniere, S. 2024. T-Rep: Representation Learning for Time Series using Time-Embeddings. In *The Twelfth International Conference on Learning Representations*.

Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 922–929.

Guo, S.; Lin, Y.; Gong, L.; Wang, C.; Zhou, Z.; Shen, Z.; Huang, Y.; and Wan, H. 2023. Self-Supervised Spatial-Temporal Bottleneck Attentive Network for Efficient Long-term Traffic Forecasting. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 1585–1596. ISSN: 2375-026X.

Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv*, abs/1606.08415: 1–10.

Huber, P. J. 1992. Robust Estimation of a Location Parameter. In *Breakthroughs in Statistics: Methodology and Distribution*, 492–518. New York, NY: Springer New York. ISBN 978-1-4612-4380-9.

Ji, J.; Wang, J.; Huang, C.; Wu, J.; Xu, B.; Wu, Z.; Zhang, J.; and Zheng, Y. 2023. Spatio-temporal self-supervised learning for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 4356–4364.

Li, Z.; Xia, L.; Xu, Y.; and Huang, C. 2023. Generative Pre-Training of Spatio-Temporal Graph Neural Networks. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Liu, J.; and Chen, S. 2024. TimesURL: Self-Supervised Contrastive Learning for Universal Time Series Representation Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(12): 13918–13926. Number: 12.

Liu, S.; Yu, H.; Liao, C.; Li, J.; Lin, W.; Liu, A. X.; and Dustdar, S. 2022. Pyraformer: Low-Complexity Pyramidal Attention for Long-Range Time Series Modeling and Forecasting. In *International Conference on Learning Representations*.

Liu, X.; Xia, Y.; Liang, Y.; Hu, J.; Wang, Y.; BAI, L.; Huang, C.; Liu, Z.; Hooi, B.; and Zimmermann, R. 2023. LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.

McDonald, G. C. 2009. Ridge regression. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(1): 93–100.

Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.

Rasp, S.; Dueben, P. D.; Scher, S.; Weyn, J. A.; Mouatadid, S.; and Thuerey, N. 2020. WeatherBench: A benchmark dataset for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11): e2020MS002203. ArXiv:2002.00469 [physics, stat].

Shabani, M. A.; Abdi, A. H.; Meng, L.; and Sylvain, T. 2023. Scaleformer: Iterative Multi-scale Refining Transformers for Time Series Forecasting. In *The Eleventh International Conference on Learning Representations*.

Shao, Z.; Zhang, Z.; Wang, F.; and Xu, Y. 2022. 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*, 1567–1577. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9385-0.

Song, C.; Lin, Y.; Guo, S.; and Wan, H. 2020a. Spatial-Temporal Synchronous Graph Convolutional Networks: A New Framework for Spatial-Temporal Network Data Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 914–921.

Song, C.; Lin, Y.; Guo, S.; and Wan, H. 2020b. Spatial-Temporal Synchronous Graph Convolutional Networks: A New Framework for Spatial-Temporal Network Data Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 914–921.

Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and ZHOU, J. 2024. TimeMixer: Decomposable Mul-

tiscale Mixing for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.

Woo, G.; Liu, C.; Sahoo, D.; Kumar, A.; and Hoi, S. 2022. CoST: Contrastive Learning of Disentangled Seasonal-Trend Representations for Time Series Forecasting. In *International Conference on Learning Representations*.

Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2023a. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *The Eleventh International Conference on Learning Representations*.

Wu, Q.; Yang, C.; Zhao, W.; He, Y.; Wipf, D.; and Yan, J. 2023b. DIFFormer: Scalable (Graph) Transformers Induced by Energy Constrained Diffusion. In *The Eleventh International Conference on Learning Representations*.

Yue, Z.; Wang, Y.; Duan, J.; Yang, T.; Huang, C.; Tong, Y.; and Xu, B. 2022. TS2Vec: Towards Universal Representation of Time Series. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8): 8980–8987. Number: 8.

Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are transformers effective for time series forecasting? In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’23/IAAI’23/EAAI’23*. AAAI Press. ISBN 978-1-57735-880-0.

Zerveas, G.; Jayaraman, S.; Patel, D.; Bhamidipaty, A.; and Eickhoff, C. 2021. A Transformer-based Framework for Multivariate Time Series Representation Learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD ’21*, 2114–2124. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-8332-5.

Zhang, K.; Wen, Q.; Zhang, C.; Cai, R.; Jin, M.; Liu, Y.; Zhang, J. Y.; Liang, Y.; Pang, G.; Song, D.; and Pan, S. 2024. Self-Supervised Learning for Time Series Analysis: Taxonomy, Progress, and Prospects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–20. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. In *The Eleventh International Conference on Learning Representations*.

Zheng, Q.; and Zhang, Y. 2023. TAGnn: Time Adjoint Graph Neural Network for Traffic Forecasting. In *Database Systems for Advanced Applications: 28th International Conference, DASFAA 2023, Tianjin, China, April 17–20, 2023, Proceedings, Part I*, 369–379. Springer-Verlag. ISBN 978-3-031-30636-5.

Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12): 11106–11115. Number: 12.

Zhou, J.; Lu, X.; Xiao, Y.; Su, J.; Lyu, J.; Ma, Y.; and Dou, D. 2022. SDWPF: A Dataset for Spatial Dynamic Wind Power Forecasting Challenge at KDD Cup 2022. arXiv:2208.04360.