

DarkFarseer: Robust Spatio-temporal Kriging Under Graph Sparsity and Noise

Zhuoxuan Liang¹, Wei Li^{*1,2}, Dalin Zhang³, Ziyu Jia^{4,5}, Yidan Chen¹, Zhihong Wang¹, Xiangping Zheng^{1,2}, Moustafa Youssef⁶

¹ College of Computer Science and Technology, Harbin Engineering University, Harbin, China

² Modeling and Emulation in E-Government National Engineering Laboratory, Harbin Engineering University, Harbin, China

³ Space Information Research Institute, Hangzhou Dianzi University, Hangzhou, China

⁴ Institute of Automation, Chinese Academy of Sciences, Beijing, China.

⁵ Shanghai Key Laboratory of Data Science, Shanghai, China.

⁶ Egyptian Wireless Research Center of Excellence, The American University in Cairo, Cairo, Egypt

{zz.liang, wei.li, chen.yidan, 1084472397, xpzheng}@hrbeu.edu.cn, zhangdalin@hdu.edu.cn, jia.ziyu@outlook.com, moustafa-youssef@aucegypt.edu

Abstract

The rapid expansion of the Internet of Things (IoT) has created a growing demand for large-scale sensor deployment. However, the high cost of physical sensors limits the scalability and coverage of sensor networks, making fine-grained sensing difficult. Inductive Spatio-Temporal Kriging (ISK) addresses this challenge by introducing virtual sensors that infer measurements from physical sensors, typically using graph neural networks (GNNs) to model their relationships. Despite its promise, current ISK methods often rely on standard message-passing and generic architectures that fail to effectively capture spatio-temporal features or represent virtual nodes accurately. Additionally, existing graph construction techniques suffer from sparse and noisy connections, further hindering performance. To address these limitations, we propose DarkFarseer, a novel ISK framework with three key innovations. First, the Style-enhanced Temporal-Spatial architecture adopts a temporal-then-spatial processing scheme with a temporal style transfer mechanism to enhance virtual node representations. Second, Regional-semantic Contrastive Learning improves representation learning by aligning virtual nodes with regional component patterns. Third, the Similarity-Based Graph Denoising Strategy mitigates the influence of noisy edges by leveraging temporal similarity and regional structure. Extensive experiments on real-world datasets demonstrate that DarkFarseer significantly outperforms state-of-the-art ISK methods.

1 Introduction

The rapid advancement and widespread adoption of the Internet of Things (IoT) have significantly increased the demand for fine-grained and ubiquitous sensing. As a result, the large-scale deployment of sensors has become a common objective across various domains (Zhou et al. 2021; Liu et al. 2022; Belavadi et al. 2020). However, the high cost associated with sensor deployment presents a substantial challenge in this regard. To mitigate this limitation, Spatio-Temporal Kriging (SK) has emerged as a promising solution. SK introduces the concept of virtual sensors, aiming to estimate

their readings based on data collected from physically deployed sensors. Existing SK methods are generally categorized into two branches: Transductive SK (TSK) and Inductive SK (ISK). TSK assumes prior knowledge of virtual sensor locations during training, whereas ISK infers virtual sensor readings without such information until inference time, making it more practical and scalable for real-world applications. Building on prior work in SK (Wu et al. 2021; Xu et al. 2025; Zheng et al. 2023; Li et al. 2024a), this paper focuses on the ISK task to further explore its potential and improve its effectiveness.

Recent deep learning-based approaches (Appleby, Liu, and Liu 2020; Wu et al. 2021; Hu et al. 2023; Xu et al. 2025; Li et al. 2024a; Wei et al. 2024), which represent the state of the art (SOTA) in inductive spatio-temporal kriging (ISK), predominantly leverage Graph Neural Networks (GNNs) to model the relationships between physical and virtual sensors—conceptualized as nodes in a graph. Despite their success, most existing methods adopt generic GNN architectures without considering the unique characteristics and challenges inherent to the ISK task. This oversight results in suboptimal graph construction and network design, ultimately limiting model performance.

L1: Limitations of spatial-first structures in modeling spatial relationships. Existing GNN-based ISK models adopt a spatial-first paradigm, where spatial dependencies are modeled prior to temporal encoding. However, this approach computes node embeddings directly from raw time series data without adequate temporal abstraction, weakening the effectiveness of spatial message passing, which is a critical component for accurate virtual sensor inference in ISK tasks.

L2: Lack of specific message-passing design in GNNs for virtual nodes. Existing GNN-based ISK methods (Wu et al. 2021; Hu et al. 2023; Xu et al. 2025; Li et al. 2024a) typically adopt generic message-passing frameworks that treat all nodes uniformly, aiming to capture global spatial information across the graph. These models aggregate features from multi-hop neighbors to learn node representations and derive graph-level embeddings. However, ISK has a specific goal: accurately inferring the values of virtual nodes.

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This objective requires emphasizing the contribution of immediate neighbors rather than relying on information from the broader graph. Furthermore, simple weighted aggregation fails to capture the complex interactions between virtual nodes and their physical neighbors. Therefore, a task-specific message-passing strategy tailored to virtual node inference is crucial for improving ISK performance.

L3: Challenges from sparsity and noise in graph structures. Statistical analysis and preliminary experiments with a representative GNN-based ISK method (Wu et al. 2021) across multiple datasets reveal two major issues: edge sparsity in expert-defined graphs and edge noise in proximity-based graphs. Pairwise Connectivity Graphs (PCGs), built via expert knowledge, often suffer from severe sparsity, limiting information flow from neighbors and hindering virtual node inference (see Appendix A.2.1). In contrast, Spatial Proximity Graphs (SPGs) exhibit high edge density, which introduces noise by connecting weakly relevant neighbors, thus degrading embedding quality (see Appendix A.2.2). Both sparse and noisy graph structures adversely affect ISK performance and must be carefully addressed.

Considering these limitations, we first conduct a comprehensive empirical study to demonstrate the sparsity and noise issues (L3) in existing graph structures. To address these challenges, we propose DarkFarseer, a novel framework that mitigates these limitations and achieves new SOTA performance on ISK tasks.

To address L1 and L2, we propose the Style-enhanced Temporal-Spatial (SETS) architecture. Unlike existing methods (Hu et al. 2023), which adopt a spatial-then-temporal pipeline, SETS follows a temporal-then-spatial design, where temporal features are first extracted and then used for spatial modeling. For temporal feature extraction, we adopt the channel-independence strategy (Zeng et al. 2023), which independently captures temporal fluctuations for each neighbor of a virtual node. For spatial message passing, we focus on the 1-hop neighbors of virtual nodes and apply a style transfer mechanism (Liu et al. 2024) that transfers the temporal fluctuation patterns of physical neighbors to virtual nodes, based on temporal similarity. The use of 1-hop neighbors ensures a balance between effectiveness and computational efficiency.

To address the **sparsity issue in L3**, we propose the Regional-semantic Contrastive Learning (RSCL) approach, which leverages regional information as contrastive samples rather than relying solely on individual nodes. This enriches the representations of nodes with sparse spatial links and alleviates the impact of sparsity. Specifically, we identify highly connected subgraphs, or Biconnected Components (BCCs), and aggregate them into prototypes (Ma et al. 2023b; Li et al. 2021) to capture high-level regional semantics. Within a contrastive learning framework, RSCL encourages virtual nodes to align with their associated prototypes while pushing them away from unrelated ones, thereby enhancing their alignment with regional patterns. Moreover, unlike existing ISK CL methods (Li et al. 2024a) that treat all nodes as anchors, RSCL uses only virtual nodes as anchors and performs fine-grained contrast, leading to more discriminative representations for sparsely connected vir-

tual nodes. To address the **noise issue in L3**, we introduce the Similarity-based Graph Denoising Strategy (SGDS), which reduces the influence of noisy edges. SGDS evaluates both the temporal similarity between each virtual node and its neighbors and their similarity to regional prototypes. It then selectively downweights the connections to dissimilar neighbors, refining the graph structure and improving representation quality.

The contributions of this paper are as follows:

- We conduct an in-depth analysis of existing graph construction methods and, through preliminary experiments, reveal two limitations: sparsity and noise in graph structures that limit the effectiveness of ISK tasks.
- We propose the SETS architecture, which adopts a temporal-then-spatial scheme and incorporates a temporal style transfer mechanism, enhancing virtual node representations by effectively capturing and fusing the fluctuation characteristics of their neighbors.
- We introduce the RSCL strategy and the SGDS module to explicitly address graph sparsity and noise. These components enhance virtual node learning by leveraging region-aware contrastive learning and edge filtering.
- We evaluate the proposed framework on multiple real-world datasets, comparing it with SOTA baselines. The results demonstrate significant performance improvements.

2 Problem Definition

Spatio-temporal Kriging (SK)

Given the observations $\mathbf{X}_{t:t+T}^o \in \mathbb{R}^{N_o \times T}$ from N_o physical sensors located at $\mathcal{L}^o = \{l_i^o\}_{i=1}^{N_o}$, the objective of SK is to infer the readings $\mathbf{X}_{t:t+T}^u \in \mathbb{R}^{N_u \times T}$ for N_u virtual sensors at locations $\mathcal{L}^u = \{l_j^u\}_{j=1}^{N_u}$. These virtual sensors are spatially related to the physical ones through spatial relationships (e.g., geographical distance). The data $\mathbf{X}_{t:t+T}^o$ and $\mathbf{X}_{t:t+T}^u$ represent multivariate time series for the observed and virtual sensors, respectively. The spatial relationships between the physical and virtual sensors are captured in a set of triplets $\mathcal{R} = \{(l_i^o, l_j^u, d_{i,j}) \mid l_i^o \in \mathcal{L}^o, l_j^u \in \mathcal{L}^u\}$, where $d_{i,j}$ denotes the spatial distance between the physical sensor l_i^o and the virtual sensor l_j^u . These spatial relationships form the foundation for modeling the inference process in SK.

Inductive SK (ISK). In ISK, the model only has access to the spatial relationships between physical nodes during the training phase, but during the inference phase, it can access the complete spatial relationships between both physical and virtual nodes. Let \mathcal{F} denote the kriging model. The kriging process during training can be formulated as:

$$\hat{\mathbf{X}}_{t:t+p}^o = \mathcal{F}_{\text{train}}(\mathbf{X}_{t:t+p}^o, \mathcal{R}^o), \quad (1)$$

where p represents the time window size. $\mathcal{R}^o = \{(l_i^o, l_j^o, d_{i,j}) \mid l_i^o, l_j^o \in \mathcal{L}^o, i \neq j\}$ represents the spatial relationships of the observed sensors. The inference phase for ISK can be formulated as:

$$\hat{\mathbf{X}}_{t:t+p}^u = \mathcal{F}_{\text{infer}}(\mathbf{X}_{t:t+p}^o, \mathcal{R}), \quad (2)$$

where \mathcal{R} represents the complete spatial relationships between both observed and virtual nodes.

Graph-based ISK. By considering sensors as nodes \mathcal{V} and the reachability relationships between nodes as edges \mathcal{E} , sensors and their spatial relationships can be represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The spatial relationships R in ISK can be represented using an adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. On top of this, the training and inference phases of graph-based ISK can be expressed as:

$$\hat{\mathbf{X}}_{t:t+p}^o = \mathcal{F}_{\text{train}}(\mathbf{X}_{t:t+p}^o, \mathbf{A}^o), \text{ where } \mathbf{A}^o \in \mathbb{R}^{N_o \times N_o}, \quad (3)$$

$$\text{and } \hat{\mathbf{X}}_{t:t+p}^u = \mathcal{F}_{\text{infer}}(\mathbf{X}_{t:t+p}^o, \mathbf{A}), \text{ where } \mathbf{A} \in \mathbb{R}^{N \times N}. \quad (4)$$

3 Analysis and Motivation

In this section, we showcase the sparsity and noise issues in current graph structures. We first investigate two types of graph structures. Subsequently, we analyze the issues of inconsistent graph density (Appendix A.2.1) and noisy edges (Appendix A.2.2) present in the two graph construction methods. Finally, we obtain three observations.

Graphs for ISK

Deep learning-based kriging methods primarily rely on GNNs that necessitate constructing a graph to model the interrelationships of sensors (i.e., nodes in a graph). Heuristic-based Graphs (Jin et al. 2024) based on the distances between nodes is a common and effective approach (Wu et al. 2019; Song et al. 2020; Li et al. 2024b). There are generally two branches to constructing heuristic-based Graphs:

Pairwise Connectivity Graphs (PCGs): In this approach, the reachability and corresponding distance between two nodes are available according to expert intervention. Thus, the adjacency matrix can be built based on the pairwise distance:

$$\mathbf{A}_{i,j} = \begin{cases} d_{i,j} & \text{if node } i \text{ and } j \text{ are linked,} \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where $d_{i,j}$ is the distance between node i and j .

Spatial Proximity Graphs (SPGs): The reachability between nodes is unknown while the geographical coordinates of each node are available. Thus, it is essential to calculate the distance between any two nodes to construct a graph. The adjacency matrix can be formulated as:

$$\mathbf{A}_{i,j} = \begin{cases} \frac{\exp(-\|d_{i,j}\|_2)}{\sigma} & \text{if } \frac{\exp(-\|d_{i,j}\|_2)}{\sigma} < \varepsilon, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

where $d_{i,j}$ is the node distance, σ is a hyper-parameter to normalize the distance distribution, and ε is the threshold to control the sparsity of the adjacency matrix.

Notably, PCGs require expert intervention for pairwise node distances, while SPGs connect nodes based on geographic distances and a pre-defined threshold ε .

Motivations

We conducted comprehensive preliminary experiments and made observations to confirm our conjectures and motivate the proposal of DarkFarseer. We present the results and analysis in Appendix A., and only report three key observations as follows:

Observation I: The PCGs are constrained by sparsity, which hinders their performance. However, these limitations can be alleviated by incorporating additional topological prompts into the ISK model.

Observation II: The PCGs are subject to a density limitation: excessive edges introduce noise into the model. Thus, designing noise-reduction strategies can help mitigate this issue and enhance performance.

Observation III: Although increasing ε to make the SPGs sparser can help, noise caused by edges constructed from geographical distance still persists. In other words, the noise edges inherent to SPGs remain, regardless of the selected value for ε .

Considering these observations, it is imperative to develop a framework that can mitigate the shortcomings associated with both PCGs and SPGs. This framework should possess two key features: (1) The capability to employ topological prompts to counteract potential graph sparsity. (2) A noise reduction strategy designed to diminish noise resulting from either excessive graph density or misleading information intrinsic to the edges themselves.

4 Methodology

DarkFarseer Overview

Training Protocol. The decrement strategy (Wu et al. 2021) is the prevailing training approach for ISK. In line with this, we adopt the decrement training strategy as proposed by IGNNK (Wu et al. 2021). Specifically, during the training phase, we randomly select N_u nodes to serve as the virtual nodes to be predicted, while using the subgraph formed by the remaining $N_o - N_u$ nodes to optimize the kriging model.

Architecture. The overall architecture of DarkFarseer is illustrated in Figure 1. At its core, DarkFarseer leverages graph structures and observed node sequences to jointly learn patterns for virtual nodes. It consists of three key components: Style-enhanced Temporal-Spatial (SETS) architecture, Regional-semantic Contrastive Learning (RSCL), and Similarity-Based Graph Denoising Strategy (SGDs). The SETS architecture employs a temporal-then-spatial fashion and equips with a fluctuation style transfer strategy to represent virtual nodes, as depicted in Figure 1(a). RSCL operates on two graph views to compute the contrastive loss \mathcal{L}_{RS} . It utilizes graph topological information to create prototypes, emphasizing the relationship between virtual node patterns and regional prototype patterns. This mitigates graph sparsity, as illustrated in Figure 1(b). SGDs reduces graph noise by adjusting edge weights based on the similarity of node sequences and their prototypes. The resulting denoised graph structure, A_s^{drop} , is shown in Figure 1(c). Finally, SETS re-encodes spatio-temporal patterns using the denoised graph. DarkFarseer is optimized with a combination of prediction loss \mathcal{L}_{MSE} and contrastive loss \mathcal{L}_{RS} . The subsequent sections provide detailed explanations of each module.

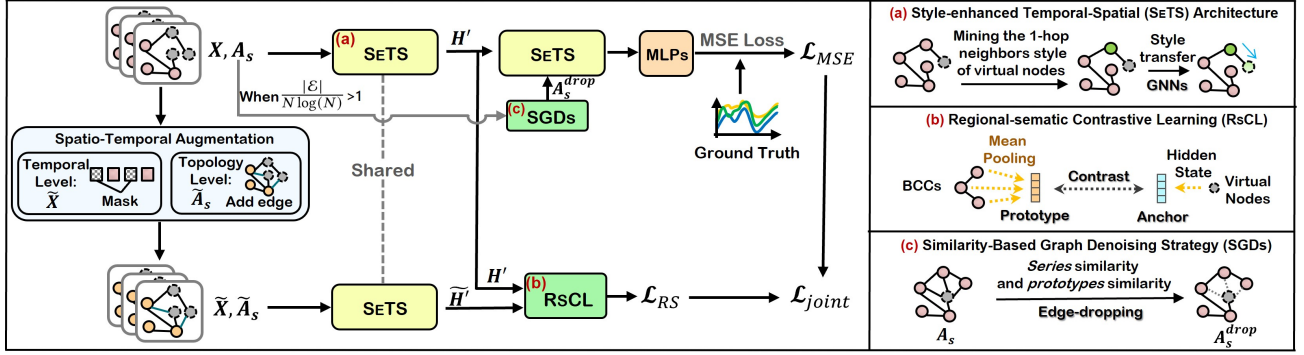


Figure 1: Overall architecture of DarkFarseer.

Style-enhanced Temporal-Spatial (SETS) Architecture

SETS first employs DLinear to independently encode each node’s seasonality and trend—leveraging channel-independence techniques (Nie et al. 2023; Han, Ye, and Zhan 2024; Dong et al. 2024)—so that the temporal dynamics of all observed neighbours are fully represented before any spatial aggregation, thus eliminating **L1**. Guided by the complete adjacency matrix \mathbf{A} , we then perform a fluctuation–style transfer: the DLinear (Zeng et al. 2023)-encoded temporal style of every virtual node’s 1-hop neighbours is propagated to the virtual node itself, injecting missing temporal knowledge to mitigate **L2**, after which spatial modules capture inter-node interactions.

Let $\mathcal{N}(i)$ represent the set of neighbors of the virtual node i . To begin, we apply DLinear to extract hidden style sequences \mathbf{S} for each corresponding time series of the neighbors. This operation produces temporal style sequences S_j for each neighbor j of virtual node i :

$$\mathbf{S}_j = \sigma(\text{DLinear}(\mathbf{X}_{s,j})), j \in \mathcal{N}(i), i \in \mathcal{V}^u, \quad (7)$$

where $\sigma(\cdot)$ is the Sigmoid activation function. Next, we fuse these style sequences with the corresponding time series data of the virtual node neighbors. We perform this fusion via element-wise multiplication:

$$\mathbf{H}_j = \mathbf{H}_j \odot \mathbf{S}_j, j \in \mathcal{N}(i), i \in \mathcal{V}^u, \quad (8)$$

where \odot denotes the Hadamard product. Here, $\mathbf{H} \in \mathbb{R}^{(N_o - N_u) \times p \times \varphi}$ represents the hidden state, with p being the number of time steps and φ the hidden dimension.

We introduce the MPNN (Cini, Marisca, and Alippi 2022) as the GNN backbone, which is expressed as:

$$\mathbf{H}' = \text{MPNN}(\mathbf{H}, \mathbf{A}_s), \quad (9)$$

where \mathbf{H}' is the updated hidden states after message passing.

Regional-semantic Contrastive Learning (RSCL)

Building on **Observation I**, we design the RSCL to enrich the representation of virtual nodes and address the limitations of sparse graphs.

For ISK tasks, we adopt two augmentation methods: (1) Temporal Augmentation, where each sequence element

within a time window is masked with a probability $p_t = 0.2$; and (2) Topology Augmentation, where edges are added between unconnected nodes with a probability $p_s = 0.003$.

CL design involves selecting three types of objects: anchors, positive samples, and negative samples.

Anchor Selection. Anchors in CL highlight target features. In ISK, we set virtual nodes as the only anchors, instead of all nodes as in (Li et al. 2024a), which reduces complexity and keeps the model focused on their temporal patterns.

Biconnected Component (BCC). BCCs are subgraphs with high local dependency, capable of representing regional information (Ma et al. 2023a). Following (Ma et al. 2023a), we set a threshold μ , and only edges with values greater than μ are used for solving BCC. In other words, μ is used to control the sparsity of the BCC. Subsequently, we use the Tarjan algorithm (Tarjan 1972) to identify the BCCs of the given graph. Mathematically, we represent each node i in the graph and its corresponding BCC C_j using $\{(i, C_j) \mid i \in \mathcal{V}, i \in C_j, j \in \{1, 2, \dots, m\}\}$, where m is the number of solved BCCs.

Positive and Negative Samples Selection. To fully leverage topological prompts, we decide to combine Prototypical Contrastive Learning (Li et al. 2021) and BCC. Inspired by (Ma et al. 2023b; Ji et al. 2023), we construct soft cluster centers as prototypes through BCC. Specifically, for each BCC C_i , we calculate the average hidden state of all nodes within C_i as the prototype, which can be formalized as:

$$\mathbf{P}_i = \frac{1}{|C_i|} \sum_{j \in C_i} \tilde{\mathbf{H}}'_j, \quad (10)$$

where $\tilde{\mathbf{H}}$ is the hidden state of augmented view. This contrast connects virtual node patterns to regional BCC patterns. However, constructing this contrast requires considering two scenarios: (1) A virtual node (anchor) exists in two or more BCCs; (2) The prototypes of other BCCs corresponding to the remaining nodes in the BCC containing the virtual node cannot be used as negative samples.

For **scenario (1)**, as shown in Figure 2, the virtual node 6 appears in both \mathbf{P}_3 and \mathbf{P}_4 . \mathbf{P}_3 and \mathbf{P}_4 should have similar patterns to node 6, so \mathbf{P}_3 and \mathbf{P}_4 should be considered positive samples. That is, all prototypes corresponding to the BCCs where the virtual node is located should be considered **positive samples**.




Figure 2: Illustrations of Contrastive Scenario (1) and (2).

For **scenario (2)**, in Figure 2, the virtual node 6 is located in the BCC corresponding to prototype \mathbf{P}_3 (yellow area). This BCC includes node 5, which is also located in the BCC corresponding to \mathbf{P}_2 (blue area). This means node 5 shares part of the pattern of \mathbf{P}_3 and part of the pattern of \mathbf{P}_2 . Thus, \mathbf{P}_2 would become a false negative sample. Therefore, when selecting negative samples, we filter out these types of prototypes to avoid generating *false negative samples*. On the other hand, \mathbf{P}_1 does not share any nodes with \mathbf{P}_3 or \mathbf{P}_4 , so \mathbf{P}_1 can be used as a **negative sample**.

Considering the above two scenarios, the InfoNCE (Oord, Li, and Vinyals 2018) for a single virtual node i can be expressed as:

$$CL(h_i) = \frac{1}{|P_i^+|} \sum_{\mathbf{P}_k \in P_i^+} -\log \frac{\exp(s(h_i, \mathbf{P}_k)/\tau)}{\exp(s(h_i, \mathbf{P}_k)/\tau) + \sum_{\mathbf{P}_j \in P_i^-} \exp(s(h_i, \mathbf{P}_j)/\tau)}, \quad (11)$$

where h_i is the hidden state of virtual node i . $s(\cdot, \cdot)$ is the cosine similarity function. τ is the temperature hyperparameter. P_i^+ is the set of prototypes that should be positive samples (**scenario (1)**) of virtual node i . P_i^- is the set of prototypes that should be negative samples (**scenario (2)**) of virtual node i . The total contrastive loss that the model needs to optimize is:

$$\mathcal{L}_{RS} = \frac{1}{|\mathcal{V}^o|} \sum_{i \in \mathcal{V}^o} CL(h_i). \quad (12)$$

Training Objective. Unlike previous works (Wu et al. 2021) that reconstruct all nodes to train the ISK model, DarkFarseer focuses on virtual nodes by training with a randomly selected sequence of \mathcal{V}^u nodes, optimizing both MSE (Mean Squared Error) loss and contrastive loss \mathcal{L}_{RS} :

$$\mathcal{L}_{joint} = \frac{1}{|\mathcal{V}^u|p} \sum_{i \in \mathcal{V}^u} \sum_{t=1}^p (x_{i,t}^o - \hat{x}_{i,t})^2 + \eta \mathcal{L}_{RS}, \quad (13)$$

where η is the weight for the contrastive loss \mathcal{L}_{RS} .

Similarity-based Graph Denoising Strategy (SGDs)

Combining **Observation II** and **Observation III**, we highlight the need for a noise reduction to address noisy edges in dense PCGs and SPGs. Thus, in this section, we propose a simple yet effective graph denoising strategy for ISK.

GNNs are sensitive to structure noise (Gilmer et al. 2017). To address this, inspired by (Dai et al. 2024; Ju et al. 2024), we propose SGDs, which optimizes graph structures and reduces the impact of noisy edges to improve the robustness of ISK models. Our focus is on edges related to virtual

nodes. Specifically, for each virtual node $i \in \mathcal{V}^u$, we assess the sequence similarity and prototype similarity between it and each of its neighboring nodes $j \in \mathcal{N}(i)$, obtaining the weighted similarity $\gamma_{i,j}$:

$$\gamma_{i,j} = \alpha s(\mathbf{P}_i, \mathbf{P}_j) + (1 - \alpha) s(\mathbf{H}'_i, \mathbf{H}'_j), \quad (14)$$

where $\alpha = \frac{1}{\mathcal{D}}$ balances the influence of different similarities, with $\mathcal{D} = \frac{|\mathcal{E}|}{N \log(N)}$ measuring graph sparsity. This ensures prototype similarity is prioritized in sparse graphs, while node feature similarity is emphasized in dense graphs to evaluate hidden edges. Next, we select the Bottom- K γ values (weakest connections) and set the corresponding edges in \mathbf{A}_s to ω , obtaining \mathbf{A}_s^{drop} :

$$\mathbf{A}_s, i, j = \omega, \quad (15)$$

where ω is a low-intensity value, and K for virtual node i is defined as $\lfloor |\mathcal{N}(i)| \times \beta \rfloor$, with β being the edge dropping probability $0 < \beta < 1$. Here, i and j represent the indices of the two nodes in the Bottom- K γ . This approach balances fine-grained temporal pattern similarity and regional pattern similarity between nodes, optimizing edge selection and enhancing graph quality. Furthermore, to allow DarkFarseer to apply SGDs adaptively, the SGDs module is activated only when $\mathcal{D} > 1$.

Complexity Analysis

In summary, DarkFarseer demonstrates good scalability *as its time and space complexities grow linearly rather than exponentially with graph size N* .

Time Complexity. For SETS, the time complexity of Equation 7 is $\mathcal{O}(p)$, and Equation 8 is implemented using a mask, resulting in a time complexity of $\mathcal{O}(1)$. Therefore, the total time complexity of SETS is $\mathcal{O}(p) + \mathcal{O}(1) = \mathcal{O}(p)$. During the training phase, for RSCL, calculating the prototypes for positive and negative samples for each virtual node incurs a cost of $\mathcal{O}(mN_u)$, assuming each BCC is utilized. Computing the \mathcal{L}_{RS} for N_u virtual nodes requires $\mathcal{O}(N_u)$. Therefore, the overall time complexity of RSCL is $\mathcal{O}(mN_u) + \mathcal{O}(N_u) = \mathcal{O}(mN_u)$. For SGDs, assuming that the average degree of nodes is d and ignoring the complexity of calculating the Bottom- K , it requires $\mathcal{O}(dN_u)$. In summary, disregarding the GNN, the time complexities of DarkFarseer during the training and inference phases are $\mathcal{O}(p + (m + d)N_u)$ and $\mathcal{O}(p + dN_u)$, respectively.

Space Complexity. For DarkFarseer, a batch of sequences is input with the shape $[B, \varphi, N', p]$, where B represents the batch size and N' is the number of neighboring nodes corresponding to the virtual nodes. SETS utilizes DLinear (Equation 7) consuming memory $\mathcal{O}(B\varphi N'p)$. In RSCL and SGDs, obtaining the prototype of a BCC using mean pooling incurs a cost of $\mathcal{O}(B\varphi p)$. In RSCL, computing the InfoNCE loss requires $\mathcal{O}(B\varphi p)$ in Equation 11. In SGDs, calculating γ requires $\mathcal{O}(B\varphi p)$ in Equation 14. Therefore, disregarding the GNN, the overall space complexity of DarkFarseer during the training and inference phases is $\mathcal{O}(B\varphi N'p + B\varphi p)$.

Datasets	Metrics	MEAN	KNN	IGNNK	DualSTN	INCREASE	KCP	KITS	DarkFarseer
PEMS03	MAE	79.78	94.804	<u>62.028±1.076</u>	71.819±1.039	77.238±0.554	79.485±0.060	86.814±1.635	59.344±0.845
	RMSE	101.791	114.96	<u>91.228±0.392</u>	106.941±1.309	106.917±0.122	107.688±0.125	129.958±3.883	86.561±2.459
	MRE	0.457	0.977	<u>0.371±0.006</u>	0.426±0.006	0.460±0.004	0.476±0.001	0.505±0.010	0.355±0.005
PEMS04	MAE	89.241	91.036	<u>70.424±0.670</u>	70.794±0.538	83.344±0.911	88.114±3.892	72.588±0.845	67.994±0.570
	RMSE	116.288	110.426	<u>97.934±0.976</u>	100.378±0.630	116.856±0.431	117.769±1.528	109.299±1.304	96.892±0.520
	MRE	0.406	0.443	<u>0.335±0.003</u>	0.333±0.003	0.395±0.005	0.419±0.018	0.319±0.004	<u>0.323±0.002</u>
PEMS-BAY	MAE	4.694	4.947	<u>3.842±0.011</u>	4.999±0.007	3.857±0.027	N/A	4.042±0.185	3.796±0.044
	RMSE	7.882	8.07	<u>6.420±0.014</u>	8.202±0.002	6.902±0.103	N/A	7.144±0.300	6.403±0.038
	MRE	0.075	0.079	<u>0.061±0.000</u>	0.079±0.000	<u>0.061±0.000</u>	N/A	0.064±0.003	0.060±0.000
AIR36	MAE	19.499	20.696	<u>18.860±0.310</u>	21.015±0.736	21.984±0.110	22.979±0.090	18.151±0.029	17.555±0.361
	RMSE	31.902	32.256	<u>29.105±0.320</u>	33.322±0.748	33.527±0.188	35.136±0.174	37.814±0.069	27.933±0.351
	MRE	0.229	0.277	<u>0.227±0.003</u>	0.252±0.009	0.265±0.001	0.278±0.001	0.238±0.001	0.211±0.004
USHCN	MAE	4.776	4.912	<u>3.713±0.031</u>	3.998±0.041	<u>3.205±0.077</u>	3.887±0.001	4.778±0.004	3.199±0.053
	RMSE	6.557	6.493	<u>5.339±0.029</u>	5.703±0.016	<u>5.406±0.147</u>	6.014±0.001	6.874±0.016	4.683±0.048
	MRE	0.664	0.822	<u>0.488±0.004</u>	0.529±0.006	<u>0.422±0.010</u>	0.511±0.000	0.659±0.001	0.421±0.007

Table 1: Performance comparison of ISK over five datasets, with the results reported as mean±std. The top and second-best results are emphasized in **bold** and underlined text. N/A denotes that augmentation (Li et al. 2024a) cannot work on PEMS-BAY.

Variants	PEMS04			PEMS-BAY			AIR36		
	MAE	RMSE	MRE	MAE	RMSE	MRE	MAE	RMSE	MRE
w/o SETS	74.594±0.591	105.669±1.592	0.355±0.002	3.806±0.109	6.385±0.114	0.060±0.001	18.759±0.544	29.288±0.459	0.226±0.006
w/o TF	69.182±0.529	98.496±0.464	0.329±0.002	3.921±0.060	6.445±0.078	0.062±0.000	18.004±0.651	28.377±0.693	0.216±0.007
w/o RsCL	68.244±0.467	97.889±0.916	0.325±0.002	3.827±0.016	6.417±0.053	0.061±0.000	18.125±0.911	28.568±0.417	0.218±0.010
w/o SGDs	-	-	-	3.889±0.105	6.472±0.036	0.062±0.001	18.572±1.419	29.149±1.108	0.223±0.017
DarkFarseer	67.994±0.570	96.892±0.520	0.323±0.002	3.796±0.044	6.403±0.038	0.060±0.000	17.555±0.361	27.933±0.351	0.211±0.004

Table 2: Ablation studies on PEMS04, PEMS-BAY and AIR36.

5 Experiments and Results

Datasets. We select five datasets from different domains (transportation, air quality, and precipitation), varying in the number of nodes and time steps: PEMS03, PEMS04, PEMS-BAY, AIR-36, and USHCN, to validate the generalization ability of the model. Among them, PEMS03 and PEMS04 belong to sparse PCGs, PEMS-BAY is a dense PCG, while AIR-36 and USHCN are SPGs.

Baselines. We select two traditional methods: MEAN and KNN, as well as five deep learning-based baselines: IGNNK (Wu et al. 2021) (AAAI’21), DualSTN (Hu et al. 2023) (TNNLS’23), INCREASE (Zheng et al. 2023) (WWW’23), KCP (Li et al. 2024a) (AISTATS’24) and KITS (Xu et al. 2025) (AAAI’25).

Implementation Details. We set the random seed to 0 for selecting virtual nodes with an observed-to-virtual nodes ratio of 3:1. The dataset is split into training, validation, and test sets in a ratio of 6:2:2. Experiments are conducted four times, and the averaged results are reported. The time window p is set to 24. We normalize the training, validation, and test sets using the global statistics of the observed nodes from the training set. The adjacency matrices constructed for all datasets are normalized using Equation 6, where ε is set to $+\infty$ for PEMS03, PEMS04 and PEMS-BAY, and to 0.1 for AIR36 and USHCN. We use Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Relative Error (MRE) as evaluation metrics. Lower values are better. The code and appendix are available at <https://github.com/DKRG-HEU/DarkFarseer>.

Performance Comparison. We compare DarkFarseer with seven baseline methods across five real-world datasets spanning multiple domains. As shown in Table 1, overall, DarkFarseer achieves the best performance. These results underscore the model’s strong generalization. On **graph-sparse datasets** (PEMS03 and PEMS04), DarkFarseer consistently outperforms all baselines, including IGNNK. This performance boost is largely attributed to the RsCL, which enriches virtual node representations by contrasting them against regional graph components, thereby alleviating the limitations imposed by sparse graph connectivity. On **graph-dense dataset** (PEMS-BAY), DarkFarseer again achieves SOTA. The SGDs plays a key role here by effectively reducing graph noise, enabling more meaningful message-passing. For **SPG datasets** (AIR36 and USHCN), where noise arises from geographic proximity, DarkFarseer maintains superior performance. The denoising capability of SGDs proves especially valuable in these settings. Notably, USHCN, with 1,218 nodes, is a very large-scale graph, and DarkFarseer’s strong performance on it shows its scalability.

Ablation Study. We perform an ablation study to assess each key component in DarkFarseer. Specifically, we create the following variants: **w/o SETS**: Removes the SETS architecture. **w/o TF**: Replaces the temporal-first structure with a spatial-first one (cf. L1). **w/o RsCL**: Removes the RsCL. **w/o SGDs**: Removes the SGDs. The results on the PEMS04, PEMS-BAY, and AIR36 datasets, presented in Table 2, demonstrate the importance of each component. Note that SGDs is not activated on PEMS04.




Figure 3: Impact of BCC sparsity levels μ , contrastive loss weight η , and edge-dropping rates β on PEMS-BAY.




Figure 4: Evaluation in various scenarios on PEMS-BAY.

Robustness Against Diverse Masking Scenarios. By default, we focus on random masking with a 25% missing rate—a prevalent evaluation protocol (Li et al. 2024a; Li, Riedel, and Beigl 2025) that involves proportionally selecting virtual nodes at random. Here, we set up different masking scenarios: (1) We validate the robustness under a 25% masking ratio by selecting three mask distributions. (2) We conduct validation under a higher masking ratio (50%). Masking ratios above 50% are not considered, as the performance degrades dramatically beyond this point.

We conducted experiments on PEMS-BAY dataset with multiple baselines. The overall results are reported in Figure 4. As shown in Figure 4(a), which aggregates results from multiple random mask distributions, DarkFarseer consistently achieves the best performance. This demonstrates the strong generalization of our method. Furthermore, Figure 4(b) presents the performance under a 50% masking ratio. At higher masking ratios, models must better utilize observed nodes to reconstruct virtual ones. In this challenging context, DarkFarseer still achieves optimal performance.

Parameter Study. We conducted parameter studies on the key parameters of RSCL and SGDs using PEMS-BAY. We analyze the impact of each parameter as follows: **BCC sparsity μ :** The μ determines sparsity of BCCs, and its value affects the RSCL. Figure 3(a) presents the results concerning μ . DarkFarseer’s performance shows low sensitivity to variations in μ . The optimal performance is observed at $\mu = 0.7$, with both excessively large and small μ values leading to suboptimal results. This suggests that for dense graphs (even with noise), an intermediate BCC sparsity is preferred over extremely sparse ones, indicating a need for balanced topological information rather than maximal dispersion. **RSCL weight η :** η determines the weight of RSCL in the loss. We perform a grid search for η within a reason-

able range. Notably, we do not strictly select η with equal spacing. Figure 3(b) reports the results. Overall, as η varies, DarkFarseer’s MAE and RMSE fluctuate minimally. This indicates that DarkFarseer is insensitive to the selection of η . **Edge-dropping rate β :** β determines the proportion of edges from the virtual node that are downgraded in SGDs. Figure 3(c) reports the results. Overall, it can be observed that as β increases, both MAE and RMSE show a significant upward trend. This aligns with the intuition that neglecting too many edges leads to a decline in the performance of DarkFarseer, which relies on GNNs.

6 Related Work

Kriging is a prominent geostatistical method used for spatial interpolation. ISK has recently sparked a wave of research (Appleby, Liu, and Liu 2020; Zhang et al. 2022; Hu et al. 2023; Xu et al. 2025; Li et al. 2024a). Deep learning-based ISK methods can be broadly categorized into two paradigms according to the scale of transfer: *partial transfer* and *global transfer*. Global transfer aims to migrate the temporal patterns of all nodes to virtual nodes. The most representative approaches employ GNNs. For example, (Li et al. 2024a) utilizes GraphSAGE (Hamilton, Ying, and Leskovec 2017). Currently, the vast majority of methods rely on global transfer. Partial transfer generally involves migrating based on the k-nearest neighbors of virtual nodes. For instance, (Zheng et al. 2023) aggregates information from heterogeneous k-neighbors. Nevertheless, *the quality of the graph structure significantly affects the effectiveness of graph-based ISK, yet few works have addressed this issue.*

7 Conclusion

In this study, we tackle key challenges hindering GNN-based ISK: generic network designs that are ill-suited for virtual node inference, and poor graph quality stemming from sparsity and noise. To this end, we propose DarkFarseer, which pioneers a temporal-first structure through SETS, using style transfer to capture spatio-temporal dependencies around virtual nodes. It further addresses graph quality issues with two innovations: RSCL leverages regional patterns to enrich representations, while the SGDs prunes noisy connections. Extensive experiments on multiple datasets show DarkFarseer achieves SOTA performance.

Acknowledgments

This research was supported by Natural Science Foundation of Heilongjiang Province, grant number LH2023F020, Supporting Fund of Intelligent Internet of Things and Crowd Computing, grant number B25029, National Natural Science Foundation of China, grant number 62502113, Project (ZR2025QC640) supported by Shandong Provincial Natural Science Foundation, Hainan Provincial Natural Science Foundation of China, grant number 625QN383, and Project (25-1-1-107-zyyd-jch) supported by Qingdao Municipal Natural Science Foundation.

References

- Appleby, G.; Liu, L.; and Liu, L.-P. 2020. Kriging convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3187–3194.
- Belavadi, S. V.; Rajagopal, S.; Ranjani, R.; and Mohan, R. 2020. Air quality forecasting using LSTM RNN and wireless sensor networks. *Procedia Computer Science*, 170: 241–248.
- Cini, A.; Marisca, I.; and Alippi, C. 2022. Filling the Gaps: Multivariate Time Series Imputation by Graph Neural Networks. In *International Conference on Learning Representations*.
- Dai, E.; Zhao, T.; Zhu, H.; Xu, J.; Guo, Z.; Liu, H.; Tang, J.; and Wang, S. 2024. A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. *Machine Intelligence Research*, 1–51.
- Dong, J.; Wu, H.; Zhang, H.; Zhang, L.; Wang, J.; and Long, M. 2024. Simtm: A simple pre-training framework for masked time-series modeling. *Advances in Neural Information Processing Systems*, 36.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*, 1263–1272. PMLR.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Han, L.; Ye, H.-J.; and Zhan, D.-C. 2024. The capacity and robustness trade-off: Revisiting the channel independent strategy for multivariate time series forecasting. *IEEE Transactions on Knowledge and Data Engineering*.
- Hu, J.; Liang, Y.; Fan, Z.; Liu, L.; Yin, Y.; and Zimmermann, R. 2023. Decoupling long-and short-term patterns in spatiotemporal inference. *IEEE Transactions on Neural Networks and Learning Systems*.
- 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.
- Jin, M.; Koh, H. Y.; Wen, Q.; Zambon, D.; Alippi, C.; Webb, G. I.; King, I.; and Pan, S. 2024. A survey on graph neural networks for time series: Forecasting, classification, imputation, and anomaly detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ju, W.; Yi, S.; Wang, Y.; Xiao, Z.; Mao, Z.; Li, H.; Gu, Y.; Qin, Y.; Yin, N.; Wang, S.; et al. 2024. A survey of graph neural networks in real world: Imbalance, noise, privacy and ood challenges. *arXiv preprint arXiv:2403.04468*.
- Li, C.; Riedel, T.; and Beigl, M. 2025. Feature Deviation Embedding Improves Graph Structure Learning for Spatial Interpolation. In *Proceedings of the 2025 SIAM International Conference on Data Mining (SDM)*, 356–365. SIAM.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. 2021. Prototypical Contrastive Learning of Unsupervised Representations. In *International Conference on Learning Representations*.
- Li, Z.; Nie, Y.; Li, Z.; Bai, L.; Lv, Y.; and Zhao, R. 2024a. Non-Neighbors Also Matter to Kriging: A New Contrastive-Prototypical Learning. In *International Conference on Artificial Intelligence and Statistics*, 46–54. PMLR.
- Li, Z.; Xia, L.; Xu, Y.; and Huang, C. 2024b. GPT-ST: generative pre-training of spatio-temporal graph neural networks. *Advances in Neural Information Processing Systems*, 36.
- Liu, J.; Huang, H.; Cao, J.; and He, R. 2024. ZePo: Zero-Shot Portrait Stylization with Faster Sampling. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 3509–3518.
- Liu, M.; Zeng, A.; Chen, M.; Xu, Z.; Lai, Q.; Ma, L.; and Xu, Q. 2022. Scinet: Time series modeling and forecasting with sample convolution and interaction. *Advances in Neural Information Processing Systems*, 35: 5816–5828.
- Ma, Q.; Zhang, Z.; Zhao, X.; Li, H.; Zhao, H.; Wang, Y.; Liu, Z.; and Wang, W. 2023a. Rethinking sensors modeling: Hierarchical information enhanced traffic forecasting. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 1756–1765.
- Ma, Y.; Song, Z.; Hu, X.; Li, J.; Zhang, Y.; and King, I. 2023b. Graph component contrastive learning for concept relatedness estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 13362–13370.
- 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*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Song, C.; Lin, Y.; Guo, S.; and Wan, H. 2020. 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.
- Tarjan, R. 1972. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2): 146–160.
- Wei, T.; Lin, Y.; Guo, S.; Lin, Y.; Zhao, Y.; Jin, X.; Wu, Z.; and Wan, H. 2024. Inductive and adaptive graph convolution networks equipped with constraint task for spatial-temporal traffic data kriging. *Knowledge-Based Systems*, 284: 111325.

- Wu, Y.; Zhuang, D.; Labbe, A.; and Sun, L. 2021. Inductive graph neural networks for spatiotemporal kriging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4478–4485.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; and Zhang, C. 2019. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 1907–1913.
- Xu, Q.; Long, C.; Li, Z.; Ruan, S.; Zhao, R.; and Li, Z. 2025. Kits: Inductive spatio-temporal kriging with increment training strategy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12945–12953.
- Zeng, A.; Chen, M.; Zhang, L.; and Xu, Q. 2023. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11121–11128.
- Zhang, Y.; Li, A.; Li, J.; Han, D.; Li, T.; Zhang, R.; and Zhang, Y. 2022. SpecKriging: GNN-based secure cooperative spectrum sensing. *IEEE Transactions on Wireless Communications*, 21(11): 9936–9946.
- Zheng, C.; Fan, X.; Wang, C.; Qi, J.; Chen, C.; and Chen, L. 2023. Increase: Inductive graph representation learning for spatio-temporal kriging. In *Proceedings of the ACM Web Conference 2023*, 673–683.
- 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. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.