CTR-Net: Continuous Tensor Ring Network for Off-the-Grid 5D Data Reconstruction

Qingfang Wang [1] **Dawei Liu** [1] **Yisi Luo** [2] **Xiaokai Wang** [1] **Wenchao Chen** [1] **Deyu Meng** [2]

Abstract

Due to cost and sampling constraints, real-world scientific data are often incomplete, posing significant challenges for data interpretation. As a result, interpolation becomes a fundamental task in scientific data processing. However, these constraints also lead to off-the-grid observations, making it difficult to directly apply conventional methods. To address this, gridding strategies are commonly used to map irregular samples onto a regular grid, but this process inevitably introduces approximation errors. Meanwhile, the rapid growth of data volume in large-scale surveys places significant demands on computational efficiency. To address these issues, we propose a Continuous Tensor Ring Network (CTR-Net), which integrates Tensor Ring (TR) decomposition with Implicit Neural Representation (INR) for large-scale off-the-grid data reconstruction. By combining continuous coordinate modeling with low-rank priors, CTR-Net effectively captures the intrinsic structure of high-dimensional data. A LoRA-based progressive learning scheme enables scalable reconstruction over large survey areas. To account for distinct temporal–spatial sampling characteristics, we introduce a decoupling strategy for temporal and spatial modeling, improving computational efficiency. Experimental results show that CTR-Net achieves a favorable balance between reconstruction accuracy and computational efficiency for large-scale 5D seismic data reconstruction.

1. Introduction

Modern scientific discovery increasingly relies on the acquisition, processing, and interpretation of large-scale multidimensional data, such as seismic exploration data (Trad, 2009), remote sensing and hyperspectral images (Liu et al., 2023; Zhang et al., 2018), dynamic medical imaging (Lustig et al., 2007), and climate and oceanographic observations (Beckers & Rixen, 2003). Although these applications are governed by different physical mechanisms, their data typically manifest as complex physical fields distributed over multiple coordinate dimensions, including space, time, frequency, angle, offset, wavelength, or other acquisition variables. From a mathematical perspective, such data can be naturally represented as high-dimensional arrays or tensors, providing an effective framework for modeling multivariate correlations in scientific and engineering problems.

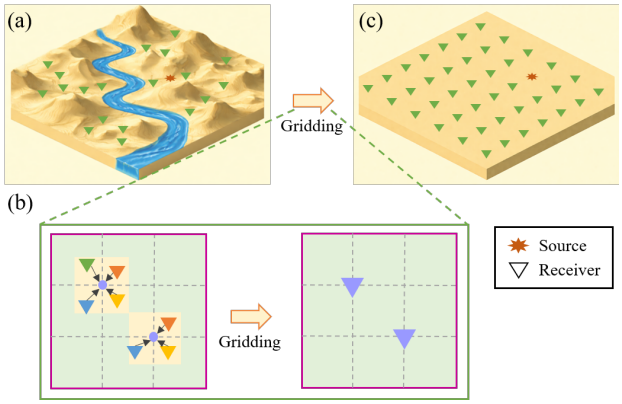


Figure 1: Illustration of off-the-grid sampling and gridding process. (a) Real off-the-grid sampling; (b) illustration of gridding; (c) ideal on-the-grid sampling.

In practice, fully sampled multidimensional scientific data are often difficult or even impossible to acquire. Unlike data defined on regular pixel (Liu et al., 2018) or voxel (Shao et al., 2023) grids, real-world high-dimensional observations are typically constrained by acquisition cost and measurement conditions, making dense, regular, and complete sampling infeasible. Complex acquisition environments further lead to off-the-grid and incomplete observations, as shown in Fig. 1(a). As a result, scientific data often exhibit strong irregularity and high missing rates, making robust reconstruction challenging.

To handle such irregularly sampled data, conventional in-

[1]School of Information and Communications Engineering, Xi'an Jiaotong University. [2]School of Mathematics and Statistics, Xi'an Jiaotong University. Correspondence to: Dawei Liu <dawei.liu@xjtu.edu.cn>.

ICML 2026 Workshop on Connecting Low-rank Representations in AI. Seoul, South Korea. Copyright 2026 by the author(s).

dustrial workflows typically rely on gridding. Specifically, observations collected at irregular spatial locations (Fig. 1(a)) are mapped onto a unified regular grid (Fig. 1(c)) via weighted averaging, followed by interpolation to fill missing samples. Although this procedure facilitates storage and numerical computation, it essentially projects a continuous physical field onto a discrete grid, inevitably introducing approximation errors that may accumulate and propagate through subsequent processing stages.

Moreover, in large-scale scientific applications, the rapid growth of data volume—often reaching TB scale or beyond—places significant demands on computational efficiency. Due to GPU memory limitations, large-scale data are often reconstructed using a window-wise strategy (Kreimer et al., 2013). However, this introduces new challenges: variations in structural features, signal characteristics, and noise distributions across different windows hinder globally consistent reconstruction when trained independently, while also significantly increasing computational cost due to redundant training across regions.

Another important but often overlooked characteristic is that different dimensions exhibit substantially different sampling densities and physical meanings. Existing reconstruction methods typically adopt fully coupled modeling (Chen et al., 2016b; Liu et al., 2022), which exploits global correlations but fails to account for the intrinsic differences between densely sampled temporal dimensions and sparsely sampled spatial dimensions. Directly coupling these heterogeneous dimensions not only increases the complexity of the inverse problem but also introduces additional computational overhead, and may degrade robustness under high missing rates and highly irregular sampling geometries.

To address these issues, we model large-scale off-the-grid data as a continuous parametric function and use 5D seismic data as a representative testbed. Seismic data are organized along time and multiple acquisition-related dimensions (e.g., spatial location, offset, and azimuth), capturing rich multi-dimensional correlations while remaining computationally tractable. 5D seismic reconstruction involves challenges including off-the-grid sampling, high missing rates, large data volumes, and complex spatial-temporal characteristics. To address these challenges, the main contributions of this work are summarized as follows:

1. A CTR-Net framework is proposed. It integrates the low-rank constraint of TR decomposition with the continuous coordinate representation of INR for modeling large-scale off-the-grid data.
2. A progressive curriculum learning strategy is proposed to reduce redundant training in window-wise processing. By combining LoRA with parameter slots, the model performs lightweight parameter updates while keeping the core parameters stable, which reduces training time across different data windows.
3. A temporal-spatial decoupled learning strategy is proposed to improve reconstruction efficiency. It separately models densely sampled temporal dimensions and sparsely sampled spatial dimensions, thereby avoiding the inefficiency of point-wise INR queries in large-scale scenarios.

2. Related Work

Traditional methods: Traditional methods can be broadly categorized into sparsity-based (Gilbert et al., 2014; Chen et al., 2016a) and low-rank methods. The low-rank methods involves matrix (Candes & Plan, 2010) and tensor completion (Liu et al., 2012) methods. Among them, singular value decomposition (SVD)-based methods are particularly prominent. The Cadzow method (Cadzow, 2002; Oropeza & Sacchi, 2011) transforms reconstruction into a low-rank matrix completion problem through Hankel matrix embedding, enabling unified denoising and interpolation. With the advancement of high-dimensional acquisition technologies, low-rank matrix approaches have naturally extended to tensor-based formulations. More recently, tensor-based methods, such as Tucker (Liu et al., 2025), Tensor Train (TT) (Sedighin et al., 2020), and Tensor Ring (TR) decomposition (Liu et al., 2022), have demonstrated strong capability in capturing global correlations through low-rank constraints. However, these methods are inherently defined on discrete grids, which limits their applicability to irregularly sampled data and continuous signal representation.

Implicit neural representations: Implicit Neural Representations (INRs) have recently emerged as a powerful framework for continuous signal modeling (Mildenhall et al., 2021). By learning a mapping from spatial coordinates to signal values, INRs provide a compact and flexible representation of high-dimensional data. Early INR models (Chen et al., 2021) typically rely on positional encoding to enhance high-frequency representation. Sinusoidal representation networks (Sitzmann et al., 2020) significantly improve the modeling of high-frequency details and have been applied to off-the-grid data reconstruction(Li et al., 2024). Because INRs directly model continuous coordinates, they are naturally well-suited for off-the-grid data, avoiding the discretization errors introduced by conventional gridding. Nevertheless, existing INR-based methods generally employ direct point-wise querying, which fails to fully exploit the intrinsic low-dimensional structure of high-dimensional data and becomes computationally inefficient for large-scale datasets.

Structured low-rank neural representations: To jointly exploit structural constraints and the expressive power of neural networks, a class of hybrid approaches has recently

emerged, which integrate low-rank priors with neural representations(Li et al., 2023). Early methods typically couple low-rank models with neural networks in a straightforward manner, but are limited to small-scale datasets and often require patch-wise processing (Liu et al., 2025). Subsequently, Luo et al. (2023) combined Tucker decomposition with INR, removing the dependence on patch-wise processing and successfully extending the framework to 3D seismic data reconstruction (Wang et al., 2026). However, the computational efficiency of existing methods for large-scale irregular data reconstruction still requires further improvement.

LoRA: In large-scale data applications, computational efficiency and scalability are critical. Window-wise training alleviates GPU memory limitations but introduces repeated training across different data windows, thereby increasing time cost. Recently, parameter-efficient learning methods, such as low-rank adaptation (LoRA) (Hu et al., 2022; Dettmers et al., 2023; Mao et al., 2025), update only a small number of additional parameters while keeping the backbone stable. This enables fast adaptation and provides a new solution to this problem. Given that different windows of the same data still exhibit strong structural similarity, LoRA can preserve shared structures during adaptation and fit local variations with only a small number of parameter updates, thereby improving computational efficiency.

3. Notations and Preliminaries

3.1. Notation

In this paper, the following notations are defined. Lowercase boldface letters (e.g., $\mathbf{x}$) denote vectors, uppercase boldface letters (e.g., $\mathbf{X}$) denote matrices, and uppercase calligraphic letters (e.g., $\boldsymbol{\mathcal{X}}$) denote higher-order tensors. Let $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$ denote a N-th order tensor. An element of tensor $\boldsymbol{\mathcal{X}}$ is represented as $\boldsymbol{\mathcal{X}}(i_1, i_2, \ldots, i_N)$, where $(i_1, i_2, \ldots, i_N)$ are the indices of the corresponding entries. Given two tensors $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ and $\boldsymbol{\mathcal{B}} \in \mathbb{R}^{I_3 \times I_4 \times I_5}$, their contraction along the third mode of $\boldsymbol{\mathcal{A}}$ and the first mode of $\boldsymbol{\mathcal{B}}$ is defined as $\boldsymbol{\mathcal{A}} \times \boldsymbol{\mathcal{B}} \in \mathbb{R}^{I_1 \times I_2 \times I_4 \times I_5}$. Here the operator "$\times$" denotes this contraction over the common mode. Furthermore, the ℓ_2 norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is defined as $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^{n} x_i^2}$, and the Frobenius norm of a tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$ is defined as $\|\boldsymbol{\mathcal{X}}\|_F = \sqrt{\sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \cdots \sum_{i_N=1}^{I_N} |\boldsymbol{\mathcal{X}}(i_1, i_2, \ldots, i_N)|^2}$.

To formally describe the 5D data, we represent a sample from the tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_t \times I_x \times I_y \times I_o \times I_a}$ using the coordinate vector:

$$\mathbf{v} = [v_t, v_x, v_y, v_o, v_a], \tag{1}$$

where $\mathbf{v}$ denotes a single point coordinate, and v_t, v_x, v_y, v_o, v_a correspond to the temporal dimension t and spatial dimensions x, y, o, and a, respectively. In addition, we use $\mathbf{v} = [v_t, v_x, v_y, v_o, v_a]$ for off-the-grid coordinates, and $\boldsymbol{\xi} = [\xi_t, \xi_x, \xi_y, \xi_o, \xi_a]$ for on-the-grid coordinates.

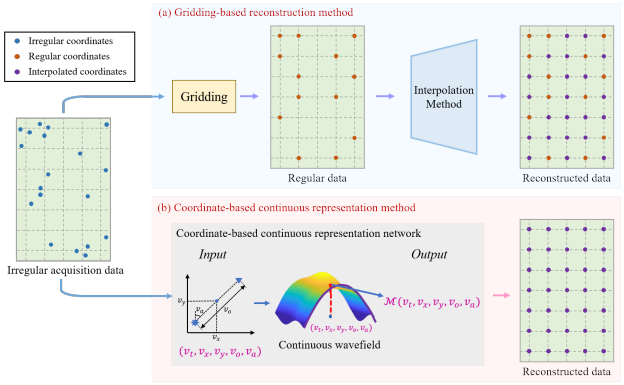


Figure 2: Illustration of reconstruction methods for off-the-grid data. (a) Reconstruction using the gridding-based method; (b) Reconstruction using the coordinate-based continuous representation method.

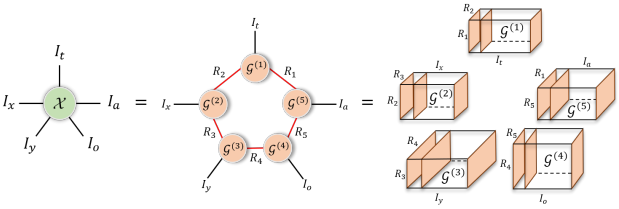


Figure 3: Illustration of the Tensor Ring (TR) decomposition.

3.2. Gridding-Based Discrete Reconstruction

Real scientific data are typically sampled off the grid, conventional methods map the irregularly sampled data onto a regular grid before reconstruction through weighted averaging, as shown in Fig. 2(a).

Let $\boldsymbol{\mathcal{D}} \in \mathbb{R}^{I_t \times I_x \times I_y \times I_o \times I_a}$ denote the gridded, incompletely sampled 5D data, and let $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{I_t \times I_x \times I_y \times I_o \times I_a}$ denote the fully sampled tensor. Their relationship can be expressed as:

$$\boldsymbol{\mathcal{D}} = \boldsymbol{\mathcal{P}} \odot \boldsymbol{\mathcal{X}}, \tag{2}$$

where $\odot$ denotes the Hadamard product and $\boldsymbol{\mathcal{P}} \in \mathbb{R}^{I_t \times I_x \times I_y \times I_o \times I_a}$ is the sampling tensor. Each element of $\boldsymbol{\mathcal{P}}$ is defined as:

$$\boldsymbol{\mathcal{P}}_{\xi_t,\xi_x,\xi_y,\xi_o,\xi_a} = \begin{cases} 0, & \text{if } \boldsymbol{\mathcal{D}}_{\xi_t,\xi_x,\xi_y,\xi_o,\xi_a} \text{ is not observed,} \\ 1, & \text{if } \boldsymbol{\mathcal{D}}_{\xi_t,\xi_x,\xi_y,\xi_o,\xi_a} \text{ is observed,} \end{cases} \tag{3}$$

where $\boldsymbol{\mathcal{D}}_{\xi_t,\xi_x,\xi_y,\xi_o,\xi_a}$ denotes the element of the 5D data $\boldsymbol{\mathcal{D}}$ indexed by $(\xi_t, \xi_x, \xi_y, \xi_o, \xi_a)$. To recover the complete tensor, we represent $\boldsymbol{\mathcal{X}}$ using a parametric model:

$$\boldsymbol{\mathcal{X}} = \mathcal{M}_\vartheta(\boldsymbol{\mathcal{D}}), \tag{4}$$

where $\mathcal{M}_\vartheta$ denotes the reconstruction model parameterized

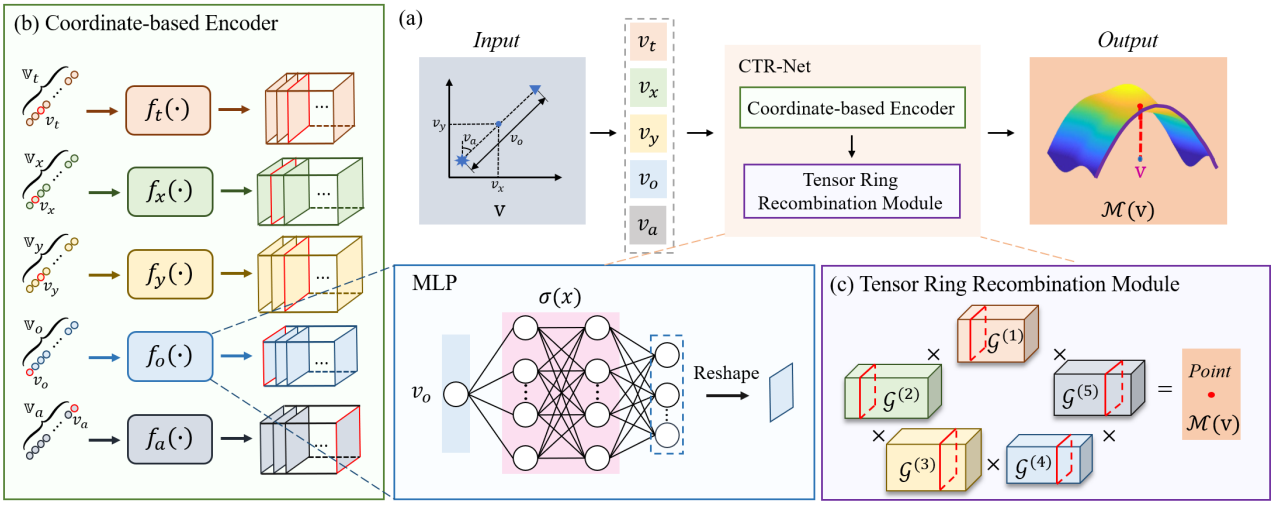


Figure 4: Illustration of the CTR-Net method. (a) Overall framework of CTR-Net, consisting of two main modules; the coordinate encoder and the tensor ring module; (b) Coordinate-based encoder, which parameterizes the core tensors by processing a single coordinate input with an MLP, followed by element reshaping to generate the core matrices; (c) Tensor Ring Recombination, where five core tensors are cyclically contracted to produce the amplitude at a single coordinate.

by ϑ. The model is optimized by:

$$\min_{\vartheta} \mathcal{L}(\vartheta) = \min_{\vartheta} \|\boldsymbol{\mathcal{P}} \odot \mathcal{M}_{\vartheta}(\boldsymbol{\mathcal{D}}) - \boldsymbol{\mathcal{D}}\|_F^2, \tag{5}$$

where $\mathcal{L}(\vartheta)$ is the reconstruction loss.

On the regular grid, 5D data exhibit strong multidimensional correlations and redundancy. Low-rank methods are particularly effective at exploiting these intrinsic structures. Among them, TR decomposition (Zhao et al., 2016) has been widely adopted due to its powerful high-dimensional representation capability. TR represents a high-order tensor using a sequence of cyclically connected third-order core tensors, as shown in Fig. 3, enabling flexible modeling of inter-dimensional dependencies.

For a 5D tensor, the TR representation is given by:

$$\begin{aligned}\mathcal{M}_{\vartheta}(i_t, i_x, i_y, i_o, i_a) = \mathrm{Tr}\{&\mathbf{G}^{(1)}(i_t) \times \mathbf{G}^{(2)}(i_x) \\ &\times \mathbf{G}^{(3)}(i_y) \times \mathbf{G}^{(4)}(i_o) \times \mathbf{G}^{(5)}(i_a)\},\end{aligned} \tag{6}$$

where $\mathbf{G}^{(n)}(i_{v_n}) = \boldsymbol{\mathcal{G}}^{(n)}(:, i_{v_n}, :) \in \mathbb{R}^{R_n \times R_{n+1}}$ denotes the i_{v_n}-th lateral slice of the core tensor $\boldsymbol{\mathcal{G}}^{(n)}$.

3.3. Coordinate-Based Continuous Reconstruction

Although TR decomposition and other discrete reconstruction methods have achieved impressive performance for regularly sampled data, it is only applicable to on-the-grid data. For off-the-grid data, an explicit gridding step is required beforehand, which inevitably introduces approximation errors and limits reconstruction accuracy.

Coordinate-based continuous reconstruction, represented by INRs (Gao et al., 2025; Sitzmann et al., 2020), offers an alternative for off-the-grid data reconstruction. As illustrated in Fig. 2(b), INR directly learns from irregularly sampled coordinates without requiring a prior gridding step. Once trained, the model can predict signal values at arbitrary target locations, enabling direct reconstruction from irregular samples to a regular grid.

The point-wise INR formulation can be written as:

$$\mathcal{M}_{\vartheta}(\mathbf{v}) = f_{\vartheta}(\mathbf{v}), \tag{7}$$

where f_{ϑ} denotes an MLP network and $\mathbf{v}$ represents the coordinate. By learning a continuous mapping from coordinates to amplitudes, INR avoids the information loss and distortion introduced by gridding. However, because INR implicitly models the signal as a continuous function without explicit global structural constraints, its reconstruction may become unstable under noisy conditions.

4. Method

4.1. Continuous Tensor Ring Network

The essence of discrete TR decomposition lies in modeling cross-dimensional correlations through the cyclic contraction of a sequence of low-order core tensors, $\boldsymbol{\mathcal{G}}^{(n)} \in \mathbb{R}^{R_n \times I_n \times R_{n+1}}$. This formulation reduces the storage complexity of a fifth-order tensor from $\mathcal{O}(I^5)$ to $\mathcal{O}(IR^2)$, where, for simplicity, we assume $I_n = I$ and $R_n = R$ for all n. Thereby effectively alleviating the curse of dimensionality.

To overcome the inherent limitation of discrete TR on regu-

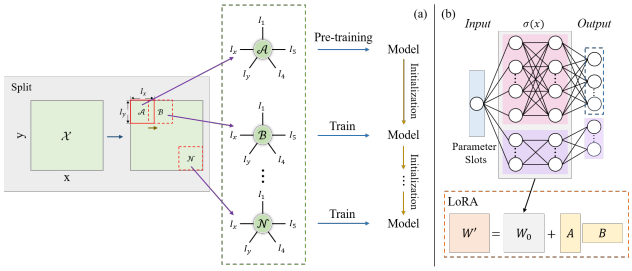


Figure 5: Progressive curriculum learning strategy based on LoRA. (a) Window-wise reconstruction of large-scale 5D data and the curriculum learning process of the model across different regions; (b) Parameter update process based on LoRA and parameter slots.

lar grids, we generalize the discrete index i_n to a continuous coordinate $v_n \in \Omega_n$, where Ω_n denotes the continuous domain of the n-th dimension. Accordingly, the discrete slice matrix $\mathbf{G}^{(n)}(i_{v_n}) \in \mathbb{R}^{R_n \times R_{n+1}}$ is extended to a continuous kernel matrix function $f_n(v_n) : \Omega_n \to \mathbb{R}^{R_n \times R_{n+1}}$.

This leads to a continuous low-rank tensor mapping from multidimensional coordinates to amplitudes (i.e., the proposed CTR-Net):

$$\begin{aligned}\mathcal{M}_\vartheta(\mathbf{v}) &= \mathrm{Tr}\left(\prod_{n=1}^{5} f_n(v_n)\right) \\ &= \mathrm{Tr}(f_t(v_t) \circ f_x(v_x) \circ f_y(v_y) \circ f_o(v_o) \circ f_a(v_a)), \quad (8)\end{aligned}$$

where $\mathbf{v} = [v_t, v_x, v_y, v_o, v_a]$ denotes the input coordinate vector, $\mathrm{Tr}(\cdot)$ represents the matrix trace, and $\circ$ denotes matrix composition in the continuous domain.

Since INRs naturally model continuous functions, each kernel function $f_n(v_n)$ is parameterized by a multilayer perceptron (MLP), as showed in Fig. 4. Specifically:

$$f_n(v_n) = \mathbf{W}_L \sigma(\mathbf{W}_{L-1} \cdots \sigma(\mathbf{W}_1 v_n)), \quad (9)$$

where $\sigma(\cdot)$ denotes the nonlinear activation function, $\{\mathbf{W}_l\}_{l=1}^{L}$ are the learnable weight matrices. In this work, we adopt the sinusoidal activation (Sitzmann et al., 2020) due to its ability to represent high-frequency details.

4.2. Progressive Curriculum Learning Strategy

Despite its strong representational capability, this continuous formulation still faces significant optimization and computational challenges in large-scale settings.

In practice, large-scale 5D data are reconstructed in windows due to GPU memory limits (Fig. 5(a)). Each window (i.e., region) is treated as an independent data segment; however, this leads to inconsistencies across regions and increases computational cost due to repeated training. To address this, we propose a progressive curriculum learning strategy. The model is first trained on a simple, high-SNR base region to obtain pretrained parameters ϑ^*, then gradually transferred to more complex regions. As shown in Fig. 5(b), a lightweight LoRA-based fine-tuning mechanism is introduced.

For the linear mapping layers in the network, the cross-region update can be expressed as:

$$\mathbf{W}' = \mathbf{W}_0 + \Delta\mathbf{W} = \mathbf{W}_0 + \mathbf{AB}, \quad (10)$$

where the core weight matrix $\mathbf{W}_0$ is inherited from the pretrained model ϑ^* and remains frozen during reconstruction in new regions. The low-rank increment $\Delta\mathbf{W}$ is formed by the product of two matrices $\mathbf{A} \in \mathbb{R}^{d \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times k}$, where $r \ll \min(d, k)$. The increment $\Delta\mathbf{W}$ is assigned to lightweight parameter slots in the network to capture region-specific geological variations. During training, only the low-rank increment $\Delta\mathbf{W}$ is updated, while the core parameters remain unchanged.

Overall, this framework supports knowledge transfer across regions, providing a scalable and consistent solution for large-scale data.

4.3. Temporal-Spatial Decoupled Learning Strategy

In TR decomposition, core tensors model data structure via cyclic products, where each element reflects interactions across dimensions. As shown in Fig. 6, INRs represent five core tensor slices that are cyclically multiplied to obtain the amplitude at a given coordinate.

Traditional INRs adopt point-wise queries, predicting one value per coordinate. However, real-world scientific datasets are often TB-scale, leading to substantial I/O overhead, redundant computation, and low GPU utilization. Moreover, different dimensions differ significantly in sampling density and physical meaning. Directly coupling dense temporal coordinates with sparse, irregular spatial coordinates increases the inverse problem complexity.

To address this issue, we exploit the variable-separation property of CTR-Net and decouple temporal and spatial dimensions based on their distinct sampling characteristics. Taking seismic data as an example, it exhibits high and uniformly sampled temporal resolution but sparse spatial sampling, allowing each time series to be treated as the basic computational unit, thereby shifting from point-wise queries to parallel time-series representation. Similar characteristics also exist in other scientific data such as acoustic signals, where temporal measurements are densely sampled while spatial sensor distributions are typically sparse. The main steps are as follows.

(1) Temporal–spatial decoupling and spatial feature reuse. All time samples in a time series share the same spatial coordinates, allowing decoupled computation. For

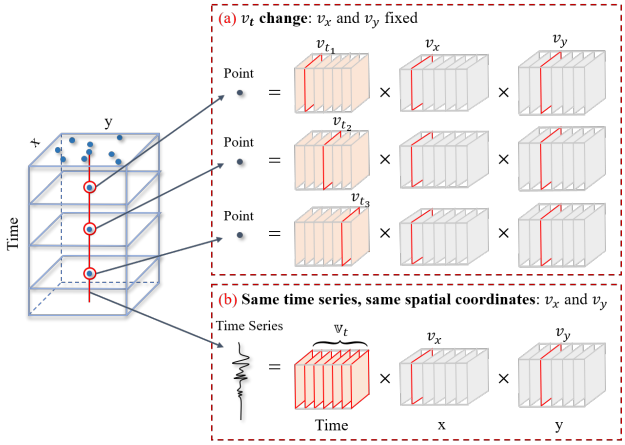


Figure 6: Comparison of the point-wise query and time-series mapping paradigms in the continuous tensor ring framework. (a) In the conventional point-wise query mode, spatial core matrices are repeatedly computed for each time sample, leading to substantial redundancy. (b) In the time-series mapping mode, samples within a time-series share the same spatial coordinates and are vectorized along the time dimension, greatly improving inference efficiency.

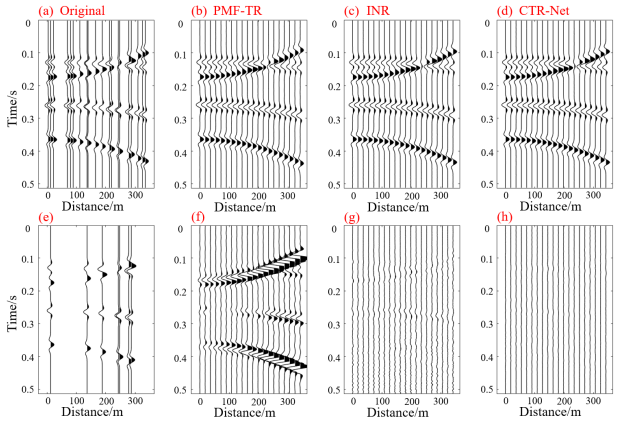


Figure 7: Reconstruction results of different methods on clean off-the-grid synthetic data with a 50% missing rate. (a) synthetic off-the-grid data; (e) observed data with missing time-series; (b)–(d) reconstruction results of PMF-TR, INR, and CTR-Net, respectively; (f)–(h) corresponding reconstruction errors. The error amplitudes are magnified by 10.

each spatial location, the corresponding spatial core tensors (i.e., features of v_x and v_y in Fig. 6) are computed once and cached.

(2) Time-series generation via tensor contraction. The cached spatial features are then contracted with the time core tensor spanning the full time axis (Fig. 6(b)), enabling one forward pass to generate an entire seismic time-series rather than a single amplitude.

The training objective minimizes the difference between the predicted and observed time-series:

$$\min_{\vartheta} \mathcal{L}(\vartheta) \quad = \quad \min_{\vartheta} \sum_{j=1}^{N_p} (\left\| \mathcal{M}_{\vartheta}(\mathbf{v}^j) - \mathbf{u}(\mathbf{v}^j) \right\|_2^2), \quad (11)$$

where $\mathbf{v}^j = \left[v_x^j, v_y^j, v_o^j, v_a^j\right]$ denotes the spatial coordinate, $\mathbf{u}(\mathbf{v}^j)$ represents the observed time-series at coordinate $\mathbf{v}^j$, and N_p is the number of observed time-series. Notably, the learning is unsupervised, with the wavefield represented by optimizing network parameters without external data. During inference, the trained model is evaluated on grid coordinates $\boldsymbol{\xi}^j$ to reconstruct amplitudes $\mathcal{M}_{\vartheta}(\boldsymbol{\xi}^j)$.

5. Experiments

We evaluate the proposed method on off-the-grid 5D synthetic seismic data and real seismic data. To better evaluate the performance, we compare our method with two methods, the discrete TR-based PMF-TR method (Liu et al., 2022) and the point-wise INR method (Li et al., 2024). The PMF-TR method is implemented in MATLAB, whereas the INR-based methods are implemented in PyTorch and trained on an NVIDIA GeForce RTX 3090 GPU.

5.1. Synthetic Off-the-Grid Data

We generated an off-the-grid 5D seismic dataset, and the reconstructed regular-grid data size is $256 \times 20 \times 20 \times 10 \times 10$. The first dimension corresponds to time, while the remaining four dimensions correspond to spatial axes. The temporal sampling interval is 2 ms, and the spatial sampling interval is 18 m. The dataset contains four reflection events. Experiments are conducted under missing rates ranging from 0.1 to 0.9 for both noise-free and noisy cases. The noisy data are obtained by adding additive white Gaussian noise (AWGN) to the clean data, resulting in an SNR of -0.02 dB. The reconstruction results are shown in Table 1 and Table 2. Since PMF-TR performs poorly on clean data, we do not test it on noisy data.

We visualize reconstruction results on clean data with a missing rate of 0.5. Fig. 7 shows a slice at (:, :, 5, 2, 2). Ideally, regularly sampled data should be evenly spaced, but Fig. 7(a) exhibits clear irregularity. Although PMF-TR results appear visually smooth, the SNR remains low due to waveform distortion and loss of signal details introduced by gridding. In contrast, CTR-Net avoids gridding and directly models irregular data, preventing resampling-induced information loss. As a result, it preserves more waveform features and achieves higher reconstruction accuracy and robustness than conventional INR methods.

We further compare computational efficiency under different

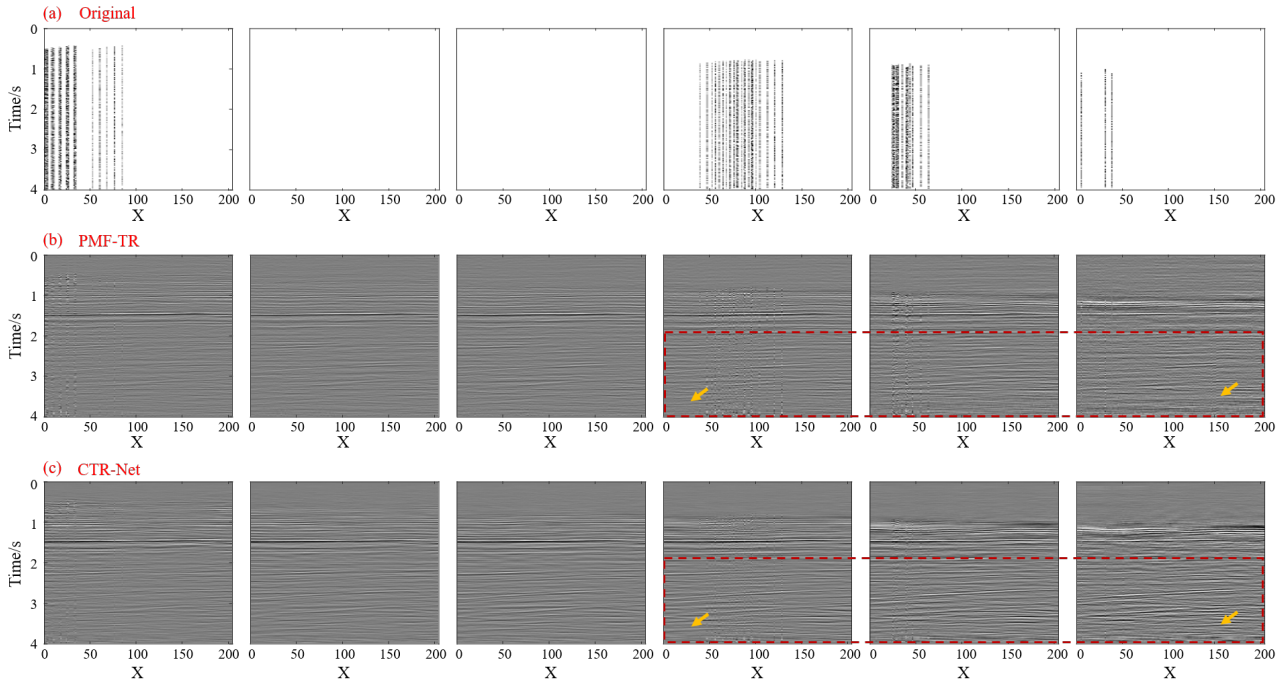


Figure 8: Reconstruction results of different methods on the real dataset. (a) Observed data; (b) and (c) reconstruction results of PMF-TR and CTR-Net, respectively. In the region marked by the red box, the proposed method produces clearer reflection events. Meanwhile, as indicated by the yellow arrow, the data recovered by CTR-Net is not significantly affected by bottom noise.

Table 1: Reconstruction results for clean off-the-grid synthetic seismic data (unit: dB).

Missing rate	PMF-TR	INR	CTR-Net
0.1	8.72	42.85	**53.96**
0.2	9.03	41.96	**52.47**
0.3	9.31	40.84	**50.89**
0.4	9.47	39.76	**49.52**
0.5	9.59	38.90	**48.30**
0.6	10.02	37.88	**47.42**
0.7	10.43	36.75	**46.71**
0.8	9.87	34.28	**41.85**
0.9	9.21	31.82	**36.69**

Table 2: Reconstruction results for noisy off-the-grid synthetic seismic data (-0.02 dB) (unit: dB).

Missing rate	INR	CTR-Net
0.1	19.84	**22.36**
0.2	19.12	**21.41**
0.3	18.46	**20.55**
0.4	17.82	**19.76**
0.5	17.21	**19.09**
0.6	16.63	**18.32**
0.7	15.87	**17.47**
0.8	14.14	**15.46**
0.9	11.90	**13.06**

training strategies. Conventional INR uses point-wise training and takes about 2h11m. In contrast, CTR-Net adopts time-series learning strategy and completes training in about 36m. This indicates that time-series learning significantly reduces computational cost and improves efficiency.

5.2. Real Off-the-Grid Data

We conduct experiments on a real 5D seismic dataset, and the reconstructed regular-grid data size is $2001 \times 204 \times 51 \times 11 \times 8$, with a missing rate of 67.2%. The sampling parameters are $\Delta x = 10$ m, $\Delta y = 10$ m, $\Delta o = 400$ m, and $\Delta a = 45°$. Due to the large data size, PMF-TR cannot reconstruct the full dataset directly. Therefore, PMF-TR is applied in a window-wise manner. In contrast, CTR-Net is trained on the entire dataset.

Fig. 8 shows the sections (:, :, 7, 3:8, 1). Off-the-grid acquisition results in only a small number of valid time-series, and the data are contaminated by significant noise, which greatly increases the difficulty of reconstruction. Nevertheless, both low-rank-based methods still recover the overall data structure reasonably well. However, in the region marked by the red box, the proposed method produces clearer and more

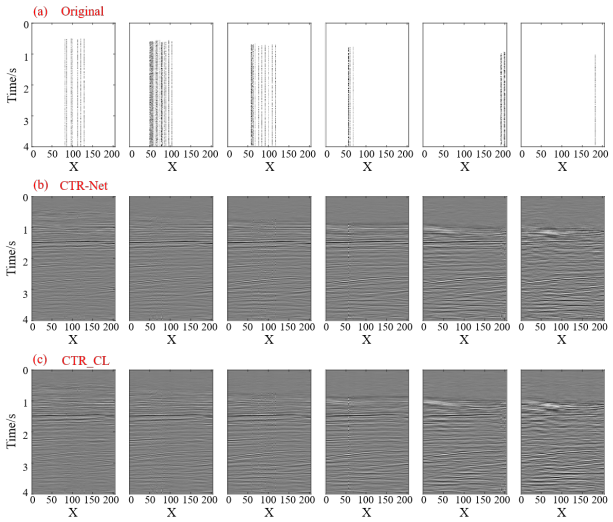


Figure 9: Reconstruction results on the same real dataset with and without the curriculum learning strategy. (a) Observed data; (b) reconstruction without a pretrained model (300 epochs); (c) reconstruction with the curriculum learning strategy (100 epochs).

continuous reflection events. In addition, at the locations indicated by the yellow arrows, the CTR-Net results are not affected by obvious bottom noise. The reconstructed events show better continuity and consistency.

5.3. Ablation Studies

5.3.1. Curriculum Learning Strategy

We use a previously trained CTR-Net model on seismic data as the initialization and introduce an curriculum learning strategy. We then train the model on another 5D seismic dataset from the same survey area, with the same reconstructed data size. Fig. 9 shows the sections (:, :, 7, 3:8, 1) and presents the results obtained using two different training strategies. Fig. 9(b) presents the result without a pre-trained model, where the network is trained from scratch for 300 epochs. Fig. 9(c) shows the result with curriculum learning, where only 100 epochs are required. The reconstruction quality of the two methods is similar, indicating that the proposed curriculum learning strategy improves training efficiency while maintaining reconstruction performance.

To further quantify the efficiency, we compare the training time with and without curriculum learning. Without curriculum learning, the training takes about 3h45m. With curriculum learning, the training time is reduced to about 1h20m, achieving more than a 50% improvement. These results show that the proposed mechanism effectively reduces redundant training cost while preserving model performance.

5.3.2. Temporal-Spatial Decoupled Learning Strategy

To evaluate the efficiency of the proposed time-series learning strategy, we compare it with conventional point-wise training on a synthetic clean dataset.

Experimental results show that point-wise training requires approximately 2h33m, whereas time-series training completes in only 36m, achieving an overall speedup of approximately $4.3\times$. This substantial acceleration is primarily attributed to the fact that time-series learning generates an entire seismic time-series in a single forward pass, thereby greatly reducing the number of network evaluations and improving computational resource utilization.

5.3.3. Effect of Tensor Decomposition

To further validate the superiority of TR decomposition in continuous tensor representation, we replace the TR module in CTR-Net with Tensor Train (TT) (Oseledets, 2011) and Tucker (Kolda & Bader, 2009) decompositions, resulting in CTT and CTucker. The results are shown in Table 3. All other settings are kept identical for a fair comparison.

Table 3: Reconstruction SNR (dB) using different tensor decomposition methods on clean synthetic data.

Missing Rate	CTucker	CTT	CTR
0.5	43.52	45.96	**48.30**
0.7	41.78	44.19	**46.71**
0.9	31.64	34.21	**36.69**

As shown in Table 3, CTR consistently achieves the best reconstruction performance across all missing rates, followed by CTT, while CTucker performs the worst. This superiority can be attributed to the cyclic structure of TR, which effectively alleviates boundary rank constraints and provides greater representational flexibility.

6. Conclusions

Experimental results demonstrate that the proposed CTR-Net achieves accurate and stable reconstruction of high-dimensional (e.g., 5D) scientific data under off-the-grid and extremely sparse sampling conditions. Compared with conventional low-rank and grid-based methods, it produces fewer aliasing artifacts, better preserves structural continuity, and more faithfully recovers details. The LoRA-based approach improves large-scale training efficiency. Meanwhile, the temporal-spatial decoupled strategy significantly improves computational efficiency, enabling practical processing of large-scale data. Overall, CTR-Net provides an effective, efficient, and scalable solution for off-the-grid data reconstruction.

7. Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Beckers, J.-M. and Rixen, M. EOF calculations and data filling from incomplete oceanographic datasets. *Journal of Atmospheric and oceanic technology*, 20(12):1839–1856, 2003.

Cadzow, J. A. Signal enhancement-a composite property mapping algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(1):49–62, 2002.

Candes, E. J. and Plan, Y. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

Chen, H., He, B., Wang, H., Ren, Y., Lim, S. N., and Shrivastava, A. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34: 21557–21568, 2021.

Chen, X., Kane, D. M., Price, E., and Song, Z. Fourier-sparse interpolation without a frequency gap. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 741–750. IEEE, 2016a.

Chen, Y., Zhang, D., Jin, Z., Chen, X., Zu, S., Huang, W., and Gan, S. Simultaneous denoising and reconstruction of 5-D seismic data via damped rank-reduction method. *Geophysical Journal International*, 206(3):1695–1717, 2016b.

Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.

Gao, W., Liu, D., Chen, W., Sacchi, M. D., and Wang, X. NeRSI: Neural implicit representations for 5D seismic data interpolation. *Geophysics*, 90(1):V29–V42, 2025.

Gilbert, A. C., Indyk, P., Iwen, M., and Schmidt, L. Recent developments in the sparse fourier transform: A compressed fourier transform for big data. *IEEE Signal Processing Magazine*, 31(5):91–100, 2014.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

Kreimer, N., Stanton, A., and Sacchi, M. D. Tensor completion based on nuclear norm minimization for 5D seismic data reconstruction. *Geophysics*, 78(6):V273–V284, 2013.

Li, J., Liu, D., Trad, D., and Sacchi, M. D. Robust unsupervised 5D seismic data reconstruction on regular and irregular grids. *Geophysics*, 89(6):V537–V549, 2024.

Li, S., Cheng, L., Zhang, T., Zhao, H., and Li, J. Striking the right balance: Three-dimensional ocean sound speed field reconstruction using tensor neural networks. *The Journal of the Acoustical Society of America*, 154(2):1106–1123, 2023.

Liu, D., Sacchi, M. D., and Chen, W. Efficient tensor completion methods for 5-D seismic data reconstruction: Low-rank tensor train and tensor ring. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–17, 2022.

Liu, D., Wang, Q., You, N., Sacchi, M. D., and Chen, W. Filling the gap: Enhancing borehole imaging with a tensor neural network. *Geophysics*, 90(3):D71–D83, 2025.

Liu, J., Musialski, P., Wonka, P., and Ye, J. Tensor completion for estimating missing values in visual data. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):208–220, 2012.

Liu, N., Li, W., Wang, Y., Tao, R., Du, Q., and Chanussot, J. A survey on hyperspectral image restoration: From the view of low-rank tensor approximation. *Science China Information Sciences*, 66(4):140302, 2023.

Liu, Y., Chen, X., Wang, Z., Wang, Z. J., Ward, R. K., and Wang, X. Deep learning for pixel-level image fusion: Recent advances and future prospects. *Information fusion*, 42:158–173, 2018.

Luo, Y., Zhao, X., Li, Z., Ng, M. K., and Meng, D. Low-rank tensor function representation for multi-dimensional data recovery. *IEEE transactions on pattern analysis and machine intelligence*, 46(5):3351–3369, 2023.

Lustig, M., Donoho, D., and Pauly, J. M. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

Mao, Y., Ge, Y., Fan, Y., Xu, W., Mi, Y., Hu, Z., and Gao, Y. A survey on lora of large language models. *Frontiers of Computer Science*, 19(7):197605, 2025.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

Oropeza, V. and Sacchi, M. Simultaneous seismic data denoising and reconstruction via multichannel singular spectrum analysis. *Geophysics*, 76(3):V25–V32, 2011.

Oseledets, I. V. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

Sedighin, F., Cichocki, A., Yokota, T., and Shi, Q. Matrix and tensor completion in multiway delay embedded space using tensor train, with application to signal reconstruction. *IEEE Signal Processing Letters*, 27:810–814, 2020.

Shao, R., Zheng, Z., Tu, H., Liu, B., Zhang, H., and Liu, Y. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16632–16642, 2023.

Sitzmann, V., Martel, J., Bergman, A., Lindell, D., and Wetzstein, G. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

Trad, D. Five-dimensional interpolation: Recovering from acquisition constraints. *Geophysics*, 74(6):V123–V132, 2009.

Wang, Q., Liu, D., Sacchi, M. D., Wang, X., Meng, D., and Chen, W. STISR: A Stacked Tucker Implicit Seismic Reparameterization Framework for 3D Seismic Data Denoising. *Petroleum Science*, 2026.

Zhang, Q., Yuan, Q., Zeng, C., Li, X., and Wei, Y. Missing data reconstruction in remote sensing image with a unified spatial–temporal–spectral deep convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 56(8):4274–4288, 2018.

Zhao, Q., Zhou, G., Xie, S., Zhang, L., and Cichocki, A. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*, 2016.