

Latent assimilation with implicit neural representations for unknown dynamics

Zhuoyuan Li^{a,*}, Bin Dong^{b,c}, Pingwen Zhang^{d,a}

^a School of Mathematical Sciences, Peking University, Beijing 100871, China

^b Beijing International Center for Mathematical Research, Peking University, Beijing 100871, China

^c Center for Machine Learning Research, Peking University, Beijing 100871, China

^d School of Mathematics and Statistics, Wuhan University, Wuhan 430072, China

ARTICLE INFO

Keywords:

Data assimilation
Implicit neural representation
Spherical harmonics
Unstructured data modeling
Uncertainty estimation

ABSTRACT

Data assimilation is crucial in a wide range of applications, but it often faces challenges such as high computational costs due to data dimensionality and incomplete understanding of underlying mechanisms. To address these challenges, this study presents a novel assimilation framework, termed Latent Assimilation with Implicit Neural Representations (LAINR). By introducing Spherical Implicit Neural Representations (SINR) along with a data-driven uncertainty estimator of the trained neural networks, LAINR enhances efficiency in the assimilation process. Experimental results indicate that LAINR holds a certain advantage over existing methods based on AutoEncoders, both in terms of accuracy and efficiency.

1. Introduction

Data assimilation (DA) has increasingly become an essential tool across a multitude of disciplines, including but not limited to weather forecasting, climate modeling, oceanography, ecology, and even economics [39,40,55,51,54]. It aims to integrate information from various data sources, such as in-situ measurements, satellite, and radar observations [41], with existing physical models that describe a system's underlying dynamics. The objective of DA is to obtain the best estimate of the system states while accounting for uncertainties in both observations and physical models.

From a mathematical standpoint, DA offers an efficient and practical solution to a specific type of optimal control problem [5]

$$\begin{cases} \dot{\mathbf{u}}(t) = \mathcal{M}(\mathbf{u}(t), t) + \boldsymbol{\varepsilon}^M(t), & 0 \leq t \leq T, \\ \mathbf{u}(0) = \mathbf{u}_0 \end{cases} \quad (1)$$

with the objective function

$$\mathbf{J}(\mathbf{u}_0) = \mathcal{L}_B(\mathbf{u}_0, \mathbf{u}^b) + \int_0^T \mathcal{L}_O(\mathcal{H}(\mathbf{u}(t)), \mathbf{y}(t), t) dt + \int_0^T \mathcal{L}_E(\boldsymbol{\varepsilon}^M(t)) dt \quad (2)$$

* Corresponding author.

E-mail addresses: zy.li@stu.pku.edu.cn (Z. Li), dongbin@math.pku.edu.cn (B. Dong), pzhang@pku.edu.cn (P. Zhang).

to be minimized. Here, the system dynamics is denoted by \mathcal{M} , and \mathcal{H} stands for the observation operator. \mathcal{L}_B and \mathcal{L}_O penalize the discrepancy between the system states and the background estimate \mathbf{u}^b and the observation data $\mathbf{y}(t)$, respectively, while \mathcal{L}_E penalizes the model error $\epsilon^M(t)$. The goal is to find the optimal control \mathbf{u}_0 as the initial system state that minimizes the objective function $\mathbf{J}(\mathbf{u}_0)$. For instance, in the realm of modern numerical weather prediction (NWP), $\mathbf{u}(t)$ may consist of the meteorological variables of interest governed by atmospheric dynamics, then DA frameworks are employed to determine the optimal system states for further prediction. In practice, dealing with the optimization problem (1) defined on an infinite-dimensional functional space is challenging. Therefore, the problem is typically reformulated by discretizing on both the spatial and temporal dimensions, which leads to minimizing

$$\mathbf{J}(\mathbf{x}_0) = \mathcal{L}_B(\mathbf{x}_0, \mathbf{x}^b) + \sum_k \mathcal{L}_O(\mathcal{H}_k(\mathbf{x}_k), \mathbf{y}_k) + \sum_k \mathcal{L}_E(\epsilon_k^M) \quad (3)$$

subject to

$$\mathbf{x}_{k+1} = \mathcal{M}_k(\mathbf{x}_k) + \epsilon_k^M, \quad k = 0, 1, \dots, \quad (4)$$

where the subscripts k denote temporal indices. Here, $\mathbf{x} = \{\mathbf{u}(\mathbf{p})\}_{\mathbf{p} \in S}$ represents a discrete sampling of \mathbf{u} with a sampling set S . To model the discrete dynamics, a numerical approximation \mathcal{M}_k for integrating $\mathcal{M}(\cdot, t)$ in (1) is typically employed, which inherently encompasses the integration steps required for transitioning between successive time steps. Similarly, the accumulated error ϵ_k^M is obtained from $\epsilon^M(t)$ in (1). Besides, \mathcal{H}_k stands for the observation operator discretized on the sampling set. In this work, we refer to such downsampled version \mathbf{x} as the ‘‘physical state’’ in contrast to the ‘‘latent state’’ \mathbf{z} to be introduced later, and one should not confuse the physical state \mathbf{x} with the system state \mathbf{u} since the physical state does not contain all the physical features of the system state in general.

Early efforts for DA primarily revolved around techniques like nudging and optimal (statistical) interpolation methods. Subsequently, variational methods (3/4D-Var) have been proposed based on the optimal control theory. A key advantage of the variational approaches lies in their ability to make meteorological fields comply with the dynamical equations of NWP models, while simultaneously minimizing the discrepancies between simulations and observations. However, these methods are often computationally intensive, particularly for high-dimensional systems. Additionally, the necessity to derive the corresponding tangent linear models and adjoint models becomes another challenge, especially for complex dynamics.

Another advanced development of DA frameworks is the proposal of Kalman Filter (KF) methods [27,5], which facilitate sequential assimilation of data, and thus offer benefits for real-time applications. Literature such as [74] has shown that Kalman filtering is equivalent to variational assimilation in some limits. Unfortunately, the challenge of high dimensionality persists in that the scales of the computational grid and the underlying matrices are linked to the state dimension. Some approximate or suboptimal Kalman filters have been developed to partially alleviate the issue, but at the expense of confining the assimilated states within a low-dimensional subspace, potentially harming the overall performance.

Moreover, existing DA frameworks rely on accurate dynamical cores. Nevertheless, capturing small-scale physical processes or edge effects is still a challenge, especially where the physical knowledge is limited and no well-designed parameterizations are available, which will probably lead to significant and undetermined model errors. Another notable obstacle arises when dealing with highly non-linear systems. The inherent difficulties in optimizing non-convex problems and the inconsistency with the linear assumptions of the Kalman filter point to a need for more advanced DA techniques.

One approach to handle the high-dimensionality issue is Reduced-Order-Models (ROMs), which establish a lower-dimensional parameterization space and then shift the DA framework to the latent space. Classical ROMs usually employ linear trial subspaces for model reduction, and they require the subspace of dimension much larger than the intrinsic dimension of the solution manifold to ensure high accuracy, especially for highly non-linear problems [56]. Based on the Kolmogorov n -width [63], [24] has proposed a continuous non-linear version

$$d_n(S) = \inf_{\substack{P: S \rightarrow \mathbb{R}^n \\ Q: \mathbb{R}^n \rightarrow S}} \sup_{f \in S} \|f - Q(P(f))\| \quad (5)$$

to quantify the intrinsic complexity of a compact functional space K , which together with the recent work [44,31,82] shows the theoretical superiority of non-linear ROMs over linear ones.

Machine Learning (ML), and specifically Deep Learning (DL) has emerged as an effective data-driven tool to approximate high-dimensional mappings thanks to its strengths to extract intricate structures in high-dimensional data [43], and consequently, neural networks are used not only to establish non-linear ROMs, whose training processes depend on reconstruction loss functions similar to (5) where the mappings P and Q are translated as parameterized encoders and decoders, but also to learn the corresponding reduced or latent dynamics [34,67,16] from the encoded data. Provided with a well-trained encoder-decoder pair denoted by \mathcal{E}_θ and \mathcal{D}_φ , respectively, we may obtain a latent state \mathbf{z}_k for each \mathbf{x}_k from (3)(4) such that

$$\mathbf{z}_k = \mathcal{E}_\theta(\mathbf{x}_k), \quad \mathbf{x}_k = \mathcal{D}_\varphi(\mathbf{z}_k) + \epsilon_k^{\text{E-D}} \quad (6)$$

with minor reconstruction error $\epsilon_k^{\text{E-D}}$. Furthermore, if we additionally succeed in attaining a latent transition $\mathcal{G}_{\psi,k}$ such that

$$\mathbf{z}_{k+1} = \mathcal{G}_{\psi,k}(\mathbf{z}_k) + \epsilon_k^L \quad (7)$$

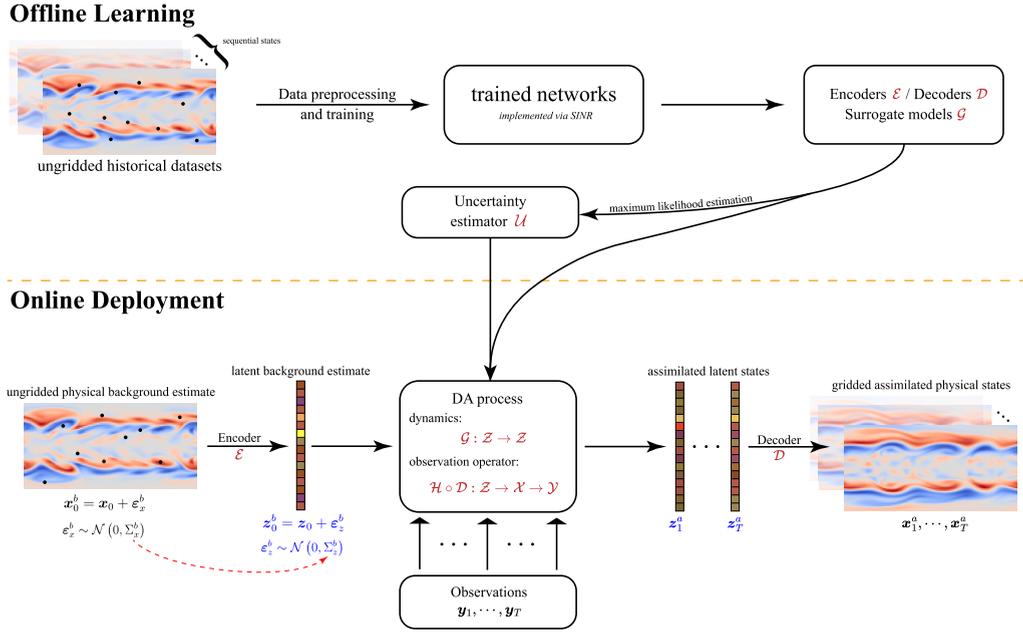


Fig. 1. The pipeline of the LAINR framework.

with minor prediction error ϵ_k^L , we may solve the assimilation problem by minimizing the objective function

$$J(z_0) = \mathcal{L}_B(z_0, z^b) + \sum_k \mathcal{L}_O(H_k(D_\varphi(x_k)), y_k), \tag{8}$$

on a lower-dimensional space where all the latent states $\{z_k\}_k$ reside. Such a transition for DA from physical space to latent space is referred to as the latent assimilation (LA) [1]. The latent dynamical model $\mathcal{G}_{\psi,k}$ is usually data-driven, and no expert knowledge is required. By optimizing on a lower-dimensional space, LA has the benefit of reducing both the storage and computational demands compared with traditional methods.

In recent years, the rapid development of DL has had a transformative impact on various domains, including computer vision, natural language processing, and more recently scientific computing. The applications of DL in solving complex physical problems, such as solving differential equations [64,79,71], discovering underlying physical laws [47,12,77], and simulating spatio-temporal dynamics [48,45], have attracted significant attention and achieved promising results. In the specific area of atmospheric sciences, DL has proven itself highly effective in handling a variety of tasks, ranging from weather nowcasting and forecasting to post-processing and data assimilation. Specifically, the success of forecast models like FourCastNet [58], Pangu-Weather [11], GraphCast [42] and FengWu [15] focusing on learning surrogate dynamics has gradually dissolved the initial doubts [70] regarding DL’s capability to handle complex dynamics.

In this work, we present the Latent Assimilation with Implicit Neural Representations (LAINR) framework (Fig. 1). Our framework is designed to assimilate data from systems with partially or even completely unknown dynamics. By proposing a new variant of the implicit neural representation, namely the Spherical Implicit Neural Representation (SINR), and a comprehensive assimilation procedure for learning surrogate models and estimating corresponding uncertainties, we aim to make a meaningful contribution to the ongoing exploration of DL and ML techniques in data assimilation and atmospheric sciences.

2. Related works and contributions

Reduced-order models The concept of Reduced-Order Models (ROMs) is not a new idea for data assimilation. As introduced previously, dimensionality reduction techniques are employed to represent high-dimensional vectors with lower-dimensional ones. Classical ROMs usually rely on linear methods like Proper Orthogonal Decomposition (POD) [14,4,3], Dynamic Mode Decomposition (DMD) [69] or wavelet decomposition [76] to extract system’s dominant modes from historical physical states, then the corresponding parameterizations $z \in \mathcal{Z}$ for each $x \in \mathcal{X}' \subseteq \mathcal{X}$ are the projection coefficient vectors onto each mode. Recently, DL advancements have offered alternative approaches to construct or enhance ROMs. AutoEncoders [8,84] are one of the most common methods to learn a non-linear, lower-dimensional representation of the data, and the learned representations can then be used to study and understand the behavior of the system at a lower computational cost [44,31]. Meanwhile, Recurrent Neural Networks (RNNs) [46] including the Long Short-Term Memory (LSTM) networks [37] and the Gated Recurrent Unit (GRU) networks [21] have also been employed to model time-dependent problems and capture the temporal correlations in the dataset.

Table 1

Comparison of different methods for latent assimilation. The method names are borrowed from the original papers.

method	encoder-decoder	latent dynamics	scalability	efficiency	flexibility
ETKF-Q-L [62]	fully-connected	ReZero	low	high	no
RNN-ETKF [60]	RNN	RC	medium	high	no
NIROM-DA [59]	POD/PCA	LSTM	low	medium	no
GLA [20]	POD/PCA + Conv1d	LSTM	low	medium	no
LAINR (ours)	SINR	Neural ODE	high	medium	yes

Latent assimilation Latent assimilation has also benefited from the advent of DL, particularly in constructing non-linear encoder-decoder mappings for assimilation in latent space. For instance, [62] selects simple fully-connected deep neural networks for both the encoder and the decoder, with ReZero [7] as a variant of ResNet [35] to enhance the stability of learning the latent dynamics. In spite of its effectiveness of dimensionality reduction, it is faced with scalability limitations since the fully-connected architecture can hardly be applied to multi-dimensional cases due to the unaffordable costs of storage and training for the huge number of network parameters. Recurrent Neural Networks (RNNs) can serve as an alternative way for merging the encoder-decoder structure and evolving latent dynamics simultaneously as proposed in [60]. The training of RNNs takes advantage of Reservoir Computing (RC) [49] to make the forward propagation faster, but as suggested in [2,60], the RC approaches are more suitable for short-term prediction tasks. Meanwhile, creating patches when scaling up to high-dimensional problems will probably result in disagreement between neighborhoods. To effectively capture long-term dependencies in sequential data, LSTM architectures have been employed in [59,20], but they face the same issues as those of [62] that the model is difficult to scale up to handle high-dimensional problems. Finally, all the architectures mentioned above require the input data to lie on certain fixed grids, and the inevitable interpolation procedures probably bring additional errors. Table 1 summarizes the comparative strengths and weaknesses of these methods as well as our proposed framework introduced later.

Implicit neural representations Implicit Neural Representations (INRs) have emerged as an effective tool for encoding high-dimensional data, geometrical objects, or even functions in a compact and parameterized manner. Diverging from common neural networks representing data in the form of grids or graphs, INRs model them as continuous functions. Such approach has shown its superiority in a variety of applications, including shape modeling [19,57], image processing [73,9,26], and spatio-temporal dynamics [38,16,83]. One key advantage of INRs lies in their agnosticism to data structures, leading to a high degree of flexibility. Moreover, similar to AutoEncoders, the non-linearity inherent to neural networks enables them to capture non-linear correlations, which have more potential than classical linear ROMs. In this work, we propose a new variant of INRs specially designed for 2D spherical data, namely the Spherical Implicit Neural Representations (SINRs), which aims to capture the complex dynamics of the system in a low-dimensional latent space based on the spherical harmonics.

The key aspects of our contributions can be summarized below.

- **A Novel Mesh-free DA Framework:** We present a new assimilation framework named Latent Assimilation with Implicit Neural Representation (LAINR), which is mesh-free and applicable to multi-dimensional unknown dynamics.
 - *SINRs.* We have formulated a specialized variant of INRs called the Spherical Implicit Neural Representations (SINRs) to handle 2D spherical data. SINRs are capable of capturing the system’s complex dynamics at a lower cost by embedding the physical states into a low-dimensional latent space while ensuring high accuracy and reliable convergence guarantee.
 - *Uncertainty Estimators.* A simple but effective uncertainty qualification method has been developed to estimate uncertainties associated with both the latent surrogate model and the encoding process, leveraging maximum likelihood estimation (MLE) techniques.
- **Interoperability with Existing DA Algorithms:** LAINR has been engineered for compatibility to work seamlessly with existing data assimilation algorithms, thereby enhancing the performances and extending their applications. This feature encourages wider adoption for future advancements in the field of DA.
- **Experimental Validation:** We have conducted experiments on multi-dimensional cases, including an ideal shallow-water model and a real meteorological dataset (ERA5) to demonstrate the superiority of LAINR compared to existing AutoEncoder-based methods. The effectiveness of LAINR in capturing spatio-temporal dynamics, seamlessly integrating with existing DA algorithms and handling unstructured and unseen data is emphasized in the empirical findings.

The remainder of this paper is structured as follows. First, an overview of the existing LA approaches as well as their limitations is introduced in Section 3, which is followed by Section 4 that presents our proposed LAINR framework in detail. Section 5 describes the experimental setup and exhibits the model performances. Finally, Section 6 concludes the main features of our LAINR framework and outlines some of the potential avenues for future research.

3. Latent assimilation

Latent Assimilation (LA) is essentially a specialized form of Data Assimilation (DA) methodology that makes use of dimensionality reduction of the assimilation space in order to save computational and storage costs. It characterizes the high-dimensional state



(a) Classical DA that operates on the physical space

(b) LA that operates on the latent space

Fig. 2. Illustrations of classical DA and LA frameworks with colored forward operators (red) and observation operators (blue). The temporal indices here are omitted for simplicity. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

(physical) space \mathcal{X} , with a more compact, lower-dimensional parameterization (latent) space \mathcal{Z} . By establishing a well-designed encoder-decoder pair capable of transforming data between the physical space and the latent space with high accuracy, LA shifts the physical states, physical dynamics, and the assimilation mechanism from \mathcal{X} onto \mathcal{Z} . Same as conventional DA methodologies, LA also requires a forward propagation operator within the latent space \mathcal{Z} and an associated observation operator that links the latent space \mathcal{Z} to the observational space \mathcal{Y} (Fig. 2). Additionally, the uncertainties should also be accounted for within these mappings if possible.

3.1. Embeddings from physical spaces to latent spaces

LA framework explores suitable latent embeddings with the assumption that a low-dimensional parameterization exists for the high-dimensional physical states. Let $\mathcal{X} = \mathbb{R}^n$ and $\mathcal{Z} = \mathbb{R}^m$ be the physical state space and the latent space, respectively, where n is considerably larger than m . The goal is to identify an encoder $\mathcal{E} : \mathcal{X} \rightarrow \mathcal{Z}$ and a decoder $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{X}$ that function as appropriate inverse mappings between the low-dimensional submanifold $\mathcal{X}' \subseteq \mathcal{X}$ encapsulating all possible physical states, and the latent space \mathcal{Z} . Here, we use the notation $\mathbf{u} : \mathbb{R}^d \rightarrow \mathbb{R}^c$ for the target system state to be assimilated, where d signifies the spatial dimension of the dynamical system (usually 2 or 3), and c refers to the number of features under consideration, such as horizontal wind, temperature, humidity, and so forth. For a fixed sampling set $S \subseteq \mathbb{R}^d$, several strategies can be established to construct ROMs for the physical states or the snapshots $\mathbf{x}_k = \{\mathbf{u}_k(\mathbf{p})\}_{\mathbf{p} \in S} \in \mathbb{R}^{|S| \times c} = \mathbb{R}^n = \mathcal{X}$ corresponding to the system states \mathbf{u}_k . In this section, we review some prior approaches, including the classical linear ROMs and DL-based ROMs usually implemented via AutoEncoder structures.

3.1.1. Classical linear ROMs

Linear Reduced-Order Models (ROMs) for representing system states have been extensively studied in the literature. Here, the word “linear” means that by fixing a collection of basis vectors $\{\boldsymbol{\phi}_j\}_j$ of the underlying functional space, the snapshot \mathbf{x} for each fixed time step is approximated as a linear combination of the basis vectors

$$\mathbf{x} \approx \sum_{j=1}^m c_j \boldsymbol{\phi}_j, \tag{9}$$

where c_j denotes the coefficient of the j -th basis vector and m is the reduced dimension. In this context, the snapshot \mathbf{x} is encoded as a lower-dimensional vector $\mathbf{c} = (c_1, \dots, c_m)^\top$. We only review some classical approaches here and choose not to dive deeply since they are not the focus of this work.

Proper Orthogonal Decomposition (POD), more specifically, the snapshot POD [50,72], is a well-known procedure to extract an optimal basis from an ensemble of observations, which is determined by minimizing the empirical average of the associated projection error for each snapshot. For simplicity we introduce it in the context of scalar fields: given a collection of numerical and/or experimental observations $\{\mathbf{u}_k : \mathbb{R}^d \rightarrow \mathbb{R}\}_{k=1}^K \subseteq L^2$ with temporal index $k = 1, 2, \dots, K$, POD solves the optimization problem

$$\boldsymbol{\phi} = \underset{\|\boldsymbol{\phi}\|_{L^2}=1}{\operatorname{argmin}} \frac{1}{K} \sum_{k=1}^K |\langle \mathbf{u}_k, \boldsymbol{\phi} \rangle|^2, \tag{10}$$

where $\langle \cdot, \cdot \rangle$ stands for the inner product on the L^2 space. A necessary condition for (10) can be derived by calculus of variations, which leads to the following eigenvalue problem

$$\frac{1}{K} \sum_{k=1}^K \mathbf{u}_k(s) \int \mathbf{u}_k(s') \boldsymbol{\phi}(s') ds' = \lambda \boldsymbol{\phi}(s), \tag{11}$$

and it remains to use Fourier modes [10] or apply eigenvalue decompositions to the correlation matrix of the snapshots in a discrete sense [22] to derive the solution. Once the mapping $\boldsymbol{\phi}$ is achieved, $\boldsymbol{\phi}_j$ is obtained by evaluating on the same sampling set $S \subseteq \mathbb{R}^d$ for each solution $\boldsymbol{\phi}$.

Dynamic Mode Decomposition (DMD) is another reduction approach based on the linearity assumption that

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k, \quad k = 0, 1, \dots, \quad (12)$$

where \mathbf{A} is a time-invariant linear operator. Such model reduction technique makes use of temporal correlations of the underlying dynamics, and the assumption (12) amounts to a linear tangent approximation for non-linear dynamics [69]. By denoting $\mathbf{B}_1 = (\mathbf{x}_0, \dots, \mathbf{x}_{K-1})$ and $\mathbf{B}_2 = (\mathbf{x}_1, \dots, \mathbf{x}_K)$, the singular-value decomposition (SVD) of $\mathbf{B}_1 = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ helps to derive the dynamic modes $\phi_j = \mathbf{U}\mathbf{v}_j$, where \mathbf{v}_j is the j -th eigenvector of the matrix $\mathbf{U}^T\mathbf{B}_2\mathbf{V}\mathbf{\Sigma}^{-1}$.

For linear ROMs, each snapshot \mathbf{x} is linearly parameterized in the latent space \mathcal{Z} as the coefficient vector \mathbf{c} . Taking POD as an example, the corresponding encoder and decoder can be formulated as

$$\mathcal{E}^{\text{POD}}(\mathbf{x}) = \mathbf{V}\mathbf{x}, \quad \mathcal{D}^{\text{POD}}(\mathbf{z}) = \mathbf{V}^T\mathbf{z}, \quad (13)$$

where \mathbf{V} has its rows being the eigenvectors or the so-called dominant modes of the correlation matrix. Such linear ROMs are not only straightforward to implement but also offer simple linear bijective mappings between \mathcal{Z} and \mathcal{X}' , which simplify the derivation of the latent dynamics on \mathcal{Z} as long as the physical dynamics are provided. However, the reconstruction accuracy is heavily dependent on the number of basis functions, and as the system grows in complexity and non-linearity, many more basis functions are required to capture the small-scale structures, which diminishes the intended benefit of dimensionality reduction.

3.1.2. AutoEncoders

AutoEncoder (AE), consisting of an encoder and a decoder, is a common architecture in deep learning and serves as a viable non-linear method for dimensionality reduction. With the encoder and the decoder

$$\mathcal{E}^{\text{AE}}(\mathbf{x}) = f_\theta(\mathbf{x}), \quad \mathcal{D}^{\text{AE}}(\mathbf{z}) = g_\varphi(\mathbf{z}) \quad (14)$$

implemented via neural networks, AutoEncoder is capable of extracting lower-dimensional latent representations from high-dimensional input data. Usually, the encoder progressively decreases the data dimension with its multi-layered structure, and by mirroring the encoder's architecture, the decoder takes the encoded representations and then reconstructs the original input data.

The objective of an AutoEncoder is to minimize the empirical reconstruction error by employing a loss function to measure the divergence between the input and the reconstructed data. The optimization

$$\theta, \varphi = \underset{\theta, \varphi}{\operatorname{argmin}} \sum_{k=1}^K \mathcal{L}(\mathbf{x}_k, g_\varphi(f_\theta(\mathbf{x}_k))) \quad (15)$$

promotes the learning of compressed meaningful data representation, where the function \mathcal{L} measures the discrepancy between the inputs \mathbf{x}_k and the reconstructed states $g_\varphi(f_\theta(\mathbf{x}_k))$. Once both the encoder f_θ and the decoder g_φ are trained, $\mathbf{z}_k = f_\theta(\mathbf{x}_k)$ is the corresponding latent representation of \mathbf{x}_k , and the decoder g_φ can be used to reconstruct the original data \mathbf{x}_k from \mathbf{z}_k conversely.

Compared with classical linear ROMs, AutoEncoders are capable of capturing non-linear relationships in the data, thereby allowing for the possibility of reducing complex dynamical system in physical space to simpler latent dynamics. Nonetheless, several limitations persist. Unlike classical linear ROMs, AutoEncoders do not offer an optimal representation since the minimizer is not guaranteed to be obtained. They may also encounter scalability challenges when dealing with high-dimensional problems as the number of grid points will increase exponentially concerning the state dimension and resolutions. Such complexity could lead to difficulties in designing effective structures. Moreover, AutoEncoders are generally restricted to operating on inputs with fixed structural attributes, limiting their utility in real-world applications that frequently involve partially missing data, ungridded data, or data with varying resolutions.

3.2. Evolution of latent dynamics

Establishing a suitable surrogate model for the evolution of latent dynamics forms another crucial part of the LA framework, which not only accounts for the system's inherent non-linearity but also incorporates the auto-regressive nature of the latent dynamics evolution, regardless of whether or not the explicit form of the physical dynamics is provided. We simplify our discussion by assuming time-invariant latent dynamics, which is approximated by a parameterized surrogate model \mathcal{G}_Ψ such that

$$\mathbf{z}_k \mapsto \mathcal{G}_\Psi(\mathbf{z}_k) \approx \mathbf{z}_{k+1}, \quad k = 0, 1, \dots. \quad (16)$$

The goal is to minimize the discrepancies between subsequent latent states, \mathbf{z}_{k+1} and $\mathcal{G}_\Psi(\mathbf{z}_k)$. While linear regression serves as a straightforward method for constructing the latent surrogate model, its performance is limited in describing highly non-linear dynamics. As an alternative, the neural network offers a common data-driven approach to capturing the non-linear relationships in the data and approximating the latent dynamics.

The choices for the structure of \mathcal{G}_Ψ can be quite diverse, depending on the specific characteristics of the system state's temporal dynamics. For instance, Long Short-Term Memory (LSTM) [37] and Gated Recurrent Units (GRUs) [21] networks are specially designed for sequential data, whose structures enable the network to remember or forget past states based on the context, and the work of [59,20] has followed such way. ResNet [35] is another approach to modeling the sequential data, which coincides with the

forward Euler method as a classical numerical integrator for solving differential equations. Embedding the physical dynamics into a latent space naturally leads to latent dynamics

$$\frac{dz}{dt} = F(z) \quad (17)$$

with an unknown operator F . The equation can then be discretized as

$$z_{k+1} \approx z_k + F(z_k)\Delta t \approx z_k + f_\psi(z_k)\Delta t =: \mathcal{G}_\psi(z_k), \quad (18)$$

where a neural network f_ψ is employed to approximate the unknown F , and the trainable parameters are denoted by ψ . ResNet-based architectures have been adopted to learn the latent dynamics [62] via ReZero [7] to enhance the performance. However, such discretization methods are subject to temporal resolution and step size constraints, and applying them directly to real-world assimilation deployment with varying time steps may introduce additional errors.

3.3. Assimilation in latent spaces

Both the LA frameworks and classical DA frameworks require observed data $\{y_k\}_k$, observation operators $\{\mathcal{H}_k\}_k$ for each time step and an initial background estimate x_0 for start-up. They also need certain noise modeling mechanism, which plays a crucial role in tracking error evolution and subsequently deriving the assimilated state. The main difference is that LA can learn the latent dynamics from a historical dataset, and thus more suitable when the physical dynamics are costly or unavailable. Once the encoder-decoder mappings \mathcal{E} , \mathcal{D} and the surrogate model \mathcal{G} have been trained, a similar assimilation routine can be executed on the latent spaces. Next, we show the procedure of LA by taking the classical Ensemble Kalman Filter (EnKF) as an example.

First, we establish the encoder $\mathcal{E} : \mathcal{X}' \rightarrow \mathcal{Z}$ and the decoder $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{X}'$ according to equations such as (13)(14). Concurrently, a surrogate model $\mathcal{G} : \mathcal{Z} \rightarrow \mathcal{Z}$ gets trained as described in Section 3.2 such that

$$\mathcal{D} \circ \mathcal{E} \approx \text{identity mapping on } \mathcal{X}'; \quad \mathcal{D} \circ \mathcal{G} \circ \mathcal{E} \approx \mathcal{M}. \quad (19)$$

In other words, the encoder \mathcal{E} and the decoder \mathcal{D} are nearly inverse mappings such that the diagram in Fig. 2b commutes up to a small error. After that, the encoder gives the latent background estimate by defining $z_0^b = \mathcal{E}(x_0^b)$, which is then used to generate ensemble members for the initial state $\{z_0^{a,j}\}_j$. Here, we denote the ensemble index by j . Applying perturbed forward propagation in the latent space gives

$$z_k^{b,j} = \mathcal{G}_k(z_{k-1}^{a,j}) + \epsilon_k^{L,j}, \quad \epsilon_k^{L,j} \sim \mathcal{N}(0, \Sigma_k^L), \quad (20)$$

where Σ_k^L indicates the covariance matrix of the error for the latent surrogate model S_k at the k -th step. As for the analysis phase, the Kalman gain matrix can be calculated via

$$\mathbf{K}_k = \mathbf{P}_k^b \mathbf{H}_k^\top (\mathbf{H}_k \mathbf{P}_k^b \mathbf{H}_k^\top + \Sigma_k^o)^{-1}. \quad (21)$$

The matrix \mathbf{P}_k^b is the covariance matrix for the background ensembles $\{z_k^{b,j}\}_j$, and the matrix \mathbf{H}_k is the tangent linear model corresponding with the latent observation operator $\mathcal{H}'_k = \mathcal{H}_k \circ \mathcal{D}$ mapping from the latent space to the observation space. Finally, the analyzed ensembles can be updated via

$$z_k^{a,j} = z_k^{b,j} + \mathbf{K}_k \mathbf{d}_k^j, \quad (22)$$

where the innovation vector for the j -th ensemble is computed as

$$\mathbf{d}_k^j = y_k^o + \epsilon_k^{o,j} - \mathcal{H}'_k(z_k^{b,j}), \quad \epsilon_k^{o,j} \sim \mathcal{N}(0, \Sigma_k^o), \quad (23)$$

with the observation data and the covariance matrix denoted by y_k^o and Σ_k^o , respectively. A simple decoding mapping

$$x_k^a = \mathcal{D}(z_k^a) \quad (24)$$

on the ensemble mean z_k^a can help obtain the analyzed physical state x_k^a whenever needed.

4. LAINR framework

The central contribution of our work lies in the introduction of the Latent Assimilation with Implicit Neural Representations (LAINR) framework, an extension of the previously proposed LA framework. In LAINR, we borrow the idea of Implicit Neural Representations (INRs) to take advantage of its flexibility. More precisely, we establish the Spherical Implicit Neural Representations (SINRs) utilizing the spherical harmonics defined on the 2D sphere to build the encoder-decoder mappings. Besides, LAINR incorporates Neural Ordinary Differential Equations (Neural ODEs) to model the latent dynamics as a continuous-time system. Importantly, the framework also offers reliable uncertainty estimations for the trained networks, which are then integrated into the assimilation cycle. The overall framework has been depicted in Fig. 1.

Table 2
Comparison of different methods for latent embedding.

method	encoder \mathcal{E}	decoder \mathcal{D}	linearity	flexibility
POD	$\mathbf{x} \mapsto \mathbf{V}^T \mathbf{x}$	$\mathbf{z} \mapsto \mathbf{V} \mathbf{z}$	linear	fixed grid
AE	$\mathbf{x} \mapsto f_\theta(\mathbf{x})$	$\mathbf{z} \mapsto g_\varphi(\mathbf{z})$	non-linear	fixed grid
INR	$\mathbf{x} \mapsto \operatorname{argmin}_{\mathbf{z}} \ \mathbf{I}(\cdot; \mathbf{z}, \boldsymbol{\varphi}) - \mathbf{u}\ _S$	$\mathbf{z} \mapsto \{\mathbf{I}(\mathbf{p}; \mathbf{z}, \boldsymbol{\varphi})\}_{\mathbf{p} \in S}$	non-linear	mesh-free

LAINR inherits the advantages of the LA framework, which allows for the integration of data-driven encoder-decoder pairs and latent surrogate models that work seamlessly with existing well-developed DA algorithms. No explicit formula for the physical dynamics is required since the latent surrogate model is also data-driven and trained from scratch. Therefore, the framework supports the application to a wide range of problems and provides versatility in assimilating complex system states. Moreover, LAINR is much more flexible compared with the previous AutoEncoder-based LA framework in that SINRs approximate the underlying physical fields with continuous mappings and thus can handle unstructured data more naturally. The continuous modeling of the latent dynamics allows for irregular time steps, which provides a more suitable tool for practical real-world applications.

4.1. Spherical implicit neural representations

4.1.1. Implicit neural representations

Implicit Neural Representations (INRs) adopt a distinct approach to data modeling that deviates from traditional structure-based methods such as the AutoEncoder approach described earlier in Section 3.1.2. Rather than capturing state features on discrete positions, INRs operate on coordinate-based information and model the states of interest (signals or fields) as continuous mappings, which provides a substantial advantage when working with continuous states that are sampled only at discrete grid points or mesh intersections.

The core idea of INR is to use a neural network $\mathbf{I}(\cdot; \mathbf{z}; \boldsymbol{\varphi})$ to approximate a given system state $\mathbf{u} : \mathbb{R}^d \rightarrow \mathbb{R}^c$, where \mathbf{z} stands for the corresponding latent representations and $\boldsymbol{\varphi}$ symbolizes the trainable network parameters. The objective is to minimize the discrepancy between $\mathbf{I}(\cdot; \mathbf{z}; \boldsymbol{\varphi})$ and \mathbf{u} , measured with an appropriate measure μ defined on a predefined sampling set S . The network parameters $\boldsymbol{\varphi}$ can be interpreted as the consistent state characteristics and thus remain invariant across the dataset, which is to say, all the system states of the same dataset are recovered in the form of $\{\mathbf{I}(\cdot; \mathbf{z}_k; \boldsymbol{\varphi})\}_{k=1}^K$ as mappings, and the latent representation \mathbf{z}_k contains the temporal variations among different system states.

The training of INRs is straightforward. Let

$$\mathbf{x}_k = \{\mathbf{u}_k(\mathbf{p})\}_{\mathbf{p} \in S_k}, k = 1, 2, \dots, K \quad (25)$$

be a sequence of snapshots we wish to encode, where each \mathbf{x}_k is made up of samples for the system state \mathbf{u}_k on the finite sampling set S_k at the k -th step. As an example, one may choose the L^2 norm with the measure μ_k being a sum of Dirac delta functions $\mu_k = |S_k|^{-1} \sum_{\mathbf{p} \in S_k} \delta_{\mathbf{p}}$, which leads to a minimization problem for the function

$$\operatorname{argmin}_{\mathbf{z}_k, \boldsymbol{\varphi}} \sum_{k=1}^K \|\mathbf{I}(\cdot; \mathbf{z}_k; \boldsymbol{\varphi}) - \mathbf{u}\|_{L^2(\mu_k)}^2 = \sum_{k=1}^K |S_k|^{-1} \sum_{\mathbf{p} \in S_k} |\mathbf{I}(\mathbf{p}; \mathbf{z}_k; \boldsymbol{\varphi}) - \mathbf{u}(\mathbf{p})|^2. \quad (26)$$

After the training process, any latent representation \mathbf{z} can be decoded as the system state evaluated on an arbitrary sampling set S as

$$\mathcal{D}^{\text{INR}}(\mathbf{z}) = \mathbf{x} = \{\mathbf{I}(\mathbf{p}; \mathbf{z}; \boldsymbol{\varphi})\}_{\mathbf{p} \in S}. \quad (27)$$

To obtain the corresponding latent representation \mathbf{z}' for a given possibly unseen observations \mathbf{x}' on S' , we need to define the encoder mapping, which is much more challenging than the decoder given above. The encoding from $\mathbf{x}' = \{\mathbf{u}'(\mathbf{p})\}_{\mathbf{p} \in S'}$ to \mathbf{z}' is not explicitly defined and involves solving an optimization problem

$$\mathcal{E}^{\text{INR}}(\mathbf{x}') = \mathcal{E}^{\text{INR}}\left(\{\mathbf{u}'(\mathbf{p})\}_{\mathbf{p} \in S'}\right) = \mathbf{z}' = \operatorname{argmin}_{\mathbf{z}} \|\mathbf{I}(\cdot; \mathbf{z}; \boldsymbol{\varphi}) - \mathbf{u}'\|_{S'}. \quad (28)$$

The high cost of the optimization problem during encoding is a major concern in most INR-based methodologies. However, our LAINR framework remains efficient since the optimization needs to be executed only for the background estimate as explained later. Meanwhile, the flexibility that the sampling set S in the training and the deployment process may be different and unstructured is particularly suitable for data assimilation scenarios since the real observation data are often spatially and temporally unevenly distributed. Table 2 compares INR with classical ROMs.

4.1.2. Spherical filters

Recall that a classical L -layer deep neural network f is defined by the following update formula

$$\begin{aligned} \gamma^{(0)} &= \mathbf{x}, \\ \gamma^{(l)} &= \sigma(\mathbf{W}^{(l)}\gamma^{(l-1)} + \mathbf{b}^{(l)}), \quad l = 1, 2, \dots, L, \\ \mathbf{f}(\mathbf{x}) &= \gamma^{(L)}, \end{aligned} \quad (29)$$

where each layer produces an output by applying an affine mapping followed by a non-linear activation function. Inspired by the success of SIREN [73] and Fourier Feature Networks [75] which make use of the trigonometric functions as the activation function and the position embeddings, Multiplicative Filter Networks (MFNs) [29] propose the multiplicative updates

$$\begin{aligned} \gamma^{(0)} &= \mathbf{g}_0(\mathbf{x}), \\ \gamma^{(l)} &= (\mathbf{W}^{(l)}\gamma^{(l-1)} + \mathbf{b}^{(l)}) \odot \mathbf{g}_l(\mathbf{x}), \quad l = 1, 2, \dots, L, \\ \mathbf{f}(\mathbf{x}) &= \mathbf{W}^{(L)}\gamma^{(L)} + \mathbf{b}^{(L)} \end{aligned} \quad (30)$$

to represent the network as linear combinations of basis functions. Here, the symbol \odot stands for the element-wise multiplication. For instance, by implementing the filters $\mathbf{g}_l(\mathbf{x})$ as sinusoidal functions $\mathbf{g}_l(\mathbf{x}) = \sin(\boldsymbol{\omega}_l \mathbf{x} + \boldsymbol{\phi}_l)$, which is referred to as the ‘‘Fourier Network’’, the network output $\mathbf{f}(\mathbf{x})$ can always be expanded as a linear combination of sinusoidal bases.

Theorem 4.1. *Each coordinate of the output of (30) is given by a finite linear combination of sinusoidal bases*

$$f_j(\mathbf{x}) = \sum_{i=1}^T \bar{\alpha}_i^{(j)} \sin(\bar{\boldsymbol{\omega}}_i^{(j)} \mathbf{x} + \bar{\boldsymbol{\phi}}_i^{(j)}) + \bar{\beta}^{(j)} \quad (31)$$

for some coefficients $\bar{\alpha}_i^{(j)}$, $\bar{\boldsymbol{\omega}}_i^{(j)}$, $\bar{\boldsymbol{\phi}}_i^{(j)}$ and $\bar{\beta}^{(j)}$ only dependent on the network parameters [29].

Nevertheless, no convergence guarantee exists for the MFN structure with sinusoidal filters since no explicit constraint is imposed on the frequencies $\bar{\boldsymbol{\omega}}_i^{(j)}$ and thus the learned sinusoidal functions are not necessarily orthogonal. Meanwhile, MFN does not take advantage of the intrinsic geometry of the input space and always assumes that the input data can be extended to the whole Euclidean space.

Consequently, to deal with the input data lying on the 2D sphere, which is typically encountered in meteorological applications, we propose the spherical filter as follows

Definition 4.2 (spherical filters). Let $\mathbf{p} = (\theta, \phi) \in \mathbb{S}^2$ be a point of the 2D sphere. A spherical filter $\mathbf{g} : \mathbb{S}^2 \rightarrow \mathbb{R}^h$ of degree D , shift s and hidden dimension h is defined as

$$\mathbf{g}(\mathbf{p}) = \Xi \left(Y_{D+s}^{-D}(\mathbf{p}), Y_{D-1+s}^{-(D-1)}(\mathbf{p}), \dots, Y_{0+s}^0(\mathbf{p}), \dots, Y_{D-1+s}^{D-1}(\mathbf{p}), Y_{D+s}^D(\mathbf{p}) \right)^\top, \quad \mathbf{p} \in \mathbb{S}^2, \quad (32)$$

where both D and s are non-negative integers and $\Xi \in \mathbb{R}^{h \times (2D+1)}$ is a coefficient matrix with real entries. The function Y_ℓ^m is the real spherical harmonics derived from the complex version

$$Y_{\ell,m}(\theta, \varphi) = \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos \theta) e^{im\varphi} \quad (33)$$

via

$$Y_\ell^m = \begin{cases} Y_{\ell,0}, & m = 0, \\ \frac{1}{\sqrt{2}} (Y_{\ell,m} - (-1)^m Y_{\ell,-m}), & m < 0, \\ \frac{1}{\sqrt{2}} (Y_{\ell,-m} + (-1)^m Y_{\ell,m}), & m > 0. \end{cases} \quad (34)$$

P_ℓ^m is the associated Legendre polynomial with the convention

$$P_\ell^{-m} = (-1)^m \frac{(\ell-m)!}{(\ell+m)!} P_\ell^m, \quad m > 0, \quad (35)$$

and θ and ϕ are the polar and azimuthal angles of the spherical coordinate, respectively.

Remark. The spherical harmonics are orthogonal basis functions on \mathbb{S}^2 , making them a natural choice for the basis filter functions. We anticipate that this orthogonality, which has also been proved later in Theorem 4.9, will aid in establishing the convergence of our SINR structure. It's worth noting that in the MFN structure (30) with sinusoidal filters, both the frequencies and phase shifts of the sinusoidal functions may vary. However, we opt to fix the spherical harmonics and only train the coefficient matrix Ξ in

(32) in our proposed spherical filters. The main reason is that prior works typically handle data lying on a certain Euclidean space, where sinusoidal functions can be easily generalized to having arbitrary real frequencies and phase shifts, especially when boundary behaviors are of little concern. Conversely, for physical states defined on the sphere \mathbb{S}^2 , an intrinsic periodical boundary condition is introduced by the spherical geometry. Under such circumstances, the collection of spherical harmonics $\{Y_{\ell,m}\}_{\ell,m}$ can hardly be indexed by real numbers for training, as the integer parameters ℓ, m originate from the Legendre polynomial degrees.

Next, we focus on deriving certain convergence results for our SINR model. Fortunately, similar to the multiplication formula of trigonometric functions, any product of two spherical harmonics is also a finite sum of spherical harmonics.

Lemma 4.3. For any $|m_1| \leq \ell_1$ and $|m_2| \leq \ell_2$, we have

$$Y_{\ell_1}^{m_1} Y_{\ell_2}^{m_2} = \sum_{\ell=0}^{\ell_1+\ell_2} \sum_{|m| \leq \ell} C_{\ell_1, \ell_2, \ell}^{m_1, m_2, m} Y_{\ell}^m, \tag{36}$$

where the real coefficient $C_{\ell_1, \ell_2, \ell}^{m_1, m_2, m}$ only depends on $\ell_1, \ell_2, \ell, m_1, m_2$ and m .

Proof. By the definition of the real spherical harmonics (34), we have for $m > 0$

$$\begin{aligned} Y_{\ell}^m &= Y_{\ell,0}, \\ Y_{\ell}^m &= \sqrt{2}(-1)^m \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_{\ell}^m \cos m\varphi, \\ Y_{\ell}^{-m} &= \sqrt{2}(-1)^m \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_{\ell}^m \sin m\varphi. \end{aligned} \tag{37}$$

The collection $\{Y_{\ell}^m\}_{|m| \leq \ell}$ forms a basis of the space of real spherical harmonics of degree ℓ by [6]. As a result, the multiplication $Y_{\ell_1}^{m_1} Y_{\ell_2}^{m_2}$ of degree $(\ell_1 + \ell_2)$ is a finite linear combination of real spherical harmonics of degree no more than $(\ell_1 + \ell_2)$. \square

Different from the original MFN architecture, we add a bypass connection from each layer to the output layer in order to retain the spherical harmonics terms for lower degree ℓ . Consequently, we propose a new variant of INRs specially designed for the 2D sphere, which is referred to as the Spherical Implicit Neural Representations (SINRs), with the detailed update formulas given below.

Definition 4.4 (Spherical Implicit Neural Representations (SINRs)). A SINR $f : \mathbb{S}^2 \rightarrow \mathbb{R}^c$ with L layers is given by

$$\begin{aligned} \gamma^{(0)} &= \mathbf{g}_0(\mathbf{p}), \quad \mathbf{p} \in \mathbb{S}^2, \\ \gamma^{(l)} &= (\mathbf{W}^{(l)} \gamma^{(l-1)} + \mathbf{b}^{(l)}) \odot \mathbf{g}_l(\mathbf{p}), \quad l = 1, 2, \dots, L, \\ f(\mathbf{p}) &= \sum_{l=0}^L (\tilde{\mathbf{W}}^{(l)} \gamma^{(l)} + \tilde{\mathbf{b}}^{(l)}), \end{aligned} \tag{38}$$

where for each layer $\mathbf{g}_l(\mathbf{p})$ is a spherical filter of degree D , shift l and hidden dimension h following the definition of (32).

Note that for each layer, the weights and biases $\mathbf{W}^{(l)}$, $\mathbf{b}^{(l)}$, $\tilde{\mathbf{W}}^{(l)}$ and $\tilde{\mathbf{b}}^{(l)}$, along with the coefficient matrix of the spherical filter $\mathbf{g}_l(\mathbf{p})$, contribute to the trainable network parameters, while the number of layers L , the degree D and the hidden dimension h are hyperparameters. Similar to the Fourier Network introduced by MFN, the output of SINR can also be expanded as a linear combination of spherical harmonics.

Proposition 4.5 (expansion of SINRs). Each coordinate of the output of (38) is given by a finite linear combination of real spherical harmonics. Formally, we have

$$f_j(\mathbf{p}) = \sum_{\ell=0}^{(2D+L)(L+1)/2} \sum_{|m| \leq \ell} \tilde{\alpha}_{\ell,m}^{(j)} Y_{\ell}^m(\mathbf{p}) + \tilde{\beta}^{(j)} \tag{39}$$

for some coefficients $\tilde{\alpha}_{\ell,m}^{(j)}$ and $\tilde{\beta}^{(j)}$ only dependent on the network parameters.

Proof. It suffices to show that for each layer $\gamma^{(l)}$, all of its coordinates can be expressed in a similar way

$$\gamma_j^{(l)} = \sum_{\ell=0}^{(2D+l)(l+1)/2} \sum_{|m| \leq \ell} \alpha_{\ell,m}^{(j,l)} Y_{\ell}^m(\mathbf{p}) + \beta^{(j,l)}. \tag{40}$$

We prove it by induction. For $l = 0$, the property is trivial by the definition of $\mathbf{g}_0(\mathbf{p})$. Suppose that (40) holds for l , it immediately follows by the update

$$\boldsymbol{\gamma}^{(l+1)} = (\mathbf{W}^{(l+1)}\boldsymbol{\gamma}^{(l)} + \mathbf{b}^{(l+1)}) \odot \mathbf{g}_l(\mathbf{p}) \quad (41)$$

that all of its coordinates can be expressed by a linear combination of products of spherical harmonics $Y_{\ell_1}^{m_1} Y_{\ell_2}^{m_2}$ where

$$|m_1| \leq \ell_1 \leq \frac{1}{2}(2D+l)(l+1), \quad |m_2| \leq \ell_2 \leq D+l. \quad (42)$$

By the multiplication formula of spherical harmonics (Lemma 4.3), the products can be expressed by linear combinations of spherical harmonics Y_{ℓ}^m with

$$\ell \leq \frac{1}{2}(2D+l)(l+1) + D+l = \frac{1}{2}(2D+l+1)(l+2), \quad (43)$$

which implies that the property holds for $l+1$. \square

Conversely, the linear subspace spanned by the spherical harmonics can be represented by a SINR with the following proposition.

Proposition 4.6 (representability of SINRs). *Any linear combination of spherical harmonics Y_{ℓ}^m with $|m| \leq D$, $0 \leq \ell - |m| \leq L$ can be represented accurately by a SINR (38) with L layers and degree D (32) as long as the hidden dimension $h \geq 2D+1$.*

Proof. By setting $\mathbf{W}^{(l)} = \mathbf{0}$, $\tilde{\mathbf{b}}^{(l)} = \mathbf{0}$ and $\mathbf{b}^{(l)} = \mathbf{1}$ for each layer, the output of (38) is given by

$$\begin{aligned} \mathbf{f}(\mathbf{p}) &= \sum_{l=0}^L \tilde{\mathbf{W}}^{(l)} \boldsymbol{\gamma}^{(l)} = \sum_{l=0}^L \tilde{\mathbf{W}}^{(l)} \mathbf{g}_l(\mathbf{p}) \\ &= \sum_{l=0}^L \tilde{\mathbf{W}}^{(l)} \Xi^{(l)} \left(Y_{D+l}^{-D}(\mathbf{p}), Y_{D+l}^{-(D-1)}(\mathbf{p}), \dots, Y_{0+l}^0(\mathbf{p}), \dots, Y_{D+l}^{D-1}(\mathbf{p}), Y_{D+l}^D(\mathbf{p}) \right)^{\top}. \end{aligned} \quad (44)$$

For any linear combination of spherical harmonics Y_{ℓ}^m with $|m| \leq D$, $0 \leq \ell - |m| \leq L$ written as

$$\sum_{m=-D}^D \sum_{\ell=|m|}^{|m|+L} a_{\ell,m} Y_{\ell}^m(\mathbf{p}) = \sum_{l=0}^L \sum_{m=-D}^D a_{|m|+l,m} Y_{|m|+l}^m(\mathbf{p}), \quad (45)$$

clearly there exist infinitely pairs of $(\tilde{\mathbf{W}}^{(l)}, \Xi^{(l)})$ such that

$$\tilde{\mathbf{W}}^{(l)} \Xi^{(l)} = \text{diag} \{ a_{D+l,-D}, a_{D+l,-D+1}, \dots, a_{D+l,0}, \dots, a_{D+l,D-1}, a_{D+l,D} \} \quad (46)$$

if $h \geq 2D+1$, which completes our proof. \square

Remark. In fact, the linear combinations mentioned in Proposition 4.6 only make up a very small proportion of the set of all possible SINRs. Readers may refer to Appendix A for another more general class of functions that can be represented by SINRs. We would like to emphasize that one of the advantages of our SINR model is the possibility of the explorations of larger degree ℓ with small degrees by using the graded multiplication trick, and such a feature is beneficial for a data-driven model to learn the representation of the high-frequency components of the input data.

Corollary 4.7. *Denote by S_T the set of all linear combinations of spherical harmonics Y_{ℓ}^m with $(m, \ell) \in T$. Let*

$$\begin{aligned} T &= \{(m, \ell) \mid |m| \leq D, 0 \leq \ell - |m| \leq L\}, \\ T' &= \left\{ (m, \ell) \mid |m| \leq \frac{1}{2}(2D+L)(L+1) \right\}. \end{aligned} \quad (47)$$

Then the set of all possible SINRs with L layers, hidden dimensions $h \geq 2D+1$ and degree D (32) denoted by S_{SINR} satisfies

$$S_T \subseteq S_{\text{SINR}} \subseteq S_{T'}. \quad (48)$$

Proof. The conclusion follows immediately from Proposition 4.5 and Proposition 4.6. Readers may refer to Fig. 3 for visualization of the inclusion relationships. \square

By combining the expansion of the SINR architecture and its representability, we can conclude that the SINR architecture acts as a universal approximator on the 2D-sphere with the aid of the linear subspace spanned by the spherical harmonics.

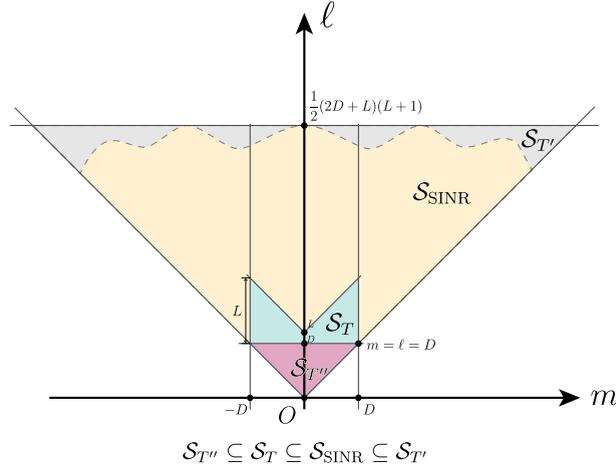


Fig. 3. Illustration of the inclusion relationships projected into the spectral modes (m, ℓ) . Note that the range for S_{SINR} is only for reference.

Lemma 4.8 (convergence of Fourier-Laplacian series). Let \mathcal{P}_d be the projection operator onto the space spanned by the spherical harmonics Y_ℓ^m with $|m| \leq \ell = d$, then the Fourier-Laplacian series of $f \in L^2(\mathbb{S}^2)$

$$\sum_{d=0}^N \mathcal{P}_d f \rightarrow f \text{ in } L^2(\mathbb{S}^2) \text{ as } N \rightarrow +\infty, \tag{49}$$

and the convergence is uniform if f is smooth enough [6].

Theorem 4.9 (convergence of SINRs). For any $f \in L^2(\mathbb{S}^2)$,

$$\lim_{D, L \rightarrow \infty} \inf_{f' \in S_{SINR}} \|f' - f\|_{L^2(\mathbb{S}^2)} = 0. \tag{50}$$

Furthermore, if $f \in C^\infty(\mathbb{S}^2)$,

$$\lim_{D, L \rightarrow \infty} \inf_{f' \in S_{SINR}} \|f' - f\|_{C(\mathbb{S}^2)} = 0. \tag{51}$$

Proof. The proof is a direct consequence of the inclusion $S_{T''} \subseteq S_T \subseteq S_{SINR}$, where

$$T'' = \{(m, \ell) \mid |m| \leq \ell \leq \min(D, L)\} \tag{52}$$

by Corollary 4.7 and Lemma 4.8. \square

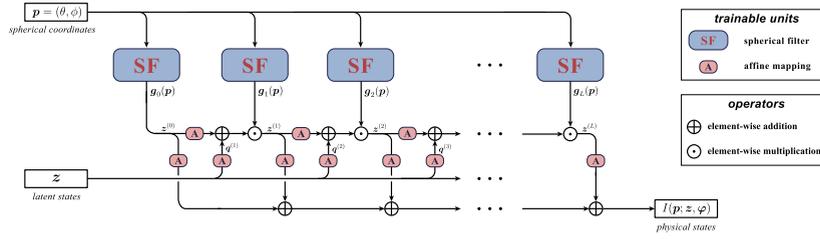
Remark. Qualitatively speaking, the rate of convergence is tied to the smoothness of the target function f . As the degree D and the number of layers L increase, the smoother the function is, the faster $S_{T''}$ converges as well as S_{SINR} . Readers may refer to [6] for detailed discussions. Generally, uniform convergence is expected since most natural physical fields exhibit smooth behaviors.

4.2. Modulation adjustment

Directly treating the network parameters $\mathbf{W}^{(l)}$, $\mathbf{b}^{(l)}$, $\tilde{\mathbf{W}}^{(l)}$ and $\tilde{\mathbf{b}}^{(l)}$ together with the scalar matrix of the spherical filter for each layer of our SINR architecture (38) as the latent representation \mathbf{z} can hardly achieve our purpose of efficient reduction of dimensions. Therefore, we propose to use the modulation adjustment, which is also a common technique of many existing INR networks. The idea is to use a shared base network across different snapshots to model universal physical structure, with modulations modeling the variation specific to individual time steps [61,53]. Experiments have indicated that a simple modulation shift [25] for each layer is sufficient for our tasks. The modulation adjustment for the SINR architecture (38) is

$$\begin{aligned} \gamma^{(0)} &= \mathbf{g}_0(\mathbf{p}), \quad \mathbf{p} \in \mathbb{S}^2, \\ \gamma^{(l)} &= (\mathbf{W}^{(l)} \gamma^{(l-1)} + \mathbf{b}^{(l)} + \mathbf{q}^{(l)}(\mathbf{z})) \odot \mathbf{g}_l(\mathbf{p}), \quad l = 1, 2, \dots, L, \\ I(\mathbf{p}; \mathbf{z}, \boldsymbol{\varphi}) &= \sum_{l=0}^L \left(\tilde{\mathbf{W}}^{(l)} \gamma^{(l)} + \tilde{\mathbf{b}}^{(l)} \right), \end{aligned} \tag{53}$$

SINR



Spherical Filter

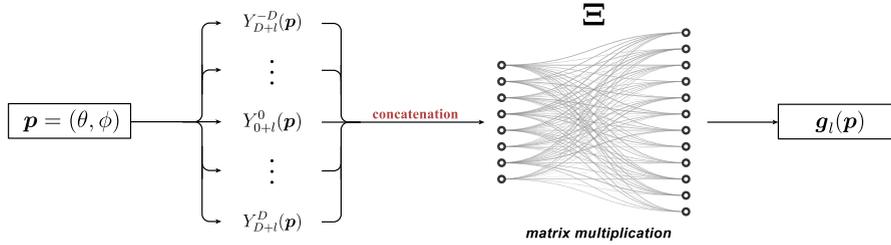


Fig. 4. Visualization of the SINR architecture with modulation adjustments and the spherical filters.

where $q^{(l)}(z)$ stands for a trainable affine mapping of the latent representation z for the l -th layer, and the parameters are not shared across different layers. The original weights and biases of the SINR structure as well as the trainable parameters of the modulation adjustment $q^{(l)}$ together make up the optimization variable φ in (26). Note that only the latent representation z varies across different snapshots, while the network parameters φ remain consistent. Therefore, the dimension of the latent space is reduced to the size of z . Fig. 4 visualizes the detailed structure of a SINR with modulation adjustments.

4.3. Continuous-time latent surrogate model

To address the challenges posed by temporally uneven distributions of observation data, we model the latent dynamics as a continuous-time dynamical system

$$\frac{dz}{dt} = F(z) \approx f_{\psi}(z) \tag{54}$$

with the set of trainable parameters denoted by ψ , and we have chosen the Neural ODE structure [17,18]

$$z_{k+1} = z_k + \int_{t_k}^{t_{k+1}} f_{\psi}(z) dt \tag{55}$$

for the implementation. It's worth noting that such structures are suitable for handling time-series data, as they offer a data-efficient way to parameterize a trajectory over time and naturally deal with arbitrary time steps. These advantages make Neural ODEs a fitting choice for modeling the latent dynamics in our LAINR framework in that the time steps are usually irregular for realistic assimilation processes.

One may have noticed that same as Section 3.2, we only consider learning the temporal latent dynamics that come from the original physical dynamics. Other constraints are set aside, such as boundary conditions, conservation laws and divergence-free conditions if they exist. One of the reasons is that we expect these additional constraints to have been implicitly encoded within the learned lower-dimensional submanifolds, which is encouraged by the optimization process requiring the encoder-decoder mapping to obtain high reconstruction accuracy at low cost. The other reason is that the trained latent surrogate model will be inserted into assimilation cycles afterward, and the defects of the predicted latent states may be corrected by the observation data to some extent.

4.4. Modeling of uncertainties

There are in general two types of uncertainty we need to model in the assimilation process. One is the prediction uncertainties of the surrogate model responsible for latent dynamics, which is inevitable due to the very nature of network training. The other is the uncertainty of the latent background estimate, which characterizes the latent uncertainty transformed from the physical background by the encoder mapping. We assume that the encoder mapping is perfect, which means that the optimization problem (28) is solved with a negligible margin of error. Experimental results have validated the practical effectiveness of such simplification.

4.4.1. Uncertainty estimator for the latent surrogate model

Upon successful training of the latent surrogate model, let $\hat{\mathbf{z}}_{k+1} = \mathcal{G}(\mathbf{z}_k)$ be the predicted latent state for the $(k+1)$ -th time step based on the k -th latent state \mathbf{z}_k . We model the prediction uncertainty as a Gaussian distribution. To put it formally,

$$\mathcal{G}(\mathbf{z}_k) = \hat{\mathbf{z}}_{k+1} = \mathbf{z}_{k+1} + \boldsymbol{\varepsilon}_{k+1}^L, \quad \boldsymbol{\varepsilon}_{k+1}^L \sim \mathcal{N}(0, \boldsymbol{\Sigma}^L), \quad (56)$$

where $\boldsymbol{\Sigma}^L$ is the covariance matrix of the uncertainty $\boldsymbol{\varepsilon}_{k+1}^L$ of the latent surrogate model \mathcal{G} , and we assume it to be time-independent.

We have to point out, however, that our choice of using Gaussian models is far from optimality. We only aim to argue for the necessity of such a mechanism in our LAINR framework, which is capable of estimating the uncertainty of the learned latent surrogate model and thus benefiting the assimilation process. The Gaussian error modeling for the surrogate model depends on the model's adequate representability for the latent dynamics, and we assume that the model error only results from unknown random forcing or random variation of parameters. Nonetheless, such an assumption does not necessarily hold in general. Literature such as [28] has employed specialized and advanced models to characterize the possible systematical biases of the surrogate model, which is not the main focus of our current work.

To train the uncertainty estimator, our LAINR framework assumes that the $\boldsymbol{\Sigma}^L$ is diagonal, which seems to be quite arbitrary, but experimental results have surprisingly shown its effectiveness (Section 5.2.4). We parameterize $(\boldsymbol{\Sigma}^L)^{1/2} = \exp \mathbf{D}$, and $\exp \mathbf{D}$ stands for the element-wise exponential of a diagonal matrix \mathbf{D} to ensure its positive-definiteness. For convenience, we refer to the matrix \mathbf{D} as our uncertainty estimator for the latent surrogate model hereafter. The matrix \mathbf{D} is initialized with zeros and then optimized during training, where the target loss function is derived as follows. With the independence assumption, the likelihood of the uncertainties $\{\boldsymbol{\xi}_k\}_k$ across the whole dataset reads

$$\mathbb{P}(\{\boldsymbol{\xi}_k\}_k) = \prod_k (2\pi)^{-m/2} \det(\boldsymbol{\Sigma}^L)^{-1/2} \exp\left(-\frac{1}{2} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1})^\top (\boldsymbol{\Sigma}^L)^{-1} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1})\right). \quad (57)$$

Consequently, we train the uncertainty estimator \mathbf{D} by minimizing the negative log-likelihood

$$\begin{aligned} \mathcal{L}_{\text{MLE}} &= \sum_k \frac{1}{2} \log \det(\boldsymbol{\Sigma}^L) + \frac{1}{2} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1})^\top (\boldsymbol{\Sigma}^L)^{-1} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1}) \\ &= \sum_k \log \det(\exp \mathbf{D}) + \frac{1}{2} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1})^\top (\exp \mathbf{D})^{-1} (\exp \mathbf{D})^{-1} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1}) \\ &= \sum_k \left(\sum_j \mathbf{D}_{jj} + \frac{1}{2} \left\| (\exp \mathbf{D})^{-1} (\mathcal{G}(\mathbf{z}_k) - \mathbf{z}_{k+1}) \right\|^2 \right) \end{aligned} \quad (58)$$

up to a constant by the method of maximum likelihood estimation (MLE), where the outer summation runs through all the latent representations \mathbf{z}_k , and the inner summation $\sum_j \mathbf{D}_{jj}$ represents the trace of \mathbf{D} .

4.4.2. Uncertainty of the latent background estimate

Consider a physical background estimate

$$\mathbf{x}_0^b = \mathbf{x}_0 + \boldsymbol{\varepsilon}_x^b, \quad \boldsymbol{\varepsilon}_x^b \sim \mathcal{N}(0, \boldsymbol{\Sigma}_x^b) \quad (59)$$

given as a Gaussian perturbation of the true physical state \mathbf{x}_0 . The encoder mapping \mathcal{E} transforms it into a latent background estimate

$$\mathbf{z}_0^b = \mathcal{E}(\mathbf{x}_0^b) = \mathcal{E}(\mathbf{x}_0 + \boldsymbol{\varepsilon}_x^b) = \mathcal{E}(\mathbf{x}_0) + \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \boldsymbol{\varepsilon}_x^b + o(\boldsymbol{\varepsilon}_x^b) = \mathbf{z}_0 + \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \boldsymbol{\varepsilon}_x^b + o(\boldsymbol{\varepsilon}_x^b), \quad (60)$$

where $\mathbf{J}_{\mathcal{E}}$ is the Jacobian matrix of \mathcal{E} . By omitting the high-order term, we make the following approximation

$$\mathbf{z}_0^b \approx \mathbf{z}_0 + \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \boldsymbol{\varepsilon}_x^b = \mathbf{z}_0 + \boldsymbol{\varepsilon}_z^b, \quad \boldsymbol{\varepsilon}_z^b \sim \mathcal{N}(0, \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \boldsymbol{\Sigma}_x^b \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0)^\top). \quad (61)$$

Since the encoder mapping \mathcal{E} is associated with an optimization process, no explicit formula for the Jacobian matrix $\mathbf{J}_{\mathcal{E}}$ is available. Therefore, we choose to make use of the decoder mapping \mathcal{D} in that the chain rule gives

$$\mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \mathbf{J}_{\mathcal{D}}(\mathbf{z}_0) = \mathbf{J}_{\mathcal{E}}(\mathcal{D}(\mathbf{z}_0)) \mathbf{J}_{\mathcal{D}}(\mathbf{z}_0) = \mathbf{J}_{\mathcal{E} \circ \mathcal{D}}(\mathbf{z}_0) = \mathbf{I}_m. \quad (62)$$

First, for any physical background estimate \mathbf{x}_0^b , LAINR calculates the latent background estimate $\mathbf{z}_0^b = \mathcal{E}(\mathbf{x}_0^b)$. It follows that the Jacobian $\mathbf{J}_{\mathcal{D}}(\mathbf{z}_0)$ is approximated by $\mathbf{J}_{\mathcal{D}}(\mathbf{z}_0^b)$ with the aid of automatical differentiation on the decoder mapping \mathcal{D} . The Moore-Penrose inverse of $\mathbf{J}_{\mathcal{D}}(\mathbf{z}_0^b)$ is used to approximate $\mathbf{J}_{\mathcal{E}}(\mathbf{x}_0)$ afterwards in that they are both the left inverse of $\mathbf{J}_{\mathcal{D}}(\mathbf{z}_0^b)$, which is then employed to construct the covariance matrix $\mathbf{J}_{\mathcal{E}}(\mathbf{x}_0) \boldsymbol{\Sigma}_x^b \mathbf{J}_{\mathcal{E}}(\mathbf{x}_0)^\top$ for $\boldsymbol{\varepsilon}_z^b$.

We need to address that the uncertainty of the latent background estimate is only used for start-up, and all the subsequent assimilation cycles are conducted within the latent space. The uncertainty $\boldsymbol{\Sigma}_x^b$ for the physical state is commonly accessible since the background state usually comes from the forecast starting from the past analysis state. Besides, the uncertainty $\boldsymbol{\Sigma}_x^b$ does not need to be very accurate since, in the context of data assimilation, it is the consistent observations at different time steps, rather than the initial conditions that calibrate the assimilated model state to keep evolving closely to the true state.

4.5. The proposed LAINR framework

LAINR shares some similarities with the LA framework described in Section 3.3. It requires observed data $\{y_k\}_k$, observation operators $\{H_k\}_k$ and a background estimate x_0 to initiate the assimilation process. In contrast to the LA framework, any unstructured historical dataset encoding the physical dynamics is sufficient for LAINR to obtain a latent surrogate model for the dynamics, and LAINR additionally provides the corresponding uncertainty estimations for the trained networks. The overall procedure for the deployment of our LAINR framework is summarized as follows.

- (i) Train the encoder-decoder mapping implemented via the SINR architecture as detailed in Section 4.1 with (26). The resulting encoder $\mathcal{E} : \mathcal{X}' \rightarrow \mathcal{Z}$ and decoder $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{X}'$ are given by (28)(27). Besides, the latent surrogate model $\mathcal{G} : \mathcal{Z} \rightarrow \mathcal{Z}$ is implemented via Neural ODE (Section 4.3) such that

$$D \circ \mathcal{E} \approx \text{the identity mapping on } \mathcal{X}', \quad D \circ \mathcal{G} \circ \mathcal{E} \approx \mathcal{M}. \quad (63)$$

In other words, the encoder \mathcal{E} and the decoder \mathcal{D} are inverse mappings such that the diagram in Fig. 2 (b) commutes up to a small error.

- (ii) Use the uncertainty estimator to evaluate the uncertainty covariance matrix Σ^L for the latent surrogate model \mathcal{G} . The uncertainty estimator is trained by the method of maximum likelihood estimation (MLE) using the latent representation $\{z_k\}_k$ of the training dataset. (Section 4.4.1)
- (iii) Determine the latent background estimate $z_0^b = \mathcal{E}(x_0^b)$ and calculate its associated uncertainty covariance matrix. (Section 4.4.2)
- (iv) Implement the assimilation process on the latent space with an existing well-developed DA algorithm, where \mathcal{G} serves as the latent forward propagation operator and $\mathcal{H}' = \mathcal{H} \circ \mathcal{D}$ acts as the observation operator. The observation noise of $\mathcal{H}' = \mathcal{H} \circ \mathcal{D}$ remains the same as that of \mathcal{H} . The uncertainty estimations are given in the previous two steps.
- (v) Retrieve the assimilated system states $\{x_k^a\}_k$ from their assimilated latent representations via the decoder $x_k^a = \mathcal{D}(z_k^a)$.

5. Experiments

To assess the capability of our proposed LAINR framework, we have conducted multiple experiments to make comparisons with existing works based on AutoEncoders. Additionally, we have studied the flexibility of our LAINR framework when handling various inconsistent data structures as well.

First, the test cases we have adopted in this paper are introduced, consisting of an ideal shallow-water dataset and a more realistic ERA5 reanalysis dataset to explore the model applicability. After that, the model configurations including the structures and the hyper-parameters along with the metrics for training and evaluation are exhibited, which are then followed by the details of the training process.

The detailed experiments are presented in Section 5.2, where we have compared our LAINR framework with the existing AutoEncoder-based methodologies. Both the training and assimilation processes are based on a regular latitude-longitude grid (Fig. 5a). The comparisons are made in terms of the reconstruction, the prediction as well as the assimilation accuracy. All the experimental results have demonstrated the superiority of our LAINR framework over the existing AutoEncoder-based methodologies.

Furthermore, LAINR has the flexibility to work with diverse types of sampling sets with little adjustment, and we have examined the flexibility of our LAINR framework in Section 5.3 by testing on a reconstruction task with only rare available data, and a zero-shot assimilation task where the training and testing datasets come from different grids (Fig. 5b). Our LAINR framework has shown its large potential in handling incomplete data and unseen sampling sets.

5.1. Test cases and model configurations

5.1.1. Test cases

We choose multi-dimensional partial differential equations (PDEs) as our test cases instead of chaotic ordinary differential equations (ODEs) such as the Lorenz-96 model adopted in the previous works [62,60] to study the performances of our LAINR framework. PDEs are more common in nature and technology, and they typically possess much more complex physical dynamics than ODEs. For this study, both the ideal shallow-water model as well as the more realistic ERA5 reanalysis data have been established as the test cases in this paper.

The first one is the ideal shallow-water model, which is a simplified version of the primitive equations of the atmosphere. The second one is the ERA5 reanalysis datasets, which are the latest global atmospheric reanalysis datasets produced by the European Centre for Medium-Range Weather Forecasts (ECMWF). A brief description of the two test cases is given below, and readers may refer to Appendix B for detailed dataset configurations.

It is also important to mention that similar to most of the previous DL-based assimilation frameworks such as [62,60,59,20], our LAINR framework and the AutoEncoder baselines do not assume any prior knowledge of the underlying dynamics. We only extract partial physical features from each PDE system to build the datasets used for model training to impose a much more difficult setting.



(a) The grid G (dark & light blue) is used for both training and assimilation when we make comparisons with AutoEncoder-based methods. The dark blue points represent the possible observed locations S_{obs} during assimilation.

(b) The grid G (dark blue) is used for training and the staggered grid G' (dark & light red) is used for assimilation to study the flexibility of LAINR. The dark red points represent the possible observed locations S'_{obs} during assimilation.

Fig. 5. Grids for training and assimilation.

Shallow-water dataset We consider the following spherical shallow-water equation on the 2D sphere

$$\begin{aligned} \frac{d\mathbf{u}}{dt} &= -f\mathbf{k} \times \mathbf{u} - g\nabla h + \nu\Delta\mathbf{u}, \\ \frac{dh}{dt} &= -h\nabla \cdot \mathbf{u} + \nu\Delta h \end{aligned} \quad (64)$$

as a simplification of the primitive equations of atmospheric flow with the material derivative denoted by $\frac{d}{dt}$. The vorticity field $w = \nabla \times \mathbf{u}$ together with the thickness of the fluid layer h becomes the system states to be assimilated. ν stands for the viscosity, and the parameters f , g , Ω are all fixed and in consistent with the real earth surface. 20 trajectories starting from different initial states have been generated and downsampled onto a 128×64 regular latitude-longitude grid with $\Delta t \approx 1$ real hour. All the trajectories have been truncated into 240 time steps, and 10% of them are separately used for testing purposes, which means the shapes of the training and testing datasets are $(18, 240, 128, 64, 2)$ and $(2, 240, 128, 64, 2)$, respectively.

ERA5 dataset The ERA5 dataset [36] is the latest global atmospheric reanalysis dataset produced by the European Centre for Medium-Range Weather Forecasts (ECMWF), which is generated by assimilating observations from across the world into a delicate numerical weather prediction model. In our experiments, we do not make comparisons with the DA algorithms employed in the generation of the ERA5 dataset, but treat the dataset as the ground truth to provide a more realistic dynamical system than the previous ideal shallow-water dataset. We select the Z500 (geopotential at 500 hPa) and the T850 (temperature at 850 hPa) fields to construct our dataset, and we still refer to it as the ‘‘ERA5 dataset’’ for convenience. The choice of Z500 and T850 aligns with those of the early attempts [23,68,81,66] of DL-based weather forecasting, which in our setting, become the physical features we aim to assimilate. The spatial resolution of the Z500 and T850 fields is 128×64 with a 1-hour time interval. We have extracted 10 trajectories of a whole year for training, and 2 additional trajectories for testing purposes, which produce datasets of size approximately $(10, 365 \times 24, 128, 64, 2)$ and $(2, 365 \times 24, 128, 64, 2)$, respectively.

5.1.2. Model configurations

AutoEncoder In the previous work [62], the encoder-decoder mappings have been implemented with fully-connected neural networks, but such design choice raises potential scalability issues for our two-dimensional test cases, considering the prohibitively large parameter volume. Directly applying fully-connected networks could lead to a parameter volume of roughly $(2 \times 128 \times 64)^2 \times d \approx 2.7 \times 10^8 \times d$ for a network of depth d with 2-channel input of size 128×64 , which is impractical, particularly as this volume grows with the scale of the real deployment. Consequently, to make a systematical comparison with the AutoEncoder-based LA framework and circumvent the scalability concerns, we have opted to leverage convolutional neural networks (CNNs) to establish our baselines. The decision aligns with the natural suitability of CNNs for tasks that exhibit local spatial or temporal correlations.

We have chosen the Convolutional AutoEncoder (CAE) structure as employed in [1] for assimilating the room CO_2 concentration, and the AEflow [33] structure proposed in the field of turbulent flow simulation. With the latent dimension fixed, we vary the kernel size $K \in \{3, 5, 7\}$ and the number of hidden channels $C \in \{16, 32, 64\}$ for the CAE structure, and the kernel size $K \in \{3, 5, 7\}$ and the number of residual blocks $N_{\text{res}} \in \{4, 8, 12\}$ for the AEflow structure. To make the convolutions compatible with the boundary condition, we utilize periodic paddings along the zonal directions and replication paddings at the poles. See Appendix C.1 for detailed descriptions.

INR As introduced in Section 4.1, we have chosen the SINR structure as our INR implementation due to its reliable convergence guarantee. By empirically tuning the hyperparameters, we fix the hidden dimensions $h = 128$ for each state variable and $D = L = 8$ for our implementation of SINR in the following experiments.

5.1.3. Metrics

To show the advantages of our LAINR framework, it is important to quantify the performances by either the rooted-mean-square error (RMSE) evaluated on a pre-defined quasi-uniform skipped latitude-longitude grid [80], or the rooted-mean-square error (RMSE) weighted by the cosine values of latitudes. We choose to employ the latter one since it is more compatible with the L^2 norm on the sphere and more flexible when applied to unstructured data. The weighted RMSE is formulated as

$$\mathcal{L}_{\text{sqrt}}(\mathbf{x}, \mathbf{x}') = \sqrt{\frac{1}{\sum_{p \in G} \cos \phi(p)} \sum_{p \in G} \|\mathbf{x}(p) - \mathbf{x}'(p)\|^2 \cos \phi(p)}, \quad (65)$$

where $\phi(p)$ stands for the latitude of the position p , and all the RMSEs mentioned hereafter are computed as above unless otherwise specified. Besides, we also define an ‘‘unrooted’’ version

$$\mathcal{L}(\mathbf{x}, \mathbf{x}') = \frac{1}{\sum_{p \in G} \cos \phi(p)} \sum_{p \in G} \|\mathbf{x}(p) - \mathbf{x}'(p)\|^2 \cos \phi(p) \quad (66)$$

for network training, and the details can be found in the following subsections.

5.1.4. Training process

We divide the training process into 2 stages, namely the *pre-training* stage and the *fine-tuning* stage. The pre-training stage is designed to obtain a good latent representation for all of the training snapshots, and the fine-tuning stage focuses mainly on the reconstruction of the latent dynamics.

Pre-training We need to emphasize that the trainable parameters for the AutoEncoder baselines are the network parameters of both their encoders and decoders. On the contrary, no explicit encoding mapping is needed during the INR training, and the trainable parameters are made up of the decoder parameters as well as the latent representations. The reconstruction loss is used to measure the discrepancy between the original data and the data reconstructed from the latent representations. Naturally, we define the reconstruction loss for the AutoEncoder baselines as

$$L_{\text{rec}}^{\text{AE}}(\mathbf{x}) = \mathcal{L}(\mathbf{x}, D_{\varphi} \circ \mathcal{E}_{\theta}(\mathbf{x})). \quad (67)$$

However, since the latent representations are also trainable parameters for the INR models, the corresponding reconstruction loss is a little different:

$$L_{\text{rec}}^{\text{INR}}(\mathbf{x}, \mathbf{z}) = \mathcal{L}(\mathbf{x}, D_{\varphi}(\mathbf{z})). \quad (68)$$

Consequently, for the training dataset $\{\mathbf{x}_k\}_{k=1}^K$, the pre-training stage solves

$$\operatorname{argmin}_{\theta, \varphi} \frac{1}{K} \sum_{k=1}^K L_{\text{rec}}^{\text{AE}}(\mathbf{x}_k) = \operatorname{argmin}_{\theta, \varphi} \frac{1}{K} \sum_{k=1}^K \mathcal{L}(\mathbf{x}_k, D_{\varphi} \circ \mathcal{E}_{\theta}(\mathbf{x}_k)) \quad (69)$$

and

$$\operatorname{argmin}_{z_k, \varphi} \frac{1}{K} \sum_{k=1}^K L_{\text{rec}}^{\text{INR}}(\mathbf{x}_k, \mathbf{z}_k) = \operatorname{argmin}_{\theta, \varphi} \frac{1}{K} \sum_{k=1}^K \mathcal{L}(\mathbf{x}_k, D_{\varphi}(\mathbf{z}_k)) \quad (70)$$

for the AutoEncoder baselines and the INR models, respectively.

Fine-tuning To train a latent surrogate model to approximate the unknown latent dynamics, we define the multi-step prediction errors for the s -th step as

$$L_{\text{pred}}^{\text{AE},(s)}(\mathbf{x}_k) = \left\| \mathcal{E}_{\theta}(\mathbf{x}_{k+s}) - \mathcal{G}_{\psi}^s \circ \mathcal{E}_{\theta}(\mathbf{x}_k) \right\|^2, \quad L_{\text{pred}}^{\text{INR},(s)}(\mathbf{z}_k) = \left\| \mathbf{z}_{k+s} - \mathcal{G}_{\psi}^s(\mathbf{z}_k) \right\|^2 \quad (71)$$

with the common mean-squared errors on the latent space. The choice of s is a bit tricky, as a smaller s will force the models to learn the next-step prediction behavior and ignore the long-term dynamics, while a larger s may result in difficulties in training. Nevertheless, in our test cases, we have found that a fixed s is sufficient for the surrogate models to learn effective latent dynamics used for assimilation. The training process of the latent surrogate models is formulated as

$$\operatorname{argmin}_{\psi} \frac{1}{K} \sum_{k=1}^K L_{\text{pred}}^{\text{AE},(s)}(\mathbf{x}_k) = \operatorname{argmin}_{\psi} \frac{1}{K} \sum_{k=1}^K \left\| \mathcal{E}_{\theta}(\mathbf{x}_{k+s}) - \mathcal{G}_{\psi}^s \circ \mathcal{E}_{\theta}(\mathbf{x}_k) \right\|^2 \quad (72)$$

and

$$\operatorname{argmin}_{\psi} \frac{1}{K} \sum_{k=1}^K L_{\text{pred}}^{\text{INR},(s)}(\mathbf{z}_k) = \operatorname{argmin}_{\psi} \frac{1}{K} \sum_{k=1}^K \left\| \mathbf{z}_{k+s} - \mathcal{G}_{\psi}^s(\mathbf{z}_k) \right\|^2 \quad (73)$$

for the AutoEncoder baselines and the INR models, respectively. Note that in the fine-tuning stage, the optimization problems (69)(70) are solved as well, but with a smaller learning rate.

The detailed training procedures for the AutoEncoder models and the INR models are summarized in Algorithm 1 and Algorithm 2, respectively, where we demonstrate the optimization with Stochastic Gradient Descent (SGD) with a consistent learning rate η for simplicity.

Algorithm 1 training for AutoEncoders

```

1: Input: Training data  $\{\mathbf{x}_k\}_k$ , epochs  $T$ 
2: Output: trained models  $\mathcal{E}_\theta, D_\varphi, G_\psi$ 
3: Initialize  $\mathcal{E}_\theta, D_\varphi$  and  $G_\psi$ .

4: for  $t = 1$  to  $T$  do
5:   for  $\mathbf{x}_k$  in  $\{\mathbf{x}_k\}_k$  do
6:     ** forward propagation **
7:     Calculate the reconstruction loss  $L_{\text{rec}}^{\text{AE}}(\mathbf{x}_k)$ .
8:     Calculate the prediction loss  $L_{\text{pred}}^{\text{AE}(s)}(\mathbf{x}_k)$ .
9:   end for
10:  ** update model weights **
11:   $\theta \leftarrow \theta - \eta \nabla_\theta \sum_k (L_{\text{rec}}^{\text{AE}}(\mathbf{x}_k) + L_{\text{pred}}^{\text{AE}(s)}(\mathbf{x}_k))$ 
12:   $\varphi \leftarrow \varphi - \eta \nabla_\varphi \sum_k (L_{\text{rec}}^{\text{AE}}(\mathbf{x}_k) + L_{\text{pred}}^{\text{AE}(s)}(\mathbf{x}_k))$ 
13:   $\psi \leftarrow \psi - \eta \nabla_\psi \sum_k L_{\text{pred}}^{\text{AE}(s)}(\mathbf{x}_k)$ 
14:  Update  $s$ .
15: end for

```

Algorithm 2 training for INRs

```

1: Input: Training data  $\{\mathbf{x}_k\}_k$ , epochs  $T$ 
2: Output: trained models  $D_\varphi, G_\psi$ 
3: Initialize  $D_\varphi, G_\psi$  and  $\{\mathbf{z}_k\}_k$ .

4: for  $t = 1$  to  $T$  do
5:   for  $\mathbf{x}_k$  in  $\{\mathbf{x}_k\}_k$  do
6:     ** train the encoder-decoder **
7:     Calculate the reconstruction loss  $L_{\text{rec}}^{\text{INR}}(\mathbf{x}_k, \mathbf{z}_k)$ .
8:     Calculate the prediction loss  $L_{\text{pred}}^{\text{INR}(s)}(\mathbf{z}_k)$ .
9:     ** update only the latent vectors **
10:     $\mathbf{z}_k \leftarrow \mathbf{z}_k - \eta \nabla_{\mathbf{z}_k} L_{\text{rec}}^{\text{INR}}(\mathbf{x}_k, \mathbf{z}_k)$ 
11:   end for
12:  ** update model weights **
13:   $\varphi \leftarrow \varphi - \eta \nabla_\varphi \sum_k L_{\text{rec}}^{\text{INR}}(\mathbf{x}_k, \mathbf{z}_k)$ 
14:   $\psi \leftarrow \psi - \eta \nabla_\psi \sum_k L_{\text{pred}}^{\text{INR}(s)}(\mathbf{z}_k)$ 
15:  Update  $s$ .
16: end for

17: ** output the encoder **
18: The encoder  $\mathcal{E}$  is given implicitly as follows:
19:  $\mathbf{x} \mapsto \mathbf{z} = \text{argmin}_{\mathbf{z}'} \mathcal{L}(\mathbf{x}_k, D_\varphi(\mathbf{z}'))$ .

```

5.2. Performance comparisons

5.2.1. Exploring suitable encoder-decoder structures

Our experiments start with an exploration of suitable encoder-decoder structures for the shallow-water dataset, by making a comparative analysis of the reconstruction capabilities of the three distinct ROMs: CAE, AEflow and SINR. The assessment has been conducted independently of the surrogated models assigned to handle latent dynamics before the fine-tuning stage.

We set the latent dimension for the AutoEncoder baselines (CAE and AEflow) as 1024 to obtain minimal modifications from the original proposed structure [1,33] and that for the INR encoders as 400, which presents a more challenging setup. Furthermore, to perform an ablation study on the superiority of our proposed SINR architecture among other existing INR models, we additionally implement

- **FourierNetCartes** [83], a Fourier Network structure [29] fed with the 3D cartesian coordinates,
- **FourierNetLatlon**, a Fourier Network structure [29] fed with the 2D latitude-longitude coordinates, and
- **SINRNoSkip**, a SINR structure without the skip connections, where all the terms of the last summation in (38) are removed except for the last layer.

To make a fair comparison with our proposed framework, we have also conducted a grid search (Appendix C.2) on the hyperparameters (the network depth, the convolutional kernel size, etc.) for the previous CAE and AEflow architectures. The optimal network configurations are fixed and used for all the following experiments.

The reconstruction errors evaluated on the testing datasets are summarized in Fig. 6, where the numbers of latent dimensions are appended to the method names for clarification. Even with a higher latent dimension, the AutoEncoder baselines have obtained relatively lower accuracy, indicating their inefficiency in capturing and restoring the physical features. In contrast, INR-based methods operate on a more compressed latent space of merely 400 dimensions, which is less than half the latent dimensionality utilized by the AutoEncoders, but they can still achieve superior performances. Besides, our SINR architecture outperforms the other INR models, demonstrating the effectiveness of the spherical filters as well as the skip connections, and Fig. 6 has shown that the spherical filters tend to play a much more significant role than the skip connection mechanism. Upon observing the compelling accuracy and efficiency of the SINR structure, it becomes clear that our choice of SINR for the LAINR framework is well-justified.

5.2.2. Embedding the physical dynamics

Embedding a physical dynamical system into a latent space involves not only a precise encoding of the snapshot at each time step but also an accurate discovery of the latent dynamics, which is necessary for the assimilation process in the latent space. To evaluate both the consistency and the stability of the surrogate models for latent dynamics, the multi-step prediction RMSE

$$L_{\text{sqr}}^{(s)} = \frac{1}{K} \sum_{k=1}^K \mathcal{L}_{\text{sqr}}(\mathbf{x}_{k+s}, D_\varphi \circ G^s \circ \mathcal{E}_\theta(\mathbf{x}_k)) \quad (74)$$

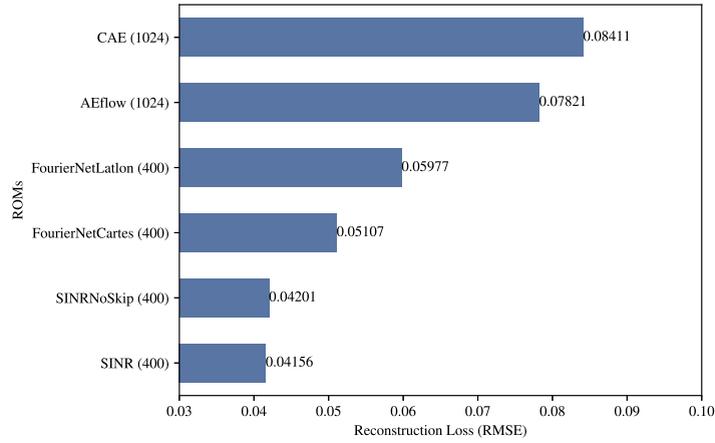


Fig. 6. Comparison of the capabilities for low-dimensional reconstructions on the shallow-water dataset.

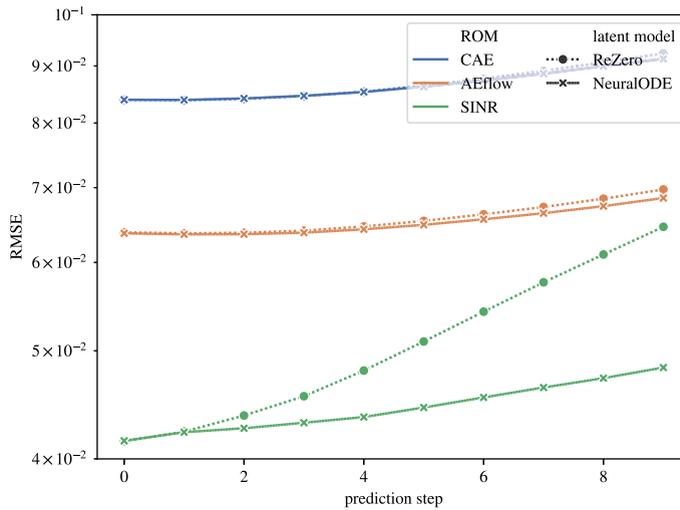


Fig. 7. Comparison of the multi-step prediction errors on the physical space for the shallow-water dataset.

is utilized, where k is the time index for the state \mathbf{x} and s indicates the number of prediction steps. We have investigated two distinct surrogate models for the latent dynamics:

- **ReZero** [7]. We have followed the work of [62], which implements the ReZero structure as a discrete-time dynamical model. The ReZero structure is implemented with 5 residual blocks, and the activation function is fixed as the LeakyReLU [52] except for the final block.
- **NeuralODE** [17,18]. The NeuralODE model interprets the latent dynamics as a continuous-time dynamical system, where the surrogate model is implemented via a fully-connected network and trained with an RK4 integrator via the TorchDiffEq [17,18] package.

We have fixed $s = 1$ to define the loss function for training the latent surrogate models. Fig. 7 exhibits the comparison of the increments in prediction errors with respect to the prediction steps. Compared with CAE, AEflow shows relatively better performances probably due to its higher reconstruction accuracy as visualized in Fig. 6, while our SINR model presents significantly lower prediction errors. Meanwhile, by comparing the two latent surrogate models, it can be concluded that NeuralODE is more capable of capturing the latent dynamics than ReZero especially when working with SINR. In contrast, they do not show significant differences in the prediction errors when paired with the AutoEncoder baselines, which implies that their failure to obtain good latent representations probably leads to much more complicated latent dynamics.

5.2.3. Working together with existing DA algorithms

The effectiveness of a DA algorithm is linked to the quality of the forward integrator, but an accurate surrogate model alone does not guarantee better assimilation performances, owing to the potential incompatibility between the surrogate model’s non-linearity

Table 3
Configurations for assimilation on the shallow-water dataset.

variables	values	count
σ_z^b	0.01, 0.003, 0.001	3
σ^m	0.1, 0.03, 0.01, 0.003, 0.001, 0.0003, 0.0001	7
inflation	1.02, 1.05, 1.10	3
constants	$\sigma^o = 0.1, \sigma_x^b = 0.1, N = 64$	-

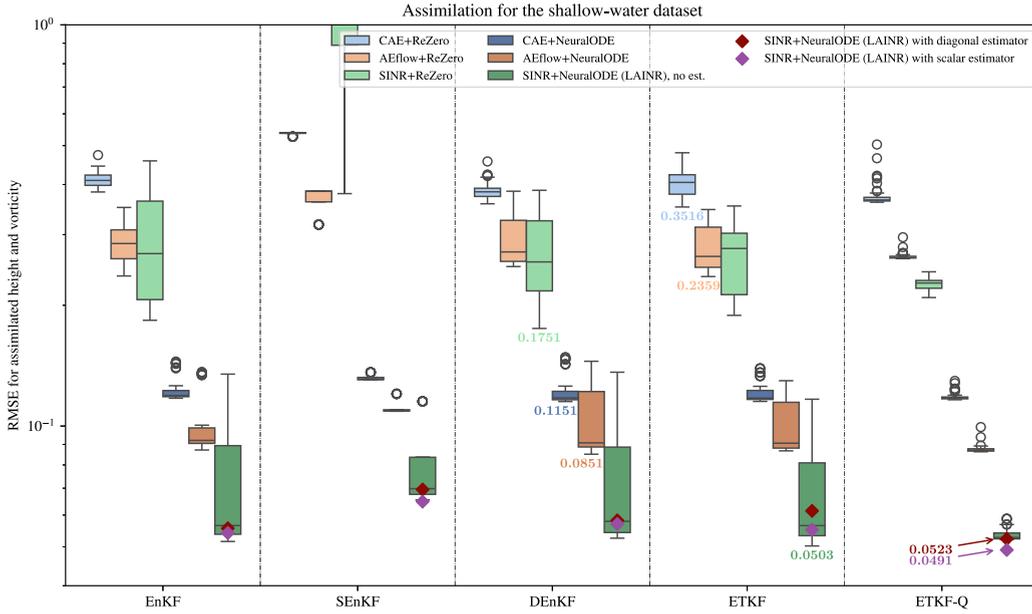


Fig. 8. Comparison of assimilation performances on the shallow-water dataset.

behaviors and the assimilation algorithms. To address such concern, we check the compatibility of our LAINR framework, alongside the AutoEncoder baselines, when paired with existing assimilation algorithms.

We choose several variants of the Kalman filter method for testing, and we would like to emphasize that similar experiments can be extended to other types of assimilation algorithms including the variational methods (3/4D-Var). The ensemble size is fixed as $N = 64$. Detailed implementations for the Kalman filter (KF) algorithms (EnKF, SEnKF, DEnKF, ETKF, ETKF-Q) involved in our experiments are all exhibited in Appendix D.

The observation operator $\mathcal{H} : G \rightarrow S_{\text{obs}}$ is set as a random sampling on the grid G of size $128 \times 64 \times 2$ as illustrated in Fig. 5a with the number of the observations $|S_{\text{obs}}| = 1024$ (6.25% of the total number of grid points). All observations are perturbed with independent Gaussian noise of variance $(\sigma^o)^2 = 0.1^2$, which is approximately 10% of the ground truth since the datasets are normalized to zero mean and standard deviation in advance. The initial physical background estimate is set as a perturbed observation of the ground truth via \mathcal{H} with a Gaussian noise $\sigma_x^b = 0.1$. Additionally, the noise for the latent surrogate model is assumed as Gaussian of covariance matrix $(\sigma^m)^2 I_m$ without any uncertainty estimator for now, and the covariance for the uncertainty of the latent background estimate is set as $(\sigma_z^b)^2 I_m$.

To work with existing DA algorithms, the latent background estimate σ_z^b , the uncertainty σ^m of the latent models as well as the inflation needs to be specified. To study more thoroughly the behaviors of the DL models when applied to assimilation, we have conducted experiments for each model with different KF algorithms under $3 \times 7 \times 3 = 63$ individual configurations (Table 3). See Fig. 8 for the assimilation performances on the shallow-water dataset, where each box contains 63 results. For each model, the best RMSE across all KFs (315 configurations in total) is marked in bold below the corresponding box.

Notably, NeuralODE significantly outperforms ReZero when combined with either the AutoEncoder baselines or our SINR model, although such huge differences have not been observed previously on the prediction tasks when evaluating the fine-tuning performances. It can be inferred that modeling the latent dynamics with a continuous-time model helps stabilize the assimilation cycle. When comparing different encoder-decoder mappings, our SINR outperforms the AutoEncoder baselines (CAE and AEflow) when working together with KFs under different assumptions on the model noise σ^m and the latent background estimate σ_z^b . Such superiority of INR emphasizes its robustness across various configurations. The comparison of optimal performances further validates the supremacy of the SINR model. It achieves considerably lower RMSEs than the AutoEncoder methods, reinforcing the idea that the LAINR framework is a superior choice for latent assimilation tasks.

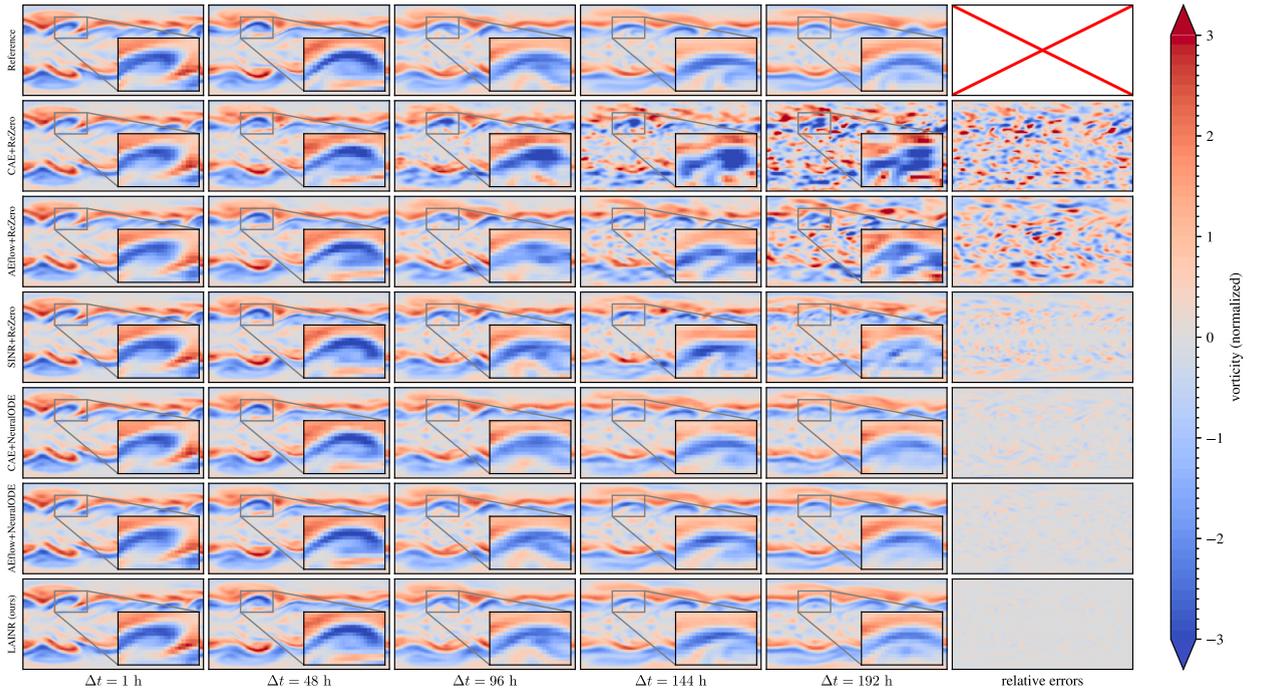


Fig. 9. Assimilation vorticity fields for the shallow-water dataset.

5.2.4. Uncertainty estimations

To study the effectiveness of the uncertainty estimator in our LAINR framework, we have run additional experiments with the model noise given by the following estimators (Section 4.4.1):

- **scalar**: the estimator D is a scalar matrix with a unique trainable parameter;
- **diagonal**: the estimator D is a diagonal matrix with all of its diagonal entries being trainable.

Meanwhile, the latent background estimates are provided as described in Section 4.4.2. For each KF, the optimal performances of LAINR (SINR+NeuralODE), where only the inflation varies, combined with uncertainty estimators (abbreviated as u.e.) are marked on the box plot (Fig. 8).

When working with different KFs, the uncertainty estimator serves as a robust method that eliminates the need for empirical tuning of the model noise level σ^b as well as the uncertainty for the latent background estimate σ_z^b . As illustrated in Fig. 8, the scalar estimator uniformly outperforms the diagonal estimator, and it obtains better performances than all the empirically-tuned configurations when coupled with SEnKF and ETKF-Q, while it also achieves nearly optimal results across other KFs. Such optimality may also reveal the isotropy of the LAINR prediction errors, which indicates that the SINR model has done well in the model reduction task during training. It is worth noting that the uncertainty estimator offers an effective data-driven tool to estimate the uncertainties of the trained networks, which can be particularly advantageous in situations where time or computational resources are limited, or where expert knowledge for parameter-tuning is costly or unavailable.

5.2.5. Visualization of the assimilation processes

By fixing the optimal choices of latent surrogate models, KFs as well as the assimilation configurations for each encoder-decoder mapping, we have visualized the evolution of the assimilated vorticity fields of the shallow-water dataset along with the ground truth via Fig. 9, where a zoom-in view for each snapshot is attached for illustration purpose. The absolute errors for the last time step based on the ground truth (reference) have been attached to the last column as well. Fig. 10 plots the detailed RMSEs for each assimilation cycle.

The AutoEncoder-based methods are unable to capture the physical dynamics during assimilation cycles when coupled with ReZero, and the assimilated states tend to diverge after 48 hours. Our SINR model is able to keep tracking the dynamics until the assimilation time exceeds 144 hours (6 days) because of its high accuracy in model reduction. When using NeuralODE for the latent surrogate model, all the ROMs remain stable as the assimilation time steps increase, which does not show any divergence of errors. However, some small features of the vorticity fields have been missed by the AutoEncoders after 144 hours (6 days). Our LAINR still persistently recovers most of the physical state with relatively lower assimilation errors.

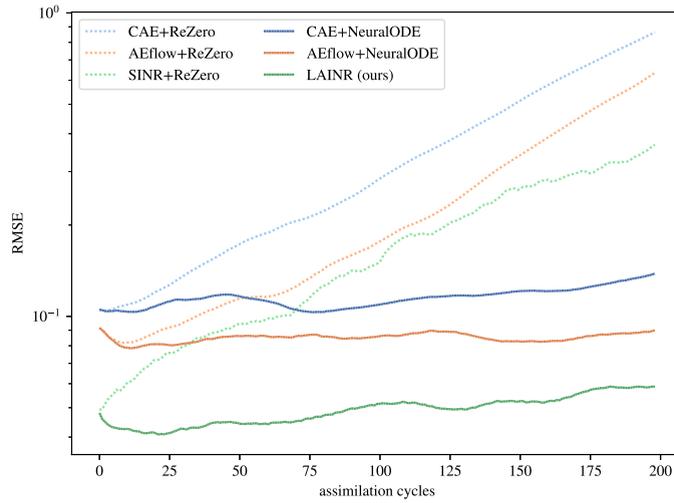


Fig. 10. Increments of RMSEs during assimilation.

5.2.6. Testing on a more realistic case

To explore the capability of our proposed LAINR framework when dealing with real data, we have conducted similar comparative experiments but on a much more realistic ERA5 dataset. The study involves CAE, AEflow and our SINR as the encoder-decoder mappings. We have chosen not to advance other architectures due to their inferiority shown in the previous ablation study. Meanwhile, we have fixed the latent surrogate model as NeuralODE in that it has outperformed ReZero for both prediction and assimilation tasks, and thus we do not specify it hereafter.

Same as before, our experiments start from training an encoder-decoder mapping aiming at accurately reconstructing the state variables from lower-dimensional latent representations. Then it follows that each model is used to learn the underlying latent dynamics. Fig. 11a displays the prediction errors (truncated from above) by computing the prediction RMSE losses on the testing dataset, and the corresponding reconstruction accuracy can be inferred by the performances at the 0-th steps.

In contrast to the shallow-water case, CAE outperforms AEflow in both the reconstruction and the prediction tasks in the ERA5 case. Meanwhile, our LAINR still achieves the best reconstruction accuracy as well as the lowest prediction errors, which is consistent with the previous results on the shallow-water dataset. Fig. 11b exhibits the subsequent assimilation performances of the trained models when collaborating with various KFs. The configurations of the assimilation remain the same as those described in Section 5.2.3, which means each box contains 63 distinct parameter combinations.

Despite AEflow's relatively poor performances in the reconstruction and the prediction tasks when compared with CAE displayed in Fig. 11a, it has shown comparable performances for the assimilation tasks, and it even surpasses CAE if the assimilation parameters are properly tuned (see some of the outliers in the box plot). Such a phenomenon supports the idea that higher reconstruction and prediction accuracy do not necessarily lead to better assimilation performances, and the compatibility between the surrogate models and the DA algorithms also counts. Our LAINR continues to show superior performances among all the configurations, and the uncertainty estimators turn out to be effective as well, which further strengthens the reliability of LAINR for real deployment.

5.3. Mesh-free flexibility

One of the main features of our LAINR framework is its flexibility in handling inputs from partially observed or even irregular sampling sets. To demonstrate this, we have performed additional experiments on a sampling set different from what we have adopted for training.

5.3.1. Unstructured reconstruction

To illustrate the effectiveness of our LAINR framework for unstructured cases, we have picked a snapshot among the testing data of the shallow-water dataset, and then randomly masked out a subset from the collection of the grid points to provide unstructured observations. For the AutoEncoder-based models (CAE and AEflow), an interpolated field is first computed and fed into as the input, where the interpolation is implemented via an RBF interpolator from the SciPy package [78] with the cubic kernel.¹ After that, an encoding-decoding process is performed to evaluate the accuracy of the latent representations z^{CAE} and z^{AEflow} in order to emulate what an AutoEncoder-based model will probably do. On the contrary, the masked field can be directly sent to the SINR model for encoding, and the only difference from the gridded case is that the number of the summation terms of (26) is reduced. The accuracy

¹ <https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.RBFInterpolator.html>.

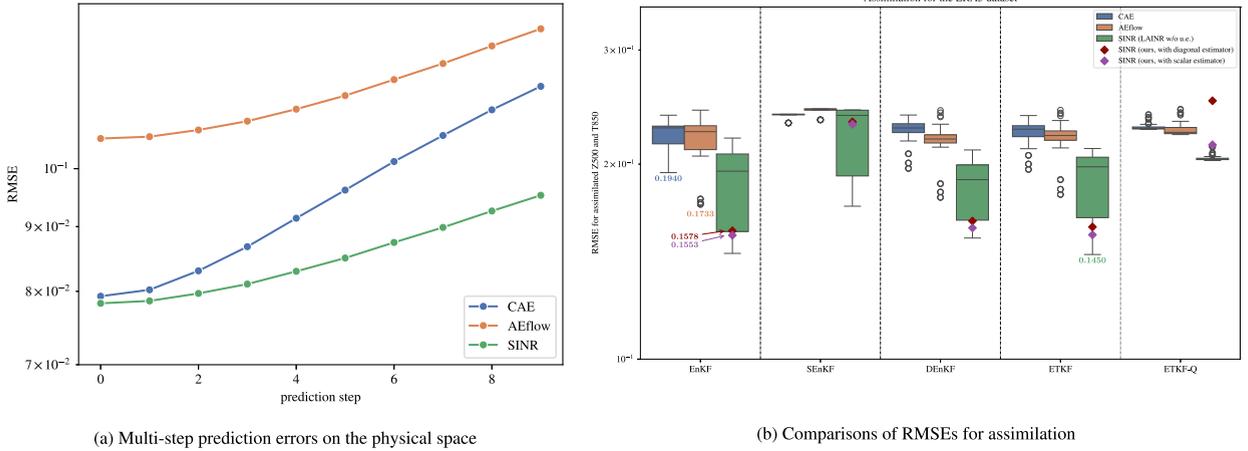


Fig. 11. Experiments on the ERA5 dataset (both truncated from above).

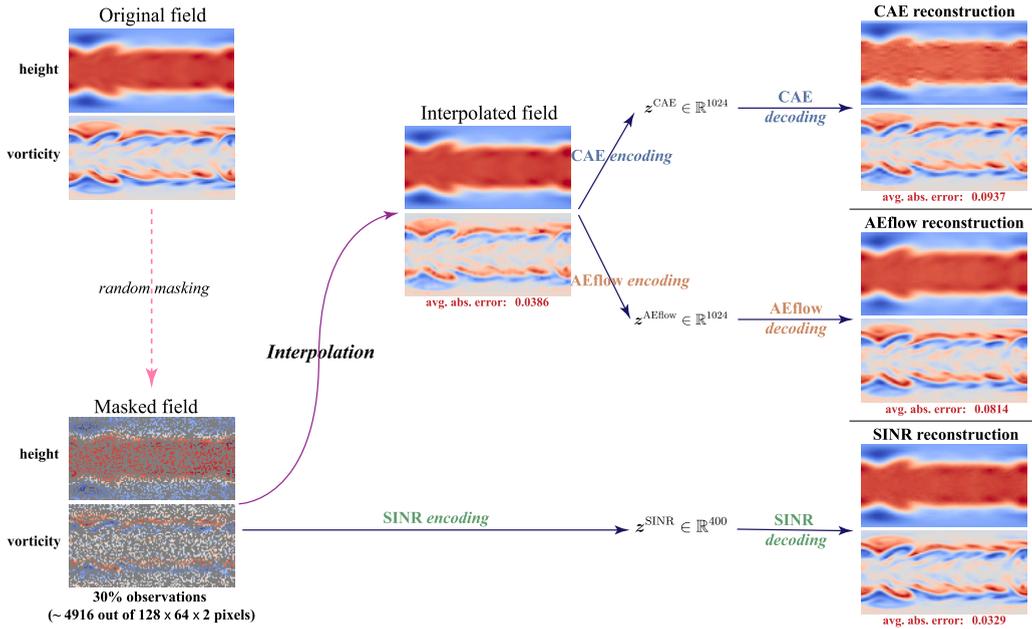


Fig. 12. Reconstruction results from a random sampling set (30% observations).

of the latent representation z^{SINR} is also implied by the reconstruction error after the decoding. See Fig. 12 for the comparisons of the reconstruction results with 30% random observations (few) and Fig. 13 with 3% observations (rare).

Under different observation ratios, the AEflow model performs slightly better than CAE, but both of the AutoEncoder baselines have failed to reconstruct the masked field accurately, especially when the observations are limited. On the contrary, our SINR model has shown consistent superiority, and it is still able to recover the system states with rare observations, where nearly all of the physical features have been masked out.

In addition, we have pushed our SINR model to the limit when the observations are extremely rare. We have conducted similar experiments with 2%, 1%, 50, 20 and 10 observations for each feature respectively, and the results are shown in Fig. 14. To our surprise, the SINR model can still recover the general fluid structures before the number of observations drops to 20 or smaller, which is only approximately 0.3% of the total number of grid points. This is strong evidence of the robustness of our SINR model for unstructured observations, and it also demonstrates the potential of our LAINR framework for real-world applications where the observations usually fail to be complete or structured.

5.3.2. Zero-shot assimilation

Previous experiments are all based on the same regular latitude-longitude grid for both training and assimilation to make comparisons with AutoEncoder baselines. To emphasize the flexibility of our proposed LAINR framework for different sampling sets as

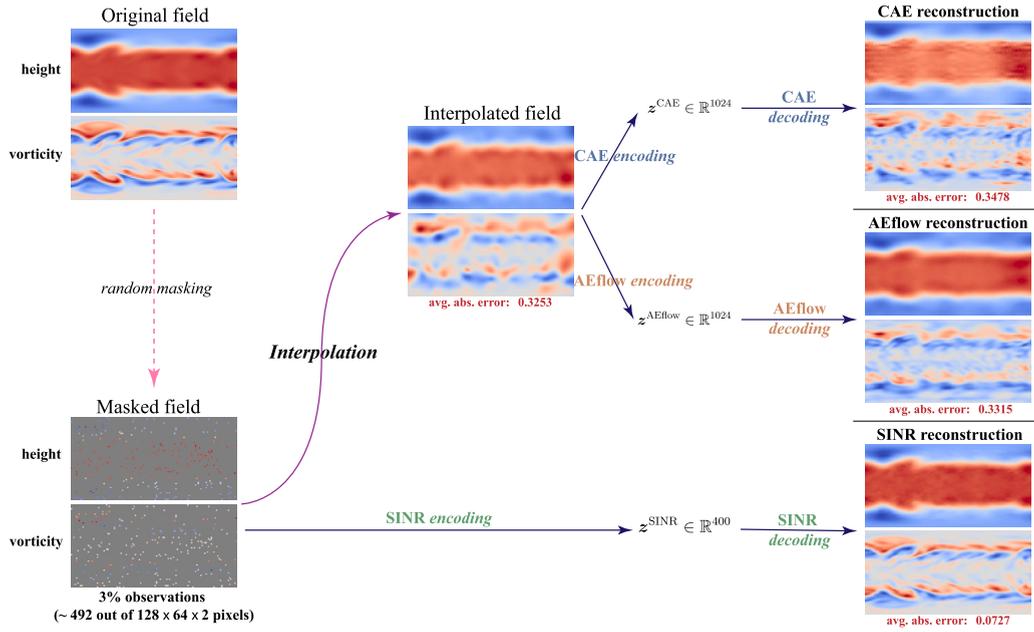


Fig. 13. Reconstruction results from a random sampling set (3% observations).

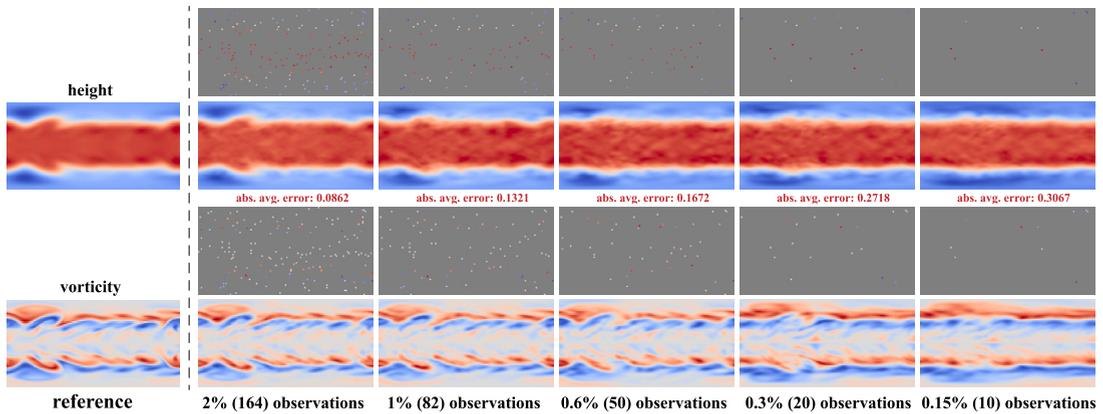


Fig. 14. Explore the limit of SINR for rare observations.

inputs, we have performed the same assimilation procedures but on a staggered grid G' as illustrated in Fig. 5b. Once the SINR model together with NeuralODE has been well-trained on the grid G , we apply LAINR to assimilate the observations on the grid G' . Note that under such circumstances, no observations from G used for training are provided, and only random observations from a subset of G' ($|S'_{\text{obs}}| = 1024$) are available for the assimilation process. Here, the grid G for training and the grid G' do not intersect, and we refer to such setting as the *zero-shot assimilation*.

The parameter configurations of the assimilation remain the same as those of Section 5.2.3. See Fig. 15 for the on-grid assimilation (on G , copied from Fig. 8) and the zero-shot assimilation (on G'). As shown in the box plot, when working with various KFs, the zero-shot performance of our LAINR framework is comparable to the on-grid performance, indicating the robustness of LAINR on zero-shot tasks when inconsistent spatially irregular observations may appear.

6. Conclusion and future work

In this study, we have proposed a latent assimilation framework called Latent Assimilation with Implicit Neural Representations (LAINR), which offers a new integration of machine learning techniques and data assimilation concepts. Besides, our LAINR framework opens up new possibilities for processing and assimilating complex data systems by using the reduce-order technique. Apart from the theoretical analysis of LAINR, in comparison to existing AutoEncoder-based models, experiments have shown that LAINR surpasses the previous models in terms of practical performance.

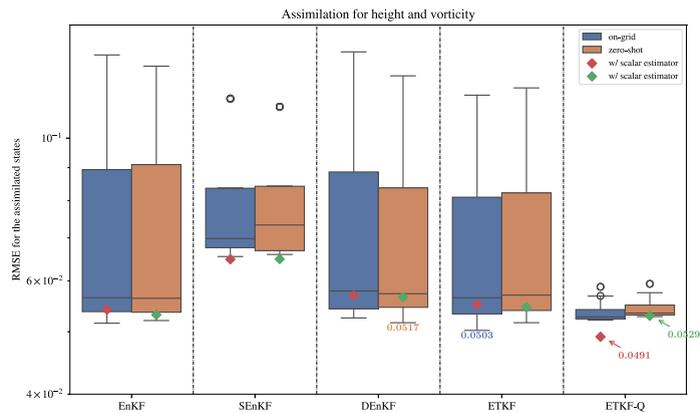


Fig. 15. Comparisons of RMSEs for zero-shot assimilation on the staggered grid.

The main features of our LAINR framework can be summarized as follows. Same as previous AutoEncoder-based models, LAINR's capability for non-linear embedding introduces the possibility that the complex dynamics of physical states can be transformed into simpler latent dynamics for learning, but the difference lies in the fact that LAINR exhibits a clear superiority in scalability. Since the architecture of LAINR uses coordinates as inputs, it effortlessly accommodates increasing dimensionality and resolution unlike most classical models, which can hardly avoid re-designing the whole structure due to their reliance on the input structures.

One of the biggest advantages of LAINR is perhaps its flexibility in handling inputs from partially observed or ungridded data, which becomes more suitable under potential hardware limitations when complete observations are not always guaranteed such as satellite or radar observations. Unlike the AutoEncoder baselines treating the state field as a structured grid of state features, LAINR stores the entire state field as a continuous mapping and such a special perspective provides a more accurate description of physical features, aiding in the learning of model-reduction mappings. Meanwhile, the availability of a continuous field may help establish more complicated observation operators such as the solution of radiative-transfer models in that the interpolation is no longer necessary. As a consequence, its utility scope is wider in real-world scenarios compared with AutoEncoder-based architectures.

Nonetheless, it is also important to acknowledge that areas where further research and development are needed, including the testing of LAINR on real observation datasets to provide further validations of its effectiveness and robustness. Besides, future work should explore more advanced uncertainty estimators as well. The SINR presented in the LAINR framework acts as a novel approach to construct an INR mapping from integer-indexed basis functions, and a similar idea can be applied to other geometrical structures other than the 2D sphere. Finding a suitable collection of orthogonal basis functions for general Riemann manifolds to obtain the convergence guarantee may become a future direction.

Reproducibility

All the codes for our proposed LAINR framework and the corresponding experiments are available online at GitHub (<https://github.com/zylipku/LAINR>) including the data generation for the shallow-water datasets and the full list of assimilation results.

CRediT authorship contribution statement

Zhuoyuan Li: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Conceptualization. **Bin Dong:** Writing – review & editing, Supervision, Software, Methodology, Funding acquisition. **Pingwen Zhang:** Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

I have shared a link to my source codes for the corresponding algorithms and data generation at the end of the article.

Acknowledgements

Zhuoyuan Li and Pingwen Zhang are supported in part by the National Natural Science Foundation of China (No. 12288101). Bin Dong is supported in part by the New Cornerstone Investigator Program.

Appendix A. An alternative for the SINR representability

In Proposition 4.6, we have deduced a subset S_T of the set of all possible SINRs, S_{SINR} . However, we have to point out that there exist many more functions excluded by S_T that can be represented by a SINR.

Denote by $\mathbf{D}^{(l)} = \text{diag} \{d_1^{(l)}, \dots, d_h^{(l)}\}$ the diagonal matrix with entries $(d_1^{(l)}, \dots, d_h^{(l)})^\top = \mathbf{g}_l(\mathbf{p})$ as defined in (38). Consequently, the update formulas (38) can be transformed into the following

$$\begin{aligned} \boldsymbol{\gamma}^{(0)} &= \mathbf{D}^{(0)} \mathbf{1}, \\ \boldsymbol{\gamma}^{(l)} &= \mathbf{D}^{(l)} (\mathbf{W}^{(l)} \boldsymbol{\gamma}^{(l-1)} + \mathbf{b}^{(l)}), \quad l = 1, 2, \dots, L, \\ \mathbf{f}(\mathbf{p}) &= \sum_{l=0}^L (\tilde{\mathbf{W}}^{(l)} \boldsymbol{\gamma}^{(l)} + \tilde{\mathbf{b}}^{(l)}). \end{aligned} \quad (\text{A.1})$$

By using basic linear algebra, we can expand the update formulas (A.1) into

$$\begin{aligned} \boldsymbol{\gamma}^{(0)} &= \mathbf{D}^{(0)} \mathbf{1}, \\ \boldsymbol{\gamma}^{(1)} &= \mathbf{D}^{(1)} \mathbf{W}^{(1)} \mathbf{D}^{(0)} \mathbf{1} + \mathbf{D}^{(1)} \mathbf{b}^{(1)}, \\ \boldsymbol{\gamma}^{(2)} &= \mathbf{D}^{(2)} \mathbf{W}^{(2)} \mathbf{D}^{(1)} \mathbf{W}^{(1)} \mathbf{D}^{(0)} \mathbf{1} + \mathbf{D}^{(2)} \mathbf{W}^{(2)} \mathbf{D}^{(1)} \mathbf{b}^{(1)} + \mathbf{D}^{(2)} \mathbf{b}^{(2)}, \\ &\vdots \\ \boldsymbol{\gamma}^{(L)} &= \mathbf{D}^{(L)} \mathbf{W}^{(L)} \mathbf{D}^{(L-1)} \dots \mathbf{W}^{(1)} \mathbf{D}^{(0)} \mathbf{1} + \sum_{l=1}^L \mathbf{D}^{(L)} \mathbf{W}^{(L)} \dots \mathbf{D}^{(l+1)} \mathbf{W}^{(l+1)} \mathbf{D}^{(l)} \mathbf{b}^{(l)}, \\ \mathbf{f}(\mathbf{p}) &= \sum_{l=0}^L (\tilde{\mathbf{W}}^{(l)} \boldsymbol{\gamma}^{(l)} + \tilde{\mathbf{b}}^{(l)}). \end{aligned} \quad (\text{A.2})$$

The proof of Proposition 4.6 essentially sets the weight matrix $\mathbf{W}^{(l)}$ and the output bias $\tilde{\mathbf{b}}^{(l)}$ for each layer l as zero to obtain a subset S_T easy for analysis, which is too rough to study the representability of the SINR architecture. To dive a bit more deeply into the representability, here we consider an alternative setting, where $\mathbf{W}^{(l)}$ is set as a matrix with its unique non-zero (i_l, i_{l-1}) -th entry as 1, and all the bias terms are set to zero. Then

$$\begin{aligned} \gamma_k^{(0)} &= d_k^{(0)}, \\ \gamma_k^{(1)} &= \delta_{k,i_1} d_{i_1}^{(1)} d_{i_0}^{(0)}, \\ \gamma_k^{(2)} &= \delta_{k,i_2} d_{i_2}^{(2)} d_{i_1}^{(1)} d_{i_0}^{(0)}, \\ &\vdots \\ \gamma_k^{(L)} &= \delta_{k,i_L} d_{i_L}^{(L)} d_{i_{L-1}}^{(L-1)} \dots d_{i_0}^{(0)}, \end{aligned} \quad (\text{A.3})$$

where k is the index of the k -th coordinates of these vectors, and δ stands for the Dirac delta function. It follows that each coordinate of the output $\mathbf{f}(\mathbf{p})$ can be written as a linear combination of the graded spherical harmonics

$$\mathbf{f}_k(\mathbf{p}) \in \text{Span} \left(d_{i_0}^{(0)}, d_{i_1}^{(1)} d_{i_0}^{(0)}, \dots, d_{i_L}^{(L)} d_{i_{L-1}}^{(L-1)} \dots d_{i_0}^{(0)} \right), \quad (\text{A.4})$$

where each $d_{i_l}^{(l)}$ is a linear combination of the spherical harmonics $Y_{D+l}^{-D}(\mathbf{p}), Y_{D-1+l}^{-(D-1)}(\mathbf{p}), \dots, Y_{0+l}^0(\mathbf{p}), \dots, Y_{D-1+l}^{D-1}(\mathbf{p}), Y_{D+l}^D(\mathbf{p})$. Specially, if one has picked a linear combination $d_{i_l}^{(l)}$ for each layer l , then their multiplication $d_{i_L}^{(L)} d_{i_{L-1}}^{(L-1)} \dots d_{i_0}^{(0)}$ can always be represented by a SINR.

Appendix B. Detailed dataset configurations

B.1. Shallow-water equations

The spherical shallow-water equation on the 2D sphere writes

$$\begin{aligned} \frac{d\mathbf{u}}{dt} &= -f \mathbf{k} \times \mathbf{u} - g \nabla h + \nu \Delta \mathbf{u}, \\ \frac{dh}{dt} &= -h \nabla \cdot \mathbf{u} + \nu \Delta h, \end{aligned} \quad (\text{B.1})$$

where $\mathbf{u} = (u, v)^\top$ is the 2D velocity field tangent to the sphere surface, and h is the thickness of the fluid layer, both of which are defined on a channel domain $(\lambda, \phi) \in [0, 2\pi] \times [-\pi/2, \pi/2]$. The Coriolis parameter $f = 2\Omega \sin \phi$ with the angular velocity Ω and the gravity g are fixed in consistency with the real earth surface.

Table C.4
Forward propagation for the CAE baseline.

	Layer	Output Shape
Encoder	Input	(-1, 2, 128, 64)
	ConvBlock-1	(-1, C, 64, 32)
	ConvBlock-2	(-1, C, 32, 16)
	ConvBlock-3	(-1, 8, 16, 8)
Decoder	Input	(-1, 8, 16, 8)
	ConvTransposeBlock-1	(-1, C, 32, 16)
	ConvTransposeBlock-2	(-1, C, 64, 32)
	ConvTransposeBlock-3	(-1, 2, 128, 64)

The model is set up by initializing two symmetric zonal flows representing typical mid-latitude tropospheric jets

$$v|_{t=0} = 0, \quad u|_{t=0} = \begin{cases} \frac{u_m}{e_n} \exp\left(\frac{1}{(\phi - \phi_0)(\phi - \phi_1)}\right), & \phi_0 < |\phi| < \phi_1, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{B.2})$$

which is slightly modified from [32] to create symmetric fields, where $\phi_0 = \pi/7$ and $\phi_1 = \pi/2 - \phi_0$ gives the boundary of the jet. $e_n = \exp(-4/(\phi_1 - \phi_0)^2)$ is the normalizer so that the maximal zonal wind u_m lies at the jet's mid-point. The thickness h is initialized as a balanced height field, which is obtained by integrating the balance equation

$$gh(\phi) = gh_0 - \int_{-\phi}^{\phi} au(\phi') (f + a^{-1}u(\phi') \tan \phi') d\phi'.$$

The constant h_0 is chosen so that the global mean of h is 10^4 . The balanced field is then perturbed by adding a localized bump

$$h'(\lambda, \phi) = \hat{h} \exp[-(\lambda/\alpha)^2 - ((\phi_2 - \phi)/\beta)^2] \cos \phi, \quad (\text{B.3})$$

where $\hat{h} = 120$, $\alpha = 1/3$, $\beta = 1/15$, and $\phi_2 = \pi/4$.

The datasets for both training and testing are generated via a sample script from the Dedalus project [13] based on Python language within a spectral framework. The maximum velocity u_m is randomly sampled from the uniform distribution on the interval (60, 80) to generate distinct 20 trajectories for separation of training and testing. Each trajectory has been recorded from the 360-th hour to the 600-th hour with a 1-hour interval to capture rich physical features, and the spatial resolution is set as 128×64 . The vorticity $w = \nabla \times \mathbf{u}$ and the thickness h are the two features for network training and assimilation.

B.2. ERA5 dataset

We have made use of the WeatherBench [65] dataset, which is a resampled version of the ERA5 dataset specially designed for DL training. We select the subset of latitude-longitude resolution $2.8125^\circ \times 2.8125^\circ$, corresponding to exactly a spatial resolution of 128×64 . The Z500 and T850 fields from 2006 to 2015 (10 years) with a 1-hour interval make up the training dataset, and the testing dataset is composed of the fields from 2017 and 2018.

Appendix C. AutoEncoder baselines

C.1. Detailed network structures

The detailed propagations for both the CAE and AEflow structures are listed in Table C.4 and Table C.5, respectively, where the batch size is replaced by “-1”. As for the activation functions, each convolutional layer is followed by a LeakyReLU [52] with a 0.2 slope except the last output layers for the two structures.

C.2. Grid search

As mentioned in Section 5.1.2, we have conducted a grid search on the configurations for the network structures and then fixed the optimal parameters to make a fair comparison with our proposed framework. The grid search results for the CAE and AEflow structures based on the shallow-water dataset are summarized in Table C.6. We fix the learning rate as 10^{-3} except for the starred results due to training divergence, where we instead set the learning rate as 10^{-4} . All the networks have been trained with 30000 epochs to avoid underfitting, and the optimal configurations are highlighted in boldface.

Table C.5
Forward propagation for the AEflow baseline.

	Layer	Output Shape
Encoder	Input	(-1, 2, 128, 64)
	Conv2d-1	(-1, 4, 128, 64)
	ResBlock-1	(-1, 4, 128, 64)
	(-1, 4, 128, 64)
	ResBlock- N_{res}	(-1, 4, 128, 64)
	Conv2d-2	(-1, 4, 128, 64)
	Skip connection from ResBlock-1	(-1, 4, 128, 64)
	ConvBlock-1	(-1, 8, 64, 32)
	ConvBlock-2	(-1, 16, 32, 16)
	ConvBlock-3	(-1, 16, 16, 8)
	Conv2d-3	(-1, 8, 16, 8)
Decoder	Input	(-1, 8, 16, 8)
	Conv2d-1	(-1, 16, 16, 8)
	ConvTransposeBlock-1	(-1, 16, 32, 16)
	ConvTransposeBlock-2	(-1, 8, 64, 32)
	ConvTransposeBlock-3	(-1, 4, 128, 64)
	ResBlock-1	(-1, 4, 128, 64)
	(-1, 4, 128, 64)
	ResBlock- N_{res}	(-1, 4, 128, 64)
	Conv2d-2	(-1, 4, 128, 64)
	Skip connection from ResBlock-1	(-1, 4, 128, 64)
	Conv2d-3	(-1, 2, 128, 64)

Table C.6
Testing RMSEs for different configurations of AutoEncoder baselines (lower is better).

AutoEncoder		$K = 3$	$K = 5$	$K = 7$
CAE	$C = 16$	0.1215	0.1114	0.1044
	$C = 32$	0.1044	0.0841	0.0979
	$C = 64$	0.1016	0.1231*	0.1060*
AEflow	$N_{\text{res}} = 4$	0.0944	0.0782	0.1404*
	$N_{\text{res}} = 8$	0.0839	0.0937	0.1080
	$N_{\text{res}} = 12$	0.0938	0.1024	0.1390*

Appendix D. Algorithms for associated Kalman filter methods

In this section, we list the detailed update rules for the assimilation algorithms involved in our experiments, including Ensemble Kalman Filter (EnKF), stochastic EnKF (SEnKF) [5], deterministic EnKF (DEnKF) [5], Ensemble Transform Kalman Filter (ETKF) [5] and ETKF-Q as a special case of IEnKS-Q [30] (see [62] for more details).

D.1. EnKF

D.1.1. Forecast step

We perform the perturbed forward propagation:

$$x_k^{b,j} = \mathcal{M}_k \left(x_{k-1}^{a,j} \right) + \varepsilon_k^{M,j}, \quad \varepsilon_k^{M,j} \sim \mathcal{N} \left(0, \Sigma_k^M \right). \quad (\text{D.1})$$

D.1.2. Analysis step

i) Obtain the ensemble mean and anomaly matrix for the ensembles:

$$x_k^b = \frac{1}{N} \sum_{j=1}^N x_k^{b,j}, \quad (\text{D.2})$$

$$X_k^b = \frac{1}{\sqrt{N-1}} \left(x_k^{b,j} - x_k^b \right)_{j=1}^N \in \mathbb{R}^{n \times N}.$$

ii) Sample and calculate the perturbed observation vectors and the corresponding innovation vectors:

$$z_k^{o,j} = y_k^o + \varepsilon_k^{o,j}, \quad \varepsilon_k^{o,j} \sim \mathcal{N} \left(0, \Sigma_k^o \right), \quad (\text{D.3})$$

$$d_k^j = z_k^{o,j} - \mathcal{H}_k \left(x_k^{b,j} \right).$$

iii) Evaluate the Kalman gain matrix:

$$K_k = X_k^b Y_k^{bT} (Y_k^b Y_k^{bT} + \Sigma_k^o)^{-1}, \quad (\text{D.4})$$

where Y_k^b is the anomaly matrix of $\left\{ \mathcal{H}_k(x_k^{b,j}) \right\}_{j=1}^N$.

iv) Update the analyzed ensembles:

$$x_k^{a,j} = x_k^{b,j} + K_k d_k^j. \quad (\text{D.5})$$

D.2. SEnKF

Same as the previous EnKF, except that the update (D.4) is replaced with

$$K_k = X_k^b Y_k^{pT} (Y_k^p Y_k^{pT})^{-1}, \quad (\text{D.6})$$

where Y_k^p is the anomaly matrix of $\left\{ y_k^{p,j} \right\}_{j=1}^N$, and $y_k^{p,j} = \mathcal{H}_k(x_k^{b,j}) - \varepsilon_k^{o,j}$.

D.3. DEnKF

Same as the previous EnKF, except that the update (D.5) is replaced with

$$\begin{aligned} d_k &= y_k^o - \mathcal{H}_k(x_k^b), \\ x_k^a &= x_k^b + K_k d_k, \\ X_k^a &= X_k^b - \frac{1}{2} K_k Y_k^b \end{aligned} \quad (\text{D.7})$$

and the ensembles $\left\{ x_k^{a,j} \right\}_{j=1}^N$ are obtained from the ensemble mean x_k^a and the anomaly matrix X_k^a .

D.4. ETKF

Same as the previous EnKF, but rather than directly evaluating the Kalman matrix in (D.4), we calculate the transforming matrix

$$T_k = \left(I_N + Y_k^{bT} (\Sigma_k^o)^{-1} Y_k^b \right)^{-1} = \left(I_N + S_k^T S_k \right)^{-1} \quad (\text{D.8})$$

for $S_k = (\Sigma_k^o)^{-1/2} Y_k^b$, and then update the analyzed ensembles with the innovation vector $d_k = y_k^o - y_k^b$ as

$$\begin{aligned} w_k^a &= \left(I_N + Y_k^{bT} (\Sigma_k^o)^{-1} Y_k^b \right)^{-1} Y_k^{bT} (\Sigma_k^o)^{-1} (y_k^o - y_k^b) = T_k (\Sigma_k^o)^{-1} d_k, \\ x_k^{a,j} &= x_k^b + X_k^b \left(w_k^a + \sqrt{N-1} T_k^{1/2} U \right), \end{aligned} \quad (\text{D.9})$$

where U is a fixed random orthogonal matrix such that $U\mathbf{1} = \mathbf{1}$.

D.5. ETKF-Q

D.5.1. Initialization

Construct $U \in \mathbb{R}^{N \times (N-1)}$ such that

$$\begin{pmatrix} \frac{1}{\sqrt{N}} \mathbf{1}_N & U \end{pmatrix} \quad (\text{D.10})$$

is orthogonal, and let

$$V = \begin{pmatrix} \frac{1}{N} \mathbf{1}_N & \frac{1}{\sqrt{N-1}} U \end{pmatrix}, \quad (\text{D.11})$$

then

$$V^{-1} = \begin{pmatrix} \mathbf{1}_N & \sqrt{N-1} U \end{pmatrix}^T. \quad (\text{D.12})$$

D.5.2. Forecast step

i) Perform the forward propagation:

$$x_k^{f,j} = \mathcal{M}_k(x_{k-1}^{a,j}). \quad (\text{D.13})$$

ii) Obtain the mean and deviation for x_k^f :

$$\left(x_k^f, \Delta_x^f\right) = \left(x_k^{f,j}\right)_j \mathcal{U}. \quad (\text{D.14})$$

iii) Compute the eigendecomposition approximately:

$$\left(\Delta_x^f \Delta_x^{f\top} + \Sigma_k^M\right) V_k \approx V_k \Lambda_k, \quad (\text{D.15})$$

where $V_k \in \mathbb{R}^{n \times (N-1)}$ and $\Lambda_k \in \mathbb{R}^{(N-1) \times (N-1)}$ are the eigenvectors and eigenvalues of $\Delta_x^f \Delta_x^{f\top} + \Sigma_k^M$.

iv) Update the deviation and the ensembles:

$$\begin{aligned} \Delta_x^b &= V_k \Lambda_k^{1/2} \\ \left(x_k^{b,j}\right)_j &= \left(x_k^f, \Delta_x^b\right) \mathcal{U}^{-1}. \end{aligned} \quad (\text{D.16})$$

D.5.3. Analysis step

i) Calculate the ensemble mean and deviation matrix from the ensembles:

$$\left(x_k^b, \Delta_x^b\right) = \left(x_k^{b,j}\right)_j \mathcal{U}. \quad (\text{D.17})$$

ii) Calculate the mean and deviation for background estimates of observations:

$$\begin{aligned} y_k^{b,j} &= \mathcal{H}_k \left(x_k^{b,j}\right), \\ \left(y_k^b, \Delta_y^b\right) &= \left(y_k^{b,j}\right)_j \mathcal{U}. \end{aligned} \quad (\text{D.18})$$

iii) Obtain the transforming matrix:

$$T_k = \left(I_{N-1} + \Delta_y^{b\top} (\Sigma_k^o)^{-1} \Delta_y^b\right)^{-1} = \left(I_{N-1} + S_k^\top S_k\right)^{-1} \quad (\text{D.19})$$

for $S_k = (\Sigma_k^o)^{-1/2} \Delta_y^b$ and the normalized innovation vector

$$\delta_k = (\Sigma_k^o)^{-1/2} (y_k^o - y_k^b). \quad (\text{D.20})$$

iv) Update the analyzed ensembles:

$$\begin{aligned} w_k^a &= \left(I_{N-1} + \Delta_y^{b\top} \Sigma_k^o \Delta_y^b\right)^{-1} \Delta_y^{b\top} (\Sigma_k^o)^{-1} (y_k^o - y_k^b) = T_k S_k^\top \delta_k, \\ x_k^a &= x_k^b + \Delta_x^b w_k^a = x_k^b + \Delta_x^b T_k S_k^\top \delta_k, \\ \Delta_x^a &= \Delta_x^b \left(I_{N-1} + \Delta_y^{b\top} (\Sigma_k^o)^{-1} \Delta_y^b\right)^{-1/2}, \\ \left(x_k^{a,j}\right)_j &= \left(x_k^a, \Delta_x^a\right) \mathcal{U}^{-1}. \end{aligned} \quad (\text{D.21})$$

References

- [1] M. Amendola, R. Arcucci, L. Mottet, C.Q. Casas, S. Fan, C. Pain, P. Linden, Y.K. Guo, Data assimilation in the latent space of a convolutional autoencoder, in: M. Paszynski, D. Kranzlmüller, V.V. Krzhizhanovskaya, J.J. Dongarra, P.M. Sloot (Eds.), Computational Science – ICCS 2021, Springer International Publishing, Cham, 2021, pp. 373–386, https://doi.org/10.1007/978-3-030-77977-1_30.
- [2] T. Arcomano, I. Szunyogh, J. Pathak, A. Wikner, B.R. Hunt, E. Ott, A machine learning-based global atmospheric forecast model, Geophys. Res. Lett. 47 (9) (2020) e2020GL087776, <https://doi.org/10.1029/2020GL087776>.
- [3] R. Arcucci, L. Mottet, C. Pain, Y.K. Guo, Optimal reduced space for variational data assimilation, J. Comput. Phys. 379 (2019) 51–69, <https://doi.org/10.1016/j.jcp.2018.10.042>.
- [4] G. Artana, A. Camilleri, J. Carlier, E. Mémin, Strong and weak constraint variational assimilations for reduced order fluid flow modeling, J. Comput. Phys. 231 (8) (2012) 3264–3288, <https://doi.org/10.1016/j.jcp.2012.01.010>.
- [5] M. Asch, M. Bocquet, M. Nodet, Data Assimilation, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2016.
- [6] K. Atkinson, W. Han, Spherical Harmonics and Approximations on the Unit Sphere: An Introduction, Lecture Notes in Mathematics, Springer, 2012.
- [7] T. Bachlechner, B.P. Majumder, H. Mao, G. Cottrell, J. McAuley, Rezero is all you need: fast convergence at large depth, in: Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, PMLR, 2021, pp. 1352–1361.
- [8] D. Bank, N. Koenigstein, R. Giryes, Autoencoders, in: Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook, 2023, pp. 353–374.
- [9] M. Bemana, K. Myszkowski, H.P. Seidel, T. Ritschel, X-fields: implicit neural view-, light- and time-image interpolation, ACM Trans. Graph. 39 (6) (2020), <https://doi.org/10.1145/3414685.3417827>.
- [10] G. Berkooz, P. Holmes, J.L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, Annu. Rev. Fluid Mech. 25 (1) (1993) 539–575, <https://doi.org/10.1146/annurev.fl.25.010193.002543>.
- [11] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, Q. Tian, Accurate medium-range global weather forecasting with 3D neural networks, Nature 619 (7970) (2023) 533–538, <https://doi.org/10.1038/s41586-023-06185-3>.

- [12] S.L. Brunton, J.L. Proctor, J.N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci.* 113 (15) (2016) 3932–3937, <https://doi.org/10.1073/pnas.1517384113>.
- [13] K.J. Burns, G.M. Vasil, J.S. Oishi, D. Lecoanet, B.P. Brown, Dedalus: a flexible framework for numerical simulations with spectral methods, *Phys. Rev. Res.* 2 (2020) 023068, <https://doi.org/10.1103/PhysRevResearch.2.023068>.
- [14] Y. Cao, J. Zhu, Z. Luo, I. Navon, Reduced-order modeling of the upper tropical Pacific Ocean model using proper orthogonal decomposition, *Comput. Math. Appl.* 52 (8) (2006) 1373–1386, <https://doi.org/10.1016/j.camwa.2006.11.012>.
- [15] K. Chen, T. Han, J. Gong, L. Bai, F. Ling, J.J. Luo, X. Chen, L. Ma, T. Zhang, R. Su, et al., Fengwu: pushing the skillful global medium-range weather forecast beyond 10 days lead, preprint, arXiv:2304.02948, 2023.
- [16] P.Y. Chen, J. Xiang, D.H. Cho, Y. Chang, G.A. Pershing, H.T. Maia, M.M. Chiramonte, K.T. Carlberg, E. Grinspun, CROM: continuous reduced-order modeling of PDEs using implicit neural representations, in: *International Conference on Learning Representations*, 2023.
- [17] R.T. Chen, Y. Rubanova, J. Bettencourt, D.K. Duvenaud, Neural ordinary differential equations, *Adv. Neural Inf. Process. Syst.* 31 (2018).
- [18] R.T. Chen, B. Amos, M. Nickel, Learning neural event functions for ordinary differential equations, in: *International Conference on Learning Representations*, 2021.
- [19] Z. Chen, H. Zhang, Learning implicit fields for generative shape modeling, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5939–5948.
- [20] S. Cheng, J. Chen, C. Anastasiou, P. Angeli, O.K. Matar, Y.K. Guo, C.C. Pain, R. Arcucci, Generalised latent assimilation in heterogeneous reduced spaces with machine learning surrogate models, *J. Sci. Comput.* 94 (1) (2023), <https://doi.org/10.1007/s10915-022-02059-4>.
- [21] J. Chung, Ç. Gülçehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, preprint, arXiv:1412.3555, 2014.
- [22] P. Cizmas, A. Palacios, T. O'Brien, M. Syamlal, Proper-orthogonal decomposition of spatio-temporal patterns in fluidized beds, *Chem. Eng. Sci.* 58 (19) (2003) 4417–4427, [https://doi.org/10.1016/S0009-2509\(03\)00323-3](https://doi.org/10.1016/S0009-2509(03)00323-3).
- [23] M.C. Clare, O. Jamil, C.J. Morcrette, Combining distribution-based neural networks to predict weather forecast probabilities, *Q. J. R. Meteorol. Soc.* 147 (741) (2021) 4337–4357, <https://doi.org/10.1002/qj.4180>.
- [24] R.A. DeVore, R. Howard, C. Micchelli, Optimal nonlinear approximation, *Manuscr. Math.* 63 (1989) 469–478, <https://doi.org/10.1007/BF01171759>.
- [25] E. Dupont, H. Kim, S.M.A. Eslami, D.J. Rezende, D. Rosenbaum, From data to functa: your data point is a function and you can treat it like one, in: *Proceedings of the 39th International Conference on Machine Learning*, 2022.
- [26] E. Dupont, H. Loya, M. Alizadeh, A. Goliński, Y.W. Teh, A. Doucet, Coin++: neural compression across modalities, *Trans. Mach. Learn. Res.* (2022).
- [27] G. Evensen, *Data Assimilation*, 2nd edn, Springer, Berlin, Heidelberg, 2009.
- [28] A. Farchi, P. Laloyaux, M. Bonavita, M. Bocquet, Using machine learning to correct model error in data assimilation and forecast applications, *Q. J. R. Meteorol. Soc.* 147 (739) (2021) 3067–3084, <https://doi.org/10.1002/qj.4116>.
- [29] R. Fathony, A.K. Sahu, D. Willmott, J.Z. Kolter, Multiplicative filter networks, in: *International Conference on Learning Representations*, 2020.
- [30] A. Fillion, M. Bocquet, S. Gratton, S. Gürol, P. Sakov, An iterative ensemble Kalman smoother in presence of additive model error, *SIAM/ASA J. Uncertain. Quantificat.* 8 (1) (2020) 198–228, <https://doi.org/10.1137/19M1244147>.
- [31] S. Fresca, L. Dede', A. Manzoni, A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized pdes, *J. Sci. Comput.* 87 (2021) 1–36, <https://doi.org/10.1007/s10915-021-01462-7>.
- [32] J. Galewsky, R.K. Scott, L.M. Polvani, An initial-value problem for testing numerical models of the global shallow-water equations, *Tellus, Ser. A Dyn. Meteorol. Oceanogr.* (2004), <https://doi.org/10.3402/tellusa.v56i5.14436>.
- [33] A. Glaws, R. King, M. Sprague, Deep learning for in situ data compression of large turbulent flow simulations, *Phys. Rev. Fluids* 5 (2020) 114602, <https://doi.org/10.1103/PhysRevFluids.5.114602>.
- [34] A. Gruber, M. Gunzburger, L. Ju, Z. Wang, A comparison of neural network architectures for data-driven reduced-order modeling, *Comput. Methods Appl. Mech. Eng.* 393 (2022) 114764, <https://doi.org/10.1016/j.cma.2022.114764>.
- [35] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [36] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R.J. Hogan, E. Hölm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, J.N. Thépaut, The era5 global reanalysis, *Q. J. R. Meteorol. Soc.* 146 (730) (2020) 1999–2049, <https://doi.org/10.1002/qj.3803>.
- [37] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780, <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [38] C.M. Jiang, S. Esmailzadeh, K. Azizzadenesheli, K. Kashinath, M. Mustafa, H.A. Tchelepi, P. Marcus, Anandkumar A. Prabhat, Meshfreeflownet: a physics-constrained deep continuous space-time super-resolution framework, in: *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 2020, pp. 1–15, <https://doi.org/10.1109/SC41405.2020.00013>.
- [39] E. Kalnay, *Atmospheric Modeling, Data Assimilation and Predictability*, Cambridge University Press, 2002.
- [40] S. Kotsuki, Y. Sato, T. Miyoshi, Data assimilation for climate research: model parameter estimation of large-scale condensation scheme, *J. Geophys. Res., Atmos.* 125 (1) (2020) e2019JD031304, <https://doi.org/10.1029/2019JD031304>.
- [41] W.A. Lahoz, P. Schneider, Data assimilation: making sense of Earth observation, *Front. Environ. Sci.* 2 (2014) 16.
- [42] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, P. Battaglia, et al., Learning skillful medium-range global weather forecasting, *Science* 382 (6677) (2023) 1416–1421, <https://doi.org/10.1126/science.adi2336>.
- [43] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444, <https://doi.org/10.1038/nature14539>.
- [44] K. Lee, K.T. Carlberg, Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders, *J. Comput. Phys.* 404 (2020) 108973, <https://doi.org/10.1016/j.jcp.2019.108973>.
- [45] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, A. Anandkumar, Fourier neural operator for parametric partial differential equations, in: *International Conference on Learning Representations*, 2021.
- [46] Z.C. Lipton, J. Berkowitz, C. Elkan, A critical review of recurrent neural networks for sequence learning, preprint, arXiv:1506.00019, 2015.
- [47] Z. Long, Y. Lu, B. Dong, Pde-net 2.0: learning pdes from data with a numeric-symbolic hybrid deep network, *J. Comput. Phys.* 399 (2019) 108925, <https://doi.org/10.1016/j.jcp.2019.108925>.
- [48] L. Lu, P. Jin, G. Pang, Z. Zhang, G.E. Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.* 3 (3) (2021) 218–229, <https://doi.org/10.1038/s42256-021-00302-5>.
- [49] M. Lukoševičius, H. Jaeger, Reservoir computing approaches to recurrent neural network training, *Comput. Sci. Rev.* 3 (3) (2009) 127–149, <https://doi.org/10.1016/j.cosrev.2009.03.005>.
- [50] J.L. Lumley, *Stochastic Tools in Turbulence*, Academic Press, 1970.
- [51] Y. Luo, K. Ogle, C. Tucker, S. Fei, C. Gao, S. LaDeau, J.S. Clark, D.S. Schimel, Ecological forecasting and data assimilation in a data-rich era, *Ecol. Appl.* 21 (5) (2011) 1429–1442, <https://doi.org/10.1890/09-1275.1>.
- [52] A. Maas, A. Hannun, A. Ng, Rectifier nonlinearities improve neural network acoustic models, in: *Proceedings of the International Conference on Machine Learning*, 2013.

- [53] I. Mehta, M. Gharbi, C. Barnes, E. Shechtman, R. Ramamoorthi, M. Chandraker, Modulated periodic activations for generalizable local functional representations, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 14214–14223, <https://doi.org/10.1109/ICCV48922.2021.01395>.
- [54] P. Nadler, R. Arcucci, Y.K. Guo, Data assimilation for parameter estimation in economic modelling, in: 2019 15th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS), IEEE, 2019, pp. 649–656, <https://doi.org/10.1109/SITIS.2019.00106>.
- [55] S. Niu, Y. Luo, M.C. Dietze, T.F. Keenan, Z. Shi, J. Li, F. Stuart Chapin III, The role of data assimilation in predictive ecology, *Ecosphere* 5 (5) (2014) art65, <https://doi.org/10.1890/ES13-00273.1>.
- [56] M. Ohlberger, S. Rave, Reduced basis methods: success, limitations and future challenges, preprint, arXiv:1511.02021, 2015.
- [57] J.J. Park, P. Florence, J. Straub, R. Newcombe, S. Lovegrove, DeepSDF: learning continuous signed distance functions for shape representation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 165–174.
- [58] J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, A. Anandkumar, FourCastNet: a global data-driven high-resolution weather model using adaptive Fourier neural operators, preprint, arXiv:2202.11214, 2022.
- [59] S. Pawar, O. San, Equation-free surrogate modeling of geophysical flows at the intersection of machine learning and data assimilation, *J. Adv. Model. Earth Syst.* 14 (11) (2022) e2022MS003170, <https://doi.org/10.1029/2022MS003170>.
- [60] S.G. Penny, T.A. Smith, T.C. Chen, J.A. Platt, H.Y. Lin, M. Goodliff, H.D.I. Abarbanel, Integrating recurrent neural networks with data assimilation for scalable data-driven state estimation, *J. Adv. Model. Earth Syst.* 14 (3) (2022) e2021MS002843, <https://doi.org/10.1029/2021MS002843>.
- [61] E. Perez, F. Strub, H. de Vries, V. Dumoulin, A.C. Courville, Film: visual reasoning with a general conditioning layer, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2018.
- [62] M. Peyron, A. Fillion, S. Gürol, V. Marchais, S. Gratton, P. Boudier, G. Goret, Latent space data assimilation by using deep learning, *Q. J. R. Meteorol. Soc.* 147 (740) (2021) 3759–3777, <https://doi.org/10.1002/qj.4153>.
- [63] A. Pinkus, Basic properties of n-widths, Springer Berlin Heidelberg, Berlin, Heidelberg, 1985, pp. 9–38, https://doi.org/10.1007/978-3-642-69894-1_2.
- [64] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707, <https://doi.org/10.1016/j.jcp.2018.10.045>.
- [65] S. Rasp, P.D. Dueben, S. Scher, J.A. Weyn, S. Mouatadid, N. Thuerey, Weatherbench: a benchmark data set for data-driven weather forecasting, *J. Adv. Model. Earth Syst.* 12 (11) (2020) e2020MS002203, <https://doi.org/10.1029/2020MS002203>.
- [66] S. Rasp, N. Thuerey, Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: a new model for weatherbench, *J. Adv. Model. Earth Syst.* 13 (2) (2021) e2020MS002405, <https://doi.org/10.1029/2020MS002405>.
- [67] F. Regazzoni, L. Dedè, A. Quarteroni, Machine learning for fast and reliable solution of time-dependent differential equations, *J. Comput. Phys.* 397 (2019) 108852, <https://doi.org/10.1016/j.jcp.2019.07.050>.
- [68] S. Scher, G. Messori, Ensemble methods for neural network-based weather forecasts, *J. Adv. Model. Earth Syst.* 13 (2) (2021), <https://doi.org/10.1029/2020MS002331>.
- [69] P.J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* 656 (2010) 5–28, <https://doi.org/10.1017/S0022112010001217>.
- [70] M.G. Schultz, C. Betancourt, B. Gong, F. Kleinert, M. Langguth, L.H. Leufen, A. Mozaffari, S. Stadler, Can deep learning beat numerical weather prediction?, *Philos. Trans. R. Soc. A, Math. Phys. Eng. Sci.* 379 (2194) (2021) 20200097, <https://doi.org/10.1098/rsta.2020.0097>.
- [71] J. Sirignano, K. Spiliopoulos, Dgm: a deep learning algorithm for solving partial differential equations, *J. Comput. Phys.* 375 (2018) 1339–1364, <https://doi.org/10.1016/j.jcp.2018.08.029>.
- [72] L. Sirovich, Turbulence and the dynamics of coherent structures part I: coherent structures, *Q. Appl. Math.* 45 (3) (1987) 561–571, <https://doi.org/10.1090/qam/910462>.
- [73] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, G. Wetzstein, Implicit neural representations with periodic activation functions, *Adv. Neural Inf. Process. Syst.* 31 (2020) 7462–7473.
- [74] O. Talagrand, 4D-VAR: four-dimensional variational assimilation, in: Advanced Data Assimilation for Geosciences: Lecture Notes of the Les Houches School of Physics: Special Issue, Oxford University Press, 2014, <https://doi.org/10.1093/acprof:oso/9780198723844.003.0001>.
- [75] M. Tancik, P.P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J.T. Barron, R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, *Adv. Neural Inf. Process. Syst.* 33 (2020) 7537–7547.
- [76] A. Tangborn, S.Q. Zhang, Wavelet transform adapted to an approximate Kalman filter system, *Appl. Numer. Math.* 33 (1) (2000) 307–316, [https://doi.org/10.1016/S0168-9274\(99\)00097-5](https://doi.org/10.1016/S0168-9274(99)00097-5).
- [77] S.M. Udrescu, M. Tegmark, Ai Feynman: a physics-inspired method for symbolic regression, *Sci. Adv.* 6 (16) (2020) eaay2631, <https://doi.org/10.1126/sciadv.aay2631>.
- [78] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nat. Methods* 17 (2020) 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.
- [79] E. Weinan, B. Yu, The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems, *Commun. Math. Stat.* 6 (1) (2018) 1–12, <https://doi.org/10.1007/s40304-018-0127-z>.
- [80] H. Weller, J. Thuburn, C.J. Cotter, Computational modes and grid imprinting on five quasi-uniform spherical c grids, *Mon. Weather Rev.* 140 (8) (2012) 2734–2755, <https://doi.org/10.1175/MWR-D-11-00193.1>.
- [81] J.A. Weyn, D.R. Durran, R. Caruana, Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere, *J. Adv. Model. Earth Syst.* 12 (9) (2020) e2020MS002109, <https://doi.org/10.1029/2020MS002109>.
- [82] Z. Ye, H. Liu, Z. Wang, B. Dong, Analysis of the decoder width for parametric partial differential equations, preprint, arXiv:2306.14390, 2023.
- [83] Y. Yin, M. Kirchmeyer, J.Y. Franceschi, A. Rakotomamonjy, P. Gallinari, Continuous pde dynamics forecasting with implicit neural representations, in: International Conference on Learning Representations, 2023.
- [84] J. Zhai, S. Zhang, J. Chen, Q. He, Autoencoder and its various variants, in: 2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC), IEEE, 2018, pp. 415–419, <https://doi.org/10.1109/SMC.2018.00080>.