

Koopman Invertible Autoencoder: Leveraging Forward and Backward Dynamics for Temporal Modeling

Kshitij Tayal Arvind Renganathan Rahul Ghosh Xiaowei Jia Vipin Kumar
University of Minnesota University of Minnesota University of Minnesota University of Pittsburgh University of Minnesota
 tayal@umn.edu renga016@umn.edu ghosh128@umn.edu xiaowei@pitt.edu kumar001@umn.edu

Abstract—Accurate long-term predictions are the foundations for many machine learning applications and decision-making processes. However, building accurate long-term prediction models remains challenging due to the limitations of existing temporal models like recurrent neural networks (RNNs), as they capture only the statistical connections in the training data and may fail to learn the underlying dynamics of the target system. To tackle this challenge, we propose a novel machine learning model based on Koopman operator theory, which we call Koopman Invertible Autoencoders (KIA), that captures the inherent characteristic of the system by modeling both forward and backward dynamics in the infinite-dimensional Hilbert space. This enables us to efficiently learn low-dimensional representations, resulting in more accurate predictions of long-term system behavior. Moreover, our method's invertibility design enforces reversibility and consistency in both forward and inverse operations. We illustrate the utility of KIA on pendulum and climate datasets, demonstrating 300% improvements in long-term prediction capability for pendulum while maintaining robustness against noise. Additionally, our method demonstrates the ability to better comprehend the intricate dynamics of the climate system when compared to existing Koopman-based methods.

I. INTRODUCTION

Temporal data, prevalent in many applications such as climate, finance, and biomedicine, present a challenging problem for accurate long-term prediction and forecasting. Recurrent Neural Networks (RNNs) have gained significant attention for their ability to model sequential data by maintaining an internal time-evolving state. However, a primary concern in training and deploying RNNs is in their degraded performance over extended time horizons, which stems from the problem of exploding and vanishing gradients [35]. This gradient instability can result in slow convergence or even hinder learning completely, making it less suitable for capturing long-term dependencies in the data. To address this issue, researchers have proposed various methods, such as constraining the weight matrix to belong to the orthogonal group [25], using unitary hidden-to-hidden matrices [20], temporal convolutional networks [3], and other solutions [23]. Despite these efforts, achieving long-term memory is still an ongoing challenge and remains an active area of research.

Additionally, applying existing temporal models directly for scientific problems presents multiple obstacles: Firstly, accurate depiction of spatial and temporal processes within physical systems necessitates a substantial amount of training data [39],

which is often scarce in real-world situations. Secondly, existing empirical models remain limited in generalizing to scenarios that look different from training data. This is because they only establish statistical connections [32] between input and the targeted system variables but do not consider the inherent characteristics of the processes involved in the target system. Lastly, the relationships learned by these models are only valid for the specific distribution of forcing variables present in the training data, limiting their ability to generalize to scenarios not covered in the training set.

In recent times, Koopman-based models have gained attention as a promising alternative approach for modeling temporal data [8]. These models are based on the Koopman operator ([22], also see related work), transforming the original nonlinear system into an infinite-dimensional linear space. Koopman operator has three distinctive properties, which make it an ideal choice for temporal modeling (i) *Linearity* [24]: The Koopman operator turns the original nonlinear system into an infinite-dimensional linear system, which simplifies the process of capturing the inherent patterns and trends in the data, which is crucial for effective temporal modeling. (ii) *Global analysis* [34]: Unlike other linearization techniques (e.g., linearization around fixed points or periodic orbits), the Koopman operator provides a global perspective, capturing the overall behavior of the system rather than just local dynamic that can enhance the generalization capabilities of the model. (iii) *Invariant properties* [40]: The eigenfunctions and eigenvalues of the Koopman operator can reveal intrinsic properties of the system that remain unchanged under the system dynamics, which helps uncover hidden structures and patterns and makes it more robust to noise. These three distinctive properties of the Koopman operator make it a powerful tool for modeling temporal data.

However, using the Koopman operator for practical computations can be challenging because it is an infinite-dimensional operator [31]. Recently, researchers have developed techniques to approximate the Koopman operator using finite-dimensional representations extracted by autoencoder-based models, e.g., the Koopman Autoencoder (KAE) [29]. These models effectively reduce the complexity by creating a low-dimensional representation space in which the Koopman operator can be suitably approximated with a linear layer that accurately captures the underlying dynamics of the system. However, because of model

architecture, the information gleaned from Koopman-based models often is primarily based on forward dynamics, which overlooks the scope to acquire knowledge from backward dynamics. The fundamental goal of the forward run is to move from the present state to the subsequent state. Conversely, a backward run aims to go from the current state to the one that preceded it. Modeling backward dynamics ensures linearity in the low-dimensional space and also regularizes the forward run to be consistent.

One straightforward strategy is to have two separate linear layers [1], each dedicated to independently modeling the forward and backward runs. However, such an approach may have a limited capacity to accurately capture the intrinsic dynamics of the process due to its need for knowledge sharing between the forward and backward states. This paper presents a method to build a long-term predictive model by leveraging the Koopman analysis to capture both forward and backward dynamics in a unified model. In particular, we model the forward and backward dynamics in the low-dimensional space using an invertible neural networks model [21], which can establish explicit invertible mappings between the input and output spaces. As a result of this integrated approach, a single layer can be trained to learn both forward and backward processes. This unified model can leverage common knowledge between the two directions and enhances the ability to capture the process dynamics fully. To the best of our knowledge, this paper is the first to present an invertibility approach for learning the Koopman operator. While analogous concepts have been utilized in the field of inverse modeling, i.e., using observable data to infer hidden characteristics, and vice versa [41], our research lays the foundation for predictive models that incorporate the essential underlying dynamics and go beyond mere trend prediction, which is of great significance in temporal modeling.

Our main contributions are as follows.

- In this work, we present Koopman Invertible Autoencoders (KIA), a novel approach that harnesses both forward and backward dynamics for learning low-dimensional representations of temporal data and illustrates its utility on the pendulum and climate datasets.
- We accurately extracted the pendulum system's dynamics, handling both clean and noisy scenarios, and achieved a remarkable 300% improvement in long-term prediction accuracy.
- We also demonstrated the capability of our model to better comprehend the intricate dynamics of the climate system relative to existing Koopman-based approaches.

II. RELATED WORK

RNNs [46] and their variants have proved indispensable in dealing with sequential data, making strides in numerous applications such as language modeling [18], speech recognition [19], and time-series prediction [27]. However, despite the versatility of RNNs, they are plagued by issues such as vanishing and exploding gradients [5], hindering their ability to model long-term dependencies. Variants of RNNs such as Long

Short-Term Memory (LSTM) [14], Gated Recurrent Unit (GRU) [10], and Quasi-Recurrent Neural Networks (QRNN) [7] have been proposed to overcome this difficulty and have achieved remarkable results. These networks address the limitations of RNN by utilizing various gating mechanisms allowing them to control the information flow, retain long-term dependencies, and mitigate vanishing/exploding gradient problems.

However, RNN-based models remain limited for long-term prediction as they rely on complex non-linear temporal structures (e.g., LSTM unit) and can easily accumulate errors over time. A recent trend in machine learning is to explore transferring physical theories from existing physics-based models to enhance the capabilities and generalization of RNNs [43]. It combines the best of both worlds: the structure and explanatory power of physics-based models and the learning and predictive capabilities of RNNs. For instance, a physics-guided recurrent neural network model (PGRNN) [16] integrated the heat transfer process with the RNN model to effectively capture the long-term dependencies in the data, a task at which traditional RNNs failed. Other works also leveraged physics to enhance the long-period prediction of turbulent flows [4]. However, these methods rely on thorough knowledge of the system's physics, which limits their applicability. Furthermore, current model architectures, both generic and physics-based, predominantly rely on using forward dynamics for training, neglecting the potential of backward dynamics. This lack of integration due to network architecture presents an overlooked opportunity for enhancing these systems. In contrast, our approach builds upon the dynamical systems theory to assimilate both forward and backward dynamics through which it learns the true underlying processes. This significantly improves long-term prediction capabilities and offers an alternative to prevailing methodologies.

The theory of Koopman operators [22], established in 1931, has recently emerged as a groundbreaking framework for systematic linearization of complex dynamical system within an infinite-dimensional Hilbert space. Dynamic Mode Decomposition (DMD) [37] and Extended Dynamic Mode Decomposition (EDMD) [44] are two popular approaches for approximating this operator; however, they face issues of computational feasibility [8]. To tackle these challenges, researchers have begun to utilize data-driven approaches through deep learning to determine the Koopman operator from observed data [29], [40]. Most data-driven Koopman operator techniques rely on autoencoders [42] to transition from nonlinear to linear Koopman subspaces in order to identify the Koopman operator. For instance, the VAMP architecture [30] utilizes a time-lagged autoencoder and a custom variational score to find Koopman coordinates on a protein folding example. Similarly, researchers [6] used Koopman operator for sequence disentanglement by assuming that the underlying complex dynamics can be represented linearly in latent space. The focus of our work is on long-term temporal modeling whereby we want to model both forward and backward dynamics, which is significantly more manageable in a linearized space as backward dynamics are the inverse of forward space. A method most relevant

to our approach is C-KAE [2], where the authors developed two separate networks for learning forward and backward dynamics, which are then trained together using consistency loss. Nevertheless, training two separate networks overlooks the connection between forward and backward dynamics and consequently loses the opportunity to exploit shared knowledge. Additionally, optimizing two neural network to be inverse of each other is computationally difficult and unstable due to their stochastic nature. Recently, a new category of neural networks, called Invertible Neural Networks (INNs), was introduced [11], [21], based on the principles of normalizing flow. INNs are bijective functions that, by design, can simultaneously be trained on both forward and backward dynamics and exploit shared knowledge, making them an ideal choice for long-term temporal modeling. To the best of our knowledge, our work represents the first attempt to employ INNs in the modeling of the Koopman operator. While there exists an orthogonal line of research work ([17], [26], [33]) that utilized INNs for modeling autoencoders, their focus is on efficient modeling of observable functions. In contrast, our primary objective is the efficient modeling of the Koopman operator and our approach could be used with any autoencoders.

III. PROBLEM FORMULATION & PRELIMINARIES

A. Problem Formulation

Temporal data, denoted as $\{x_t\}_{t=1}^T$, can be interpreted as a series of observations from a dynamical system. This could represent anything that varies over time, for example, weather predictions, stock price etc. Consider the following discrete form of the system dynamics

$$x_{t+1} = F(x_t) + r_t, \quad x \in \mathcal{M} \subset \mathbb{R}^m, \quad (1)$$

where x_{t+1} represents the state of the system at the next time step, given its current state x_t and $r_t \in \mathcal{M}$ represents deviation from the true dynamics due to e.g., measurement errors or missing values. Here, the function $F(x_t)$ is an (possibly non-linear) update rule which describes how the state of the system evolves from one time step to the next. Additionally, $\mathcal{M}$ represents a finite-dimensional manifold, embedded in a higher-dimensional Euclidean space, denoted by $\mathbb{R}^m$. In this work, we focus on multi-step forecasting task of predicting the future observations given the current observation. Formally, we seek to learn function map F such that

$$x_{t+l} = x_t \circ F^l, \quad l = 1, 2, \dots, \quad (2)$$

where x_{t+l} signifies the state of the system at a future time step and $\circ$ denotes function composition. F^l indicates that we're applying the system dynamics repeatedly, l times in total, to the current state x_t . The above model assumes that future states x_{t+l} depend only on the current observation x_t and not on the information from a sequence of previous observations (between t and $t+l$).

B. Linear Invertible Neural Network (INNs)

INNs are bijective functions with a forward mapping $\mathcal{K} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and an inverse mapping $\mathcal{K}^{-1} : \mathbb{R}^n \rightarrow \mathbb{R}^n$. These mappings can be computed in closed-form manner, building upon the foundational works by Dinh et al., [11], [21]. To construct a linearly invertible neural network, we utilize the framework of the real non-volume preserving architecture proposed by [12]. The basic unit of this network comprises a reversible bijective network, incorporating two interconnected linear coupling layers. During the forward processing stage, as described by equation (3), the input vector, denoted as u , is partitioned into u_1 and u_2 . These halves are then linearly transformed using translation functions $t_i(\cdot)$, where $i \in 1, 2$, and they are modeled by linear neural networks trained in the forward pass. The output are the concatenation $[v_1, v_2]$, which are computed as

$$\begin{aligned} v_1 &= u_1 + t_2(u_2) \\ v_2 &= u_2 + t_1(v_1) \end{aligned} \quad (3)$$

Given the output $v = [v_1, v_2]$, the above expressions are easily invertible as follows:

$$\begin{aligned} u_2 &= v_2 - t_1(v_1) \\ u_1 &= v_1 - t_2(u_2) \end{aligned} \quad (4)$$

Importantly, the mappings t_i can be arbitrarily complex functions. In our implementation, we implement bijectivity by a fully connected linear layers with linear activations and use deep INN which is composed of a sequence of these reversible blocks. One advantage of INNs is that we can train them on the well-understood forward process $\mathcal{K}$ and get the inverse $\mathcal{K}^{-1}$ for free by running them backward. This provides a unique opportunity to develop a unified model that integrates both forward and inverse dynamical processes seamlessly.

C. Koopman Operator Theory

The Koopman theory proposes a framework that allows us to study complex systems in a linear space. According to this theory, every nonlinear dynamical system can be transformed into a form where the evolution of observations can be described by an infinite-dimensional Koopman operator. This operator operates on the space of all possible measurement functions, which can be fully represented through a linear map. Formally, we define the Koopman operator as $\mathcal{K}_F^\infty : \mathcal{G}(\mathcal{M}) \rightarrow \mathcal{G}(\mathcal{M})$, where $\mathcal{G}(\mathcal{M})$ is the space of real-valued measurement functions, represented by $g : \mathcal{M} \rightarrow \mathbb{R}$. The Koopman operator maps between function spaces and transforms the observations of the state to the next time step. Mathematically, the action of the Koopman operator on a measurement function g at time t is given by:

$$\mathcal{K}_F^\infty g(x_t) = g(\mathcal{K}_F^\infty(x_t)), \quad (5)$$

where $F(x_t)$ represents the transformed state of the system to the next time step. Also, the Koopman operator is a linear





Fig. 1: The left figure depicts our network architecture (KIA). The right figure showcases the various losses employed where observations x_t are inputted and transformed into a latent representation z_t through an encoder and propagated forward and backward using INNs

operator. This means that for any real numbers α and β , the following linearity property holds:

$$\mathcal{K}_F^\infty(\alpha g_1 + \beta g_2) = \alpha \mathcal{K}_F^\infty g_1 + \beta \mathcal{K}_F^\infty g_2$$

This property ensures that the Koopman operator preserves the linear structure of the measurement functions' space.

IV. KOOPMAN INVERTIBLE AUTOENCODER (KIA)

For forecasting future observations, an intuitive route is to design a neural network that effectively learns an approximation of the mapping denoted by F . However, a model developed using this approach ignores the underlying dynamics and can be hard to analyze. To address these challenges, we propose to develop the Koopman Invertible Autoencoder (KIA), which is built upon the Koopman theory and INNs. In the following, we first provide an overview of the Koopman Autoencoder and how it incorporates the concepts from the Koopman theory. Subsequently, we will describe the integration of INNs, focusing on how they facilitate bidirectional modeling and contribute to the robustness of our proposed model.

A. Koopman Autoencoder (KAE)

Due to the infinite-dimensional nature of $\mathcal{K}_F^\infty$ in Eq. (5), the practical use of the Koopman theory requires creating a finite-dimensional approximation that is capable of encapsulating the majority of the dynamics at play. We denote finite-dimensional Koopman operator by $\mathcal{K}_F$. The primary objective of the KAE model lies in the discovery of the matrices $\mathcal{K}_F$ and the nonlinear transformation of observation to the Koopman subspace that allows for an accurate recovery of the underlying dynamics. The Koopman autoencoder consists of an encoder γ_e , which maps observations to the Koopman subspace, and a decoder γ_d , which maps the Koopman subspace back to the observation space. Both encoder and decoder are represented by neural networks. The training of the networks is conducted by minimizing the

discrepancy between the input (x_t) and its corresponding output ($\tilde{x}_t = \gamma_d \circ \gamma_e(x_t)$). This guarantees that the encoder and decoder are solely responsible for encoding and decoding processes and do not learn any dynamics. This can be formally represented as $\gamma_d \circ \gamma_e \approx \mathbb{I}$, where $\mathbb{I}$ stands for the identity function. The loss function to train this network is given by:

$$\mathcal{L}_{Recon} = \frac{1}{n} \sum_{t=1}^n \|\gamma_d \circ \gamma_e(x_t) - x_t\|_2^2, \quad (6)$$

which measures the mean squared error between the reconstructed and original data points over a dataset of size n .

Modeling Forward Dynamics: In general, $\mathcal{K}_F$ prescribes a rule to move forward in time i.e. $\hat{x}_{t+1} = \gamma_d \circ \mathcal{K}_F \circ \gamma_e(x_t)$. We can leverage this relationship for multi-step forecasting as $\gamma_d \circ \mathcal{K}_F^l \circ \gamma_e(x_t) \approx x_t \circ F^l$, i.e., we iteratively obtain forward estimates. In order to accurately represent these forward dynamics, we apply a linear invertible neural network to serve as the approximate Koopman operator, denoted as $\mathcal{K}_F$. Our encoding network, denoted as γ_e , processes observations x_t to acquire a latent representation z_t , in which the dynamics become linear through the encoder structure. During the forward evolution step (Eq 3), the input vector, referred to as z_t , is divided into two equal parts, $z_{t,1}$ and $z_{t,2}$, which are then propagated forward as follow:

$$\begin{aligned} z_{t+1,1} &= z_{t,1} + t_2(z_{t,2}) \\ z_{t+1,2} &= z_{t,2} + t_1(z_{t+1,1}) \end{aligned} \quad (7)$$

where the concatenation of output $[z_{t+1,1}, z_{t+1,2}]$ from INNs represent the forward evolution in Koopman subspace. In our tests, we noticed that our models predict as well as generalize better in multi-step forecasting, compared with the traditional auto-regressive approach that computes one step forward at a




Fig. 2: Prediction errors over a time horizon of 2000 steps for clean and noisy pendulum observations with initial conditions $\theta = 0.8$ (top row) and $\theta = 2.4$ (bottom row). The first column shows the clean results, the second column shows results with small noise, and the third column shows results with large noise. (Best seen in color)

time. Given a choice of k forward prediction steps, we define the following forward dynamical loss term:

$$\mathcal{L}_{\text{fwd}} = \frac{1}{k * n} \sum_{l=1}^k \sum_{t=1}^n \|\gamma_d \circ \mathcal{K}_F^l \circ \gamma_e(x_t) - x_{t+l}\|_2^2, \quad (8)$$

B. Bidirectional Modeling

The majority of Koopman-based networks currently available [29], [40] primarily focus on modeling forward dynamics. A benefit of preserving linearity in the latent space is that the evolution matrix $\mathcal{K}_F$ can also be exploited for backward prediction via its inverse: $\mathcal{K}_F^{-1}$, i.e., $\bar{x}_{t-1} = \gamma_d \circ \mathcal{K}_F^{-1} \circ \gamma_e(x_t)$. KAE architectures that employ the liner layer [29], [40] to learn forward dynamics cannot learn the backward dynamics due to the high computational cost of matrix inversion, which is typically $O(n^3)$. Another approach to implementing backward dynamics is through independent linear layers [2], but this may limit the capacity to capture the intrinsic dynamics due to the lack of knowledge sharing between the forward and backward states. Earlier methods have also examined the integration of backward dynamics into their non-linear model, such as in bi-directional RNN [38]. However, the inherent nonlinearities of a typical neural network make it difficult to constrain the forward and backward models.

Our tests found that models trained for forward prediction typically produce poor, backward predictions. In contrast, considering both forward and backward dynamics can contribute to more effective training of KAE. Specifically, incorporating backward dynamics ensures linearity in the low-dimensional space while regularizing the consistency with the forward run. To address this limitation, we propose modeling the forward and backward dynamics in the low-dimensional space using INNs which are invertible by design and very efficient. In contrast to existing methods that use separate structures and parameters

for backward modeling, our model allows for the direct back prediction through the INN structure (Eq. 4). Specifically, the latent representation z_t ($\gamma_e \circ x_t$) is divided into two parts, z_{t1} and z_{t2} , which are then propagated backward as follow:

$$\begin{aligned} z_{t-1,2} &= z_{t,2} - t_1(z_{t,1}) \\ z_{t-1,1} &= z_{t,1} - t_2(z_{t-1,2}) \end{aligned} \quad (9)$$

where the concatenation of output $[z_{t-1,1}, z_{t-1,2}]$ from INNs represent the backward evolution in Koopman subspace. As with forward dynamics, we noticed that our models predict as well as generalize better in multi-step backcasting. Given a choice of k backward prediction steps, we define the following backward dynamical loss term:

$$\mathcal{L}_{\text{bwd}} = \frac{1}{k * n} \sum_{l=1}^k \sum_{t=1}^n \|\gamma_d \circ \mathcal{K}_F^{-l} \circ \gamma_e(x_t) - x_{t-l}\|_2^2, \quad (10)$$

We derive our KIA framework for analyzing temporal data by integrating all the above components. Our model undergoes training by minimizing a loss function whose minimizers guarantee that we achieve an optimized autoencoder and get accurate predictions over time by effectively capturing both forward and backward dynamics. We define the complete training loss as

$$\mathcal{L} = \lambda_{\text{Recon}} \mathcal{L}_{\text{Recon}} + \lambda_{\text{fwd}} \mathcal{L}_{\text{fwd}} + \lambda_{\text{bwd}} \mathcal{L}_{\text{bwd}}, \quad (11)$$

where λ_{Recon} , λ_{fwd} , λ_{bwd} , are parameters that balance between reconstruction, forward and backward prediction. Figure 1 (left) depicts our network design with an encoder, decoder, and an INN-based Koopman module for forward and backward dynamics computation. Figure 1 (right) illustrates our loss landscape.

TABLE I: Pendulum Results: KIA model shows significant error reduction, especially in the last 100 predictions, demonstrating effective capture of long-range temporal dependencies.

Model	θ	Prediction Error (Avg)		
		All Value	First 100	Last 100
KAE	0.8	0.273 $\pm$ 0.151	0.020 $\pm$ 0.005	0.514 $\pm$ 0.014
C-KAE	0.8	0.080 $\pm$ 0.044	0.009 $\pm$ 0.001	0.154 $\pm$ 0.003
RNN	0.8	0.026 $\pm$ 0.013	0.008 $\pm$ 0.001	0.050 $\pm$ 0.001
KIA	0.8	0.010 $\pm$ 0.002	0.011 $\pm$ 0.001	0.011 $\pm$ 0.001
KAE	2.4	0.416 $\pm$ 0.236	0.024 $\pm$ 0.010	0.794 $\pm$ 0.015
C-KAE	2.4	0.270 $\pm$ 0.153	0.022 $\pm$ 0.006	0.524 $\pm$ 0.010
RNN	2.4	0.043 $\pm$ 0.023	0.008 $\pm$ 0.001	0.083 $\pm$ 0.001
KIA	2.4	0.014 $\pm$ 0.004	0.009 $\pm$ 0.001	0.022 $\pm$ 0.001

V. EXPERIMENTS

To evaluate our proposed approach for long term temporal modeling, we perform a comprehensive study on two datasets, i.e., the pendulum dataset and the sea surface temperature dataset, and compare with state of the art Koopman-based approaches as well as baseline sequential models.

A. Nonlinear Pendulum

Dataset Description: We evaluate our models on a nonlinear pendulum, whose behavior is described by a second-order ordinary differential equation (ODE), where the angular displacement from equilibrium, denoted by θ , follows the equation $\frac{d^2\theta}{dt^2} = -\frac{g}{l} \sin(\theta)$. We use standard parameters for the length, l (1), and gravity's acceleration, g (9.8 m/s²). The pendulum's behavior is different based on the initial angle. A small initial angle results in simple harmonic motion, while a larger initial angle introduces non-linearity and complexity. With this in mind, we conducted experiments using two specific initial angles for oscillations, $\theta = 0.8$ and $\theta = 2.4$, over a period ranging from 0 to 400. We also employ a random orthogonal transformation in $\mathbb{R}^{64 \times 2}$ to convert our input data points into a high-dimensional space. This transformation results in observations that reside in $\mathbb{R}^{64}$, enhancing the representation of the original data. We collected 4,000 evenly spaced points in $\mathbb{R}^2$ over the time interval $t = [0, 400]$. This dataset was divided into training (400 points), validation (1,500 points), and testing (2,100 points) sets. The model was trained on the training set, hyperparameter optimized using the validation set, and the final resulting model was evaluated on the test set.

Baselines: *Koopman Auto Encoder (KAE)* [29]:

This network model approximates the Koopman operator by only looking at its forward dynamics. *Consistent Koopman Auto Encoder (C-KAE)* [2]:

This network model builds upon the original KAE and incorporates both forward and backward dynamics by using two separate Koopman operators. Further they enforce consistency between these operators by approximating one to be inverse of the other through consistency loss. *Long Short-Term Memory (LSTM)*: We further conducted a comparison against an LSTM, a variant of a recurrent neural network (RNN). This model is designed to learn a nonlinear function $x_t \rightarrow x_{t+1}$ directly. During the inference phase, this function takes x_t as input

TABLE II: Results for pendulum with Low noise.

Model	θ	Prediction Error (Avg)		
		All Value	First 100	Last 100
KAE	0.8	0.262 $\pm$ 0.140	0.078 $\pm$ 0.001	0.508 $\pm$ 0.020
C-KAE	0.8	0.250 $\pm$ 0.106	0.082 $\pm$ 0.003	0.421 $\pm$ 0.012
RNN	0.8	0.149 $\pm$ 0.058	0.076 $\pm$ 0.001	0.253 $\pm$ 0.012
KIA	0.8	0.109 $\pm$ 0.028	0.077 $\pm$ 0.001	0.162 $\pm$ 0.010
KAE	2.4	0.500 $\pm$ 0.274	0.029 $\pm$ 0.007	0.901 $\pm$ 0.014
C-KAE	2.4	0.367 $\pm$ 0.195	0.032 $\pm$ 0.004	0.662 $\pm$ 0.010
RNN	2.4	0.113 $\pm$ 0.060	0.026 $\pm$ 0.001	0.217 $\pm$ 0.004
KIA	2.4	0.042 $\pm$ 0.012	0.026 $\pm$ 0.001	0.065 $\pm$ 0.001

TABLE III: Results for pendulum with High noise.

Model	θ	Prediction Error (Avg)		
		All Value	First 100	Last 100
KAE	0.8	1.347 $\pm$ 0.715	0.211 $\pm$ 0.007	2.337 $\pm$ 0.033
C-KAE	0.8	0.687 $\pm$ 0.261	0.224 $\pm$ 0.010	1.042 $\pm$ 0.017
RNN	0.8	0.468 $\pm$ 0.188	0.214 $\pm$ 0.005	0.794 $\pm$ 0.019
KIA	0.8	0.298 $\pm$ 0.075	0.210 $\pm$ 0.005	0.440 $\pm$ 0.023
KAE	2.4	0.757 $\pm$ 0.362	0.063 $\pm$ 0.008	1.173 $\pm$ 0.016
C-KAE	2.4	0.619 $\pm$ 0.322	0.066 $\pm$ 0.005	1.102 $\pm$ 0.021
RNN	2.4	0.227 $\pm$ 0.111	0.077 $\pm$ 0.004	0.421 $\pm$ 0.009
KIA	2.4	0.126 $\pm$ 0.039	0.076 $\pm$ 0.003	0.198 $\pm$ 0.004

to predict $x_{t+1}, x_{t+2}, \dots, x_{t+k}$ recursively. We use two LSTM layers, each with a hidden dimension of 64.

All Koopman-based models, including our proposed method, utilize the same encoder-decoder architecture. Specifically, we employ a three-layer feed-forward network (128, 64, 8) with non-linear activations. While specialized architectures like convolutional neural networks with non-linear layers are another option, previous research has shown that these architectures do not offer any advantages for Koopman-based approaches [36]. All models underwent training for 500 epochs, utilizing Adam optimizers with identical learning rates. The training, validation, and test sets were consistent across all models. An early stopping criterion was also applied, terminating training if the validation loss did not decrease for 20 consecutive epochs. Moreover, the models were trained using the same random seeds to ensure consistency.

Experimental results. To evaluate prediction performance, we use scale invariant relative forecasting error at each timestep which is $\frac{\|\hat{x}_t - x_t\|_2}{\|x_t\|_2}$. In order to assess the long-term forecasting capabilities of our models, we perform forecasts over a horizon of 2,000 steps, using 30 different initial observations. The mean and standard deviation of prediction errors for these initial observations are presented in Table I. Additionally, we provide insights into prediction capabilities across different time spans by reporting the average prediction error over all timesteps, the first 100 timesteps, and the last 100 timesteps. TABLE IV: Varying amount of training data for pendulum with $\theta_0 = 2.4$.

Model	Training Points	Prediction Error (Avg)
RNN	200	0.129 $\pm$ 0.06
KIA	200	0.052 $\pm$ 0.04
RNN	300	0.087 $\pm$ 0.06
KIA	300	0.035 $\pm$ 0.03




Fig. 3: Visualization of Prediction Trajectories on clean (top) and noisy input(bottom).

The average predictions for condition $\theta_0 = 0.8$ are consistently lower than for $\theta_0 = 2.4$, indicating that $\theta_0 = 2.4$ introduces more complexity and non-linearity. The RNN model performs better in the first 100 timesteps, suggesting its strength in capturing short-term temporal dynamics. Still, the KIA model shows a remarkable 300% improvement in long-term prediction accuracy over RNN on average. We attribute this improvement to the KIA model's ability to capture long term trends by learning inherent characteristics via mapping forward and backward dynamics in a single network. C-KAE, which considers both forward and backward, outperforms the KAE, which emphasizes the importance of studying backward dynamics but lacks consistency as it uses separate layers.

Robustness comparison: Temporal physical processes often have noisy measurements due to sensor limitations and environmental conditions. We simulated the behavior by applying additive white gaussian noise with zero mean and varying standard. Specifically, to test model robustness, we evaluated KIA with perturbed inputs using small (0.1 std) and large (0.2 std) noise levels, and the results are depicted in Table II and Table III respectively. The noise disrupts the regular oscillation of the pendulum, causing slight deviations in its motion. Consequently, the system becomes more chaotic and difficult to accurately predict, resulting in increased errors as the noise levels rise. Interestingly, the KAE and C-KAE




Fig. 4: Example SST regions of the Persian Gulf (left) and Southeast Asia (right). The Persian Gulf, has a high temperature range of 23-35°C, displaying high variability, while Southeast Asia, has a range of 26-31°C, represents a low variability region.

models shows more performance deterioration compared to the RNN and KIA models, especially over longer time horizons. Also, in table III, we observe signs of overfitting in the KAE and C-KAE models, as indicated by lower errors in the initial 100 predictions and higher errors in the last 100 predictions. Figure 2 illustrates the prediction error of the pendulum over 2,000 steps for clean and noisy observations. Comparing the RNN and KIA models, we find that KIA outperforms RNN, particularly over longer time horizons, even in the presence of noise. This highlights the robustness of KIA in effectively handling variations and disturbances in the input signals while capturing the true dynamics.

Effect of the size of the Training Data: Furthermore, to evaluate the influence of the labeled data size, we conducted experiments by utilizing different proportions of the training data on the top-performing RNN and KIA models. The results are presented in Table IV, where we compare the average prediction error obtained using subsamples of 200 and 300 instances from the training datasets. The KIA model consistently outperformed the RNN model in test error accuracy, even with fewer labeled data points, indicating its superior robustness and efficiency in using available data for precise predictions.

Visualization of Prediction Trajectories: To gain further insights, we visualize the prediction trajectories of our models for an initial condition of $\theta_0 = 2.4$. Figure 3 shows the trajectories generated using clean and noisy inputs, with the true trajectory shown in violet. Both KIA and RNN models are shown to excel in capturing the pendulum dynamics, regardless of the presence of noise. The red and green lines align closely with the true trajectory, indicating their effectiveness. However, the predicted trajectories of the KAE and C-KAE models exhibit increasing deviations from the true trajectory, especially over longer time periods. The introduction of noisy inputs further amplifies these deviations, leading to a more pronounced divergence from the actual trajectory. This highlights the advantage of using KIA which utilize both forward and backward dynamics to learn underlying dynamics.

B. Sea Surface Temperature Data

Dataset Description: In this experiment, we analyzed the NOAA Optimal Interpolation SST High Resolution dataset [15], which provides daily sea-surface temperature measurements at a spatial resolution of 0.25° . We focused on two specific regions: the Persian Gulf (Lat: 12.5° to 30° N, Long: 31.25° to 68.75° E) and Southeast Asia (Lat: -10° to 7.5° N, Long: 105° to 142.5° E) (see Fig 4). The Persian Gulf has higher sea




Fig. 5: Prediction errors over a time horizon of 180 days for Persian Gulf (left) and Southeast Asia (right)

TABLE V: Sea surface temperature prediction error statistics: mean and standard deviation.

Model	Regions	Prediction Error (Avg)
KAE	PG	0.99 ± 0.06
C-KAE	PG	1.04 ± 0.06
ConvLSTM	PG	1.93 ± 0.10
Mean	PG	0.86 ± 0.04
Persistence	PG	2.13 ± 0.05
KIA	PG	0.73 ± 0.05
KAE	SE	0.76 ± 0.06
C-KAE	SE	0.75 ± 0.04
ConvLSTM	SE	1.08 ± 0.03
Mean	SE	0.84 ± 0.03
Persistence	SE	1.10 ± 0.02
KIA	SE	0.67 ± 0.04

surface temperatures and greater seasonal variations compared to Southeast Asia, making it more nonlinear from a modeling perspective. We used five years of data, with three years for training, one year for validation, and one year for testing. The performance of our models was evaluated using Celsius MAE as the error metric. This metric calculates the average absolute difference, in degrees Celsius, between the predicted and the actual temperatures, regardless of their direction (positive or negative). Forecasting climate patterns is an inherently difficult task [9] due to the intricate interactions of a complex system. However, the input dynamics exhibit non-stationary periodic structures [2], suggesting the use of Koopman-based methods.

Baselines: We incorporate all Koopman-based baselines from Pendulum experiments, along with Persistence, Climate Mean, and replace LSTM with ConvLSTM. Persistence and Climate Mean are commonly used as baseline methods for climate predictions, offering different advantages based on the

TABLE VI: Mean error of K-day ahead sea surface temperature prediction.

Model	Regions	Prediction Error (Avg)				
		1 day	7 day	14 day	21 day	30 day
KAE	PG	0.59	0.61	0.65	0.69	0.77
C-KAE	PG	0.56	0.60	0.63	0.67	0.76
ConvLSTM	PG	0.51	0.67	0.82	0.91	0.98
Mean	PG	0.81	0.81	0.81	0.81	0.81
Persistence	PG	0.21	0.59	0.78	0.96	1.19
KIA	PG	0.48	0.57	0.63	0.67	0.74
KAE	SE	0.55	0.57	0.59	0.63	0.65
C-KAE	SE	0.54	0.56	0.59	0.62	0.64
ConvLSTM	SE	0.55	0.65	0.76	0.88	0.92
Mean	SE	0.75	0.75	0.75	0.75	0.75
Persistence	SE	0.24	0.61	0.68	0.72	0.77
KIA	SE	0.54	0.56	0.58	0.61	0.63

context and time scale. Persistence is effective for short-term forecasting, assuming stable climate conditions, while Climate Mean incorporates historical data for more reliable long-term predictions. Any useful model is expected to outperform these baselines. **Persistence:** This forecasting model assumes that future climate conditions will be similar to the current one, making identical predictions for each time step. **Climate Mean:** This model derives the average climate state from historical data of the same date. It suggests that the future climatic conditions will closely resemble the mean state derived from historical observations spanning previous years. **ConvLSTM:** is a spatiotemporal model specifically designed for data like SST, combining spatial processing with temporal dependencies using convolutional neural networks and LSTMs. It is widely adopted and used as an alternative to LSTM for spatio-temporal modeling.

Experimental results: In *Long-term Prediction Testing* (as done in [1], [36]), we generate predictions for the next 180 days on 30 different initial days of the year to evaluate the accuracy and reliability of the model over an extended period. Table V shows prediction error averages, and Fig. 6 visually compares different model outputs. The results show that our model's predictions for the 60-day, 120-day, and 180-day forecasts are closer to the ground truth compared to other models. To enhance differentiation in the model's output, we use a color scheme that represents the temperature difference with the ground truth. Grey indicates a perfect prediction, while red and blue represent warmer and colder predictions, respectively. The initial state and ground truth temperatures are in Celsius. The oval, rectangle, and triangle regions were used for clear differentiation. The visual comparison shows that our model, KIA, performs better than CKAE and KAE models, which tend to overpredict. The prediction results show that KIA can generalize to unseen climate scenarios, which is significant considering the challenges involved in predicting climate data. Furthermore, Fig 5 provides a visualization of the prediction error spanning a 180-day horizon, revealing interesting insights about different models' performance. While the persistence model is more accurate in short-term predictions, the KIA model outperforms the persistence and mean models when evaluating their overall performance. This indicates that the KIA model offers better long-term prediction capabilities, making it a more reliable choice for forecasting.

Additionally, we visualize the prediction error over the 180-day horizon in Fig 5, showing that while short-term predictions are more accurate with the persistence model, the KIA model outperforms both the persistence model and the mean model when considering overall performance.

Impact of Timescales on K-Day Ahead Predictions:

In some days of a year, accurate forecasting can be more challenging due to factors like seasonality, irregular events, or volatility. Analyzing different timescales [13] helps identify these days and assess their impact on model prediction. This information aids in resource allocation and the utilization of alternative forecasting methods during those periods. For this, we employ K day ahead forecasts for each day in a test year,




Fig. 6: We depict the future predictions for a single initial state over a horizon of 60, 120, and 180 days. The first column represents the consistent initial state, the second column shows the ground truth, and the remaining columns provide the model's predictions for days 60, 120, and 180. (best seen in color).




Fig. 7: We illustrate 30-day future forecasts for a random initial states. The first column shows the initial state, the second column presents the 30-day ahead ground truth, and the rest represent mean output, Persistence output, and KIA output.(best seen in color).

where each prediction is based on input from K days in the past. In our case, we evaluated our model for K in $\{1, 7, 14, 21$ and $30\}$ days. Table VI showcases the mean prediction error values (standard deviation is omitted for brevity). We observe that all Koopman-based models do an equally good job for short-term prediction. Also, the persistence model is the best if we are interested in a 1-day ahead forecast. However, even in the short term, from $K=7$ days onwards, our approach surpasses the persistence and mean models. Fig. 7 displays the 30-day KIA forecast for randomly selected initial states, which is visually compared with the mean and persistence models. We can observe that for $K=30$, the mean is overpredicting, persistence is underpredicting, and KIA output is considerably closer to the actual data. The prediction results are highly significant, as successfully predicting climate data and surpassing the mean and persistence models in the short term is challenging. Fig. 8 further visualizes the prediction error for the entire year for $K = 1$ and $K = 30$ day. From the $K = 30$ plot, we can identify that the later part of the years is difficult, where the models struggle to provide better predictions than the mean. In addition, from the plot, we can also identify dates like the 150th day of the year, where the mean is doing better than all the models. These results show that the KIA can capture system dynamics better than other related approaches and can be helpful for domain scientists for long-term prediction.

VI. DISCUSSION AND CONCLUSION

This paper proposes KIA, a long-term predictive model for temporal data. Our approach is built upon the principles of the Koopman theory, which enables the approximation of dynamical systems using linear evolution matrices. The key contribution of our method lies in its incorporation of both




Fig. 8: Prediction error for 1-day ahead (left) versus 30-day ahead (right) forecasts for southeast asia

forward and backward dynamics, achieved through a invertible neural network architecture that ensures consistency between these two directions. We successfully demonstrated the model's capability to accurately capture long-term trends, recover underlying dynamics, and generalize to unseen scenarios through experiments on pendulum and climate datasets. Although our improvements have been validated using these specific datasets, our architecture's versatility allows for its application in various other forecasting tasks that involve backward dynamics. For instance, consider the video captioning task [47], where the temporal order of video frames in both forward and backward directions plays a crucial role in capturing the movements of significant objects from different perspectives. Previous studies have employed graph-based networks to capture the relationships between objects in videos. In this context, our approach presents an alternative solution, thereby enabling a comprehensive understanding and captioning of the visual content. Furthermore, our method exhibits potential [28] for integration within Transformer models [45], a popular architecture widely used for sequence modeling tasks. The integration

holds promise for addressing more complex temporal problems across diverse domains, further improving the accuracy and generalizability of forecasting systems.

ACKNOWLEDGMENT

This research is supported by NSF awards 1934721 (Harnessing the Data Revolution program) and 2134904 (Advancing Deep Learning for Inverse Modeling). We would also like to acknowledge the Minnesota Supercomputing Institute for providing access to computing facilities.

REFERENCES

- [1] Omri Azencot et al. Consistent dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 18(3):1565–1585, 2019.
- [2] Omri Azencot et al. Forecasting sequential data using consistent koopman autoencoders. In *International Conference on Machine Learning*, pages 475–485. PMLR, 2020.
- [3] Shaojie Bai et al. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- [4] Tianshu Bao et al. Physics guided neural networks for spatio-temporal super-resolution of turbulent flows. In *Uncertainty in Artificial Intelligence*, pages 118–128. PMLR, 2022.
- [5] Yoshua Bengio et al. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- [6] Nimrod Berman et al. Multifactor sequential disentanglement via structured koopman autoencoders. In *The Eleventh International Conference on Learning Representations*, 2023.
- [7] James Bradbury et al. Quasi-recurrent neural networks. *arXiv preprint arXiv:1611.01576*, 2016.
- [8] Steven L Brunton et al. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2022.
- [9] Matthew Chantry et al. Opportunities and challenges for machine learning in weather and climate modelling: hard, medium and soft ai. *Philosophical Transactions of the Royal Society A*, 379(2194):20200083, 2021.
- [10] Rahul Dey et al. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE, 2017.
- [11] Laurent Dinh et al. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [12] Laurent Dinh et al. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [13] Sijie He et al. Sub-seasonal climate forecasting via machine learning: Challenges, analysis, and advances. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 169–177, 2021.
- [14] Sepp Hochreiter et al. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [15] James W Hurrell et al. A new sea surface temperature and sea ice boundary dataset for the community atmosphere model. *Journal of Climate*, 21(19):5145–5153, 2008.
- [16] Xiaowei Jia et al. Physics guided rnns for modeling dynamical systems: A case study in simulating lake temperature profiles. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pages 558–566. SIAM, 2019.
- [17] Yuhong Jin, Lei Hou, Shun Zhong, Haiming Yi, and Yushu Chen. Invertible koopman network and its application in data-driven modeling for dynamic systems. *Mechanical Systems and Signal Processing*, 200:110604, 2023.
- [18] Akos Kádár et al. Representation of linguistic form and function in recurrent neural networks. *Computational Linguistics*, 43(4):761–780, 2017.
- [19] Shigeaki Karita et al. A comparative study on transformer vs rnn in speech applications. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 449–456. IEEE, 2019.
- [20] Giancarlo Kerg et al. Non-normal recurrent neural network (nnRNN): Learning long time dependencies while improving expressivity with transient dynamics. In *Advances in Neural Information Processing Systems*, pages 13591–13601, 2019.
- [21] Durk P Kingma et al. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- [22] Bernard O Koopman. Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931.
- [23] Guokun Lai et al. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 95–104, 2018.
- [24] Yueheng Lan et al. Linearization in the large of nonlinear systems and koopman operator spectrum. *Physica D: Nonlinear Phenomena*, 242(1):42–53, 2013.
- [25] Mario Lezcano-Casado et al. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *International Conference on Machine Learning*, pages 3794–3803. PMLR, 2019.
- [26] Fan Li, Dong Liang, Jing Lian, Qidong Liu, Hegui Zhu, and Jizhao Liu. Invka: Gait recognition via invertible koopman autoencoder. *arXiv preprint arXiv:2309.14764*, 2023.
- [27] Bryan Lim et al. Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A*, 379(2194):20200209, 2021.
- [28] Xuanqing Liu et al. Learning to encode position for transformer with continuous dynamical model. In *International conference on machine learning*, pages 6327–6335. PMLR, 2020.
- [29] Bethany Lusch et al. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018.
- [30] Andreas Mardt et al. Vampnets for deep learning of molecular kinetics. *Nature communications*, 9(1):1–11, 2018.
- [31] Alexandre Mauroy et al. *Koopman operator in systems and control*. Springer, 2020.
- [32] Rodrigo F Mello et al. *Machine learning: a practical approach on the statistical learning theory*. Springer, 2018.
- [33] Yuhuang Meng, Jianguo Huang, and Yue Qiu. Physics-informed invertible neural network for the koopman operator learning. *arXiv preprint arXiv:2306.17396*, 2023.
- [34] Igor Mezić. Analysis of fluid flows via spectral properties of the koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.
- [35] Razvan Pascanu et al. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318, 2013.
- [36] Julian Rice et al. Analyzing koopman approaches to physics-informed machine learning for long-term sea-surface temperature forecasting. *arXiv preprint arXiv:2010.00399*, 2020.
- [37] Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.
- [38] Mike Schuster et al. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [39] Sima Siami-Namini et al. The performance of lstm and bilstm in forecasting time series. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 3285–3292. IEEE, 2019.
- [40] Naoya Takeishi et al. Learning Koopman invariant subspaces for dynamic mode decomposition. In *Advances in Neural Information Processing Systems*, pages 1130–1140, 2017.
- [41] Kshitij Tayal et al. Invertibility aware integration of static and time-series data: An application to lake temperature modeling. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pages 702–710. SIAM, 2022.
- [42] Rui Wang et al. Koopman neural operator forecaster for time-series with temporal distributional shifts. In *The Eleventh International Conference on Learning Representations*, 2023.
- [43] Jared Willard et al. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4):1–37, 2022.
- [44] Matthew O Williams et al. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25:1307–1346, 2015.
- [45] Thomas Wolf et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45, 2020.
- [46] Yong Yu et al. A review of recurrent neural networks: Lstm cells and network architectures. *Neural computation*, 31(7):1235–1270, 2019.
- [47] Junchao Zhang et al. Object-aware aggregation with bidirectional temporal graph for video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8327–8336, 2019.