

KALIKO: Kalman-Implicit Koopman Operator Learning For Prediction of Nonlinear Dynamical Systems

Albert H. Li*, Ivan Dario Jimenez Rodriguez*, Joel W. Burdick, Yisong Yue, Aaron D. Ames

Abstract—Long-horizon dynamical prediction is fundamental in robotics and control, underpinning canonical methods like model predictive control. Yet, many systems and disturbance phenomena are difficult to model due to effects like nonlinearity, chaos, and high-dimensionality. Koopman theory addresses this by modeling the linear evolution of embeddings of the state under an infinite-dimensional linear operator that can be approximated with a suitable finite basis of embedding functions, effectively trading model nonlinearity for representational complexity. However, explicitly computing a good choice of basis is nontrivial, and poor choices may cause inaccurate forecasts or overfitting. To address this, we present Kalman-Implicit Koopman Operator (KALIKO) Learning, a method that leverages the Kalman filter to *implicitly* learn embeddings corresponding to latent dynamics without requiring an explicit encoder. KALIKO produces interpretable representations consistent with both theory and prior works, yielding high-quality reconstructions and inducing a globally linear latent dynamics. Evaluated on wave data generated by a high-dimensional PDE, KALIKO surpasses several baselines in open-loop prediction and in a demanding closed-loop simulated control task: stabilizing an underactuated manipulator’s payload by predicting and compensating for strong wave disturbances.

I. INTRODUCTION

Dynamical modeling and prediction are fundamental problems in robotics and control, forming the bedrock for techniques like model predictive control, state estimation, and more. However, obtaining accurate analytical models in unstructured settings can be challenging, especially for nonlinear systems that may be subject to high-dimensional and/or chaotic spatiotemporal phenomena like turbulence or surface waves. This motivates a large body of work on fitting predictive models from data [1], [2], [3], [4].

One appealing approach is Koopman operator theory, which models the *linear* evolution of state embeddings¹ under the infinite-dimensional *Koopman operator* \mathcal{K} . In this lifted space, nonlinear dynamics become linear, effectively trading nonlinearity for representational complexity [2]. In practice, one builds finite approximations of the Koopman operator by *explicitly* parameterizing embedding functions via black-box neural networks [5], [3], fixed “dictionaries” of basis functions (e.g., polynomials) [6], or hybrid approaches that learn these dictionaries [7] (see Fig. 1A). However, selecting embeddings that capture an approximately invariant subspace is difficult; poor choices produce spurious modes or overfitted models and often necessitate multi-objective losses with carefully tuned weights [5], [8], [3].

* Equal contribution. A. H. Li is supported by Dow Award 227027AW. Authors are all with Caltech.

¹We use the term “embedding” rather than “observable” to avoid confusion with “measurements” and “observations” in Kalman filtering.

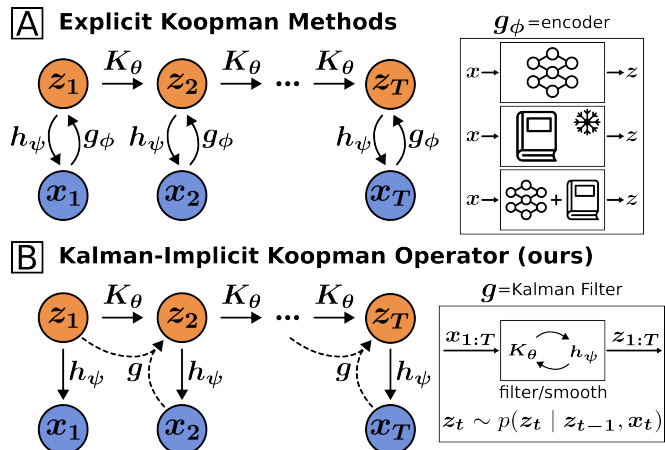


Fig. 1: (A) Most methods for learning the linear Koopman dynamics K_θ are *explicit*, parameterizing an encoder g_ϕ that maps data x to high-dimensional Koopman embeddings z . Examples include neural networks, fixed dictionaries of basis functions, or learnable dictionaries. (B) KALIKO (ours) instead *implicitly* learns Koopman embeddings by letting g be a Kalman filter and smoother, which is composed of only two models: the latent dynamics K_θ and decoder h_ψ . We show that these implicitly-recovered representations are highly interpretable and yield strong open-loop prediction and closed-loop control performance.

This work presents an alternative approach: *Kalman-Implicit Koopman Operator* (KALIKO) Learning. Unlike these explicit approaches, KALIKO reframes the problem of encoding data x into embeddings z as a Bayesian state estimation problem that infers latent states $z_{1:T}$ from data trajectories $x_{1:T}$. Thus, KALIKO *implicitly* represents an encoder via an (extended) Kalman filter, allowing it to circumvent the challenge of parameterizing an embedding function (see Fig. 1B). We show that despite its implicit structure, KALIKO learns highly interpretable dynamics consistent with Koopman theory and prior works (e.g., eigenfunctions for orbits, limit-cycles, and equilibria), achieves superior prediction performance over several baselines, and translates its capabilities to a challenging closed-loop stabilization task.

A. Related Work

Computing Koopman Embeddings. A standard data-driven method is *Dynamic Mode Decomposition* (DMD) [9], [2], which fits a least-squares linear map between state snapshots, projecting \mathcal{K} onto the span of embedding coordinates of those snapshots. As DMD operates on raw measurements, it cannot in general realize a nonlinear change of coordinates; *Extended DMD* (EDMD) addresses this by regressing in a lifted space built from a *dictionary* of embedding functions, like polynomials, which must be carefully constructed [10],

[2]. Time-delay coordinates provide another practical lifting technique: Delay-DMD constructs Hankel embeddings from past samples to capture approximately Koopman-invariant structure [11]. Yet another method is *Kernel DMD* (KDMD), which implements EDMD *implicitly* by only learning inner products between the embedded data, avoiding explicit enumeration of the dictionary [3], [12]. Lastly, data-driven methods attempt to learn these embeddings, either by learning dictionary elements or by training a deep autoencoder alongside a latent dynamics model [5], [13], [8], [3], [7].

These diverse approaches share canonical failure modes. Even for linear systems, a single spurious basis function can induce severe overfitting [3]. In the data-driven setting, practitioners counter this with heavy regularization [3], [8], complex multi-term losses [5], [14], and/or highly-sensitive weight initializations and hyperparameters [14]. Conversely, convergence of learned dictionaries is often unclear [8] and poor choices of the above can just as easily trap solutions in bad local minima, leading to underfitting [14], [15].

Kalman Filters and Dynamical Learning. Learnable Kalman filters (KFs) jointly train a latent dynamics model and decoder that maximize the likelihood of data trajectories. The decoder $p(\mathbf{x}_t \mid \mathbf{z}_t)$ is trained together with the dynamics $p(\mathbf{z}_t \mid \mathbf{z}_{t-1})$ so that new data points \mathbf{x}_t accurately update the posterior belief over the latent state \mathbf{z}_t . Unlike autoencoder-based approaches that learn an encoder/decoder independently of the dynamics, KFs couple representation and prediction by *interleaving* prediction and measurement updates (Sec. II-B), which encourages embeddings suitable for both long-horizon prediction and reconstruction.

Prior work has leveraged this insight to learn latent dynamics for image sequences [16], [17] by *compressing* high-dimensional observations into a low-dimensional latent state. In line with Koopman theory, KALIKO instead *lifts* the dimension to find an embedding space where the latent dynamics are approximately linear. Moreover, almost all prior works learn filters for state estimation [18], [19], [20], [21] or simple open-loop prediction tasks [16], [17]. In this work, we show that KALIKO’s predictions can also be used in a closed-loop control task.

Kalman Filters and Koopman. Many works have combined KFs with Koopman methods. Most commonly, KFs are used for latent state estimation *after* learning some linear dynamics [22], [23], [24]. Others have used KFs for spectral analysis [25], [26]. Closest to our work are system identification or operator learning methods, but unlike KALIKO, assume the existence of embeddings *a priori* [27] or train an encoder rather than a decoder [28].

Time Series Prediction. The machine learning community has used methods like recurrent neural networks [29] to model sequential data and learn dynamical systems [30], especially those that are hard to model, like continuum robot dynamics [31], human motion [32], and video [33]. More recently, transformer-based models have been applied to many phenomena like weather, traffic patterns, electricity usage, etc. [34]. It has been shown that LLM backbones like GPT2 are effective for time series prediction [35], prompting

research on time series tokenization [36].

The recent work “Koopman-Kalman Enhanced Variational Autoencoder” (K2VAE) at first appears similar to KALIKO, but several details distinguish them. First, K2VAE learns an explicit encoder. Second, K2VAE uses its KF to drive learning error to 0 with a dynamics \mathbf{A} distinct from \mathcal{K} , while KALIKO uses the KF directly on the learned operator. Lastly, KALIKO uses several techniques not present in K2VAE that are crucial for performance (see Sec. III-C and V-D). We compare to K2VAE in Sec. V-C.

B. Contributions

Our chief contribution is KALIKO, a method that uses the Kalman filter to *implicitly* learn a Koopman embedding space governed by globally linear latent dynamics. Second, we show that KALIKO’s embeddings produce faithful reconstructions, are highly interpretable, and align with theoretical expectations. Lastly, on complex wave data generated by a high-dimensional PDE, KALIKO outperforms several baselines on both an open-loop prediction task and a challenging simulated closed-loop stabilization task where it must stabilize an underactuated manipulator’s payload by predicting and compensating for wave disturbances.

II. MATHEMATICAL PRELIMINARIES

Let $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ be the state of a discrete-time nonlinear system which generates the data of interest,

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t), \quad (1)$$

and $\mathbf{z} \in \mathcal{Z} \subseteq \mathbb{R}^m$ be the state of some corresponding desired linear latent dynamics

$$\mathbf{z}_{t+1} = \mathbf{K}_\theta \mathbf{z}_t, \quad \mathbf{x}_t = \mathbf{h}_\psi(\mathbf{z}_t), \quad (2)$$

where θ, ψ are learnable parameters and $\mathbf{h} : \mathcal{Z} \rightarrow \mathcal{X}$ decodes the latent state back to the data domain. We use “encoder” or “embedding function” to refer to any operation $\mathbf{g} : \mathcal{X} \rightarrow \mathcal{Z}$ mapping data \mathbf{x} to latents \mathbf{z} .

A. Koopman Theory

We provide a brief overview of Koopman theory, deferring a detailed treatment to [2]. Let $\mathcal{G}(\mathcal{X})$ be a vector space of embedding functions on data $\mathbf{x} \in \mathcal{X}$. The *Koopman operator* $\mathcal{K} : \mathcal{G}(\mathcal{X}) \rightarrow \mathcal{G}(\mathcal{X})$ acts by composition,

$$(\mathcal{K}\mathbf{g})(\mathbf{x}) = \mathbf{g}(\mathbf{f}(\mathbf{x})), \quad \mathbf{g} \in \mathcal{G}(\mathcal{X}), \quad (3)$$

and is linear even if \mathbf{f} is nonlinear. Since \mathcal{K} is generally infinite-dimensional, in practice we seek finite approximations $\mathbf{K} \in \mathbb{R}^{m \times m}$ of \mathcal{K} as in (2), and we say \mathbf{g} is *\mathcal{K} -invariant* if the following closure condition is satisfied:

$$\mathbf{g}(\mathbf{f}(\mathbf{x})) = \mathbf{K}\mathbf{g}(\mathbf{x}), \quad \forall \mathbf{x}. \quad (4)$$

That is, the latents \mathbf{z} evolve linearly on the range of \mathbf{g} . When (4) fails, \mathbf{K} introduces *spurious modes* that accumulate multi-step prediction error due to the closure residual [3], [6], [2]. Thus, computing and analyzing invariant embeddings is central to Koopman analysis, which we revisit in Sec. IV.

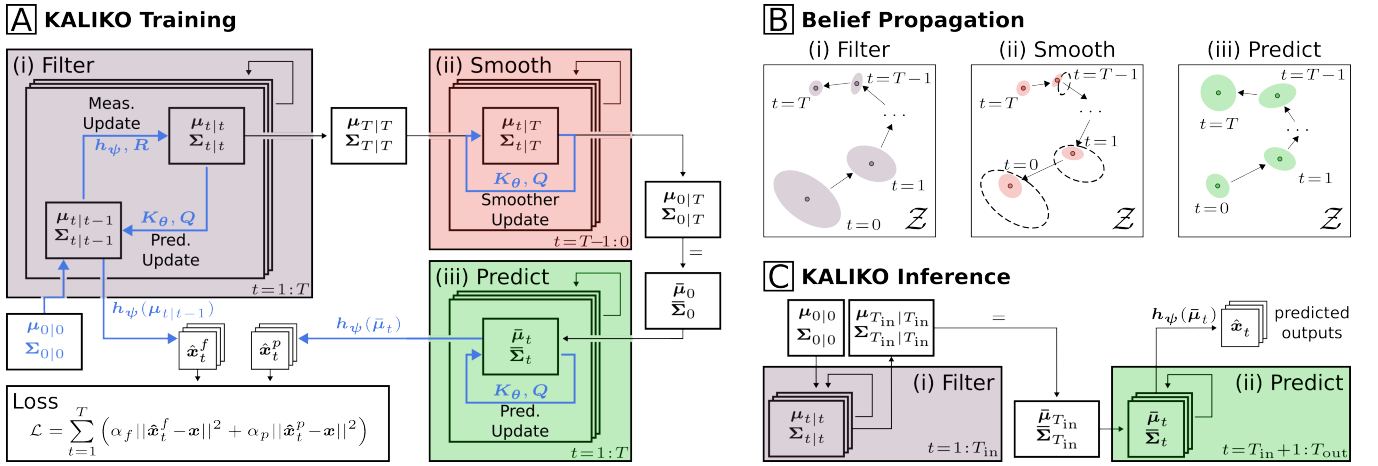


Fig. 2: (A) **KALIKO Training**. (i) Filter for T steps, outputting a sequence of filtered distributions $\{(\mu_{t|t-1}, \Sigma_{t|t-1})\}_{t=1}^T$ used for the loss. (ii) Run a backward smoother from the final belief $(\mu_{T|T}, \Sigma_{T|T})$, yielding the posterior initial belief $(\mu_{0|T}, \Sigma_{0|T})$. (iii) Roll this belief forward for T steps to get predicted beliefs $\{(\bar{\mu}_t, \bar{\Sigma}_t)\}_{t=1}^T$. Filtered and predicted means are decoded to observations and the model is trained end-to-end with a reconstruction loss. Learnable modules/parameters and their operations are highlighted in blue. (B) **Belief Propagation**. The filter’s latent belief becomes more certain with more steps. The smoother reuses future evidence to calibrate the initial belief, seeding a strong prediction rollout during training. (C) **KALIKO Inference**. KALIKO filters on T_{in} data points then rolls out its latent dynamics for T_{out} steps. The T_{out} predicted latents are decoded into the final predicted trajectory.

B. Kalman Filtering

We now review the mechanics of Kalman filters (KFs), deferring details to [37]. KFs recursively update a posterior belief $p(z_t | z_{t-1}, x_t)$ over the state given new measurements and its previous state. We assume Gaussian noise with covariance Q corrupts the dynamics such that

$$p(z_t | z_{t-1}) = \mathcal{N}(Kz_{t-1}, Q), \quad (5)$$

and similarly, that Gaussian noise with covariance R corrupts the decoder such that

$$p(x_t | z_t) = \mathcal{N}(h(z_t), R). \quad (6)$$

Starting from a prior belief over $z_0 \sim \mathcal{N}(\mu_{0|0}, \Sigma_{0|0})$, the KF iteratively updates its belief by interleaving two updates. First, the *prediction update* propagates the belief at time $t-1$ through the dynamics:

$$\mu_{t|t-1} = K\mu_{t-1|t-1}, \quad (7)$$

$$\Sigma_{t|t-1} = K\Sigma_{t-1|t-1}K^\top + Q. \quad (8)$$

Second, the *measurement update* adjusts the posterior by accounting for the likelihood of observing x_t :

$$G_t = \Sigma_{t|t-1}C_t^\top (C_t\Sigma_{t|t-1}C_t^\top + R)^{-1}, \quad (9)$$

$$\mu_{t|t} = \mu_{t|t-1} + G_t(x_t - h(\mu_{t|t-1})), \quad (10)$$

$$\Sigma_{t|t} = (I - G_tC_t)\Sigma_{t|t-1}, \quad (11)$$

where C_t is the decoder’s Jacobian, $C_t := \frac{\partial}{\partial z} h(\mu_{t|t-1})$.

The Kalman filter causally updates its belief distribution by observing only past data. However, after observing T data points, it is often useful to update *past* beliefs conditioned on future observations using *Kalman smoothing*. The Rauch-Tung-Striebel smoother caches values from the filtering pass

to recursively compute smoothed beliefs from time T to 0:

$$G_t^s = \Sigma_{t|t}K^\top (\Sigma_{t+1|t}^\top)^{-1}, \quad (12)$$

$$\mu_{t|T} = \mu_{t|t} + G_t^s(\mu_{t+1|T} - \mu_{t+1|t}), \quad (13)$$

$$\Sigma_{t|T} = \Sigma_{t|t} + G_t^s(\Sigma_{t+1|T} - \Sigma_{t+1|t})(G_t^s)^\top. \quad (14)$$

In summary, filtering and smoothing require the dynamics K , decoder h , prior belief $(\mu_{0|0}, \Sigma_{0|0})$, dynamics and measurement covariances Q and R , and a sequence of data points $x_{1:T}$. In return, they yield filtered and smoothed belief distributions over the corresponding latent states $z_{1:T}$, so we can view the Kalman filter/smoothen as an implicit encoder over trajectories, $g : x_{1:T} \mapsto z_{1:T}$.

III. THE KALMAN-IMPLICIT KOOPMAN OPERATOR

Our method has two goals: (i) producing high-quality predictions under latent linear dynamics K , and (ii) accurately decoding them with $x = h_\psi(z)$. Observe that inaccurate dynamics K degrade the prediction update (7)–(8), and an inaccurate decoder h_ψ degrades the measurement update (9)–(11). Thus, a performant KF requires both accurate K and h_ψ . This motivates our method, Kalman-implicit Koopman operator (KALIKO) learning, which trains a KF end-to-end with a single reconstruction objective to jointly achieve both goals. In contrast, autoencoding Koopman methods typically separate prediction and reconstruction [5], [3], [12], which may overemphasize one at the expense of the other.

A. Training

KALIKO is trained with the replay overshooting method [17], which consists of three stages (see Fig. 2A). Given a data trajectory $x_{1:T}$, the *filtering stage* interleaves the prediction and measurement updates in (7)–(8) and (9)–(11) respectively to recover the sequences of belief distributions $\{(\mu_{t|t-1}, \Sigma_{t|t-1})\}_{t=1}^T$ and $\{(\mu_{t|t}, \Sigma_{t|t})\}_{t=1}^T$. The *smoothing stage* then applies (12)–(14) to these beliefs in reverse




Fig. 3: **KALIKO Reconstructions.** KALIKO closely reconstructs the vector field of nonlinear systems without an encoder using globally linear latent dynamics. Shown are some ground truth trajectories in black (solid) and their reconstructions in orange (dashed). A denser heatmap of the reconstruction error is shown in blue (darker denotes more error).

to recover an initial belief at $t = 0$ conditioned on all measurements, $(\boldsymbol{\mu}_{0|T}, \boldsymbol{\Sigma}_{0|T})$. Lastly, the *prediction stage* rolls out T prediction updates (7)-(8) from the smoothed initial condition *without measurement updates* to recover a sequence of predicted beliefs $\{(\bar{\boldsymbol{\mu}}_t, \bar{\boldsymbol{\Sigma}}_t)\}_{t=1}^T$.

KALIKO learns \mathbf{K}_θ and \mathbf{h}_ψ along with the prior $(\boldsymbol{\mu}_{0|0}, \boldsymbol{\Sigma}_{0|0})$ and covariances (\mathbf{Q}, \mathbf{R}) end-to-end via the loss

$$\mathcal{L} = \sum_{t=1}^T \alpha_f \underbrace{\|\mathbf{h}_\psi(\boldsymbol{\mu}_{t|t-1}) - \mathbf{x}_t\|_2^2}_{\text{filtering}} + \alpha_p \underbrace{\|\mathbf{h}_\psi(\bar{\boldsymbol{\mu}}_t) - \mathbf{x}_t\|_2^2}_{\text{prediction}}, \quad (15)$$

where the first term corresponds to the filtered² reconstructions of the data and the second to the predicted reconstructions. The coefficients $\alpha_f, \alpha_p \in \mathbb{R}_{>0}$ control the relative weights for filtering and prediction respectively, which we set to 1 in all experiments.

The *filtering loss* is analogous to an autoencoding loss, as it fuses measurement consistency (“decoding”) via (9)-(11) with latent inference via (7)-(8) (“encoding”) [17]. The *prediction loss* ensures accurate open-loop long-horizon predictions, which mirrors inference-time execution. We remark that [17] maximizes the log-likelihood of the data while we found that a simple MSE loss was more computationally-efficient with no performance degradation.

B. Inference

At inference time, KALIKO uses T_{in} data points to predict the next T_{out} steps. First, the filter consumes $\mathbf{x}_{1:T_{\text{in}}}$ to obtain the posterior belief $(\boldsymbol{\mu}_{T_{\text{in}}|T_{\text{in}}}, \boldsymbol{\Sigma}_{T_{\text{in}}|T_{\text{in}}})$. T_{out} prediction steps are then executed from this filtered belief, yielding a sequence of predicted latent beliefs $\{(\bar{\boldsymbol{\mu}}_t, \bar{\boldsymbol{\Sigma}}_t)\}_{t=T_{\text{in}}+1}^{T_{\text{in}}+T_{\text{out}}}$. The predicted means are each decoded to obtain a predicted trajectory: $\hat{\mathbf{x}}_t = \mathbf{h}_\psi(\bar{\boldsymbol{\mu}}_t)$. See Fig. 2C for a visualization.

²We use the pre-measurement update mean $\boldsymbol{\mu}_{t|t-1}$, as in [17].

C. Additional Implementation Details

Delay Embedding. Inspired by the use of delay embeddings in Koopman methods [11], KALIKO parameterizes \mathbf{K}_θ as the dynamics on a delay-embedded latent state, yielding a sparse block-companion matrix form

$$\underbrace{\begin{bmatrix} \boldsymbol{\zeta}_{t+1} \\ \vdots \\ \boldsymbol{\zeta}_{t+n_z-1} \\ \boldsymbol{\zeta}_{t+n_z} \end{bmatrix}}_{\mathbf{z}_{t+1}} = \underbrace{\begin{bmatrix} 0 & I & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & I \\ \Gamma_1 & \Gamma_2 & \dots & \dots & \Gamma_{n_z} \end{bmatrix}}_{\mathbf{K}} \underbrace{\begin{bmatrix} \boldsymbol{\zeta}_t \\ \vdots \\ \boldsymbol{\zeta}_{t+n_z-2} \\ \boldsymbol{\zeta}_{t+n_z-1} \end{bmatrix}}_{\mathbf{z}_t},$$

where the only learnable parameters are the blocks $\Gamma_1, \dots, \Gamma_{n_z}$, and each $\boldsymbol{\zeta}_t$ is a sub-latent state of dimension d_z , which is stacked into the full delay-embedded latent state $\mathbf{z}_t = (\boldsymbol{\zeta}_t, \dots, \boldsymbol{\zeta}_{t+n_z-1}) \in \mathbb{R}^{n_z \cdot d_z}$.

Temporal Chunking. KALIKO temporally chunks the full-resolution data into groups of n_x states, each corresponding to a sub-latent vector as follows:

$$\boldsymbol{\zeta}_t \leftrightarrow (\boldsymbol{\xi}_\tau, \dots, \boldsymbol{\xi}_{\tau+n_x-1}), \quad (16)$$

where τ is the original time index satisfying $t = \lfloor \tau/n_x \rfloor$ and $\boldsymbol{\xi}_\tau$ is a raw data point. Thus, \mathbf{h} decodes n_z sub-latent vectors into a contiguous sequence of $n_z \cdot n_x$ raw data points:

$$\mathbf{h}(\underbrace{\boldsymbol{\zeta}_{t:(t+n_z-1)}}_{\mathbf{z}_t}) = \underbrace{\boldsymbol{\xi}_{\tau:(\tau+n_x \cdot n_z-1)}}_{\mathbf{x}_t}. \quad (17)$$

Temporal chunking greatly speeds up training, as it reduces the number of steps in each training stage by a factor of n_x , until the cost of memory and linear system solves associated with a large data dimension n outweighs the speedup.

Decoder Design. The decoder \mathbf{h}_ψ is a convolutional neural network that interleaves depthwise convolutions across the delay dimension (mixing $\boldsymbol{\zeta}$'s) with GELU-activated MLPs, which allows \mathbf{h} to mix sub-latent vectors across many time steps to decode long sequences of raw observations. We found two such blocks sufficient for all experiments.

IV. ANALYZING KALIKO EMBEDDINGS

A. Reconstruction Performance

As discussed in Sec. II-B, the Kalman filter/smoothen can be viewed as an encoder on sequences of data $\mathbf{x}_{1:T}$. We study KALIKO’s autoencoding capabilities on four systems for visualization purposes: the Van der Pol oscillator, pendulum, Duffing oscillator, and a two-limit cycle Hopf-Bautin system.

Van der Pol	$\begin{cases} \dot{x}_1 = x_1 - x_1^3/3 - x_2 \\ \dot{x}_2 = x_1 \end{cases}$ (18)
Pendulum	$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = \sin(x_1) \end{cases}$ (19)
Duffing	$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = x_1 - \delta x_2 - x_1^3 \end{cases}$ (20)
Hopf-Bautin (Polar)	$\begin{cases} \dot{r} = r e^{(1-2r^2)/2} (r^2 - r^4 - 3/16) \\ \dot{\theta} = 1 \end{cases}$ (21)

TABLE I: Systems used for embedding analysis. We learn discrete-time dynamics despite presenting these systems in continuous time for clarity.

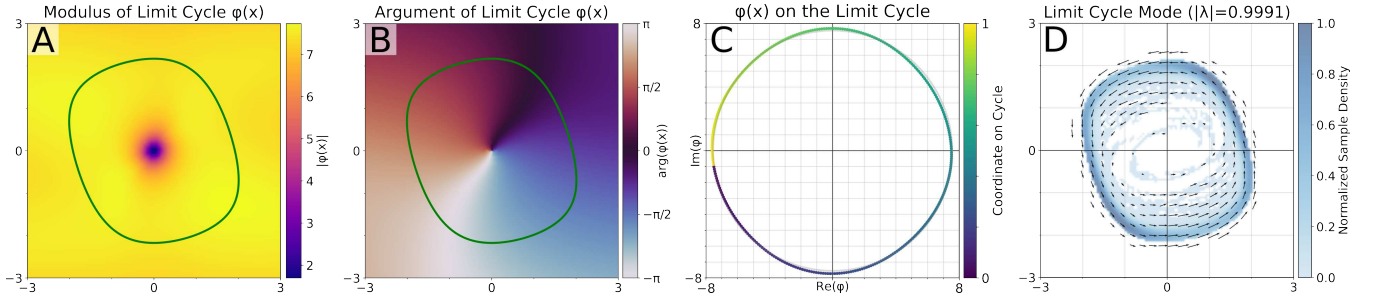


Fig. 4: **Limit Cycle Eigenfunction for the VDP System.** (A/B) Without parameterizing an encoder, KALIKO implicitly recovers an eigenfunction associated with the Van der Pol system’s limit cycle (overlaid in green) with $|\lambda| \approx 1$. (C) Evaluating $\varphi(x)$ along the limit cycle traces out a nearly perfect circle in the complex plane, showing that KALIKO entirely captures the limit cycle dynamics into this eigenfunction. (D) The Koopman mode associated with this eigenvalue corresponds to a vector field that drives trajectories onto the cycle.

We plot reconstructed trajectories on a heatmap of reconstruction errors (Fig. 3). KALIKO accurately reconstructs the original data ($\text{MAE} < 10^{-2}$), verifying that the Kalman filter and smooth function as an effective encoder g . Moreover, all models used identical hyperparameters, demonstrating KALIKO’s robustness to features of nonlinear systems like fixed points, limit cycles, and multiple attractors.

B. Analyzing KALIKO’s Eigenfunctions

Sec. II-A introduced a key property, \mathcal{K} -invariance, in (4). Observe that the *eigenfunctions* φ of \mathcal{K} satisfy (4), since

$$\lambda\varphi(\mathbf{x}_t) = (\mathcal{K}\varphi)(\mathbf{x}_t), \quad \lambda \in \mathbb{C}, \quad (22)$$

thus giving coordinates on which \mathcal{K} acts linearly, making them central objects of study in the literature [2]. A left eigenpair (λ, \mathbf{w}) of the approximation \mathbf{K} of \mathcal{K} (and g) implies an eigenfunction $\varphi(\mathbf{x}) := \mathbf{w}^\top \mathbf{g}(\mathbf{x})$, because

$$\begin{aligned} (\mathcal{K}\varphi)(\mathbf{x}_t) &= \varphi(\mathbf{f}(\mathbf{x}_t)) = \varphi(\mathbf{x}_{t+1}) = \mathbf{w}^\top \mathbf{g}(\mathbf{x}_{t+1}) \\ &= \mathbf{w}^\top \mathbf{K} \mathbf{g}(\mathbf{x}_t) = \lambda \mathbf{w}^\top \mathbf{g}(\mathbf{x}_t) = \lambda \varphi(\mathbf{x}_t). \end{aligned}$$

Thus, we examine the eigenfunctions of \mathbf{K}_θ via eigendecomposition, allowing us to interpret the dynamics captured by KALIKO’s implicit encoder, the Kalman filter and smoother. We consider two case studies.

Van der Pol. The Van der Pol (VDP) system (18) has a globally-stable limit cycle (Fig. 3, top left). We identify a complex eigenpair (λ, \mathbf{w}) of \mathbf{K}_θ with $|\lambda| \approx 1$, implying a steady-state oscillatory latent mode. We approximate the eigenfunction $\varphi(\mathbf{x})$ by evaluating $\mathbf{w}^\top \mathbf{g}(\mathbf{x})$ over latent beliefs in the plane, plotting VDP’s limit cycle (green) against the modulus and argument of $\varphi(\mathbf{x})$ in Fig. 4A/B.

We can now check whether the limit cycle is described by $\varphi(\mathbf{x})$. First, we parameterize the (unit-modulus) eigenvalue as $\lambda = e^{i\omega}$ and let \mathbf{z}_0 be an arbitrary fixed latent element of the corresponding eigenspace. Observe that the corresponding latent trajectory must have constant modulus:

$$|\mathbf{z}_t| = |\lambda^t \cdot \mathbf{z}_0| = |e^{it\omega} \cdot \mathbf{z}_0| = |\mathbf{z}_0|.$$

Moreover, the argument along the latent trajectory must evolve linearly (mod 2π):

$$\arg(\mathbf{z}_t) = \arg(\lambda^t \mathbf{z}_0) = t\omega + \arg(\mathbf{z}_0) \pmod{2\pi}.$$

Thus, if $\varphi(\mathbf{x})$ evaluated on the limit cycle traces out a single revolution of a circle in the complex plane, we can conclude that KALIKO *encodes the limit cycle dynamics entirely in this mode*, which is confirmed in Fig. 3C.

We also examine the time-invariant directions in \mathcal{X} given by the associated *right eigenvector* of \mathbf{K}_θ , or the *Koopman mode* [2]. To do so, we encode a set of trajectories, project them onto the limit cycle eigenspace, and decode them, resulting in Fig. 4D. The Koopman mode shows the attraction of the limit cycle in a neighborhood around it, verifying that KALIKO has learned VDP’s dominant dynamical behavior.

Duffing Oscillators. The undamped Duffing oscillator (UDO, $\delta = 0$ in (20)) has a saddle at the origin and two neutrally stable wells at $(\pm 1, 0)$. Because the oscillation frequency depends on energy (i.e., continuous spectrum [5]), no finite embedding space can recover a global phase eigenfunction; instead, invariants (e.g., energy) correspond to $\lambda \approx 1$. Consistent with this, KALIKO recovers an eigenfunction φ that is nearly constant along each closed orbit and varies across energy levels (Fig. 5A), capturing UDO’s energy invariance. Adding damping (DDO, $\delta = 1$) eliminates the continuous spectrum, leaving the saddle at the origin and making $(\pm 1, 0)$ asymptotically stable. Here, KALIKO identifies a contractive mode with $|\lambda| < 1$ whose eigenfunction φ tends to 0 at the three equilibria and whose phase clearly depicts the attractive spirals (Fig. 5B–D), in line with Koopman analyses in prior work [3, Fig. 5].

Takeaway. Despite using an *implicit* encoder, KALIKO reproduces the expected Koopman structure on these canonical systems (unit-modulus phase on the VDP limit cycle, energy-like invariants for UDO, and contractive modes near DDO equilibria) using *identical hyperparameters* with no system-specific tuning. These results illustrate that KALIKO’s implicit encoding does not preclude typical Koopman analysis; rather, it facilitates it by stably recovering expressive, interpretable representations without sensitive handcrafted features or model tuning.

V. EXPERIMENTS

A. Simulated Ocean Data

We now study the predictive capacity of KALIKO and its application to a closed-loop stabilization control problem. To do so, we simulate the motion of a barge in various ocean




Fig. 5: **Eigenfunctions for the Undamped and Damped Duffing Oscillator.** (A) KALIKO recovers an eigenfunction $\varphi(x)$ with $\lambda \approx +1$ capturing the invariance of systems with continuous spectra like the *undamped Duffing oscillator* (UDO). The value of $\varphi(x)$ is nearly constant along each of three distinct cycles (green) with little variation (shown are mean and stdev). Across cycles, the value clearly changes, corresponding to energy invariance. (B) Conversely, KALIKO also reconstructs the *damped Duffing oscillator* (DDO) with no adjustments. (C/D) KALIKO recovers an eigenfunction capturing the attractive “spiral” behavior about the wells at $(\pm 1, 0)$, reproducing the results from explicit methods in prior work [3, Fig. 5].

conditions, yielding thousands of trajectories in $SE(3)$. The data are generated using `capytaine` [38], which solves a linear-flow potential PDE in the frequency domain using the boundary element method. The irregular sea state is generated from the JONSWAP spectrum [39] by randomly sampling the significant wave height, peak period, and peak enhancement. For a set of discrete frequencies, we solve radiation and diffraction problems to assemble the hydrodynamic load on the moving body, which we then convert to a time series of motion by inverse Fourier transform. All parameter ranges were chosen such that the computed solutions were numerically stable, yielding realistic solutions. For more detail on hydrodynamics, we refer the reader to [40].

B. Baselines and Methodological Details

We compare KALIKO to both Koopman-specific and general time series prediction baselines.

Koopman Baselines. We have 3 Koopman baselines. The *Deep Koopman net* [5] is an autoencoding-based method that has the ability to parameterize continuous eigenvalue spectra. It is trained for long-horizon predictions with a multistep prediction consistency loss along with autoencoding losses, giving a canonical baseline for Koopman autoencoders. The *Koopman-Kalman Variational Autoencoder (K2VAE)* [41] claims state-of-the-art performance in probabilistic time series prediction. K2VAE predicts a “residual” global Koopman operator K_{glo} added to a locally-fitted operator derived from DMD, K_{loc} , yielding the model $K = K_{\text{loc}} + K_{\text{glo}}$. It uses an auxiliary Kalman filter to estimate and drive errors to 0. Since K2VAE assumes that K_{loc} is a good prior about which to learn a residual, we implement a simple baseline where we fit a *local DMD* model at each time step on the context T_{in} using delay embeddings as the Koopman embeddings. Several works have explored variations of this idea in an MPC-style context for modeling effects like fluid flows [42], [43]. We used the open-source implementations and settings of the baselines for fair comparison.

Time Series Baselines. We have 2 general time series baselines. *Recurrent neural networks (RNNs)* [29] are classic autoregressive models well-suited for dynamical prediction. RNNs can be viewed as “generalized” filters, making them a natural baseline: they ingest data, update an internal state,

and can generate prediction sequences. *PatchTST* [34] is a transformer that consumes data as temporal “patches”, inspired by vision transformers [44]. PatchTST is a popular method for general time series modeling, as it requires little domain-specific knowledge to yield reasonable predictions.

Data Preprocessing. For all methods, we map all poses in $SE(3)$ to $\mathfrak{se}(3)$, the tangent space at the identity, using the Log map. This makes the data Euclidean and enables geometrically consistent pose predictions. Second, we normalize the data with a running estimate of the mean and standard deviation along each axis, which is updated during training. At inference time, we fix these statistics to prevent leakage.

C. Prediction Performance

In *open-loop prediction*, each model is supplied an input context of $T_{\text{in}} = 128$ data points and predicts $T_{\text{out}} = 64$ steps. If applicable, methods are trained for 164k steps with a batch size of 32, and we report prediction errors in $\mathfrak{se}(3)$ averaged over all times, dimensions, and batches.

Method	MSE (\downarrow)	MAE (\downarrow)
Deep Koopman	0.0019	0.025
K2VAE	0.0015	0.026
Local DMD	0.0005	0.010
RNN	0.0003	0.010
PatchTST	0.0004	0.011
KALIKO (ours)	0.0002	0.008

TABLE II: Performance on open-loop wave prediction.

In Table II, KALIKO outperforms all methods using a global linear latent prediction model. In contrast, the RNN and PatchTST strongly leverage nonlinear models, and K2VAE and local DMD rely heavily on locally-fitted models rather than a global linear operator. Surprisingly, local DMD outperforms both Deep Koopman and K2VAE, which we attribute to the sensitivity of autoencoder-based Koopman methods to their learned embeddings, as noted in Sec. I-A, and K2VAE’s noisy predictions (see Fig. 6).

D. Ablations

Delay Length. KALIKO uses a delay length $n_z = 4$ for its latent dynamics. We perform two ablations: letting $n_z = 1$ (no delay) and increasing to $n_z = 6$. We find that the delay is crucial for performance, as the MAE nearly doubles




Fig. 6: **Control Results.** KALIKO outperforms all baselines on the control task with the least variance and near-oracle performance. Insets show one dimension of a hold-out trajectory. Gray denotes input context, black denotes ground truth prediction data. (i) K2VAE’s poor performance can be explained by its overly-noisy predictions (blue, dashed). (ii) KALIKO’s smooth prediction trajectories enable much more stable closed-loop control (pink, dashed).

without it. Conversely, using a higher delay does not improve performance, justifying the choice $n_z = 4$.

Decoder Architecture. We replaced KALIKO’s convolutional decoder with a simple MLP that performs no depth-wise mixing, instead mapping each sub-latent vector ζ_t to a distinct temporal chunk. We find this only slightly harms performance, showing KALIKO’s robustness to architectural design decisions.

Learnable Priors. We evaluate KALIKO’s performance when fixing the prior to $(0, I)$. This substantially degrades performance, which we hypothesize is because the scale and anisotropy of the prior are key for allowing KALIKO to learn the most expressive possible embedding space geometry.

Ablation	MSE (\downarrow)	MAE (\downarrow)
$n_z = 4 \rightarrow n_z = 1$	0.0007	0.015
$n_z = 4 \rightarrow n_z = 6$	0.0002	0.008
Conv \rightarrow MLP Decoder	0.0003	0.009
Learnable \rightarrow Fixed Latent Prior	0.0005	0.012
KALIKO (Original)	0.0002	0.008

TABLE III: Ablations: effect on KALIKO’s prediction performance.

E. Control Performance

We evaluate all methods on a simulated control task: stabilizing a cable-suspended payload with a barge-mounted crane (see Fig. 7). The task is challenging because the barge and crane are severely underactuated: the free-body motion is driven by the ocean waves, and the crane’s payload height is controllable only via a damped winch. We need not model any control dynamics, since the crane is controlled with sampling-based MPC, which selects controls by choosing high-performance candidates under forward simulations of the system. Thus, the predictive models in Sec. V-C suffice.

We evaluate all methods on 100 random wave simulations, feeding their predictions to an MPPI controller implemented using the open-source toolbox `judo` [45]. We calibrate performance against two other methods: an **oracle** supplied with ground-truth future barge motion, and a **nominal** stabilizing controller that ignores the wave disturbances. Since the controller is stochastic, the oracle allows us to measure the error induced by its sampling noise. We report mean absolute error in the payload target position (see Fig. 6).

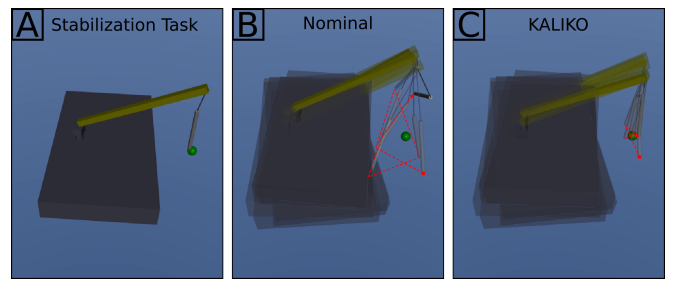


Fig. 7: (A) We aim to stabilize the tip of the payload to a goal location (green) by compensating for motion induced by realistic waves. A crane controls the payload by rotating its base and boom and/or reeling its winch. (B) The nominal controller poorly regulates the payload, as seen in this 3 second snapshot (red dashed lines). (C) KALIKO stabilizes the payload at near-oracle levels in the presence of the same waves.

KALIKO outperforms all methods with near-oracle performance and the tightest variance, showing that its superior open-loop prediction translates to high-performance closed-loop control. However, prediction performance is not a perfect indicator of control performance, as evidenced by local DMD outperforming the RNN. We attribute this gap to the role of state estimation, which is as crucial as prediction in control. Trained explicitly for filtering, KALIKO maintains an accurate latent belief in closed loop, whereas autoencoding-based Koopman methods and the RNN are trained only for prediction, and their internal states drift. Conversely, local DMD and PatchTST, which are not latent-variable methods, perform well because they operate directly on (patches of) raw data; with reasonably accurate short-horizon predictions and fast enough feedback, the controller can effectively stabilize the system.

VI. CONCLUSION

We introduced Kalman-Implicit Koopman Operator (KALIKO) learning, which leverages Kalman filtering to implicitly compute Koopman embeddings and linear latent dynamics without explicitly parameterizing an encoder. KALIKO learns highly interpretable representations, achieves superior open-loop prediction, and outperforms all baselines on a challenging closed-loop control task, with near-oracle

performance. In sum, KALIKO turns Bayesian filtering into a practical Koopman encoder, yielding performant embeddings that require no handcrafting or expert domain knowledge.

REFERENCES

- [1] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual review of fluid mechanics*, 52(1):477–508, 2020.
- [2] Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern koopman theory for dynamical systems. *SIAM Review*, 64(2):229–340, 2022.
- [3] Samuel E. Otto and Clarence W. Rowley. Linearly-recurrent autoencoder networks for learning dynamics, 2019.
- [4] Lu Shi, Masih Haseli, Giorgos Mamakoukas, Daniel Bruder, Ian Abraham, Todd Murphey, Jorge Cortes, and Konstantinos Karydis. Koopman operators in robot learning, 2025.
- [5] Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1), November 2018.
- [6] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [7] Yuhong Jin, Lei Hou, and Shun Zhong. Extended dynamic mode decomposition with invertible dictionary learning. *Neural Networks*, 173:106177, 2024.
- [8] Enoch Yeung, Soumya Kundu, and Nathan Hodas. Learning deep neural network representations for koopman operators of nonlinear dynamical systems. In *2019 American Control Conference (ACC)*, pages 4832–4839, 2019.
- [9] Clarence W. Rowley, Igor Mezic, Shervin Bagheri, Philipp Schlatter, and Dan S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641:115–127, 2009.
- [10] Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.
- [11] Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, Eurika Kaiser, and J. Nathan Kutz. Chaos as an intermittently forced linear system. *Nature Communications*, 8(1), May 2017.
- [12] Yuhuang Meng, Jianguo Huang, and Yue Qiu. Koopman operator learning using invertible neural networks. *Journal of Computational Physics*, 501:112795, March 2024.
- [13] Naoya Takeishi, Yoshinobu Kawahara, and Takehisa Yairi. Learning koopman invariant subspaces for dynamic mode decomposition, 2018.
- [14] Jack W. Miller, Charles O’Neill, Navid C. Constantinou, and Omri Azencot. Eigenvalue initialization and regularization for koopman autoencoders, 2022.
- [15] Indranil Nayak, Ananda Chakrabarti, Mrinal Kumar, Fernando L. Teixeira, and Debdipta Goswami. Temporally-consistent koopman autoencoders for forecasting dynamical systems. *Scientific Reports*, 15(1), July 2025.
- [16] Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. A disentangled recognition and nonlinear dynamics model for unsupervised learning, 2017.
- [17] Albert H. Li, Philipp Wu, and Monroe Kennedy. Replay overshooting: Learning stochastic latent dynamics with the extended kalman filter. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 852–858, 2021.
- [18] Tuomas Haarnoja, Anurag Ajay, Sergey Levine, and Pieter Abbeel. Backprop kf: Learning discriminative deterministic state estimators, 2017.
- [19] Michelle A. Lee, Brent Yi, Roberto Martín-Martín, Silvio Savarese, and Jeannette Bohg. Multimodal sensor fusion with differentiable filters, 2020.
- [20] Alina Kloss, Georg Martius, and Jeannette Bohg. How to train your differentiable filter. *Autonomous Robots*, 45(4):561–578, May 2021.
- [21] Adrien Corenflos, James Thornton, George Deligiannidis, and Arnaud Doucet. Differentiable particle filtering via entropy-regularized optimal transport, 2021.
- [22] Amit Surana and Andrzej Banaszuk. Linear observer synthesis for nonlinear systems using koopman operator framework. *IFAC-PapersOnLine*, 49(18):716–723, 2016. 10th IFAC Symposium on Nonlinear Control Systems NOLCOS 2016.
- [23] Lingyun Zeng, S.M.Hadi Sadati, and Christos Bergeles. Koopman operator-based extended kalman filter for cosserrat rod wrench estimation. In *2023 International Symposium on Medical Robotics (ISMR)*, pages 1–7, 2023.
- [24] Peng Huang, Ketong Zheng, and Gerhard Fettweis. Data-driven koopman operator-based error-state kalman filter for enhanced state estimation of quadrotors in agile flight. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10983–10989, 2024.
- [25] Ningxin Liu, Shuigen Liu, Xin T. Tong, and Lijian Jiang. Estimate of koopman modes and eigenvalues with kalman filter, 2024.
- [26] Zhexuan Zeng, Jun Zhou, Yasen Wang, and Zuwei Ping. Koopman spectral analysis from noisy measurements based on bayesian learning and kalman smoothing, 2025.
- [27] Guan Xiaoqing, You Wang, Xiaomeng Kang, Wei Yao, Jin Zhang, and Guang Li. An online system identification algorithm for spherical robot using the koopman theory. *IEEE Robotics and Automation Letters*, PP:1–8, 01 2025.
- [28] Zi Cong Guo, Vassili Korotkine, James R. Forbes, and Timothy D. Barfoot. Koopman linearization for data-driven batch state estimation of control-affine systems, 2021.
- [29] Alex Sherstinsky. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306, March 2020.
- [30] Nima Mohajerin and Steven L. Waslander. Multi-step prediction of dynamic systems with recurrent neural networks, 2018.
- [31] Abbas Tariverdi, Venkatasubramanian Kalpathy Venkiteswaran, Michiel Richter, Ole Elle, Jim Tøresen, Kim Mathiassen, Sarthak Misra, and Ørjan Martinsen. A recurrent neural-network-based real-time dynamic model for soft continuum manipulators. *Frontiers in Robotics and AI*, 8:631303, 03 2021.
- [32] Jianjing Zhang, Hongyi Liu, Qing Chang, Lihui Wang, and Robert X. Gao. Recurrent neural network for motion trajectory prediction in human-robot collaborative assembly. *CIRP Annals*, 69(1):9–12, 2020.
- [33] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels, 2019.
- [34] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers, 2023.
- [35] Tian Zhou, PeiSong Niu, Xue Wang, Liang Sun, and Rong Jin. One fits all: power general time series analysis by pretrained lm, 2023.
- [36] Sabera Talukder, Yisong Yue, and Georgia Gkioxari. Totem: Tokenized time series embeddings for general time series analysis, 2025.
- [37] Simo Särkkä. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2013.
- [38] Matthieu Ancellin and Frédéric Dias. Capytaine: a Python-based linear potential flow solver. *Journal of Open Source Software*, 4(36):1341, apr 2019.
- [39] Klaus F. Hasselmann, Tim P. Barnett, Evert Bouws, H. C. Carlson, David Edgar Cartwright, Kurt Enke, James Alfred Ewing, Hans Gienapp, Dieter E. Hasselmann, P. Kruseman, Arend Meerburg, Peter M. Müller, Dirk J. Olbers, Karl Richter, Wolfgang Sell, and Hans Walden. Measurements of wind-wave growth and swell decay during the joint north sea wave project (jonswap). 1973.
- [40] John Nicholas Newman. *Marine Hydrodynamics*. MIT Press, Cambridge, MA, 1977.
- [41] Xingjian Wu, Xiangfei Qiu, Hongfan Gao, Jilin Hu, Bin Yang, and Chenjuan Guo. k^2 vae: A koopman-kalman enhanced variational autoencoder for probabilistic time series forecasting, 2025.
- [42] Anqi Bao, Eduardo Gildin, Abhinav Narasingam, and Joseph S. Kwon. Data-driven model reduction for coupled flow and geomechanics based on dmd methods. *Fluids*, 4(3), 2019.
- [43] Qiuqi Li, Chang Liu, and Yifei Yang. Localized dynamic mode decomposition with temporally adaptive partitioning, 2025.
- [44] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [45] Albert H. Li, Brandon Hung, Aaron D. Ames, Jiuguang Wang, Simon Le Cleac’h, and Preston Culbertson. Judo: A user-friendly open-source package for sampling-based model predictive control, 2025.