Temporal stability in reduced order model prediction of sea states: A surrogate model case study

Freja H. Petersen[1,2] **Allan Engsig-Karup**[2] **Jesper Mariegaard**[1] **Rocco Palmitessa**[1]

Abstract

Physics-based hydrodynamic models provide accurate forecasts of sea surface elevations and currents but are often too computationally demanding for real-time or ensemble predictions. This study investigates linear reduced-order machine learning surrogates as efficient alternatives, focusing on methods to ensure temporal stability in long-term forecasts. Two models are compared: a partly differentiable approach based on proper orthogonal decomposition and linear regression (PODLR), and a fully differentiable linear Koopman autoencoder (LKAE). The effects of eigenvalue constraints and temporal unrolling during training are evaluated using a large-scale hydrodynamic dataset from the Øresund region. Results show that PODLR suffers from severe stability issues, which are effectively mitigated by eigenvalue constraints and temporal unrolling, achieving high accuracy relative to the full hydrodynamic model. The LKAE is inherently stable, with temporal unrolling reducing forecast errors by nearly 50%. These findings highlight that exploiting differentiable structures in machine learning surrogates enables robust and computationally efficient hydrodynamic forecasting, allowing year-long simulations to be performed within seconds.

1 Introduction

Predicting sea surface elevations and currents in coastal waters using physics-based numerical modeling is a key tool for short-term emergency responses, such as evacuations, and long-term planning, like the construction of protective infrastructure [1, 2, 3]. While physics-based models are generally accurate, they can be computationally demanding, limiting their use for ensemble forecasts for uncertainty quantification, climate scenario modeling as well as real-time predictions. This motivates the development of data-driven and machine learning (ML) methods for producing efficient surrogates of complex hydrodynamic models. Generally, we consider a nonlinear hydrodynamic system described by the discrete, partly auto-regressive timestep

$$\boldsymbol{x}_{k+1} = \mathcal{F}(\boldsymbol{x}_k, \boldsymbol{u}), \quad k = 0, 1, ..., T-1, \tag{1}$$

where $\boldsymbol{x}$ is the system state, and $\boldsymbol{u}$ is a potential external input such as wind, and k is the discrete time step. When such physical systems exhibit transient dynamics or depend on external inputs, they can be particularly difficult to model. The auto-regressive timestep can lead to numerical instabilities over time when approximated by a data-driven model. The reason is two-fold: 1) often trained on one-step training data, the ML-model parameters are not adapted to a multi-step setting, and 2) there is a distribution shift between training and testing data, because the attractors in the training data don't exactly coincide with those in the testing data. One approach to resolve this issue is to use temporal unrolling during training, where the model is exposed to more of the transient dynamics during training, in order to reduce the distribution shift [4]. Another approach for linear operators involves constraining the magnitude of the operator eigenvalues to be numerically less than or equal to one [5, 6]. These two approaches will be explored in this case study.

[1]DHI, Agern Allé 5, 2970 Hørsholm, DK, [2]DTU, Bygning 324, 2800 Kongens Lyngby, DK

1st Workshop on Differentiable Systems and Scientific Machine Learning @ EurIPS 2025

1.1 Case study and objectives

This paper investigates temporally stable reduced order surrogate models for predicting sea surface elevations and currents in the waters of Øresund between Denmark and Sweden. The data is produced by a physics-based hydrodynamic numerical model solving the shallow-water equations [7]. The model data appears on a flexible mesh in two spatial dimensions of appr. 3300 elements, with three sea state variables, resulting in a state dimension of appr. 10000 with 30 minute time resolution. The data is divided into a training and test set, consisting of a one-year forecast in the year 2021 and 2022, respectively. The hydrodynamic model uses boundary conditions and wind as input, both of which are assumed available in a modeling context and will therefore also be used as input to the surrogate. The simulation time for one year of simulation using the hydrodynamic model was 18 minutes on a NVIDIA A40 GPU.

Two types of ML-models will be investigated: A proper orthogonal decomposition (POD) with a linear regression for temporal predictions (PODLR), and a Koopman autoencoder using a linear neural autoencoder (LKAE). The Koopman Autoencoder (KAE) [8, 9] approximates a *Koopman invariant subspace* enabling nonlinear dynamics to be represented by a linear operator, while also reducing the dimensionality via the neural autoencoder. The LKAE parameters are all trainable in a fully-differentiable end-to-end setting, whereas the PODLR uses a closed-form solution for finding the latent space efficiently, while the regression model parameters are trainable. In this context, differentiability is defined in terms of the ability to apply automatic differentiation across the ML-model parameters for gradient-based optimization, meaning that the PODLR is only partly differentiable. The main objective is to investigate the following two questions for the two models with different differentiability properties:

1. What is the effect on long-term stability and accuracy of applying eigenvalue constraints for the linear operators during training, and what is the effect of temporal unrolling?
2. Do the effects differ for a fully differentiable model as opposed to a partly differentiable model?

Constraining eigenvalues of the linear operators is a hard constraint on the numerical auto-regression performed by the ML-surrogate, which ensures numerical stability. On the other hand, temporal unrolling is an implicit approach that exposes the models to longer series of the transients while training, which is expected to improve both numerical stability and accuracy.

The key contribution in this study is the exploration of techniques for differentiable reduced order ML-models that can improve stability in long-term forecasting with application in many areas. However, the practical application to emulate a large-scale hydrodynamic model underlines the potential in utilizing exactly these techniques for ensuring more reliable data-driven surrogate models for producing fast forecasts. To our knowledge, no similar study has been conducted previously for this or similar hydrodynamic applications.

2 Methodology

For a discrete time step k, the sea states can be represented by $\boldsymbol{x}_k \in \mathbb{R}^{N_x}$, and correspondingly the external inputs (boundary conditions and wind) can be represented by $\boldsymbol{u}_k \in \mathbb{R}^{N_u}$. We now seek to determine a surrogate model on the form:

$$\boldsymbol{x}_{k+1} \approx \boldsymbol{\Phi}_2 \mathcal{L}(\boldsymbol{\Phi}_1 \boldsymbol{x}_k, \boldsymbol{\Psi} \boldsymbol{u}_k, \boldsymbol{\Psi} \boldsymbol{u}_{k+1}), \quad k = 0, 1, ..., T-1 \tag{2}$$

where $\boldsymbol{\Phi}_1$ and $\boldsymbol{\Phi}_2$ are linear mappings of $\boldsymbol{x}$ to and from a latent space of reduced dimension, $\boldsymbol{\Psi}$ maps the external inputs to a latent space, and $\mathcal{L} : \mathbb{R}^{N_{\tilde{x}}+2N_{\tilde{u}}} \to \mathbb{R}^{N_{\tilde{x}}}$ is a linear regression model mapping the states from one time step to the next, using an external input with both a lagged and a lead input, $\boldsymbol{u}_k$ and $\boldsymbol{u}_{k+1}$. In this case, $\mathcal{L}$ is simply an $(N_{\tilde{x}} + 2N_{\tilde{u}}) \times N_{\tilde{x}}$ matrix, where the latent dimensions are $N_{\tilde{x}} = 15$ and $N_{\tilde{u}} = 50$.

Training of PODLR: The POD modes correspond to the right singular vectors of $\boldsymbol{X} = [\boldsymbol{x}_1, \boldsymbol{x}_2, ..., \boldsymbol{x}_M]$, where M is the number of time steps in the training data. They are found using a randomized singular value decomposition algorithm. Since the modes are orthonormal they are easily invertible, and $\boldsymbol{\Phi}_2 = \boldsymbol{\Phi}_1^{-1} = \boldsymbol{\Phi}_1^T$. The linear regression is simply a matrix of trainable parameters as described for the LKAE below with $\alpha_1 = 0$.

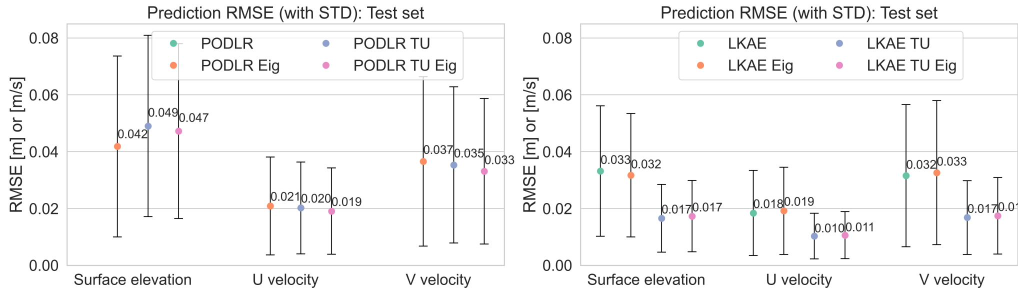

Figure 1: RMSE of a 1-year forecast in the test period with PODLR (left) and LKAE (right). The suffix 'Eig' means eigenvalue constraints, and 'TU' means temporal unrolling.

Training of LKAE: The autoencoders are linear neural networks without activation functions and with biases, and the linear regression is simply a linear layer in the middle. The time-steps are treated as the samples and the optimization is based on one-step predictions as a default. The optimization is performed in a PyTorch framework using the Adam optimizer, and it aims at minimizing the loss function:

$$L = \alpha_1 L_{recon} + \alpha_2 L_{pred}, \tag{3}$$

where L_{recon} is the mean squared error (MSE) of the reconstruction $\mathbf{\Phi}_2\mathbf{\Phi}_1\boldsymbol{x}_k$ compared to $\boldsymbol{x}_k$, L_{pred} is the MSE of the prediction (2) compared to $\boldsymbol{x}_{k+1}$, and $\alpha_i, i = 1, 2$ are the weights, in this implementation set to 1. A validation set is used for early stopping.

Eigenvalue constraints: The eigenvalues are soft-constrained by adding the penalty $\alpha_3 L_\lambda$ to the loss term (3), where L_λ is calculated using the ReLU function: $L_\lambda = \sum_{i=1}^{N_L} ReLU(|\lambda_i| - 1)$, where N_L is the dimension of the latent space (usually very small) and λ_i is the ith eigenvalues of the Koopman operator, $\mathcal{L}$, in (2). α_3 is set to 1. The penalty sums together the moduli of all eigenvalues larger than 1, pushing the eigenvalues below 1, ensuring temporal stability. In practice, all trained operators in this study have eigenvalues below 1, when $\alpha_3 = 1.0$

Temporal unrolling (TU): When applied, the optimization is based on several auto-regressive time-steps performed in the latent space, and the loss (3) is based on all time-steps. This way, the parameter optimization with backpropagation will be based on a gradient unrolling through the temporal steps. There is a trade-off between the accuracy achieved from exploring more unrolled time-steps and the stability of the computed gradients. The batch-size is set to the TU-length, such that one iterative forecast is performed per batch.

Model assessment: The predictive performances of the models are assessed by the root-mean-squared-error (RMSE) across time (1 year of forecast) and space given by $\sqrt{\frac{1}{N_x T}\sum_{i=1}^{N_x}\sum_{j=1}^{T}(X_{i,j} - \hat{X}_{i,j})^2}$, where $\hat{X}_{i,j}$ is the surrogate approximation of the sea state in the ith spatial element at time step j. The standard deviation on the absolute errors across space and time is also included in the results.

3 Results and discussion

The optimal temporal unrolling length is determined by testing lengths from 5 time steps up to 160. The error of the one-year forecast (evaluated in the training period) decreases for increasing unrolling length, but reaches a saturation point around 20 to 40 time steps. An unrolling length of 40 is chosen for the results presented here.

Figure 1 shows the RMSEs $\pm$ standard deviations of the one year test forecast for the two models for each sea state, with "U velocity" and "V velocity" being the u- and v- components of the currents in vector form, that is $u = s\cos(\theta)$ and $v = s\sin(\theta)$, when s is the current speed and θ the current direction. The errors are given in the original state units, meters for the surface elevations and meters per second for the currents. For reference, the mean and (maximum-minimum) values for the states are 0.14 (3.0) m, -0.0 (2.2) m/s and 0.019 (2.9) m/s, respectively. For one model, the patterns across states are more or less the same and will be described together. The leftmost dot describes the model without any additions (note that PODLR is unstable, and is therefore not present in the figure), the next dot is the model with eigenvalue constraints, the third is with temporal unrolling and the last is with both.

The basic LKAE is more stable than PODLR: Different initializations were tested for both models, and the LKAE was found in all cases to converge to a stable temporal operator with numerical values of the eigenvalues below 1. On the other hand, there are severe stability issues for PODLR. This is also the case, when the linear regression is found by least squares linear regression. The end-to-end differentiable training of LKAE, together with the separation of reconstruction and prediction errors in (3) ensures that the model adapts to a latent space where the dynamic is linear. This is different from the partially differentiable PODLR, where the latent space is fixed.

Using eigenvalue constraints improves stability in PODLR: Controlling the magnitude of the eigenvalues ensures a stable model, which also has decent accuracy, with a RMSE of 0.042 m, 0.021 m/s and 0.037 m/s for the three state variables. For LKAE, the eigenvalue constraint makes no difference, since the model is already stable.

Temporal unrolling improves accuracy in LKAE and stability in PODLR: The RMSE approximately halves across states, when temporal unrolling is used in LKAE, e.g. from 0.033 m to 0.017 m for surface elevation. For PODLR, temporal unrolling has the same effect as eigenvalue constraints: it improves stability in the linear regression, so that the PODLR surrogate is stable for a forecast of an entire year, despite that the temporal unrolling length was only 40 time steps (20 hours). This indicates, that the model is exposed to enough of the temporal dynamic during training, and this shows the advantage of a differentiable linear regression over a closed-form solution.

There is no remarkable effect of using both: There are only slight improvements for PODLR and no visible improvement for LKAE when both temporal unrolling and eigenvalue constraints are used.

While especially the accuracy improvements with temporal unrolling for LKAE are significant, a caveat of the method is that it is more computationally expensive to train, e.g. 200 vs. 740 seconds for 100 epochs (NVIDIA L4 GPU). The reason is the iterative predictions during training that cannot be parallelized in the same way as a batch-wide one-step prediction. Further, the gradient computation is more expensive. As opposed to this, the eigenvalue constraint does not increase the computational cost; since the latent dimension is small, the computation of eigenvalues is negligible in terms of costs. Therefore, combined with the results in Figure 1, it would be preferable to only use eigenvalue constraint for PODLR, whereas for LKAE the temporal unrolling is attractive due to the accuracy improvements, especially since temporal unrolling affects only training, not inference.

The main goal of the data-driven surrogate model is to accelerate forecasts. The PODLR and LKAE were chosen for their simplicity and lightweight design, yielding fast training and near-instant inference. Their small latent spaces enable rapid predictions, with most computational cost coming from the encoder–decoder transformations. Both models can complete a one-year forecast in seconds on a laptop CPU, including projection to and from latent space. Combined with their low errors, this makes them well-suited for ensemble and climate scenario forecasts.

4 Conclusion

This case study examined the effects of eigenvalue constraints and temporal unrolling during training of linear reduced-order machine learning (ML) surrogates of a hydrodynamic model. The objective was to improve the temporal stability of long-term forecasts of sea surface elevations and currents by comparing a partly differentiable ML-model based on proper orthogonal decomposition and linear regression (PODLR) with a fully differentiable linear Koopman autoencoder (LKAE). The PODLR exhibited severe stability issues due to its fixed autoencoder component, but both temporal unrolling and eigenvalue constraints substantially improved its stability and maintained high accuracy relative to the hydrodynamic model. In contrast, the LKAE was inherently more stable and did not benefit from eigenvalue constraints, although temporal unrolling reduced forecast errors by nearly half. These findings demonstrate that exploiting the differentiable structure of ML-based surrogates enables techniques than enhance long-term stability and reliability. Such approaches have significant potential in hydrodynamic modeling, where surrogates like PODLR and LKAE can produce year-long forecasts within seconds, opening the door to ensemble-based simulations and efficient long-term climate scenario analyses.

References

[1] José Vieira, Jakob Føns, and Giovanni Cecconi. "Statistical and hydrodynamic models for the operational forecasting of floods in the Venice Lagoon". In: *Coastal Engineering* 21.4 (1993), pp. 301–331. ISSN: 0378-3839. DOI: https://doi.org/10.1016/0378-3839(93)90012-W.

[2] Tom Bogaard et al. "The Early-Warning System for incoming storm surge and tide in the Republic of Mauritius". In: *EGU General Assembly*. Vol. 18. Geophysical Research Abstracts, 2016.

[3] DHI. *Protecting Venice from high water levels*. (visited: 2024-07-31). 2024. URL: https://www.dhigroup.com/projects/protecting-venice-from-high%5C%5C-water-levels.

[4] Bjoern List et al. *Differentiability in Unrolled Training of Neural Physics Simulators on Transient Dynamics*. en. 2024. DOI: 10.48550/arXiv.2402.12971.

[5] Shaowu Pan and Karthik Duraisamy. "Physics-Informed Probabilistic Learning of Linear Embeddings of Non-linear Dynamics With Guaranteed Stability". In: *SIAM Journal on Applied Dynamical Systems* 19.1 (2020), pp. 480–509. ISSN: 1536-0040. DOI: 10.1137/19M1267246. (Visited on 07/24/2025).

[6] Jinho Choi, Sivaram Krishnan, and Jihong Park. "Koopman Autoencoder Via Singular Value Decomposition for Data-Driven Long-Term Prediction". In: *2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP)*. ISSN: 2161-0371. 2024, pp. 1–6. DOI: 10.1109/MLSP58920.2024.10734828.

[7] DHI. *Hydrodynamic model of Øresund: MIKE21 model setup, outputs and observation data*. Version 1.0. Zenodo, Dec. 2024. DOI: 10.5281/zenodo.14160710. URL: https://doi.org/10.5281/zenodo.14160710.

[8] B. O. Koopman. "Hamiltonian Systems and Transformation in Hilbert Space". In: *Proceedings of the National Academy of Sciences* 17.5 (1931), pp. 315–318. DOI: 10.1073/pnas.17.5.315.

[9] Samuel E. Otto and Clarence W. Rowley. *Linearly-Recurrent Autoencoder Networks for Learning Dynamics*. 2019. URL: https://arxiv.org/abs/1712.01378.