

TOWARDS EFFICIENT AND STABLE OCEAN STATE FORECASTING: A CONTINUOUS-TIME KOOPMAN APPROACH

Rares Grozavescu

Department of Engineering
University of Cambridge
Cambridge, UK
{rg625}@cam.ac.uk

Pengyu Zhang

Department of Engineering
University of Cambridge
Cambridge, UK
{pz281}@cam.ac.uk

Mark Girolami

Department of Engineering
University of Cambridge
Cambridge, UK;
Alan Turing Institute
London, UK

Etienne Meunier

INRIA
Paris, France

ABSTRACT

We investigate the Continuous-Time Koopman Autoencoder (CT-KAE) as a lightweight surrogate model for long-horizon ocean state forecasting in a two-layer quasi-geostrophic (QG) system. By projecting nonlinear dynamics into a latent space governed by a linear ordinary differential equation, the model enforces structured and interpretable temporal evolution while enabling temporally resolution-invariant forecasting via a matrix exponential formulation. Across 2083-day rollouts, CT-KAE exhibits bounded error growth and stable large-scale statistics, in contrast to autoregressive Transformer baselines which exhibit gradual error amplification and energy drift over long rollouts. While fine-scale turbulent structures are partially dissipated, bulk energy spectra, enstrophy evolution, and autocorrelation structure remain consistent over long horizons. The model achieves orders-of-magnitude faster inference compared to the numerical solver, suggesting that continuous-time Koopman surrogates offer a promising backbone for efficient and stable physical-machine learning climate models.

1 INTRODUCTION

Ocean forecasting underpins understanding of Earth’s climate, yet high-fidelity Global Climate Models (GCMs) (Marshall et al., 1997; Madec et al., 2015) are computationally expensive, limiting large ensemble simulations. Recent data-driven surrogates based on large autoregressive architectures (Pathak et al., 2022; Lam et al., 2023; Bi et al., 2023; Rajabi-Kiasari et al., 2025) have shown strong short-term predictive skill but often suffer from instability over long rollout horizons.

We explore a complementary physics-structured alternative based on **Koopman Theory** (Koopman, 1931). Specifically, we apply the **Continuous-Time Koopman Autoencoder (CT-KAE)** (Grozavescu et al., 2026) to a two-layer quasi-geostrophic (QG) model. By evolving the latent representation via a linear Ordinary Differential Equation (ODE), the model enforces explicit operator structure while enabling temporally resolution-invariant forecasting through a matrix exponential formulation.

Rather than maximizing short-term pixel-wise accuracy, we prioritize long-horizon dynamical consistency: bounded error growth and preservation of bulk statistical invariants. The central question we address is whether explicit linear structure in latent space improves stability without sacrificing computational efficiency.

2 THE QUASI-GEOSTROPHIC MODEL

The Quasi-Geostrophic (QG) model have been used extensively in oceanography for modelling midlatitude oceanic circulation (Medjo, 2000; Majda & Wang, 2006).

2.1 GOVERNING EQUATIONS AND PHYSICS

We utilize a two-layer version of the QG model where the prognostic variable is potential vorticity, q_m , for layers $m \in \{1, 2\}$. The governing equations relate q_m to the streamfunction ψ_m :

$$q_1 = \nabla^2 \psi_1 + F_1(\psi_2 - \psi_1), \quad q_2 = \nabla^2 \psi_2 + F_2(\psi_1 - \psi_2), \quad (1)$$

where $F_1 = k_d^2/(1 + \delta)$, $F_2 = \delta F_1$, and $k_d^2 = (f_0^2/g')(H_1 + H_2)/(H_1 H_2)$. Here, H_m is the layer depth, and f_0 is the Coriolis frequency. The fluid velocity is related to ψ_m via $\mathbf{u}_m = (u_m, v_m) = (-\partial_y \psi_m, \partial_x \psi_m)$. The temporal evolution of the system is governed by:

$$\partial_t q_m + \mathbf{J}(\psi_m, q_m) + \beta \partial_x \psi_m + \partial_y Q_m \partial_x \psi_m + U_m \partial_x q_m = \mathcal{D}_m \quad (2)$$

where $\mathbf{J}(A, B) = \partial_x A \partial_y B - \partial_y A \partial_x B$ is the horizontal Jacobian, Q_m is background potential vorticity, and U_m is the background zonal mean flow velocity. The term \mathcal{D}_m represents small-scale dissipation, with the lower layer additionally incorporating Ekman bottom drag such that $\mathcal{D}_2 = -r_{ek} \nabla^2 \psi_2 + \text{ssd}$.

2.2 RELATED WORK ON QG AND MACHINE LEARNING

Machine learning approaches to QG dynamics broadly fall into two categories: subgrid-scale (SGS) parameterization and full state forecasting.

SGS Parameterization. Prior work has focused on learning closure terms representing unresolved processes. CNN-based approaches (Bolton & Zanna, 2019) and stability-aware refinements (Frezat et al., 2022; Ross et al., 2023) have demonstrated promising results for coarse-resolution modeling.

Full State Forecasting. Deep learning models including CNNs, GANs, Transformers, and reduced-order hybrids have been applied to turbulent flow prediction (Deo et al., 2023; Lee & You, 2019; Hemmasian & Farimani, 2023; Chattopadhyay et al., 2020; Besabe et al., 2025). More recently, diffusion-based approaches such as the Thermalizer (Pedersen et al., 2025) have targeted long-horizon stability.

In contrast, our approach enforces explicit linear structure and conduct temporal evolution in latent space, promoting stability while reducing inference cost to a matrix–vector multiplication.

3 METHODOLOGY

The CT-KAE projects high-dimensional QG states into a latent space governed by a linear continuous-time dynamical system before decoding back to the physical domain. The architecture consists of a dual-stream encoder, a latent ODE module, and a decoder (Figure 1).

Encoder and Decoder. Let $x_t \in \mathbb{R}^{H \times W \times C}$ denote the ocean state at time t , comprising channels for streamfunction ψ and potential vorticity q . CNN-based encoders are used, with a present encoder $\mathcal{E}_{\text{pres}}$ and a history encoder $\mathcal{E}_{\text{hist}}$ which take in x_t and x_{t-1} respectively, producing latent variable $z_t = \frac{1}{2}(\mathcal{E}_{\text{pres}}(x_t) + \mathcal{E}_{\text{hist}}(x_{t-1}))$. The decoder \mathcal{D} reconstructs the physical state, $\hat{x}_t = \mathcal{D}(z_t)$.

Continuous-Time Latent Dynamics. In contrast to discrete-time KAEs (Lusch et al., 2018), which propagate the latent state via a fixed-step mapping ($z_{t+1} = \mathbf{K}z_t$), the CT-KAE governs the latent evolution through a linear ODE:

$$\frac{dz}{dt} = \mathbf{K}z, \quad (3)$$

where a fourth-order Runge-Kutta (RK4) scheme is employed for numerical integration. The state at any future time $t + \tau$ can also be computed via the matrix exponential:

$$z(t + \tau) = \exp(\mathbf{K}\tau)z(t). \quad (4)$$




Figure 1: **CT-KAE architecture.** High-dimensional QG states x_t are encoded into a latent representation z_t using a dual-stream encoder that incorporates both present and historical information. The latent state evolves according to a linear continuous-time ODE $\dot{z} = \mathbf{K}z$, enabling resolution-invariant temporal querying via the matrix exponential. The decoder reconstructs physical fields from latent space.

This approach allows the model to be queried at arbitrary time. We empirically evaluate this resolution-invariant capability in Section 5.

Baseline: ViT (AR). We compare against a pure autoregressive Vision Transformer with 12 encoder layers (128-dim embeddings, 8 heads) operating on 8×8 patches. The model performs stepwise rollouts by recursively feeding predicted frames as input. Training uses AdamW (LR 2×10^{-4} , batch size 32) with the same data split as CT-KAE.

While effective at short-term prediction, the autoregressive nature of ViT leads to compounding errors over long horizons, as illustrated in Section 5. This provides a strong baseline for assessing the long-term stability benefits of the CT-KAE’s continuous-time latent dynamics.

Training. During training, models are optimized using short rollouts of 10 steps. Long-horizon behavior is therefore entirely emergent and not explicitly optimized during training. This setup allows us to evaluate whether stability arises from architectural structure rather than training heuristics.

4 DATASET

We generate data from a two-layer QG model following (Pedersen et al., 2025). The flow is driven by a background zonal mean flow ($U_1 = 0.025$ m/s, $U_2 = 0$ m/s) with physical parameters $\beta = 1.5 \times 10^{-11} \text{ m}^{-1} \text{ s}^{-1}$ and linear bottom drag $r_{ek} = 5.787 \times 10^{-7} \text{ s}^{-1}$. Ground truth simulations are generated with a 3600s integration step, then downsampled to 64×64 and subsampled every 5 steps for an effective resolution $\Delta t = 5h$. The system is integrated for 40,000 days to reach a stationary turbulent regime. We train on short 10-step trajectory segments and evaluate on unseen initial conditions using 2083-day ($\approx 10,000$ steps) rollouts without reinitialization. All fields are normalized by the training-set mean and standard deviation prior to model ingestion.

5 RESULTS

Because the QG system is chaotic, deterministic trajectory alignment becomes ill-posed at multi-year lead times. We therefore evaluate CT-KAE along two axes: (i) forecasting fidelity, from short-term accuracy to long-horizon invariants, and (ii) operational advantages, namely computational efficiency and temporal resolution invariance.

5.1 LONG-HORIZON EMULATION AND PHYSICAL INVARIANTS

We compare long-term reconstruction accuracy using RMSE (Table 1). CT-KAE achieves slightly lower error than the autoregressive ViT baseline (0.938 ± 0.031 vs. 0.993 ± 0.033), indicating that latent linear structure does not compromise predictive fidelity. At 2083-day lead times, anomaly

correlation (ACC) is near zero for both models (Table 1). This is expected: the decorrelation time scale of the QG system ($\mathcal{O}(10)$ days) is vastly exceeded, rendering deterministic trajectory alignment physically meaningless. Under such conditions, fidelity must be assessed through stability and preservation of statistical invariants rather than pointwise agreement.

To quantify exponential divergence, we define the long-horizon error growth rate:

$$\lambda = \frac{1}{T} \log \frac{\|x_T - \hat{x}_T\|}{\|x_0 - \hat{x}_0\|}. \quad (5)$$

A positive λ indicates exponential error amplification under recursive rollout, while a negative λ implies contraction toward a bounded invariant set. The autoregressive ViT exhibits positive growth ($\lambda = 0.0217$), consistent with gradual exponential error amplification under recursive rollout. In contrast, CT-KAE yields $\lambda = -0.0267$, demonstrating bounded long-horizon behavior. Notably, both models are trained only on 10-step segments; the stability of CT-KAE emerges from its continuous-time operator structure rather than explicit long-horizon optimization.

Beyond the Lyapunov time, fidelity is assessed via preservation of the invariant measure, quantified through kinetic energy spectra, enstrophy evolution, and temporal autocorrelation. Figure 2 shows that CT-KAE maintains large-scale vortical structures while mildly attenuating high-frequency modes. The KE spectrum confirms preservation of bulk energy at low wavenumbers without artificial amplification. Figure 3(b) shows bounded MSE growth over 10,000 steps, and enstrophy remains controlled.

We quantify energy stability using normalized drift,

$$\text{Drift} = \frac{E_T - E_0}{E_0}, \quad (6)$$

CT-KAE exhibits controlled negative drift (-49.4% KE, -34.6% enstrophy), consistent with mild dissipation, whereas ViT shows positive drift indicative of unphysical energy growth. Notably, the autoregressive ViT remains numerically stable over 10,000 steps and achieves competitive long-horizon RMSE, indicating that large Transformer models can learn partially self-regularizing dynamics. However, its positive error growth rate and energy drift suggest less controlled asymptotic behavior compared to CT-KAE.

5.2 OPERATIONAL ADVANTAGES OF CONTINUOUS-TIME DYNAMICS

The matrix exponential formulation, in Equation 4 provides two practical advantages over discrete autoregressive models.

Efficiency. Inference reduces to a single latent matrix–vector multiplication, yielding sub-millisecond per-step prediction on an NVIDIA RTX4090 GPU, approximately $300\times$ faster than the pseudo-spectral solver. Computational cost scales with latent dimension d , not spatial resolution, enabling efficient ensemble forecasting.

Temporal Resolution Invariance. Training is performed at $\Delta t = 5\text{h}$ and evaluated at 1h and 10h without retraining. RMSE remains stable across these unseen discretizations (Appendix A.5), confirming resolution invariance inherent to the continuous-time formulation.

Table 1: Long-horizon performance at 2083-day horizon (10,000 steps). Averaged across trajectories.

| Metric | CT-KAE | ViT (AR) |
|-------------------------|----------------------|---------------|
| RMSE ↓ | 0.938 ± 0.031 | 0.993 ± 0.033 |
| ACC ↑ | 0.033 | 0.004 |
| KE Drift (%) → 0 | -49.4 | 23.3 |
| Enstrophy Drift (%) → 0 | -34.6 | 37.0 |
| Error Growth Rate ↓ | -0.0267 | 0.0217 |




Figure 2: **Long-horizon qualitative comparison (2083 days).** **Top row:** Ground truth QG potential vorticity fields. **Middle row:** CT-KAE predictions rolled out for 9999 steps ($\Delta t = 5h$) without reinitialization. **Bottom row:** Kinetic Energy (KE) spectra averaged over the rollout for ground truth (blue) and CT-KAE (orange). The KE spectrum demonstrates preservation of bulk energy at low wavenumbers, while high-wavenumber modes are gradually attenuated. This controlled dissipation contributes to long-horizon stability, though it also indicates partial loss of fine-scale turbulent structures. Improving the capture of high-frequency dynamics while maintaining stability is a promising direction for future work.




Figure 3: **Stability over long horizons (20 trajectories).** (a) CT-KAE preserves large-scale persistence patterns. (b) MSE remains bounded over 10,000 rollout steps. (c) Enstrophy evolves without runaway amplification.

6 CONCLUSION

We show that structured continuous-time latent dynamics enable stable long-horizon emulation of chaotic ocean flows at negligible computational cost. By enforcing linear operator structure in latent space, CT-KAE achieves bounded error growth and preservation of bulk physical statistics over multi-year rollouts. Rather than optimizing short-term trajectory alignment, the model approximates the invariant statistical structure of the system. The resulting combination of stability, efficiency, and

temporal resolution invariance makes continuous-time Koopman surrogates a promising foundation for scalable physical–machine learning climate models.

REFERENCES

- Lander Besabe, Michele Girfoglio, Annalisa Quaini, and Gianluigi Rozza. Data-driven reduced order modeling of a two-layer quasi-geostrophic ocean model. *Results in Engineering*, 25:103691, 2025.
- Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.
- Thomas Bolton and Laure Zanna. Applications of deep learning to ocean data inference and subgrid parameterization. *Journal of Advances in Modeling Earth Systems*, 11(1):376–399, 2019.
- Ashesh Chattopadhyay, Mustafa Mustafa, Pedram Hassanzadeh, and Karthik Kashinath. Deep spatial transformers for autoregressive data-driven forecasting of geophysical turbulence. In *Proceedings of the 10th international conference on climate informatics*, pp. 106–112, 2020.
- Indu Kant Deo, Rui Gao, and Rajeev Jaiman. Combined space–time reduced-order model with three-dimensional deep convolution for extrapolating fluid dynamics. *Physics of Fluids*, 35, 2023.
- Hugo Frezat, Julien Le Sommer, Ronan Fablet, Guillaume Balarac, and Redouane Lguensat. A posteriori learning for quasi-geostrophic turbulence parametrization. *Journal of Advances in Modeling Earth Systems*, 14(11):e2022MS003124, 2022.
- Rares Grozavescu, Pengyu Zhang, Mark Girolami, and Etienne Meunier. Koopman autoencoders with continuous-time latent dynamics for fluid dynamics forecasting, 2026.
- AmirPouya Hemmasian and Amir Barati Farimani. Reduced-order modeling of fluid flows with transformers. *Physics of Fluids*, 35, 2023.
- B. O. Koopman. Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences*, 17:315–318, 1931.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- Sangseung Lee and Donghyun You. Data-driven prediction of unsteady flow over a circular cylinder using deep learning. *Journal of Fluid Mechanics*, 879:217–254, 2019.
- Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9, 2018.
- Gurvan Madec et al. Nemo ocean engine. 2015.
- Andrew Majda and Xiaoming Wang. *Nonlinear dynamics and statistical theories for basic geophysical flows*. Cambridge University Press, 2006.
- John Marshall, Alistair Adcroft, Chris Hill, Lev Perelman, and Curt Heisey. A finite-volume, incompressible navier stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research: Oceans*, 102(C3):5753–5766, 1997.
- T Tachim Medjo. Numerical simulations of a two-layer quasi-geostrophic equation of the ocean. *SIAM journal on numerical analysis*, 37(6):2005–2022, 2000.
- Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Four-castnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- Chris Pedersen, Laure Zanna, and Joan Bruna. Thermalizer: Stable autoregressive neural emulation of spatiotemporal chaos. *arXiv preprint arXiv:2503.18731*, 2025.

Saeed Rajabi-Kiasari, Artu Ellmann, and Nicole Delpeche-Ellmann. Sea level forecasting using deep recurrent neural networks with high-resolution hydrodynamic model. *Applied Ocean Research*, 157:104496, 2025.

Andrew Ross, Ziwei Li, Pavel Perezhogin, Carlos Fernandez-Granda, and Laure Zanna. Benchmarking of machine learning ocean subgrid parameterizations in an idealized model. *Journal of Advances in Modeling Earth Systems*, 15(1), 2023.

A APPENDIX

A.1 MODEL ARCHITECTURE

The Continuous-Time Koopman Autoencoder (CT-KAE) is designed to map high-dimensional turbulent states into a low-dimensional manifold where the dynamics can be approximated by a linear Ordinary Differential Equation (ODE). The architecture consists of three primary components: a history-aware encoder, a continuous-time latent evolution module, and a generative decoder.

Encoder. The encoder \mathcal{E} projects the physical state $x_t \in \mathbb{R}^{H \times W \times C}$ into a latent vector $z_t \in \mathbb{R}^d$. To capture the temporal derivatives necessary for defining the flow, we utilize a dual-input strategy. The encoder processes both the current state x_t and the history x_{t-1} . The backbone consists of Pre-Activation Residual Blocks equipped with Spectral Normalization to stabilize the Lipschitz constant of the mapping. To allow the model to focus on active vortex shedding regions while ignoring background noise, we incorporate a Convolutional Block Attention Module (CBAM) at the bottleneck. The final latent state is obtained via a linear projection:

$$z_t = \mathcal{E}(x_t, x_{t-1}) \tag{7}$$

Decoder. The decoder \mathcal{D} mirrors the encoder structure, utilizing pixel-shuffle upsampling blocks to reconstruct the physical field $\hat{x}_t = \mathcal{D}(z_t)$. To ensure spatial consistency and mitigate spectral bias, we employ *Coordinate Injection*, concatenating a fixed grid of (x, y) coordinates to the feature maps before the final reconstruction layer.

A.2 CONTINUOUS-TIME LATENT DYNAMICS

Unlike discrete-time recurrent models, the CT-KAE governs the latent space via a learned linear ODE. This allows for resolution-invariant forecasting and theoretical stability analysis.

Operator Decomposition. The evolution of the latent state z is defined as $\frac{dz}{dt} = \mathbf{K}z$. To ensure stable learning of the Koopman operator $\mathbf{K} \in \mathbb{R}^{d \times d}$, we decompose it into skew-symmetric and symmetric components:

$$\mathbf{K} = \mathbf{K}_{\text{skew}} + \mathbf{K}_{\text{sym}} \tag{8}$$

where $\mathbf{K}_{\text{skew}} = \frac{1}{2}(\mathbf{W} - \mathbf{W}^T)$ and $\mathbf{K}_{\text{sym}} = \frac{1}{2}(\mathbf{D} + \mathbf{D}^T)$.

- \mathbf{K}_{skew} corresponds to purely imaginary eigenvalues, capturing the oscillatory/advection dynamics (vortex rotation) while preserving energy.
- \mathbf{K}_{sym} corresponds to real eigenvalues, capturing the growth and decay of modes (dissipation).

This decomposition prevents the "exploding gradient" problem common in RNNs by allowing explicit control over the system's energy.

Integration. During training, the trajectory is integrated using a 4th-order Runge-Kutta (RK4) scheme for numerical stability over short horizons. For inference and long-term forecasting, we utilize the matrix exponential solution:

$$z(t + \tau) = \exp(\mathbf{K}\tau)z(t) \tag{9}$$

This formulation allows the model to predict the ocean state at arbitrary continuous time-steps τ , independent of the training data's temporal resolution.

A.3 PHYSICS-INFORMED TRAINING OBJECTIVE

To enforce physical consistency and prevent the model from learning trivial identity mappings, we minimize a composite loss function \mathcal{L}_{total} :

$$\mathcal{L}_{total} = \mathcal{L}_{recon} + \alpha\mathcal{L}_{pred} + \beta\mathcal{L}_{latent} + \gamma\mathcal{L}_{phys} \quad (10)$$

1. Reconstruction & Prediction ($\mathcal{L}_{recon}, \mathcal{L}_{pred}$): We calculate the Mean Squared Error (MSE) for the immediate reconstruction and the future states rolled out via the latent dynamics. We apply a gradient-weighted mask to penalize errors in high-gradient regions (e.g., shock fronts or eddy edges) more heavily than the quiescent background.

2. Latent Regularization (\mathcal{L}_{latent}): To prevent mode collapse—where the model utilizes only a fraction of the latent dimensions—we enforce a covariance whitening loss. This penalizes off-diagonal correlations in the latent batch covariance matrix, encouraging an orthogonal and information-rich latent space. Additionally, we apply an eigenvalue repulsion term to disperse the learned frequencies of \mathbf{K} .

3. Physics Consistency (\mathcal{L}_{phys}): We incorporate domain-specific constraints:

- *Sobolev Loss*: Penalizes the error in the spatial gradients ($\nabla \hat{x}$ vs ∇x) to ensure the vorticity fields remain sharp.
- *Spectral Loss*: Minimizes the error in the Fourier domain (FFT) to ensure the energy cascade of the turbulent flow is preserved across wavenumbers.

A.4 HYPERPARAMETERS

| Parameter | Value |
|--------------------------|--|
| Latent Dimension (d) | 128 |
| Batch Size | 32 |
| Optimizer | AdamW |
| Learning Rate | 2×10^{-4} |
| Weight Decay | 1×10^{-5} |
| ODE Integration | Runge-Kutta 4 (Train), Matrix Exp (Test) |
| Activation Function | SiLU (Swish) |
| Spectral Normalization | Applied to all Conv/Linear layers |

Table 2: Hyperparameters used for the CT-KAE experiments.

A.5 TEMPORAL INVARIANCE

To validate resolution invariance, we evaluate the trained model at unseen temporal discretizations (Figure 4).

A.6 LEARNED OPERATOR SPECTRUM

To analyze the stability properties of the learned Koopman operator \mathbf{K} , we compute its eigenvalue spectrum after training. The majority of eigenvalues lie in the left half of the complex plane, indicating dissipative behavior, while a small subset lies near the imaginary axis, corresponding to oscillatory modes that preserve large-scale rotational structures.

This spectral structure explains the bounded long-horizon error growth observed in Section 5. Unlike autoregressive nonlinear predictors, where instability emerges implicitly, the continuous-time formulation provides explicit control over mode growth and decay. To explicitly analyze stability, we inspect the learned operator spectrum (Figure 5).




Figure 4: **Temporal resolution invariance.** CT-KAE trained at $\Delta t = 5h$ is evaluated at different integration resolutions ($1h$ and $10h$) using the continuous-time formulation. Each row corresponds to a different query time-step. The qualitative similarity across resolutions demonstrates that latent dynamics generalize across temporal discretizations without retraining, validating the matrix exponential formulation.




Figure 5: **Eigenvalue spectrum of the learned Koopman operator \mathbf{K} .** Eigenvalues are plotted in the complex plane. The majority lie in the left half-plane (negative real part), indicating dissipative modes, while a subset lies close to the imaginary axis, corresponding to oscillatory dynamics. The learned operator spectrum exhibits no eigenvalues with positive real part, which is consistent with the empirically observed bounded latent trajectories. While this does not constitute a formal stability guarantee of the decoded dynamics, it provides structural evidence supporting long-horizon stability as from Section 5.