Forcing as Closure: A Koopman-Invariant View of HAVOK

Yuer Tang[*]
Department of Mathematics
University of California, Los Angeles
Los Angeles, CA, USA
yuertang17@g.ucla.edu

Justin M. Baker
Department of Mathematics
University of California, Los Angeles
Los Angeles, CA, USA
justin@math.ucla.edu

Andrea L. Bertozzi
Department of Mathematics
University of California, Los Angeles
Los Angeles, CA, USA
bertozzi@math.ucla.edu

Abstract

Forecasting chaotic dynamical systems from partial or noisy observations remains a fundamental challenge in computational modeling. The HAVOK (Hidden Analysis of Variance with Ordered Koopman) framework offers an elegant solution by reconstructing nonlinear flows as a linear dynamical system plus an intermittent forcing term, derived purely from delay-embedded scalar measurements. Despite its empirical success, the origin and interpretation of this forcing term remain largely heuristic. In this paper, we study the mathematical role of forcing in HAVOK through the lens of Koopman operator theory. We show that the forcing signal emerges precisely when the truncated delay-embedded coordinates fail to span a Koopman-invariant subspace. This deviation localizes nonlinear transitions and can be viewed as a data-driven closure mechanism for otherwise linear predictive models. We validate these claims through experiments on the Lorenz and Rössler systems, highlighting both the power and limitations of HAVOK in capturing essential structure in chaotic time series.

1 Introduction

Forecasting the evolution of chaotic or partially observed systems is a longstanding problem in physics, engineering, and machine learning. While traditional model-based approaches rely on full access to governing equations, recent advances in data-driven modeling aim to reconstruct dynamics directly from time series. A key idea underlying this effort is the Koopman operator, which lifts nonlinear evolution into a linear flow over observables. However, this operator lives in an infinite-dimensional space, and practical computations demand tractable approximations.

The HAVOK algorithm Brunton et al. [2017] offers one such approximation. From a single scalar observable, it builds a Hankel matrix via time-delay embedding, performs singular value decomposition (SVD), and constructs a reduced-order model in the dominant coordinates. The resulting system evolves linearly, with an additional forcing term that activates only during transient, nonlinear events such as regime shifts. Despite the intuitive appeal and empirical success of this method, the theoretical interpretation of the forcing term remains underexplored.

[*]Corresponding author

Preprint. Under review.

Motivation. What makes HAVOK particularly compelling is its ability to forecast with high accuracy from minimal input. In particular, the model's sparse forcing signal appears to identify meaningful transitions in the underlying dynamics, offering a compact lens on regime changes. This raises the question: what is the origin of this forcing? Under what conditions does it vanish, and what does it mean when it doesn't?

Contributions. In this paper, we:

- Clarify that the HAVOK forcing term arises when the delay-embedded coordinates fail to span a Koopman-invariant subspace.
- Interpret this forcing signal as a practical, data-driven correction term that captures deviations from ideal linearity in truncated systems.
- Use toy experiments on the Lorenz and Rössler systems to demonstrate HAVOK's ability to extract low-rank dynamics and localize transitions.

Our analysis provides insight into the structure of the forcing term and demonstrates HAVOK's utility in practical forecasting tasks and model reduction, even without deep theoretical derivation from formal Mori–Zwanzig machinery.

2 Background

Deterministic Dynamical System Consider an *autonomous* ordinary differential equation (ODE)

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)), \qquad \mathbf{x}(t) \in \mathcal{M} \subset \mathbb{R}^n, \tag{1}$$

where the function $\mathbf{f} : \mathcal{M} \to \mathbb{R}^n$ is Lipschitz–continuous, ensuring the existence and uniqueness of solutions. M is the state space that's a manifold. Equation 1 therefore prescribes a deterministic rule of motion at every point $\mathbf{x}$. Because no random forcing term appears, fixing an initial condition $\mathbf{x}_0$

completely determines the future trajectory—once the first frame is set, the entire "movie" is already scripted. With the chosen x_0, we can have a flow map.

Flow Map Let $\Phi : \mathcal{M} \times \mathbb{R}_{\geq 0} \to \mathcal{M}$ denote the *flow map* associated with the trajectory $\mathbf{x}(t)$ governed by Equation 1 and initial condition $\mathbf{x}_0 \in \mathcal{M}$. That is,

$$\Phi^t(\mathbf{x}_0) := \mathbf{x}(t; \mathbf{x}_0), \quad \text{with} \quad \mathbf{x}(0; \mathbf{x}_0) = \mathbf{x}_0.$$

This map describes the evolution of the system from its initial state $\mathbf{x}_0$ forward by time t. To fully reconstruct trajectories, the flow must satisfy the *semigroup property*:

$$\Phi^{t+s} = \Phi^t \circ \Phi^s \quad \text{for all } t, s \geq 0,$$

which ensures temporal consistency. In this sense, Φ^t acts as a *time–t transport map* that advances every initial state $\mathbf{x}_0$ to its location after t units of time.

Koopman flow. Classical analysis follows the trajectory $\mathbf{x}(t)$ in state space. Koopman theory instead tracks observables: scalar functions that report a quantity of interest. Equip a Banach (often Hilbert) space

$$\mathcal{G} \subseteq \{g : \mathcal{M} \to \mathbb{C}\}, \qquad \text{with norm } \|\cdot\|_{\mathcal{G}}. \tag{2}$$

The Koopman operator family propagates observables in time:

$$\big[\mathcal{K}^t g\big](\mathbf{x}) = g\big(\Phi^t(\mathbf{x}_0)\big), \qquad g \in \mathcal{G},\ t \geq 0. \tag{3}$$

RHS $\mathcal{K}^t$ operates on g. So you can interpret it as: Koopman operator acts on the observables of x equals the observables of the flow of x. So loosely $\mathcal{K}^t$ acts as the flow on the Koopman operator level. Because composition with Φ^t is linear in g, each $\mathcal{K}^t : \mathcal{G} \to \mathcal{G}$ is linear even when $\mathbf{f}$ is nonlinear. Koopman operator transforms nonlinear dynamics in state space into linear dynamics in function space. The family $\{\mathcal{K}^t\}_{t\geq 0}$ also obey the semigroup law $\mathcal{K}^{t+s} = \mathcal{K}^t\mathcal{K}^s$. Thus we showed that Koopman Operator is a flow on the measurement space.

Single-Observation Time Series Let $\Phi^t : M \to M$ be the flow and let $g : M \to \mathbb{R}$ be a scalar sensor. The assumption of h is generic and couples all degrees of freedom.

$$s_n \;=\; g\big(\Phi^{n\Delta t}(x_0)\big) \tag{4}$$

where Δt is the sampling interval. Each s_n is one observation. Collectively $\{s_n\}_{n=0}^{N}$ is a single-observation time series. We do not observe the full state $\Phi^{n\Delta t}(x_0)$ but only this scalar output.

Delay reconstruction. Following Takens, choose an embedding dimension m and an integer lag τ(Takens [1981]). Form the m-dimensional delay vector

$$\mathbf{s}_n = e(s_n) = \big(s_n,\, s_{n-\tau},\, \ldots,\, s_{n-(m-1)\tau}\big) \;\in\; \tilde{\mathcal{A}} \subset \mathbb{R}^m,$$

Induced Flow in Delay ReconstructionKantz and Schreiber [2003]

$$\begin{array}{ccc}
x_n \in \mathcal{A} \subset \Gamma & \xrightarrow{\;F\;} & x_{n+1} \in \mathcal{A} \subset \Gamma \\
\downarrow g & & \downarrow g \\
s_n \in \mathbb{R} & & s_{n+1} \in \mathbb{R} \\
\downarrow e & & \downarrow e \\
\tilde{s}_n \in \tilde{\mathcal{A}} \subset \mathbb{R}^m & \xrightarrow{\;G\;} & \tilde{s}_{n+1} \in \tilde{\mathcal{A}} \subset \mathbb{R}^m
\end{array}$$

The notation $\Phi^t(\mathbf{x}_0) := \mathbf{x}(t;\mathbf{x}_0)$ represents the flow of the original dynamical system on the state space X, starting from the initial condition $\mathbf{x}_0$. It tells us where the system state will be at time t if it starts at $\mathbf{x}_0$.

Now, when we construct a delay embedding via the map

$$e : x \mapsto \big(g(x),\, g(\Phi^{-\tau}(x)),\, \ldots,\, g(\Phi^{-(m-1)\tau}(x))\big),$$

the flow on state space is transformed into a flow on the delay-embedded space $\tilde{\mathcal{A}} \subset \mathbb{R}^m$.

This gives rise to the induced flow G on delay space, defined via $G\big(e(x)\big) = e\big(\Phi^1(x)\big)$, or in terms of the scalar measurements, $G\big(e(s_n)\big) = e(s_{n+1})$.

Interpretation: The map G plays the same role in delay space as the true flow Φ^t does in state space—it advances trajectories one step forward in time. Thus, G is structurally analogous to Φ^t, but it acts on delay vectors rather than state vectors.

Koopman Invariance A linear subspace $\mathcal{V} \subseteq \mathcal{G}$ is *Koopman-invariant* if

$$\mathcal{K}^t\mathcal{V} \subseteq \mathcal{V}, \quad \forall\, t \geq 0,$$

equivalently $\mathcal{L}\mathcal{V} \subseteq \mathcal{V}$ for the generator $\mathcal{L}$. Invariance means that once an observable starts in $\mathcal{V}$, its entire time evolution remains confined to that same span, so no new coordinates are generated.

Time-delay may not have koopman invariance We express the delay-coordinate evolution as $\mathbf{s}_{n+1} = K\mathbf{s}_n$, to indicate that the dynamics are linear in delay space. This means that the next delay vector $\mathbf{s}_{n+1}$ is a linear combination of the components of $\mathbf{s}_n$.

In particular, the delay vector is defined as

$$\mathbf{s}_n = \big(s_n,\; s_{n-\tau},\; s_{n-2\tau},\; \ldots,\; s_{n-(m-1)\tau}\big),$$

and serves as a basis for a finite-dimensional subspace of observables. To say that the dynamics are linear means that

$$s_{n+1} \in \operatorname{span}\{s_n,\; s_{n-\tau},\; \ldots,\; s_{n-(m-1)\tau}\}.$$

That is, the future scalar observation s_{n+1} lies within the linear span of past measurements. In this case, the delay-coordinate space is closed under the dynamics. This is when the time-delay embedding have koopman invariance. However, if the dynamics is not linear, then it doesn't have koopman invariant, then we need to add additional embedding or forcing term, where HAVOK uses this method and we will introduce later. Brunton et al. [2016]

3 From Delay Embedding to Forcing: HAVOK

Studies like Brunton et al. (Nature 2017) show that delay embeddings can yield approximately linear models in delay space up to a forcing term Takeishi et al. [2017], Hirsh et al. [2021]

Hankle Matrix and SVD. Given a scalar measurement $y(t) = g(\mathbf{x}(t))$, $\mathbf{x}$ here is referring to the state. form the Hankel matrix. Hankel matrix is equivalent to applying a (multi-row) convolution or sliding-window operator to a time-lag vector:

$$\boldsymbol{H} = \begin{bmatrix} x(t_1) & x(t_2) & \cdots & x(t_p) \\ x(t_2) & x(t_3) & \cdots & x(t_{p+1}) \\ \vdots & \vdots & \ddots & \vdots \\ x(t_q) & x(t_{q+1}) & \cdots & x(t_m) \end{bmatrix} = \boldsymbol{U\Sigma V}^*.$$

The columns of U and V from the SVD are arranged hierarchically by their ability to model the columns and rows of $\mathbf{H}$. You can relate to Principle trajectories. After you do the SVD, then you can do the trucation to reduce the dimension and capture the important information. Specifically you can truncate to the first r singular values. Denote the corresponding right singular vectors by $\boldsymbol{v} = [\, v_1,\ \ldots,\ v_{r-1},\ v_r \,]$, and define the resolved coordinate $\hat{\boldsymbol{v}}(t) = [\, v_1(t), \ldots, v_{r-1}(t) \,]^\top$ and the unresolved coordinate $z(t) = v_r(t)$.

Approximate Koopman Invariance from Hankel SVD. The low-rank approximation of the Hankel matrix,

$$\mathbf{H} \approx \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^*,$$

provides a data-driven measurement system using delay coordinates. According to Takens' embedding theorem, delay embeddings reconstruct the state space dynamics from a time series of a single observable $x(t)$, allowing one to construct a phase-space representation via time-lagged vectors.

Each row of the Hankel matrix $\mathbf{H}$ is a shifted version of the previous one, which implicitly captures the system's temporal evolution. The Koopman operator $\mathcal{K}$ advances observables forward in time via

$$\mathcal{K}g(x_k) = g(x_{k+1}),$$

acting linearly on functions of the state. If the system evolves on an attractor, the dynamics map the attractor onto itself, making it invariant to the flow. Hence, time evolution through the Koopman operator corresponds closely to shifting through the rows of $\mathbf{H}$.

Since the columns of $\mathbf{U}_r$ define dominant intrinsic coordinates from data, and since time evolution of these coordinates corresponds to the Koopman operator acting on observables, the subspace spanned by these SVD modes is nearly preserved under Koopman evolution. That is, for dominant modes v_i,

$$\mathcal{K}v_i \approx \text{linear combination of } \{v_j\},$$

so the system appears to evolve linearly in these coordinates. This is why the low-rank delay coordinates form an approximately Koopman-invariant measurement system.

So you can rewrite $\mathbf{H}$ with koopman operator K:

$$\mathbf{H} = \begin{bmatrix} x(t_1) & \mathcal{K}x(t_1) & \cdots & \mathcal{K}^{p-1}x(t_1) \\ \mathcal{K}x(t_1) & \mathcal{K}^2x(t_1) & \cdots & \mathcal{K}^p x(t_1) \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{K}^{q-1}x(t_1) & \mathcal{K}^q x(t_1) & \cdots & \mathcal{K}^{m-1}x(t_1) \end{bmatrix}. \tag{5}$$

Why Koopman Invariance Fails. The Koopman operator $\mathcal{K}$ acts on observables—functions of the system state—rather than on the state space itself. As such, it is an operator on an *infinite-dimensional* function space. In contrast, when we construct a finite-dimensional approximation of the dynamics (e.g., using the singular value decomposition (SVD) of a Hankel matrix), we obtain a truncated, low-rank model:

$$\mathbf{H} \approx \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^*.$$

This truncation yields a rank-r approximation that spans only a small subspace of the full Koopman-invariant function space.

Because this subspace is incomplete, it is generally not invariant under the action of $\mathcal{K}$. That is, for any of the dominant modes $v_i \in \text{span}(\mathbf{U}_r)$, we typically have:

$$\mathcal{K} v_i \notin \text{span}(\{v_1, \ldots, v_r\}),$$

so the Koopman evolution "spills out" of the resolved subspace.

This misalignment forces the system to express part of the dynamics outside the low-rank span. In the HAVOK framework, this residual behavior is captured by a forcing term. Specifically, the dynamics of the resolved modes $\mathbf{v}(t) \in \mathbb{R}^{r-1}$ are governed by:

$$\dot{\mathbf{v}}(t) = A\mathbf{v}(t) + Bv_r(t),$$

where $v_r(t)$ corresponds to the unresolved component of the Koopman action. If the span $\{v_1, \ldots, v_{r-1}\}$ were truly Koopman-invariant,then the dynamics would be closed and linear, and the forcing term $v_r(t)$ would vanish. In practice, the forcing term captures the projection error—energy that spills into the orthogonal complement of the truncated Koopman-invariant subspace.

4 Toy Experiment for Illustration

Toy Demonstration on Lorenz Attractor. The Lorenz system is an ideal candidate for HAVOK because its dynamics exhibit a dual structure: regular oscillations within lobes that are well approximated by a linear model, and rare lobe-switching events that violate Koopman invariance. Delay embedding followed by low-rank SVD cleanly separates these regimes: the resolved coordinates capture linear evolution, while the single unresolved mode captures nonlinearity and memory effects. HAVOK leverages this structure to forecast future states and detect regime shifts, making the Lorenz system both a technically valid and conceptually illuminating test case.

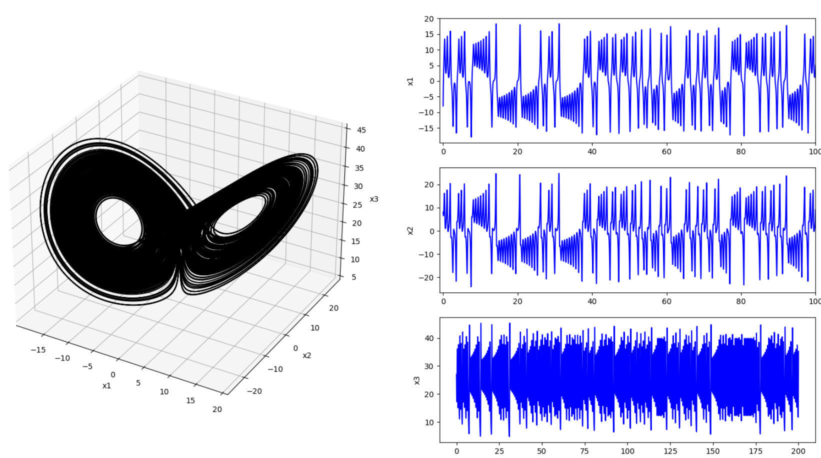

Figure 1: Left: Learned state transition matrix A from the HAVOK model, governing the linear evolution of the delay-embedded coordinates $\hat{\mathbf{v}}(t) = [v_1, \ldots, v_{r-1}]^\top$. Its near-tridiagonal structure reflects the local coupling between time-delay coordinates. Right: Control vector B, which injects the effect of the forcing term $v_r(t)$ into the dynamics. The forcing term predominantly influences the last few coordinates, corresponding to modes where nonlinear transitions (e.g., lobe switches) are most visible.

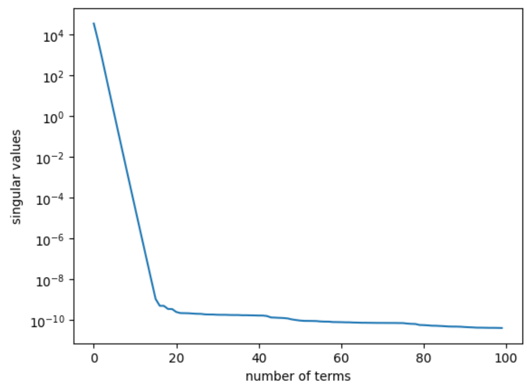

Figure 2: Singular values of the Hankel matrix constructed from a scalar observable of the Lorenz system. The rapid decay indicates that the dynamics can be well-approximated by a low-rank subspace. The first $r-1$ modes are retained for linear modeling, while the r^{th} mode is used as the forcing signal in the HAVOK framework.

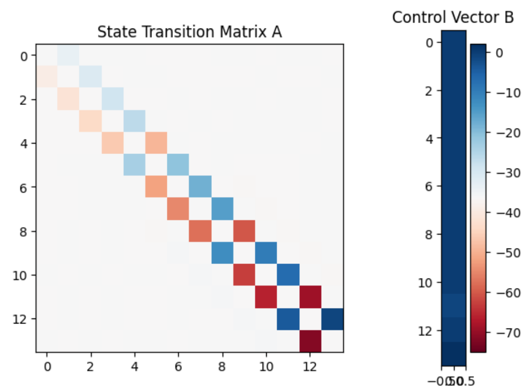

Figure 3: Left: State transition matrix A learned from the top $r-1$ singular modes in the HAVOK decomposition. Its near-tridiagonal structure suggests locally coupled linear dynamics in the delay-embedded coordinates. Right: Control vector B, indicating how the unresolved mode $v_r(t)$ (interpreted as forcing) influences the resolved dynamics. The forcing acts most strongly on the final few modes, highlighting their role in capturing nonlinear transitions.

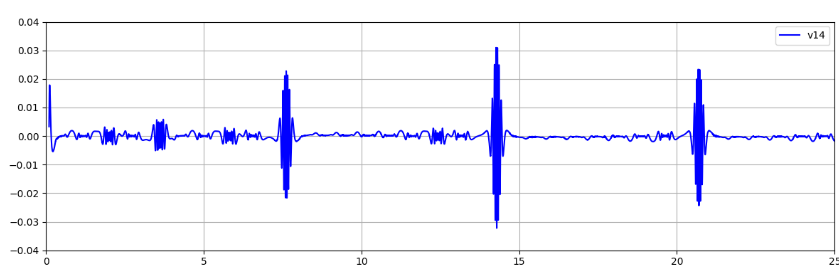

Figure 4: presents two key outputs of the HAVOK model applied to the Lorenz system. The top panel shows the evolution of the final delay coordinate $v_{14}(t)$, which serves as the forcing input to the linear dynamics. This component captures departures from Koopman-invariant behavior and becomes active primarily during lobe-switching events—nonlinear transitions between attractor wings. These brief, high-magnitude bursts indicate when the linear model alone fails to describe the dynamics.

The bottom panel compares the original system state $x_0(t)$ (solid blue) with the HAVOK model's prediction $\hat{x}_0(t)$ (dashed red). The near-perfect overlap demonstrates that the HAVOK model, despite being a low-rank linear approximation with a single forcing channel, accurately reconstructs the nonlinear Lorenz dynamics. This validates the method's ability to extract predictive structure from chaotic time series.

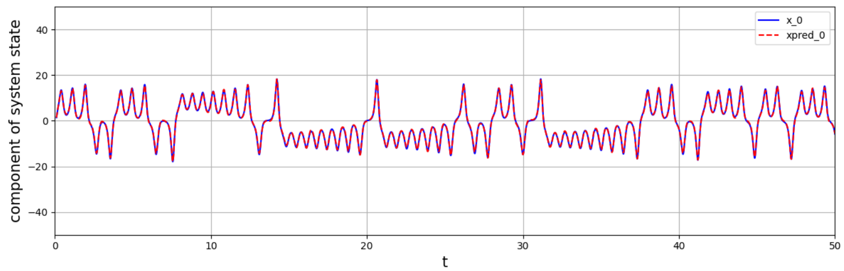

Figure 5: Time series of the forcing term $v_{14}(t)$, indicating nonlinear regime transitions. Bottom: Comparison between true signal $x_0(t)$ and HAVOK prediction $\hat{x}_0(t)$, showing excellent reconstruction accuracy.

5 Generalization: Applying HAVOK to the Rössler System

Demonstrating Generalizability with the Rössler System. To test the generality of the HAVOK framework, we apply it to the Rössler system—another well-known chaotic system governed by a three-dimensional continuous-time ODE. Compared to the Lorenz attractor, the Rössler system features simpler spiral dynamics with intermittent vertical excursions in the z-direction, making it a useful contrasting case for validating the robustness of data-driven decompositions.

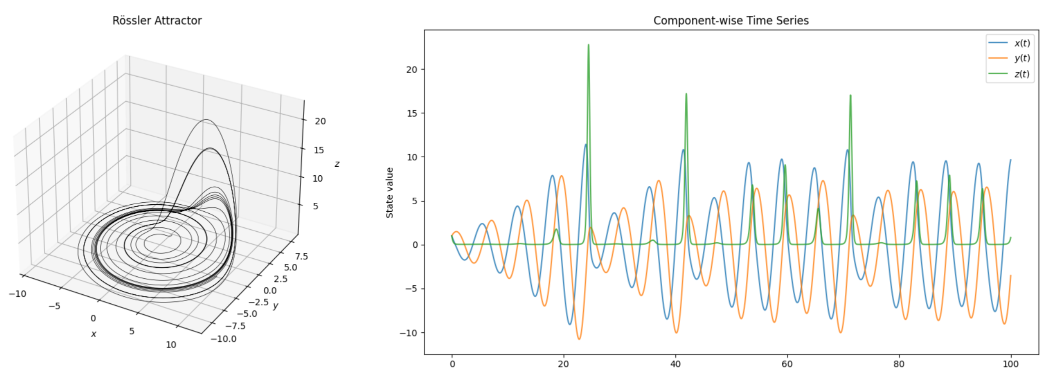

Figure 6: Rössler attractor and its component-wise time series. The spiral structure in 3D and the smooth oscillatory behavior of $x(t)$ and $y(t)$ contrast with the intermittent spikes in $z(t)$.

Using the same HAVOK pipeline with a single observed variable, we construct a time-delay embedded Hankel matrix and apply singular value decomposition. The decay of singular values, shown in Figure 7, reveals a clear spectral gap—indicating that the underlying dynamics can be captured in a low-rank subspace.

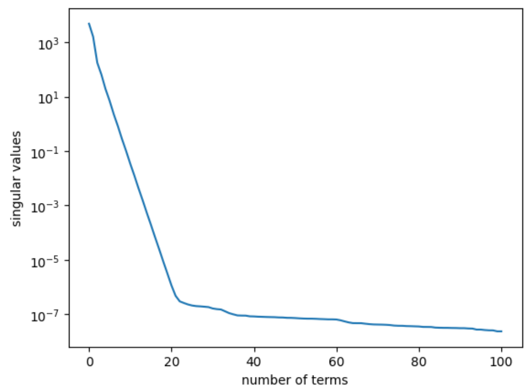

Figure 7: Spectral decay of singular values in the delay-embedded Hankel matrix. A rapid drop confirms the presence of dominant linear modes.

From the SVD basis, we estimate the state transition matrix A and the forcing vector B. As seen in Figure 8, A exhibits near-tridiagonal structure, characteristic of shift dynamics, while B places forcing on a higher mode $v_{14}(t)$, capturing rare deviations from Koopman invariance.

Weaker Resolved Dynamics in the Rössler System. Comparing the HAVOK decompositions of the Lorenz and Rössler systems, we observe that the state transition matrix $\mathbf{A}$ in the Rössler case is notably lighter in magnitude. This suggests that the resolved coordinates $[v_1, \ldots, v_{r-1}]$ carry weaker intrinsic dynamics and are insufficient to linearly capture the system's evolution on their own. Interestingly, the control vector $\mathbf{B}$, which modulates the unresolved forcing $v_r(t)$, remains similarly strong across both systems. This indicates that, despite the Rössler system's more regular and continuous oscillations, the nonlinear corrections from the forcing term remain essential. Thus, while Lorenz requires strong forcing to account for abrupt regime shifts, Rössler depends on forcing to augment a weaker resolved structure, underscoring the broader role of the forcing term in stabilizing low-rank approximations of chaotic flows.

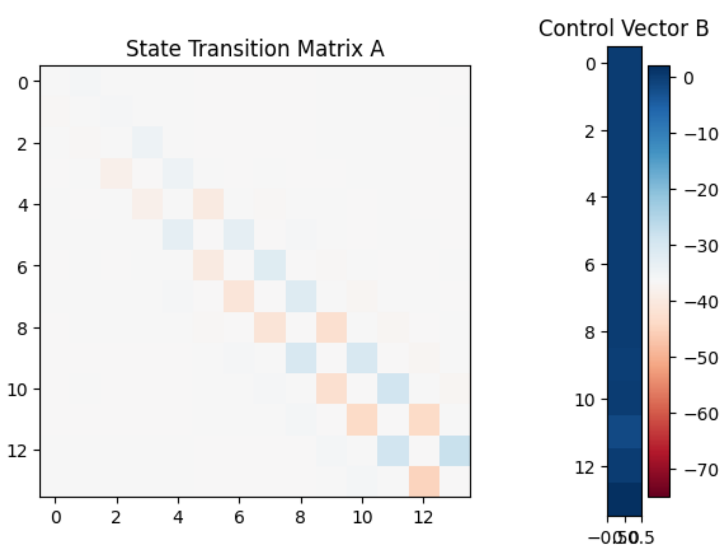

Figure 8: Left: State transition matrix A showing strong local coupling. Right: Control vector B used for forcing in the highest retained mode.

The extracted forcing term $v_{14}(t)$ is shown in Figure 9. As with the Lorenz system, it activates sparsely during rare events—typically during large excursions in $z(t)$. This supports the interpretation of forcing as a memory-closure term compensating for finite-dimensional truncation.

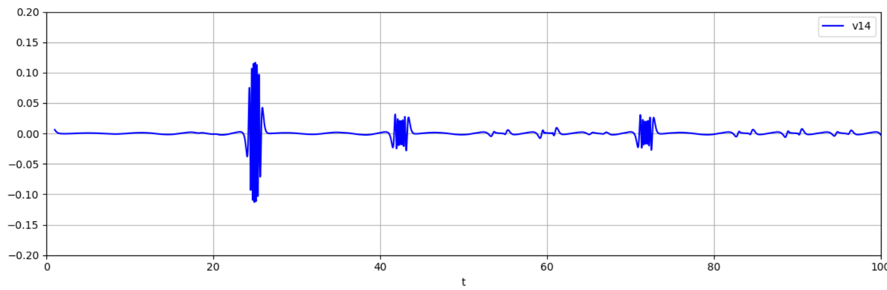

Figure 9: The forcing signal $v_{14}(t)$ shows sparse bursts aligned with rare transitions, supporting the HAVOK perspective that forcing corrects for non-Markovianity induced by model reduction.

Finally, we evaluate short-horizon prediction performance of the HAVOK model using A and B. As seen in Figure 10, the predicted trajectory $\hat{x}_0(t)$ closely matches the true $x(t)$, verifying the fidelity of the reduced model in the resolved linear subspace.

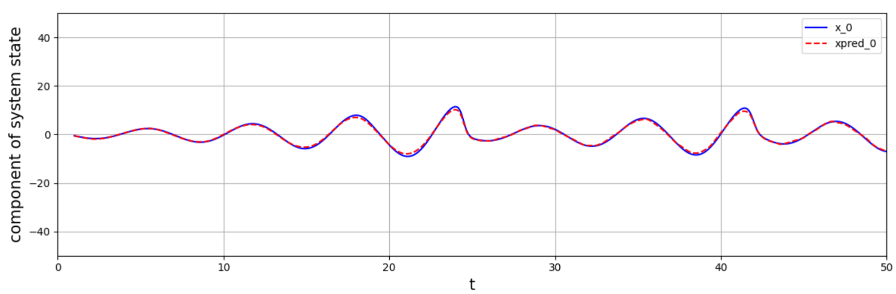

Figure 10: HAVOK forecast of the resolved variable $x(t)$ against the true signal. The model reproduces the local structure with high fidelity.

Interpretation. By applying HAVOK to the Rössler system, we demonstrate the method's flexibility and relevance beyond a single class of chaotic dynamics. Despite the system's different structure from Lorenz—featuring smoother local oscillations and a single-lobe attractor—the HAVOK decomposition remains effective. This confirms that delay-based low-rank moeling is not only system-specific but reflects more general principles of approximating nonlinear flows via Koopman-aligned coordinates.

6 Limitations and Reflections

Limitation 1: Discretization and Time Delay Assumptions. A structural limitation of the HAVOK framework lies in its reliance on fixed-lag time-delay embeddings, typically using a constant lag parameter τ. This assumes uniform time sampling and a stationary system structure. However, real-world systems often involve irregular sampling or time-varying latent dynamics, making the assumption of a constant τ too rigid. The current formulation of HAVOK does not account for variable time lags or adaptive embeddings, which limits its applicability to more complex, temporally heterogeneous datasets.

Limitation 2: Overfitting to the Lorenz System. While HAVOK demonstrates impressive performance on the Lorenz system, this very success may mask its broader limitations. The Lorenz attractor possesses features—such as low intrinsic dimensionality, clean lobe-switching events, and relatively smooth spectral decay—that align well with HAVOK's assumptions. This raises the concern that the model may be tuned too closely to the specific traits of the Lorenz system. To assess generalizability, it is crucial to apply HAVOK to a wider range of systems, including those with more rugged or high-dimensional chaotic structures. For example, preliminary results with the Rössler system reveal a weaker resolved linear component, suggesting that HAVOK may struggle when intrinsic structure is more diffuse or nonlinear dynamics dominate.

Acknowledgement

This short report was written independently over the course of two days, alongside a recently submitted NeurIPS project with Justin and Harris titled *Coherent Memory Structures in Neural Fields*. While that submission was my main focus, I became intrigued by the HAVOK framework and took this opportunity to explore it more deeply. This report serves as a personal exercise to clarify my understanding and begin articulating my own perspective on forecasting and model reduction.

What I find especially compelling about HAVOK is its ability to forecast future dynamics using a nearly linear representation, punctuated only by a sparse, interpretable forcing term. I see conceptual parallels between HAVOK and efforts to construct flexible Markovian representations within the Mori–Zwanzig formalism—specifically, models that retain just enough information from the past to remain predictive. These connections resonate with ideas from variational inference and relative entropy, where one seeks minimal but sufficient representations of uncertainty.

Looking forward, I am particularly interested in understanding when the forcing term in HAVOK (or equivalently, the memory and fluctuation terms in MZ) can be made negligible by projecting onto a Koopman-invariant subspace. In such cases, the full nonlinear dynamics reduce to a closed linear system without memory, offering both interpretability and computational efficiency. Studying HAVOK, and the associated notions of flow, Koopman operators, and invariance, has helped me approach these questions from a fresh, flow-based perspective. This work sets the stage for a longer-term goal of bridging operator-theoretic learning and statistical inference in dynamic systems.

References

Steven L Brunton, Bingni W Brunton, Joshua L Proctor, and J Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE*, 11(2):e0150171, 2016. doi: 10.1371/journal.pone.0150171.

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):19, May 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-00030-8.

Seth M Hirsh, Sara M Ichinaga, Steven L Brunton, J Nathan Kutz, and Bingni W Brunton. Structured time-delay models for dynamical systems with connections to frenet–serret frame. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 477(2257):20210097, 2021. doi: 10.1098/rspa.2021.0097.

Holger Kantz and Thomas Schreiber. *Nonlinear Time Series Analysis*. Cambridge nonlinear science series. Cambridge University Press, Cambridge, 2nd edition, 2003. Includes bibliographical references and index.

Naoya Takeishi, Yoshinobu Kawahara, and Takeshi Yairi. Learning koopman invariant subspaces for dynamic mode decomposition. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017.

Floris Takens. Detecting strange attractors in turbulence. In David Rand and Lai-Sang Young, editors, *Dynamical Systems and Turbulence, Warwick 1980*, pages 366–381, Berlin, Heidelberg, 1981. Springer. ISBN 978-3-540-38945-3. doi: 10.1007/BFb0091924.