

LINOCS: Lookahead Inference of Networked Operators for Continuous Stability

Anonymous authors

Paper under double-blind review

Abstract

Identifying latent interactions within complex systems is key to unlocking deeper insights into their operational dynamics, including how their elements affect each other and contribute to the overall system behavior. For instance, in neuroscience, discovering neuron-to-neuron interactions is essential for understanding brain function; in ecology, recognizing the interactions among populations is key for understanding complex ecosystems. Such systems, often modeled as dynamical systems, typically exhibit noisy high-dimensional and non-stationary temporal behavior that renders their identification challenging. Existing dynamical system identification methods often yield operators that accurately capture short-term behavior but fail to predict long-term trends, suggesting an incomplete capture of the underlying process. Methods that consider extended forecasts (e.g., recurrent neural networks) lack explicit representations of element-wise interactions and require substantial training data, thereby failing to capture interpretable network operators. Here we introduce **Lookahead-driven Inference of Networked Operators for Continuous Stability (LINOCS)**, a robust learning procedure for identifying hidden dynamical interactions in noisy time-series data. LINOCS integrates several multi-step predictions with adaptive weights during training to recover dynamical operators that can yield accurate long-term predictions. We demonstrate LINOCS’ ability to recover the ground truth dynamical operators underlying synthetic time-series data for multiple dynamical systems models (including linear, piecewise linear, time-changing linear systems’ decomposition, and regularized linear time-varying systems) as well as its capability to produce meaningful operators with robust reconstructions through various real-world examples.

1 Introduction

Uncovering the dynamics underlying high-dimensional time-series data is crucial for deciphering the fundamental principles that govern temporally evolving systems. This is apparent across significant scientific domains, including neuroscience (where neurons or ensembles interact over time (Vaadia et al., 1995; D’Aleo et al., 2019; Mudrik et al., 2023)), immunology (where cells regulate immune responses (Savill et al., 2002)), and ecology (where understanding population interactions yields insights into ecosystem dynamics (Stein et al., 2013)). Hence, scientific research necessitates the development of procedures adept at learning dynamic operators that can accurately capture the non-linear and non-stationary evolution of real systems.

Existing approaches for dynamical systems identification, though, often rely on either “black-box” deep learning methods, which while powerful, yield uninterpretable representations, or on simple learning procedures that maximize reconstruction between consecutive samples, and thus fail to accurately predict the system’s behavior for long scales. Specifically, common dynamical systems identification models often rely on optimizing the dynamics by minimizing the prediction error for each time point based on projecting the preceding one through the dynamics. Consequently, when using such procedures to learn the operators, post-learning long-term predictions of the system’s values (by iteratively estimating the system’s state at the next time point) often result in undesired divergence away from the system. This difficulty in long-term predictions, importantly, implies that operators recovered by these models based on local cost functions may

not capture the underlying system correctly. The challenge in identifying such underlying operators, therefore, lies in the need to incorporate long-term predictions directly into the learning procedure, which can be especially challenging in cases where the dynamics are non-stationary, non-linear, or otherwise constrained in ways that reflect real-world system behavior.

To address this challenge, we present a learning procedure that introduces Lookahead Inference of Networked Operators for Continuous Stability (LINOCS). LINOCS bridges the gap between minimizing reconstruction costs based on single time-step projections (which often result in operators that quickly diverge in long-term forecasts), and optimizing multi-step training that relies on reconstructions from a past time point (which can lead to unstable predictions). LINOCS achieves this by integrating adaptively re-weighted multi-step reconstructions into the dynamics inference. LINOCS also avoids relying on massive amounts of data (like RNNs, for example, require). LINOCS’ re-weighting progressively builds up the cost over training iterations, simultaneously considering several multi-step reconstruction terms for identifying operators that enable stable, long-term reconstruction post-training. We demonstrate the effectiveness and adaptability of LINOCS across a range of dynamical systems models, including linear, switched linear, decomposed systems, and smoothly linear time-varying systems, achieving significantly improved accuracy in operator identification and long-term predictions. Our contributions in this paper notably include:

- We propose LINOCS, a novel learning procedure that incorporates re-weighting multi-step predictions into the cost for operator identification.
- We demonstrate that applying LINOCS improves the ability to recognize ground-truth operators.
- We show LINOCS’ efficacy across a diverse range of dynamical systems, including linear, periodically linear, linear time-varying (LTV), and decomposed linear.
- Finally, we demonstrate LINOCS’ ability to work on real-world brain recordings, resulting in better long-term reconstruction compared to baselines.

2 Background and Terminology

Consider a system with p interacting elements (e.g., neurons in the brain) whose time-changing state $\mathbf{X} \in \mathbb{R}^{p \times t}$ evolves over discrete time points $t = 1 \dots T$ as $\mathbf{x}_{t+1} = g(\mathbf{x}_t, \mathbf{b}_t, t)$, where $\mathbf{x}_t \in \mathbb{R}^p$ refers to the state at time t , $\mathbf{b}_t \in \mathbb{R}^p$ represents an external input or driving force at time t , and g is a function $g : \mathbb{R}^p \times \mathbb{R}^p \times \mathbb{Z} \rightarrow \mathbb{R}^p$. For example, \mathbf{x}_t can represent the time-evolving activity of p recorded neurons over T time points in neuroscience applications, or \mathbf{x}_t can represent the activation levels of p immune cells when modeling the immune system.

In this paper we focus on linear, piece-wise linear, and linear time-varying systems. Specifically, we limit our analysis to functions $g(\cdot)$ that can be written as:

$$\mathbf{x}_{t+1} = g(\mathbf{x}_t, \mathbf{b}_t, t) := \mathbf{A}_t \mathbf{x}_t + \mathbf{b}_t, \tag{1}$$

where $\mathbf{A}_t \in \mathbb{R}^{p \times p}$ represent the transition matrices at each time t . Our focus on locally linear dynamics is supported by the fact that even highly nonlinear functions can be well approximated over small time intervals using local linearization (Khalil, 2001; Sastry, 2013). Importantly, this formulation’s advantage lies in its retention of the ability to easily extract the system’s pairwise interactions, including non-stationary changes in \mathbf{A}_t and \mathbf{b}_t over time. Specifically, any operator entry $[\mathbf{A}_t]_{i,j}$ for $i, j = 1 \dots p$ at every $t = 1 \dots T$ can be interpreted as the effect of element j on element i at time t .

In practice, however, robustly recovering operators that can accurately describe the system’s evolution in non-stationarity and non-linear settings, faces numerous computational challenges. Chiefly, if we adopt a naive approach to identify the operators as: $\hat{\mathbf{A}}_t, \hat{\mathbf{b}}_t = \arg \min_{\mathbf{A}_t, \mathbf{b}_t} \|\mathbf{x}_t - \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t\|_2^2$ for every $t = 1 \dots T$, this problem is statistically unidentifiable. Specifically, the problem has $p^2 + p$ unknowns for each time point, but only p equations.

One approach to improve inference in these settings is to introduce additional structure via a prior over \mathbf{A}_t and \mathbf{b}_t , that constrains the solution space and is often grounded in application-driven assumptions. For

instance, in many scientific settings, it is reasonable to assume that interactions change smoothly over time. Therefore, adding a temporal smoothness constraint on \mathbf{A} and \mathbf{b} (e.g., $\|\mathbf{A}_t - \mathbf{A}_{t-1}\|_F^2$ and $\|\mathbf{b}_t - \mathbf{b}_{t-1}\|_2^2$) can be beneficial for both interpretability and accuracy. Additionally, the inclusion of such constraints can be crucial, particularly in noisy settings, to prevent overfitting. The addition of such constraints transforms the problem to:

$$\widehat{\mathbf{A}}_t, \widehat{\mathbf{b}}_t = \arg \min_{\mathbf{A}_t, \mathbf{b}_t} \|\mathbf{x}_{t+1} - \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t\|_2^2 + \mathcal{R}(\mathbf{A}_t, \mathbf{b}_t), \quad (2)$$

where $\mathcal{R}(\mathbf{A}_t, \mathbf{b}_t)$ can represent regularization on the dynamics operators.

While these solutions may offer good short-term predictions (i.e., predicting \mathbf{x}_t from \mathbf{x}_{t-1}) with minor errors, they often struggle to fully reconstruct the dynamics over longer time-scales due to the build-up of estimation errors over multiple predictions (e.g., if $\mathbf{x}_t | \mathbf{x}_{t-1}$ has a variance of v_t , then $\mathbf{x}_t | \mathbf{x}_{t-K}$ will have a variance of $\sum_{k=0}^{K-1} v_{t-k}$ due to the property of variance summation). This issue of multiple potential solutions for the operators is further complicated by the practical consideration that we typically can only observe noisy observations of \mathbf{x}_t ($\tilde{\mathbf{x}}_t = \mathbf{x}_t + \boldsymbol{\epsilon}_t$, where $\boldsymbol{\epsilon} \in \mathbb{R}^{p \times T}$ represents some noise). Noisy observations further complicate the accurate identification of robust dynamics since if $\widehat{\mathbf{A}}_t$ is obtained from $\arg \min_{\mathbf{A}_t} \|\mathbf{x}_{t+1} - \mathbf{A}_t \tilde{\mathbf{x}}_t\|_2^2$, the distance between the real and the estimated operators $\{\|\mathbf{A}_t - \widehat{\mathbf{A}}_t\|_F^2\}_{t=1}^T$ may increase with $\boldsymbol{\epsilon}_t$.

As the accuracy of a dynamical system's fit to data is often evaluated based on its ability to accurately predict future values (Tabar, 2019), the inability to capture long-term prediction suggests that the learned operators may have limited capacity to fully describe the system. Consequently, we define below three prediction styles for dynamical systems assessment that we will adhere to throughout this work:

- **1-Step Prediction ($\mathbf{x}_{t+1} | \mathbf{x}_t$):** 1-Step Prediction involves using the state at each time point t to estimate the state at the next time point ($t + 1$).
- **Iterative Multi-Step Prediction (IMS) of order $K \in \mathcal{R}$ ($\mathbf{x}_{t+k} | \mathbf{x}_t$):** IMS involves iteratively forecasting one-step ahead values and using these forecasts as inputs for further one-step ahead forecasts for K times (i.e., $\forall k \in [0, K - 1]$, $\tilde{\mathbf{x}}_{t+k} | \tilde{\mathbf{x}}_{t+k-1}$, where $\tilde{\mathbf{x}}_{t-1} := \tilde{\mathbf{x}}_{t-1}$). Here, we will notate an IMS prediction of order K by $\tilde{\mathbf{x}}_t^K$. We chose to name this term IMS as to be consistent with the literature (Chevillon, 2007).
- **Full Lookahead Prediction ($\mathbf{x}_k | \mathbf{x}_0$):** This method enhances IMS by forecasting the state at each time point \mathbf{x}_t starting from the initial observations ($\tilde{\mathbf{x}}_0$). It achieves this by sequentially applying transition matrices to the estimation from the previous time point $\tilde{\mathbf{x}}_t | \tilde{\mathbf{x}}_{t-1}$, starting from $\tilde{\mathbf{x}}_0$, resulting in: $\tilde{\mathbf{x}}_t = \widehat{\mathbf{A}}_{t-1} \dots \widehat{\mathbf{A}}_0 \tilde{\mathbf{x}}_0 \quad \forall t = 1 \dots T$.
(Note: the formula above is presented without \mathbf{b}_t for simplicity, though they may be included).

Importantly, IMS Prediction and Full Lookahead Prediction often result in instability due to the accumulation of errors in the sequential reconstructing process (Fig. 1A).

2.1 Specific models considered in this work:

Of particular interest in this work is improving the model fit of a core set of linear dynamical systems with different temporal constraints on the system evolution: 1) Time-Invariant Linear Dynamical Systems (LDS), 2) Switched Linear Dynamical Systems (SLDS), 3) decomposed Linear Dynamical Systems (dLDS), and 4) regularized Linear Time-Varying (LTV) Dynamical Systems.

Time-Invariant Linear Dynamical Systems (LDS). In linear systems analysis, the evolution of a general state \mathbf{X} over $T + 1$ time points can be typically represented as $\mathbf{x}_{t+1} = \mathbf{A} \mathbf{x}_t + \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{p \times p}$ is the time-invariant dynamics matrix and $\mathbf{b} \in \mathbb{R}^{p \times 1}$ remain constant over time. One common method for determining \mathbf{A} (and \mathbf{b} if it is assumed that an unknown offset exists) involves a 1-step optimization approach that includes applying least squares across all time points. This entails solving

$$\widehat{\mathbf{A}}, \widehat{\mathbf{b}} = \arg \min_{\mathbf{A}, \mathbf{b}} \|\widetilde{\mathbf{X}}_{[:,1:T]} - \mathbf{A} \widetilde{\mathbf{X}}_{[:,0:T-1]} - [\mathbf{1}]_{1 \times T} \otimes \mathbf{b}\|_F^2, \quad (3)$$

where $\widetilde{\mathbf{X}}_{[:,1:T]}$ and $\widetilde{\mathbf{X}}_{[:,0:T-1]}$ represents the noisy observations of the state from the second time point ($t = 1$) up to the last time point (T) and from the first time point ($t = 0$) up to $T - 1$, respectively, and $[1]_{1 \times T} \otimes \mathbf{b}$ represents the horizontal concatenation of the column vector \mathbf{b} horizontally T times. Here, \mathbf{A} captures the average influence from \mathbf{x}_{t-1} to \mathbf{x}_t for all $t = 1 \dots (T + 1)$. This setting is advantageous in terms of “network” interpretability, however often cannot capture the complexity of real-world time-series which is non-linear and non-stationary.

Switching Linear Dynamical Systems (SLDS). Switched Linear Dynamical Systems (SLDS) (Ackerson and Fu, 1970; Bar-Shalom and Li, 1990; Hamilton, 1990; Ghahramani and Hinton, 1996; Murphy, 1998; Fox et al., 2008; Linderman et al., 2017) aim to provide interpretable representations of dynamics by identifying linear operators that govern periods of linear behavior, with the system transitioning between these operators over time. Variations of SLDS include, e.g., recurrent SLDS (rSLDS), which introduces an additional dependency between discrete switches and the previous state’s location in space (Linderman et al., 2017); and tree-structured recurrent SLDS, which extends rSLDS by incorporating a generalized stick-breaking procedure (Nassar et al., 2018).

While SLDS models often involve transitioning from an observed to a latent low-dimensional space, here we chose to focus on the case where switches occur within the observation space, essentially enforcing the transition to the latent space to be the identity operator. If we denote $\widetilde{\mathbf{X}} \in \mathbb{R}^{N \times T}$ as the noisy observations subjected to *i.i.d* Gaussian noise, SLDS models the evolution of $\widetilde{\mathbf{x}}_t$ using a set of J discrete states ($j = 1 \dots J$), each state j associated with its own linear dynamical system \mathbf{f}_j . These discrete states switch between them abruptly at certain time points following an HMM model. During each inter-switch period, if the system is in the j -th discrete state, SLDS models the evolution of the state linearly as $\mathbf{x}_t = \mathbf{f}_j \mathbf{x}_{t-1} + \mathbf{b}_j$, where \mathbf{f}_j represents the linear transition matrix for the j -th discrete state and \mathbf{b}_j denotes a constant offset term for that discrete state. SLDS can be trained by an alternating set of steps between the dynamics learning and the HMM update of the operators. SLDS tackles the crucial task of capturing non-stationarities while preserving interpretability, but it inherently lacks the capability to distinguish between multiple co-occurring processes or overlapping subsystems.

decomposed Linear Dynamical Systems (dLDS). The Decomposed Linear Dynamical Systems (dLDS, Mudrik et al. (2024)) model relaxed the time-invariant or piecewise constant limitation of LDSs and SLDSs to support the discovery of co-occurring processes while maintaining interpretability. Here, for simplicity, we focus on the case where the dynamics evolution is described directly in the observation space, while the full model presented in (Mudrik et al., 2024) supports learning the dynamics within an identified latent state. Specifically, dLDS models the dynamics evolution $\widetilde{\mathbf{x}}_t = \mathbf{A}_t \widetilde{\mathbf{x}}_{t-1}$ using a sparse time-changing decomposition of linear dynamical operators such that $\mathbf{A}_t = (\sum_{j=1}^J \mathbf{f}_j c_{jt})$, resulting in $\widetilde{\mathbf{x}}_t = (\sum_{j=1}^J \mathbf{f}_j c_{jt}) \widetilde{\mathbf{x}}_{t-1}$. These dynamical operators ($\{\mathbf{f}_j\}_{j=1}^J$) are global, i.e., not time dependent, and hence are interpretable globally. However, their time-changing weights (c_t) enable modeling non-stationary and non-linear dynamics (Fig. 2 right). Notably, dLDS is trained through an Expectation-Maximization (EM) procedure where the global dynamics operators $\{\mathbf{f}_j\}_{j=1}^J$ and their time-changing coefficients $\{c_{jt}\}_{t=1}^T$ are updated iteratively to maximize the posteriors as:

$$\{\widehat{c}_t\}_{t=1}^T = \arg \max_{\{c\}} P(\{c_t\}_t | \widetilde{\mathbf{X}}, \{\mathbf{f}_j\}) \quad (4)$$

$$\{\widehat{\mathbf{f}}_j\}_{j=1}^J = \arg \max_{\{\mathbf{f}_j\}} P(\{\mathbf{f}_j\} | \widetilde{\mathbf{X}}, \{c_t\}). \quad (5)$$

Interestingly, dLDS can also capture linear or switching behaviors described earlier, by fixing the dLDS coefficients over time (for linear behavior, Fig. 2 far left) or supporting abrupt change of coefficients in specific time points (for switching behaviors, Fig. 2 middle left). As dLDS estimates the parameters for each time t solely based on the values of the preceding state at time $t - 1$, it does not address the issue of inaccurate long-term prediction due to accumulated deviations.

Smooth or Sparse Linear Time-Varying Systems (LTV). In this paper, we refer to LTV systems that can be described by: $\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t$ for all $t = 1 \dots T$. We further assume that a regularization $\mathcal{R}(\mathbf{A}_t)$ may be applied to the operators $\{\mathbf{A}_t\}_{t=1}^T$. This regularization can be inspired by the application (e.g., smoothness

of operators over time, $\|\mathbf{A}_t - \mathbf{A}_{t-1}\|_F^2 < \epsilon_2$ or operator sparsity $\|\text{vec}(\mathbf{A}_t)\|_0 < \epsilon_1$) and mitigates the ill-posed nature of finding \mathbf{A}_t separately for each time point.

2.1.1 Prior relevant approaches:

Theoretical literature on long-term prediction instability traces back to Cox (1961) and Klein (2019), who respectively introduced exponential smoothing and direct estimation of distant future states. Subsequent studies, including Findley (1983; 1985); Weiss (1991); Tiao and Xu (1993); Lin and Granger (1994); Kang (2003), evaluated the effectiveness of dynamical-systems identification methods in yielding long-term predictions. Specifically, the first approach focuses on identifying dynamical operators by minimizing the reconstruction error of projecting the state from one time point to the next, which can subsequently be used for long-term predictions through IMS. The second method, direct forecasting, aims to identify a mapping function $\mathbf{F}_{kt} : \mathbb{R}^p \rightarrow \mathbb{R}^p$ that predicts states further into the future by $\mathbf{x}_{t+k} = \mathbf{F}_{kt}(\mathbf{x}_t)$, thereby skipping the explicit identification of intermediate dynamic operators. While direct estimation naturally results in more stable long-term predictions compared to 1-step optimization, it fails to provide an interpretable “network” meaning to the operators (e.g., in neuroscience, understanding the brain’s interactions entails discerning the time-changing fast transitions from \mathbf{x}_t to \mathbf{x}_{t+1}). In fact, when Marcellino et al. (2006) compared between iterated and direct estimations using macroeconomic data, they found that in contrast to previous assumptions, iterated forecasts outperform direct forecasts, especially when models can select long-lag specifications—raising questions about the appropriate approach for learning dynamical operators.

More recently, Venkatraman et al. (2015) proposed a general approach called DAD that reuses training data to build a no-regret learner with multi-step prediction. However, DAD includes “fixing” and updating the model itself based on every step within the multi-step prediction. Moreover, the authors presented it as a general abstract non-specific approach to consider without specific implementation details. More importantly, DAD does not discuss the possibility for priors over the operators (e.g. temporal smoothness) during the training nor did they consider the need to find operators that are not only expressive but also interpretable. Other learning procedures, including full forward and/or backward passes through e.g., Backpropagation Through Time (BPTT) as in Recurrent Neural Networks (RNNs), can partially handle long-term prediction instability as well as Lookahead extensions to RNNs, including Unni et al. (2023); Xiao et al. (2023); Yeung et al. (2019), that leverage the Koopman operator for improved long-term prediction. However, these methods often require extensive learning data and remain uninterpretable in the “network” sense, as non-linear state measurements in Koopman operators can obscure understanding of pairwise state interactions.

An additional approach to understanding dynamical systems involves identifying a sparse set of functions that jointly decompose the observations. For example, SINDy (Sparse Identification of Nonlinear Dynamics, Brunton et al. (2016)) utilizes a data-driven approach to discover governing fundamental equations from data using sparse regression. Although SINDy and its extensions (e.g., Kaheman et al. (2020)) are promising for discovering governing data equations, such representation does not provide explicit insight into the time-changing interactions between the state elements.

3 LINOCS

In LINOCS we aim to learn the unknown dynamic operators $\{\widehat{\mathbf{A}}_t\}_{t=1}^T$ by integrating several multi-step predictions simultaneously into the inference procedure. This approach yields not only a more accurate full-lookahead post-learning reconstruction but also operators that are more closely aligned with the ground truth. Particularly, for every $t = 1 \dots T$, LINOCS finds the most likely estimate of $\{\mathbf{A}_t, \mathbf{b}_t\}$ given $K+1$ ($K \in \mathbb{Z}_{\geq 0}$ hyper-parameter) multi-step reconstructions of orders $k = 0 \dots K$ with different weights $\{\mathbf{w}_k\}_{k=0}^K$:

$$\widehat{\mathbf{A}}_t = \arg \min_{\mathbf{A}_t} \sum_{k=0}^K w_k \|\tilde{\mathbf{x}}_{t+1} - \mathbf{A}_t \tilde{\mathbf{x}}_t^k\|_F^2, \quad (6)$$

where $\tilde{\mathbf{x}}_t, \tilde{\mathbf{x}}_{t+1}$ are the observations at time t and $t+1$, respectively. Let $\tilde{\mathbf{x}}_{t+1}^k$ ($k = 0 \dots K$) denote the recursive rule for predicting the state at time $t+1$, starting from the observations at $t-k$, where $\tilde{\mathbf{x}}_{t-k}$ is set

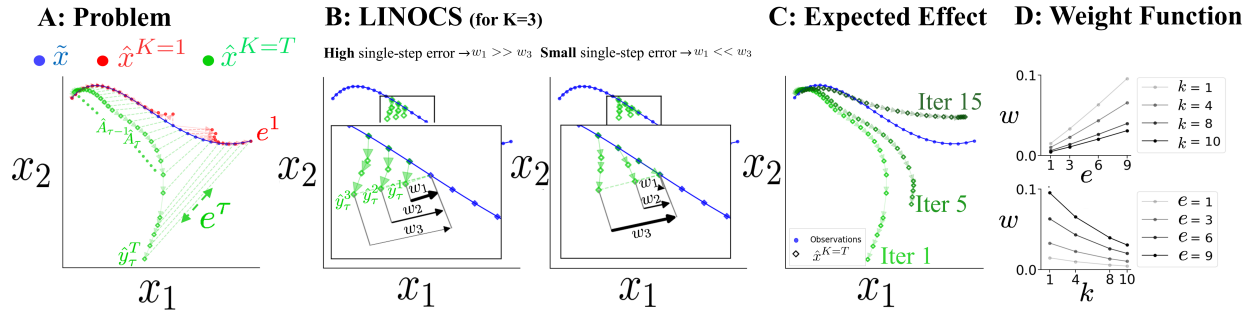


Figure 1: **Illustration of the problem and approach.** **A:** Models that perform well on 1-step prediction (i.e., prediction order 0, in red), often fail in higher orders reconstruction (e.g., order- T , in green). **B:** LINOCS (e.g., for training order $K = 3$) integrates weighted multi-step reconstructions for all $k = 0 \dots K$ orders. It adapts the weights of these reconstruction orders to prioritize minimizing large errors at lower orders before addressing higher orders. The system gradually increases the weight of the most effective lookahead reconstruction until convergence conditions are met. **C:** LINOCS improves long-term reconstruction during training iterations. **D:** Weights of different lookahead training orders (k). $w_k : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a function of the order k and the k -th order reconstruction error e . Top: Illustration of an exemplary effect of e on w when fixing k . Bottom: Illustration of an exemplary effect of k on w when fixing e .

to the observations at time $t - k$, i.e., $\hat{\mathbf{x}}_{t-k} := \tilde{\mathbf{x}}_{t-k}$. Particularly, the $(k + 1)$ -th order multi-step prediction is defined by:

$$\hat{\mathbf{x}}_{t+1}^k = \hat{\mathbf{A}}_t \hat{\mathbf{A}}_{t-1} \hat{\mathbf{A}}_{t-2} \dots \hat{\mathbf{A}}_{t-k} \tilde{\mathbf{x}}_{t-k}. \quad (7)$$

The weights $\{w_k\}_{k=0}^K$ associated with the orders $k = 0 \dots K$ are dynamically adjusted throughout the inference process (Fig. 1B). This adjustment considers both the order number (k) and the current reconstruction error related to that order, e_k (e.g., the ℓ_2 norm, $e_k = \|\tilde{\mathbf{x}}_t - \hat{\mathbf{x}}_t^k\|_2^2$). Unlike other multi-step methods (e.g., Venkatraman et al. (2015)), LINOCS adapts the weights of the reconstruction orders to prioritize the minimization of large errors in lower orders before considering higher orders (Fig. 1B). Specifically, it gradually increases the weight of the best lookahead reconstruction until convergence conditions are satisfied. In our implementations, the weights can be chosen from a list of built-in choices such as uniform, linearly decreasing, and exponentially decreasing weights. Additionally, our framework allows custom weight functions that suit their specific needs. In the experiments presented in this paper, we concentrate on showcasing three specific options for the weights:

- Adapting the weights to sequentially introduce higher multi-step reconstruction orders once the error for each preceding order falls below a designated threshold, while continuing to maintain the activation of lower orders. Specifically, in the initial iterations, only $w_0 > 0$, with all other weights $w_j = 0$ for $j \in [1, K]$. As the error of each step’s reconstruction falls below this threshold, the subsequent weight w_{j+1} is activated. For instance, if w_0 is the last activated weight and the error for the first (1-step) reconstruction falls below the threshold, w_1 becomes active ($w_1 > 0$), and this sequence of activation continues for higher-order weights as each subsequent step achieves the required accuracy.
- Constant weights over iterations with an exponential decay over k , defined as $w_k = \exp^{-\sigma k}$ for some scalar $\sigma \in \mathbb{R}_{>0}$.
- A weight function that considers both k and e_k , exhibiting a monotonic decrease in k and an increase in e , with k decreasing faster than e increases (Fig. 1 D).

Importantly, throughout this paper, we distinguish between two concepts: training order and prediction order. We denote “training order” (K_{train}) as the maximum order considered *during* LINOCS training.




Figure 2: **Example time-varying LDS systems A:** The baseline linear time-invariant dynamical system will present dynamical operator constant over time. **B:** Switching linear dynamical systems (SLDS) jump between different linear operators that are time-invariant between jumps. **C:** “Pseudo-Switching” dynamics is similar to SLDS with the inclusion of smoother transitions between periods of constant linear dynamics. **D:** The decomposed linear dynamical systems (dLDS) model is a generalization of SLDS to sparse time-changing linear combinations of linear operators. dLDS can also model negative coefficients.

Throughout this work “1-step” optimization specifically refers to the use of the 1-step cost ($\|\mathbf{x}_t - \mathbf{A}_t \mathbf{x}_{t-1}\|_2^2$), while excluding higher orders, *during* training. In contrast, prediction order refers to *post-training* predictions that involve iteratively propagating the identified operators for K_{pred} steps into the future.

Here, we demonstrate the contribution of LINOCS for accurate long-scale predictions in four types of systems: 1) time-invariant linear; 2) switching linear; 3) decomposed linear; and 4) LTV systems. Importantly, in our experiments, we assume that we observe the underlying system under additive *i.i.d* Gaussian noise conditions, however LINOCS can be easily adjusted to other noise statistics.

3.1 LINOCS for Linear Dynamics

We first present the LINOCS learning rule for the simplest case of time-invariant linear dynamical systems (TI-LDS). Let $\widetilde{\mathbf{X}} \in \mathbb{R}^{p \times T}$ be the observations of state \mathbf{X} , such that $\widetilde{\mathbf{X}} = \mathbf{X} + \eta$, with η being an *i.i.d* Gaussian noise ($\eta \sim \mathcal{N}(0, \sigma^2)$). In the TI-LDS case, \mathbf{X} evolves linearly as $\mathbf{x}_{t+1} = \mathbf{A} \mathbf{x}_t + \mathbf{b}$ for all $t = 1 \dots T$, where $\mathbf{b} \in \mathbb{R}^{p \times 1}$ is a constant offset.

In this case, LINOCS estimates \mathbf{A} and \mathbf{b} by

$$\widehat{\mathbf{A}}, \widehat{\mathbf{b}} = \arg \min_{\mathbf{A}, \mathbf{b}} \sum_{k=0}^K \mathbf{w}_k \|\widetilde{\mathbf{x}}_{t+1} - \mathbf{A}^{k+1} \widetilde{\mathbf{x}}_{t-k} - \sum_{j=0}^k \mathbf{A}^j \mathbf{b}\|_2^2,$$

where K is a hyperparameter that dictates the maximum reconstruction order, and the set $\{\mathbf{w}_k\}_{k=0}^K$ can be either pre-defined or automatically adapted over training based on each order error.

3.2 LINOCS for Switching Linear Dynamical Systems (SLDS)

For SLDS (Fig. 2, middle-left), we integrate LINOCS into the SLDS operator inference stage (to infer $\{\mathbf{f}\}_{j=1}^J$, $\{\mathbf{b}\}_{j=1}^J$, see Sec. 2.1) using the SSM framework proposed by Linderman et al. (2020). We maintain the existing SLDS approach to estimating switch times that delineate the boundaries of the linear periods between switches (i.e., for this stage, we have kept it as implemented by Linderman et al. (2020)). To find the operators within these identified linear periods, we integrate the learning rule for the TI-LDS case described above in Section 3.1.

3.3 LINOCS for decomposed Linear Dynamical Systems (dLDS)

For dLDS, as in SLDS, we incorporate LINOCS into the dynamical systems update step. Let $\widehat{\mathbf{x}}_{t+1}^k$ denote the k -th order reconstruction of \mathbf{x}_{t+1} , calculated by iteratively propagating the dLDS reconstruction $k + 1$ times, starting from \mathbf{x}_{t-k} . Furthermore, let $\mathbf{x}_{t+1} \approx \sum_{j=1}^J \widehat{\mathbf{c}}_j \widehat{\mathbf{f}}_j \mathbf{x}_t$, where $\widehat{\mathbf{c}}_t$ represents the current estimate of the dLDS coefficients and $\{\widehat{\mathbf{f}}_j\}_{j=1}^J$ denotes the current estimate of the basis operators. We can now write

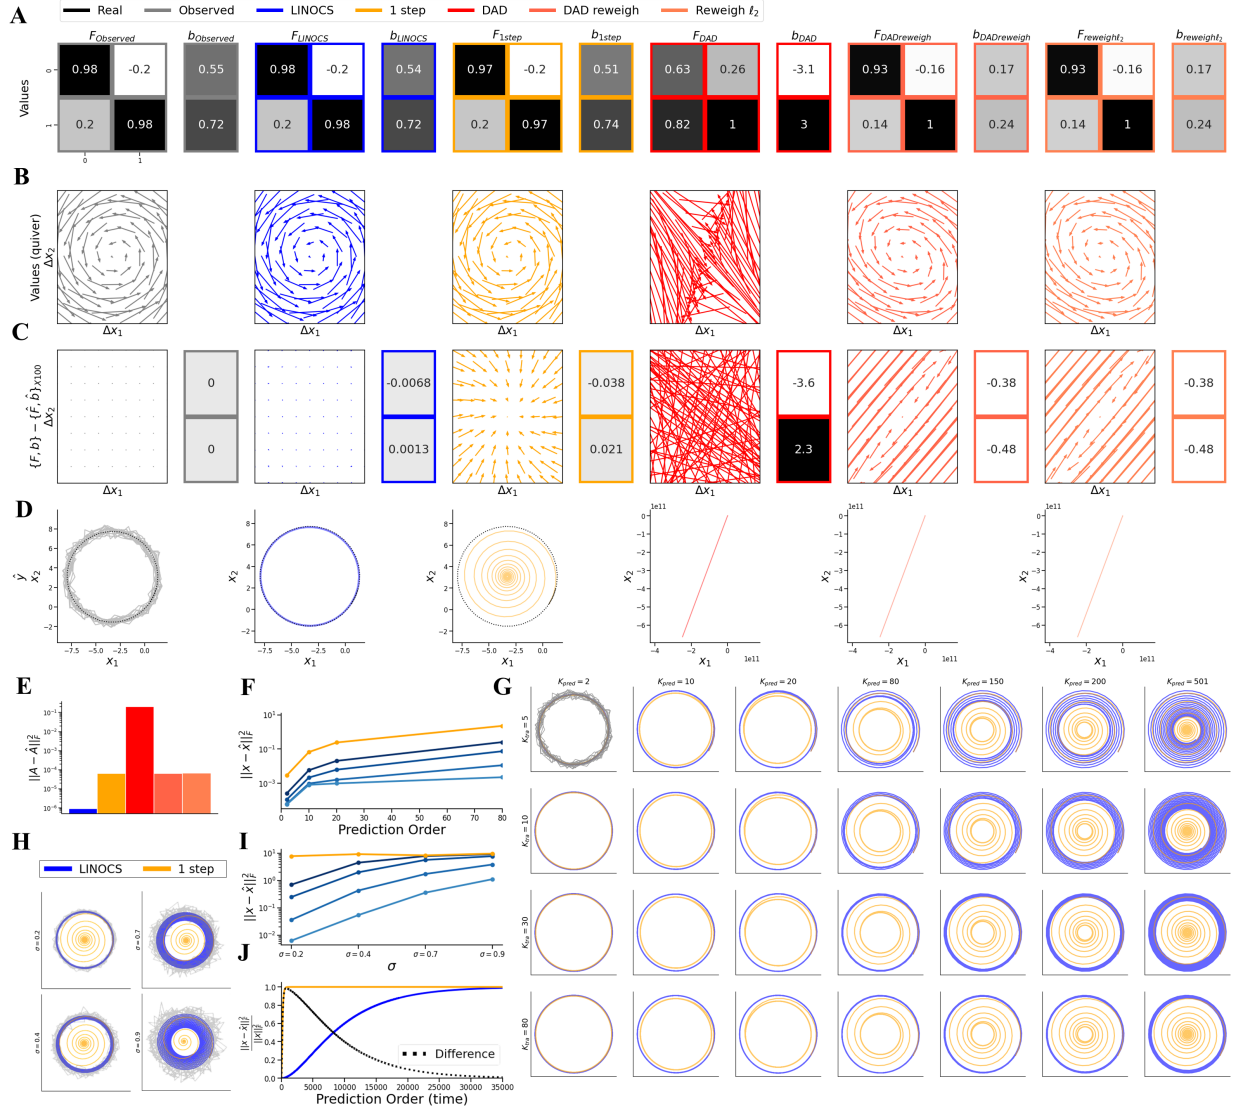


Figure 3: **Linear System Experiment.** **A:** Real vs. identified operators and offsets. **B:** Quiver plots of real and identified operators. **C:** Highlighted differences in effects between real operators and inferred operators showing how small differences in dynamic operators gain prominence during lookahead reconstruction (calculation details in Section A.1). **D:** Full lookahead reconstructions (ground truth vs. baselines) show swift convergence to the origin for the 1-step optimization (yellow) and divergence for DAD-based results (three most right subplots). **E:** Frobenius norm of the differences between the ground truth operators (A) and the identified operators (\hat{A}). **F:** MSE under increasing prediction orders. For all orders, LINOCS achieves better (lower) MSE compared to 1-step optimization. **G:** Full lookahead post-training predictions using operators identified by 1-step optimization (yellow) vs. the predictions using those identified by LINOCS (blue) under various training orders (rows) and prediction orders (columns). **H, I:** LINOCS reconstruction compared to 1-step optimization under increasing noise levels demonstrates its robustness. **J:** Propagating the identified operators until reaching a relative reconstruction error of ~ 1 . LINOCS identifies operators that can accurately predict $\sim 35,000$ time points, much higher than 1-step training that decay immediately.

the k -th order reconstruction (\hat{x}_{t+1}^k) as

$$\hat{x}_{t+1}^k \leftarrow \left(\sum_{j=1}^J \hat{c}_{jt} \hat{f}_j \right) \left(\sum_{j=1}^J \hat{c}_{j(t-1)} \hat{f}_j \right) \cdots \left(\sum_{j=1}^J \hat{c}_{j(t-k)} \hat{f}_j \right) \tilde{x}_{t-k},$$

as follows from the recursive update rule $\hat{\mathbf{x}}_t^k \leftarrow (\sum_{j=1}^J \hat{c}_{jt} \hat{\mathbf{f}}_j) \hat{\mathbf{x}}_{(t-1)}^{k-1}$, where $\hat{\mathbf{x}}_{t-1}^0 := \tilde{\mathbf{x}}_{t-1}$.

To effectively integrate LINOCS into dLDS, we incorporate multi-step predictions into the training procedure of dLDS itself. Specifically, in each iteration, we start with the update of the dynamics coefficients \mathbf{c}_t . Second, we define $\mathbf{F}_{x_t^k} \in \mathbb{R}^{p \times J}$ as the horizontal concatenation

$$\mathbf{F}_{x_t^k} := [\mathbf{f}_1 \mathbf{x}_t^k, \mathbf{f}_2 \mathbf{x}_t^k, \dots, \mathbf{f}_J \mathbf{x}_t^k] \text{ for all } k = 0 \dots K.$$

Next, we define a new matrix $\tilde{\mathbf{F}}_{x_t}^K$ that extends \mathbf{F}_{x_t} to all $K + 1$ reconstructions stacked vertically, resulting in $\tilde{\mathbf{F}}_{x_t}^K \in \mathbb{R}^{(K+1)p \times J}$

$$\tilde{\mathbf{F}}_{x_t}^K = \begin{bmatrix} w_0 \mathbf{F}_{x_t} \\ w_1 \mathbf{F}_{x_t^1} \\ \vdots \\ w_{K-1} \mathbf{F}_{x_t^{K-1}} \end{bmatrix},$$

where w_k is the weight of the k -th multi-step order. This matrix can then be used to infer to coefficients (\mathbf{c}_t) while considering different reconstruction orders with varying weights.

To mirror this concatenated matrix of dynamics that represents multiple time-steps, we further define a matching concatenated state vector $(\tilde{\mathbf{x}}_{t+1})_{vert} \in \mathbb{R}^{p(K+1) \times 1}$ by $(\tilde{\mathbf{x}}_{t+1})_{vert} := w \otimes \tilde{\mathbf{x}}_{t+1}$ where $w = [w_0; w_1; \dots; w_K] \in \mathbb{R}^{(K+1) \times 1}$. I.e., $(\tilde{\mathbf{x}}_{t+1})_{vert}$ is obtained by vertically stacking $K + 1$ times the observations $\tilde{\mathbf{x}}_{t+1} \in \mathbb{R}^{p \times 1}$ at time $t + 1$ weighted by their corresponding w_k values (resulting in $(\tilde{\mathbf{x}}_{t+1})_{vert} = [w_0 \tilde{\mathbf{x}}_{t+1}; w_1 \tilde{\mathbf{x}}_{t+1}; \dots; w_K \tilde{\mathbf{x}}_{t+1}]$).

The coefficients, \mathbf{c}_t , are thus updated in every iteration by minimizing the squared ℓ_2 norm

$$\hat{\mathbf{c}}_t = \arg \min_{\mathbf{c}_t} \|(\tilde{\mathbf{x}}_{t+1})_{vert} - \tilde{\mathbf{F}}_{x_t}^K \mathbf{c}_t\|_2^2. \quad (8)$$

Note that $[\tilde{\mathbf{F}}_{x_t}^K \mathbf{c}_t] \in \mathbb{R}^{(K+1)p \times 1}$ produces a vector of estimates of \mathbf{x}_{t+1} computed from all different $K + 1$ past states. This way, the estimator in Equation (8) seeks the \mathbf{c}_t vector that best predicts \mathbf{x}_{t+1} considering all $K + 1$ lookaheads.

One additional modification we make (compared to the original learning of dLDS as presented by Morduk et al. (2024)) is that rather than updating each \mathbf{f}_j using gradient descent, we infer the dLDS' basis dynamics operators $\{\mathbf{f}_j\}_{j=1}^J$ by fully and directly minimizing the cost. Specifically, let $\mathbf{F}_{all} := [1]_{1 \times J} = [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_J] \in \mathbb{R}^{p \times pJ}$, be the concatenated matrix of all $\{\mathbf{f}_j\}_{j=1}^J$.

Also, let $(\mathbf{x}_c)_t := ([1]_{J \times 1} \otimes \mathbf{x}_t) \circ (\mathbf{c}_t \otimes [1]_{p \times 1}) \in \mathbb{R}^{pJ \times 1}$, where \otimes denoted the Kronecker product and \circ denotes element-wise multiplication, and let $\mathbf{X}_c \in \mathbb{R}^{pJ \times T}$ be the horizontal concatenation of all $(\mathbf{x}_c)_t$ (for $t = 0 \dots T - 1$).

With these definitions, the dLDS operators $\{\mathbf{f}_j\}_{j=1}^J$ are updated by

$$\hat{\mathbf{F}}_{all} = \arg \min_{\mathbf{F}_{all}} \|\tilde{\mathbf{X}}_{:,2} - \mathbf{F}_{all}(\mathbf{X}_c)\|_F^2, \quad (9)$$

with each $\{\mathbf{f}_j\}_{j=1}^J$ being extracted from $\hat{\mathbf{F}}_{all}$ and normalized to a Frobenius norm of 1.

3.4 LINOCS for regularized Linear Time-Varying Systems (LTV)

Finally, we focus on the more general case of regularized linear time varying systems that are not necessarily switched or decomposed. Particularly, we focus on two types of regularizations, 1) that the operators change smoothly over time, i.e., $\|\mathbf{A}_t - \mathbf{A}_{t-1}\|_F^2$ is small, and 2) that the operators are sparse, i.e., $\|\mathbf{A}_t\|_0$ is small.

For the LTV case, we apply LINOCS to find the time-changing operators $\{\mathbf{A}_t\}_{t=1}^T$ by iteratively integrating multi-step reconstruction with the appropriate regularization. The operators are initialized with a regularized 1-step optimization

$$\hat{\mathbf{A}}_t = \arg \min_{\mathbf{A}_t} \|\tilde{\mathbf{x}}_t - \mathbf{A}_t \tilde{\mathbf{x}}_{t-1}\|_F^2 + \mathcal{R}(\mathbf{A}_t).$$

We integrate LINOCS into the estimation process by iteratively updating the operator estimates one at a time. Specifically, during each round of updates, we loop over every time point $t = 0 \dots T$ and, holding all other operators $\{\mathbf{A}_\tau\}_{\tau \neq t}$ fixed at their former estimates, and update \mathbf{A}_t by

$$\hat{\mathbf{A}}_t = \arg \min_{\mathbf{A}_t} \sum_{k=0}^K \left[w_k \sum_{k_i=1}^k \|\tilde{\mathbf{x}}_{t-k_i+k} - \hat{\mathbf{A}}_{t-k_i+k-1} \hat{\mathbf{A}}_{t-k_i+k-2} \cdots \mathbf{A}_t \cdots \hat{\mathbf{A}}_{t-k_i+1} \hat{\mathbf{A}}_{t-k_i} \tilde{\mathbf{x}}_{t-k_i}\|_2^2 \right] + \mathcal{R}(\mathbf{A}_t)$$

where $k = 0 \dots K$ denotes the order of the reconstruction and $t - k_i$ denotes the starting point of the reconstruction. The weights w^k are set as in all other models.

4 Results

To showcase LINOCS' ability to capture the dynamics in multiple models, we applied LINOCS to the above systems under diverse settings. The hyper-parameters used in each experiment are summarized in Section A.2.

4.1 LINOCS more accurately identifies ground truth linear systems under noisy observations

We first test LINOCS' ability to robustly learn time-invariant linear dynamical systems from noisy observations. We then simulate the dynamics $\mathbf{A} \in \mathbb{R}^{2 \times 2}$ as a rotational transition operator and a random offset $\mathbf{b} \in \mathbb{R}^{2 \times 1}$, where each $\mathbf{b}_i \sim \text{Uniform}(0, 1)$. We build the synthetic state $\mathbf{x}_t \in \mathbb{R}^{2 \times 1}$ as $\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{b}$ starting from random initial value $\mathbf{x}_t \in \text{Uniform}(0, 1)^{2 \times 1}$, such that the noisy observations are $\tilde{\mathbf{x}} = \mathbf{x} + \eta$, where $\eta \sim \mathcal{N}(0, 0.3^2)$ (Fig. 3A,B).

We compare the learned operators using LINOCS with four baselines. First we compare to traditional 1-step optimization (Eqn. (3)). We further compare the linear LINOCS to our implementations of the conceptual framework presented in DAD (Venkatraman et al., 2015), as it is the approach closest to LINOCS in terms of integrating multi-step predictions into model training. Our implementation of DAD integrates expert and non-expert demonstrations for model training, inspired by the Dataset-Aggregation (Dagger) approach (Ross et al., 2011). Specifically we test three implementations of DAD. For each implementation we initialized the transition matrix (\mathbf{A}_{init}) and the offset (\mathbf{b}_{init}) using the optimal estimate from 1-step optimization. We then train the model through 100 iterations where at each iteration, we used our latest estimates of \mathbf{A} and \mathbf{b} to perform full lookahead reconstruction, starting from time $t = 0$. We then update our estimates of \mathbf{A} and \mathbf{b} using the optimal 1-step optimization, and tested all three options outlined below:

- *DAD with full model update:*
At each iteration, we update \mathbf{A} and \mathbf{b} based on the lookahead reconstruction of the state ($\hat{\mathbf{x}}_t$) calculated based on the last operators estimate. Namely, $\{\hat{\mathbf{A}}_{iter+1}, \hat{\mathbf{b}}_{iter+1}\} = \arg \min_{\{\mathbf{A}, \mathbf{b}\}} \frac{1}{T} \sum_{t=0}^{T-1} \|\tilde{\mathbf{x}}_{t+1} - (\mathbf{A}\hat{\mathbf{x}}_t + \mathbf{b})\|_2^2$.
- *Dagger-inspired DAD (reweighted DAD):*
For reweighted DAD, we estimate \mathbf{A} and \mathbf{b} at each iteration using both the observations and the lookahead reconstruction from the last estimates of \mathbf{A} and \mathbf{b} . In particular, let $[\hat{\mathbf{x}}_t, \tilde{\mathbf{x}}_t] \in \mathbb{R}^{p \times 2}$ be a horizontal concatenation of the lookahead reconstruction and of the observations at time t . Then we iteratively solve: $\{\hat{\mathbf{A}}_{iter+1}, \hat{\mathbf{b}}_{iter+1}\} = \arg \min_{\{\mathbf{A}, \mathbf{b}\}} \frac{1}{T} \sum_{t=0}^{T-1} \|[\hat{\mathbf{x}}_{t+1}, \tilde{\mathbf{x}}_{t+1}] - (\mathbf{A}[\hat{\mathbf{x}}_t, \tilde{\mathbf{x}}_t] + \mathbf{b})\|_2^2$.
- *Dagger with ℓ_2 constraint (reweighted DAD with ℓ_2):*
This option is solved similarly to the reweighted DAD, with the addition of the Frobenius norm ($\|\cdot\|_F$) on the operators (\mathbf{A}) and on (\mathbf{b}) during training.

We find that LINOCS identifies operators that yield accurate dynamics in long-time scale predictions (Fig. 3D). The other methods we tested, including 1-step optimization (Fig. 3C,E cyan) and DAD-based implementations (Fig. 3D, reds), instead decay to zero (away from the real system), indicating a less accurate estimation of the dynamics. The improved accuracy of the operators identified by LINOCS (Fig. 3E)

becomes apparent when examining the effects of the operators’ estimation errors (Fig. 3C). These errors are larger for the other methods, showcasing that those methods accumulate more errors over shorter time spans than LINOCS’ does.

We next investigated the effect of the training order in LINOCS on long-term reconstruction. We trained LINOCS on the noisy observations with increasing training orders (e.g. 5, 10, 30, 80) and then tested the performance under multiple prediction orders (Fig. 3G). Compared to the baselines tested, LINOCS exhibits increased performance even with very low training orders (e.g., 5), with higher orders resulting in almost perfect reconstruction (Fig. 3G bottom-right subplot). Additionally, exploring LINOCS’s robustness to noise reveals that, unlike one-step reconstruction, LINOCS is robust even under very high levels of noise (Fig. 3G, H, blue). The resulting MSE compared to the ground truth dynamics is much lower in LINOCS, even under very high σ noise levels (Fig. 3H blue vs. orange-red).

When examining the duration for which LINOCS remains robust without converging, we observe that our approach accurately predicts approximately 35,000 time points into the future before deviating from the real system and decaying to 0—demonstrating stability over exceptionally long time scales (Fig. 3J).

Additionally, when examining LINOCS’ long-term prediction robustness to increasing noise levels introduced during training, we find it to be robust to even extreme Gaussian noise levels, $\sigma = 0.9$ (Fig. 3D, blue), in contrast to 1-step optimization (Fig. 3 D, orange-red).

We further tested LINOCS on linear systems with structured noise (Fig. 12) as well as on a simulation of 3-dimensional cylinder (Fig. 13), yielding similar results. For structured noise, we modeled the observation as $\tilde{\mathbf{x}} = \mathbf{x} + \sigma \sin(\gamma t)$ with $\sigma = 0.5$ and $\gamma = 3$ for $t = 1 \dots 501$. Unlike other methods, LINOCS found operators that led to accurate long-term predictions (Fig. 12 D). Moreover, when examined under increasing training and prediction orders, we found that LINOCS is robust for long-term predictions, even for full lookahead reconstructions ($k_{\text{pred}} = 501$, Fig. 12 E,H). When evaluating its robustness to increasing noise levels (σ), we found that even for very high noise levels ($\sigma = 0.9$), LINOCS achieved much more robust results than 1-step optimization (Fig. 12 F,G). Additionally, when exploring how far into the future it enables robust reconstruction before converging, we found that it is capable of full lookahead for approximately 70,000 time points—a testament to its ability to find more robust operators that can independently describe the system (Fig. 12 I).

For the 3D cylinder case (Fig. 13), with Gaussian noise ($\sigma = 0.4$), we similarly demonstrate that LINOCS recovers more accurate operators, leading to significantly more robust long-term predictions and enabling full recovery of the process (Fig. 13 A,B,C,D), both under increasing prediction orders (Fig. 13 E, G) and noise (σ) levels (Fig. 13 H), as well as exhibiting an impressive ability to reconstruct lookahead predictions (starting from \mathbf{x}_0) for very long periods (approximately 70,000 time points) before converging to similar error as 1-step does (Fig. 13 I). In contrast, 1-step optimization yields high-error within a few IMS prediction orders.

4.2 LINOCS identifies accurate interactions in switching systems

We next tested LINOCS-driven SLDS as detailed in Section 3.2 on simulated data comprising of $J = 3$ discrete states. The transition operators for each of the distinct states was set to a 3×3 rotational matrix oriented in a different direction. Additionally, the offset for each state ($\mathbf{b}_j \in \mathbb{R}^{3 \times 1}$) was set to a random vector drawn from a uniform distribution between 0 and 1 (Fig. 16).

Notably, since the method is invariant to the order of the operators, to compare the identified operators to the ground truth operators, we sorted the operators using the “linear sum assignment” problem (SciPy’s implementation, by Crouse (2016)), with the cost function being the Frobenius norm between each pair of \mathbf{f} s (ground truth vs. estimated for each model). As baselines, we compare the results of LINOCS-augmented SLDS with standard SLDS and recurrent SLDS (Linderman et al., 2016) with varying numbers of iterations.

When comparing LINOCS-SLDS to the baselines (Fig. 4), LINOCS consistently outperformed the other approaches across multiple metrics including operator recovery (Fig. 4C,E), switching times recovery (Fig. 4A,D), and dynamics reconstruction (Fig. 4B, Fig. 14B). In particular, LINOCS-SLDS accurately identified switching times, whereas classical SLDS and rSLDS tended to introduce additional redundant switches (Fig. 4A). Moreover, the discrepancies between the ground truth operators and those identified by




Figure 4: **Results on switching systems.** **A:** Active discrete states for LINOCS (blue) compared to baselines, including SLDS and rSLDS with 10 or 100 training epochs. **B:** Correlation between the ground truth dynamics (\mathbf{x}) and the full-lookahead reconstructed dynamics ($\hat{\mathbf{x}}$). **C:** Correlation between the ground truth operators (\mathbf{f}) and the identified operators ($\hat{\mathbf{f}}$). **D:** Correlation between the ground truth coefficients (\mathbf{c}) and the identified coefficients ($\hat{\mathbf{c}}$). **E:** Difference between the ground-truth sub-dynamics ($\hat{\mathbf{F}}$) and reconstructed basis dynamics by different models. LINOCS was able to achieve sub-dynamics that are much closer to the ground truth than the other baselines.

LINOCS (Fig. 4E, right-most four columns) were substantially smaller than the differences observed with classical SLDS/rSLDS (Fig. 4 E, left columns), as evidenced by the higher correlations between LINOCS’ operators and the ground truth (Fig. 4 C). Furthermore, when examining the eigenvalues of the identified operators compared to the ground truth (Fig. 15), the eigenspectrum derived from the LINOCS-driven solver closely resembled the ground truth eigenspectrum more than the classical SLDS and rSLDS cases, highlighting the effectiveness of LINOCS in capturing the underlying dynamics.

4.3 LINOCS finds dLDS operators that yield accurate dLDS lookahead predictions

Next, we applied LINOCS to dLDS, as described in Section 3.3. First we generated ground-truth data that represent a “pseudo-switching” (Fig. 2) process—i.e. linear dynamics that switch more smoothly (in our case between $J = 3$ systems) compared to SLDS where operators switch abruptly. This creates overlap periods where two dynamical systems are active at once as they trade off (Fig. 17). LINOCS-dLDS demonstrated significantly improved stability in full lookahead reconstruction compared to single-step dLDS (Fig. 5). Notably, training with orders approximately greater than 35 ($K_{\text{train}} > 35$) on our synthetic dataset (containing 1000 time points) resulted in high-quality full reconstruction (Fig. 5A). Additionally, when comparing MSE and correlation of the time-evolving operator $\mathbf{F}_t = \sum_{j=1}^J c_{jt} \mathbf{f}_j$ to the ground truth, we observed a monotonic decrease in MSE with increasing maximal LINOCS training orders (K_{train}), while the correlation showed a monotonic increase (Fig. 5B). Moreover, when comparing the largest eigenvalue of \mathbf{F}_t over time between the ground truth operator and the identified operators, the high-order LINOCS-dLDS achieved a structure much closer to the ground truth (Fig. 18).

Interestingly, although the 1-step reconstruction exhibited good performance even for non-LINOCS or low-order LINOCS-dLDS (Fig. 5D left), the importance of LINOCS became more apparent in full lookahead reconstruction (Fig. 5D right). This underscores the necessity of multi-step reconstructions for accurately estimating dynamical system dynamics, where errors might be obscured when only assessing the single-step reconstructions. It also highlights the importance of integrating multiple orders simultaneously during training.

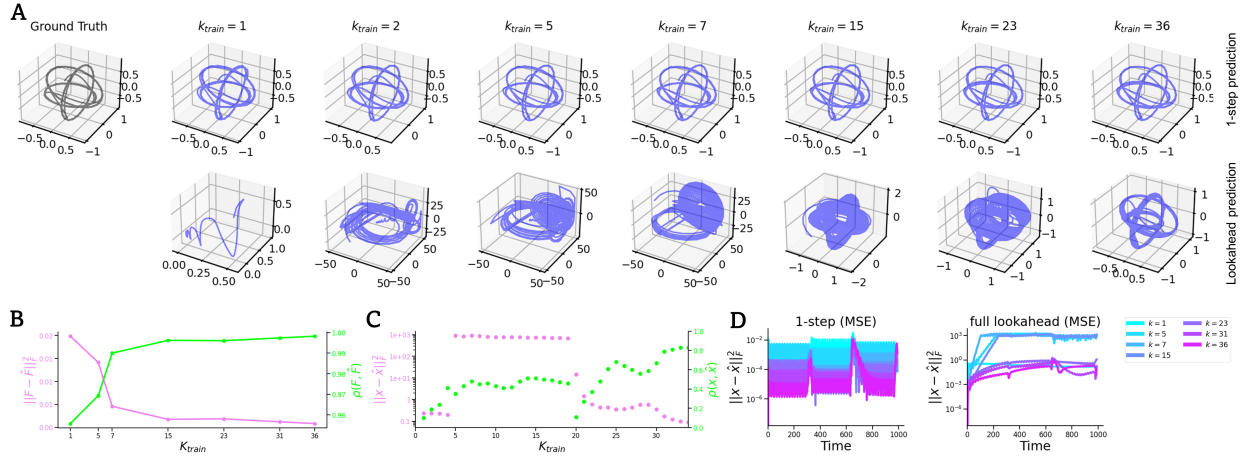


Figure 5: **Decomposed linear dynamical systems results.** **A:** Ground truth dynamics compared to 1-step (top) and full lookahead (bottom) reconstructions for non-LINOCs dLDS ($K_{train} = 1$) and LINOCs-dLDS with different training orders. **B:** MSE (pink) and correlation (green) between ground truth operators and the operators identified by LINOCs under different orders. **C:** MSE (pink) and correlation (green) between ground truth dynamics and full lookahead reconstructions using the different LINOCs training orders. **D:** Local MSE for 1-step (left) and for full lookahead reconstruction (right) over the time points of the dynamics.

We further extended our study to encompass more nuanced dLDS scenarios, exhibiting prolonged time scales and recurring patterns of identical active operators across distinct intervals (Fig. 19). We found analogous enhancements of LINOCs over the traditional 1-step dLDS implementation. Specifically, LINOCs demonstrates robust accurate long-term predictions, including full lookahead prediction (Fig. 6B), in contrast to 1-step optimization, which yield high lookahead error (Fig. 6C, last three subplots). Furthermore, also for this more complex example, upon comparing the identified time-varying transition operators $F_t = \sum_j^J c_{jt} f_j$ to the ground truth, LINOCs revealed operators with eigenvalues significantly more correlated with the real operators’ evaluations (Fig. 6D, E) compared to the 1-step optimization results. Additionally, when comparing the operators themselves against the ground truth, those identified by LINOCs exhibited higher correlation and smaller MSE with the ground truth compared to these identified by 1-step dLDS (Fig. 6F,G,H).

4.4 LINOCs finds interactions that yield robust lookahead predictions in time-varying systems

To test the applicability of LINOCs to more general LTV systems, we implemented LINOCs-LTV to capture the chaotic behavior of the Lorenz attractor (Sec. A.3) through a smoothly changing LTV approximation (Fig. 7). We compared LINOCs-LTV with several other LTV solvers with varying constraints, including smoothness and sparsity. Unlike methods relying on 1-step optimization, LINOCs, despite similar regularization constraints, achieved superior full lookahead reconstruction (Fig. 7A bottom).

Also here, while different methods performed satisfactorily in the 1-step (post-training) prediction (Fig. 7A top, B red, C red), disparities emerged in higher-orders lookahead predictions where alternative methods failed. While all methods, including LINOCs, achieved commendable 1-step reconstruction, LINOCs demonstrated a markedly lower full lookahead error (Fig. 7B green, 5 most right bar pairs) and superior correlation with the ground truth (Fig. 7C green, five most right bar pairs).

In addition, we analyzed operators identified across various training iterations of LINOCs to assess their proficiency in achieving lookahead reconstruction (Fig. 20). For this analysis, we used the Lorenz attractor with 900 time points with intervals of 0.1/9 arbitrary units (a.u.), and applied a smoothness constraint with a weight of $\lambda = 0.1$. We observed that over training iterations, LINOCs adaptively influenced the predicted lookahead dynamics to gradually converge towards the ground truth dynamics (Fig. 20 C), with a monotonic decrease MSE (Fig. 20A, B).

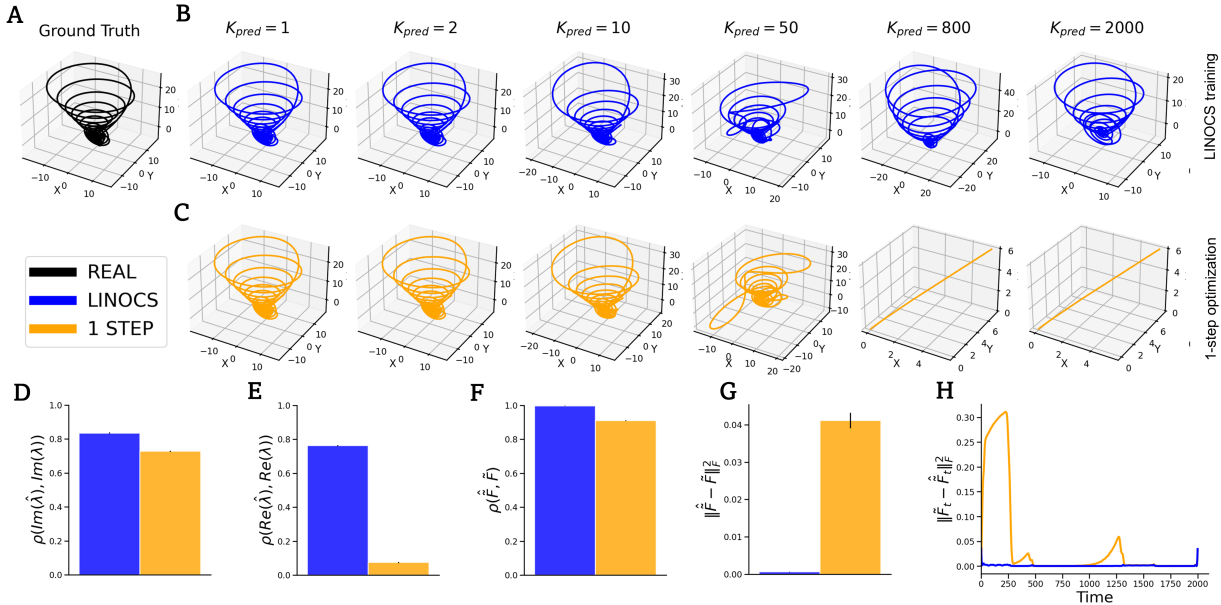


Figure 6: **Additional decomposed linear dynamical systems results.** **A:** Ground truth dynamics compared to **(B)** LINOCS results and **(C)** 1-step optimization results, for increasing prediction orders. **D,E:** Correlation between the eigenvalues of the ground truth transition matrix ($\mathbf{F}_t = \sum_j^J c_{jt} \mathbf{f}_j$) and the eigenvalues of the one identified by LINOCS. **D:** imaginary part; **E:** real part. Results display the average correlation over time. Eigenvalues were matched using the “linear sum assignment problem” (Scipy’s Crouse (2016)). **F,G:** Comparing the identified time-changing transition matrix ($\hat{\mathbf{F}}_t$) identified by LINOCS vs. 1-step optimization in terms of correlation (**F**) and MSE (**G**). **H:** Comparing the MSE of the identified $\hat{\mathbf{F}}_t$ over time.

When analyzing which time points of the dynamics contributed to higher MSEs in the full post-training lookahead prediction, we noticed that early training iterations tended to produce higher full lookahead prediction errors at later time points of the dynamics (Fig. 20B, top right). However, over subsequent iterations, the effect of LINOCS managed to mitigate the accumulation of errors at these late time points (Fig. 20B, bottom right).

4.5 LINOCS finds robust interactions in real-world neural data

Finally, we applied LINOCS to real-world dataset described by Kyzar et al. (2024), which consists of high density electrode array of populations of single units in the human medial temporal and medial frontal lobes while subjects were engaged in a screening task. We applied linear LINOCS, SLDS and dLDS-LINOCS, and LTV-LINOCS to a single recording session that includes recordings from five brain areas (amygdala left and right, cingulate cortex, hippocampus, pre-supplementary motor area). All the dynamical systems models were trained on the firing rate data, which we inferred from the spike-sorted electrophysiology via a Gaussian kernel convolution.

We investigated several LINOCS models to showcase their distinct characteristics. First, we examined the linear case for each brain area individually and explored the mean field interactions between areas (Fig. 8). Importantly, while typical real-world brain dynamics are assumed to be non-linear and non-stationary, our aim in starting with the linear model was to demonstrate how LINOCS can identify the fundamental background neural interactions under linear assumptions and check how its identified interactions defer from these identified by the 1-step approach. We first applied the linear LINOCS on the firing rate activity from all neurons within each region to identify between-region interactions. We observed that LINOCS identified different linear interactions within areas compared to the 1-step approximation. Drawing from our

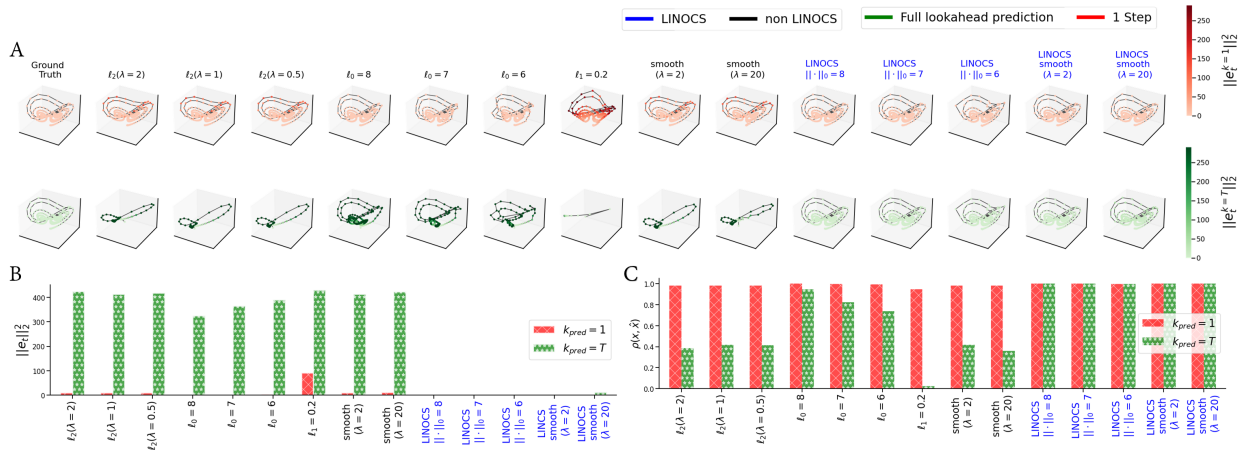


Figure 7: **LTV approximation of the Lorenz attractor.** **A:** 1-step post-training prediction (top, pink-red) vs. full lookahead prediction (bottom, green) for different baselines, including LINOCS with various smoothness levels or ℓ_0 regularization. Color indicates the local MSE. **B:** Squared ℓ_2 of the error between the ground truth and 1-step prediction (red) vs. full lookahead predictions (starry green), for the different methods. **C:** Correlations between the ground truth and 1-step prediction (red) vs. full lookahead predictions (starry green), for the different methods.

conclusions based on synthetic data linear results, this suggests that LINOCS may better capture the linear approximation of brain activity than the common 1-step optimization (Fig. 8 A).

Then, we also applied the linear LINOCS on the mean activity of each region, and found that when examining the full lookahead reconstructions (Fig. 8 C) the 1-step optimization, in contrast to LINOCS, decayed to zero activity due to small accumulated deviations in operator values. In contrast, LINOCS managed to maintain activity closer to the average values of the dynamics. However, due to linear enforcement, neither approach could capture fluctuations in dynamics. Moreover, the full lookahead reconstruction error for LINOCS-linear was overall much smaller compared to the classical 1-step (Fig. 8B).

We next applied LINOCS-SLDS with three discrete states and compared it with regular SLDS using the same number of iterations. LINOCS identified operators that exhibit slight differences compared to those found by classical SLDS (Fig. 9B vs Fig. 21; Fig. 22 A vs. B) as well as slightly different switching patterns between the two approaches (Fig. 22 C,D). These operators resulted in significantly more robust lookahead predictions. Specifically, differences are evident in both connection presence, weights, and distribution among global operators. For example, in the ‘‘Amygdala left’’ region, both classical SLDS and LINOCS-driven SLDS identify a connection from neuron 10 to neuron 2 as part of \mathbf{f}_3 , albeit with varying weights. Additionally, both methods identify connections from 5 to 3 (in \mathbf{f}_1 for classical SLDS and in \mathbf{f}_2 for LINOCS-SLDS) as well as from 6 to 3 (in \mathbf{f}_2 for classical SLDS and in \mathbf{f}_1 for LINOCS-SLDS), but with differing weights. Similar discrepancies are observed in other regions. Furthermore, LINOCS-SLDS and classical-SLDS each identify connections that the other overlooks; for instance, in the ‘‘Amygdala right’’ region, LINOCS-SLDS identifies a strong connection from 11 to 6, whereas classical SLDS does not. Conversely, classical SLDS identifies a connection from 3 to 6 (in \mathbf{f}_3), which LINOCS-SLDS does not recognize.

Importantly, the operators found by LINOCS enable full lookahead reconstruction without diverging, in contrast to regular classical SLDS that diverge to extreme values in full lookahead prediction (Fig. 9C). Moreover, the reconstruction error for the full lookahead prediction was overall much smaller for LINOCS-SLDS compared to the classical SLDS (Fig. 9A). These observations suggest that if the real neural process follows switched dynamics, LINOCS may capture the underlying dynamics more effectively, as inferred from our analysis of the synthetic case.

We observed similar patterns using dLDS-LINOCS, which revealed underlying global brain interactions potentially fundamental to brain function (Fig. 10 A). When examining their dynamic activations (\mathbf{c}_t), we

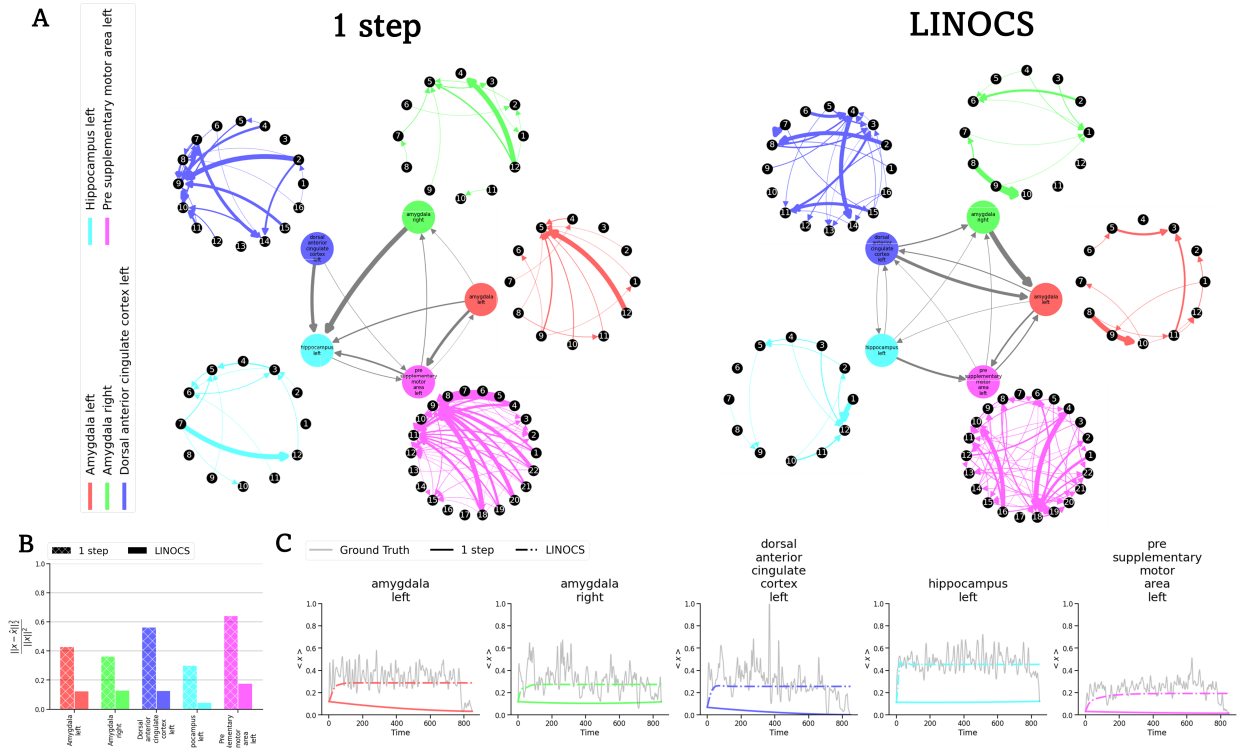


Figure 8: **Application of linear-LINOCS to multi-region neural recordings.** **A:** The per-region and between-region mean field linear dynamics operators identified by 1-step linear optimization vs. LINOCS with a linear time-invariant system. Each within-region network describes the linear operator $\hat{\mathbf{A}}_{region}$ derived by applying LINOCS to the firing rate matrix constrained to include only the neurons from that region. **B:** Full-lookahead reconstruction error for 1-step linear optimization vs. LINOCS-linear approach. **C:** Ground truth mean activity per region compared to the lookahead prediction trace of the mean field activity by 1-step linear optimization vs. LINOCS-linear.

noted a “background” interaction consistently active, with slight modulations over time (Fig. 10 B, brown), alongside gradually changing activities of other interactions (Fig. 10 B, gray-blue-purple). Importantly, these results provided lookahead predictions that did not decay and maintained a high correlation with the observations (Fig. 10 C).

Finally, employed the LTV-LINOCS on all neurons from all regions simultaneously while imposing a smoothness constraint on consecutive operators (with regularization of $\lambda = 0.1$ on $\|\mathbf{A}_t - \mathbf{A}_{t-1}\|_2^2$, Fig. 11). Our findings reveal that LINOCS identifies operators capable of producing full lookahead reconstructions without divergence, closely approximating observed data. Comparative analysis against 1-step optimization with various smoothness levels (Fig. 11A,B,E) underscores the superiority of LINOCS in achieving reconstructions faithful to the data. Additionally, examination of error evolution over time suggests a monotonic increase in error for non-LINOCS approaches (Fig. 11C). Moreover, we observed notable discrepancies between the operators identified by LINOCS and the baselines (Fig. 11D). These results highlight the efficacy of LTV-LINOCS in capturing complex temporal dynamics in real world data while maintaining data fidelity. Overall, we showed that in all these real-world neural versions, LINOCS was able to recover more robust descriptions of the dynamic evolution for the long run, which, based on our synthetic results, may imply that these are closer to the real unknown interactions.

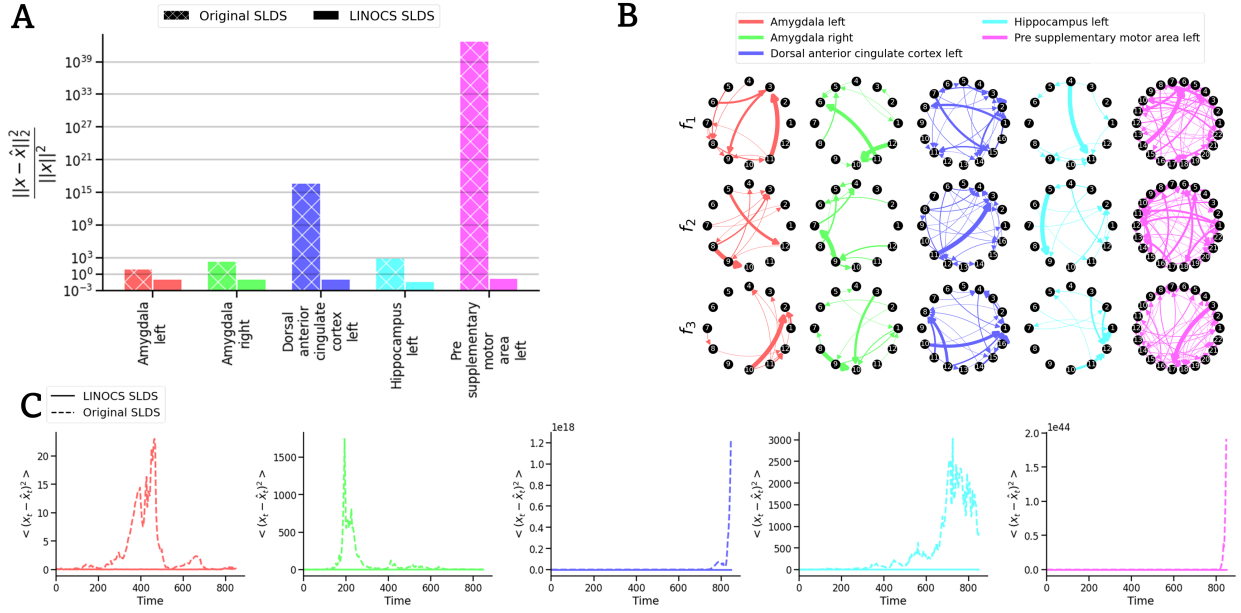


Figure 9: **SLDS results on real data.** **A:** Relative error of the full lookahead prediction compared to the ground truth, for both LINOCS-SLDS vs. classical SLDS, for each brain area. **B:** The networks identified by LINOCS-SLDS (see Fig. 21 for the networks identified by the classical-SLDS). **C:** Lookahead reconstruction using LINOCS-SLDS (full curve) vs. classical SLDS (dashed curve). Classical-SLDS diverge to extreme values in the full lookahead reconstruction.

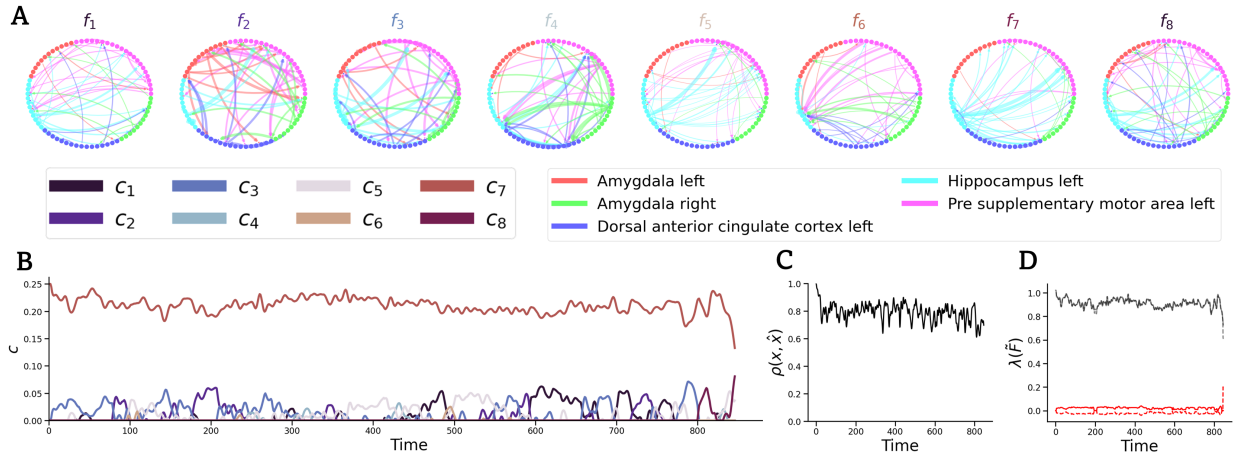


Figure 10: **dLDS-LINOCS results on the real neural data.** **A:** The identified operators $\{f_j\}_{j=1}^{J=8}$. **B:** The identified sparse coefficients c_t . **C:** Correlation between the full-lookahead reconstruction results and the observations. **D:** The two largest eigenvalues of the time-varying transition operator $\tilde{F}_t = \sum_{j=1}^8 c_{jt} f_j$. Black: real part. Red: imaginary part.

5 Discussion

In this paper we introduced LINOCS (Lookahead Inference of Networked Operators for Continuous Stability), a learning procedure to improve stability and accuracy of dynamical system inference that leverages lookahead estimation. By iteratively integrating re-weighted multi-step reconstructions with additional constraints on the operators, LINOCS enables robust inference of networked operators in dynamical systems, even in the presence of noise and nonlinearity.

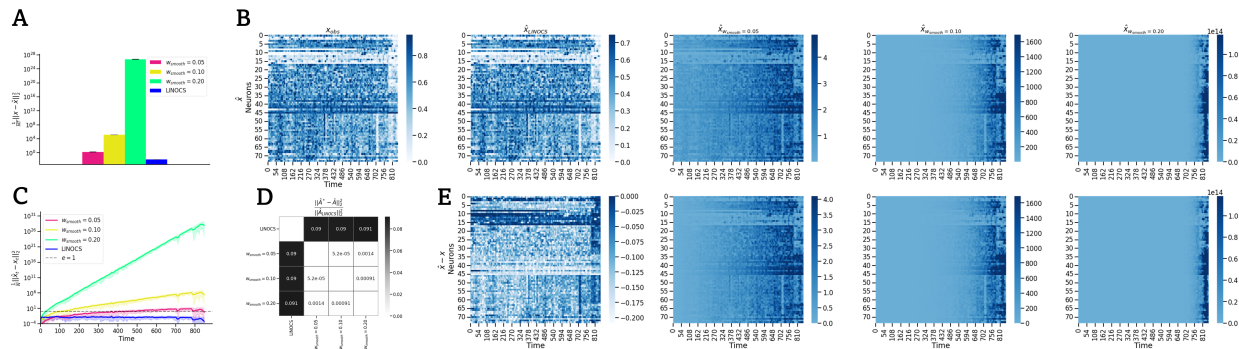


Figure 11: **LTV-LINOCs results on real neural data.** **A:** MSE between the ground truth data and the full lookahead reconstructions. **B:** The ground truth data (left) compared to the predictions produced by LINOCs-LTV (2nd left) and 1-step optimization with increasing smoothness constraints. **C:** MSE of the reconstruction over time points, compared to 1-step optimization with different smoothness levels. **D:** Frobenius norm of the differences between \mathbf{A} s identified by the different models, normalized by the magnitude of the operator identified by LINOCs. **E:** Difference between the observations and the full lookahead reconstruction for LINOCs-LTV vs. 1-step optimization with increasing smoothness levels.

Our experimental results highlight LINOCs’ effectiveness across various dynamical systems, including Linear Systems (LDSs), Switched Linear Systems (SLDS), decomposed LDSs (dLDS) (Mudrik et al., 2024), and Linear Time-Varying Systems (LTV) in both simulation and real-world neural data. LINOCs not only achieves more precise full lookahead reconstruction compared to baseline methods but also successfully retrieves ground truth operators in synthetic data, signifying its superior capacity to capture underlying systems. These findings suggest that LINOCs holds greater potential than alternative approaches for accurately identifying unknown hidden interactions also in real-world data, where the real underlying interactions are often obscured but pivotal for robust scientific interpretation.

Looking ahead, several promising avenues exist for future directions and extensions, including applying LINOCs to improve the reconstruction robustness of highly non-linear systems and integrate it to advance robust RNN training. Particularly, if incorporating LINOCs to deep networks, the integration of multi-step reconstructions into the networks’ training, may help mitigate issues such as vanishing or exploding gradients. Additionally, extending LINOCs to handle non-Gaussian noise could enhance its applicability to a wider range of real-world scenarios.

References

- Guy A. Ackerson and K. S. Fu. On State Estimation in Switching Environments. *IEEE Transactions on Automatic Control*, AC-15(1):10–17, 1970. ISSN 15582523. doi: 10.1109/TAC.1970.1099359.
- Yaakov Bar-Shalom and Xiao-Rong Li. *Estimation and tracking*. Artech House, Boston, MA, 1990.
- 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.
- Guillaume Chevillon. Direct multi-step estimation and forecasting. *Journal of Economic Surveys*, 21(4):746–785, 2007.
- David R Cox. Prediction by exponentially weighted moving averages and related methods. *Journal of the Royal Statistical Society: Series B (Methodological)*, 23(2):414–422, 1961.
- David F Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016.
- Raina D’Aleo, Adam Rouse, Marc Schieber, and Sridevi V Sarma. Quantifying interactions between neural populations during behavior using dynamical systems models. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1960–1964. IEEE, 2019.
- David F Findley. On the use of multiple models for multi-period forecasting. In *Proceedings of Business and Economic Statistics, American Statistical Association*, pages 528–531, 1983.
- David F Findley. Model selection for multi-step-ahead forecasting. *IFAC Proceedings Volumes*, 18(5):1039–1044, 1985.
- Emily Fox, Erik Sudderth, Michael Jordan, and Alan Willsky. Nonparametric bayesian learning of switching linear dynamical systems. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2008/file/950a4152c2b4aa3ad78bdd6b366cc179-Paper.pdf>.
- Zoubin Ghahramani and Geoffrey E. Hinton. Switching state-space models. Technical report, King’s College Road, Toronto M5S 3H5, 1996.
- James D. Hamilton. Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45(1-2):39–70, 1990. ISSN 03044076. doi: 10.1016/0304-4076(90)90093-9.
- Kadierdan Kaheman, J Nathan Kutz, and Steven L Brunton. Sindy-pi: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A*, 476(2242):20200279, 2020.
- In-Bong Kang. Multi-period forecasting using different models for different horizons: an application to us economic time series data. *International Journal of Forecasting*, 19(3):387–400, 2003.
- Hassan K Khalil. *Nonlinear systems*, 2001.
- Lawrence R Klein. An essay on the accuracy of economic prediction. In *Paul Samuelson and the Foundations of Modern Economics*, pages 29–69. Routledge, 2019.
- Michael Kyzar, Jan Kamiński, Aneta Brzezicka, Chrystal M Reed, Jeffrey M Chung, Adam N Mamelak, and Ueli Rutishauser. Dataset of human-single neuron activity during a sternberg working memory task. *Scientific Data*, 11(1):89, 2024.
- Jin-Lung Lin and Clive WJ Granger. Forecasting from non-linear models in practice. *Journal of Forecasting*, 13(1):1–9, 1994.

- Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial Intelligence and Statistics*, pages 914–922. PMLR, 2017.
- Scott Linderman, Benjamin Antin, David Zoltowski, and Joshua Glaser. SSM: Bayesian Learning and Inference for State Space Models. <https://github.com/lindermanlab/ssm>, 2020.
- Scott W Linderman, Andrew C Miller, Ryan P Adams, David M Blei, Liam Paninski, and Matthew J Johnson. Recurrent switching linear dynamical systems. *arXiv preprint arXiv:1610.08466*, 2016.
- Massimiliano Marcellino, James H Stock, and Mark W Watson. A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of econometrics*, 135(1-2):499–526, 2006.
- Noga Mudrik, Gal Mishne, and Adam S Charles. Sibblings: Similarity-driven building-block inference using graphs across states. *arXiv preprint arXiv:2306.04817*, 2023.
- Noga Mudrik, Yenho Chen, Eva Yezerets, Christopher J Rozell, and Adam S Charles. Decomposed linear dynamical systems (dlDs) for learning the latent components of neural dynamics. *Journal of Machine Learning Research*, 25(59):1–44, 2024.
- K. P. Murphy. Switching Kalman filters. Technical report, Technical Report, UC Berkeley, 1998.
- Josue Nassar, Scott W Linderman, Monica Bugallo, and Il Memming Park. Tree-structured recurrent switching linear dynamical systems for multi-scale modeling. *arXiv preprint arXiv:1811.12386*, 2018.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Shankar Sastry. *Nonlinear systems: analysis, stability, and control*, volume 10. Springer Science & Business Media, 2013.
- John Savill, Ian Dransfield, Chris Gregory, and Chris Haslett. A blast from the past: clearance of apoptotic cells regulates immune responses. *Nature Reviews Immunology*, 2(12):965–975, 2002.
- Richard R Stein, Vanni Bucci, Nora C Toussaint, Charlie G Buffie, Gunnar Rättsch, Eric G Pamer, Chris Sander, and Joao B Xavier. Ecological modeling from time-series inference: insight into dynamics and stability of intestinal microbiota. *PLoS computational biology*, 9(12):e1003388, 2013.
- Rahimi Tabar. *Analysis and data-based reconstruction of complex nonlinear dynamical systems*, volume 730. Springer, 2019.
- George C Tiao and Daming Xu. Robustness of maximum likelihood estimates for multi-step predictions: the exponential smoothing case. *Biometrika*, 80(3):623–641, 1993.
- Vinit S Unni, Ashish Mittal, Preethi Jyothi, and Sunita Sarawagi. Improving rnn-transducers with acoustic lookahead. *arXiv preprint arXiv:2307.05006*, 2023.
- E Vaadia, I Haalman, M Abeles, Hagit Bergman, Y Prut, Hi Slovin, and AMHJ Aertsen. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514):515–518, 1995.
- Arun Venkatraman, Martial Hebert, and J Bagnell. Improving multi-step prediction of learned time series models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Andrew A Weiss. Multi-step estimation and forecasting in dynamic models. *Journal of Econometrics*, 48(1-2):135–149, 1991.
- Yongqian Xiao, Zixin Tang, Xin Xu, Xinglong Zhang, and Yifei Shi. A deep koopman operator-based modelling approach for long-term prediction of dynamics with pixel-level measurements. *CAAI Transactions on Intelligence Technology*, 2023.

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. IEEE, 2019.

A Appendix

A.1 calculation details for operators differences (Fig. 3C)

We computed the operator differences illustrated in Figure 3C using the expression:

$$((\widehat{\mathbf{A}} - \mathbf{I})\mathbf{x} - (\mathbf{A} - \mathbf{I})\mathbf{x}) \times \text{factor}$$

Here, $\widehat{\mathbf{A}}$ represents the operators identified by the different methods, \mathbf{A} denotes the ground-truth operators, and "factor" is a scalar used for visualization purposes only (identical for all methods).

A.2 Hyperparamters used in experiments

Table 1: Hyperparameter settings for DAD baseline in linear experiment

Parameter	Value	Additional Info
seed	0	random seed
T	500	number of time points
w_{ℓ_2}	0	weight of ℓ_2 regularization on dynamics
w_{decay}	1	decay of regularization coefficient over iterations
$N_{iterations}$	100	number of iterations
A_init_type	step	initialize A with 1-step optimization
reweight	False	whether to reweight the observations and the lookahead during training.

Table 2: Hyperparameter settings for "DAD reweigh" baseline in linear experiment

Parameter	Value	Additional Info
seed	0	random seed
T	500	number of time points
w_{ℓ_2}	0	weight of ℓ_2 regularization on dynamics
w_{decay}	1	decay of regularization coefficient over iterations
$N_{iterations}$	100	number of iterations
A_init_type	'step'	initialization type for matrix A
reweight	True	whether to reweight the observations and the lookahead during training.

Table 3: Hyperparameter settings for "reweigh ℓ_2 " baseline in linear experiment

Parameter	Value	Additional Info
seed	0	random seed
T	500	number of time points
w_{ℓ_2}	1	weight of ℓ_2 regularization on dynamics
w_{decay}	1	decay of regularization coefficient over iterations
$N_{iterations}$	100	number of iterations
A_init_type	'step'	initialization type for matrix A
reweight	True	whether to reweight the observations and the lookahead during training.

Table 4: Hyperparameter settings for LINOCS in linear experiment

Parameter	Value	Additional Info
K	80	maximum training order
with offset	True	whether to look for an offset term \mathbf{b}
cal_offset	True	whether to calculate offset
weights_style	exponential	weight style for orders
σ_w	exponential parameter for the weights	0.01
infer_b_way	'each'	approach to infer the onset. Based on each order.
K_b	20	maximum lookahead training order for the offsets \mathbf{b}
weights_style_b	'exponential'	weights style for the offsets \mathbf{b}

Table 5: Hyperparameter settings for dLDS experiment

Parameter	Value	Additional Info
K	50	maximum lookahead order
additional_update	True	whether to include an additional update step
ℓ_1 _init	0	value of ℓ_1 on the coefficients for the 1-st iteration
max_iters	200	maximum number of iterations
$\ F\ _2$	0	ℓ_2 norm on the basis of dLDS dynamics
$freq_{update_F}$	5	frequency of updating $\{\mathbf{f}_j\}_{j=1}^J$
ℓ_2 decay	0.99	decay of the ℓ_2 norm on the coefficients
ℓ_1 decay	0.9999,	decay of the ℓ_1 norm on the coefficients
$w_{smoothdecay}$	1.01,	decay of smoothing weight on coefficients
l_smooth_time	1.1,	regularization weight on coefficients smoothing \mathbf{c}_t
ℓ_2	0,	ℓ_2 on the coefficients
$w_{smooth_{time}}$	0.1,	smoothness on \mathbf{c}_t over time
σ_{noisy_c}	0.05	std of the noise to add to coefficients during training
ℓ_1	2.5,	ℓ_1 regularization on \mathbf{c}_t
decor	False	wether to decorrelate the basis dynamics
max_interval_k	1	maximum interval to increase k during training
to_norm_F	True	whether to normalize the basis dynamics
\mathbf{x}_0	[0.2160895, 0.97627445, 0.00623026]	ground truth initial state (at $t = 0$)
J	3	number of basis dynamics operators (\mathbf{f}_s)

A.3 Lorenz equations

The Lorenz attractor follows:

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x), \\ \frac{dy}{dt} &= x(\rho - z) - y, \\ \frac{dz}{dt} &= xy - \beta z,\end{aligned}$$

Table 6: Hyperparameter setting for the LTV-Lorenz experiment

Parameter	Value	Additional Info
K	5	maximum lookahead order
w_{smooth}	2 (or 20)	smoothness weight
$estimate_{thres}$	2	threshold to increase k value
<code>with_future_cost_or</code>	True	consider both fast and future reconstructions
$with_{future}$	True	wether to apply time smoothness to future weights
$w_{smooth_{future}}$	0.05	weights future smoothness
$weights_{style}$	'uni'	Uniform weight style.
$\ w\ _F$	0	Frobenius norm weights
<code>init_style</code>	'step'	initialize based on optimal 1-step
max_{iters}	40	maximum number of iterations
$error_{thres}$	8	threshold to increase error
$with_{reverse}$	False	whether to include reverse update
N_{zeros}	0 or 1 or 2 or 3	how many zeros (relevant only to the sparse networks)

where x , y , and z represent the state variables, and σ , ρ , and β are the system parameters, set to $\sigma = 10$, $\rho = 28$, and $\beta = \frac{8}{3}$.

B Appendix: Neural data additional information

We further demonstrated LINOCS on multiple additional linear settings including data with structured noise (i.e., $\tilde{\mathbf{x}} = \mathbf{x} + \sigma \sin(3t)$, with $\sigma = 0.5$, Fig. 12) and a 3-dimensional cylinder (Fig. 13). We found that LINOCS was able to identify the ground truth dynamics for different training orders and under increasing noise levels also under these settings.

C Code and Data Availability

All algorithm implementation and figure creation codes will be shared over GitHub upon publication, and are attached as a supplement to the submission. The human neural recordings data is available at (Kyzar et al., 2024).

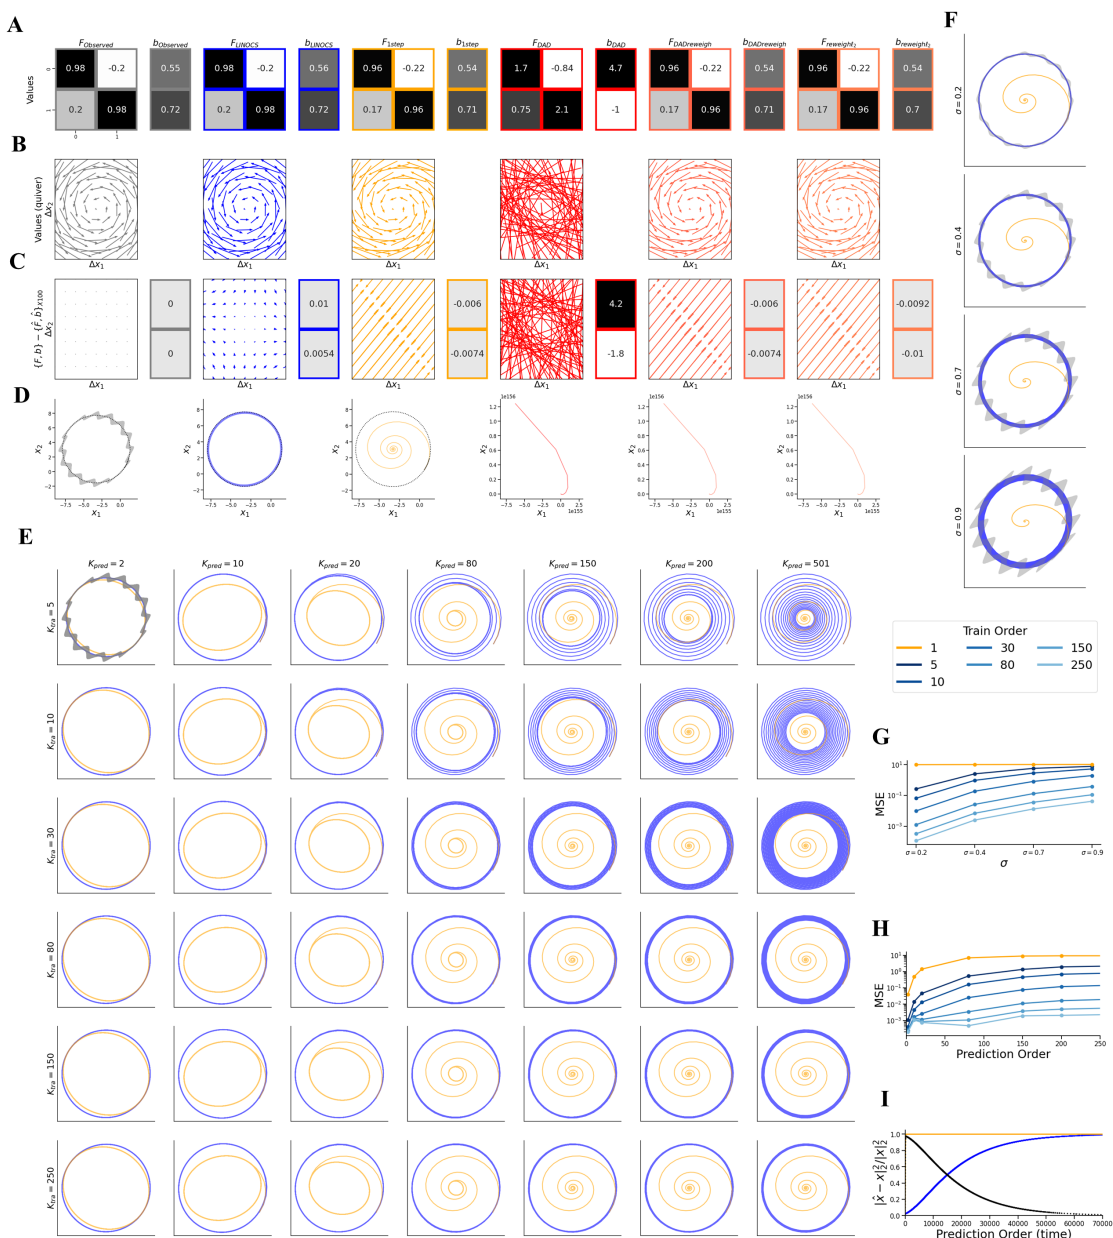


Figure 12: **Linear System under Structured Noise.** **A:** Real vs. identified operators and offsets. **B:** Quiver plots of real and identified operators present patterns that appear similar, rendering it challenging to discern differences when examined in isolation. **C:** The differences in effects between real operators and inferred operators highlight how minor distinctions in dynamic operators gain prominence during lookahead reconstruction (calculation details in A.1). **D:** Full lookahead reconstruction (ground truth vs. baselines) shows swift convergence to the circle’s center for the one-step optimization results due to small differences in dynamic values (mid-yellow subplot) and divergence for DAD-based results (three most-right subplots). **E, H:** MSE under increasing prediction orders. LINOCS achieves better (lower) MSE compared to 1-step optimization. **F, G:** LINOCS reconstruction compared to 1-step optimization under increasing noise values reveals that LINOCS maintains good reconstruction even under extreme noise conditions. **I:** By propagating identified operators until a relative reconstruction error of ~ 1 , LINOCS enables future predictions of $\sim 70,000$ time points, contrasting with immediate convergence in one-step optimization. Black indicates error differences between one-step optimization and LINOCS.

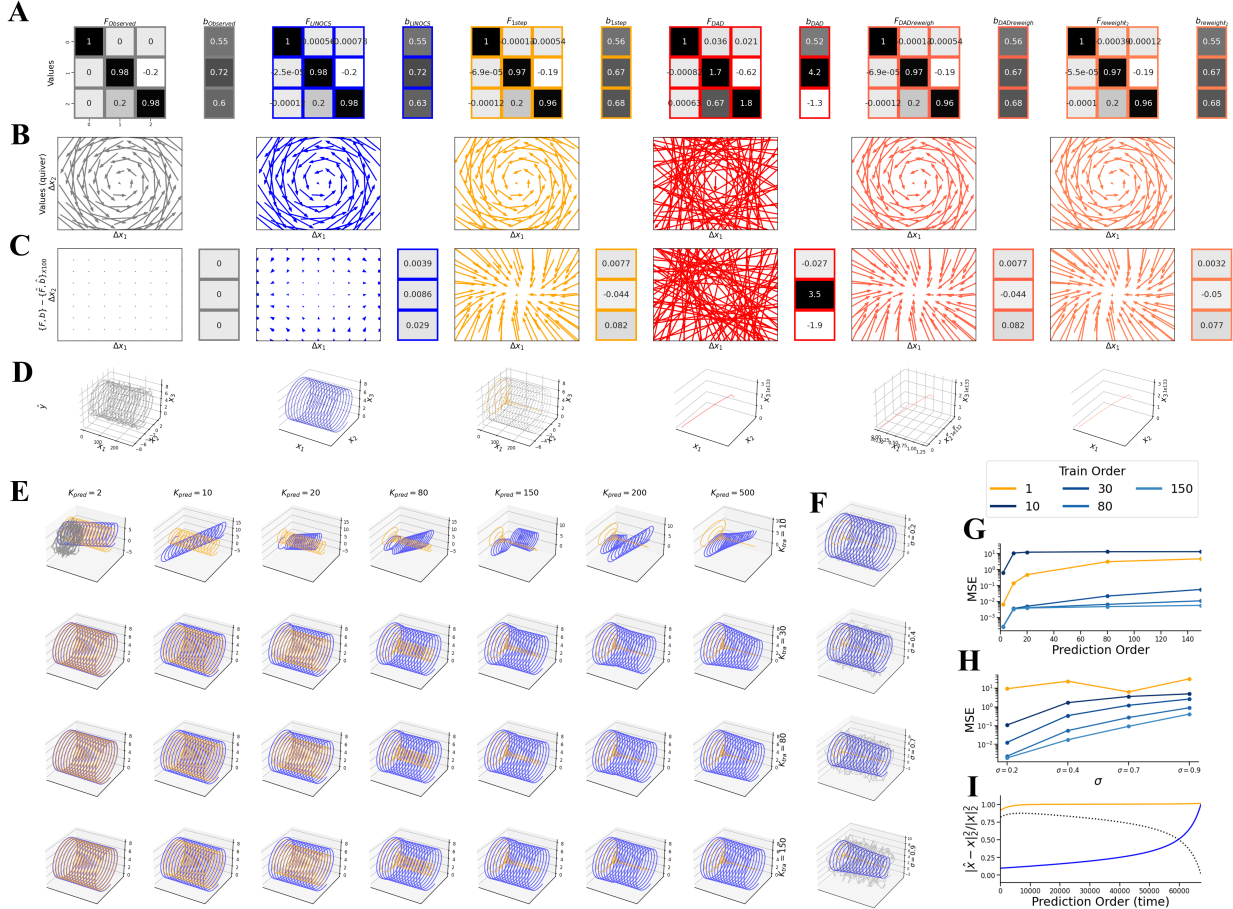


Figure 13: **Linear System for 3D Cylinder.** **A:** Real vs. identified operators and offsets. **B:** Quiver plots of real and identified operators present patterns that appear similar, rendering it challenging to discern differences when examined in isolation. **C:** The differences in effects between real operators and inferred operators highlight how minor distinctions in dynamic operators gain prominence during lookahead reconstruction (calculation details in A.1). **D:** Full lookahead reconstruction (ground truth operators vs. baselines) shows swift convergence to the cylinder’s center for the 1-step optimization results due to small differences in dynamic values (yellow) and divergence for DAD-based results (three most-right subplots). **E, G:** MSE under increasing prediction orders. LINOCS achieves better (lower) MSE compared to 1-step optimization with perfect full lookahead reconstruction under high-enough training order (**E** right bottom). **F, H:** LINOCS reconstruction compared to 1-step optimization under increasing noise values reveals that LINOCS maintains good reconstruction even under extreme noise conditions. **I:** Propagating the identified operators until reaching a relative reconstruction error of ~ 1 shows that LINOCS identifies operators that enable a future prediction of $\sim 70,000$ time points before converging, unlike one-step optimization that converges immediately. black: error difference between one-step optimization and LINOCS.

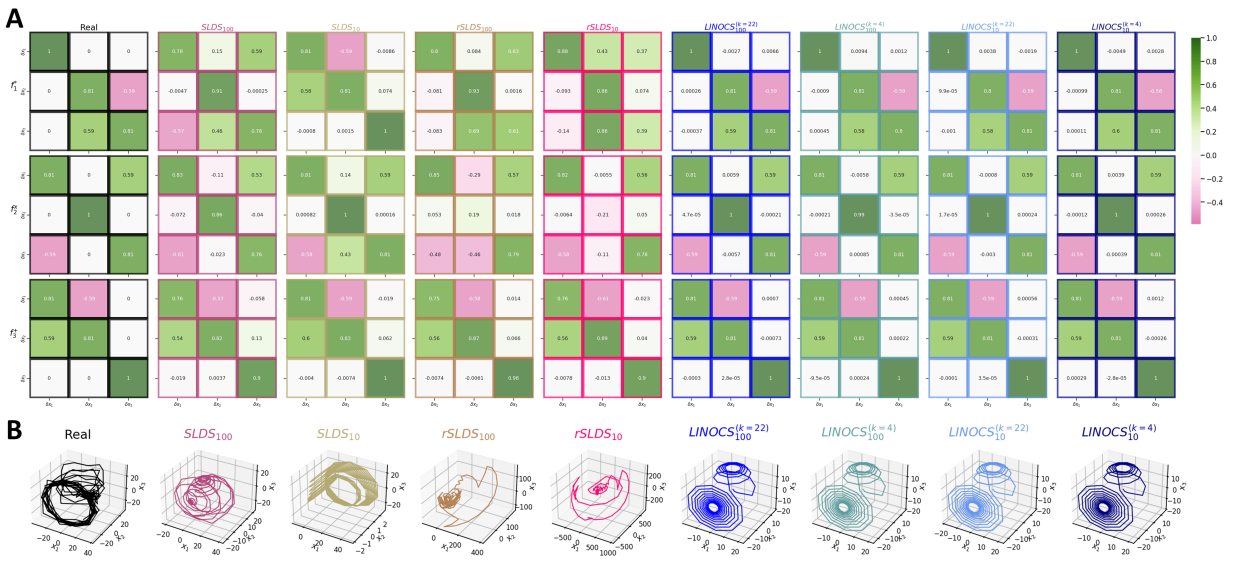


Figure 14: Operators and reconstructions by the LINOCS-driven SLDS compared to classical SLDS. **A**: Dynamical operators identified by LINOCS-driven SLDS. **B**: Full lookahead reconstruction.




Figure 15: **Eigenvalues of identified operators by the different systems compared to the eigenvalues of the ground truth operators.** Rows represent different methods, with LINOCS (with different parameter combination) in the four last rows. Columns represent the three eigenvalues of each of the three different 3×3 linear operators. LINOCS enabled the identification of almost perfect eigenvalues while the other methods found at least one wrong eigenvalue per operator, explaining the decaying/divergence of their reconstruction.




Figure 16: **Ground Truth operators and coefficients for the SLDS experiment.** **A:** The ground truth basis dynamics operators $\{f_j\}_{j=1}^J$ consist of rotational matrices oriented in various directions. **B:** Ground truth operators' coefficients (c). **C:** Ground truth state x .




Figure 17: **Ground Truth operators and coefficients for the “pseudo-switching” dLDS experiment.** **A:** The ground truth basis dynamics operators $\{f_j\}_{j=1}^J$ consist of rotational matrices oriented in various directions. **B:** Ground truth operators’ coefficients (c). **C:** Ground truth state \mathbf{x} .




Figure 18: Real and imaginary parts of the eigenvalues of the time-changing operator F_t captured by LINOCS-driven dLDS




Figure 19: **Ground Truth operators and coefficients for the second dLDS experiment.** **A:** The ground truth basis dynamics operators $\{f_j\}_{j=1}^J$ consist of rotational matrices oriented in various directions. **B:** Ground truth operators’ coefficients (c). **C:** Ground truth state \mathbf{x} .

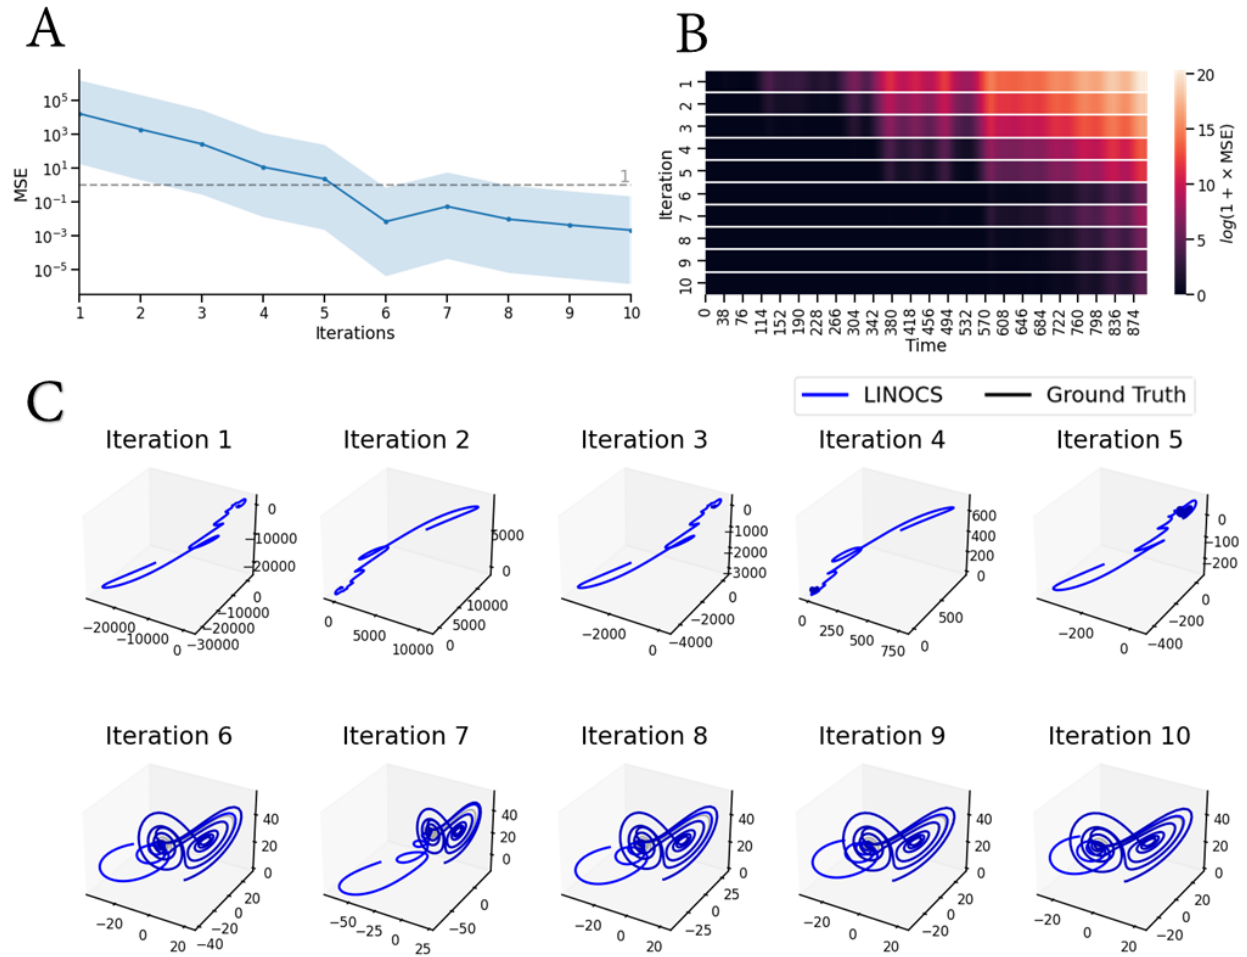


Figure 20: **Demonstration of LINOCS' effect over iterations for Lorenz attractor with 900 time points.** **A:** MSE over iterations (curve corresponds to median values; shade represents 25%-75% percentiles over time). **B:** MSE over time points over training iterations. **C:** Full lookahead reconstruction based on operators identified under different iterations.




Figure 21: **Classical SLDS results (non LINOCS) on the real world data.** The identified networks ($\{\mathbf{f}_j\}_{j=1}^J$) by the non-LINOCS SLDS code (Linderman et al., 2020), for each region in the real world data (Kyzar et al., 2024).

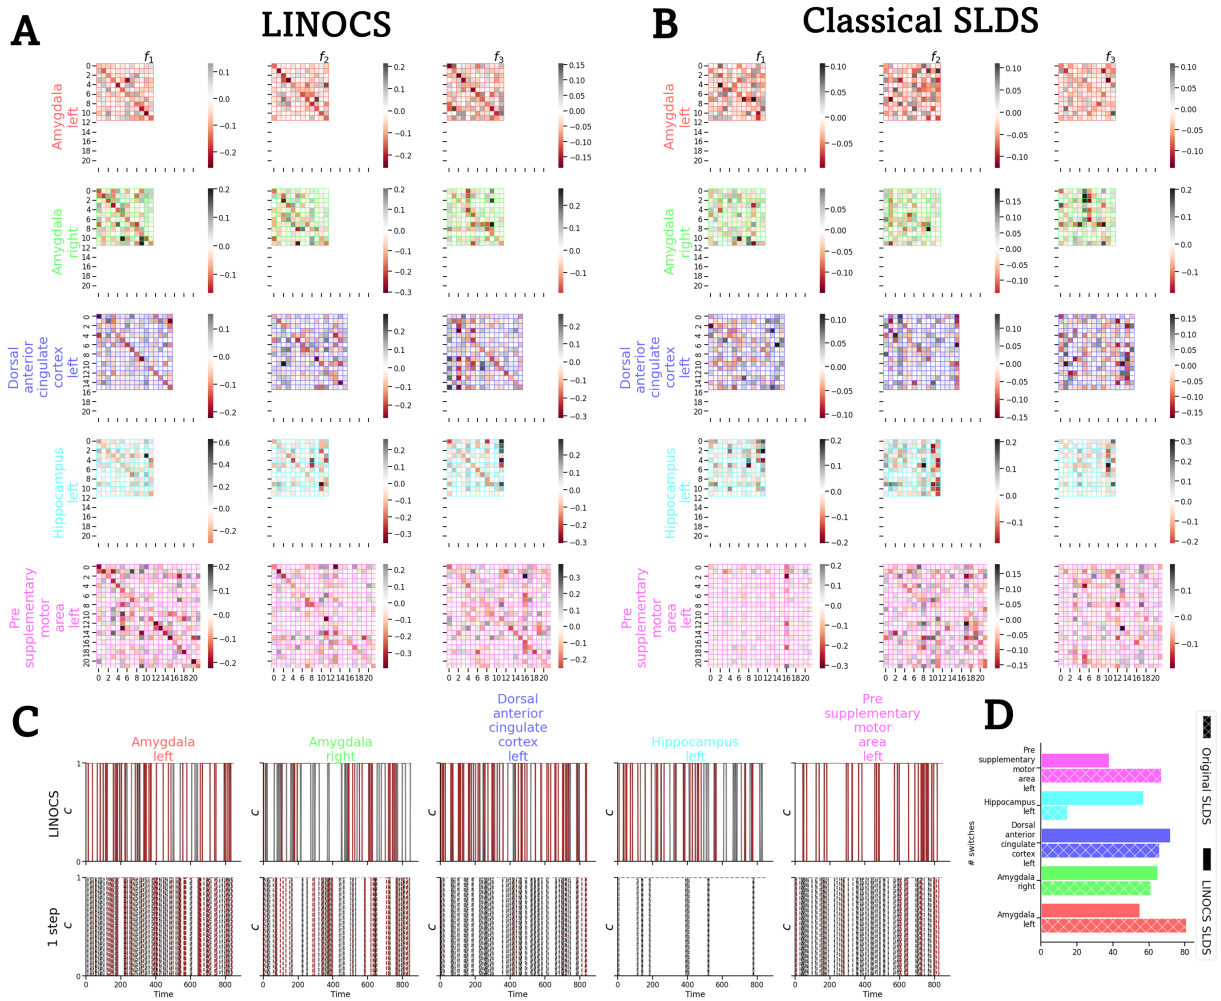


Figure 22: **SLDS results on real data.** **A:** Identified operators per region by LINOCS. **B:** Identified operators per region by classical SLDS. **C:** Switch times by LINOCS-SLDS vs classical SLDS. **D:** Number of switches for LINOCS-SLDS vs. classical SLDS.