

OPEN ACCESS

EDITED BY

Mei Liu,
Multi-scale Medical Robotics Center Limited,
China

REVIEWED BY

Amrit Kashyap,
Charité University Medicine Berlin, Germany
Jason Moore,
NYU Grossman School of Medicine,
United States

*CORRESPONDENCE

Jochen Ditterich
✉ jditterich@ucdavis.edu

RECEIVED 23 February 2025

ACCEPTED 21 April 2025

PUBLISHED 12 May 2025

CITATION

Wu J, Asamoah B, Kong Z and Ditterich J
(2025) Exploring the suitability of
piecewise-linear dynamical system models for
cognitive neural dynamics.
Front. Neurosci. 19:1582080.
doi: 10.3389/fnins.2025.1582080

COPYRIGHT

© 2025 Wu, Asamoah, Kong and Ditterich.
This is an open-access article distributed
under the terms of the [Creative Commons
Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use,
distribution or reproduction in other forums is
permitted, provided the original author(s) and
the copyright owner(s) are credited and that
the original publication in this journal is cited,
in accordance with accepted academic
practice. No use, distribution or reproduction
is permitted which does not comply with
these terms.

Exploring the suitability of piecewise-linear dynamical system models for cognitive neural dynamics

Jiemin Wu¹, Boateng Asamoah², Zhaodan Kong^{3,4} and Jochen Ditterich^{2,4,5*}

¹Department of Computer Science, University of California, Davis, Davis, CA, United States, ²Center for Neuroscience, University of California, Davis, Davis, CA, United States, ³Department of Mechanical and Aerospace Engineering, University of California, Davis, Davis, CA, United States, ⁴Center for Neuroengineering and Medicine, University of California, Davis, Davis, CA, United States, ⁵Department of Neurobiology, Physiology and Behavior, University of California, Davis, Davis, CA, United States

Dynamical system models have proven useful for decoding the current brain state from neural activity. So far, neuroscience has largely relied on either linear models or non-linear models based on artificial neural networks (ANNs). Piecewise linear approximations of non-linear dynamics have proven useful in other technical applications. Moreover, such explicit models provide a clear advantage over ANN-based models when the dynamical system is not only supposed to be observed, but also controlled, in particular when a controller with guarantees is needed. Here we explore whether piecewise-linear dynamical system models (recurrent Switching Linear Dynamical System or rSLDS models) could be useful for modeling brain dynamics, in particular in the context of cognitive tasks. These models have the advantage that they can be estimated not only from continuous observations like field potentials or smoothed firing rates, but also from sparser single-unit spiking data. We first generate artificial neural data based on a non-linear computational model of perceptual decision-making and demonstrate that piecewise-linear dynamics can be successfully recovered from these observations. We then demonstrate that the piecewise-linear model outperforms a linear model in terms of predicting future states of the system and associated neural activity. Finally, we apply our approach to a publicly available dataset recorded from monkeys performing perceptual decisions. Much to our surprise, the piecewise-linear model did not provide a significant advantage over a linear model for these particular data, although linear models that were estimated from different trial epochs showed qualitatively different dynamics. In summary, we present a dynamical system modeling approach that could prove useful in situations, where the brain state needs to be controlled in a closed-loop fashion, for example, in new neuromodulation applications for treating cognitive deficits. Future work will have to show under what conditions the brain dynamics are sufficiently non-linear to warrant the use of a piecewise-linear model over a linear one.

KEYWORDS

cognitive neural dynamics, dynamical system models, non-linear dynamics, perceptual decision-making, piecewise-linear

1 Introduction

Applications like brain-computer/machine interfaces for controlling prostheses or responsive neuromodulation for therapeutic purposes require estimating the current brain state from neural activity (Amunts et al., 2016; Kriegeskorte and Douglas, 2018; Ruff et al., 2018; Hurwitz et al., 2021). Dynamical system models improve the quality of the state estimate by reducing the problem's dimensionality, eliminating noise, and taking advantage of knowledge about how the brain state tends to evolve over time (Rabinovich et al., 2006; Kao et al., 2015; John et al., 2022; Nozari et al., 2024). Dynamical system models, however, do not only allow estimating the brain state, they also support predictions of how the brain state is likely to progress in the near future, which is essential for model-based control, i.e., planning state trajectories for adjusting the brain state through stimulation (Shanechi, 2019; Kim and Bassett, 2020; Srivastava et al., 2020; Moxon et al., 2023). Intelligent implantable stimulators based on such a strategy could open up new avenues for the treatment of, for example, cognitive deficits resulting from neurological and psychiatric disorders (Rao et al., 2018; Sani et al., 2018; Scangos et al., 2021a,b, 2023).

Linear dynamical system models have desirable mathematical properties and seem adequate for decoding movement intentions from the motor cortex (Kao et al., 2015), but cognitive functions might be associated with more complex dynamics and require non-linear models. Neuroscience has largely focused on artificial neural networks for learning non-linear dynamics (Pandarinath et al., 2018; Schulz et al., 2020; Abbaspourazad et al., 2024; Lee et al., 2024), but these black-box models are not ideal for model-based control. Technical applications in engineering often rely on linear approximations of non-linear dynamics (Doyle et al., 2013; Gu et al., 2015; Yang et al., 2018, 2019; Åström and Murray, 2021). The dynamics are still considered sufficiently linear in a local neighborhood in state space, but the systems are allowed to switch/transition to a different dynamic mode based on how the state evolves. In other words, a globally non-linear model can be sufficiently approximated by a piecewise linear model, a collection of locally linear models (Le Quang, 2013; Xu and Xie, 2014). Piecewise linear models (a particular class of hybrid systems) together with linear dynamical models are two of the main pillars of modern control (Borrelli et al., 2017; Hespanha, 2018). *Here we ask whether such piecewise linear dynamical system models could be useful for applications in neuroscience.*

Dynamical system models are often estimated from continuous observations of neural activity (local field potentials or smoothed firing rates). *Here we also address the question of whether these models could be estimated from sparse firing patterns of individual neurons, as they might be observed in higher-order cortical areas related to cognitive processing.*

We approach these questions by first generating synthetic neural data based on a piecewise linear computational model of perceptual decision-making inspired by the model proposed by Wong and Wang (2006). We generate both continuous observations as well as single-neuron spike trains based on Poisson processes, whose rates are controlled by the state of the dynamical system. We demonstrate that the piecewise-linear or recurrent switching linear dynamical system (rSLDS) models

(Linderman et al., 2017) can be successfully recovered from the synthetic data. Furthermore, we show that the estimated rSLDS models provide a significantly better explanation for the observed data than linear dynamical system (LDS) models and, more importantly, also make significantly better predictions for future states of the system. Finally, we apply our approach to a publicly available dataset collected from non-human primates performing a perceptual decision task.

2 Materials and methods

We start by introducing the rSLDS model, explaining the theory behind it and how it is applied to neural data. We then outline the synthetic data generation process, which is essential for testing our models before applying them to actual neural data. Next, we present the evaluation metrics we use to measure the performance of the different models. Our overall approach is illustrated in Figure 1.

2.1 rSLDS model

Recurrent Switching Linear Dynamical Systems (rSLDS) (Linderman et al., 2017) are an advanced class of piecewise linear dynamic models designed to decompose complex non-linear time-series data into a series of discrete segments (modes), each governed by a simpler linear dynamical system (LDS). The rSLDS framework can be conceptualized as a composite of multiple LDS models, where transitions between these models are determined by the latent state itself.

Unlike traditional linear dynamical system (LDS) models, recurrent switching linear dynamical system (rSLDS) models offer a piecewise-linear approximation of non-linear dynamics by dividing the state space into multiple regions, each governed by its own linear dynamics. This approach allows rSLDS to capture more complex dynamics while maintaining the mathematical tractability and interpretability of linear systems. In contrast, artificial neural networks (ANNs), although powerful for modeling highly non-linear relationships, often act as black-box models, making it difficult to interpret their internal workings or apply them effectively in model-based control scenarios. The transparency and structure of rSLDS models facilitate not only better understanding of the underlying neural dynamics but also enable precise computation of necessary stimuli for achieving desired brain states through existing control algorithms. Thus, rSLDS presents a balanced solution between capturing the complexity of neural systems and providing actionable insights for therapeutic applications.

Given an rSLDS model comprising K distinct LDS models, or modes, and time-series data observed over T time steps, the model's formulation is as follows:

For each time step $t \in T$, a *discrete* latent state $z_t \in \{1, 2, \dots, K\}$ is defined to signify the LDS model currently in effect that only depends on the last state, exhibiting Markovian characteristics. This state evolves according to the logistic regression model with a weight matrix $\mathbf{R}_{z_t} \in \mathbb{R}^{(K-1) \times H}$ and a bias vector $\mathbf{r}_{z_t} \in \mathbb{R}^{K-1}$:

$$z_{t+1} \mid z_t, \mathbf{x}_t \sim \pi_{\text{SB}}(\mathbf{v}_t), \quad \mathbf{v}_t = \mathbf{R}_{z_t} \mathbf{x}_t + \mathbf{r}_{z_t}, \quad (1)$$

2.2 Evaluation metrics

To test model quality, we used the following three metrics: (1) the evidence lower bound (ELBO), (2) the coefficient of determination (R^2), and (3) the Mean Euclidean Distance (MED).

The ELBO, synonymous with the variational lower bound, serves as a metric for the log-likelihood of observational data within the framework of variational Bayesian models. It is instrumental in comparing the performance of various probabilistic models across a consistent dataset, with a larger ELBO being indicative of better model performance.

The coefficient of determination, denoted as R^2 , quantifies the proportion of variance in the observed data that can be explained by a model. R^2 is typically used for one-dimensional, continuous data. Since we are dealing with observation vectors, we adopted the multi-dimensional R^2 measure proposed by Jones and Meyer (2024), which is based on the Euclidean distances between observed and predicted vectors. For multi-dimensional Poisson-distributed spike count data, we extended the deviance-based R^2 measure for one-dimensional Poisson-distributed data proposed by Cameron and Windmeijer (1996), assuming statistical independence between the observed spike counts for different neurons. We define the deviance term $D(y_{n,t}, \hat{\mu}_{n,t})$ as:

$$D(y_{n,t}, \hat{\mu}_{n,t}) = y_{n,t} \cdot \log\left(\frac{y_{n,t}}{\hat{\mu}_{n,t}}\right) - (y_{n,t} - \hat{\mu}_{n,t}) \quad (5)$$

where $y_{n,t}$ is the observed spike count for neuron n in a time bin at time point t . $\hat{\mu}_{n,t}$ is the model's predicted rate for the same neuron and time point. Using this deviance term, the deviance-based R^2 for Poisson-distributed data is given by:

$$R_{\text{DEV},P}^2 = 1 - \frac{\sum_{n=1}^N \sum_{t=1}^T D(y_{n,t}, \hat{\mu}_{n,t})}{\sum_{n=1}^N \sum_{t=1}^T D(y_{n,t}, \bar{y}_n)} \quad (6)$$

where \bar{y}_n is the mean spike count for neuron n across all time points, N is the number of observed neurons, and T is the total number of time bins.

It is important to note that, while R^2 can approach values close to unity in the context of assessing model quality, reflecting the extent to which measurement noise limits its value, the scenario is different for Poisson-distributed observations. Even for a model that accurately captures the underlying rates of the Poisson processes, there remains an inherent stochastic variability in the observed counts, which cannot be captured by the model and therefore imposes a ceiling on the attainable R^2 values. To address this, we calculate the expected value of the numerator in the R^2 computation, assuming a Poisson process with rates as determined by the model, i.e., assuming that $y_{n,t}$ is Poisson-distributed with mean $\hat{\mu}_{n,t}$:

$$R_{\text{Expected,DEV},P}^2 = 1 - \frac{\sum_{n=1}^N \sum_{t=1}^T E_{y_{n,t} \sim \text{Poisson}(\hat{\mu}_{n,t})}[D(y_{n,t}, \hat{\mu}_{n,t})]}{\sum_{n=1}^N \sum_{t=1}^T D(y_{n,t}, \bar{y}_n)} \quad (7)$$

A model that can explain all systematic (i.e., other than Poisson) variability in the data is therefore expected to have a $R_{\text{DEV},P}^2$ that matches $R_{\text{Expected,DEV},P}^2$.

Finally, the mean Euclidean distance (MED) is utilized to measure the average discrepancy between the model's predictions \hat{y}_t and the observed data y_t , and is expressed as:

$$\text{MED} = \frac{1}{T} \sum_{t=1}^T \|y_t - \hat{y}_t\|_2 \quad (8)$$

This metric directly assesses the model's accuracy in predicting the observed phenomena.

2.3 Synthetic data

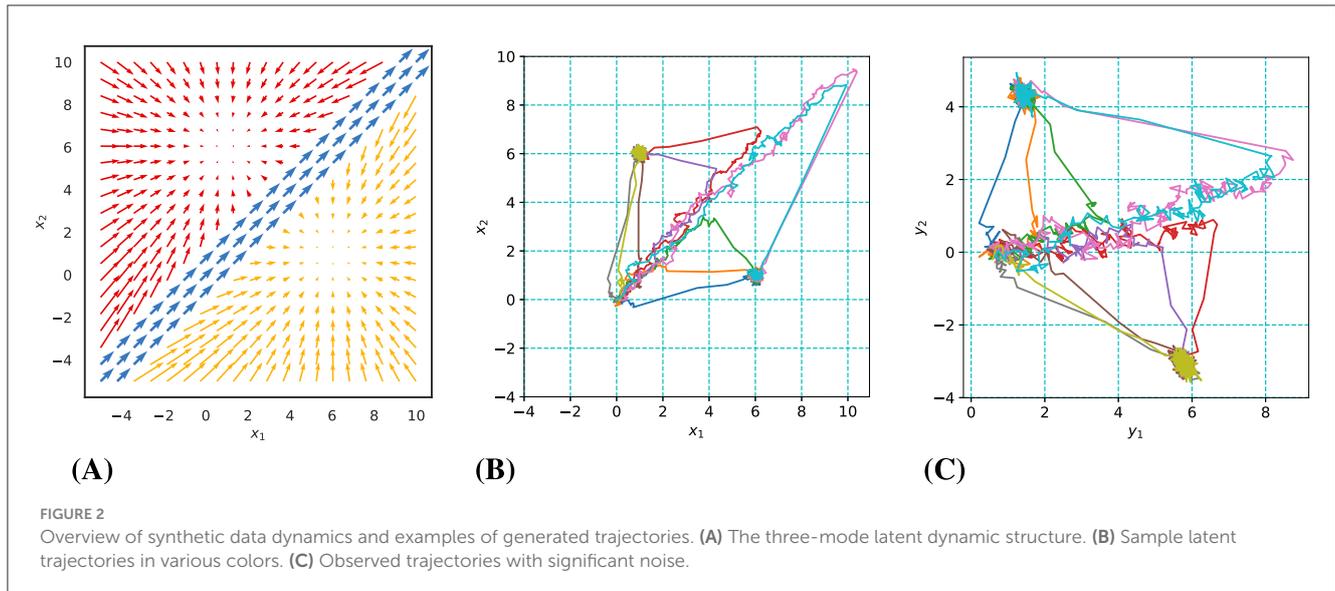
We first generated synthetic data on the basis of a piecewise linear adaptation of the Wong and Wang decision-making model (Wong and Wang, 2006), designed for decision-making between two choice options. This simulation assumed an equal provision of sensory evidence for both options on average. Since we can access the underlying dynamics, this synthetic data allowed us to test the piecewise approach.

2.3.1 Latent dynamics

The model underlying our synthetic data has a two-dimensional state space consisting of three modes (Mode 1: blue, Mode 2: red, Mode 3: yellow) as illustrated in Figure 2A. Trials initiate in Mode 1, characterized by the accumulation of noisy sensory evidence for both alternatives. This noisy evidence directs the state predominantly along the diagonal. When enough excess evidence has been accumulated in favor of one of the two choices, the system transitions into either Mode 2 or 3. These latter modes are each defined by a point attractor, which the system state converges toward. The transition from Mode 1 to the next is determined by whether the system state has crossed the boundary of the region of Mode 2 or 3. State transitions are governed by Equation 2, where z_t represents the active mode from the set $\{1, 2, 3\}$. Trials start at the origin (0,0) and proceed into the first quadrant, with mode switches occurring when the absolute difference between state coordinates exceeds one unit. The point attractors of Mode 2 and 3 are located at (1,6) and (6,1), respectively, and sample trajectories are depicted in Figure 2B.

2.3.2 Observations

Employing the three-mode latent dynamic structure, we generated synthetic datasets following Gaussian and Poisson distributions, as detailed in Equations 3, 4. Some example Gaussian observation trajectories (for two-dimensional observations) are shown in Figure 2C. Parameters \mathbf{C} and \mathbf{d} were adjusted to achieve a range of average firing rates. Similar to the situation in our real dataset, the firing rate varied across neurons. For example, in the case of an average firing rate of 2 spikes/bin across neurons, the firing rates of individual neurons ranged from about 0.1 spikes/bin (10th percentile) to about 5 spikes/bin (90th percentile). Each dataset comprised 250 trials, with 200 designated for model identification and 50 reserved for performance evaluation in a testing period. This process allowed us to ascertain the predictive capabilities of our model through a controlled assessment.



2.4 Visualization of higher-dimensional states

We employed Principal Component Analysis (PCA) to project these dynamics onto a 2-D embedding space for a visual representation of the inferred high-dimensional latent dynamics. PCA is a method of linear dimensionality reduction that effectively maps high-dimensional data into a lower-dimensional space, optimizing for preserving maximum data covariance and maintaining large pairwise distances.

To ensure that the choice information embedded within the inferred high-dimensional latent dynamics, denoted as $x \in \mathbb{R}^L$, is retained as much as possible after dimension reduction via PCA, we derived the normalized difference response matrix D using the following computation (Galgali et al., 2023):

$$D = 0.5 \times (\bar{x}_{choice1} - \bar{x}_{choice2}) \quad (9)$$

Here, $\bar{x}_{choice1}, \bar{x}_{choice2} \in \mathbb{R}^{T \times L}$ represent the trial-averaged latent dynamics. Each row within matrix D corresponds to the average observation difference between the two choices at a given time step. These row vectors serve as samples from which a PCA model can be fitted, yielding the final 2-D embedding of the latent dynamics, denoted as x_{2D} .

2.5 Choice decoding

To evaluate the preservation of choice information after dimensionality reduction, we decoded the animal's choice using a Support Vector Machine (SVM) for each time step. For each time step t , we combined the latent states from trajectories corresponding to choice 1 and choice 2, and then split this combined dataset into a training set (90%) and a testing set (10%). We trained an SVM classifier on the training set to

classify the latent states into two classes based on their ground truth labels, performing this training for both the original 8-D latent states and the 2-D embedding latent states. After training, we applied the SVM to the testing set to predict the choice labels and computed the predictive accuracy, which is the proportion of correctly classified test samples. By comparing the predictive accuracies obtained from classifying the original 8-D latent states and the 2-D embedding, we assessed the effect of the dimensionality reduction.

2.6 Model comparison

To provide a clear and thorough comparison of the LDS and rSLDS models' performance on the synthetic dataset, we designed a series of analyses focused on three critical areas: (1) estimating models, (2) making predictions, and (3) subsequent analysis of results.

2.6.1 Estimating models

Prior to initiating our data analysis, we systematically partitioned the synthetic dataset into three subsets: a training set, a validation set, and a test set, adhering to a 7:1:2 ratio. Given the stochastic nature of model initialization and the utilization of the variational EM algorithm, there exists an inherent variability in the training outcomes, even when the model configurations and training data remain constant. To mitigate the risk of models converging to suboptimal solutions, particularly under conditions characterized by high noise levels and a limited dimensionality of observed variables, we implemented a strategy of training 40 instances of both the LDS and rSLDS models concurrently for each analysis setup. Subsequently, we identified the model variant that demonstrated the most robust performance, as indicated by the highest R^2 value (smallest prediction error) on the validation set, for advancement to the subsequent predictive tasks and comparative result analysis.

2.6.2 Making predictions

Following the training period, each model was tasked with forecasting observations (Gaussian observations or spike counts) for the next 10 time points, leveraging historical data up to the current time point for every trial within the test set. These predictive outputs were then juxtaposed with their corresponding ground truth observation data and aggregated to facilitate the computation of the evaluation metrics: the coefficient of determination (R^2) and Mean Euclidean Distance (MED).

2.6.3 Analysis of results

The analysis of results involved a detailed examination of model performance across different hyperparameter settings, particularly focusing on the effects of varying the number of modes and latent dimensions. A Wilcoxon signed-rank test (Woolson, 2005) was applied to the R^2 and MED metrics to assess the significance of differences between the models. Additionally, we visualized the inferred latent trajectories and computed discrepancy scores relative to the ground truth model.

For the computation of discrepancy scores in relation to the ground truth model, it is essential to first map the estimated model into the same state space as the ground truth model. This alignment is necessary because, even if the estimated model converges to an optimal solution, it can still be an arbitrary linear transformation of the original model that generated the data. This arises due to the observation equation introducing a separate linear transformation between the hidden states and the observations. Therefore, before comparing the models, they must be aligned through a common linear transformation.

We denote the true model parameters as A_z , b_z , C , d , R , and r , and the model parameters that were estimated from the observations as \hat{A}_z , \hat{b}_z , \hat{C} , \hat{d} , \hat{R} , and \hat{r} . Our goal is to find the linear transformation that establishes the link between a location \hat{x}_n in the state space of the estimated model and the corresponding location x_n in the original state space:

$$\hat{x}_n = Fx_n + g \quad (10)$$

The transformation matrix F and the bias vector g can be found using the original observation equation determined by C and d and the estimated observation equation determined by \hat{C} and \hat{d} , as the observations have to be the same. F and g can then be used to map the estimated model parameters back into the original state space, resulting in A'_z , b'_z , R' , and r' . F and g are determined by:

$$F = (\hat{C})^{-1}C, \quad g = (\hat{C})^{-1}(d - \hat{d}) \quad (11)$$

If \hat{C} is not invertible, the Moore-Penrose pseudo-inverse can be used instead. With F and g determined, the other estimated parameters are transformed as follows:

$$C' = \hat{C}F = C, \quad d' = \hat{C}g + \hat{d} = d \quad (12)$$

$$A'_z = F^{-1}\hat{A}_zF, \quad b'_z = F^{-1}(\hat{A}_zg + \hat{b}_z - g) \quad (13)$$

$$R' = \hat{R}F, \quad r' = \hat{R}g + \hat{r} \quad (14)$$

These transformations allow us to align the estimated parameters with the ground truth model for a direct comparison. Using these aligned parameters, we computed the discrepancy score, which included differences in attractor locations, switching boundaries, and eigenvalues of A_z . Based on these scores, models were classified as “excellent” (score < 0.5), “good” (0.5 to 2), or “poor” (score > 2), providing a quantitative assessment of model fidelity to the true parameters.

The same type of transformation was also used for aligning different models in section *Models for different task periods and Model initialization*.

To allow for a comparison of prediction errors across models using a Wilcoxon signed-rank test, we computed the R^2 and MED distributions on a trial-by-trial basis for the test set.

3 Results

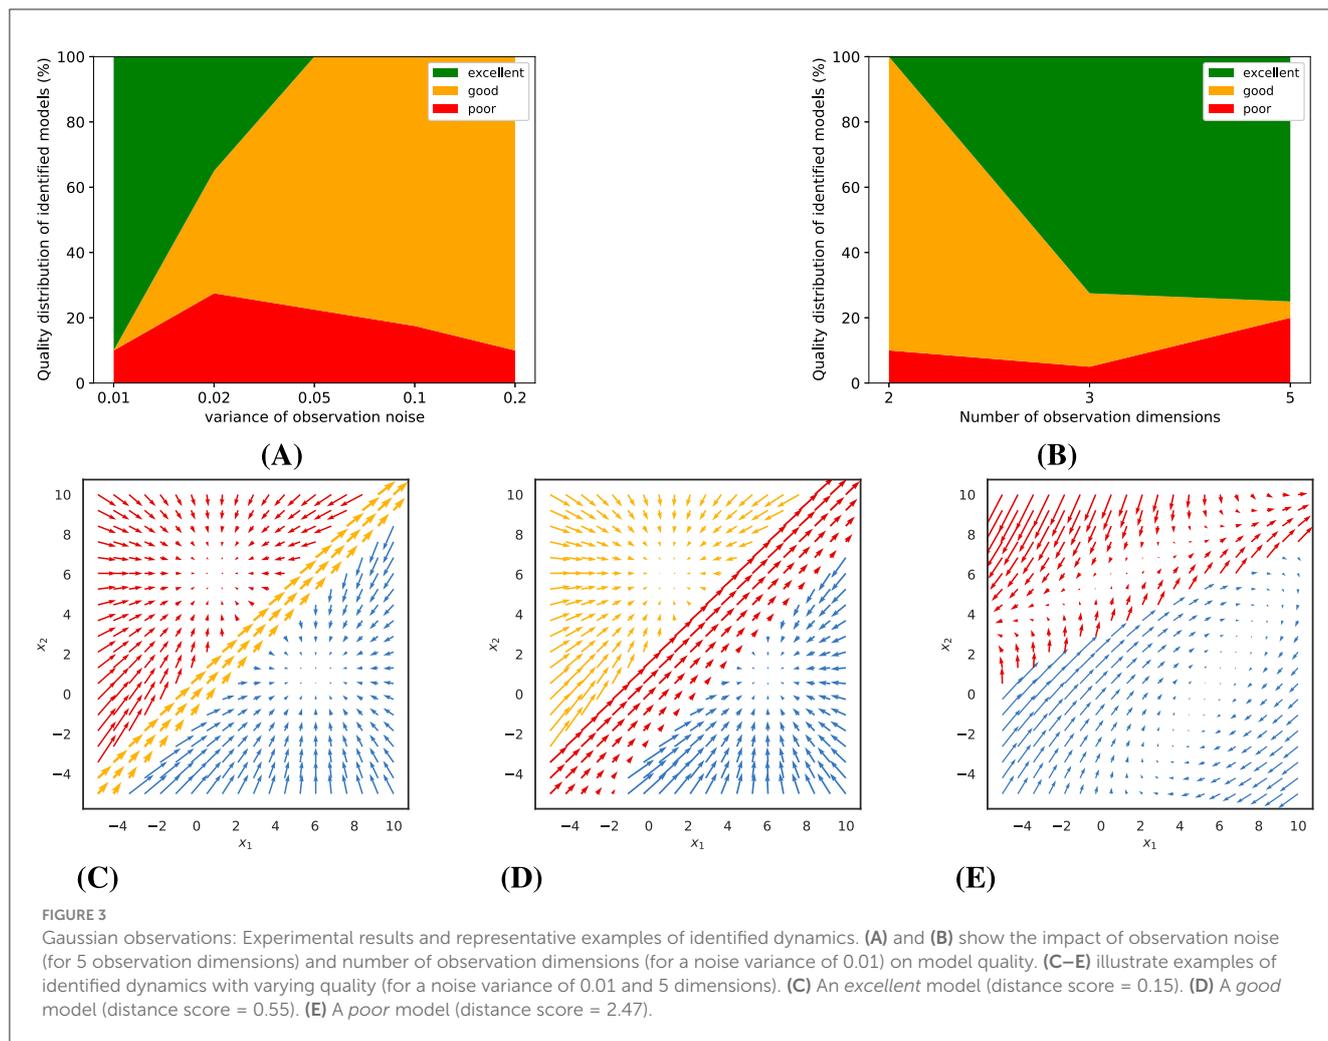
3.1 Analysis of synthetic data

3.1.1 Gaussian observations

Our model fits are based on variational expectation maximization, which does not guarantee convergence to an optimal solution. We therefore always fitted a larger number of models of the same type with somewhat different initialization, report the distribution of the quality of these models, and then used the best one for further analysis (see section 2.6.1 for details). For Gaussian observations with low noise levels, we observed that a multitude of piecewise-linear model fits achieved excellent convergence, yet not all models were equally successful. The probability of identifying a superior model diminished with increasing observation noise, as depicted in Figure 3A. Conversely, increasing the number of simultaneous observations, or the observation dimensions, significantly elevated the likelihood of discovering a model of high quality, as illustrated in Figure 3B.

Building upon these findings, we categorized the trained models into three distinct tiers based on their computed distance scores relative to the ground truth dynamics: poor, good, and excellent. To provide a visual context to this classification, Figures 3C–E presents examples of dynamic vector fields of models that correspond to each of these categories.

To test whether the rSLDS model provided a more useful description of the dynamics in terms of being able to predict how the system would evolve in the near future compared to a LDS model, we tasked both models with predicting how the (hidden) model state and the associated observations would evolve over the next 10 time steps and compared these predictions with the ground truth data. Figure 4 elucidates the comparative performance between the LDS and rSLDS models on Gaussian data in observation space (top) and latent state space (bottom), highlighting the distinct advantages of the rSLDS model in terms of both the coefficient of determination (R^2) and Mean Euclidean Distance (MED) metrics (between predicted and observed vectors). The results from a Wilcoxon signed-rank test strongly indicate a significant advantage of the rSLDS model over the LDS model (smaller prediction errors), with a p-value of < 0.001 for all



measures and all prediction horizons, underscoring the robustness of the rSLDS model in accurately capturing the dynamics of the synthetic datasets.

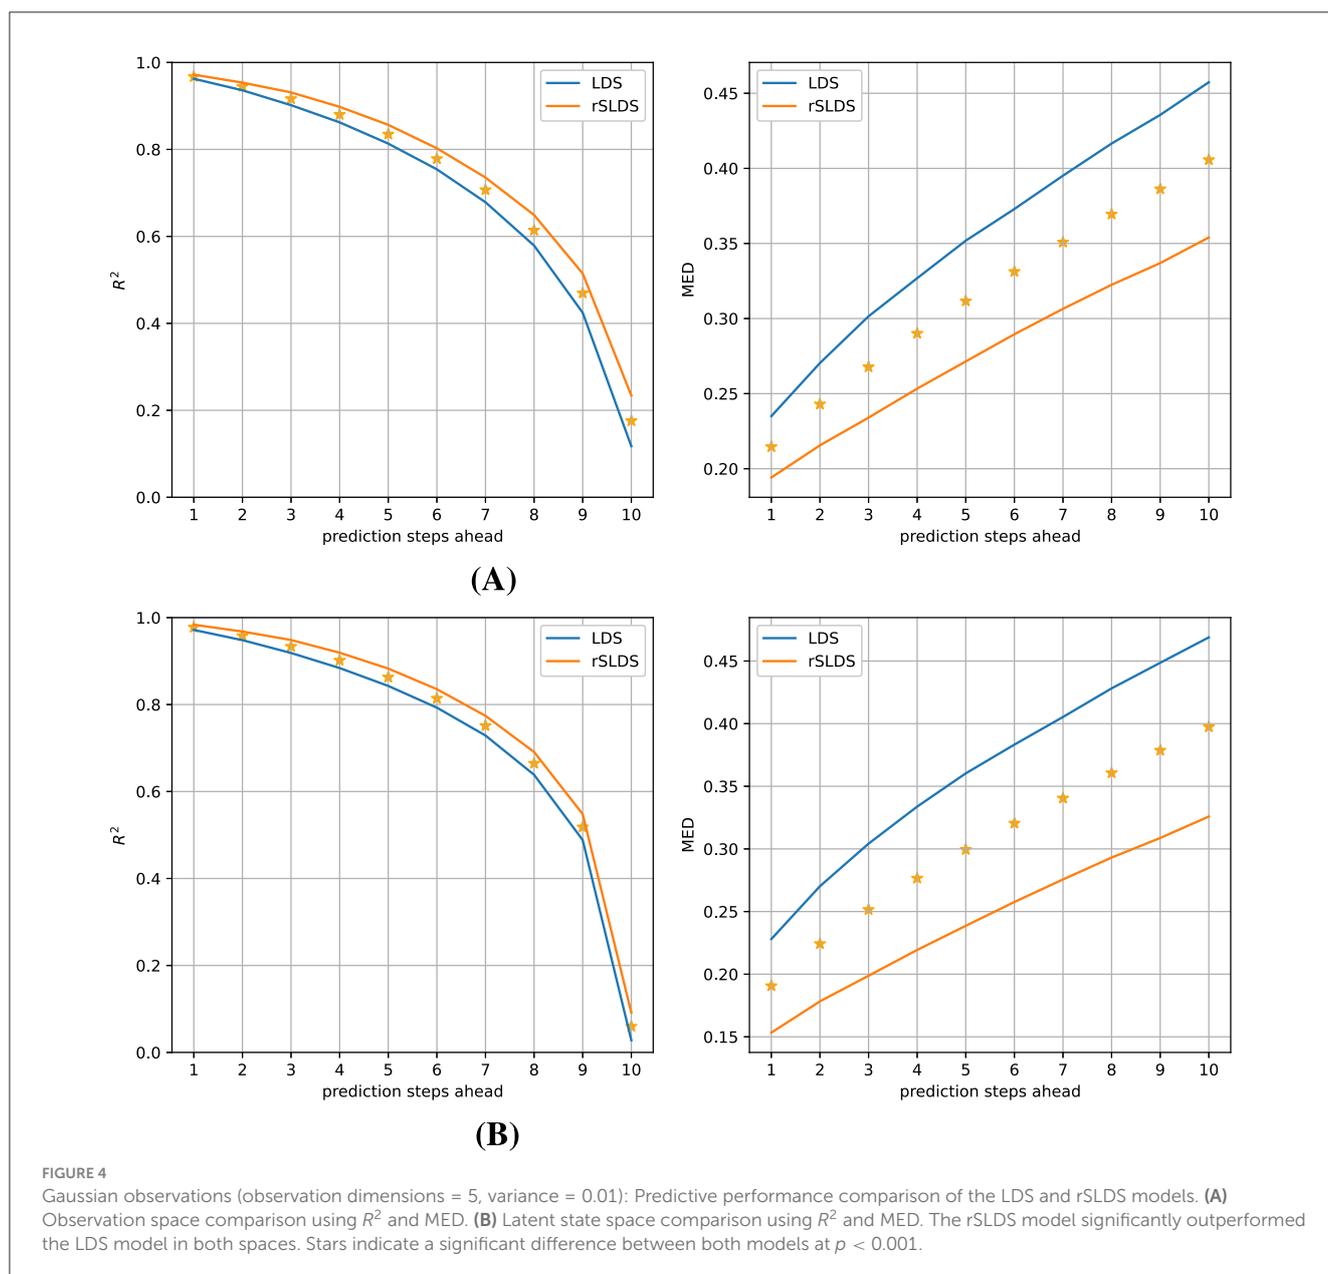
3.1.2 Poisson observations

When dynamical system models have to be estimated from single-unit spiking data, the typical approach is to count spikes in bins of some chosen width, apply some smoothing, and treat the result as continuous observations. However, when dealing with neurons with low firing rates and/or when a good temporal resolution is desired, i.e., when the time bins should be rather narrow, the spike count data can become pretty sparse (with a substantial number of zero counts), and this approach doesn't work anymore. We therefore wanted to test whether it is still possible to estimate rSLDS models from these data, treating them as Poisson rather than continuous observations. Since the variance of a Poisson distribution is identical to its mean, the inherent observation noise will typically be larger than what we have been considering in the case of Gaussian observations.

As a consequence, as can be seen in [Figure 5](#), the number of excellent fitted models is much smaller. The quality of the fits depends on the expected number of spikes per bin and on the

number of observed neurons. The higher the Poisson rate, the more likely it is to obtain a good or perhaps even excellent model. About 20% of the fitted models were still good for an expected spike count of 0.5 spikes per bin. We would not recommend applying the technique to much lower rates. For an expected spike count of 2 spikes per bin, a few excellent models were only observed for a large number of observations. Higher firing rates (e.g., 4 or 8 spikes/bin) improved model quality, with excellent fits increasing to 8% and 15%, respectively. At the same time, the number of poor models also increased with the number of observations (neurons), likely due to the increasing number of model parameters, which increases the risk of getting stuck in suboptimal solutions. As mentioned before, it is therefore important to always fit a larger number of models and to select the best one. Similar to [Figures 3C–E](#), [Figures 5C–E](#) present examples of models with varying quality, which were estimated from the synthetic Poisson data (with an average rate of 2 spikes/bin and 200 observation dimensions).

When creating our synthetic dataset, we used a model with two hidden dimensions and three modes. When fitting models with various numbers of hidden dimensions and modes to the data, the Evidence Lower Bound (ELBO) indicated that, when starting with a 2-dimensional model with one mode, adding a 2nd and a 3rd mode to the model improved the quality of the fit much more than adding more hidden dimensions (see [Figure 6](#)).

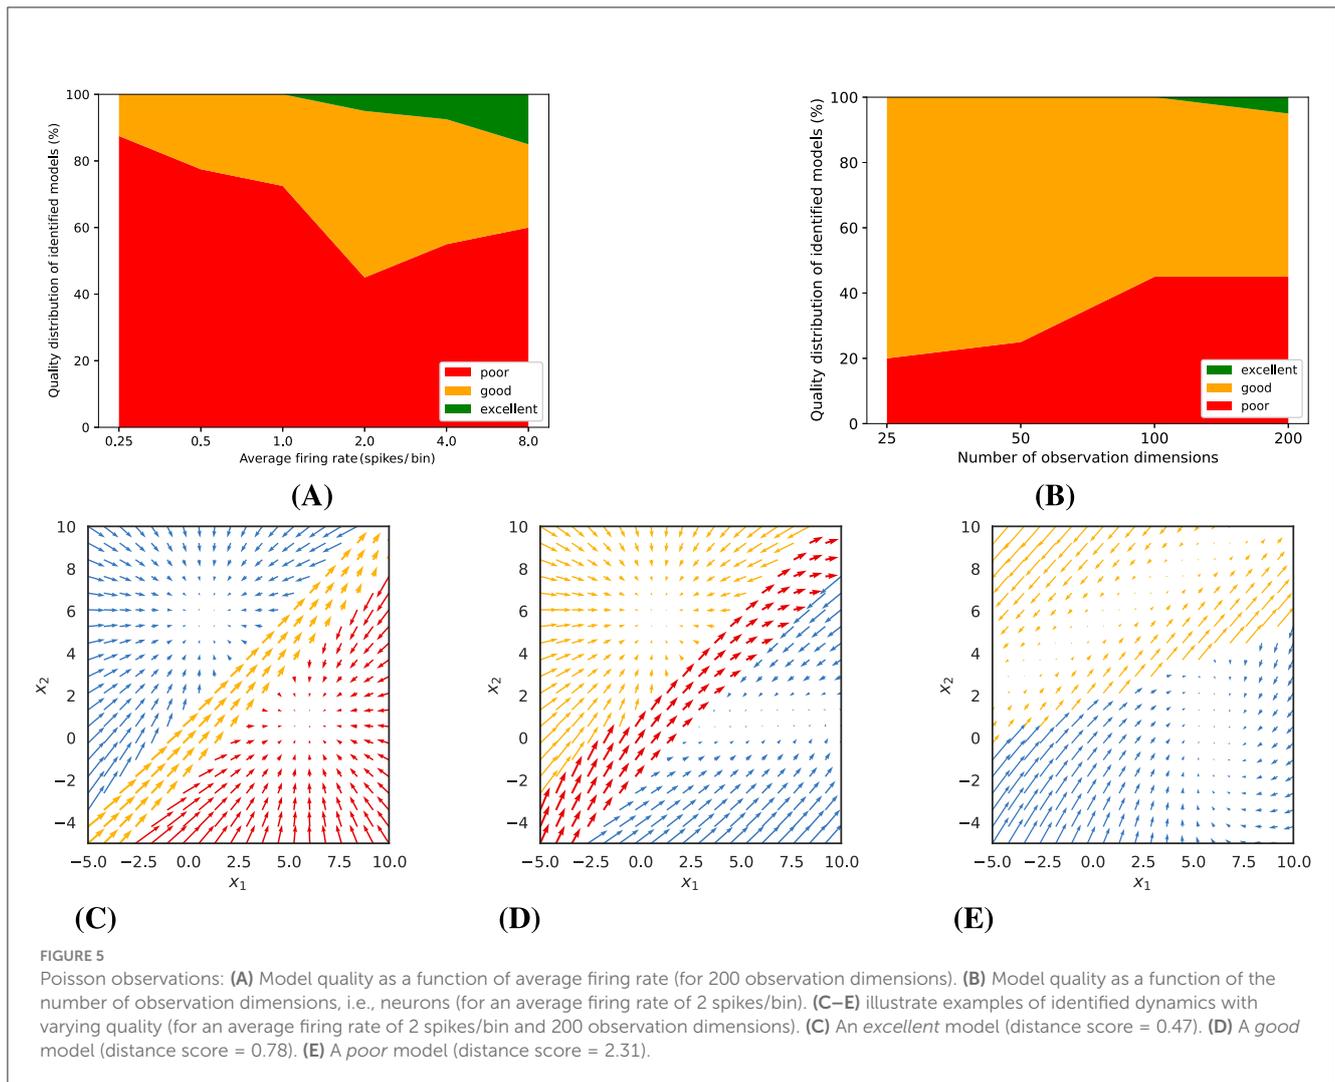


Despite the larger challenge of estimating a good rSLDS model from the Poisson data compared to the Gaussian data earlier, its predictive accuracy still surpasses that of an LDS model, as illustrated in Figure 7.

An analysis leveraging the coefficient of determination (R^2) and Mean Euclidean Distance (MED) metrics demonstrates that the rSLDS model significantly outperforms the LDS model in forecasting up to at least 7 time steps ahead. Due to the inherent variability of the Poisson observations, the observed R^2 values are much lower than in the earlier case of Gaussian observations. To still get an idea of the quality of the estimated models, we calculated what the R^2 value was expected to be, if the model had perfectly captured the Poisson rates, which we refer to as the “Expected R^2 ” (see Section 2), plotted in green in Figure 7A. The results suggest

that our rSLDS model is close to optimal for predictions one and two time steps ahead.

Since we know the ground-truth hidden states underlying our synthetic data, we can also calculate the prediction error in the latent state space, which is shown in Figure 7B. Since this eliminates the inherent variability of the Poisson observations, the observed R^2 values are much larger, and the significant advantage of the rSLDS model compared to the LDS model is more obvious and statistically significant for all analyzed prediction horizons. This indicates that the rSLDS model outperforms the LDS model in terms of decoding and predicting system states from Poisson observations. We will not be able to perform this latter analysis for the real neural data, as the ground-truth hidden states are unknown.



3.2 Analysis of real neural data from perceptual decision-making

To test our approach with real neural data, we took advantage of a publicly available dataset of ~ 200 simultaneously recorded neurons in prefrontal cortex of monkeys performing a perceptual decision task. The “Kiani dataset” was derived from a direction discrimination task involving macaque monkeys trained to perform a task, in which they viewed a patch of randomly moving dots for 800 ms, followed by a variable-length delay period (Kiani et al., 2014). At the end of the delay, a “Go” cue prompted the monkeys to report their perceived motion direction by making a saccadic eye movement to one of two targets (T1 or T2). We analyzed either data from the whole trial or from two particular time periods: the “Decision Period,” which extends from the onset of the dots to the “Go” cue, representing the time during which the monkey makes its decision, and the “Motor Period,” which was defined as starting 100 ms before the “Go” cue and ending 150 ms after the saccade onset, indicating the window in which the monkey prepares and executes its eye

movement response. Neural activity was simultaneously recorded from hundreds of single- and multi-neuron units in the prearcuate gyrus (area 8Ar) using a 96-channel multi-electrode array, which covered a $4\text{ mm} \times 4\text{ mm}$ area of the cortical surface. This setup allowed for the dynamic decoding of the decision variable and the prediction of the monkeys’ choices based on neural population responses.

In this section, we delve into the analysis of the Kiani dataset, aiming to uncover insights into the neural dynamics underlying the Decision and Motor Periods. We begin by comparing the predictive performance of two models, the linear dynamical system (LDS) and the recurrent switching linear dynamical system (rSLDS), across varying levels of model complexity, using data from the full trials. For the rSLDS model, we utilize a 2-mode configuration, as our preliminary tests with three and four modes did not yield significantly different results. We then explore the Decision and Motor Periods separately to address whether the underlying dynamics are different or the same. Subsequently, we examine how well the animal’s choice can be decoded from estimated hidden states.

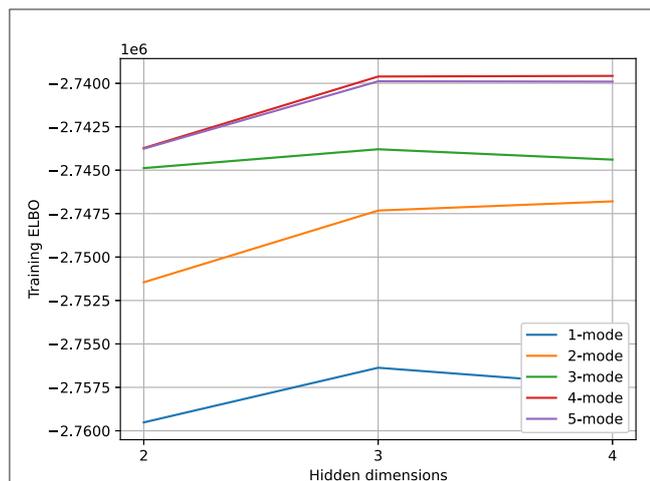


FIGURE 6
Poisson observations: Impact of varying the number of modes and hidden dimensions of the rSLDS model on the Training ELBO (ability of the model to explain the data). Adding a second and a third mode provides more benefit than adding a third or fourth hidden dimension.

3.2.1 LDS vs. rSLDS models

The number of hidden dimensions is a pivotal hyperparameter for both linear and piecewise linear models in modeling neural data. To evaluate the efficacy of these models in capturing the latent dynamics encoded in the neural data, we conducted a comparative analysis using the T33 session from the Kiani dataset, the session with the best behavioral performance. The dataset was segmented into 80 ms-long time bins, a duration chosen based on the average firing rate of the recorded neurons, which is ~ 7 spikes per second [ranging from about 2 spikes/s (10th percentile) to about 12 spikes/s (90th percentile)]. This results in an expected spike count per bin of about 0.6, a value for which we anticipate being able to estimate good models, based on our insights from the synthetic data.

We estimated models with hidden dimensions varying from 2 to 24, aiming to predict the spike count of the next time bin based on the current one. The predictive accuracy of these models was again measured using the coefficient of determination (R^2) and Mean Euclidean Distance (MED) metrics. As illustrated in Figure 8, the comparative results show that the performance metrics for both LDS and rSLDS models improve with an increasing number of hidden dimensions. Notably, as the number of hidden dimensions increases, the R^2 values approach the expected value for Poisson-distributed observations with the rates predicted by the model, indicating that our largest models (24 hidden dimensions) capture almost as much variability as they possibly can for Poisson observations.

Our results on synthetic data, which were inspired by a popular non-linear computational model of perceptual decision-making, have shown that the rSLDS model was able to capture these nonlinearities and provide a better explanation for the data than a linear model. This led us to expect that the rSLDS model should also achieve superior predictive performance over the LDS model for the Kiani dataset. However, contrary to these expectations, rSLDS models with two modes did not significantly exceed the LDS models

in performance (see Figure 8). The same qualitative observation was made when using data from the C42 session (data not shown). To rule out that the missing benefit of additional modes was not just due to the rSLDS models not having converged to an optimal solution, we conducted further analyses.

3.2.2 Models for different task periods and model initialization

First, we fitted separate linear models based on data from the Decision Period and the Motor Period. To be able to visualize the results, we use models with two hidden dimensions. The resulting vector fields, plotted in a common state space (see Section 2) are shown on the left side of Figure 9. The dynamics during the decision period were found to be stable, while those in the motor period were characterized by instability. To maximize the chances of a 2-mode rSLDS model being able to capture this difference in dynamics for different task periods, we initialized the two modes with these two linear models. The boundary between the two modes was determined as follows: we calculated the latent trajectories for all trials during different task periods and identified the midpoints of these trajectories for each phase. The perpendicular bisector of the line segment connecting these two midpoints was then used as the boundary between the two modes (see middle panel of Figure 9).

The result after optimizing the model using the full trial data is shown in the top-right panel of Figure 9. The unstable dynamics have disappeared, and the result is virtually identical to a linear model fitted to the same data (bottom-right panel), again plotted in a common state space.

Considering that the Decision Period is substantially longer than the Motor Period, we questioned whether the Variational EM algorithm might favor a model that more accurately represents the decision period, potentially neglecting the motor period due to its relatively minor contribution in terms of data points. To address this, we constructed a dataset containing only the final part of the Decision Period and the entire Motor Period, equalizing their durations. Despite these modifications, our findings mirrored those presented in Figure 9. Consequently, we conclude that a piecewise linear model with a common observation equation cannot capture the observed change in dynamics from the decision period to the motor period. We will revisit this observation later in Discussion.

3.2.3 Decoding accuracy

Since the neural activity was recorded while monkeys were making perceptual decisions, we also wanted to assess how well an upcoming choice could be decoded from estimated hidden states, depending on which dynamical system model was used. Figure 8 had shown that adding hidden dimensions (within the studied range of up to 24 dimensions) continually increased the model's ability to capture and predict the neural activity. This prompted the question how the ability to decode the choice would depend on the dimensionality of the hidden state space. For each of the models, we first decoded the latent trajectories associated with each experimental trial. Subsequently, for each time step, we trained a Support Vector Machine (SVM) algorithm (Cortes and Vapnik, 1995) on the training dataset. After establishing

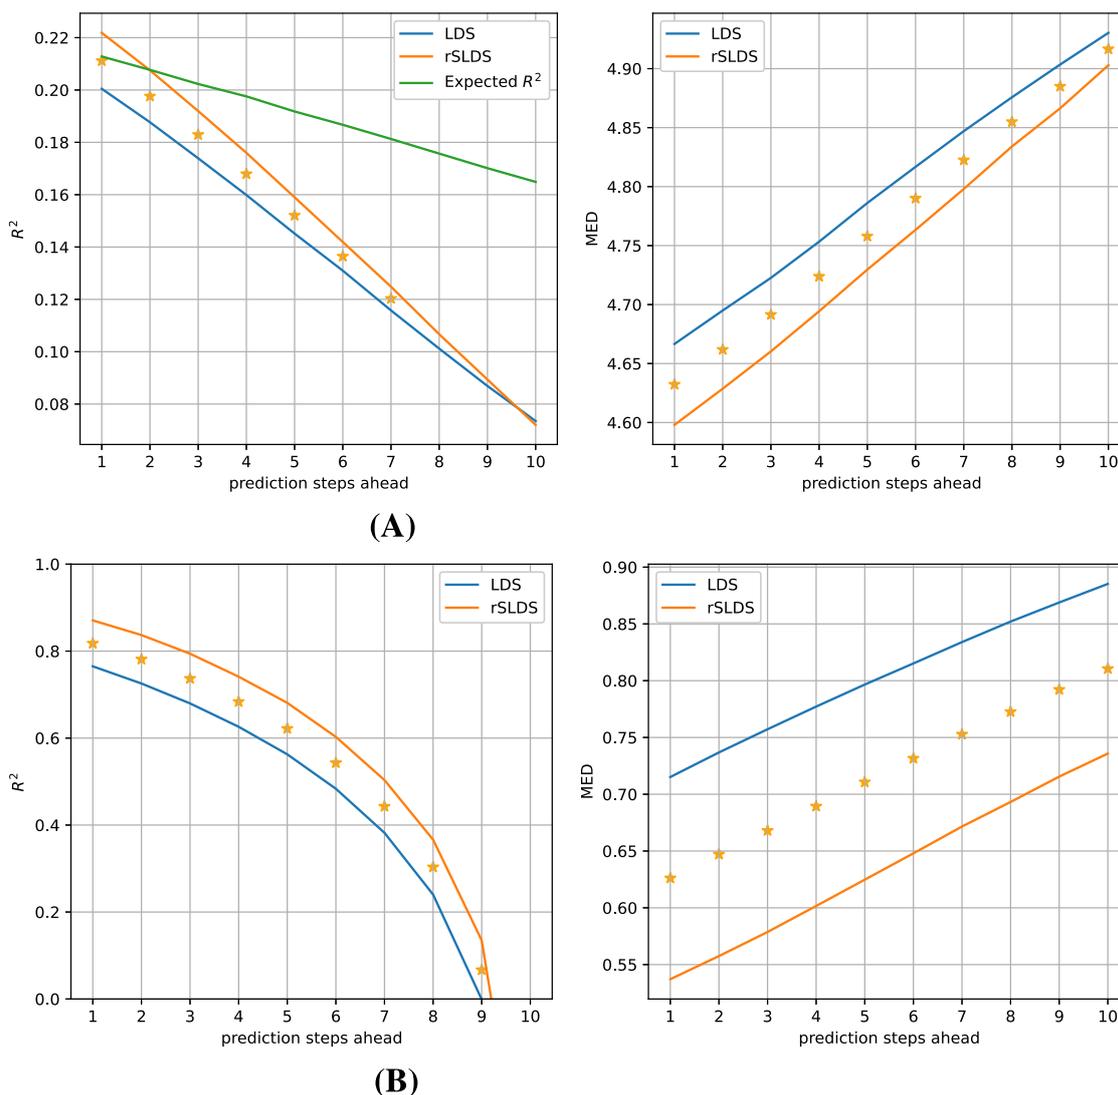


FIGURE 7

Predictive performance of the models for Poisson observations (observation dimensions = 200, average firing rate = 2). (A) Comparison in observation space using R^2 and MED. (B) Comparison in latent state space using R^2 and MED. The rSLDS model significantly outperformed the LDS model in both the observation and latent state spaces. Stars indicate a significant difference between both models at $p < 0.001$.

the SVM model, we evaluated its performance on a separate test set by measuring the accuracy of choice classification as a function of time.

Figure 10 illustrates that the rSLDS model's choice-decoding accuracy progressively increases across a broad range of hidden dimensions, highlighting the dependency of choice information preservation on the model's dimensionality. With a limited number of hidden dimensions (H), such as 2 to 4, the decoding accuracy remains limited and only deviates from chance toward the end of the decision trial, when the animal reports its choice, long after the motion viewing period. For 6 or 8 dimensions, the ability to decode choice starts earlier in the trial. Finally, for 10 dimensions or more an accuracy of about 80% is reached toward the end of the motion viewing period, which then further increases and approaches 100% toward the end of the trial.

3.2.4 High-dimensional dynamics embedding

While an expanded number of hidden dimensions enhances choice decoding accuracy and augments the model's predictive capabilities, it introduces challenges in visualizing the dynamics. To address this issue, we started with an 8-dimensional model and determined a two-dimensional projection of the hidden states that would maximize the separation between the states associated with both possible choices (see Section 2 for details). Figure 11A presents the average latent trajectories in this two-dimensional space for both possible choices. The trajectories start at the same location, but diverge as time progresses, and end at very different locations. To assess the efficacy of this 2D embedding approach, we recomputed the choice decoding accuracy on the embedded trajectories. The results, depicted in Figure 11B, closely resemble those derived from the original 8-dimensional trajectories, indicating that almost all

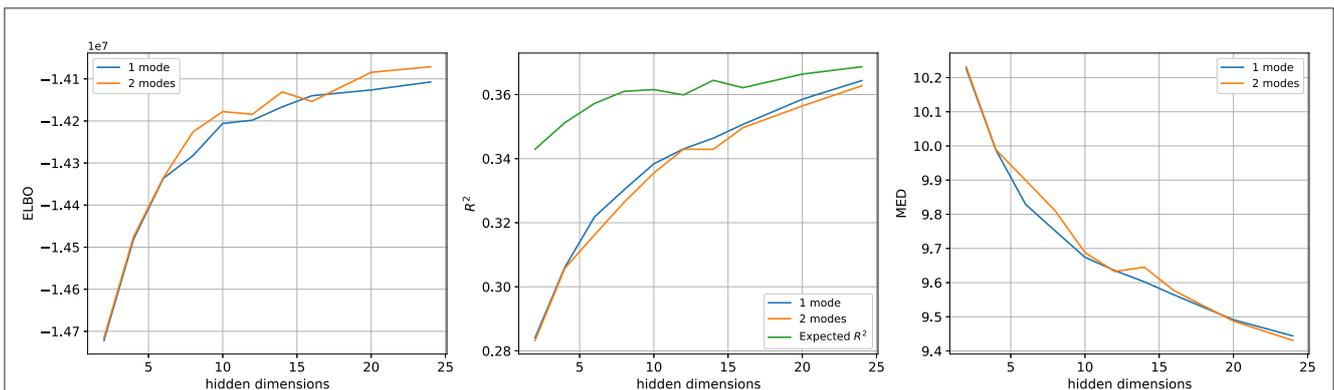


FIGURE 8
Performance comparison of 1-mode LDS and 2-mode rSLDS models across a range of hidden dimensions on the T33 dataset. The performance of both models improves with additional hidden dimensions and eventually tends to plateau. No clear advantage of the rSLDS model is observed across all three evaluation metrics compared to the LDS model.

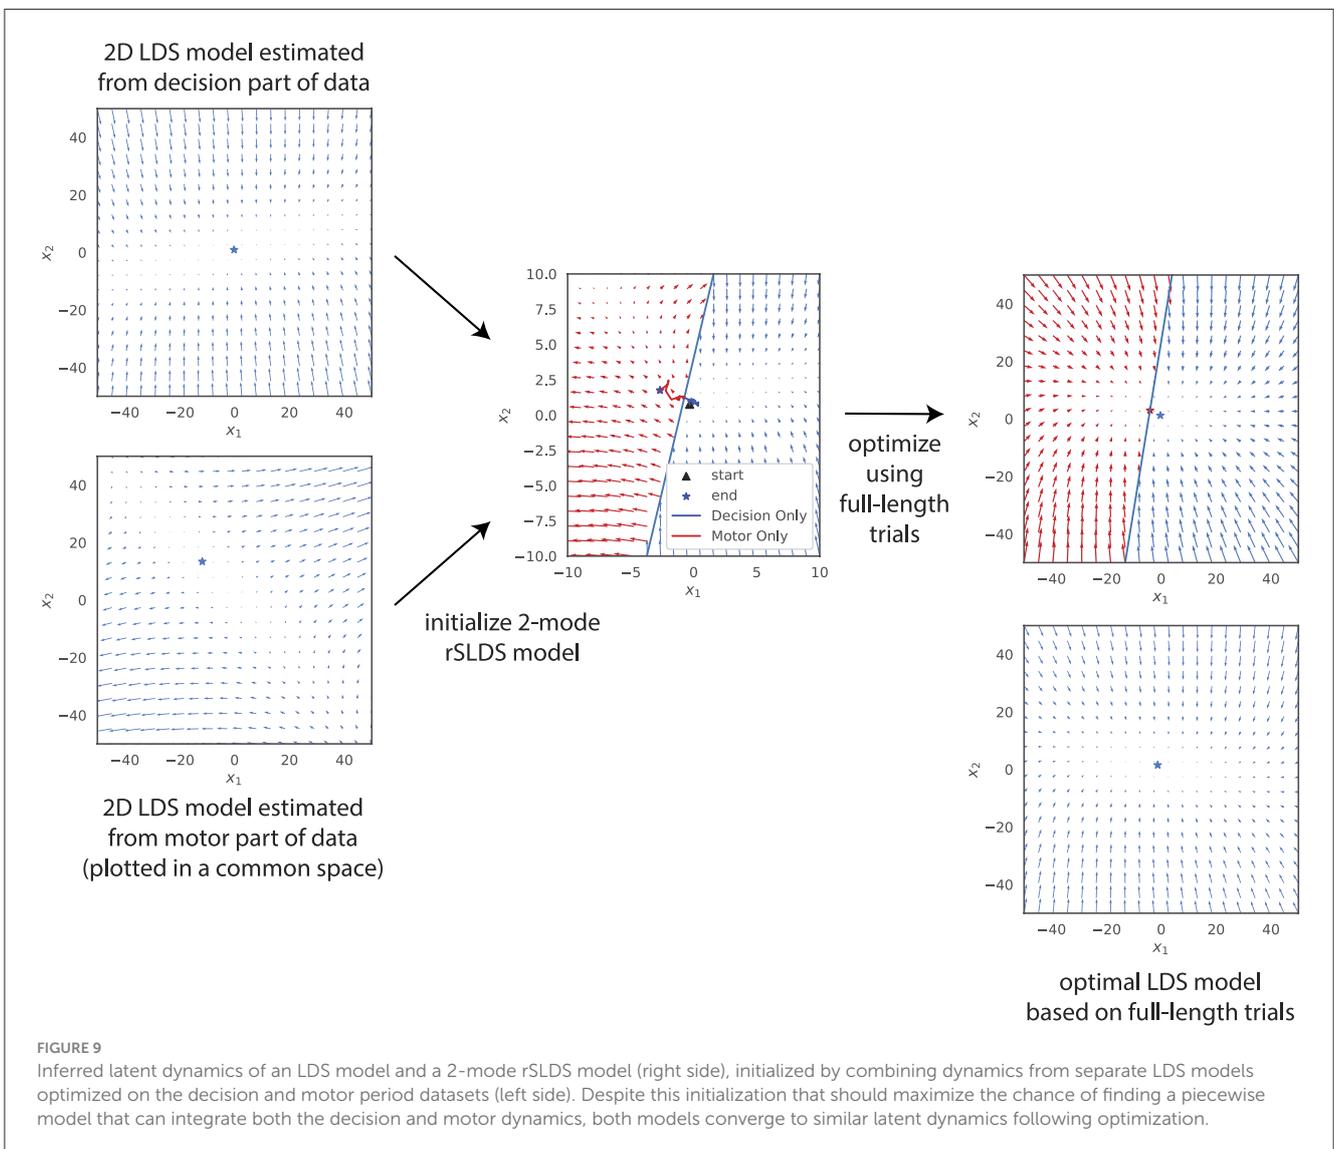
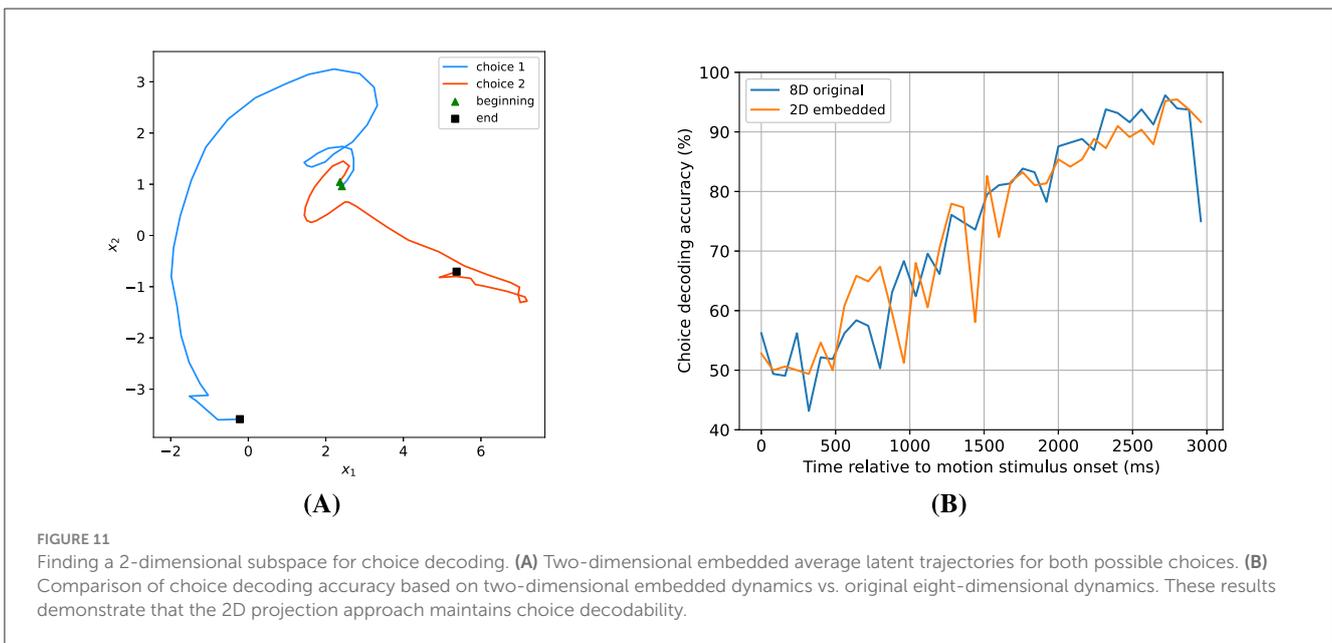
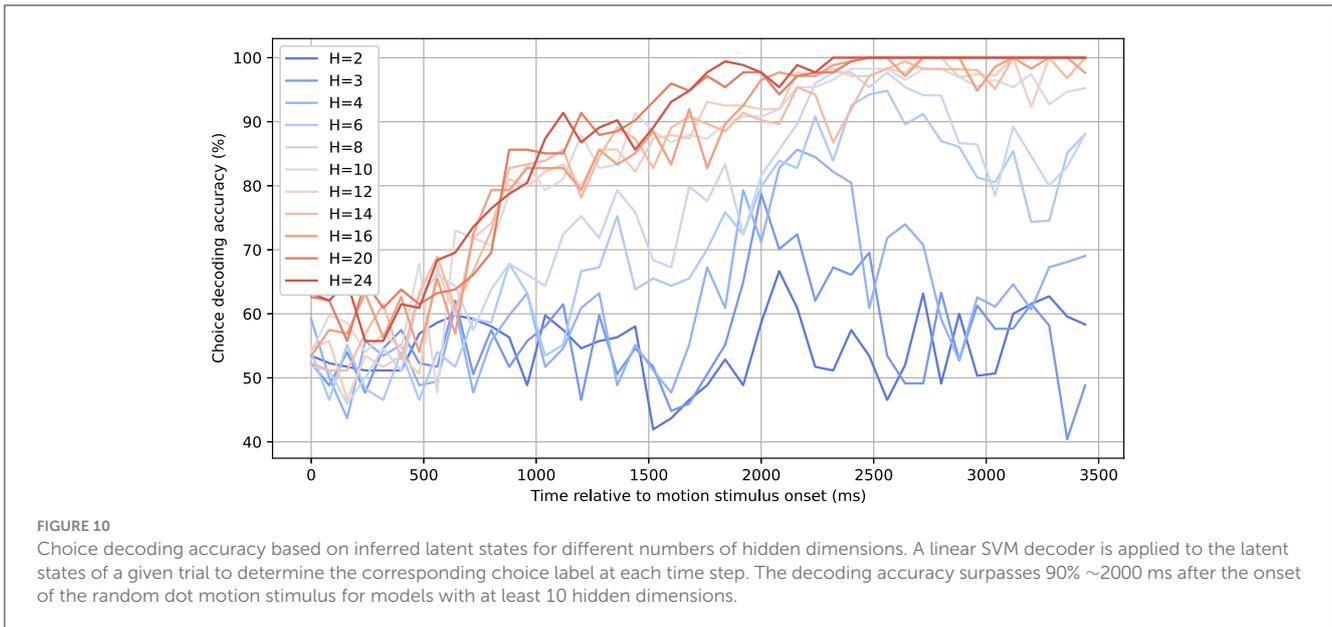


FIGURE 9
Inferred latent dynamics of an LDS model and a 2-mode rSLDS model (right side), initialized by combining dynamics from separate LDS models optimized on the decision and motor period datasets (left side). Despite this initialization that should maximize the chance of finding a piecewise model that can integrate both the decision and motor dynamics, both models converge to similar latent dynamics following optimization.



choice information is contained in a particular two-dimensional subspace of the higher-dimensional hidden state space.

4 Discussion

4.1 Successful estimation of piecewise-linear models from synthetic neural data and superior prediction performance

Using synthetic neural data generated by a piecewise-linear adaptation of the computational model proposed by Wong and Wang (2006) for perceptual decision-making, we were able to demonstrate that the non-linear structure of the dynamical system

could be successfully estimated from the data. Not surprisingly, the quality of the estimated models depended on how many signals were observed and the amount of noise on these observations. For continuous observations, the estimated models tended to be of good quality, as long as the number of observed signals matched at least the dimensionality of the hidden state space and the observation noise was not too large. Estimating the model structure from Poisson-distributed single-unit spiking data was more challenging, but still possible. These observations are inherently more noisy, due to the stochastic nature of the Poisson process, and potentially quite sparse, if the firing rates of the observed neurons are rather low. Good models could still be obtained as long as the number of observed neurons was at least an order of magnitude larger than the dimensionality of the hidden state space and the average expected spike count per time bin

wasn't much lower than 0.5. Since the model estimation based on Variational Expectation Maximization can get stuck in suboptimal solutions, we recommend fitting a larger number of models and selecting the best one.

Perhaps more importantly, the estimated piecewise-linear dynamical models were able to make significantly more accurate predictions for how the state of the system would evolve in the near future than just a linear dynamical system model. This could give these piecewise-linear models a significant advantage when considering applications like closed-loop neural modulation, where state predictions can be used for model-based control.

Being able to estimate these models not only from continuous observations like field potentials or smoothed firing rates, which requires sacrificing temporal resolution, but also from Poisson-distributed observations, i.e., sparser single-unit spiking data, is a clear advantage of this type of model. Most alternative dynamic system models require continuous, normally-distributed observations, although there are some alternative models that can also be estimated from Poisson spike trains (e.g., Pandarinath et al., 2018; Kim et al., 2021; Schimel et al., 2022).

4.2 Piecewise-linear models did not outperform their linear counterparts when applying the approach to a publicly available dataset recorded from prefrontal cortex of monkeys performing perceptual decisions

To our surprise, we did not see the same benefit when applying our approach to a publicly available dataset of about 200 simultaneously recorded neurons in prefrontal cortex of monkeys performing perceptual decisions. Standard linear dynamical system models and piecewise-linear models provided virtually identical ability to explain the neural data and prediction performance. This does not imply that the same would have to be true for other datasets obtained from other brain areas or when performing different tasks, but it certainly warrants some discussion.

One possibility would be that, for some reason, the structure of the dynamical system, although compatible with our piecewise-linear model, could not be estimated from the available data. Given that the number of available neurons was at least an order of magnitude larger than the dimensionality of the hidden state space and given that we chose a time bin width of 80 ms, which resulted in an average expected spike count per bin of ~ 0.6 , a value that allowed robust model estimation when using our synthetic data, this seems unlikely.

Another possibility would be that the neural dynamics are not sufficiently non-linear to warrant fitting a non-linear model. Some recent observations in the literature provide some support for this idea. For example, Nozari et al. (2024) analyzed resting-state neural activity (both fMRI and intracranial EEG data) and fitted a wide range of both linear and non-linear dynamical models. The model best capturing the data turned out to be a linear autoregressive model, suggesting that the dynamics were not sufficiently non-linear to give a non-linear model an advantage. Galgali et al. (2023) analyzed the same dataset by Kiani et al. that was also used

for our analysis here and made a very similar observation: The neural dynamics during a particular task period, like the decision formation period and the saccadic eye movement period, were well-described by a linear dynamical model.

Finally, it could be that the dynamical system, although non-linear, is incompatible with certain assumptions made by our model. Why could the different linear dynamics during the different task periods not be captured by a piecewise linear model with multiple modes? The rSLDS model we have been using is designed to approximate one larger non-linear flow field through breaking it down into multiple locally linear modes or regions in state space. There is still one common observation equation that applies to all modes/regions. The model switches to a different mode due to the location in state space changing, and the observations (firing pattern of the neurons) have to change accordingly. In contrast, the change in linear dynamics when transitioning between different task periods might happen without a major change in the firing pattern/location in state space and might therefore be better described by a time-variant linear system rather than a non-linear system. Future research will have to show whether time-variant models can provide a better description of the neural dynamics.

We found that, up to the explored maximum of 24 dimensions, the quality of our LDS models kept increasing when adding more hidden dimensions and that a model with a minimum of 8 to 10 dimensions had to be chosen to be able to decode the animal's choice reliably from the hidden state throughout the trial. At first glance, this might seem at odds with previous observations in the literature that the neural dynamics underlying perceptual decisions appear to be surprisingly low-dimensional. However, we were also able to demonstrate that the choice information largely appears to be contained in a two-dimensional subspace of the higher-dimensional hidden state space. This suggests that the choice encoding itself is still low-dimensional and that the additional dimensions are required to capture variability in the neural state due to factors other than choice.

4.3 Conclusions

We have set out to evaluate the suitability of piecewise-linear dynamical system models for capturing cognitive neural dynamics. We were able to demonstrate that piecewise-linear dynamical system (rSLDS) models could be successfully estimated from synthetic neural data generated by a non-linear computational model of perceptual decision-making. Importantly, the rSLDS model significantly outperformed a standard LDS model in terms of predicting the state of the dynamical system in the near future. Given the straightforward mathematical structure of a piecewise-linear dynamical system model, it would be particularly useful for model-based control applications, e.g., closed-loop neuromodulation for the treatment of cognitive deficits resulting from neurological or psychiatric disorders, when nonlinearities of the dynamics need to be captured.

Much to our surprise, when applying our modeling approach to a publicly available dataset of ~ 200 simultaneously recorded neurons in prefrontal cortex of monkeys performing perceptual decisions, we did not find a significant advantage of rSLDS

models over standard LDS models, suggesting that the neural dynamics were not sufficiently non-linear at any given time to warrant the use of a non-linear model. Future research will have to show whether this finding applies to a wide range of cognitive neural datasets, including different types of neural activity (single-unit spiking, multi-unit activity, LFPs, iEEG) and recordings from different brain areas and during different tasks, or whether it was due to specific features of the particular dataset that was used for our analysis, and other datasets will show a benefit of using piecewise-linear dynamic models, similar to our synthetic data.

Insights gained from the application of piecewise-linear dynamic models to scientific and engineering systems suggest that these models can offer clear advantages over linear dynamical systems, particularly when the underlying dynamics exhibit multiple equilibria (as seen in our synthetic dataset), non-linear switching behavior, threshold effects, or operate across distinct regimes. Representative examples include: (1) gene regulatory networks, where gene expression is switched “off” below a certain concentration threshold of a transcription factor and “on” above it (Gouzé, 1998); (2) neuronal spiking dynamics, where the membrane potential increases approximately linearly until it reaches a threshold and then resets, producing a spike (Izhikevich, 2003); (3) legged locomotion, in which robots experience hybrid dynamics due to alternating stance and swing phases during ground contact (Lygeros et al., 2003). In each of these cases, piecewise-linear models not only better capture the essential dynamics but have also been shown to outperform linear models in both accuracy and control performance.

Author’s note

A preprint of this manuscript was posted on bioRxiv (Wu et al., 2025).

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found at: <https://www.cns.nyu.edu/kianilab/Datasets.html>.

References

- Åström, K. J., and Murray, R. (2021). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ: Princeton University Press.
- Abbaspourazad, H., Erturk, E., Pesaran, B., and Shanechi, M. M. (2024). Dynamical flexible inference of nonlinear latent factors and structures in neural population activity. *Nat. Biomed. Eng.* 8, 85–108. doi: 10.1038/s41551-023-01106-1
- Amunts, K., Ebell, C., Müller, J., Telefont, M., Knoll, A., and Lippert, T. (2016). The human brain project: creating a european research infrastructure to decode the human brain. *Neuron* 92, 574–581. doi: 10.1016/j.neuron.2016.10.046
- Borrelli, F., Bemporad, A., and Morari, M. (2017). *Predictive Control for Linear and Hybrid Systems*. Cambridge: Cambridge University Press. doi: 10.1017/9781139061759
- Cameron, A. C., and Windmeijer, F. A. (1996). R-squared measures for count data regression models with applications to health-care utilization. *J. Bus. Econ. Stat.* 14, 209–220. doi: 10.1080/07350015.1996.10524648
- Cortes, C., and Vapnik, V. (1995). Support-vector networks. *Mach. Learn.* 20, 273–297. doi: 10.1007/BF00994018
- Doyle, J. C., Francis, B. A., and Tannenbaum, A. R. (2013). *Feedback Control Theory*. Chelmsford, MA: Courier Corporation.
- Galgali, A. R., Sahani, M., and Mante, V. (2023). Residual dynamics resolves recurrent contributions to neural computation. *Nat. Neurosci.* 26, 326–338. doi: 10.1038/s41593-022-01230-2
- Gouzé, J.-L. (1998). Positive and negative circuits in dynamical systems. *J. Biol. Sys.* 6, 11–15. doi: 10.1142/S0218339098000054
- Gu, S., Pasqualetti, F., Cieslak, M., Telesford, Q. K., Yu, A. B., Kahn, A. E., et al. (2015). Controllability of structural brain networks. *Nat. Commun.* 6:8414. doi: 10.1038/ncomms9414
- Hespanha, J. P. (2018). *Linear Systems Theory*. Princeton, NJ: Princeton University Press. doi: 10.23943/9781400890088
- Hurwitz, C., Kudryashova, N., Onken, A., and Hennig, M. H. (2021). Building population models for large-scale neural recordings: opportunities and pitfalls. *Curr. Opin. Neurobiol.* 70, 64–73. doi: 10.1016/j.conb.2021.07.003

Author contributions

JW: Formal analysis, Software, Visualization, Writing – original draft. BA: Investigation, Writing – review & editing. ZK: Conceptualization, Funding acquisition, Project administration, Supervision, Writing – review & editing. JD: Conceptualization, Funding acquisition, Methodology, Project administration, Supervision, Writing – review & editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This work was supported by the National Science Foundation (award 2024526).

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Netw.* 14, 1569–1572. doi: 10.1109/TNN.2003.820440
- John, Y. J., Sawyer, K. S., Srinivasan, K., Müller, E. J., Munn, B. R., and Shine, J. M. (2022). It's about time: linking dynamical systems with human neuroimaging to understand the brain. *Netw. Neurosci.* 6, 960–979. doi: 10.1162/netn_a_00230
- Jones, T., and Meyer, M. J. (2024). A New (old) Goodness of Fit Metric for Multidimensional Outcomes. Available online at: <https://github.com/TommyJones/r-squared-short-paper/blob/main/r-squared-short.pdf> (accessed May 1, 2025).
- Kao, J. C., Nuyujukian, P., Ryu, S. I., Churchland, M. M., Cunningham, J. P., and Shenoy, K. V. (2015). Single-trial dynamics of motor cortex and their applications to brain-machine interfaces. *Nat. Commun.* 6:7759. doi: 10.1038/ncomms8759
- Kiani, R., Cueva, C. J., Reppas, J. B., and Newsome, W. T. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biol.* 24, 1542–1547. doi: 10.1016/j.cub.2014.05.049
- Kim, J. Z., and Bassett, D. S. (2020). “Linear dynamics and control of brain networks,” in *Neural Engineering* (Springer, Cham), 497–518. doi: 10.1007/978-3-030-43395-6_17
- Kim, T. D., Luo, T. Z., Pillow, J. W., and Brody, C. D. (2021). “Inferring latent dynamics underlying neural population activity via neural differential equations,” in *Proceedings of the 38th International Conference on Machine Learning*, 139.
- Kriegeskorte, N., and Douglas, P. K. (2018). Cognitive computational neuroscience. *Nat. Neurosci.* 21, 1148–1160. doi: 10.1038/s41593-018-0210-5
- Le Quang, T. (2013). *Piecewise Affine Dynamical Systems: Well-Posedness, Controllability and Stabilizability*. Groningen: University of Groningen.
- Lee, W.-H., Karpowicz, B. M., Pandarinath, C., and Rouse, A. G. (2024). Identifying distinct neural features between the initial and corrective phases of precise reaching using autolufs. *J. Neurosci.* 44:e1224232024. doi: 10.1523/JNEUROSCI.1224-23.2024
- Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. (2017). “Bayesian learning and inference in recurrent switching linear dynamical systems,” in *Artificial Intelligence and Statistics* (Ft. Lauderdale, FL: PMLR), 914–922.
- Lygeros, J., Johansson, K. H., Simic, S. N., Zhang, J., and Sastry, S. S. (2003). Dynamical properties of hybrid automata. *IEEE Trans. Automat. Contr.* 48, 2–17. doi: 10.1109/TAC.2002.806650
- Moxon, K., Kong, Z., and Ditterich, J. (2023). “Brain-machine interfaces: from restoring sensorimotor control to augmenting cognition,” in *Handbook of Neuroengineering* (Springer: New York), 1343–1380. doi: 10.1007/978-981-16-5540-1_36
- Nozari, E., Bertolero, M. A., Stiso, J., Caciagli, L., Cornblath, E. J., He, X., et al. (2024). Macroscopic resting-state brain dynamics are best described by linear models. *Nat. Biomed. Eng.* 8, 68–84. doi: 10.1038/s41551-023-01117-y
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* 15, 805–815. doi: 10.1038/s41592-018-0109-9
- Rabinovich, M. I., Varona, P., Selverston, A. I., and Abarbanel, H. D. (2006). Dynamical principles in neuroscience. *Rev. Mod. Phys.* 78, 1213–1265. doi: 10.1103/RevModPhys.78.1213
- Rao, V. R., Sellers, K. K., Wallace, D. L., Lee, M. B., Bijanzadeh, M., Sani, O. G., et al. (2018). Direct electrical stimulation of lateral orbitofrontal cortex acutely improves mood in individuals with symptoms of depression. *Current Biol.* 28, 3893–3902. doi: 10.1016/j.cub.2018.10.026
- Ruff, D. A., Ni, A. M., and Cohen, M. R. (2018). Cognition as a window into neuronal population space. *Annu. Rev. Neurosci.* 41, 77–97. doi: 10.1146/annurev-neuro-080317-061936
- Sani, O. G., Yang, Y., Lee, M. B., Dawes, H. E., Chang, E. F., and Shanechi, M. M. (2018). Mood variations decoded from multi-site intracranial human brain activity. *Nat. Biotechnol.* 36, 954–961. doi: 10.1038/nbt.4200
- Scangos, K. W., Khambhati, A. N., Daly, P. M., Makhoul, G. S., Sugrue, L. P., Zamanian, H., et al. (2021a). Closed-loop neuromodulation in an individual with treatment-resistant depression. *Nat. Med.* 27, 1696–1700. doi: 10.1038/s41591-021-01480-w
- Scangos, K. W., Makhoul, G. S., Sugrue, L. P., Chang, E. F., and Krystal, A. D. (2021b). State-dependent responses to intracranial brain stimulation in a patient with depression. *Nat. Med.* 27, 229–231. doi: 10.1038/s41591-020-01175-8
- Scangos, K. W., State, M. W., Miller, A. H., Baker, J. T., and Williams, L. M. (2023). New and emerging approaches to treat psychiatric disorders. *Nat. Med.* 29, 317–333. doi: 10.1038/s41591-022-02197-0
- Shimel, M., Kao, T.-C., Jensen, K. T., and Hennequin, G. (2022). “ilqr-vae: control-based learning of input-driven dynamics with applications to neural data,” in *The Tenth International Conference on Learning Representations*. doi: 10.1101/2021.10.07.463540
- Schulz, M.-A., Yeo, B. T., Vogelstein, J. T., Mourao-Miranada, J., Kather, J. N., Kording, K., et al. (2020). Different scaling of linear models and deep learning in ukbiobank brain images versus machine-learning datasets. *Nat. Commun.* 11:4238. doi: 10.1038/s41467-020-18037-z
- Shanechi, M. M. (2019). Brain-machine interfaces from motor to mood. *Nat. Neurosci.* 22, 1554–1564. doi: 10.1038/s41593-019-0488-y
- Srivastava, P., Nozari, E., Kim, J. Z., Ju, H., Zhou, D., Becker, C., et al. (2020). Models of communication and control for brain networks: distinctions, convergence, and future outlook. *Netw. Neurosci.* 4, 1122–1159. doi: 10.1162/netn_a_00158
- Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328. doi: 10.1523/JNEUROSCI.3733-05.2006
- Woolson, R. F. (2005). *Wilcoxon Signed-Rank Test*. Encyclopedia of Biostatistics (Hoboken, NJ: Wiley), 8. doi: 10.1002/0470011815.b2a15177
- Wu, J., Asamoah, B., Kong, Z., and Ditterich, J. (2025). Exploring the suitability of piecewise-linear dynamical system models for cognitive neural dynamics. *bioRxiv*. doi: 10.1101/2025.01.14.633062
- Xu, J., and Xie, L. (2014). *Control and Estimation of Piecewise Affine Systems*. Amsterdam: Elsevier. doi: 10.1533/9781782421627.17
- Yang, Y., Connolly, A. T., and Shanechi, M. M. (2018). A control-theoretic system identification framework and a real-time closed-loop clinical simulation testbed for electrical brain stimulation. *J. Neural Eng.* 15:66007. doi: 10.1088/1741-2552/aad1a8
- Yang, Y., Sani, O. G., Chang, E. F., and Shanechi, M. M. (2019). Dynamic network modeling and dimensionality reduction for human ecog activity. *J. Neural Eng.* 16:56014. doi: 10.1088/1741-2552/ab2214