Targeted Neural Dynamical Modeling

Anonymous Author(s)
Affiliation
Address
email

Abstract

Latent dynamics models have emerged as powerful tools for modeling and interpreting neural population activity. Recently, there has been a focus on incorporating simultaneously measured behaviour into these models to further disentangle sources of neural variability in their latent space. These approaches, however, are limited in their ability to capture the underlying neural dynamics (e.g. linear) and in their ability to relate the learned dynamics back to the observed behaviour (e.g. no time lag). To this end, we introduce Targeted Neural Dynamical Modeling (TNDM), a nonlinear state-space model that jointly models the neural activity and external behavioural variables. TNDM decomposes neural dynamics into behaviourally relevant and behaviourally irrelevant dynamics; the relevant dynamics are used to reconstruct the behaviour through a flexible linear decoder and both sets of dynamics are used to reconstruct the neural activity through a linear decoder with no time lag. We implement TNDM as a sequential variational autoencoder and validate it on recordings taken from the premotor and motor cortex of a monkey performing a center-out reaching task. We show that TNDM is able to learn low-dimensional latent dynamics that are highly predictive of behaviour without sacrificing its fit to the neural data.

1 Introduction

Recent progress in high-density, microelectrode array technology now allows for recording from hundreds to thousands of neurons with the precision of single spikes [10]. Despite the apparent high dimensionality of these datasets, neural activity is often surprisingly well-explained by low-dimensional latent dynamics [3, 21, 5, 7]. Extracting these dynamics from single trials is crucial for understanding how neural activity relates to a behavioural task or stimulus [15].

Latent variable models (LVMs) are a natural choice for capturing low-dimensional structure from neural activity as they can learn to map a few latent variables to arbitrarily complicated response structure in the activity. Already, there exist a number of LVMs that have been successfully applied to neural data ranging from simple non-temporal models such as principal components analysis (PCA) [4] to complex state-space models such as LFADS [15]. In these models, the goal is to learn a set of latent factors that best explain neural variability. As such, there is no guarantee that the different sources of variability present in the population activity will be disentangled in the latent space (e.g. behaviour, arousal, thirst, etc.) [22, 9].

To better partition sources of neural variability in the latent space, some LVMs have been developed that incorporate an external behaviour into the generative process [12, 14, 17, 28]. These models, however, do not model temporal dependencies between the latent states which is essential for modeling neural activity. Recently, a novel state-space model termed preferential subspace identification (PSID) was developed that jointly models neural activity and behaviour with a shared set of dynamics [22]. When applied to neural activity recorded in the premotor cortex (PMd) and primary motor cortex (M1) of a monkey during a 3D reaching task, PSID was able to extract latent dynamics that were more

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

predictive of behaviour than those learned by other approaches. Despite the strength and simplicity of this approach, it suffers from two main drawbacks. First, PSID is a linear state-space model and cannot capture nonlinear dynamics which are thought to underlie phenomena such as rhythmic motor patterns [20, 8] or decision making [19]. Second, PSID assumes that behaviourally relevant dynamics explain both the neural activity and behaviour with no time lag. This limits the ability of PSID to capture more complex temporal relationships between the latent dynamics and the behaviour.

In this work, we introduce Targeted Neural Dynamical Modeling (TNDM), a nonlinear state-space model that jointly models neural activity and behaviour. Similarly to PSID, TNDM decomposes neural activity into behaviourally relevant and behaviourally irrelevant dynamics and uses the relevant dynamics to reconstruct the behaviour and both sets of dynamics to reconstruct the neural activity. Unlike PSID, TNDM does not constrain the latent dynamics at each time step to explain behaviour at each time step and instead allows for any linear relationship (constrained to be causal in time) between the relevant dynamics and the behaviour of interest. We further encourage partitioning of the latent dynamics by imposing a disentanglement penalty on the distributions of the initial conditions of the relevant and irrelevant dynamics. To perform efficient inference of the underlying nonlinear dynamics, TNDM is implemented as a sequential variational autoencoder (VAE) [11, 26]. We compare TNDM to PSID and to LFADS, a nonlinear state-space model that just models neural activity, to illustrate that TNDM extracts more behaviourally relevant dynamics without sacrificing its fit to the neural data. We validate TNDM on neural population recordings taken from the premotor and motor cortex of a monkey during a center-out reaching task. In this analysis, we find that the behaviourally relevant dynamics revealed by TNDM are lower dimensional than those of other methods.

2 Background/Related work

Notatation. Let $x \in \mathbb{N}^{N \times T}$ be the observed spike counts and let $y \in \mathbb{R}^{B \times T}$ be the observed behaviour during a single-trial.[1] We define the unobserved latent dynamics in a single trial as $z \in \mathbb{R}^{M \times T}$ where $M < N$. For TNDM, as with PSID, it is important to distinguish between behaviourally relevant z_r and behaviourally irrelevant z_i latent dynamics. The behaviourally relevant latent dynamics z_r summarize the variability in the neural activity associated with the observed behaviour while the behaviourally irrelevant latent dynamics z_i explain everything else in the neural data (Figure 1a). We assume that each of the unobserved, single-trial dynamics can be partitioned into these relevant and irrelevant dynamics $z \coloneqq \{z_r, z_i\}$.

State-space models for neural data. There are a number of state-space models that have been developed and applied to neural population activity. The expressivity of these models range from simple linear dynamical models [24, 2, 18] to more complex nonlinear models where the latent dynamics are parameterized by recurrent neural networks (RNNs) [15, 23]. For this work, there are two state-space models that are most relevant: LFADS and PSID.

LFADS, or latent factor analysis via dynamical systems, is a state-of-the-art nonlinear state-space model for neural data. In LFADS, the latent dynamics are generated by sampling high-dimensional initial conditions g_0 from some distribution p_{g_0} and then evolving g_0 with a deterministic RNN f_θ. A linear mapping W_z is then applied to the high-dimensional dynamics g_t to transform them into the low-dimensional dynamical factors z. These dynamical factors are transformed into spike counts by mapping each time point to a rate parameter r of a Poisson distribution using a weight matrix W_r followed by an exponential nonlinearity. The generative process is then: $g_0 \sim p_{g_0}, g_t = f_\theta(g_{t-1}), z_t = W_z(g_t), r_t = \exp(W_r(z_t)), x_t \sim \text{Poisson}(x_t|r_t)$. The initial conditions g_0 are inferred from x with an RNN encoder network q_ϕ. The model is trained by optimizing the evidence lower-bound (ELBO) of the marginal log-likelihood. While LFADS provides an excellent fit to the neural data, it inevitably mixes different sources of neural variability in the latent dynamics z as there is no constraints imposed to disentangle these dynamics. Extracting disentangled representation is especially challenging with deep generative models and is still an active area of research [13].

PSID is a linear dynamical model that aims to partition different sources of neural variability in the learned latent dynamics. In PSID, the latent dynamics are partitioned into behaviourally relevant and behaviourally irrelevant dynamics $z \coloneqq \{z_i, z_r\}$. The dynamical evolution of z is defined by

[1] In this work, we assume that behaviour is temporal and has the same time length as recorded neural activity (e.g. hand position). TNDM can be extended to discrete/non-temporal behaviours (e.g. reach direction).

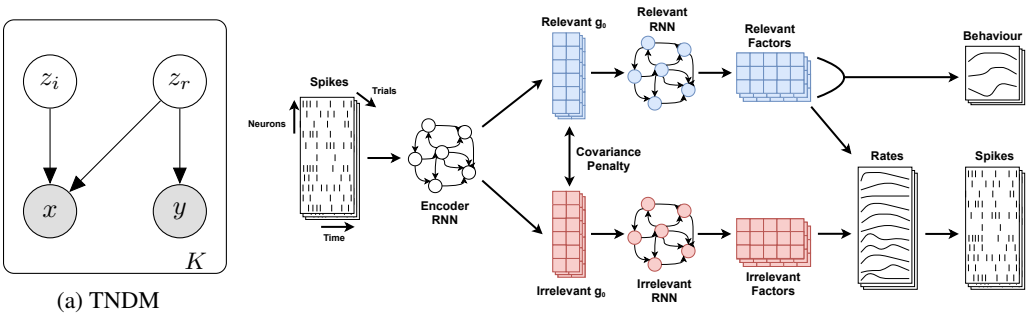

(a) TNDM graphical model

(b) TNDM architecture

Figure 1: (a) The latent space of TNDM is partioned into irrelevant z_i and relevant z_r dynamics that give rise to the jointly observed neural activity and behaviour (we assume there are K trials in the dataset). (b) TNDM utilizes a sequential variational autoencoding approach to ammortize inference of the relevant and irrelevant initial conditions g_0. The initial conditions are passed through two separate RNNs that generate relevant and irrelevant dynamical factors which are used to reconstruct the neural activity and behaviour. The behaviour is reconstructed from the relevant factors using a flexible linear decoder which can capture complex temporal relationships.

a transition matrix A which is block diagonal (i.e. it evolves the relevant and irrelevant dynamics separately) along with Gaussian noise term w_z. z is transformed into the observed firing rates x by mapping each time point to the mean of a Gaussian distribution with a weight matrix W_x and noise term w_x. The behaviourally relevant dynamics z_r at each time point are transformed into the observed behaviour y with a weight matrix W_y. The state space model for PSID is then defined as: $z_t \sim \mathcal{N}(z_t|A(z_{t-1}), w_z), x_t \sim \mathcal{N}(x|W_x(z_t), w_x), y_t = W_y(z_r)$. PSID uses a novel two-stage identification approach to extract z_r and z_i separately. In comparison to LFADS, PSID was shown to extract more behaviourally relevant latent dynamics from long neural recordings which could better predict the behaviour of interest. Despite the analytical simplicity of this approach, it suffers from a few drawbacks. First, it can only model linear dynamics. Second, the relevant dynamics at each time step z_{rt} must be mapped one-to-one to the behaviour y_t. This imposes a strong structural constraint on the relevant dynamics which may hamper their ability to explain neural variability.

3 Model

In this work, we introduce Targeted Neural Dynamical Modeling (TNDM). TNDM is a nonlinear state-space model which jointly models neural activity and an observed behaviour. Crucially, TNDM learns to reconstruct both the population activity and behaviour by disentangling the behaviourally relevant and behaviourally irrelevant dynamics that underlie the neural activity.

Generative model. A diagram of TNDM's graphical model is shown in Figure 1a. Simarily to PSID, we assume that the observed neural activity x and behaviour y in each trial is generated by two sets of latent dynamics z_i and z_r.

The generative process of TNDM is defined below:

$$\begin{aligned} &g_{i0} \sim p_{g_{i0}}, g_{r0} \sim p_{g_{r0}},\ g_{it} = f_{\theta i}(g_{it-1}),\ g_{rt} = f_{\theta r}(g_{rt-1}) \\ &z_{it} = W_{iz}(g_{it}),\ z_{rt} = W_{rz}(g_{rt}),\ r_t = \exp(W_r((z_{it}, z_{rt}))) \\ &x_t \sim \text{Poisson}(x_t|r_t),\ y \sim \mathcal{N}(y|C_y(z_r), I) \end{aligned} \tag{1}$$

In the above equation, $p_{g_{i0}}$ and $p_{g_{r0}}$ are the distributions over the initial conditions for the behaviourally irrelevant and behaviourally relevant dynamics, respectively. Similarly to LFADS, we parameterize the nonlinear transition functions $f_{\theta i}$ and $f_{\theta r}$ using RNNs[2]. W_{iz} and W_{rz} are weight matrices that map the high-dimensional dynamics g_{it} and g_{rt} to the low-dimensional relevant and

[2]To implement TNDM, we primarily adapt the original Tensorflow [1] implementation of LFADS from `https://github.com/lfads/models` (Apache License 2.0).

irrelevant dynamics at each time step. We assume that the observed neural activity and behaviour are conditionally independent given the relevant latent dynamics. An important feature of our generative model is that although neural activity is reconstructed from the latent dynamics at each time step (i.e. one-to-one), we let the relevant factors reconstruct the behaviour with a more flexible linear decoder C_y that allows for time lags. We discuss the behaviour decoding in the following paragraph.

Behaviour Decoding. As mentioned above, PSID utilizes a linear weight matrix that maps the relevant latent dynamics at each time step to the behaviour at each time step, i.e $y_t = W_y(z_{r_t})$. This parameterization does not allow for modeling any long-term dependencies/correlations, therefore, it severely limits the ability of the relevant dynamics to simultaneously explain neural activity and behaviour . To demonstrate the drawbacks of this one-to-one behaviour decoder, we show that while training TNDM using this decoder leads to accurate behaviour prediction, the reconstruction of the neural activity noticeably decreases. In fact, this issue gets exacerbated in models with non-linear dynamics as the expressivity of the underlying RNNs, along with the inflexible one-to-one behaviour decoding, leads the relevant dynamics to simply learn to replicate the behaviour of interest. These results are summarized in Supplement 1.

To overcome the limitations of this one-to-one approach, we instead allow the relevant latent dynamics to reconstruct the behaviour through any learned linear casual relationship. To this end, we introduce a linear weight matrix C_y with dimensionality $n_{z_r}T \times BT$ where n_{z_r} is the number of relevant dynamics, B is the number of behaviour dimensions, and T is the time length of a single trial. To transform the relevant factors z_r into behaviours using C_y, we concatenate each dimension of z_r in time to form a 1D vector Z_r with length $n_{z_r}T$ and then perform the operation $Y = C_y Z_r$ where Y is the resulting concatenated behaviour. As Y is a 1D vector with length BT, we can reshape Y to recover the reconstructed behaviour $\hat{y}$. Importantly, we do not allow acausal connections in C_y (i.e. the lower triangular components of each of the dynamics to behaviour blocks are set to zero during training). For an example weight matrix C_y, see Figure 4a. In comparison to a simple one-to-one temporal mapping, we find that our flexible, causal linear decoder allows the relevant latent dynamics to both reconstruct the measured behaviour and capture neural variability.

Inference To extract the latent dynamics z_r and z_i from the neural activity x, we learn to approximate the posterior over the initial conditions of the dynamics g_{r0} and g_{i0} and then we learn the RNN mapping from the initial conditions to the latent dynamics as together they deterministically define z_r and z_i. To approximate the true posterior $p(g_{r0}, g_{i0}|x, y)$, we implement TNDM as a sequential VAE and define our variational posterior as the product of two independent multivariate Gaussian distributions with diagonal covariances. The variational parameters of each Gaussians are computed with a shared encoder network e_{ϕ_0} followed by separate linear transformations:

$$q_{\Phi_r}(g_{r0}|x)q_{\Phi_i}(g_{i0}|x) = \mathcal{N}(\mu_{\phi_{r_1}}(e_{\phi_0}(x)), {\sigma^2}_{\phi_{r_2}}(e_{\phi_0}(x))) \cdot \mathcal{N}(\mu_{\phi_{i_1}}(e_{\phi_0}(x)), {\sigma^2}_{\phi_{i_2}}(e_{\phi_0}(x))) \tag{2}$$

The inference networks for the behaviourally relevant and the behaviourally irrelevant initial conditions are parameterized by $\Phi_r = \{\phi_0, \phi_{r_1}, \phi_{r_2}\}$ and $\Phi_i = \{\phi_0,\ \phi_{i_1}, \phi_{i_2}\}$, respectively. It is important to note that TNDM's variational posterior only depends on the neural activity x. This approximation forces the learned initial conditions to come from the observed activity and allows for decoding of unseen behaviours after training. Generating the dynamics z from the neural activity x using this approximate posterior can be performed as shown in equation 1.

The ELBO for the observed data from a single trial is therefore defined as:

$$\text{ELBO}(x, y) = -\text{KL}\left[q_{\Phi_r} \,\|\, p_{g_{r_0}}\right] - \text{KL}\left[q_{\Phi_i} \,\|\, p_{g_{i_0}}\right] + \mathbb{E}_{q_{\Phi_r} q_{\Phi_i}}[\log p_{\theta_1}(x|g_i, g_r) p_{\theta_2}(y|g_r)] \tag{3}$$

where p_{θ_1} and p_{θ_2} are the observation models defined in equation 1 for neural activity and behaviour, respectively.

Disentangling the latent dynamics While the factorisation in the variational posterior helps, q_{Φ_r} and q_{Φ_i} can still become entangled during training without any additional constraints on their

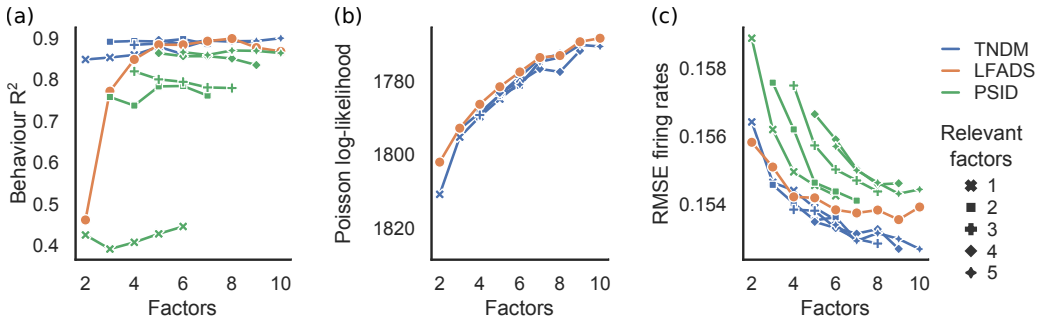

Figure 2: Summary of the behaviour and activity reconstruction accuracy for TNDM, LFADS, and PSID fit to neural recordings from the monkey primary motor cortex (M1) during a center-out reaching task. Each plot shows performance as a function of the total number of latent factors. (a) Coefficient of determination (R^2) for measured and reconstructed behaviour (hand position). For TNDM, the reconstruction performed by the behavour decoder (with only the relevant factors) while for LFADS and PSID, a ridge regression had to be used to decode the behaviour ex post factor (PSID's behaviour decoding fails completely on this dataset). (b) Poisson log-likelihood for the activity reconstruction per single trial. (c) Root mean square error (RMSE) between the predicted and actual ground-truth firing rates. Averaging was performed across all trials with the same movement direction. The train and test sets were pooled together in this computation to get more robust estimates of the rates from the experimental data.

information content. To this end, we introduce a penalty to limit sharing of information between the two random variables. As mutual information is hard to compute for high-dimensional variables and the experimental data often has a very limited number of trials to estimate these distributions reliably, we instead penalize the mean of the sample cross-correlations between $\mathbf{g}_{0_r}$ and $\mathbf{g}_{0_i}$. We refer to this penalty as $Q(q_{\Phi_r}, q_{\Phi_i})$ and we adjust its weight with a hyperparameter λ_Q (see Supplement 5 for an ablation study of this penalty). The final objective function for TNDM is then:

$$J(x, y) = -\,\mathrm{KL}\left[q_{\Phi_r} \,\|\, p_{g_{r_0}}\right] - \mathrm{KL}\left[q_{\Phi_i} \,\|\, p_{g_{i_0}}\right] + \mathbb{E}_{q_{\Phi_r} q_{\Phi_i}}[\log p_{\theta_1}(x|g_i, g_r)] + \lambda_b \mathbb{E}_{q_{\Phi_r}}[\log p_{\theta_2}(y|g_r)] + \lambda_Q Q(q_{\Phi_r}, q_{\Phi_i}) \tag{4}$$

where λ_b is an additional hyperparameter introduced to balance the behavioural likelihood with the neural likelihood (see Supplement 3 for hyperparameter details). While TNDM is not the first VAE to jointly model two observed variables with a partitioned latent space [27], it is distinguished by its unique objective function and penalties, its RNN-based architecture, its flexible linear decoder, and its novel application to neural activity and behaviour.

4 Application to M1 neural recordings during reach

To verify that TNDM is able to extract interpretable, behaviourally relevant latent trajectories without sacrificing its fit to the observed neural activity, we apply it to data gathered from a previously published monkey reaching experiment [6]. In this experiment, the monkey is trained to perform a center-out reaching task with eight outer targets. On a go cue, the monkey moves a manipulandum along a 2D plane to the presented target and receives a liquid reward upon success. Spiking activity from M1 and PMd along with the 2D hand position are recorded during each trial.

For our analysis, we train each TNDM on single-session data gathered from one of the six trained monkeys. The data consist of two paired datasets: PMd activity paired with hand position and M1 activity paired with hand position. We show results for the M1 recordings in the main text and the results for the PMd recordings in Supplement 2. The neural activity is counted in 10ms bins and the behaviour is also measured every 10ms. We align the behaviour to the spikes for both datasets by taking the activity starting with movement onset. We set the length of the neural activity to be the minimum time until the target is reached across all trials. As one of our baselines, PSID cannot model spike count data, we smooth the spike counts with a Gaussian kernel smoother before applying PSID (width 200ms). We also find that the behavioural prediction of PSID (which relies on a Kalman

filtering approach) suffers greatly if there is not adequate time for the Kalman filter to 'burn in'. So to enable a more fair comparison, we allow PSID to train on data that starts 100ms before movement onset. Out of the 176 trials from the experiment, we use 80% for training (136 trials). We hold out the remaining 34 trials to test the models.

Models We compare the performance of TNDM, PSID, and LFADS on the M1 data by training each model with different numbers of latent factors. For TNDM and PSID, we train models with all combinations of 1-5 relevant latent dynamics and 1-5 irrelevant dynamics (e.g. 3 relevant and 2 irrelevant). For LFADS, we train models with the number of latent dynamics ranging from 2-10. As TNDM and LFADS are both implemented as sequential variational autoencoders, we fix the architectures to be same for the two methods (64 units in the generators and encoder). We fix all shared hyperparameters to be the same between the two methods except for the dropout (TNDM requires more regularization due to the behavioural loss). For a detailed explanation of the hyperparameters and architectures used in these experiments, please see Supplement 3.

Fit to behaviour The performance of TNDM, LFADS, and PSID are summarised in Figure 2 and a full results table is provided in Supplement 4. First, we measure the ability of the learned dynamics to reconstruct behaviour by computing the coefficient of determination (R^2) between the decoded and ground-truth behaviour (Figure 2a; evaluated on test data). For TNDM, the output of the behaviour decoder, which only relies on the relevant factors, is reported. For LFADS, we regress *all* learned dynamical factors to the behaviour with a linear model (i.e. ridge regression[1]). Note that the latter approach for regressing LFADS factors is more flexible than the decoder in TNDM which is also linear but constrained to be causal. Behaviour decoding accuracy for TNDM with this regression, however, is only slightly higher than the reported model decoder results (see Supplement 4). Surprisingly, we could not get good behaviour prediction for PSID using its native behaviour decoder. Similarly to LFADS, we had to train a ridge regression to predict the behaviour from the relevant dynamical factors (see Supplement 6 for a more detailed analysis of PSID's results).

As can be seen in the figure, LFADS behavioural reconstruction performance saturates at around six factors ($R^2 \approx 0.89$) and with just four factors yields a respectable behavioural fit of $R^2 = 0.84$. This indicates that LFADS factors, which are constrained only by neural activity, are interpretable in terms of externally measured variables. TNDM achieves saturating performance with just three latent factors ($R^2 \approx 0.89$) where *only two are constrained to be relevant* for behavioural reconstruction. In fact, all TNDM models with three or more factors (where at least two are constrained to be relevant) have similar behaviour reconstruction accuracy. In comparison to TNDM, LFADS achieves a behaviour reconstruction of just $R^2 = 0.77$ for three latent factors. For PSID, we find that its behaviour decoder is unable to reconstruct the behaviour. The ex post facto regression of PSID's latent dynamics to the behaviour performs better, saturating at $R^2 = 0.78$ for 5 latent factors. Overall, TNDM's result on the behaviour decoding implies that the dimensionality of the behaviourally relevant dynamics for this 2D task are lower-dimensional than previously thought.

Fit to neural activity Do the additional constraints and penalties of TNDM affect the accuracy of neural activity reconstruction? This is an important question to answer as the learned latent dynamics are only meaningful if they also explain the observed activity well. We measure the neural reconstruction in two ways. First, we compute the Poisson likelihood of the observed activity per trial for both LFADS and TNDM (Figure 2b; evaluated on test data). Second, we extract condition-averaged firing rates from the ground-truth spike counts and measure the RMSE between these rates and each model's condition-averaged inferred firing rates. We compute these condition-averaged rates over the full dataset (i.e. the test and train splits) so that we have enough trials to get reasonable rate estimates (Figure 2c).

Surprisingly, we find that TNDM's neural reconstruction is nearly identical to that of LFADS, indicating that the partitioning of the latent dynamics and the additional constraints have virtually no effect on neural activity reconstruction for this dataset (Figure 2b,c). Instead, TNDM and LFADS both show a gradual improvement of neural activity reconstruction as a function of the number of factors. The only deviation from this trend is TNDM with one relevant and one irrelevant factor. This is not surprising, however, as much of the neural variability is explained by the behaviour of interest; only allowing one latent dynamical factor to explain the neural variability related to behaviour (while

[1]The ridge regression we utilize for this analysis is from scikit-learn [16] with default parameters.

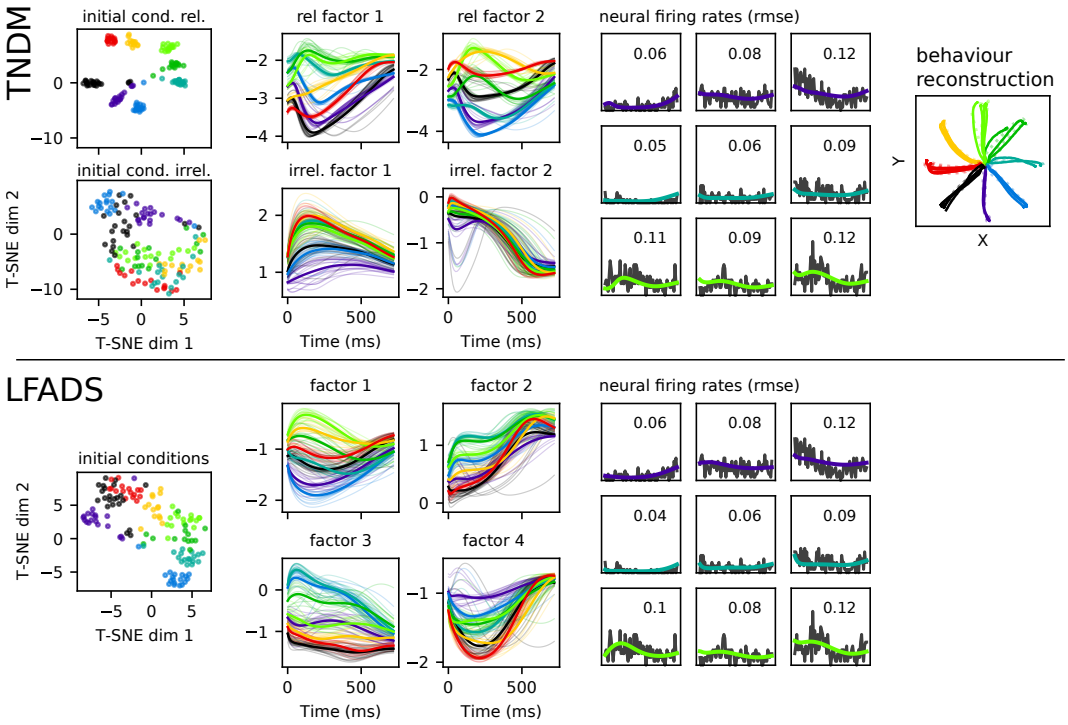

Figure 3: We visualize each component of the generative process for TNDM (top) and LFADS (bottom) after training. On the far left, we visualize the inferred initial conditions for each method after reducing the dimension to 2 with T-SNE. As can be seen, TNDM's inferred initial conditions show a clear distinction between behaviourally relevant and behaviourally irrelevant information whereas LFADS inferred initial conditions mix this information together. Next, we show the condition-averaged inferred latent dynamical factors (along with the single-trial factors) for each method to demonstrate that there is a clear distinction between the behaviourally relevant and behaviourally irrelevant factors in TNDM but not in LFADS. Finally, we show neural activity reconstruction for both methods (numbers are RMSE between data and prediction) and behaviour reconstruction for TNDM to illustrate that TNDM provides an excellent fit to the neural data despite the partitioned latent space and behavioural prediction.

simultaneously enforcing disentanglement between the relevant and irrelevant dynamics) will cause TNDM's neural activity reconstruction to suffer. PSID provides a worse fit to the neural data which is expected given that it is constrained to learn linear dynamics. An unexpected result for PSID, however, can be seen in Figure 2c; PSID appears to be heavily reliant on the irrelevant dynamics to get a good fit to the neural data. This is surprising given that we expect a good portion of the variability in the neural activity to be task dependent and may indicate that the relevant dynamics are not being correctly captured.

Interpretation of the learned dynamics In Figure 3, we visualize each component of the generative process of TNDM (2 relevant and 2 irrelevant factors) and LFADS (4 factors) after training both models on the M1 dataset. As can be seen in the figure, there appears to be a clear separation between the relevant and irrelevant initial condition distributions in TNDM that is less apparent in the mixed latent space of LFADS. In fact, the relevant initial conditions of TNDM seem to precisely capture the reach direction of the behaviour. Despite the noticeable differences between the dynamics of LFADS and TNDM, their ability to infer the underlying firing rates from this dataset are nearly identical (as illustrated in Figure 3).

Looking at the learned dynamical factors for TNDM, one can see that the relevant dynamics are more clearly distinguished by reach condition and there is much less variance in the learned trajectories than those of LFADS. At the same time, the relevant TNDM factors do not trivially re-capitulate the behaviour dynamics, indicating that the dual constraint of behaviour and neural activity unmasks a more complicated relationship between the two. This relationship can be analysed by visualizing the

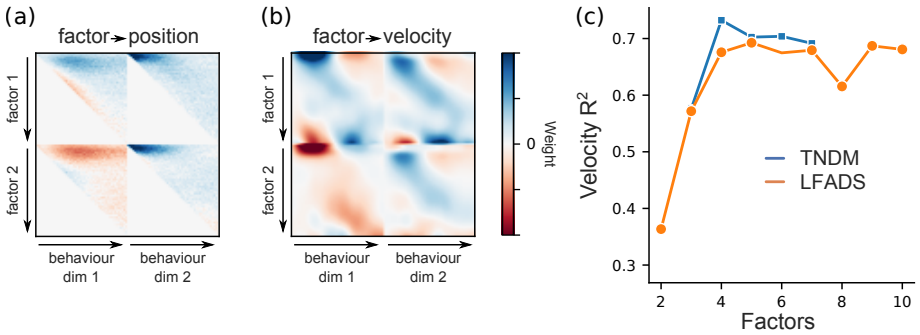

Figure 4: (a) Visualization of the weights of the TNDM behaviour decoder C_y that transforms two relevant latent factors into behaviour (hand position). The behaviours are aligned horizontally and the factors vertically. The upper-triangular structure reflects causal decoding, i.e. factors can only influence future behaviour. The model had two relevant and two irrelevant factors. (b) Weights obtained by using ridge regression to predict movement velocity (x and y components) from the relevant factors. A diagonally banded structure can be observed indicating a (delayed) identity transformation. Unlike in (a), this matrix is not constrained to be causal. (c) Decoding accuracy for velocity obtained using ridge regression for LFADS and TNDM with two relevant factors and a varying number of irrelevant factors.

learned weights of TNDM's behaviour decoder as shown in Figure 4a (for two relevant factors and two irrelevant factors). In this weight matrix, each time point of the behaviour receives contributions from a broad time interval of preceeding factor activity. This corresponds to a temporal integration of the factors and suggests that the relevant factors represent information about movement velocity. Indeed, velocity can be decoded well from the relevant factors using a simple ridge regression (both for TNDM and LFADS, Figure 4c). The learned coefficients of this ridge regression for TNDM have a diagonally banded structure that corresponds to a delayed identity transformation (Figure 4b), which is not visible for the LFADS factors (not illustrated). Taken together, these results suggest that M1 neural dynamics are related to velocity of hand movements in this task. Interestingly, we find that velocity decoding peaks at two relevant factors for TNDM and is less discernible when this number is increased, indicating that the addition of more relevant factors may spread this velocity information across the factors in a non-linear way which cannot be recovered by the ridge regression (this is result is not visualized in Figure 4). It also illustrates that the TNDM's behaviour prediction saturation point (two relevant factors) has perhaps the most interpretable latent space of all the trained TNDM models.

The irrelevant factors in TNDM show task-aligned dynamics that do not depend strongly on the task type, but are rather homogeneous (see Figure 3). For instance, over the course of each trial, irrelevant factor 2 has a large initial fluctuation followed by a steady increase in its absolute value over time until around 600ms where it tapers off (around when the monkey reaches the target destination). As this factor is agnostic to the reach direction, this may reflect dynamics associated with execution of these movements more generally.

5 Discussion

In this work, we introduce TNDM, a nonlinear state-space model designed to disentangle the behaviourally relevant and behaviourally irrelevant dynamics underlying neural activity. We evaluated TNDM on neural population recordings from PMd and M1 paired with a 2D center-out reach behaviour. We showed that TNDM was able to extract low-dimensional latent dynamics that were much more predictive of behaviour than those of current state-of-the-art state-space models without sacrificing any fit to the neural activity. This led us to the interpretation that the dimensionality of the neural activity associated with the 2D reaching task is potentially lower than previously thought and may be associated with the velocity of the hand.

Although the initial results presented for TNDM are quite promising, the method has a few limitations that should be addressed. First, we find that some random initializations of the model (i.e. seeds) can lead to biologically implausible oscillations in the learned latent dynamics. While more work needs

to be done to understand what gives rise to these oscillations, it could be related to the weighting between the behavioural and neural likelihoods or to the capacity of the model. For the experiments in the main paper, however, we did not run into any issues with the learned dynamics. A second limitation of TNDM is whether the disentanglement penalty between the relevant and irrelevant dynamics is sufficient. Although the covariance penalty works well in practice (on the presented datasets), disentangling sources of variation using deep generative models is still an open problem [13]. Similarly to PSID, TNDM could be implemented in multiple separate stages which may allow for better disentanglement of the relevant and irrelevant dynamics [25]. Third, the linear causal decoder we introduce for behaviour reconstruction is relatively parameter inefficient; the number of parameters scales quadratically with time and dynamics/behaviour dimension. We imagine that a more parameter efficient decoder could be utilized that still retains the desirable properties of our current decoder. Lastly, it can be challenging to determine the 'correct' latent dimensionalities for TNDM. For the datasets we present, we found that performing a wide sweep over a number of relevant and irrelevant dimensionalities and then choosing the relevant dimensionality where the behaviour prediction saturates is a potential recipe for finding an interpretable latent space.

In future work, we plan to train TNDM with higher-dimensional behaviours such as joint angle or electromyography (EMG) recordings. We also plan to extend TNDM to non-temporal/discrete behaviours which are of interest in behavioural neuroscience (e.g. decision making). Finally, we hope to extend TNDM such that it can model dynamics with external inputs from another brain region, i.e. non-autonomous dynamics.

6 Broader Impact

In this work, we develop an expressive and interpretable model of neural population activity. As such, we imagine that TNDM will be useful for answering important questions about neural function (e.g. how the motor cortex gives rise to behaviour). We also imagine that the ability of TNDM to accurately model both the neural activity and the behaviour of interest will be of interest for the brain-computer interface community. We believe that TNDM (or the principles behind it) can be used to improve behaviour decoding from neural activity. We also hope that TNDM inspires more research into deep generative models of neural activity that incorporate in external variables of interest. Currently, only a few models exist that go beyond modeling just the neural activity. We hope that the development of TNDM will lead to more interest and research in this area.

A possible negative societal impact of TNDM is that, like all deep neural network models, it requires a relatively large amount of compute and has a noticeable carbon footprint.

References

[1] *Abadi Martín, Agarwal Ashish, Barham Paul, Brevdo Eugene, Chen Zhifeng, Citro Craig, Corrado Greg S., Davis Andy, Dean Jeffrey, Devin Matthieu, Ghemawat Sanjay, Goodfellow Ian, Harp Andrew, Irving Geoffrey, Isard Michael, Jia Yangqing, Jozefowicz Rafal, Kaiser Lukasz, Kudlur Manjunath, Levenberg Josh, Mané Dandelion, Monga Rajat, Moore Sherry, Murray Derek, Olah Chris, Schuster Mike, Shlens Jonathon, Steiner Benoit, Sutskever Ilya, Talwar Kunal, Tucker Paul, Vanhoucke Vincent, Vasudevan Vijay, Viégas Fernanda, Vinyals Oriol, Warden Pete, Wattenberg Martin, Wicke Martin, Yu Yuan, Zheng Xiaoqiang*. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015. Software available from tensorflow.org.

[2] *Buesing Lars, Macke Jakob H, Sahani Meneesh*. Spectral learning of linear dynamics from generalised-linear observations with application to neural population data // Advances in Neural Information Processing Systems 25: 26th Conference on Neural Information Processing Systems (NIPS 2012). 2013. 1691–1699.

[3] *Churchland Mark M., Cunningham John P., Kaufman Matthew T., Foster Justin D., Nuyujukian Paul, Ryu Stephen I., Shenoy Krishna V.* Neural Population Dynamics during Reaching // Nature. VII 2012. 487, 7405. 51–56.

[4] *Cunningham John P, Byron M Yu*. Dimensionality reduction for large-scale neural recordings // Nature neuroscience. 2014. 17, 11. 1500–1509.

[5] *Elsayed Gamaleldin F, Cunningham John P*. Structure in neural population recordings: an expected byproduct of simpler phenomena? // Nature neuroscience. 2017. 20, 9. 1310.

[6] *Gallego Juan A, Perich Matthew G, Chowdhury Raeed H, Solla Sara A, Miller Lee E*. Long-term stability of cortical population dynamics underlying consistent behavior // Nature neuroscience. 2020. 23, 2. 260–270.

[7] *Gallego Juan A, Perich Matthew G, Miller Lee E, Solla Sara A*. Neural manifolds for the control of movement // Neuron. 2017. 94, 5. 978–984.

[8] *Hall Thomas M, Carvalho Felipe de, Jackson Andrew*. A common structure underlies low-frequency cortical dynamics in movement, sleep, and sedation // Neuron. 2014. 83, 5. 1185–1199.

[9] *Hurwitz Cole, Kudryashova Nina, Onken Arno, Hennig Matthias H*. Building population models for large-scale neural recordings: opportunities and pitfalls // arXiv preprint arXiv:2102.01807. 2021.

[10] *Jun James J, Steinmetz Nicholas A, Siegle Joshua H, Denman Daniel J, Bauza Marius, Barbarits Brian, Lee Albert K, Anastassiou Costas A, Andrei Alexandru, Aydın Çağatay, others* . Fully integrated silicon probes for high-density recording of neural activity // Nature. 2017. 551, 7679. 232–236.

[11] *Kingma Diederik P, Welling Max*. Auto-encoding variational bayes // arXiv preprint arXiv:1312.6114. 2013.

[12] *Kobak Dmitry, Brendel Wieland, Constantinidis Christos, Feierstein Claudia E, Kepecs Adam, Mainen Zachary F, Qi Xue-Lian, Romo Ranulfo, Uchida Naoshige, Machens Christian K*. Demixed principal component analysis of neural population data // Elife. 2016. 5. e10989.

[13] *Locatello Francesco, Bauer Stefan, Lucic Mario, Raetsch Gunnar, Gelly Sylvain, Schölkopf Bernhard, Bachem Olivier*. Challenging common assumptions in the unsupervised learning of disentangled representations // international conference on machine learning. 2019. 4114–4124.

[14] *McInnes Leland, Healy John, Melville James*. Umap: Uniform manifold approximation and projection for dimension reduction // arXiv preprint arXiv:1802.03426. 2018.

[15] *Pandarinath Chethan, O'Shea Daniel J, Collins Jasmine, Jozefowicz Rafal, Stavisky Sergey D, Kao Jonathan C, Trautmann Eric M, Kaufman Matthew T, Ryu Stephen I, Hochberg Leigh R, others* . Inferring single-trial neural population dynamics using sequential auto-encoders // Nature methods. 2018. 15, 10. 805–815.

[16] *Pedregosa F., Varoquaux G., Gramfort A., Michel V., Thirion B., Grisel O., Blondel M., Prettenhofer P., Weiss R., Dubourg V., Vanderplas J., Passos A., Cournapeau D., Brucher M., Perrot M., Duchesnay E.* Scikit-learn: Machine Learning in Python // Journal of Machine Learning Research. 2011. 12. 2825–2830.

[17] *Perich Matthew G, Conti Sara, Badi Marion, Bogaard Andrew, Barra Beatrice, Wurth Sophie, Bloch Jocelyne, Courtine Gregoire, Micera Silvestro, Capogrosso Marco, others* . Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior // BioRxiv. 2020.

[18] *Pfau David, Pnevmatikakis Eftychios A, Paninski Liam*. Robust learning of low-dimensional dynamics from large neural ensembles. // Advances in Neural Information Processing Systems. 2013.

[19] *Rabinovich Mikhail I, Huerta Ramón, Varona Pablo, Afraimovich Valentin S*. Transient cognitive dynamics, metastability, and decision making // PLoS Computational Biology. 2008. 4, 5. e1000072.

[20] *Russo Abigail A, Khajeh Ramin, Bittner Sean R, Perkins Sean M, Cunningham John P, Abbott Laurence F, Churchland Mark M*. Neural trajectories in the supplementary motor area and motor cortex exhibit distinct geometries, compatible with different classes of computation // Neuron. 2020. 107, 4. 745–758.

[21] *Sadtler Patrick T, Quick Kristin M, Golub Matthew D, Chase Steven M, Ryu Stephen I, Tyler-Kabara Elizabeth C, Byron M Yu, Batista Aaron P*. Neural constraints on learning // Nature. 2014. 512, 7515. 423–426.

[22] *Sani Omid G., Abbaspourazad Hamidreza, Wong Yan T., Pesaran Bijan, Shanechi Maryam M.* Modeling Behaviorally Relevant Neural Dynamics Enabled by Preferential Subspace Identification // Nat Neurosci. XI 2020. 1–10.

[23] *She Qi, Wu Anqi*. Neural dynamics discovery via gaussian process recurrent neural networks // Uncertainty in Artificial Intelligence. 2020. 454–464.

[24] *Smith Anne C, Brown Emery N*. Estimating a state-space model from point process observations // Neural Comput. 2003. 15, 5. 965–991.

[25] *Srivastava Akash, Bansal Yamini, Ding Yukun, Hurwitz Cole, Xu Kai, Egger Bernhard, Sattigeri Prasanna, Tenenbaum Josh, Cox David D, Gutfreund Dan*. Improving the Reconstruction of Disentangled Representation Learners via Multi-Stage Modelling // arXiv preprint arXiv:2010.13187. 2020.

[26] *Sussillo David, Jozefowicz Rafal, Abbott LF, Pandarinath Chethan*. Lfads-latent factor analysis via dynamical systems // arXiv preprint arXiv:1608.06315. 2016.

[27] *Whiteway Matthew R, Biderman Dan, Friedman Yoni, Dipoppa Mario, Buchanan E Kelly, Wu Anqi, Zhou John, Noel Jean-Paul R, Cunningham John P, Paninski Liam, others* . Partitioning variability in animal behavioral videos using semi-supervised variational autoencoders // bioRxiv. 2021.

[28] *Zhou Ding, Wei Xue-Xin*. Learning identifiable and interpretable latent models of high-dimensional neural activity using pi-VAE // arXiv preprint arXiv:2011.04798. 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] **See Discussion section**.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] **See broader impact section**.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] **The datasets we use are large and not publicly available yet. We plan to release our codebase upon acceptance.**
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] **In the supplement and in the main text we specify parameters and splits.**
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] **We provide standard deviations for our results in the supplementary tables**.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] **The models we train are relatively small and compute should not be an issue.**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] **See footnote 1**
   (b) Did you mention the license of the assets? [Yes] **See footnote 1**
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]