
Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation

Justin Jude¹ Matthew G. Perich² Lee E. Miller³ Matthias H. Hennig¹

Abstract

Neural population activity relating to behaviour is assumed to be inherently low-dimensional despite the observed high dimensionality of data recorded using multi-electrode arrays. Therefore, predicting behaviour from neural population recordings has been shown to be most effective when using latent variable models. Over time however, the activity of single neurons can drift, and different neurons will be recorded due to movement of implanted neural probes. This means that a decoder trained to predict behaviour on one day performs worse when tested on a different day. On the other hand, evidence suggests that the latent dynamics underlying behaviour may be stable even over months and years. Based on this idea, we introduce a model capable of inferring behaviourally relevant latent dynamics from previously unseen data recorded from the same animal, without any need for decoder recalibration. We show that unsupervised domain adaptation combined with a sequential variational autoencoder, trained on several sessions, can achieve good generalisation to unseen data and correctly predict behaviour where conventional methods fail. Our results further support the hypothesis that behaviour-related neural dynamics are low-dimensional and stable over time, and will enable more effective and flexible use of brain computer interface technologies.

1. Introduction

In the brain, stimuli and behaviour can be decoded from the activity of populations of neurons, and it is well estab-

¹School of Informatics, University of Edinburgh, Edinburgh, Scotland, EH8 9AB ²Université de Montréal and Mila, Montréal, QC, Canada H3C 3J7 ³Feinberg School of Medicine, Northwestern, Chicago, IL 60611. Correspondence to: Justin Jude <justin-jude@me.com>.

lished that correlations or co-variations between neurons are a key ingredient in neural population codes (Saxena & Cunningham, 2019). There has been considerable success developing methods for decoding external variables from recordings of even modestly sized populations of 10s or 100s of neurons (Hurwitz et al., 2021a), raising hopes that brain computer interfaces (BCIs) can be an effective assistive technology for severely disabled patients. However, a decoder, once trained, requires stable recordings to perform well. Over the course of days and weeks, the signals recorded from implanted extracellular probes will inevitably change and drift due to factors such as impedance changes, gliosis and probe and brain movement (Chestek et al., 2011). Non-invasive systems such as electromyography (EMG) sensors will not be worn permanently and positioned slightly differently every time, creating even stronger variations in recorded signals. Moreover, the activity of individual neurons can change considerably over similar time scales due to neural plasticity (Rule et al., 2019). Together these fluctuations will lead to degradation of decoder performance over time, thus to be effective, frequent recalibration of BCI systems would be inevitable.

Given the limited long-term stability of recorded neural signals, reports of relatively stable behaviour decoding over days with the same decoder may seem surprising (Chestek et al., 2007). Recent work by Gallego et al. (2020) however showed that some aspects of the population activity of cortical neurons remain very stable even over months and years. Specifically, this study showed that neural population activity in the primary motor cortex is highly restricted to and evolves along a low-dimensional manifold that is stable even when single neuron activity constantly fluctuates.

Low-dimensional neural dynamics can be effectively extracted from neural population activity with latent variable models (Hurwitz et al., 2021a). These models use an often small number of latent variables (or factors) together with an appropriate observation model that relates latent variables to the recorded activity. Importantly, the latent variables in such models often predict stimuli or behaviour very well even when they were only optimised to reproduce neural activity (Hurwitz et al., 2021a). Nonlinear state space models such as LFADS are particularly powerful in predicting

single trial activity and behaviour in test data (Pandarinath et al., 2018).

Therefore, instabilities in neural recordings can be successfully compensated for by re-training the part of a model that translates neural activity into the latent dynamics, which are assumed stable over time (Farshchian et al., 2019; Dabagia et al., 2022; Degenhart et al., 2020; Wen et al., 2021; Herrero-Vidal et al., 2021). As a behaviour decoder, this can be more data-efficient than re-training a decoder from scratch, but still requires regular interventions. Here we ask if it is possible to recover the latent dynamics without any re-training.

Our approach uses a domain adaptation inspired solution. Sources of session to session variability in neural recordings are shown in Figure 1 and include existing neurons lost from recording electrodes, existing neurons replaced by unseen neurons, and all recording electrodes shifting systematically due to probe array shift. This variability is substantial enough that each recording session can be constituted as a separate domain.

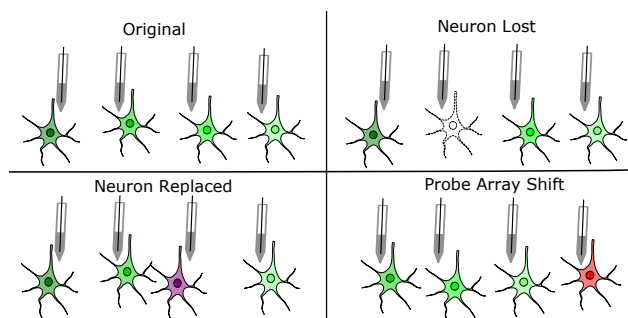


Figure 1. Causes of session-to-session variability in recordings from neural populations. Neurons from the original recording session can be lost from the recording array, original neurons can be replaced by unseen neurons and the entire probe array can shift, causing a systematic change in neuron position. In addition, spike sorting can induce variability as the signal to noise ratio of individual neurons changes between sessions. Domain adaptation of many varied sessions enables our model to learn these sources of variability.

We treat each recording session as a separate domain, each of which can be used to predict the same set of behaviours. The model is optimised using both recorded activity and behaviour to recover the same latent variables irrespective of the domain so it is capable of predicting behaviour correctly for a previously unseen session without need for re-calibration. In contrast, latent variable models without domain adaptation fail to generalise to unseen data, and instead partition the latent space into distinct parts corresponding to the individual recording sessions. We test this

model with long-term recordings from the primate motor cortex during a reach task and show that, provided sufficient training data, it can predict behaviour well for previously unseen sessions. BCI decoders that can generalise well to unseen sessions or subjects without any re-training have not yet been demonstrated. We believe this is the first work to show such cross-session decoder generalisation without recalibration.

2. Related Work

This issue of neural stability is investigated by Gallego et al. (2020) where the dynamics of a set of a single animal’s M1 cortex neurons are recorded from over many days. The authors find that the underlying dynamics of these neurons over time are indeed reconcilable. Principal component analysis (PCA) is used to reduce the dimensionality of the neural activity on each day, and these variables are then aligned using canonical correlation analysis (CCA). After alignment, neural activity is regenerated for up to 16 days with close similarity and accurate decoding of behaviour.

Farshchian et al. (2019) take this approach a step further and utilise an adversarial approach with a non-linear model (ADAN) to directly align neural activity over many days in order to accurately predict EMG during movement. A discriminator network is trained in a similar fashion to LFADS, tasked with autoencoding neural activity from day 0. A generator neural network is optimised to align neural population activity to that recorded at day 0. The autoencoding discriminator is tasked with maximising the alignment loss.

Sussillo et al. (2016) build a robust decoder capable of utilising large amounts of training data and maintaining decoding performance in the face of recording condition changes such as neuron turnover. Herrero-Vidal et al. (2021) introduces a robust probabilistic approach for neural alignment in a common low dimensional manifold. Their method is applied to recordings from the mouse olfactory bulb, revealing low-dimensional population dynamics that are odour specific and have consistent structure across animals. Following the same idea, Wen et al. (2021) uses adversarial generative modelling to generate large amounts of synthetic spike data from just the behaviour of a separate recording session or subject, mimicking the spike data of that session/subject. Together with this generated synthetic spike data and a small amount of real spike data from the unseen session, the authors are able to achieve relatively good behaviour decoding accuracy on the held out session. This model is more data efficient than the previously mentioned approaches. Nevertheless, in all these cases data from all recording sessions or animals is required for good behaviour decoding, and the models are not capable of generalisation to unseen sessions.

Hurwitz et al. (2021b) combines ideas from Pandarinath et al. (2018) and Sani et al. (2021) to jointly model the neural activity and external behavioural variables by separating the latent space into behaviourally relevant and behaviourally irrelevant components; 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. This work shows that an LFADS-like model can jointly model neural activity and associated behaviour or movement, hence potentially isolate invariant behaviour-related dynamics that can be used for cross-session decoding.

Domain adaptation broadly aims to predict classes from labelled data of a similar nature, albeit from differing sources or domains. The method relevant to this work is by Ganin & Lempitsky (2015), who use a negative gradient between a domain classifier and feature extractor in order to coerce the feature extractor to produce domain invariant features from which a label predictor can infer data classes reliably. This method of domain unification is unsupervised.

An application of domain adaptation to correct for variability in experimental data by Gonschorek et al. (2021) used an autoencoder model and a domain classifier to align two-photon imaging data across experiments. The authors successfully align their recording sessions but they do not test efficacy on unseen sessions. They also explicitly use experimental session ID as domains and show efficacy on non-sequential data in this respect. In this work we show that it is beneficial to not explicitly use session/experiment ID for domain adaptation but to instead use neural patterns directly to align recording sessions for high dimensional sequential data.

In this work we model each session of neural recording as a separate domain and predict behaviour from all of these sessions simultaneously. Domain-invariant latent variables are obtained using the paradigm of unsupervised domain adaptation via a negative gradient, which are then optimised to reconstruct the observed behaviour.

3. Model

This model is based on the hypothesis that behaviour y is encoded in a stable latent space with variables z , and that the two are related linearly as $y = f(z)$. Equally, neural activity x is related to the latent variables through a simple function, and as in related models we choose a linear read-out with a Poisson link function to generate non-negative firing rates (Pandarinath et al., 2018). However, this function will differ between recording sessions (or domains) d as we expect to observe different neurons in each session, and the activity of neurons may change over time. The problem is thus to find the correct encoding function $z = g(x)$ to transform neural

activity into the latent space which then allows decoding of behaviour. As explained above, re-training this part of the model for each session can successfully align different sessions. Here we show that this can be achieved without the need for re-training.

Specifically, as proposed by Pandarinath et al. (2018) we assume that the latent dynamics evolve autonomously provided a set of initial conditions z_i that are modelled as Gaussian random variables. These latent variables are produced for each trial by an encoder network consisting of bidirectional Gated Recurrent Units (Cho et al., 2014) (GRU). They are used to simultaneously predict behaviour, and to reconstruct the original trial-specific neural activity. We apply recurrent and kernel regularisation to the encoder GRU to enable better generalisation to unseen sessions.

A further bidirectional GRU is used as a decoder for neural reconstruction and a final separate GRU is used to predict behaviour from the generated latent variables. Training is based on a mean squared error loss for behaviour and Poisson likelihood for neural activity. Importantly, we reverse the backpropagation gradient between the neural reconstruction decoder and the encoder. This gradient reversal layer leads to maximisation of the neural reconstruction loss in the encoder network while, at the same time, the neural decoder network is adversarially optimised to minimise neural reconstruction loss. This implicitly encourages the encoder to generate latent variables which are not separated by session of data collection.

The behaviour decoder meanwhile forces the encoder to generate latent variables which are differentiated by behaviour. Ultimately, this produces a latent space separable by behaviour but not by session of data collection. The complete model is illustrated in Figure 2.

The model is trained using real neural activity which corresponds to consistent behaviours (movement directions in a centre-out reach task, see below). The generative process of our model is as follows:

$$z_i = W_{enc}(\text{GRU}_{\theta_{enc}}(x_{i,1:T})), \quad (1)$$

$$g_{1:T} = \text{GRU}_{\theta_{dec}}(z_i), \quad (2)$$

$$b_{1:T} = \text{GRU}_{\theta_{beh}}(z_i), \quad (3)$$

$$r_t = \exp(W_{rate}(W_{fac}(g_t))), \quad (4)$$

$$\bar{x}_t \sim \text{Poisson}(r_t), \quad (5)$$

$$\bar{y}_t = W_{beh}(b_t) \quad (6)$$

where θ_{enc} , θ_{dec} , θ_{beh} are the parameters of the GRUs used to encode spike trains into latent variables, decode spike trains from the generated latent variables, and to decode behaviour from the latent variables respectively. W_{enc} , W_{fac} , W_{rate} and W_{beh} are non-linear layers which produce latent

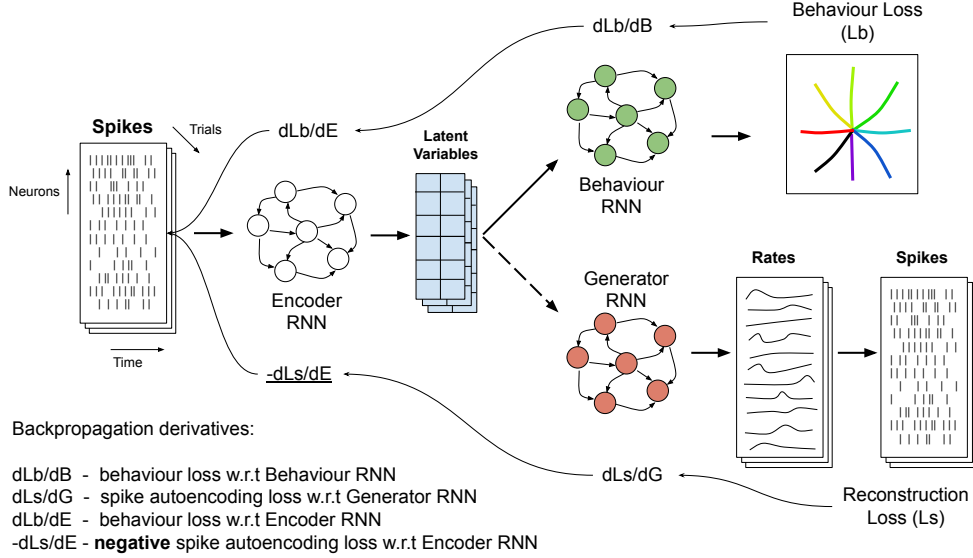


Figure 2. Our model (SABLE) consists of a sequential variational autoencoding approach combined with a sequential behaviour decoder. Notably, we implement a reverse gradient layer between the neural decoder and encoder GRUs. The encoder can then learn to extract the invariant latent dynamics determining behaviour from data obtained in different sessions with variability in the recorded neural activity.

variables, neural activity factors, generate firing rates and predict behaviour respectively at each time step per trial.

At each training iteration the following three losses are optimised using Adam (Kingma & Ba, 2015) asynchronously:

$$L_{rec} = - \sum_{t=1}^t \log(\text{Poisson}(x_{i,t}|r_t)) \quad (7)$$

$$L_{beh} = \frac{1}{T} \sum_{t=1}^t (y_{i,t} - \bar{y}_{i,t})^2 \quad (8)$$

$$L_{kl} = D_{KL}[\text{GRU}_{\theta_{enc}}(z_i|x_i)||\mathcal{N}(0, I)] \\ = -\frac{1}{2} [\log(z_{i,\sigma}^2) - z_{i,\mu}^2 - z_{i,\sigma}^2 + 1] \quad (9)$$

where y_i is the true behaviour per trial and \bar{y}_i is the predicted behaviour. The loss in Eq. 7 is maximised by the encoder network and minimised by the neural decoder network (and not applicable to the behaviour decoder network). This adversarial training is the most crucial aspect of our model. As the encoder maximises the neural reconstruction loss throughout training, it produces increasingly spike pattern-invariant latent variables.

Behaviour loss (Eq. 8) is minimised by both the encoder and behaviour decoding network while the Kullback–Leibler (KL) divergence loss (between a multivariate standard Gaussian distribution and the encoder generated latent variables) is minimised by just the encoder network. Thus the total error for all parameters in the model across all training trials

can be summarised as:

$$E(\theta_{enc}, W_{enc}, \theta_{dec}, W_{fac}, W_{rate}, \theta_{beh}, W_{beh}) = \\ \sum_{i=1..N} L_{beh}^i(\theta_{enc}, W_{enc}, \theta_{beh}, W_{beh}) \\ + L_{rec}^i(\theta_{dec}, W_{fac}, W_{rate}) \\ + \lambda_{kl} L_{kl}^i(\theta_{enc}, W_{enc}) - \lambda_r \sum_{i=1..N} L_{rec}^i(\theta_{enc}, W_{enc}) \quad (10)$$

where λ_{kl} is the weight of KL divergence and λ_r is the weight of the reverse gradient applied to the encoder RNN. λ_{kl} rises exponentially as training progresses while λ_r decays exponentially (thereby increasing session invariance over training). We denote our model Stable Alignment of Behaviour through spike-invariant Latent Encoding (SABLE).

SABLE does not require specific hyperparameter tuning for either monkey tested in Section 6, however, in subjects or experimental setups where neural drift is more variable between recording sessions, tuning recurrent dropout and kernel regularisation values may be beneficial for optimal behaviour decoding performance from unseen session trials.

4. Data

4.1. M1 neural recordings during reach task

We verify that SABLE is able to predict behaviour from unseen neural activity by applying it to data from a previ-

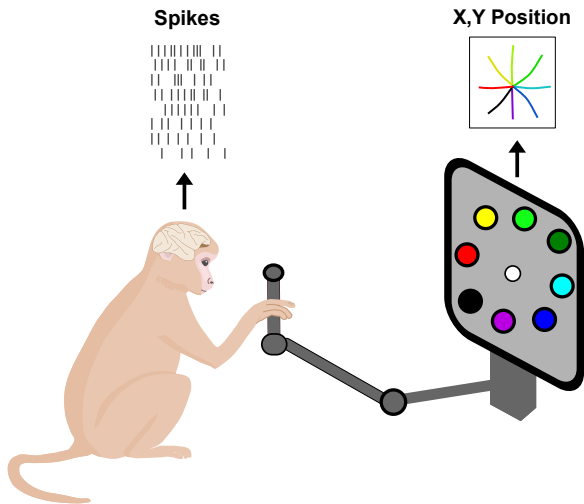


Figure 3. Experimental setup: In each trial one randomly chosen target direction (indicated by one of 8 coloured circles) appears on screen, and the monkey is instructed to control the cursor (white circle) by moving the manipulandum. The monkey moves the cursor to the target location after a go cue. The collected data for each trial consists of the neural spikes and monkey hand position across all timesteps. Our model is tasked with predicting hand position from neural spikes at each timestep.

ously published experiment (Gallego et al., 2020). In this experiment, two monkeys are trained to perform a center-out reach task towards eight outer targets. On a go cue, each monkey moves a manipulandum along a 2D plane to guide a cursor on a screen to the target location (Figure 3). On successful trials a liquid reward is given. Spiking activity from the motor cortex (M1) along with the 2D hand position are recorded during each trial. Spike trains are converted into spike counts in 10ms bins, and behaviour variables are used at the same resolution. Only successful trials are used, all trials are aligned to movement onset and cut from movement onset to the shortest reach time across all trials.

For our analysis, we train SABLE on many consecutive days of recorded data and test on a subsequent held out day of recordings for each monkey. In total there are 13 near consecutive days of recordings for monkey 1 and 6 near-consecutive days of recordings for monkey 2, with fewer recorded neurons and timesteps for monkey 2 overall. Each day for each monkey consists of one recording session.

5. Models for comparison

We compare the ability of SABLE to predict behaviour from sessions of unseen spike data against existing methods and against a variation of our own model. We look at the following existing models: LFADS (Pandarinath et al., 2018) and RAVE+ (Gonschorek et al., 2021). We also compare

against our own model where we do not reverse the gradient between the encoder and decoder, which we denote SABLE-noREV. In addition, we compare against a baseline RNN (GRU) with a linear readout layer optimised to predict movement from spiking data without autoencoding.

LFADS has been shown to have good efficacy at neural reconstruction across trials and sessions with some separation of behaviour in its latent space in previous work. We implement RAVE+ as an autoencoding model with GRUs for the encoder and decoder as our data are time series, and treat recording sessions as separate domains. As with our own model, the encoder is tasked with generating a small number of latent variables following a multivariate standard Gaussian distribution from neural data while the decoder is tasked with reconstructing the data from the latent variables. We use a non-linear layer as a domain classifier on the latent space between the encoder and decoder and implement a negative gradient between this classifier and the encoder network, thus encouraging the encoder to produce session-invariant latent variables. For all models we use the same regularisation techniques in the encoder or predictor as we do for SABLE to maximise generalisation.

For LFADS and RAVE+ we use a separately trained GRU to predict behaviour from the latent space of these models. We do not include ADAN (Farshchian et al., 2019) or the generative adversarial model by Wen et al. (2021) as both require at least some training data from held out session or subject to be effective. Implementation details of SABLE and all comparison models can be found in the Appendix (Section B).

6. Results

6.1. Application to motor cortex neural recordings during a reach task

We train all models on varying numbers of training sessions and for both monkeys, testing behaviour (2D hand position) prediction on intermediate and subsequent held out recording sessions. Our results, summarised in Figure 4, show that SABLE is capable of generalising to unseen data provided a sufficient number of training sessions are provided. In all cases tested SABLE outperforms the comparison models. For example, decoding accuracy for SABLE on an unseen intermediate session for monkey 1 with 12 training sessions is 0.91, which exceeds all other models by at least 0.25. For comparison, the RNN decoder typically yields an accuracy of around 0.92 on held-out data when trained and tested on a single session, indicating that SABLE can achieve saturation performance on unseen data. RAVE+ has a relatively high decoding performance when a large number of training sessions are used, likely because its domain adaptation paradigm removes some session specific variance in this

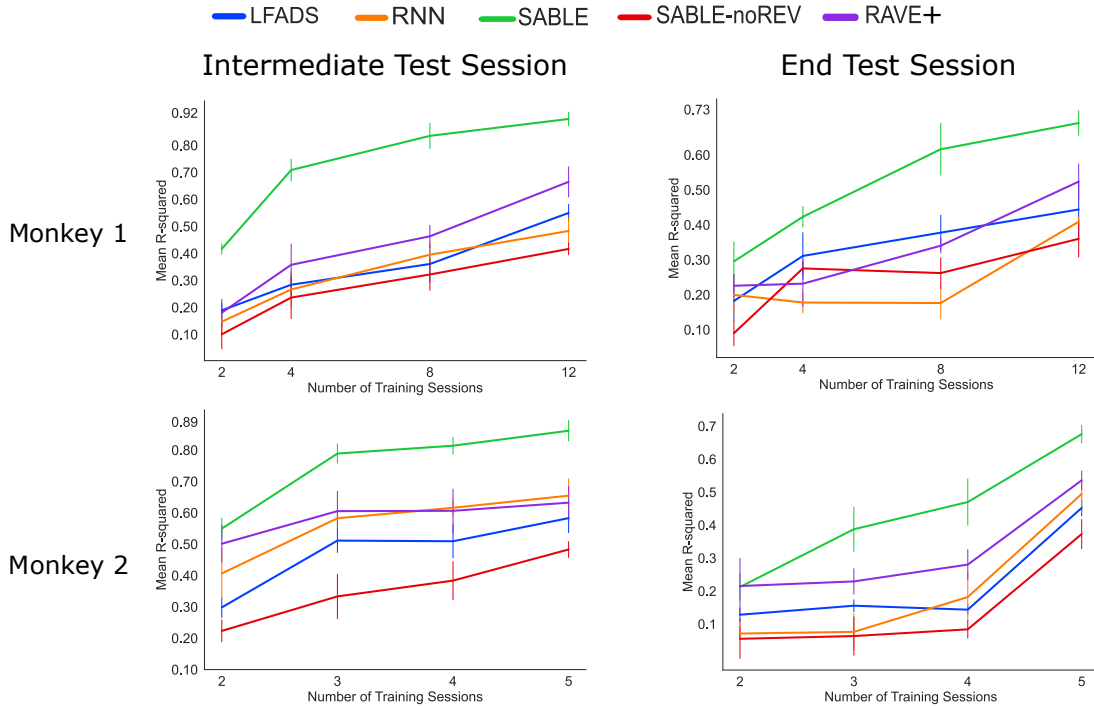


Figure 4. Behaviour prediction performance when testing all models on a completely unseen recording session. We report the mean r-squared between the inferred and true x,y positions. Each model is tested on a held out session while trained on different numbers of training sessions for both monkeys. The left column shows results for a held-out testing session which is in the chronological centre of the training sessions whereas the right column shows results for a test session recorded after all training sessions. Each test condition is run 10 times with different random seeds, with error bars showing standard deviation.

case. In contrast, SABLE-noREV and LFADS have low decoding performance across monkeys and session numbers although they gradually improve with increasing session numbers.

Comparing the performance between the two monkeys, we see generally better overall decoding performance for monkey 2 at the same number of train sessions as monkey 1, although monkey 1 has far more total training data available (12 total consecutive sessions from monkey 1 vs. 5 for monkey 2) and so has higher peak behaviour decoding performance for all models. In addition, we limit the number of neurons for each monkey to the lowest number of neurons in any given session. Therefore, monkey 1 has 42 neurons of neural data across sessions whereas monkey 2 has 16.

Next we compare the difference between test performance for all models on both monkeys for different held out test session ordering (intermediate or end). While SABLE achieves end test session decoding performance exceeding that of current methods (0.71 mean r-squared with 12 train sessions), performance on any given intermediate test session is substantially higher (0.90 mean r-squared with 12 train sessions). Moreover, the performance of SABLE decreases noticeably faster when applied to an end test session when

the number of training sessions is reduced versus an intermediate test session. This confirms that drift in recordings is gradual, not random.

6.2. Latent space analysis

Figure 5 shows T-SNE embeddings of the latent space of all autoencoding models. Each colour represents a different target direction with respect to behaviour, and embeddings of the training data are shown as circles while the test trials are shown as triangles.

The embeddings of the LFADS latent space show a clear separation that corresponds to the different training sessions. This shows that there are indeed significant differences between the sessions that are captured in the latent space and prevent generalisation to unseen sessions. Within each session cluster there is a good separation by target direction, indicating that the latent variables extract behaviourally relevant information from the neural activity. In contrast, many trials from the test session form a cluster in a region not covered by the training data, with some degree of separation by behaviour in monkey 1 (where more sessions are available for training). Here the model fails to assign these trials to meaningful latent variables as the differences in activity are

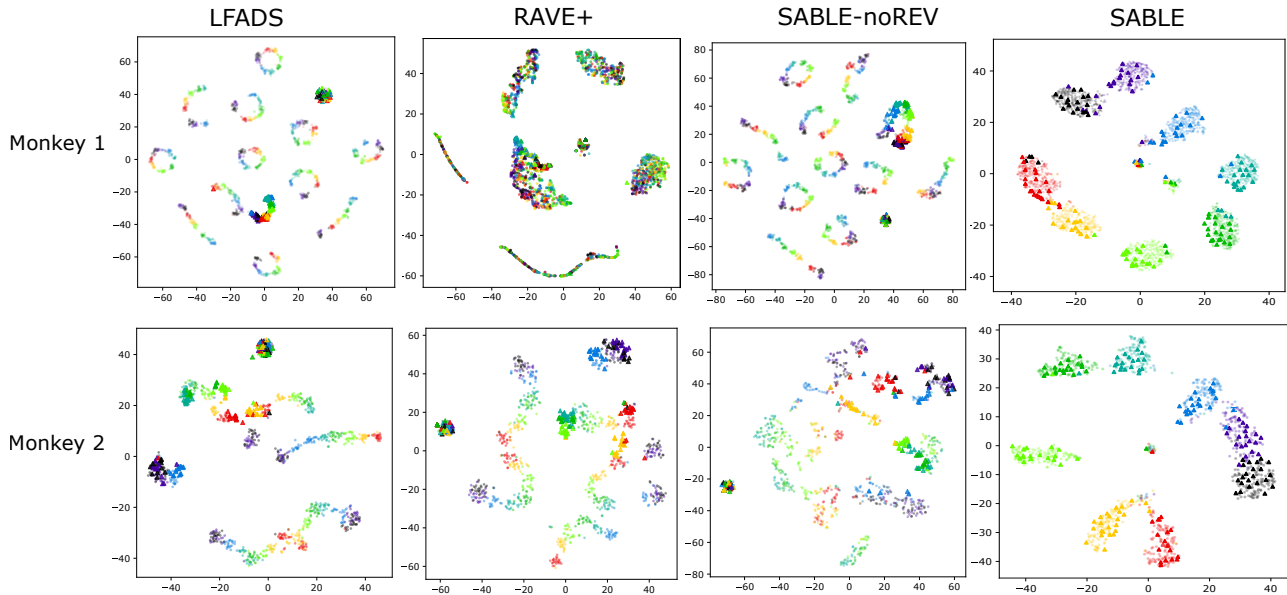


Figure 5. T-SNE embeddings of latent space for autoencoder models. In each embedding, points denoted by a circle are trials from 12 training sessions for monkey 1 and 5 training sessions for monkey 2. Points denoted by a triangle are trials from a held out intermediate test session for both monkeys. Each colour represents a target direction for the centre out reach task.

too large to be mapped appropriately. We denote trials such as these hereafter as unassigned. Some trials are assigned locations in the latent space also occupied by the training data, indicating that despite session differences occasionally a matching of unseen to training data can be achieved, again with some separation by target direction. However there is also a cluster of unseparated test trials which the encoder of the model has failed to produce meaningful latent variables for due to these test trials being too disparate from the training trials.

For LFADS applied to monkey 2, the encoder still manages to separate trials by train session, but to a lesser degree. We suspect this is due to fewer neurons being available for the sessions of monkey 2. Here there is no longer a cluster of separated test trials, instead some test trials are assigned to existing train clusters as they are fairly similar. Here again however we see a large cluster of unassigned test trials. Overall we see that LFADS clusters train trials and coinciding test trials well but its encoder cannot effectively generate latent variables for dissimilar test trials for behaviour decoding.

The picture is different for RAVE+, where the latent space seems better aligned (more session invariant) but no longer well separated by behaviour for monkey 1 (larger training set and more neurons). In contrast, for monkey 2 we see little session alignment but better behaviour separation is achieved (fewer sessions and fewer recorded neurons). In

this case there is some degree of merging of session clusters and organisation by behaviour target direction, but this is insufficient for good test behaviour decoding. For monkey 1 we see 7 clusters, 4 of which have some separation by direction for train session trials. However test trials are not separated well at all by direction. For monkey 2 there is far less clustering by session and some of these clusters separate by direction. Here there is also a large cluster of unassigned test trials. The domain adaptation method used in RAVE+ (reverse gradient based on session ID explicitly) thus does not seem to prevent trials clustering by session.

SABLE-noREV, our model without the reverse gradient, produces a result very similar to LFADS. For both monkeys, there is a cluster of well separated test trials and also many test trials that are unassigned to any cluster (either by direction or session). Therefore, using the latent space for both neural reconstruction and behaviour decoding simultaneously, as proposed by Hurwitz et al. (2021b), is not beneficial to test behaviour decoding across sessions.

Finally, applying SABLE to either monkey produces latent spaces which are very well separated by behaviour and almost entirely training-session invariant. We denote each unseen test trial as correctly classified in terms of target direction by observing whether a given trial gives a behaviour decoding r-squared of above 0.6. When applied to monkey 1 we see a small degree of misclassification of test trials by direction (13% of total test trials), but only when the correct

and incorrect target directions are adjacent to each other spatially in the task outlined in Figure 3. This confirms that more similar behaviours in a task have more similar neural patterns and may be difficult for any decoder to disentangle. There are also a small number of unclassified test trials (3% of total test trials) in the centre of the embedding plot, we suspect these may be trials with highly contrasting spiking patterns to any train trial. When applied to monkey 2, we see less misclassification (4% of total test trials) by test behaviour direction and just a couple of unclassified test trials. We also note that both SABLE embeddings are topographically similar and correspond to the spatial aspect of the movement directions of the task outlined in Figure 3.

The stark differences in latent variables seen between SABLE-noREV and SABLE are quite surprising considering the only difference between these models is the reverse gradient between neural decoder and encoder in SABLE versus a positive gradient in SABLE-noREV. This shows the importance of a negative neural reconstruction gradient in training the SABLE encoder network to generate session invariant latent variables. In addition, we suspect that SABLE’s encoder generates far fewer unassigned latent variables than the other autoencoding models due to the simpler and more behaviourally structured latent space.

6.3. Behaviour decoding

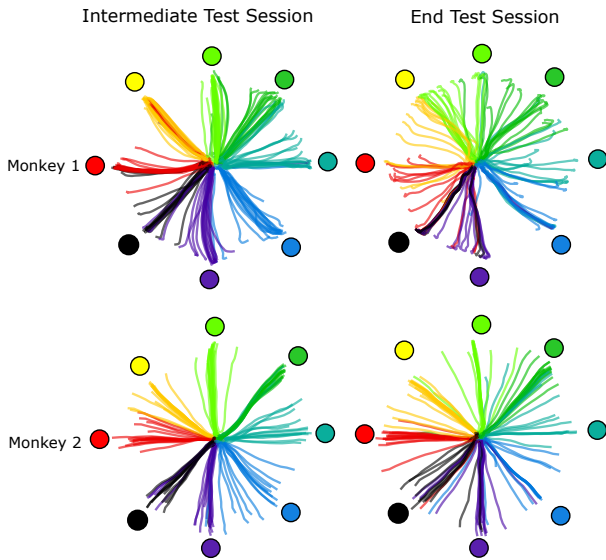


Figure 6. Predicted 2D monkey hand position of test trials using SABLE when trained on 12 train sessions for monkey 1 and 5 train sessions for monkey 2 and tested on an unseen intermediate or end test session.

Examining the decoded behaviour of monkey hand position using SABLE (Figure 6) shows good overall reconstruction of movement trajectories, especially when testing on an

intermediate test session. The intermediate test session behaviour decoding for both monkeys mirrors the SABLE T-SNE embedding in Figure 5. Test trials which are incorrectly assigned with respect to movement directions (Figure 5) are decoded correspondingly (Figure 6). Therefore, the behaviour decoder network of SABLE directly utilises latent variables in a particular cluster and decodes one particular direction of movement. Our model is thus consistent with the hypothesis outlined above.

When decoding from an end test session however, this phenomenon is less pronounced as the encoder seems to be less certain of the clusters formed in the latent space. There are more wrongly assigned test trials and the decoded movement trajectories are more spread out, leading to a lower overall mean r-squared when predicting behaviour.

6.4. Predicting behaviour from an unseen subject

Next we predict behaviour from the unseen neural data from 37 sessions of monkey 2 when SABLE is trained on 14 recording sessions of monkey 1. We use all available recording sessions available for monkey 1 spread across 3 years to train SABLE as we believe this gives the best opportunity for cross subject generalisation. However, for the held-out data we only obtain a mean r-squared of 0.03, so the model fails to generalise to a different animal. Examining the T-SNE embedding in Figure 7 shows that the trials from the training sessions of monkey 1 cluster well by movement direction but the trials from monkey 2 do not map to these direction clusters as the separated sessions of either just monkey 1 or monkey 2 do (as seen in Figure 5). Therefore it appears that the relationship between the recovered latent dynamics and behaviours differs between the two animals, and may require an extra alignment step.

7. Discussion

In this work we present a new method, SABLE, for aligning neural activity with complex temporal dynamics from different recording sessions to allow for consistent behaviour decoding across sessions. We apply it to neural recordings from primate motor cortex during a reaching task where the considerable variability between recording sessions prevents generalisation for a conventional decoder.

The model is trained as a variational autoencoder similar to LFADS (Pandarinath et al., 2018), with an additional gradient from behaviour decoder that disentangles the latent space to enable improved behaviour decoding (Hurwitz et al., 2021b). Reversing the gradient from the neural reconstruction encourages the model to ignore variability in the activity that is irrelevant for decoding activity, which in turn results in an session-invariant encoding of behaviour-relevant factors.

Unlike other domain adaptation methods, our model does not require domain labels, but is trained on a single data set that contains different experimental sessions. This is an advantage in potential BCI applications as variability may not only exist between single sessions, but also within a single session, and in addition the degree of variability may differ as well. As a result, the model still requires considerable amounts of session data for good behaviour reconstruction. We found that performance did not saturate when it was trained on 12 sessions. In contrast, good behaviour decoding with our baseline RNN model could be achieved from a single session with less than 200 trials. Yet we expect that the number of trials per session required is much less than used here.

A main limitation of this method, which may also limit its direct application in a BCI system, is that it assumes that behaviour is generated by autonomous neural dynamics which relies on specification of appropriate initial conditions that form the latent variables in the model. This approach has been shown to successfully capture neural dynamics in a range of scenarios (Pandarinath et al., 2018) and has the advantage of a relatively compact and behaviorally relevant latent encoding that supports discovering invariant features in the neural activity. A possible extension to remove this limitation may be the inclusion of a controller input that models additional temporal dynamics to better account for behavioural variability (Pandarinath et al., 2018). This extension of the latent space can be trained in the same manner and may allow modelling of more complex and variable behavioural paradigms.

We compare our model to RAVE+ (Gonschorek et al., 2021), to our knowledge currently the only other method for domain adaptation of inter-session data. RAVE+ does show some indication of alignment when sufficient individual recording sessions are available, but its latent space fails to capture behaviourally relevant structure. As a result, behaviour decoding for unseen test data is poor. We suspect that the RAVE+ fails because the temporal dynamics in our data are too variable between trials. As pointed out by the authors, RAVE+ requires consistent temporal dynamics between trials, which can be controlled in experiments where stimuli are presented, but that are rarely obtained in behavioural experiments. The other models shown here (LFADS, RNN decoder) are included to contrast domain adaptation to conventional encoders, and not as a baseline for generalisation performance.

Our results are consistent with recent reports showing that motor control is based on low-dimensional latent neural dynamics that are very stable over time despite ongoing neural drift (Gallego et al., 2020). Our model can be used to discover these latent dynamics in data with high variability. Tests we performed on synthetic data (a Lorenz system)

indicate that this approach is also successful when neural dynamics are generated from latent dynamics with random transformations (not shown).

Our finding that SABLE has a better performance for intermediate held-out sessions than for sessions at the end of a sequence of training sessions suggests that performance will likely eventually decline for long time intervals between train and test sessions. As long as the latent dynamics are stable however, we expect that training the model with more sessions will eventually stabilise generalisation performance. Taken together these results are encouraging for BCI application as they suggest highly consistent recordings may not be required for good performance as long as it is possible to recover relevant latent dynamics.

References

- Chestek, C. A., Batista, A. P., Santhanam, G., Byron, M. Y., Afshar, A., Cunningham, J. P., Gilja, V., Ryu, S. I., Churchland, M. M., and Shenoy, K. V. Single-neuron stability during repeated reaching in macaque premotor cortex. *Journal of Neuroscience*, 27(40):10742–10750, 2007.
- Chestek, C. A., Gilja, V., Nuyujukian, P., Foster, J. D., Fan, J. M., Kaufman, M. T., Churchland, M. M., Rivera-Alvidrez, Z., Cunningham, J. P., Ryu, S. I., et al. Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *Journal of Neural Engineering*, 8(4):045005, 2011.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, 2014. doi: 10.3115/v1/d14-1179.
- Dabagia, M., Kording, K. P., and Dyer, E. L. Comparing high-dimensional neural recordings by aligning their low-dimensional latent representations. *arXiv preprint arXiv:2205.08413*, 2022.
- Degenhart, A. D., Bishop, W. E., Oby, E. R., Tyler-Kabara, E. C., Chase, S. M., Batista, A. P., and Yu, B. M. Stabilization of a brain-computer interface via the alignment of low-dimensional spaces of neural activity. *Nature Biomedical Engineering*, 4(7), 2020. doi: 10.1038/s41551-020-0542-9.
- Farshchian, A., Gallego, J. A., Miller, L. E., Solla, S. A., Cohen, J. P., and Bengio, Y. Adversarial domain adaptation for stable brain-machine interfaces. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.

- Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., and Miller, L. E. Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience*, 23(2), 2020. doi: 10.1038/s41593-019-0555-4.
- Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *32nd International Conference on Machine Learning, ICML 2015*, volume 2, 2015.
- Gonschorek, D., Höfling, L., Szatko, K. P., Franke, K., Schubert, T., Dunn, B. A., Berens, P., Klindt, D. A., and Euler, T. Removing inter-experimental variability from functional data in systems neuroscience. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=lVmIjQiJJSr>.
- Herrero-Vidal, P., Rinberg, D., and Savin, C. Across-animal odor decoding by probabilistic manifold alignment. *Advances in Neural Information Processing Systems*, 34, 2021.
- Hurwitz, C., Kudryashova, N., Onken, A., and Hennig, M. H. Building population models for large-scale neural recordings: Opportunities and pitfalls. *Current Opinion in Neurobiology*, 70:64–73, 2021a.
- Hurwitz, C., Srivastava, A., Xu, K., Jude, J., Perich, M., Miller, L., and Hennig, M. Targeted neural dynamical modeling. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Kingma, D. P. and Ba, J. L. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015.
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815, 2018.
- Rule, M. E., O’Leary, T., and Harvey, C. D. Causes and consequences of representational drift. *Current Opinion in Neurobiology*, 58:141–147, 2019.
- Sani, O. G., Abbaspourazad, H., Wong, Y. T., Pesaran, B., and Shanechi, M. M. Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nature Neuroscience*, 24(1), 2021. ISSN 15461726. doi: 10.1038/s41593-020-00733-0.
- Saxena, S. and Cunningham, J. P. Towards the neural population doctrine. *Current Opinion in Neurobiology*, 55: 103–111, 2019.
- Sussillo, D., Stavisky, S. D., Kao, J. C., Ryu, S. I., and Shenoy, K. V. Making brain-machine interfaces robust to future neural variability. *Nature Communications*, 7, 2016. ISSN 20411723. doi: 10.1038/ncomms13749.
- Wen, S., Yin, A., Furlanello, T., Perich, M., Miller, L., and Itti, L. Rapid adaptation of brain–computer interfaces to new neuronal ensembles or participants via generative modelling. *Nature Biomedical Engineering*, pp. 1–13, 2021.

A. Cross-subject decoding T-SNE Embedding

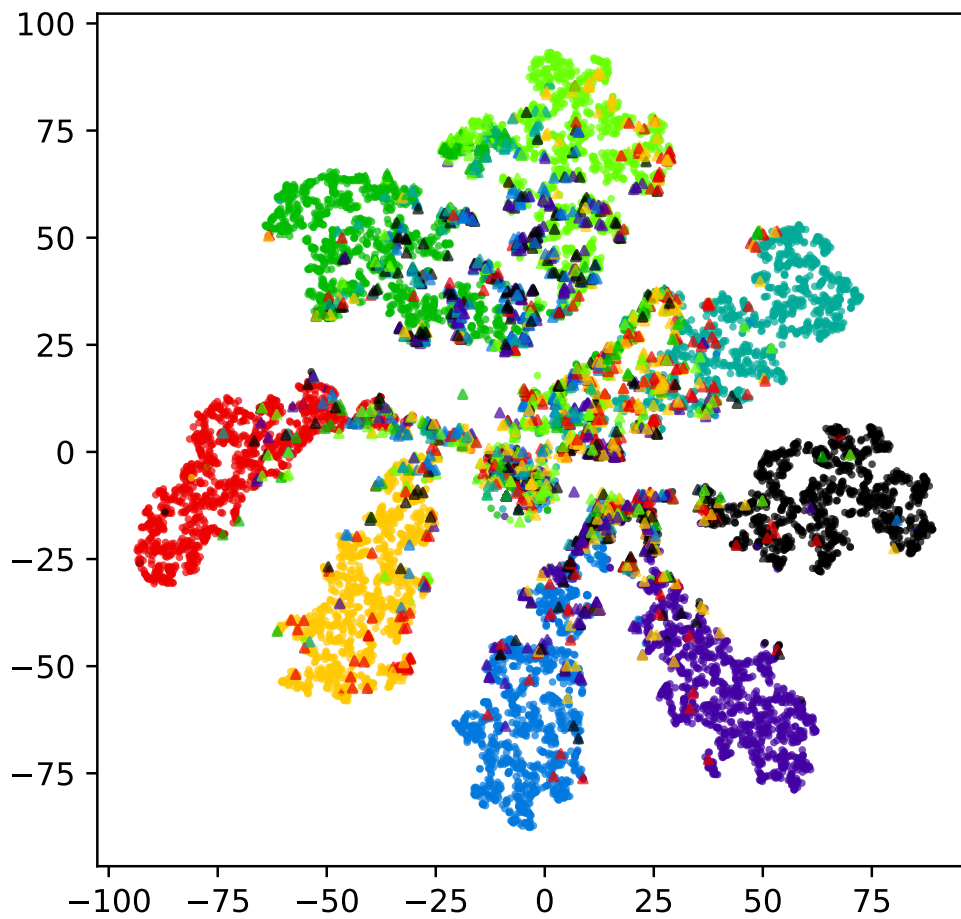


Figure 7. T-SNE embedding of SABLE latent space when training on 37 sessions of monkey 1 and testing on 14 sessions of monkey 2. Training trials are denoted by circles and test trials by triangles. Each colour denotes a particular movement direction.

B. Model Details

Below are implementation details for all models used in this paper.

SABLE		
Parameter	Value	Notes
Encoder		
- RNN Units	512 X 3	Stacked Gated Recurrent Unit
- RNN L2 Kernel Regularisation	1000	
- RNN L2 Recurrent Regularisation	1000	
- Recurrent Dropout	0.2	
- W_{enc} Units	512	Non-linear layer
- W_{enc} Dropout	0.8	
- W_{enc} L2 Regularisation	1000	
- Latent space dimension	64	
Neural Decoder		
- RNN Units	256	Gated Recurrent Unit
- RNN L2 Kernel Regularisation	0.1	
- RNN L2 Recurrent Regularisation	0.1	
- W_{fac} Units	128	Non-linear layer
- W_{fac} Dropout	0.2	
- W_{fac} L2 Regularisation	10	
Behaviour Decoder		
- RNN Units	256 X 2	Batch Normalisation on all layers Stacked Gated Recurrent Unit
- W_{beh} Units	512	Non-linear layer
- W_{beh} Dropout	0.1	
- W_{beh} L2 Regularisation	1.0	
Training		
Kullback–Leibler (KL) divergence weighting (λ_{kl})	0.01 to 10000	Rising exponentially (between encoder and neural de-coder)
Reverse Gradient weighting (λ_r)	1.0 to 0.000000001	Decaying exponentially

Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation

SABLE-noREV		
Parameter	Value	Notes
Encoder		
- RNN Units	512 X 3	Stacked Gated Recurrent Unit
- RNN L2 Kernel Regularisation	1000	
- RNN L2 Recurrent Regularisation	1000	Non-linear layer
- Recurrent Dropout	0.2	
- W_{enc} Units	512	
- W_{enc} Dropout	0.8	
- W_{enc} L2 Regularisation	1000	
- Latent space dimension	64	
Neural Decoder		
- RNN Units	256	Gated Recurrent Unit
- RNN L2 Kernel Regularisation	0.1	
- RNN L2 Recurrent Regularisation	0.1	Non-linear layer
- W_{fac} Units	128	
- W_{fac} Dropout	0.2	
- W_{fac} L2 Regularisation	10	
Behaviour Decoder		
- RNN Units	256 X 2	Batch Normalisation on all layers Stacked Gated Recurrent Unit
- W_{beh} Units	512	
- W_{beh} Dropout	0.1	Non-linear layer
- W_{beh} L2 Regularisation	1.0	
Training		
Kullback–Leibler (KL) divergence weighting (λ_{kl})	0.01 to 10000	Rising exponentially
Reverse Gradient weighting (λ_r)	N/A	Constant positive gradient of 1

LFADS		
Parameter	Value	Notes
Encoder		
- RNN Units	512 X 3	Stacked Gated Recurrent Unit
- RNN L2 Kernel Regularisation	1000	
- RNN L2 Recurrent Regularisation	1000	Non-linear layer
- Recurrent Dropout	0.2	
- W_{enc} Units	512	
- W_{enc} Dropout	0.8	
- W_{enc} L2 Regularisation	1000	
- Latent space dimension	64	
Neural Decoder		
- RNN Units	256	Gated Recurrent Unit
- RNN L2 Kernel Regularisation	0.1	
- RNN L2 Recurrent Regularisation	0.1	Non-linear layer
- W_{fac} Units	128	
- W_{fac} Dropout	0.2	
- W_{fac} L2 Regularisation	10	
Behaviour Decoder		
- RNN Units	256 X 2	Trained separately to rest of model Batch Normalisation on all layers Stacked Gated Recurrent Unit
- W_{beh} Units	512	
- W_{beh} Dropout	0.1	Non-linear layer
- W_{beh} L2 Regularisation	1.0	
Training		
Kullback–Leibler (KL) divergence weighting (λ_{kl})	0.01 to 10000	Rising exponentially
Reverse Gradient weighting (λ_r)	N/A	Constant positive gradient of 1

Robust alignment of cross-session recordings of neural population activity by behaviour via unsupervised domain adaptation

RAVE+		
Parameter	Value	Notes
Encoder		
- RNN Units	512 X 3	Stacked Gated Recurrent Unit
- RNN L2 Kernel Regularisation	1000	
- RNN L2 Recurrent Regularisation	1000	
- Recurrent Dropout	0.2	
- W_{enc} Units	512	Non-linear layer
- W_{enc} Dropout	0.8	
- W_{enc} L2 Regularisation	1000	
- Latent space dimension	64	
Neural Decoder		
- RNN Units	256	Gated Recurrent Unit
- RNN L2 Kernel Regularisation	0.1	
- RNN L2 Recurrent Regularisation	0.1	
- W_{fac} Units	128	Non-linear layer
- W_{fac} Dropout	0.2	
- W_{fac} L2 Regularisation	10	
Behaviour Decoder		Trained separately to rest of model
- RNN Units	256 X 2	Batch Normalisation on all layers
- W_{beh} Units	512	Stacked Gated Recurrent Unit
- W_{beh} Dropout	0.1	Non-linear layer
- W_{beh} L2 Regularisation	1.0	
Domain Classifier		
- Non-linear layer Units	256 X 2	Batch Normalisation
- Dropout	0.1	
- L2 Regularisation	0.001	
Training		
Kullback–Leibler (KL) divergence weighting (λ_{kl})	0.01 to 10000	Rising exponentially (between encoder and Domain Classifier)
Reverse Gradient weighting (λ_r)	1.0 to 0.000000001	Decaying exponentially