

Multi-modal Gaussian Process Variational Autoencoders for Neural and Behavioral Data

Rabia Gondur¹, Usama Bin Sikandar², Evan Schaffer³, Mikio
Aoi⁴, and Stephen Keeley⁵

¹Department of Computer and Information Sciences, Fordham University

²School of Electrical and Computer Engineering, Georgia Institute of Technology

³Zuckerman Mind Brain Behavior Institute, Columbia University

⁴Halicioğlu Data Science Institute, University of California, San Diego

⁵Department of Natural Sciences, Fordham University

Abstract

Characterizing the relationship between neural population activity and behavioral data is a central goal of neuroscience. While latent variable models (LVMs) are successful in describing high-dimensional time-series data, they are typically only designed for a single type of data, making it difficult to identify structure shared across different experimental data modalities. Here, we address this shortcoming by proposing an unsupervised LVM which extracts temporally evolving shared and independent latents for distinct, simultaneously recorded experimental modalities. We do this by combining Gaussian Process Factor Analysis (GPFA), an interpretable LVM for neural spiking data with temporally smooth latent space, with Gaussian Process Variational Autoencoders (GP-VAEs), which similarly use a GP prior to characterize correlations in a latent space, but admit rich expressivity due to a deep neural network mapping to observations. We achieve interpretability in our model by partitioning latent variability into components that are either shared between or independent to each modality. We parameterize the latents of our model in the Fourier

domain, and show improved latent identification using this approach over standard GP-VAE methods. We validate our model on simulated multi-modal data consisting of Poisson spike counts and MNIST images that scale and rotate smoothly over time. We show that the multi-modal GP-VAE (MM-GPVAE) is able to not only identify the shared and independent latent structure across modalities accurately, but provides good reconstructions of both images and neural rates on held-out trials. Finally, we demonstrate our framework on two real world multi-modal experimental settings: *Drosophila* whole-brain calcium imaging alongside tracked limb positions, and *Manduca sexta* spike train measurements from ten wing muscles as the animal tracks a visual stimulus.

1 Introduction

Recent progress in experimental neuroscience has enabled researchers to record from a large number of neurons while animals perform naturalistic behaviors in sensory-rich environments. An important prerequisite to analyzing these data is to identify how the high-dimensional neural data is related to the corresponding behaviors and environmental settings. Traditionally, researchers often employ a two-step approach, involving first dimensionality reduction on neural data followed by a post-hoc investigation of the latent structure of neural recordings with respect to experimental variables of interest such as stimuli or behavior [1–5]. Researchers may also specifically design latent variable models (LVMs) that extract low-dimensional structure in neural data using these experimental variables of interest [6–9]. However, recent deep learning advancements allow for both experimental and behavioral variables to be part of a single latent-variable model [10–13], thus opening a new avenue for unsupervised discoveries on the relationship between neural activity and behavior.

The existing approaches that jointly model neural activity and behavior are limited in that they often rely on a single latent space to describe data from both modalities [10, 11], making it difficult for practitioners to isolate latent features that are shared across and independent to neural activity or behavior. Approaches that do isolate modality specific and shared latent structure either do so with no temporal structure [14, 15, 13, 16–18], or use a relatively inflexible linear dynamical system [12]. Moreover, because of the deep neural network mapping from latents to observations, such existing

multi-modal approaches generate latent spaces that are not obviously related to experimental variables of interest, and so additional ad hoc model features are often added to aid in interpretability and analysis in neuroscience settings [10, 11, 8].

However, LVMS developed for neural data often are able to uncover interpretable features using neural activity alone in an unsupervised fashion with minimal *a priori* assumptions. One example of these is Gaussian Process Factor Analysis (GPFA), a widely used LVM in neuroscience that finds smoothly evolving latent dynamics in neural population activity, and can illustrate different aspects of neural processing [19–23]. GPFA constrains the latent space of neural activity through the use of a Gaussian Process (GP) prior. GP priors have also been adapted to regularize the latent space of variational autoencoders (GP-VAE or GP-prior VAE) with a variety of applications [24–27]. In each of these approaches, the GP prior provides a flexible constraint in the latent dimension, specifying correlations across auxiliary observations like viewing angle, lighting, or time. The GP prior often is used for out-of-sample prediction in GP-VAEs, but in GPFA is frequently used to visualize latent structure in noisy neural population activity on a trial-by-trial basis [20, 2, 21]. Here, we wish to leverage the interpretability seen in unsupervised GPFA models with the power of GP-VAEs to use with multi-modal time-series datasets in neuroscience.

We propose a model for jointly observed neural and behavioral data that partitions shared and independent latent subspaces across data modalities while flexibly preserving temporal correlations using a GP prior. We call this the multi-modal Gaussian Process variational autoencoder (MM-GPVAE). Our first innovation is to parameterize the latent space time-series in terms of a small number of Fourier frequencies, an approach that has been used before in the linear GPFA setting [28, 29, 21]. We show that the Fourier representation dramatically improves latent identification for the standard GP-VAE. Our second innovation is to augment our new Fourier GP-VAE model to describe two data modalities simultaneously by mapping the latents linearly to neural data (like in GPFA), as well as nonlinearly to another experimental variable via a deep neural network (like a GP-VAE). We leave the specific identity of this other experimental modality intentionally vague - it could be keypoints of limb positions as animals freely move or visual stimuli across time. Because this observation modality is characterized by a deep neural network, our model can be adapted to any additional experimental variable of interest.

We validate our MM-GPVAE model on a simulated dataset of a smoothly rotating and scaling MNIST digit alongside simulated Poisson neural activity, whereby the digit data and the spiking data share latent structure. We show that the MM-GPVAE model is able to recover both shared and independent latent structures while simultaneously providing accurate reconstructions of data from both modalities. Lastly, we demonstrate the utility of our model by fitting the MM-GPVAE to two real-world multi-modal neural datasets: 1) *Drosophila* (fly) whole-brain calcium imaging alongside tracked 16 limb positions, and 2) *Manduca sexta* (hawkmoth) spike train measurements from ten wing muscles alongside a continuously moving visual stimulus. In the former case, we demonstrate that the neural and shared subspaces best separate *Drosophila* behavioral conditions and in the latter case, we show distinct time-varying latent components tracking muscle and stimuli movement in the experiment. By showing the MM-GPVAE in these two domains, we demonstrate that our model is adaptable to a range of diverse experimental preparations in systems neuroscience.

2 The Gaussian Process Variational Autoencoder

The Gaussian Process variational autoencoder (GPVAE) uses high-dimensional data (e.g. images) accompanied by auxiliary information, like viewing angle, lighting, object identity, or observation time. This auxiliary information provides the indices in the GP prior latent representation, specifying correlations across the latent space and allowing for out-of-sample predictions at new auxiliary values [24–26]. However, in our case we consider continuously observed time-series data in experimental neuroscience experiments, so our auxiliary information here are evenly sampled time-bins. Because of this, we can leverage the advantages of the Fourier-domain GP representation [29, 22, 28].

Formally, consider smoothly-varying image data across timepoints, represented by the pixels-by-time matrix $\mathbf{Y} \in \mathbb{N}^{N \times T}$. For latent variable $\mathbf{z}(t) \in \mathbb{R}^P$ each latent $z_p(t)$ ($t \in \{1, 2, \dots, T\}$) evolves according to a Gaussian process, $z_p(t) \sim \mathcal{GP}(0, k_\theta(\cdot, \cdot))$, with covariance kernel k_θ . The time-by-time covariance matrix of each $z_p(t)$ is then given by the Gram matrix \mathbf{K}_θ corresponding to k_θ . We use a squared exponential (RBF) kernel for \mathbf{K} governed

by a marginal variance and length scale $\theta = \{\rho, \ell\}$ with an additive diagonal term $\alpha \mathbf{I}$ to help with inference [30, 24]. The likelihood of the image data is Gaussian with mean given by the latent values at any timepoint t passed through a deep neural network $g_\psi(\cdot)$ with parameters ψ and whose covariance is $\sigma_y^2 \mathbf{I}$.

The GP-VAE is learned using standard VAE amortized inference [31], where the parameters μ_ψ and σ_ψ^2 of a variational distribution q_ϕ are given as neural network functions of the observed data \mathbf{Y} . Here, the evidence lower bound (ELBO) is maximized with respect to the variational parameters ϕ , model parameters ψ , and GP hyperparameters. In this work, α is set to a fixed value of $1e-2$ for all experiments except for the final data analysis example, where it is set to a value of $1e-4$. While the ELBO may be expressed in a variety of ways, we will follow the form that includes the variational entropy term. For details about this approach for the standard GP-VAE, see [24].

2.1 Fourier-domain representation of the GP-VAE

We consider a version of the GP-VAE whose auxiliary variables are a Fourier frequency representation of the time domain, as opposed to timepoints sampled on a regular lattice. This allows us to parameterize a frequency representation of the latent variables that is Gaussian distributed according to a $\tilde{\mathbf{z}}(\omega) \sim \mathcal{N}(0, \tilde{\mathbf{K}}_p)$ where the original RBF GP covariance matrix \mathbf{K} may be diagonalized by $\tilde{\mathbf{K}} = \mathbf{B}\mathbf{K}\mathbf{B}^\top$. Here, \mathbf{B} is the orthonormal discrete Fourier transform matrix and $\omega \in 0, 1, 2 \dots F$ represents the frequency. The model prior can now be written as

$$p(\tilde{\mathbf{Z}} \mid \boldsymbol{\omega}, \boldsymbol{\theta}, \alpha) = \prod_{p=1}^P \mathcal{N}(\tilde{\mathbf{z}}_p \mid \mathbf{0}, \tilde{\mathbf{K}}_\theta(\boldsymbol{\omega}) + \alpha \mathbf{I}_F) \quad (1)$$

Where $\boldsymbol{\omega}$ represents a F -values long vector of Fourier frequencies, and $\tilde{\mathbf{Z}}$ represents a $P \times F$ matrix of the frequency representation of the P latent variables at the \mathcal{F} frequencies. The model likelihood retains the same form as the standard GP-VAE.

$$p(\mathbf{Y} \mid \mathbf{Z}, \phi, \sigma^2) = \prod_{t=1}^T \mathcal{N}(g_\phi(\mathbf{z}_t), \sigma^2 \mathbf{I}_T) \quad (2)$$

Where \mathbf{z}_t represents the t^{th} time index of \mathbf{Z} , a $P \times T$ matrix where each row is the time-domain representation of the p^{th} latent, given by $\mathbf{z}_p = \mathbf{B}\tilde{\mathbf{z}}_p$.

We define the variational posterior so that it factorizes over the Fourier frequencies. The variational distribution can be written as

$$q_\psi(\tilde{\mathbf{Z}}^i | \mathbf{Y}^i) = \prod_{\omega} \mathcal{N}(\tilde{\mathbf{z}}_{\omega}^i | \tilde{\boldsymbol{\mu}}_{\psi}(\mathbf{Y}^i), \text{diag}(\tilde{\boldsymbol{\sigma}}_{\psi}^2(\mathbf{Y}^i))), \quad (3)$$

where i indexes a trial of times-series images \mathbf{Y} . Training is done by batches over a subset of trials. In contrast to the standard GP-VAE [24], the Fourier representation requires that, for each trial, images at all timepoints are mapped to a single $P \times F$ -dimensional Fourier representation. We accomplish this in two steps - first, we use a deep neural network for each image at each time point \mathbf{y}_t^i , which will result in T total network embeddings for a single trial, each of dimension P . We follow that with a single linear layer each for the mean and variance of the variational distribution, $l_{\tilde{\boldsymbol{\mu}}}(\cdot), l_{\tilde{\boldsymbol{\sigma}}^2}(\cdot)$, that will map from the number of timepoints to the number of Fourier frequencies. $l_{\tilde{\boldsymbol{\mu}}}, l_{\tilde{\boldsymbol{\sigma}}^2} : \mathbb{R}^{P \times T} \rightarrow \mathbb{R}^{P \times F}$. A schematic for the Fourier domain variational and generative architecture is displayed in Figure 1(a).

There are a number of advantages to representing the GP-VAE latents and variational parameters in the Fourier domain. For one, the diagonal representation of \mathbf{K} means we avoid a costly matrix inversion when evaluating the GP prior [29, 28, 32–34]. Secondly, we can *prune* the high frequencies in the Fourier domain, effectively sparsifying the variational parameters, and hence frequencies ($F < T$), while enforcing a smooth latent representation. Though this is a free parameter in our model thus if the user would like to model steady-states they are able to easily adjust this. Lastly, the Fourier representation of the variational parameters mean we can retain both temporal correlations and the advantages of using a mean-field approximation for the variational posterior q_ϕ [21]. These advantages have been seen in simpler GP models. For more information, see [29, 21].

The Fourier-represented GP-VAE dramatically improves the ability of the GP-VAE to learn a true underlying smoothly evolving latent for high-dimension data in a non-linear model. We demonstrate this on a simulated example using an MNIST digit that rotates by a time-varying angle given by a draw from a GP with an RBF kernel. We fit our Fourier GP-VAE model as well as the standard VAE and standard GP-VAE models to these trials. In each case, we are looking to recover the true underlying generative latent

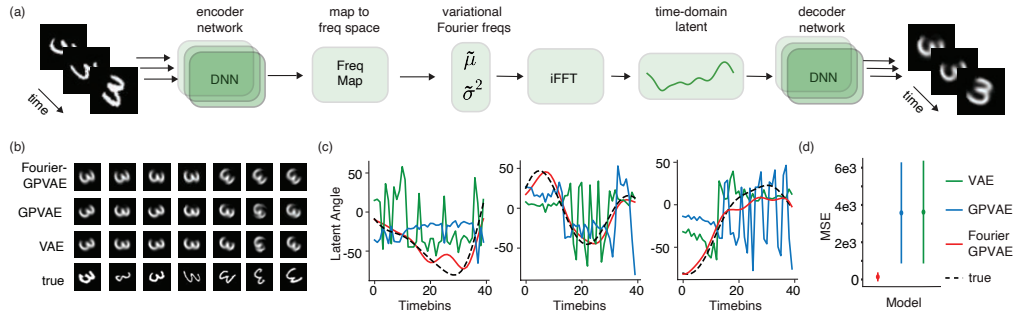


Figure 1: (a) schematic for the Fourier domain GP-VAE. All images at all timepoints for a given trial are encoded via a deep neural network into variational parameters of a pruned Fourier representation of the latent space. This Fourier representation is then mapped back into the time domain before being passed through a decoder network to give the image reconstruction at each timepoint. (b) Image reconstructions of the standard VAE, GP-VAE and Fourier domain GP-VAE. (c) Estimated latent for each model alongside the true underlying latent angle. (d) Mean squared error (MSE) of estimated latents and true latents for 60 held out trials. Error bars indicate standard error.

angle and hence we assert a one-dimensional latent dimensionality for the model. For more information on training and testing, see the supplement.

Because the Fourier variational distribution preserves temporal correlations and prunes the high-frequency components of the latents \mathbf{z} , the Fourier GP-VAE model shows much smoother underlying latent representations on held-out trials (Fig 1 (c)), though each model retains the ability to accurately reconstruct the images through the network mapping from the latent space (Fig 1(b)). When the latent space is mapped through an affine transform to align latents on held-out trials to the true latent space, the Fourier domain GP-VAE does much better uncovering the true smoothly evolving latent angle (Fig 1 (c), (d)).

3 The Multi-Modal Gaussian Process Variational Autoencoder

We now focus on extending the GP-VAE to model data of two modalities simultaneously. The examples we will emphasize here involve neural activity alongside some other experimental variable such as naturalistic movement or high-dimensional stimuli. The observations of our model are two distinct data modalities represented by matrices $\mathbf{y}_A^{(i)} \in \mathbb{R}^{N \times T}$ and $\mathbf{y}_B^{(i)} \in \mathbb{R}^{M \times T}$. Here, $t \in (1, 2 \dots T)$ denotes time-bin indices, i denotes trials, and N and M denote the dimension of the observations for each modality. We assume, as before, that all data are generated by smoothly varying low-dimensional latent variables that are now either independent to, or shared between data modalities. As before, latents are initially parameterized in a pruned Fourier representation before being mapped to the time domain. The latents are then partitioned by a loadings matrix \mathbf{W} which linearly combines the shared latent representation with the independent latents.

$$\begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \end{bmatrix} = \begin{bmatrix} \mathbf{W}_A & \mathbf{W}_{S1} & 0 \\ 0 & \mathbf{W}_{S2} & \mathbf{W}_B \end{bmatrix} \begin{bmatrix} \mathbf{z}_A \\ \mathbf{z}_S \\ \mathbf{z}_B \end{bmatrix} + \mathbf{d} \quad (4)$$

Here, \mathbf{z}_S refers to latents that are shared between regions, while \mathbf{z}_A and \mathbf{z}_B denote region-specific latent variables and \mathbf{d} represents an additive offset. The outputs after this linear mixing result in modality-specific embeddings which we call \mathbf{x}_A and \mathbf{x}_B . Here, modality A , the 'behavioral modality' is passed through a deep neural network, and modality B , the 'neural modality', is passed through a pointwise nonlinearity f to enforce non-negativity of Poisson rates. The likelihood of the data given the embedding is:

$$p(\mathbf{y}_A | \mathbf{x}_A) \sim \mathcal{N}(g_\psi(\mathbf{x}_A), \sigma_y^2 \mathbf{I}_N), \quad p(\mathbf{y}_B | \mathbf{x}_B) \sim \mathcal{P}(f(\mathbf{x}_B)), \quad (5)$$

Where the function $g_\psi(\cdot)$ represents a decoder neural network for the experimental data and we use the exponential nonlinearity for f . Learning is performed as before - the mean and variance of a mean-field Gaussian variational distribution is given by the data passed through a neural network with parameters ϕ . Here, our lower bound is similar to that found in Casale et al. (2018) with an additional term for the neural data modality.

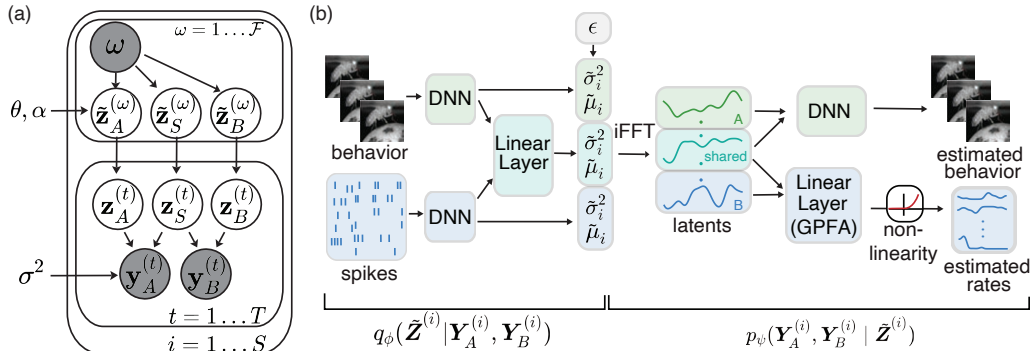


Figure 2: (a) Graphical model of the multimodal GP-VAE. A set of Fourier frequencies describe the Fourier representation of shared and independent latents across modality with a GP prior over each latent. Latents are transformed to the time domain and combined to generate data for each modality. (b) Schematic of the MM-GPVAE.

$$\begin{aligned}
 \text{ELBO} = \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\phi} & \left[\underbrace{\sum_t \log \mathcal{N}(\mathbf{y}_A | g_\psi(\mathbf{x}_A), \sigma_y^2 \mathbf{I}_N)}_{\text{Gaussian Likelihood (other modality)}} + \right. \\
 & \underbrace{\sum_t \log(\mathcal{P}(\mathbf{y}_B | f(\mathbf{x}_B)))}_{\text{Poisson Likelihood (Neural Rates)}} + \underbrace{\log p(\tilde{\mathbf{Z}} | \mathbf{T}, \boldsymbol{\theta})}_{\text{GP Prior}} + \underbrace{H(q_\phi)}_{\text{Entropy}} \left. \right]
 \end{aligned} \tag{6}$$

The expectation with respect to the variational distribution is performed by drawing stochastic samples ϵ from a standard normal and using the variational mean and variance to get a sample from q_ϕ (the so-called reparameterization trick [31]). However, because we use a mean-field Gaussian variational distribution, both the Poisson log-likelihood and the GP prior terms in the expectation can be computed in closed-form and preclude the need for sampling for these terms. For clarity, we omit that closed-form expression here. See the supplement for details.

A graphical depiction of the generative model of the MM-GPVAE is shown in figure 2(a). The hyperparameters of the model are the latent-specific kernel parameters θ that provide the GP covariance structure as well as a constant α additive diagonal offset (fixed during inference). The entire schematic of the

MM-GPVAE model, including the generative and variational distributions, is depicted in figure 2(b).

4 Experiments

4.1 Simulated data

We first assess the performance of the MM-GPVAE model using a simulated dataset of two modalities: a smoothly rotating and scaling MNIST digit, and simultaneously accompanied Poisson spike counts from 100 neurons. A total of three latents were drawn from an underlying GP kernel to generate single-dimensional shared, neural, and image subspaces. There were a total of 300 trials, each trial consisting of 60 timebins. The data was split into 80% for training and 20% for testing. For this simulated example, one latent represents an interpretable modulation of the image as it directly effects the scaling of the MNIST digit. So as the latent values change along the trial, the image scales smaller and larger continuously. Another latent, corresponding to the independent component of the neural modality, provides one component of the log-rates of the Poisson spiking data. The final latent reflects the shared variability between both modalities. Here, the latent is again interpretable with respect to the image in that this latent corresponds to the angle of the smoothly rotating MNIST digit. This shared latent also linearly combines with the neural-only subspace to provide the the log-rates of the spiking Poisson population. The MM-GPVAE trained simultaneously on both images and spikes can successfully recover the three-dimensional latent structure across these two modalities. Figure 3(a) shows the underlying true latent in a held-out trial alongside the estimated latent structure extracted from the MM-GPVAE. As before, the held-out latent space is scaled to align with the true latent space as closely as possible, as the MM-GPVAE model latent space is invariant to scaling transform. In addition to accurately extracting the latent structure, the MM-GPVAE also has the ability to accurately reconstruct both neural rates and images across the time-series from the latent space. Four example images and three example neural rates are show in Figure 3 (b). For more examples, see the supplemental materials.

The MM-GPVAE is also able to exploit information across modalities to better model the observed data. The left side of 3(c) shows the mean squared error (MSE) across pixels for reconstructed images in held-out trials. By

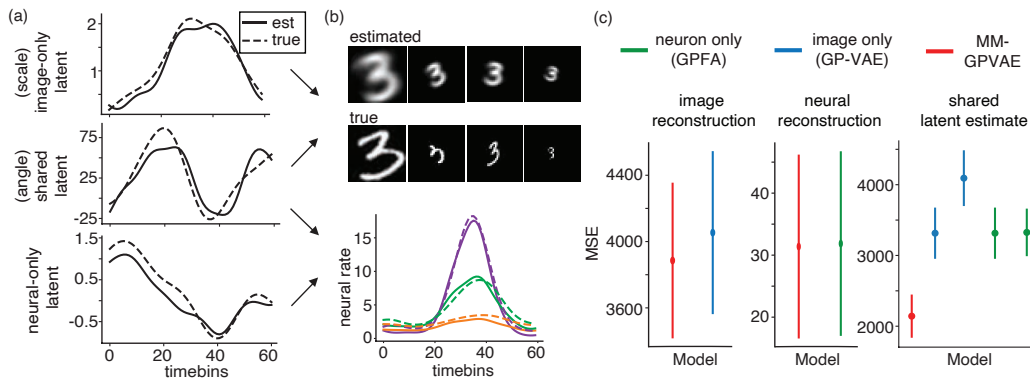


Figure 3: (a) True and estimated latents for the MM-GPVAE trained on simulated neural spiking data as well as a smoothly scaling and rotating MNIST digit. (b) Estimated neural rates on an example trial for 3 example neurons as well as 4 example reconstructed images of the MNIST digit at different angles and scales. (c) (*left*) Reconstruction accuracy from the image data trained on the images alone (GPVAE) compared to training with both modalities simultaneously (MM-GPVAE). (*middle*) Accuracy of estimated neural rates (*left*) trained on neural-activity alone (Poisson - GPFA) compared to MM-GPVAE. (*right*). Accuracy of shared latent estimated from the MM-GPVAE compared to single-modality model variants. Error bars are standard error.

leveraging latent information from the spiking modality, the MM-GPVAE is able to reconstruct the images better than a single-modality GP-VAE trained on the images alone [24]. Similarly, the held-out neural rates are better estimated when the image data is included, though the effect here is more modest than for the image modality. Finally, we also show that the shared latent variable can be more accurately identified when data from both modalities is used (Figure 3(c)). Here we assess the accuracy of the shared latent estimate on test trials when we use data from each modality alone (GPVAE and GPFA) compared to both modalities simultaneously (MM-GPVAE). As expected, data from both modalities better identifies shared structure across the data types. The settings for all three models were exactly the same for these comparisons. We only ablated the models to make it into uni-modal settings. For additional details regarding training and the neural network set-ups for this example, please see the supplemental material.

4.2 Application to fly experimental data

Next we look to evaluate the MM-GPVAE on a real-world multi-modal dataset. Here, we consider a whole-brain calcium imaging from an adult, behaving *Drosophila* [35]. We isolate 1000 calcium traces recorded using SCAPE microscopy [36] from an animal while it performs a variety of distinct behaviors fixed on a spherical treadmill. These 1000 calcium traces are dispersed widely and uniformly across the central fly brain and were randomly picked from the dataset to use in our model. Alongside the neural measurements, eight 2-d limb positions are extracted from a recorded video using the software tool DeepLabCut [37]. These simultaneously recorded calcium traces and limb position measurements are split into 318 trials of 35 time-bins sampled at 70hz. Each trial has one of 5 corresponding behavioral labels (still, running, front grooming, back grooming and abdomen bending) determined via a semi-supervised approach from the tracked limb position measurements [38]. Importantly, these behavioral labels are not used when fitting the MM-GPVAE. For additional information on how we isolate this behavioral data and neural traces, see the supplemental materials.

For this set-up, we consider a variant of the MM-GPVAE where we have removed the non-linearity for the neural data modality, and instead consider Gaussian observations with an additional parameter controlling the observation variance for the fluorescence traces (akin to the original formulation of GPFA [30]). We do this because calcium imaging traces take on continuous values and are not well-described by Poisson observations, though in principle other observation likelihoods could be used here [39]. Next, we consider a 7-dimensional latent behavioral subspace and 26 dimensional neural subspace, with 5 of these dimensions being shared across modalities. These choices were made through initial exploration of the model and examining cross-validated model performance, though the results we show are robust to a wide range of dimensionality choices for each of the subspaces.

The MM-GPVAE is able to successfully reconstruct both behavioral trajectories and calcium traces in held-out trials for these data. The top of Figure 4(c) shows the true and decoded 16 limb position measurements with 5 highlighted for clarity. The bottom part of figure 4(c) shows 6 randomly selected calcium traces alongside their model reconstructions. In each case, the MM-GPVAE can roughly capture the temporal trends in this multi-modal dataset. To analyze the shared and independent latent subspaces of these data we consider a 2d projection of each latent space, and plot the mean latent

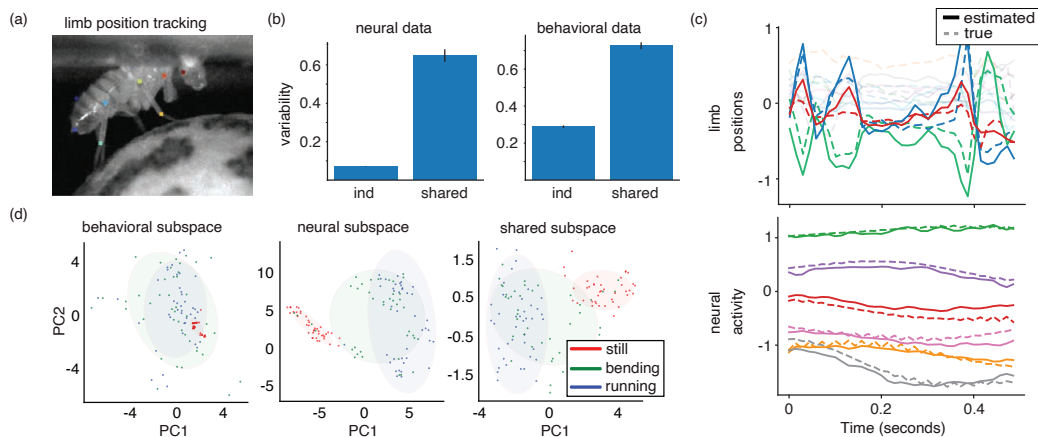


Figure 4: (a) Limb position tracking in *Drosophila* (b) (*top*) Contribution of the variability in the data across trials of the neural (*left*) and behavioral (*right*) modalities due to shared and independent subspaces (c) (*top*) Limb position estimates and true values and (*bottom*) six randomly selected calcium trace estimates and true values for a given trial. (d) Visualization of average latent value across time in the neural, behavioral and independent subspaces for 3 behavioral categories.

value in that subspace calculated each trial. We additionally color-code the trial according to the behavioral label. In these shared and independent latent subspaces, we find that the "still" behavioral state is well separated from many of the other behavioral conditions in the neural and shared subspaces. Here, we show two others (bending and running) as an illustration (Figure 4(d)). We refer the reader to the supplement for visualizing all 5 behaviors in the latent space. Lastly, we illustrate how much each subspace contributes to the overall variability of neural and behavioral data. Figure 4(b) shows the variance of the neural (left) and behavioral (right) data calculated for each trial where the reconstruction of the data is generated either from the shared subspace, or either of the independent subspaces. Here, we find that the shared subspace for each modalities contributes more to the overall variability in the data than either of the independent subspaces, suggesting a large fraction of shared variability between the data modalities.

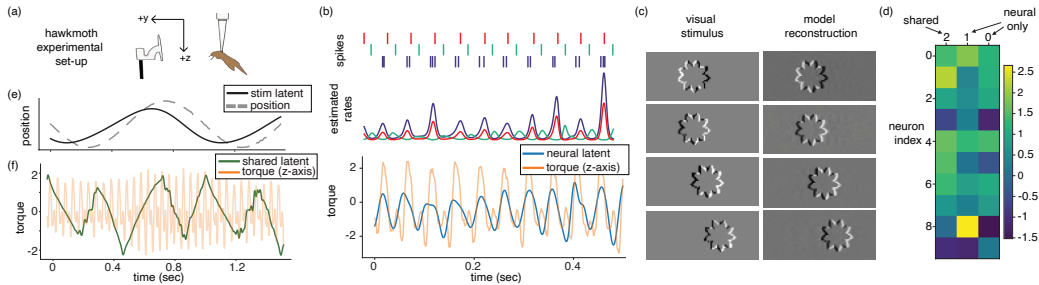


Figure 5: (a) Experimental set-up for *Manduca sexta*. Spikes from 10 muscle groups are recorded as an animal tracks a 1 Hz moving flower stimulus (b) *top*: Spikes from 3 example motor neurons in a short (half-second) segment of the recording show periodic activity at approximately 22 Hz. *middle*: estimated Poisson rates captured by the MM-GPVAE. *bottom*: The neural latent and a torque measurement from the hawkmoth. The strong ~ 22 Hz modulation reflects the flapping of the moth wings. (c) Visual stimulus and reconstructions from the MM-GPVAE (d) Weights of hawkmoth spike decoder for neural-only and shared latents. Values here indicate how strongly each individual neuron is modulated by the shared or independent latent subspaces. (e) A one dimensional latent from the visual-modality subspace closely tracks the stimulus position. (f) The shared latent between modalities plotted alongside the torque measurement.

4.3 Application to moth experimental data

Next, we evaluate the MM-GPVAE on a dataset of a hawkmoth (*Manduca sexta*) tracking a moving flower. The hawkmoth is an agile flier that is able to closely follow swiftly moving targets while hovering in midair, making it a model organism for the study of sensorimotor control [40]. The modalities in our dataset consist of a time-series of images of a visual stimulus temporally paired with electromyography (EMG) signals recorded during 20-second experimental trials. Each image is an event-based reconstruction (inspired by insect visual system) of a white floral target moving laterally during a trial as viewed in the hawkmoth’s frame of reference [41] (Fig 5(a)). The floral target moves sinusoidally at 1 Hz against a black background. Each pixel in a video frame can occupy a polarity of -1, 0 or +1 based on whether its luminosity has respectively decreased, unchanged or increased as compared to the previous frame in the temporal sequence. Temporally paired with the

floral stimulus, EMG recordings are precisely measured spike trains from the 10 major flight muscles of the hawkmoth responding to the floral target while flying on a tether and flapping at about 22 Hz. Altogether, the 10-muscle recordings form a near-complete motor program that controls its flight during target tracking response [42].

We fit the MM-GPVAE model to these data using a 2-dimensional latent space to describe the moving visual stimulus and a 3-dimensional latent space to describe the Poisson rates from the 10 motor units. Of these, one latent dimension is shared between modalities. Here, a small number of dimensions are used both because only a few dimensions were needed to be able to accurately reconstruct the data of each modality, and also because it allows for easier interpretation of the latent dynamics with respect to the experiment. We split a 20 second recording into 6 trials of equal length for model fitting. Here, it is worth noting that the trials are much longer (~ 3300 timebins) than in the previous experiments. To encourage the visual-stimulus latent space to capture the slow dynamics of the moving flower, and the neural latent space to capture the fast dynamics of the flapping wings, we initialized the length-scale (ℓ) parameter of the GP kernel at a value of 300 for the image latent, 10 for the neural latent, and 150 for the latent component shared between modalities. We find that the MM-GPVAE is able to accurately reconstruct both the neural rates (Fig 5(b)) and visual stimulus (Fig 5(c)) from its low dimensional latent space. We see that the neural rates are strongly modulated by the wing-flap oscillation at approximately 22 Hz. This is of course also seen in the neural latent space identified by the MM-GPVAE. At the bottom of Figure 5(b) we see that the first PC of the 2-dimensional latent space closely aligns with the torque measurement along the z-axis of the moth motion, a measurement previously identified to be closely associated with wing beating [42]. The MM-GPVAE is also able to reconstruct the visual stimulus data, as seen in 5(c). When viewing the latent space corresponding to the visual stimulus modality, we find that the value of the smoothly evolving latent variable closely tracks the position of the stimulus in x dimension (the only dimension along which it moves, see Figure 5(e)). Finally, the shared latent dimension shows some modulation with wing-flapping, but exhibits temporal structure not obviously related to the flapping frequency or stimulus motion. This may correspond to a longer-timescale motor behavior such as the variation in how strongly the hawkmoth is responding to track the stimulus, which could depend on the degree of its attention to the stimulus and its state of motivation. Importantly, we find that including a shared

dimension better reconstructs the data than not including one, suggesting that there is a component of neural dynamics shared with stimulus dynamics. Specifically, we find that if we remove the shared latent we see approximately a 70% drop in mse in reconstruction of the image modality and about a 5% drop for reconstruction of the rates. Because the neural rates have such a strong wing-flapping modulation, we expect the slower-moving dynamics to account for only a small amount of the variability here.

Lastly, to emphasize the added interpretability we get from the linear layer in the neural likelihood (GPFA), we plot the weights W_{s2} and W_B that map the shared and independent neural latents to the log-rate of the hawkmoth neurons. This is easiest to visualize in this example dataset because there is only a small neural population from which we record spikes. Figure 5 (d) shows each latent loading on to each of the ten neurons. Here, we z-score the weights along each dimension for easy visualization. We see that certain neurons are more heavily driven by the neural-only latents whereas others are more strongly modulated by the shared latent component. Future analysis of these data may help us further uncover how these shared and independent latent dynamics relate to the hawkmoth tracking and specific sensorimotor neurons and function in this experiment.

5 Comparison to Existing Approaches

Our work contrasts in important ways with similar recently-developed multi-modal latent variable models for neural time series experiments. One closely related model, *Targeted Neural Dynamical Modeling* (TNDM) [11], was specifically designed to non-linearly separate the neural latent representation into behaviorally relevant and irrelevant subspaces. However, this work was evaluated on a dataset with a relatively simple behavioral paradigm, and fails to reconstruct complex behavioral modalities (so is not well suited to reconstructing stimuli, for example). Another related multi-modal time-series neuroscience model, *Preferential Subspace Identification* (PSID) [12], was designed to linearly separate the neural latent representation into behaviorally relevant and irrelevant subspaces. However, PSID is restricted in its ability to describe more complex dynamics, and similarly to TNDM only performs well when behavioral or stimulus modalities have simple structure due to lack of a DNN decoder. To demonstrate the capabilities of MM-GPVAE alongside these completing methods, we evaluate the MM-GPVAE, TNDM, and PSID

on the simple dataset used in TNDM paper as well as our simulated rotating and scaling MNIST example.

We first compare our MM-GPVAE with TNDM [11] and PSID [12] using our simulated data of rotating and scaling MNIST digit ‘3’ and simultaneous Poisson neural rates to show the reconstruction performance of our model in comparison to these two competing models. Error on both data modalities as well as reconstructions of each modality on held-out trials can be found in the supplement.

To further compare these models in a real-world setting, we then implement PSID, TNDM, and MMGPVAE in a simple neuroscience example where all models are capable of recovering behavioral trajectories as well as neural rates. We use the primate reaching data used in the TNDM paper. Here, the 2D reaching task is simple, and so linear dynamics as well as our deep neural network can reconstruct both modalities well. The results from this can be seen in Figure 6. While both TNDM and PSID can provide only two subspaces in which only one can separate the 8 reaching positions (relevant subspace), MM-GPVAE can provide visualization for 3 subspaces, 2 modality specific, and 1 shared, and we can see a clear separation with both the behavioral-only subspace as well as shared subspace.

6 Conclusion

In this work, we have introduced the multi-modal Gaussian Process variational autoencoder (MM-GPVAE) to identify temporally evolving latent variables for jointly recorded neural activity and behavioral or stimulus measurements. We parameterized the model in the Fourier domain to better extract identifiable temporal structure from high-dimensional time-series data. We first demonstrated that the Fourier-mean field representation is able to better recover the true latent variables in a single-modality setting compared to previous approaches in the time domain. We then show that our multi-modal GPVAE can also accurately recover latent structure across data modalities, which we demonstrate using a smoothly rotating and scaling MNIST digit alongside simulated neural spike trains. We then show that the model can be flexibly adapted to multiple real-world multi-modal behavioral settings. We implement the MM-GPVAE on calcium imaging traces and tracked limb positions of *Drosophila* and visualize the latent embeddings of distinct fly behaviors in shared and independent latent subspaces. We also use the MM-

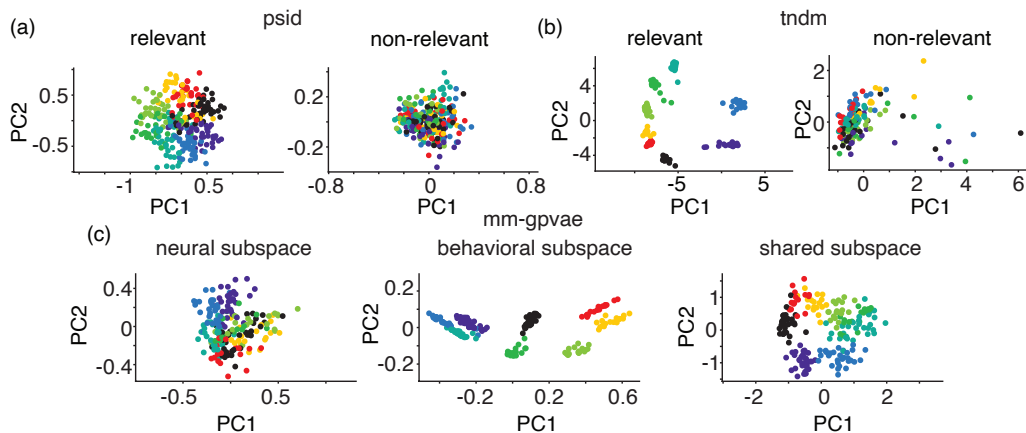


Figure 6: (a) Separation of all 8 reaching directions in the relevant subspace using PSID. (b) Separation of all 8 reaching directions in the relevant subspace using TNDM. (c) Separation of all 8 reaching directions in behavior-only and shared subspaces. Dots here indicate the mean latent value across the entire trial. The neural subspace shows no behavioral separation in the latent-space whereas the behavioral and shared subspaces show strong behavioral separation. This result closely parallels with [11, 12], which isolates behaviorally relevant and irrelevant neural subspaces. However, in contrast to [11, 12] the MM-GPVAE isolates a distinct shared subspace as well as both neural and behavioral independent subspaces from a raw unsupervised partitioning of both the behavioral and neural data.

GPVAE on spikes from hawkmoth flight muscles as the animal tracks an oscillating visual stimulus. Here, our model extracted latent variables from each modality that corresponds to distinct experimentally relevant variables - wing flapping and a much slower stimulus movement. Here, we also demonstrated that the MM-GPVAE can be adapted to long time-series data to identify latent variables at widely varying timescales.

Choice of prior: In this work we focus on the use of GP priors to describe dynamics, though many other choices could have been made. As stated above, the existing multi-modal time-series models in neuroscience using linear dynamics and RNNs have limitations when compared with the MM-GPVAE. However, other dynamical approaches such as switching linear dynamical systems or non-linear dynamics models exist in neuroscience and are powerful in single modality settings [9, 43]. Such models contain dependencies across

the latent dynamical variables, and so require reformulation and developing new inference to adapt them to multi-modal settings, and do not afford the interpretability and scalability gains seen with the Fourier approach. Thus, development of multi-modal versions of different dynamical models and evaluating their strengths and limitations alongside the MM-GPVAE remains an important avenue of future work in neuroscience.

Limitations: One limitation of our model is that it is specifically designed for temporally aligned time-series data, therefore requiring the data from both modalities to be recorded at matching time intervals or needing resampling to align them (see supplement). However, techniques such as time warping can be implemented to address this limitation. The use of the GP prior with RBF kernel and pruned Fourier representation also assumes latent dynamics are varying smoothly in time, precluding the appropriateness of our model for data with abrupt state-changes. Lastly, despite the improvements of the Fourier domain representation in our model, model performance can be sensitive to optimization step-size, covariance parameter α , as well as initial length-scale parameter values, pruning thresholds and marginal variance σ_y^2 . See supplement for more information and example code.

Broader impact: Our work provides a general model for exploratory data analysis for multi-modal time-series data. Though we develop the model with neuroscience experiments in mind, in principle the model could be adapted to other settings where the assumption of smoothly evolving latent dynamics holds across distinct data modalities. We do not anticipate any potential negative societal impacts of our work.

Additional notes on inference: All models in this manuscript were trained end-to-end in PyTorch using the ADAM optimizer. Training was done on a Macbook Pro with Apple M1 max chip and all evaluations took less than an hour to fit. All encoder and decoder neural networks were standard feedforward neural networks with ELU activation functions. Additional details on neural network architectures for the experiments can be found in the supplement. The code for MM-GPVAE will be available upon formal publication.

References

- [1] Anqi Wu, Nicholas A Roy, Stephen Keeley, and Jonathan W Pillow. Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in neural information processing systems*, 30, 2017.
- [2] Yuan Zhao and Il Memming Park. Variational latent gaussian process for recovering single-trial dynamics from population spike trains. *Neural computation*, 29(5):1293–1316, 2017.
- [3] Shreya Saxena and John P Cunningham. Towards the neural population doctrine. *Current opinion in neurobiology*, 55:103–111, 2019.
- [4] Giwon Bahg, Daniel G Evans, Matthew Galdo, and Brandon M Turner. Gaussian process linking functions for mind, brain, and behavior. *Proceedings of the National Academy of Sciences*, 117(47):29398–29406, 2020.
- [5] Felix Pei, Joel Ye, David Zoltowski, Anqi Wu, Rameed H Chowdhury, Hansam Sohn, Joseph E O’Doherty, Krishna V Shenoy, Matthew T Kaufman, Mark Churchland, et al. Neural latents benchmark’21: evaluating latent variable models of neural population activity. *arXiv preprint arXiv:2109.04463*, 2021.
- [6] Mikio C Aoi, Valerio Mante, and Jonathan W Pillow. Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making. *Nature neuroscience*, 23(11):1410–1420, 2020.
- [7] Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474):78–84, 2013.
- [8] Ding Zhou and Xue-Xin Wei. Learning identifiable and interpretable latent models of high-dimensional neural activity using pi-vae. *Advances in Neural Information Processing Systems*, 33:7234–7247, 2020.
- [9] Daniel Hernandez, Antonio Khalil Moretti, Ziqiang Wei, Shreya Saxena, John Cunningham, and Liam Paninski. Nonlinear evolution via spatially-dependent linear dynamics for electrophysiology and calcium data. *arXiv preprint arXiv:1811.02459*, 2018.
- [10] Steffen Schneider, Jin Hwa Lee, and Mackenzie Weygandt Mathis. Learnable latent embeddings for joint behavioural and neural analysis. *Nature*, pages 1–9, 2023.

- [11] Cole Hurwitz, Akash Srivastava, Kai Xu, Justin Jude, Matthew Perich, Lee Miller, and Matthias Hennig. Targeted neural dynamical modeling. *Advances in Neural Information Processing Systems*, 34:29379–29392, 2021.
- [12] Omid G Sani, Hamidreza Abbaspourazad, Yan T Wong, Bijan Pesaran, and Maryam M Shanechi. Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nature Neuroscience*, 24(1):140–149, 2021.
- [13] Jonnathan Singh Alvarado, Jack Goffinet, Valerie Michael, William Liberti III, Jordan Hatfield, Timothy Gardner, John Pearson, and Richard Mooney. Neural dynamics underlying birdsong practice and performance. *Nature*, 599(7886):635–639, 2021.
- [14] Gregory Gundersen, Bianca Dumitrascu, Jordan T Ash, and Barbara E Engelhardt. End-to-end training of deep probabilistic cca on paired biomedical observations. In *Uncertainty in artificial intelligence*, 2019.
- [15] Michael Kleinman, Alessandro Achille, Stefano Soatto, and Jonathan Kao. Gacs-korner common information variational autoencoder. *arXiv preprint arXiv:2205.12239*, 2022.
- [16] Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. *Advances in neural information processing systems*, 31, 2018.
- [17] Yuge Shi, Brooks Paige, Philip Torr, et al. Variational mixture-of-experts autoencoders for multi-modal deep generative models. *Advances in neural information processing systems*, 32, 2019.
- [18] Manuel Brenner, Georgia Koppe, and Daniel Durstewitz. Multimodal teacher forcing for reconstructing nonlinear dynamical systems. *arXiv preprint arXiv:2212.07892*, 2022.
- [19] Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Advances in neural information processing systems*, 21, 2008.
- [20] Lea Duncker and Maneesh Sahani. Temporal alignment and latent gaussian process factor inference in population spike trains. In *Advances in Neural Information Processing Systems*, pages 10445–10455, 2018.

- [21] Stephen Keeley, Mikio Aoi, Yiyi Yu, Spencer Smith, and Jonathan W Pillow. Identifying signal and noise structure in neural population activity with gaussian process factor models. *Advances in neural information processing systems*, 33:13795–13805, 2020.
- [22] Stephen Keeley, David Zoltowski, Yiyi Yu, Spencer Smith, and Jonathan Pillow. Efficient non-conjugate gaussian process factor models for spike count data using polynomial approximations. In *International Conference on Machine Learning*, pages 5177–5186. PMLR, 2020.
- [23] Edoardo Balzani, Jean Paul Noel, Pedro Herrero-Vidal, Dora E Angelaki, and Cristina Savin. A probabilistic framework for task-aligned intra-and inter-area neural manifold estimation. *arXiv preprint arXiv:2209.02816*, 2022.
- [24] Francesco Paolo Casale, Adrian Dalca, Luca Saglietti, Jennifer Listgarten, and Nicolo Fusi. Gaussian process prior variational autoencoders. *Advances in neural information processing systems*, 31, 2018.
- [25] Vincent Fortuin, Dmitry Baranchuk, Gunnar Rätsch, and Stephan Mandt. Gp-vae: Deep probabilistic time series imputation. In *International Conference on Artificial Intelligence and Statistics*, pages 1651–1661. PMLR, 2020.
- [26] Siddharth Ramchandran, Gleb Tikhonov, Kalle Kujanpää, Miika Koskinen, and Harri Lähdesmäki. Longitudinal variational autoencoder. In *International Conference on Artificial Intelligence and Statistics*, pages 3898–3906. PMLR, 2021.
- [27] Changmin Yu, Hugo Soulat, Neil Burgess, and Maneesh Sahani. Amortised inference in structured generative models with explaining away. *arXiv preprint arXiv:2209.05212*, 2022.
- [28] Christopher J Paciorek. Bayesian smoothing with gaussian processes using fourier basis functions in the spectralgp package. *Journal of statistical software*, 19(2):nihpa22751, 2007.
- [29] Mikio Aoi and Jonathan W Pillow. Scalable bayesian inference for high-dimensional neural receptive fields. *bioRxiv*, page 212217, 2017.
- [30] B Yu, J Cunningham, G Santhanam, S Ryu, K Shenoy, and M Sahani. Gaussian-process factor analysis for low-d single-trial analysis of neural population activity. In *Frontiers in Systems Neuroscience. Conference Abstract: Computational and systems neuroscience*, 2009.

- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [32] James Hensman, Nicolas Durrande, and Arno Solin. Variational fourier features for gaussian processes. *The Journal of Machine Learning Research*, 18(1):5537–5588, 2017.
- [33] Andrew Wilson and Ryan Adams. Gaussian process kernels for pattern discovery and extrapolation. In *International conference on machine learning*, pages 1067–1075. PMLR, 2013.
- [34] Jackson Loper, David Blei, John P Cunningham, and Liam Paninski. A general linear-time inference method for gaussian processes on one dimension. *The Journal of Machine Learning Research*, 22(1):10580–10615, 2021.
- [35] Evan S Schaffer, Neeli Mishra, Matthew R Whiteway, Wenze Li, Michelle B Vancura, Jason Freedman, Kripa B Patel, Venkatakaushik Voleti, Liam Paninski, Elizabeth MC Hillman, et al. Flygenvectors: the spatial and temporal structure of neural activity across the fly brain. *bioRxiv*, pages 2021–09, 2021.
- [36] Venkatakaushik Voleti, Kripa B Patel, Wenze Li, Citlali Perez Campos, Srinidhi Bharadwaj, Hang Yu, Caitlin Ford, Malte J Casper, Richard Wenwei Yan, Wenxuan Liang, et al. Real-time volumetric microscopy of in vivo dynamics and large-scale samples with scape 2.0. *Nature methods*, 16(10):1054–1062, 2019.
- [37] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 2018.
- [38] Matthew R Whiteway, Evan S Schaffer, Anqi Wu, E Kelly Buchanan, Omer F Onder, Neeli Mishra, and Liam Paninski. Semi-supervised sequence modeling for improved behavioral segmentation. *bioRxiv*, pages 2021–06, 2021.
- [39] Elad Ganmor, Michael Krumin, Luigi F Rossi, Matteo Carandini, and Eero P Simoncelli. Direct estimation of firing rates from calcium imaging data. *arXiv preprint arXiv:1601.00364*, 2016.
- [40] Jordanna DH Sprayberry and Thomas L Daniel. Flower tracking in hawkmoths: behavior and energetics. *Journal of Experimental Biology*, 210(1):37–45, 2007.

- [41] U. B. Sikandar, H. Choi, J. Putney, H. Yang, S. Ferrari, and S. Sponberg. Predicting visually-modulated precisely-timed spikes across a coordinated and comprehensive motor program. In *2023 International joint conference on neural networks (IJCNN)*. IEEE, 2023.
- [42] Joy Putney, Rachel Conn, and Simon Sponberg. Precise timing is ubiquitous, consistent, and coordinated across a comprehensive, spike-resolved flight motor program. *Proceedings of the National Academy of Sciences*, 116(52):26951–26960, 2019.
- [43] Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial Intelligence and Statistics*, pages 914–922. PMLR, 2017.

Supplementary Material

Additional comparison with Targeted Neural Dynamical Modeling (TNDM) and Preferential Subspace Identification (PSID)

In order to compare MM-GPVAE with Targeted Neural Dynamical Modeling (TNDM) [11] and Preferential Subspace Identification (PSID) [12], we evaluated both TNDM and PSID with our rotating and scaling MNIST digit '3' alongside Poisson neural rates. In Figure 7, we show the reconstruction of both the behavioral modality, the MNIST digit '3', and Poisson neural rates for each model. In addition to these, we also provide the error on both data modalities on held-out trials. Both TNDM and PSID have linear decoders to observations, thus are not capable of accurately describing the rotating digit. On top of this, we see that the MMGPVAE shows better reconstruction of the underlying smooth neural rates.

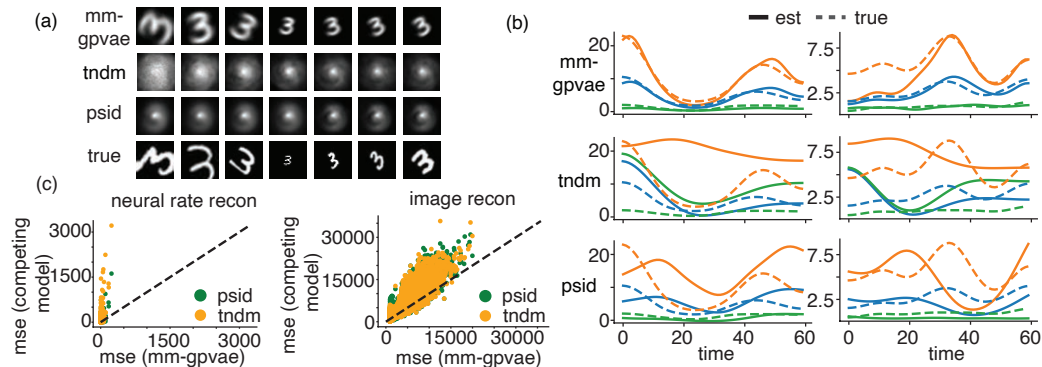


Figure 7: (a) Reconstruction of the scaling/rotating MNIST digit '3' with MM-GPVAE, TNDM and PSID. (b) Reconstruction of neural rates with MM-GPVAE, TNDM, and PSID. (c) MSE for neural rate reconstruction (left), and MSE for image reconstruction (right). Here, each dot indicates one trial mse from MMGPVAE vs a competing model. The majority of the trials errors fall above the unity line for both models, indicating overall better reconstruction with MMGPVAE.

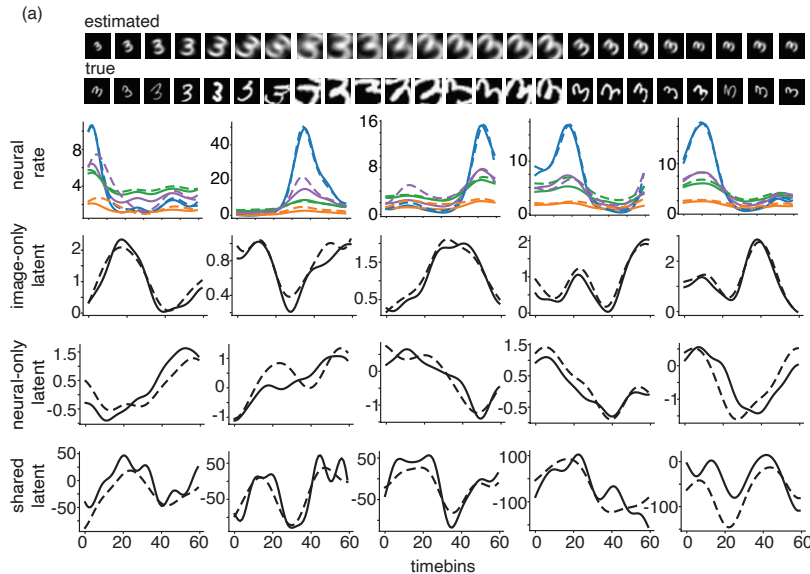


Figure 8: (a) More examples of estimated images and neural rates of the MNIST digit '3' as well as independent and shared latent estimation.

Additional evaluations of the MM-GPVAE on simulated data

To provide a more complete picture of the ability of the MM-GPVAE to both reconstruct simulated 'behavioral' and neural data, and to accurately recover the true underlying latent variables, we show additional performance evaluations here. Figure 8 shows an example of 24 reconstructed 3s from the evaluation shown in Figure 3 of the manuscript. Again, here we reconstruct a scaling and rotating MNIST digit '3' alongside 100 neural spike trains, where one latent dimension is shared across modalities. Figure 8 also shows latent reconstructions and 5 example neural rates in 5 held out trials. Using the same set-up, we additionally run the MM-GPVAE with MNIST digit '2', and show an example of 24 reconstructed images as well as 3 example neural rates (of 100) on 5 held-out trials in Figure 9.

In addition to these, we also extend figure 3(a) from our manuscript with all the latent trajectories for our simulated data comparisons, including estimates from GPFA and GPVAE. We additionally provide clarification for the 'shared' latent for our single modality comparisons (GP-VAE and

GPFA). Because in GPFA and GPVAE there is ambiguity as to which latent is “shared” (as these are unimodal models), attached Figure 10 panels (b) and (c) show the mean squared error for both of the possible latents for GPVAE and GPFA. Here, we note that the MMGPVAE is able to ‘disentangle’ the shared representation across the modalities, something that each unimodal model is unable to do. We visualize both the average mean-squared error (c) on held-out data as well as a per-trial mean-squared error scatter-plot (b) to provide a better sense of model performance.

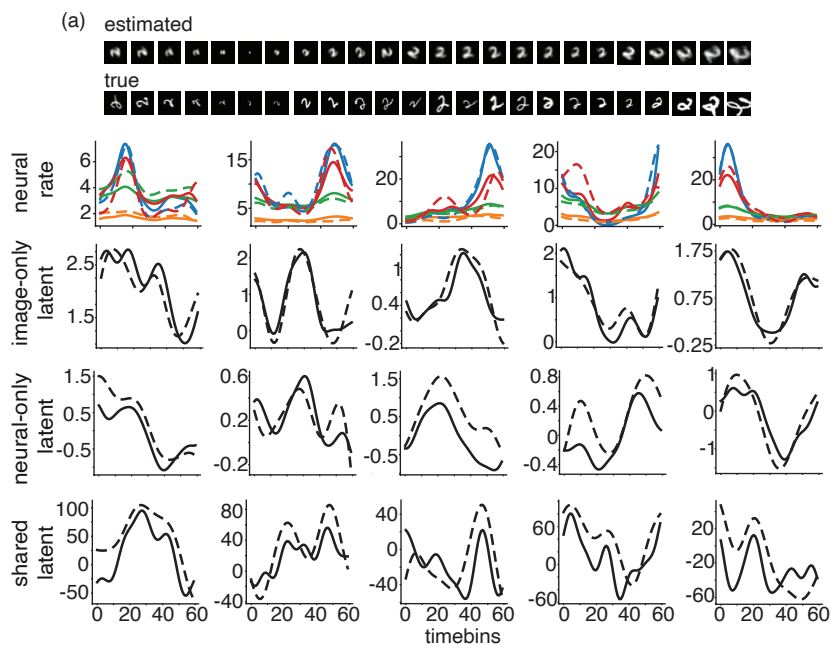


Figure 9: (a) An example of using the MNIST digit '2' in MM-GPVAE with reconstructed digits and neural rates as well as independent and shared latent estimation.

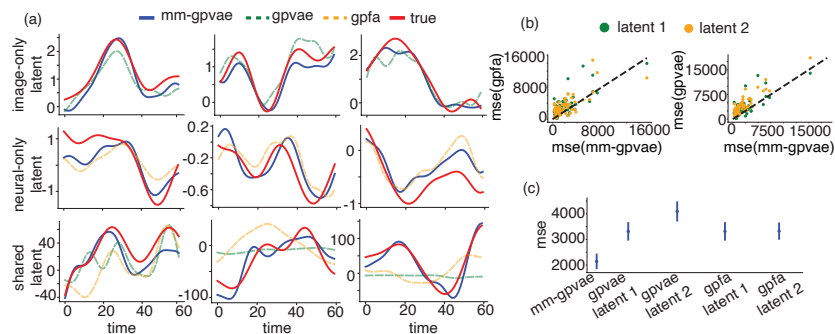


Figure 10: (a) Extension of Fig 3.(a) from the manuscript with all latent trajectories, showing the best matched latent from GPFA and GP-VAE to the true shared and independent latents. (b) MSE for both possible latents estimating the shared latent for GPFA (left) and GP-VAE (right). Each dot indicates a trial. (c) Average MSE for recovery of shared latent for the possible latents for all three models.

Neural Net Architectures

For the MM-GPVAE, both the simulated and real-world multi-modal evaluations used similar neural net architectures. However, there were some modifications of the nodes/layers that were unique to each evaluation. This was necessary as the latent dimensionality was different across our different evaluations, and certain behavioral reconstructions, especially those that were high-dimensional, required richer neural network parameterizations. The schematic for the MM-GPVAE neural network architecture in our simulated example can be found in Figure 11, and the schematics for the neural network architectures for the real-world multi-modal datasets can be found in Figure 12 (fly) and Figure 13 (hawkmoth). For our evaluation on the dataset used in [11], whose results are above, the architecture can be found in Figure 14. Note that for all experiments, each modality is encoded to its own set of variational means and variances (transformed into the Fourier domain). The encoded means and variances representing the shared latents are then summed to give the encoded shared latents means and variances. Across all evaluations, we parameterized our latents in the Fourier domain and converted back to the time domain before decoding.

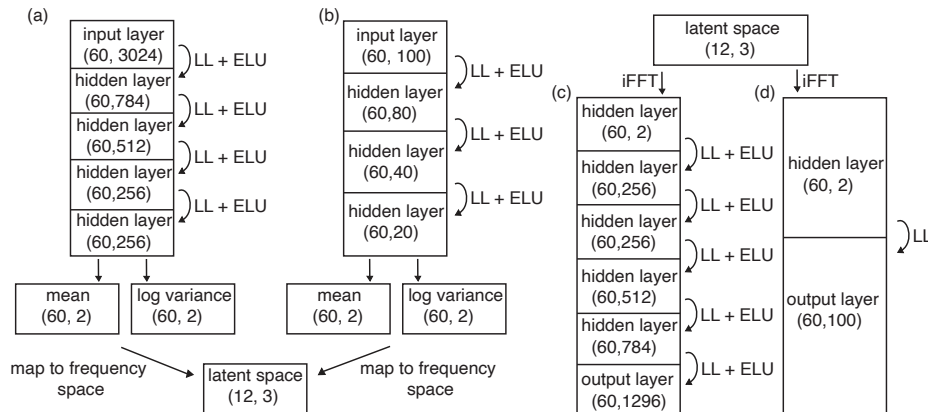


Figure 11: MM-GPVAE architecture for simulated data. (a) Encoder network of the MNIST digit. (b) Encoder network of the neural information. (c) Decoder network of the MNIST digit. (d) Decoder network of the neural information.

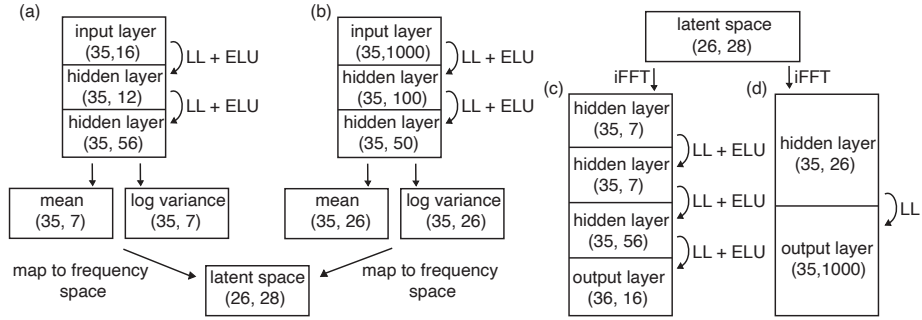


Figure 12: Neural network architecture for evaluation on fly dataset. (a) Encoder network for behavior. (b) Encoder network for neural data. (c) Decoder network for behavior. (d) Decoder network for neural data.

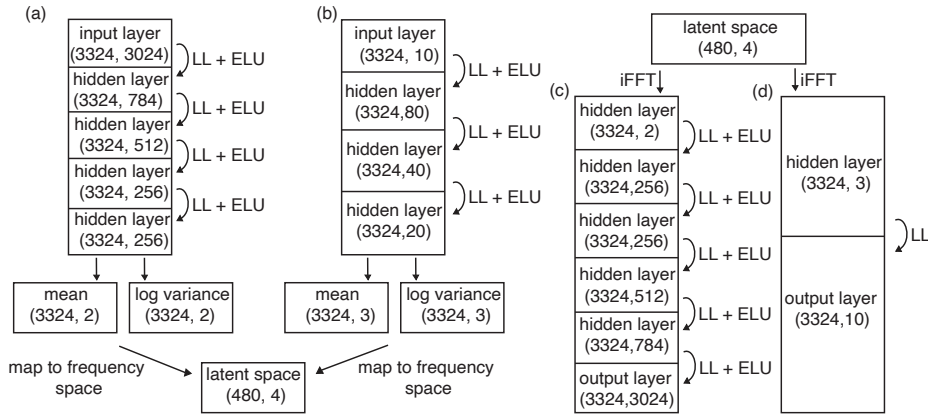


Figure 13: Neural network architecture for evaluation on hawkmoth dataset. (a) Encoder network for behavior. (b) Encoder network for neural data. (c) Decoder network for behavior. (d) Decoder network for neural data.

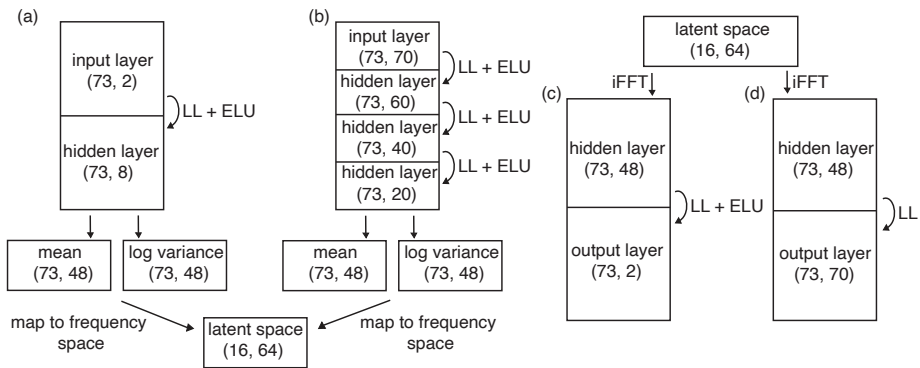


Figure 14: Neural network architecture for Monkey reaching data, evaluated above. (a) Encoder network for the behavior. (b) Encoder network for the neural activity. (c) Decoder network for the behavior. (d) Decoder network for the neural activity.

Additional information on fly experimental data

Pre-processing: We isolated 1000 raw calcium traces from [35] for use with our model for the evaluations in section 4.2 of the manuscript. This dataset also contained x,y positions from 8 tracked limbs positions [37]. The data came as one continuous recording, and at every time point there was a behavioral categorization probability for one of 6 distinct behaviors (undefined, still, running, front grooming, back grooming, abdomen bending) determined via the algorithm outlined in [38]. To split this continuous recording into trials, we found segments of the recording where, for 35 continuous samples, a single behavior was estimated at a $\geq 60\%$ probability. This generated 318 total trials where each trial was of one of 5 possible behaviors. There was no section of the recording where there was an 'undefined' behavior for 35 continuous samples, so there are no undefined behavioral trials in our analysis. To avoid the possibility of shared latent to contain dynamics that primarily drive one modality, we use a specific cross-validation procedure. Specifically, after first identifying the appropriate number of dimensions per modality, we increased the number of shared dimensions sequentially until predictive performance no longer improves, assuring we had no more shared dimensions than needed. Additionally, this is more likely to be the case when the dimensionality of the shared subspace is too large.

Additional evaluation: Though we show the time-course of the behavioral reconstruction in the main manuscript, we show a 2-d reconstruction on 2 example trials in Figure 15. Here we plot true x,y positions for 2 trails for each of the 8 limb positions alongside our model's reconstruction. We can see here that our model also captures the spatial information in the behavioral data. We also show for comparison the 2-d depiction of the shared and independent latent representation of all trials in the dataset (Figure 16), with all five of the behaviors labelled. You can see here that the 'still' behavior separates well from all other behaviors in the shared and neural subspaces.

Additional information on hawkmoth experimental data

Pre-processing: Original synthetic visual stimuli were sampled at 125 Hz and the neural and torque recordings were sampled at 10K Hz [40–42].

To prepare these data for evaluation with the MM-GPVAE, we first downsampled neural information and torque measurements to 1K Hz by binning the spike counts and averaging the torque measurements at this

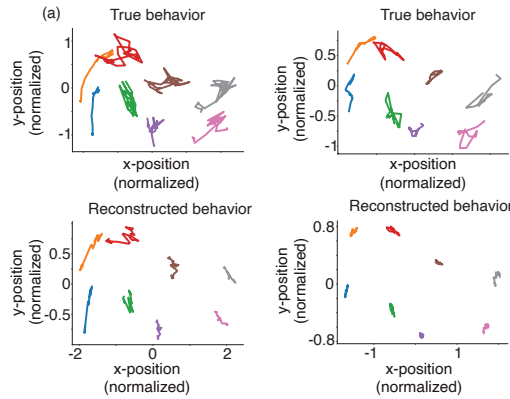


Figure 15: (a) Reconstruction of 8 fly limb positions in 2 held-out trials. Here, we see the MM-GPVAE is able to reconstruct the spatial information of the behavioral modality in the 8 tracked limb positions.




Figure 16: (a) Separation of all 5 fly behaviors in the behavior-only, neural-only, and shared subspaces. Similar to what is seen in the main manuscript, the 'still' behavior is well separated from trials with other behavioral labels in the neural and shared latent subspaces.

temporal resolution. To align these measurements with the visual stimuli, we upsampled the images 8-fold. The entire recording was 20 seconds long and was split into 6 evenly-divided trials.

The visual stimuli contained 3024 pixels and there were ten total hawkmoth motor neurons. We set the neural-independent subspace to 2-dimensional, the images-independent subspace to 1-dimensional, and an additional 1 dimension for the shared subspace. To encourage slow-evolving smooth latents in the shared and image subspaces, and faster-evolving neural latents, we initialized the length scale parameters for each latent dimension to different values.




Figure 17: (a) True spikes and estimated rates of all 10 hawkmoth neurons over 1500 time points.

The length scale was set to 10 for the neural latents, 150 for the shared latent, and 300 for the image latent. This biased the model fits to capture slower dynamics in the image subspace and the faster dynamics in the neural subspace.

Additional evaluation: The neural rates are well-captured by the MM-GPVAE for all 10 hawkmoth motor neurons, which each show strong periodicity. Figure 17 shows the strong fast-oscillation spike rates captured by the MM-GPVAE for all ten hawkmoth neurons alongside the recorded spikes for a 1.5 second period. Similar periodicity exists across the entire 20 second recording.

Additional model fitting information

All data had a 80-20 split for training and testing respectively. The Fourier frequency pruning was set to the minimum length scale of 10, 10, 3, and 16 for GP-VAE (simulated), MM-GPVAE (simulated), MM-GPVAE (fly), and MM-GPVAE (moth) respectively. GP length scales parameters were initialized to a value of 30 for all except for the hawkmoth evaluations (where initial values are indicated above), and jointly optimized with the ELBO. The covariance parameter α was set at a fixed value of $1e-2$, $1e-2$, $1e-3$, and $1e-4$ for GP-VAE (simulated), MM-GPVAE (simulated), MM-GPVAE (fly), and MM-GPVAE (moth) respectively. We additionally initialized the offsets \mathbf{d} of the neural modality to the average log-rate of the neural data.

Choice of Gaussian observation variance initialization, σ^2

We found that how we initially set the value of σ^2 could effect the performance of the MM-GPVAE, especially regarding the reconstruction of the behavioral data. Although we learned the σ^2 value jointly with the other parameters during optimization, the value of σ^2 tended to vary minimally from its initial value. Therefore, our setting of σ^2 tended to act as a reconstruction penalty, balancing the ability of the MM-GPVAE to prefer to reconstruct either the behavioral or neural data. Such a scaling term in the ELBO reconstruction has been seen in other models [11], and in the original GP-VAE this parameter was chosen carefully through a specific cross-validation approach [24]. Here, we simply set σ^2 to a value proportional to the dimensionality of the behavioral modality, which tended to balance the neural and behavioral terms in the ELBO, and generate good reconstructions for each modality. The initial values for σ^2 were 1000, 100, 1e-6, and 1 for GP-VAE (simulated), MM-GPVAE (simulated), MM-GPVAE (fly), and MM-GPVAE (moth) respectively.

Derivation of the evidence lower bound (ELBO)

Here we show the derivation of the evidence lower bound used for the MM-GPVAE. For clarity, we will start by deriving the ELBO for the MM-GPVAE just in the time domain. We set the GP prior parameters $\Theta = \{\alpha, \theta\}$, \mathbf{W} to be the loadings matrix and offsets from equation 4 in the manuscript, \mathbf{T} to be the observed timepoints, and \mathbf{Z} to be the collection of all the shared and independent latents $\mathbf{Z} = \{\mathbf{z}_a, \mathbf{z}_s, \mathbf{z}_b\}$.

$$\begin{aligned}
& \log p(\mathbf{Y}^A, \mathbf{Y}^B \mid \mathbf{T}, \phi, \mathbf{W}, \sigma^2, \Theta) = \\
& \log \int \frac{p(\mathbf{Y}^A, \mathbf{Y}^B \mid \mathbf{Z}, \phi, \mathbf{W}, \sigma^2) p(\mathbf{Z} \mid \mathbf{T}, \Theta)}{q_\psi(\mathbf{Z} \mid \mathbf{Y}^A, \mathbf{Y}^B)} q_\psi(\mathbf{Z} \mid \mathbf{Y}^A, \mathbf{Y}^B) d\mathbf{Z} \\
& \geq \int \log \left(\frac{p(\mathbf{Y}^A, \mathbf{Y}^B \mid \mathbf{Z}, \phi, \mathbf{W}, \sigma^2) p(\mathbf{Z} \mid \mathbf{T}, \Theta)}{q_\psi(\mathbf{Z} \mid \mathbf{Y}^A, \mathbf{Y}^B)} \right) q_\psi(\mathbf{Z} \mid \mathbf{Y}^A, \mathbf{Y}^B) d\mathbf{Z} \\
& = \mathbb{E}_{\mathbf{Z} \sim q_\psi} \left[\log p(\mathbf{Y}^A \mid \phi, \sigma^2, \mathbf{z}_s, \mathbf{z}_a) + \log p(\mathbf{Y}^B \mid \mathbf{W}, \mathbf{z}_s, \mathbf{z}_b) + \log p(\mathbf{Z} \mid \mathbf{T}, \Theta) \right] \\
& \quad - \int \log q_\psi(\mathbf{Z} \mid \mathbf{Y}) q_\psi(\mathbf{Z} \mid \mathbf{Y}) d\mathbf{Z} \\
& = \mathbb{E}_{\mathbf{Z} \sim q_\psi} \left[\sum_t \log \mathcal{N}(\mathbf{y}_A \mid g_\phi(\mathbf{x}_A), \sigma^2 \mathbf{I}_N) + \sum_t \log(\mathcal{P}(\mathbf{y}_B \mid f(\mathbf{x}_B))) \right. \\
& \quad \left. + \log p(\mathbf{Z} \mid \mathbf{T}, \Theta) \right] + H(q_\psi)
\end{aligned}$$

The Fourier domain representation of the ELBO only requires sampling over variational parameters in the Fourier space, but only changes the expression of the GP prior term $p(\mathbf{Z})$.

GP prior: The expectation of the GP prior term can be expressed in the Fourier domain as:

$$\begin{aligned}
\mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} \left[p(\tilde{\mathbf{Z}} \mid \Theta, \omega) \right] &= \sum_{p, \omega} \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} \left[\log \mathcal{N}(\tilde{z}_{p, \omega} \mid 0, [\tilde{\mathbf{K}}_p]_{\omega, \omega}) \right] \\
&= \frac{1}{2} \sum_{p, \omega} \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} \left[\log([\tilde{\mathbf{K}}_p]_{\omega, \omega} + \alpha) + \frac{\tilde{z}_{p, \omega}^2}{([\tilde{\mathbf{K}}_p]_{\omega, \omega} + \alpha)} \right] \\
&= \frac{1}{2} \sum_{p, \omega} \log([\tilde{\mathbf{K}}_p]_{\omega, \omega} + \alpha) + \frac{\tilde{\sigma}_{p, \omega}^2(\mathbf{Y}) + \tilde{\mu}_{p, \omega}^2(\mathbf{Y})}{([\tilde{\mathbf{K}}_p]_{\omega, \omega} + \alpha)},
\end{aligned}$$

where the double sum is due to the variational distribution q being a mean field Gaussian, and p here indexes latents and ω indexes Fourier frequencies.

Neural likelihood: Since the estimated log-rates of the neural data is a linear transform of the shared and neural latent variables, we can also evaluate the expectation of the neural-modality likelihood term in closed form. Recall that

$$\mathbf{x} = \mathbf{WZ} = \mathbf{W}\tilde{\mathbf{Z}}\mathbf{B}^\top,$$

where \mathbf{x} is the matrix of embeddings, and $\tilde{\mathbf{Z}}$ is the $P \times \mathcal{F}$ matrix of Fourier-domain latent variables. We may therefore write the embedding for measurement i at time t as

$$x_{i,t} = \mathbf{w}_i^\top \tilde{\mathbf{Z}} \mathbf{b}_t,$$

where \mathbf{w}_i is the i^{th} row of \mathbf{W} and \mathbf{b}_t^\top is the t^{th} column of \mathbf{B}^\top . If we let $\tilde{\mu}_{p,\omega} = \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi}[\tilde{z}_{p,\omega}]$, and correspondingly let $\tilde{\boldsymbol{\mu}}$ be the $P \times \mathcal{F}$ matrix of $\tilde{\mu}_{p,\omega}$'s then $E_{\tilde{\mathbf{Z}} \sim q_\psi}[x_{i,t}] = \mathbf{w}_i^\top \tilde{\boldsymbol{\mu}} \mathbf{b}_t$. We may equivalently write $x_{i,t}$ as

$$\begin{aligned} x_{i,t} &= \text{vec}(x_{i,t}) = \text{vec}(\mathbf{w}_i^\top \tilde{\mathbf{Z}} \mathbf{b}_t) \\ &= \text{vec}(\mathbf{b}_t^\top \tilde{\mathbf{Z}}^\top \mathbf{w}_i) \\ &= (\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \tilde{\mathbf{z}}, \end{aligned}$$

where $\tilde{\mathbf{z}} = \text{vec}(\tilde{\mathbf{Z}}^\top)$. We may therefore we may derive the variance of $x_{i,t}$ as

$$\begin{aligned} \text{Var}_{\tilde{\mathbf{Z}} \sim q_\psi}[x_{i,t}] &= \text{Var}_{\tilde{\mathbf{Z}} \sim q_\psi}[(\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \tilde{\mathbf{z}}] \\ &= \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi}[\text{Trace}[(\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \tilde{\mathbf{z}} \tilde{\mathbf{z}}^\top]] - (\mathbf{w}_i^\top \tilde{\boldsymbol{\mu}} \mathbf{b}_t)^2 \\ &= \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi}[\text{Trace}[(\mathbf{w}_i \otimes \mathbf{b}_t)(\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \tilde{\mathbf{z}} \tilde{\mathbf{z}}^\top]] - (\mathbf{w}_i^\top \tilde{\boldsymbol{\mu}} \mathbf{b}_t)^2 \\ &= \text{Trace}[(\mathbf{w}_i \otimes \mathbf{b}_t)(\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) (\mathbf{V} + \text{vec}(\tilde{\boldsymbol{\mu}}) \text{vec}(\tilde{\boldsymbol{\mu}})^\top)] - (\mathbf{w}_i^\top \tilde{\boldsymbol{\mu}} \mathbf{b}_t)^2 \\ &= (\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \mathbf{V} (\mathbf{w}_i \otimes \mathbf{b}_t), \end{aligned}$$

where \mathbf{V} is the diagonal posterior covariance of $\tilde{\mathbf{z}}$ whose elements are the encoded Fourier variational variances, $\tilde{\sigma}_\omega^2(\mathbf{Y})$. Therefore, we observe that under the variational posterior $x_{i,t} | \sim \mathcal{N}(m_{i,t}, v_{i,t})$, where $m_{i,t} \equiv \mathbf{w}_i^\top \tilde{\boldsymbol{\mu}} \mathbf{b}_t$ and $v_{i,t} \equiv (\mathbf{w}_i^\top \otimes \mathbf{b}_t^\top) \mathbf{V} (\mathbf{w}_i \otimes \mathbf{b}_t)$.

We note that for $\lambda_{i,t} = e^{x_{i,t}}$ follows a long-normal distribution, meaning that, for $x_{i,t} \sim \mathcal{N}(m_{i,t}, v_{i,t})$ then $\mathbb{E}[\lambda_{i,t}] = e^{m_{i,t} + \frac{1}{2}v_{i,t}}$. This allows us to specify the posterior expectation of the Poisson likelihood in closed form. Namely,

$$\begin{aligned} \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} \left[\log \mathcal{P}(y_{i,t} | f(x_{i,t})) \right] &= \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} \left[y_{i,t} \log \lambda_{i,t} + \lambda_{i,t} \right] - \log y_{i,t}! \\ &= y_{i,t} \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} [x_{i,t}] + \mathbb{E}_{\tilde{\mathbf{Z}} \sim q_\psi} [\lambda_{i,t}] + \text{const}_{\tilde{\mathbf{z}}} \\ &= y_{i,t} m_{i,t} + e^{m_{i,t} + \frac{1}{2}v_{i,t}} + \text{const}_{\tilde{\mathbf{z}}} \end{aligned}$$