Multi-Region Markovian Gaussian Process: An Efficient Method to Discover Directional Communications Across Multiple Brain Regions

Weihan Li [1] **Chengrui Li** [1] **Yule Wang** [1] **Anqi Wu** [1]

Abstract

Studying the complex interactions between different brain regions is crucial in neuroscience. Various statistical methods have explored the latent communication across multiple brain regions. Two main categories are the Gaussian Process (GP) and Linear Dynamical System (LDS), each with unique strengths. The GP-based approach effectively discovers latent variables with frequency bands and communication directions. Conversely, the LDS-based approach is computationally efficient but lacks powerful expressiveness in latent representation. In this study, we merge both methodologies by creating an LDS mirroring a multi-output GP, termed Multi-Region Markovian Gaussian Process (MRM-GP). Our work establishes a connection between an LDS and a multi-output GP that explicitly models frequencies and phase delays within the latent space of neural recordings. Consequently, the model achieves a linear inference cost over time points and provides an interpretable low-dimensional representation, revealing communication directions across brain regions and separating oscillatory communications into different frequency bands.

1. Introduction

The number of simultaneous neural recordings from various brain regions has increased recently. These recordings offer opportunities to explore the mechanisms through which inter-areal communication supports brain function (Kohn et al., 2020). Brain regions linked to sensory and cognitive functions often display interconnectedness, with signals transmitted bidirectionally and potentially simultaneously (Harris & Mrsic-Flogel, 2013; Miller et al., 2018; Wang et al., 2024). However, the high-dimensional neural recordings typically present a complex view of this concurrent communication—for example, neurons may concurrently represent overlapping neural activities within a certain region. Therefore, uncovering the interactions between different brain regions presents a challenging task.

Many statistical methodologies have been employed to address the challenge of understanding communications across multiple brain regions. Hultman et al. 2018 examined multi-region local field potential data and identified frequency-based interactions across brain regions using a Gaussian Process Factor Analysis model. Following a similar approach, Gokcen et al. 2022; 2023 suggested that latent variables can be divided into across- and within-region components. This model was applied to disentangle the concurrent and bidirectional communications across brain regions with a multi-output Squared Exponential kernel. Glaser et al. 2020 developed a switching linear dynamic system to uncover low-dimensional interactions among multiple brain regions. This method captured regions responsible for transitioning between latent states by specifying a novel transition rule.

Broadly categorized into Gaussian Process (GP) and Linear Dynamical System (LDS) classes, these methods offer distinct advantages. The GP-based approach, leveraging the robust representational capability of multi-output kernels, performs well in discovering latent variables with crucial information, such as frequencies and directional communications. Conversely, the LDS-based approach, while computationally efficient with a linear cost in time points, lacks the powerful expressiveness of GP in latent representation.

Our goal is to combine the strengths of both methodologies by constructing an LDS that mirrors a GP. Several studies have explored this connection: Hartikainen & Särkkä 2010 established a framework about converting a single-output GP with Matern or Squared Exponential kernel to an LDS, which relied on spectral factorization (Sayed & Kailath, 2001). Building upon this, Solin & Särkkä 2014 proposed the conversion for single-output periodic kernels, and Särkkä et al. 2013 extended the single-output conversions to a spatiotemporal GP. However, applying these conversions for GP-based multi-region methods is non-trivial because a gap exists in converting a multi-output GP to an LDS. One approach to bridge this gap is to assume the kernel

[1]School of Computational Science & Engineering, Georgia Institute of Technology, Atlanta, USA. Correspondence to: Anqi Wu <anqiwu@gatech.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

is separable over spatial and temporal domain (Solin et al., 2016). This allows us to create a closed-form multi-output GP-LDS conversion following the framework proposed in single-output cases.

Consequently, choosing a separable kernel becomes essential in this study. An effective option is the complex-valued multi-region kernel (Ulrich et al., 2015), specifically designed to facilitate learning latent interactions encompassing frequencies and phase delays across brain regions. However, it is important to note that the connection between an LDS and a GP with a complex-valued multi-output kernel remains unknown.

We introduce the Multi-Region Markovian Gaussian Process (MRM-GP) to model latent representations, where Markovian means the discrete state space representation of a GP. Our work establishes a connection between an LDS and a multi-output GP that explicitly discovers frequency-based latent communications and their directionality via phase delays. By doing so, we can have three advantages: (1) utilizing the powerful representational capability of kernel functions; (2) employing the efficient inference algorithm to ensure a linear computational cost over time points; (3) extending the LDS to incorporate time-varying frequencies and delays by switching states.

We test MRM-GP using multi-region spike trains and local field potential recordings. The model proves its capability to produce understandable low-dimensional representations. These representations illustrate the direction of communication flow among regions and effectively disentangle oscillatory interactions into diverse frequencies.

2. Background

We introduce the multi-region kernel, a multi-output kernel for modeling interactions across different brain regions. Then, we demonstrate how a Gaussian Process with this kernel can be employed to model latent communications across regions. It is worth noting that various mapping methods can be used to project latent representations onto neural recordings, and in this case, we opt for Factor Analysis.

2.1. Multi-Region Kernel

The complex-valued multi-region kernel proposed in (Ulrich et al., 2015) explicitly models communication frequencies and phase delays within the latent space of neural data:

$$\mathbf{K}_{pp'}(\tau) = \sum_{r=1}^{R} a_p^r a_{p'}^r \exp\left(-\frac{1}{2\sigma^2}\tau^2 + i\eta(\tau + \phi_{pp'})\right), \quad (1)$$

$$= \underbrace{\exp\left(-\frac{1}{2\sigma^2}\tau^2 + i\eta\tau\right)}_{\text{temporal}} \underbrace{\sum_{r=1}^{R} a_p^r a_{p'}^r \exp(i\eta\phi_{pp'})}_{\text{spatial}}.$$

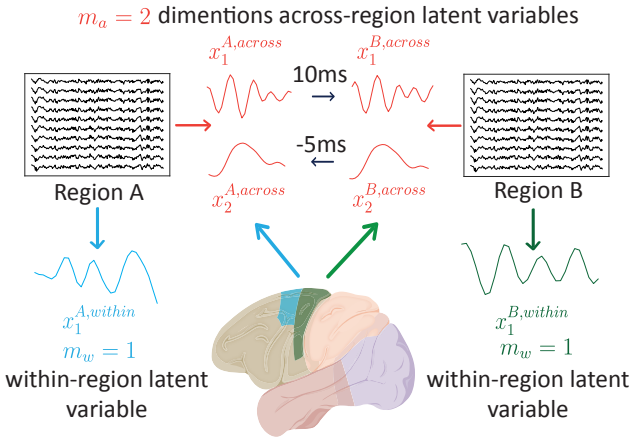

Figure 1. An example of two dimensions across-region latent variables and one dimension within-region latent variable. Brain region A and region B have bidirectional communications within different frequency bands. Each region also has a one-dimensional neural activity unrelated to the other region.

This kernel ensures separability over space and time, where p and p' are two brain regions, $\tau = t - t'$ is the time interval. In the temporal part, σ signifies the length scale, $i = \sqrt{-1}$ denotes the imagery unit, η represents the communication frequency between regions. In the spatial part, $\phi_{pp'}$ represents the phase delay between region p and p', a_p^r and $a_{p'}^r$ are amplitudes, and $R > 1$ denotes the rank number ensuring positive definiteness.

The separability is required to establish a connection between the multi-region kernel and a linear dynamic system (LDS). Moreover, the real part of this kernel $\mathrm{Re}[\mathbf{K}_{pp'}(\tau)] = \sum_{r=1}^{R} a_p^r a_{p'}^r \exp(-\frac{1}{2\sigma^2}\tau^2)\cos(\eta(\tau+\phi_{pp'}))$, denoted as the Cross-Spectral Mixture (CSM) kernel (Ulrich et al., 2015), has been shown to effectively capture frequency-based communications among various brain regions (Hultman et al., 2018). However, due to CSM's non-separability, we work with the complex-valued kernel as represented in Eq. 1 to build an LDS.

2.2. Multi-Region Gaussian Process Factor Analysis

A Gaussian Process Factor Analysis model, utilizing the real part of the multi-region kernel in Eq. 1 and named CSM-GPFA, can identify latent variables that capture frequencies and phase delays across brain regions.

Given single region neural recording $y^p \in \mathbb{R}^{n^p \times T}$, $p \in \{1, \ldots, P\}$ is the brain region index, n^p denotes the number of neurons in region p, and T represents time steps. Our goal is to find the M independent low-dimensional variables $x^p \in \mathbb{R}^{M \times T}$ for each region's neural data y^p. These variables from P regions together form as $x =$

$[x^1, \dots, x^P]^\top \in \mathbb{R}^{MP \times T}$, representing a latent representation for multi-region recordings $y = [y^1, \dots, y^P]^\top \in \mathbb{R}^{N \times T}$, where $N = n^1 + \dots + n^P$ is the total number of neurons over P regions. Besides, y is a linear mapping of x: $y = \mathbf{C}x + d + \epsilon$, where $\mathbf{C}$ is a block diagonal matrix $\mathbf{C} = \text{diag}\{\mathbf{C}^1, \dots, \mathbf{C}^p, \dots, \mathbf{C}^P\} \in \mathbb{R}^{N \times MP}$, $d \in \mathbb{R}^{N \times 1}$ is bias, and $\epsilon \sim \mathcal{N}(0, \mathbf{V})$ is Gaussian noise with $\mathbf{V} \in \mathbb{R}^{N \times N}$.

Meanwhile, a widely used assumption of x^p is to split it into across- and within-region parts (Gokcen et al., 2022): $x^p = [x^{p,a}, x^{p,w}]^\top, x^{p,a} \in \mathbb{R}^{m_a \times T}, x^{p,w} \in \mathbb{R}^{m_w \times T}$, where m_a, m_w are the number of dimensions for across- or within-region part and $m_a + m_w = M$. The across-region variables $x^{p,a}$ describe neural activity that is shared across all brain regions, meaning that for the remaining $P-1$ regions, they have the latent variables with the same frequencies and dynamics except phase delays, while the within-region variables $x^{p,w}$ describe the neural activity of region p that is not related to other regions (see Figure 1).

Consequently, we model $x^{p,a}$ and $x^{p,w}$ separately with $\mathbf{K}$ in Eq. 1. For region p, there are m_a dimensions of across-region variables, and each dimension $x^{p,a}_m \in \mathbb{R}^{T \times 1}, m \in [1, m_a]$ has spatial correlations with the remaining $P-1$ regions. So, the mth dimension across-region variables over P regions: $x^a_m = [x^{1,a}_m, \dots, x^{P,a}_m] \in \mathbb{R}^{P \times T}$ are considered as a group and modeled as the real part of a multi-output Complex Gaussian Process with $\mathbf{K}^m$. For within-region variables, each dimension $x^{p,w}_m \in \mathbb{R}^{T \times 1}, m \in [1, m_w]$ is independently modeled as the real part of a single-output Complex Gaussian Process with $\mathbf{K}(\tau)^m = \sum_{r=1}^{R} a_r^{m2} \exp(-\frac{1}{2\sigma^{m2}}\tau^2 + i\eta^m\tau)$, where $p = p'$ and $\phi_{pp'} = 0$. Furthermore, we also assume independence among different dimensions of across- and within-region variables, implying that different index m refers to distinct kernel parameters.

Unlike the approach in (Ulrich et al., 2015), which uses a mixture of frequencies, we employ a single frequency in Eq 1 to achieve frequency disentanglement. Specifically, each dimension of the across-region latent variable will have a single frequency peak. Consequently, the mixture of frequencies present in the data will be captured by multiple dimensions.

3. Method

Modeling latent variables x^p with Gaussian Process is inefficient with a $\mathcal{O}(T^3)$ time complexity. So, we want to build Markovian representations of these latent variables, indicating the state space representations of each dimension: across-region x^a_m and within-region $x^{p,w}_m$, where every Markovian representation follows a Linear Dynamical System (LDS) (Solin et al., 2016).

Spectral factorization has been used in multi-output GP cases (Zhu et al., 2023). However, for the complex-valued multi-output kernel, we need to develop a new spectral factorization-based method to first convert the complex-valued temporal part to an LDS (Section 3.1) and then use the kernel's separability to combine the complex-valued spatial part to get the final LDS (Section 3.2).

3.1. Markovian Within-Region Latent Variables

The Markovian representation of region p's m^{th} dimension within-region latent variable $x^{p,w}_m \in \mathbb{R}^{T \times 1}$ follows a discrete-time LDS structure:

$$f^{p,w}_{m,t} = \mathbf{A}^w_m f^{p,w}_{m,t-1} + q_{t-1}, \quad q_{t-1} \sim \mathcal{CN}(0, \mathbf{Q}^w_m), \quad (2)$$

where $f^{p,w}_{m,t} = [g^{p,w}_{m,t}, \frac{dg^{p,w}_{m,t}}{dt}, \dots, \frac{d^{k-1}g^{p,w}_{m,t}}{dt^{k-1}}]^T \in \mathbb{C}^{k \times T}$, denoting the complex-valued dynamics $g^{p,w}_{m,t}$ and its derivatives up to $(k-1)^{th}$ order at time t. Especially, within-region latent variable $x^{p,w}_m$ is the real part of $g^{p,w}_{m,t}$. $\mathbf{A}^w_m \in \mathbb{C}^{k \times k}$ represents the complex-valued transition matrix, and q_{t-1} is the sampling from a complex normal distribution $\mathcal{CN}(\cdot)$ with the complex-valued measurement (Hermitian) matrix $\mathbf{Q}^w_m \in \mathbb{C}^{k \times k}$.

The key question now becomes how to associate single-output $\mathbf{K}(\tau)^m = \sum_{r=1}^{R} a_r^{m2} \exp(-\frac{1}{2\sigma^{m2}}\tau^2 + i\eta^m\tau)$ with $\mathbf{A}^w_m$ and $\mathbf{Q}^w_m$. To achieve this, our approach involves two steps: forming a continuous-time LDS for within-region variables in each region through spectral factorization (Kailath et al., 2000), and subsequently transforming it into the discrete-time version as specified in Eq. 2. This linkage is new and differs from previous connections (Hartikainen & Särkkä, 2010; Solin & Särkkä, 2014) as the kernel is situated in the complex domain.

Forming a continuous-time LDS. Given single-output $\mathbf{K}^m$, the continuous-time LDS we want to form is:

$$\frac{df(t)^{p,w}_m}{dt} = \mathbf{F}^w_m f(t)^{p,w}_m + \mathbf{L}u(t), \quad (3)$$

where $f(t)^{p,w}_m = [g(t)^{p,w}_m, \frac{dg(t)^{p,w}_m}{dt}, \dots, \frac{d^{k-1}g(t)^{p,w}_m}{dt^{k-1}}]^\top \in \mathbb{C}^{k \times T}$, denoting the continuous-time version of $g^a_{m,t}$ and its derivatives up to $(k-1)^{th}$ order. $\mathbf{F}^w_m \in \mathbb{C}^{k \times k}$ is a continuous-time transition matrix, $\mathbf{L} = [0, \dots, 0, 1]^\top \in \mathbb{R}^{k \times 1}$ signifies a constant vector, and $u(t)$ denotes a single-dimensional white noise with spectral density v. We need to obtain both $\mathbf{F}^w_m$ and v from $\mathbf{K}^m$.

$\mathbf{F}^w_m$ takes a companion form of LDS (Grewal & Andrews, 2014):

$$\mathbf{F}^w_m = \begin{bmatrix} 0 & 1 & & \\ & 0 & 1 & \\ & & \ddots & 1 \\ -a_0 & \dots & -a_{k-2} & -a_{k-1} \end{bmatrix}, \quad (4)$$

where $a_0, \ldots, a_{k-1}$ are the coefficients in a stochastic differential equation that is equivalent to Eq. 3:

$$\frac{d^k g(t)_m^{p,w}}{dt^k} + a_{k-1}\frac{d^{k-1} g(t)_m^{p,w}}{dt^{k-1}} + \cdots + a_0 g(t)_m^{p,w} = u(t). \tag{5}$$

To obtain $\mathbf{F}_m^w$ and v, we first apply Fourier transform on both sides of the continuous-time LDS in Eq. 3 to achieve a frequency domain representation (see Appendix A for derivation):

$$S(\omega) = \mathbf{G}(\mathbf{F}_m^w - i\omega\mathbf{I})^{-1}\mathbf{L}v\mathbf{L}^\top(\mathbf{F}_m^w + i\omega\mathbf{I})^{-T}\mathbf{G}^\top \tag{6}$$

where $S(\omega) = \sqrt{2\pi}\sigma^m \exp(-\frac{(\eta^m-\omega)^2}{2\sigma^m})$ is the spectral density of single-output $\mathbf{K}^m$, $\mathbf{G} = [1, 0, \ldots, 0] \in \mathbb{R}^{1\times k}$ represents a constant vector, $\mathbf{I} \in \mathbb{R}^{k\times k}$ denotes an identity matrix, and, notably, $v = \sqrt{2\pi}\sigma^m$. Now, we only need to solve Eq. 6 to obtain the coefficients $a_0, \ldots, a_{k-1}$ in $\mathbf{F}_m^w$.

On the left-hand side, $S(\omega)$ follows an exponential family, which is infinitely differentiable. On the right-hand side, however, the finite number coefficients $a_0, \ldots, a_{k-1}$ in $\mathbf{F}_m^w$ determine a finite polynomial function: $\mathbf{G}(\mathbf{F}_m^w - i\omega\mathbf{I})^{-1}\mathbf{L}$. Therefore, we can only construct a finite polynomial approximation of $S(\omega)$. But one observation from Eq. 6 is that $S(\omega)$ can be factorized into two parts, i.e., a complex function multiplying its conjugate.

Previous established connections between real-valued kernel and LDS assume a symmetric $S(\omega)$, implying using a Taylor expansion to approximate $S(\omega)$ as a polynomial of ω (Hartikainen & Särkkä, 2010). However, in the case of complex-valued $\mathbf{K}^m$, $S(\omega)$ is non-symmetric due to frequency η^m, so our solution is to approximate it as a polynomial of $i\omega$:

$$\begin{aligned}\frac{1}{S(\omega)} &\approx \sqrt{\frac{\sigma^m}{2\pi}}(b_0 + b_1 i\omega + b_2(i\omega)^2 + \cdots + b_{2k}(i\omega)^{2k}),\\ &= T(i\omega),\end{aligned} \tag{7}$$

where $b_{2k} = 1$ if k is even, $b_{2k} = -1$ if k is odd, $b_0, b_2, \ldots, b_{2k-2}$ are real numbers, and $b_1, b_3, \ldots, b_{2k-1}$ are complex numbers that only have imagery parts. These coefficients' values depend on σ^m, η^m, and k. Figure 2(A-C) shows the approximation of $S(\omega)$ when $k = 2, 3, 4$, demonstrating a reliable approximation even in the case of $k = 2$. Appendix D shows the effect of k on generated samples.

Now our target becomes solving the following equation for $a_0, \ldots, a_{k-1}$:

$$\begin{aligned}T(i\omega) &= \sqrt{\frac{\sigma^m}{2\pi}}H(i\omega)H(-i\omega),\\ H(i\omega) &= a_0 + a_1 i\omega + \cdots + a_{k-1}(i\omega)^{k-1} + (i\omega)^k,\end{aligned} \tag{8}$$

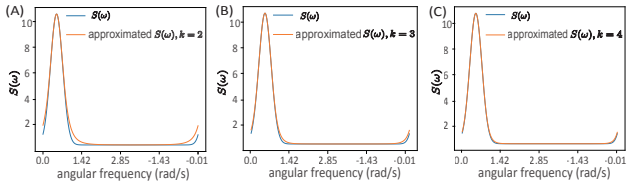

Figure 2. Approximate $S(\omega)$ with $\eta = 0.5$ rad/s, $\sigma = 4.5$ when varying the number of derivatives $k = 2, 3, 4$. $k = 2$ could provide a satisfactory approximation.

where $H(i\omega)$ is commonly referred to as the transfer function with $a_0, \ldots, a_{k-1}$ acting as its coefficients. Its reciprocal $\frac{1}{H(i\omega)}$ is the function form of $\mathbf{G}(\mathbf{F}_m^w - i\omega\mathbf{I})^{-1}\mathbf{L}$. Solving Eq. 8 is often referred to as spectral factorization, and an advantageous aspect of this factorization is that $a_0, \ldots, a_{k-1}$ are a subset of complex-valued roots of $T(i\omega)$, where they are all situated within the left-half complex plane (Kailath et al., 2000) and can be found by QR algorithm with time complexity $\mathcal{O}(k)$. See Appendix A for derivation.

Forming a discrete-time LDS. Given $\mathbf{F}_m^w$, $\mathbf{L}$ and $v = \sqrt{\frac{2\pi}{\sigma^m}}$ in Eq. 6, the computation of $\mathbf{A}_m^w$ and $\mathbf{Q}_m^w$ in Eq. 2 are as follows (Solin et al., 2016):

$$\begin{aligned}\frac{d\mathbf{P}_\infty}{dt} &= \mathbf{F}_m^w\mathbf{P}_\infty + \mathbf{P}_\infty\mathbf{F}_m^{w\,H} + \mathbf{L}v\mathbf{L}^\top = 0,\\ \mathbf{A}_m^w &= \text{expm}(\mathbf{F}_m^w\Delta t),\\ \mathbf{Q}_m^w &= \mathbf{P}_\infty - \mathbf{A}_m^w\mathbf{P}_\infty\mathbf{A}_m^{w\,H},\end{aligned} \tag{9}$$

where v is the spectral density of white noise $u(t)$, expm$(\cdot)$ represents the matrix exponential function, and Δt signifies the time interval in discrete-time LDS.

3.2. Markovian Across-Region Latent Variables

Region p's m^{th} dimension across-region latent variable $x_m^a \in \mathbb{R}^{P\times T}$ can be expressed as $x_m^a = [x_m^{1,a}, \ldots, x_m^{p,a}, \ldots, x_m^{P,a}]$. This indicates that they consist of P variables sharing the same η^m and σ^m for describing temporal features while using phase delays $\{\phi_{pp'}^m\}_{p,p'=1}^P$ to capture cross-spatial differences. It's important to note that the temporal and spatial components are separable in $\mathbf{K}^m$. Consequently, we can initially create a within-region Markovian representation, denoted as $\mathbf{A}_m^w, \mathbf{Q}_m^w$, for the temporal features in each $x_m^{p,a}$. Then, this representation is extended to the across-region Markovian representation for x_m^a through the incorporation of phase delays $\{\phi_{pp'}^m\}_{p,p'=1}^P$.

Therefore, the Markovian representation of mth dimension across-region latent variable x_m^a is:

$$\begin{aligned}f_{m,t}^a = \mathbf{A}_m^a f_{m,t-1}^a + q_{t-1}, &\quad q_{t-1} \sim \mathcal{CN}(0, \mathbf{Q}_m^a),\\ \mathbf{A}_m^a = \mathbf{I} \otimes \mathbf{A}_m^w, &\quad \mathbf{Q}_m^a = \mathbf{K}_{\text{spatial}}^m \otimes \mathbf{Q}_m^w,\end{aligned} \tag{10}$$

where $f^a_{m,t} = [g^a_{m,t}, \frac{dg^a_{m,t}}{dt}, \ldots, \frac{d^{k-1}g^a_{m,t}}{dt^{k-1}}]^\top \in \mathbb{C}^{Pk\times T}$, $x^a_{m,t}$ is the real part of $g^a_{m,t}$, $\mathbf{A}^a_m \in \mathbb{C}^{Pk\times Pk}$ is transition matrix, denoting the Kronecker product of identity matrix $\mathbf{I} \in \mathbb{R}^{P\times P}$ and $\mathbf{A}^w_m \in \mathbb{C}^{k\times k}$, and $\mathbf{Q}^a_m \in \mathbb{C}^{Pk\times Pk}$ is measurement matrix, denoting the Kronecker product of $\mathbf{K}^m$'s spatial part $\mathbf{K}^m_{\text{spatial}} = \sum_{r=1}^R a^{m,r}_p a^{m,r}_{p'} \exp(i\eta^m \phi^m_{pp'})$ and $\mathbf{Q}^w_m \in \mathbb{C}^{k\times k}$.

3.3. Multi-Region Markovian Gaussian Process

Given our assumption of independence among different dimensions of across-region variables and distinct dimensions of within-region variables, the Markovian representation for all variables $x \in \mathbb{R}^{MP\times T}$, both across- and within-region, spanning P brain regions can be expressed as:

$$f_t = \mathbf{A} f_{t-1} + q_{t-1}, \quad q_{t-1} \sim \mathcal{CN}(0, \mathbf{Q}), \tag{11}$$

where $f_t = [g_t, \frac{dg_t}{dt}, \ldots, \frac{d^{k-1}g_t}{dt^{k-1}}]^\top \in \mathbb{C}^{MPk\times T}$, x_t is the real part of g_t, and $\mathbf{A} \in \mathbb{C}^{MPk\times MPk}$, $\mathbf{Q} \in \mathbb{C}^{MPk\times MPk}$ are block diagonal matrices: $\mathbf{A} = \text{diag}\{\mathbf{A}^a_1, \ldots, \mathbf{A}^a_{m_a}, \mathbf{A}^w_1 \ldots \mathbf{A}^w_{m_w}\}$, $\mathbf{Q} = \text{diag}\{\mathbf{Q}^a_1, \ldots, \mathbf{Q}^a_{m_a}, \mathbf{Q}^w_1, \ldots, \mathbf{Q}^w_{m_w}\}$. Meanwhile, the neural recordings $y \in \mathbb{R}^{N\times T}$ can be reconstructed by $y = \mathbf{C}\,\text{Re}[\mathbf{G}f] + d + \epsilon$, with $\mathbf{C}$, d, ϵ from CSM-GPFA in Section 2.2, and $\mathbf{G}$ in Eq. 6.

3.4. Multi-Region Markovian Gaussian Process with Switching States

After the link between the multi-region Gaussian Process and linear dynamical system (LDS) is established, we can seamlessly extend the across-region discrete-time LDS in Eq. 10 to incorporate switching states.

Integrating a Hidden Markov Model (HMM) into LDS leads to Switching LDS (Fox et al., 2008), and similarly, combining HMM with MRM-GP results in Switching MRM-GP. A significant advantage of this integration is the ability to link the across-region's transition and measurement matrices with distinct, discrete states $z \in \{1, \ldots, Z\}$: $\mathbf{A}^a_z = \text{diag}\{\mathbf{A}^a_{1,z}, \ldots, \mathbf{A}^a_{m_a,z}\}, \mathbf{Q}^a_z = \text{diag}\{\mathbf{Q}^a_{1,z}, \ldots, \mathbf{Q}^a_{m_a,z}\}$, which makes it easy to accommodate time-varying frequencies and delays in across-region latent variables.

4. Inference

We have now established a connection between a Gaussian Process with a multi-region kernel and a linear dynamical system (LDS). The next step is to learn discrete states, model parameters, and latent variables.

MRM-GP, as a discrete-time LDS, affords a significant advantage: the ability to learn its parameters with a cost linear in time steps: $\mathcal{O}(T)$. To achieve this, we employ the variational Laplace EM inference algorithm proposed in the general recurrent state space framework for decision-making (Zoltowski et al., 2020).

If denoting the number of discrete states as Z, the parameters θ of MRM-GP can be categorized into two groups: (1) kernel parameters: $\{\sigma^{m,z}, \eta^{m,z}\}^{m_a,Z}_{m=1,z=1}$, $\{\sigma^{m,p}, \eta^{m,p}\}^{m_w,P}_{m=1,p=1}$, $\{\phi^{m,z}_{pp'}\}^{m_a,P,Z}_{m=1,p=1,p'=p+1,z=1}$; (2) emissions parameters: $\mathbf{C}, d, \mathbf{V}$. Additionally, the hyperparameters consist of the number of discrete states Z, the number of derivatives k, the kernel rank R, and the number of latent dimensions M. The value of Z depends on the data, k is discussed in Section 3.1 and Figure 2, M is determined through a cross-validation strategy (Section 5.2), and the rank R is consistently set to 2 to ensure positive definiteness without introducing many amplitude parameters. Besides, there is no need to learn the amplitude parameters, denoted as $\{a^{m,r}_p\}^{P,M,R}_{p=1,m=1,r=1}$, since the emissions parameter $\mathbf{C}$ fulfills a similar role in MRM-GP.

The variational Laplace EM inference algorithm alternatively updates discrete switching states $z \in \{1, \ldots, Z\}$, latent dynamics $f \in \mathbb{C}^{MPk\times T}$, and model parameters θ. The time complexity and memory storage of each step are all linear in time as follows: (1) updating z: $\mathcal{O}(Z)$, $\mathcal{O}(ZT)$; (2) updating f: $\mathcal{O}(T)$, $\mathcal{O}(2M^2P^2k^2T)$; (3) updating θ: $\mathcal{O}(ZMk)$, $\mathcal{O}(MPkT)$.

Furthermore, to avoid the calculation of the complex number when updating f in our implementations, we rewrite the complex latent dynamics f in Eq. 11 to be a joint signal in the real domain, such that the latent dynamics becomes:

$$\begin{aligned} \begin{bmatrix} f_r \\ f_i \end{bmatrix}_t &= \begin{bmatrix} \mathbf{A}_r & -\mathbf{A}_i \\ \mathbf{A}_i & \mathbf{A}_r \end{bmatrix} \begin{bmatrix} f_r \\ f_i \end{bmatrix}_{t-1} + \begin{bmatrix} q_r \\ q_i \end{bmatrix}_{t-1}, \\ \begin{bmatrix} q_r \\ q_i \end{bmatrix}_{t-1} &\sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{Q}_r & -\mathbf{Q}_i \\ \mathbf{Q}_i & \mathbf{Q}_r \end{bmatrix} \right), \end{aligned} \tag{12}$$

where $f_r, f_i, q_r, q_i, \mathbf{A}_r, \mathbf{A}_i, \mathbf{Q}_r, \mathbf{Q}_i$ are the real and imagery part of $f, q, \mathbf{A}, \mathbf{Q}$, respectively.

5. Experiments

Our code is available at `https://github.com/WeihanLikk/MRM-GP`.

Datasets. We evaluate MRM-GP on three datasets:

- **Synthetic Data**: We generate simulated data incorporating both across-region communications and within-region neural activities, along with time-varying frequencies and phase delays introduced by various states.
- **Local Field Potential Recordings (LFP)** (Siegle et al., 2021): Local Field Potential recordings from mouse's primary visual area (V1) and visual anteromedial area (VISam). The external stimulus consisted of an 8Hz drifting grating with eight orientation directions.
- **Neural Spike Trains** (Semedo et al., 2019; Zandvakili

& Kohn, 2019): Simultaneous spike trains from monkey's primary visual area (V1) and secondary visual cortex (V2). The external stimulus is a 6Hz drifting grating with eight orientation directions.

Baselines for comparison. We compare MRM-GP with two methods designed to discover the directional communications in the latent space of multi-region recordings:

- **DLAG** (Gokcen et al., 2022): A Gaussian Process Factor Analysis employs a multi-output Squared Exponential kernel. Its goal is to uncover simultaneous or bidirectional latent communications across different regions. The kernel function incorporates a time delay parameter to determine the directions for learned communications.
- **CSM-GPFA**: The Gaussian Process Factor Analysis, using a multi-region kernel as described in Section 2.2, is an extension of the model presented in (Hultman et al., 2018). This extension introduces a new classification assumption, distinguishing latent variables into across-region and within-region types.

Metrics. For every model and dataset, we fit the model on the training set, denoted as y_{train}, and test its performance on the test set y_{test}. Specifically, we randomly select some trials as y_{train}, while the remaining trials serve as y_{test}. Additionally, we randomly divide the test data y_{test} into two parts: $y_{\text{test}}^{\text{held-in}}$ with 90% neurons as held-in test data and $y_{\text{test}}^{\text{held-out}}$ with 10% neurons as held-out test data. We infer $x_{\text{test}}^{\text{held-in}}$ based on $y_{\text{test}}^{\text{held-in}}$, which is then used as the test latent variables when computing **test log-likelihood (LL)** $p(y_{\text{test}}^{\text{held-out}}|x_{\text{test}}^{\text{held-in}};\theta)$(Pei et al., 2021), serving as the final metric in our experiments. To reduce the randomness when creating $y_{\text{test}}^{\text{held-in}}$ and $y_{\text{test}}^{\text{held-out}}$, we also average $p(y_{\text{test}}^{\text{held-out}}|x_{\text{test}}^{\text{held-in}};\theta)$ over five distinct partitions.

5.1. Synthetic Data

This section aims to assess how well MRM-GP can identify switching states, latent variables, and parameters.

Experimental setup. We generate 50 independent trials for two brain regions $P = 2$, where each region has 30 neurons, $m_a = 1$ dimension across-region variables, and $m_w = 1$ dimension within-region variable. We also introduce the time-varying across-region frequencies and phase delays by two discrete states $Z = \{z_1, z_2\}$: (1) state 1: $\eta^{z_1,a} = 1.0$ rad/s, $\phi_{1,2}^{z_1} = -10$ ms, $\sigma^{z_1,a} = 10$, state 2: $\eta^{z_2,a} = 0.25$ rad/s, $\phi_{1,2}^{z_2} = 10$ ms, $\sigma^{z_2,a} = 10$. Different sign of $\phi_{1,2}$ means the change of directions. We set $\eta^w = 0.75$ rad/s and $\sigma^w = 10$ for within-region variables. For the generative and inference process, we set hyperparameters $k = 2$, $R = 2$, and compare the test log likelihood when $Z = 1, 2, 3$.

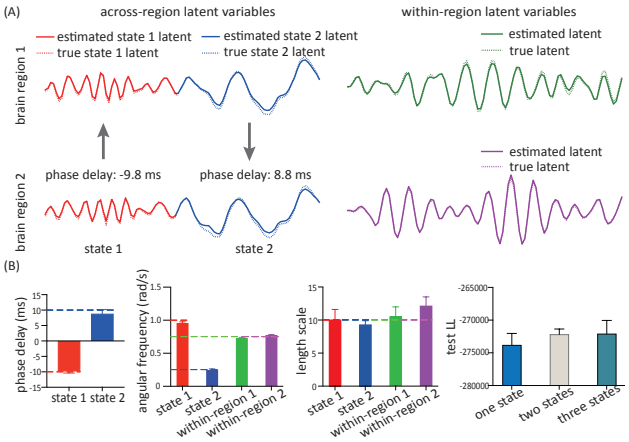

Figure 3. Applying MRM-GP to synthetic data. **(A)**: Compare the estimated latent variables and discrete states with the ground truth. MRM-GP accurately identifies two states with time-varying frequencies and phase delays, aligning with the ground truth. **(B)**: Compare learned parameters with the ground truth indicated by dashed lines. **(C)**: Examine the test LL with varying numbers of discrete states Z. The findings indicate that $Z = 2$ and $Z = 3$ exhibit similar LL, both larger than $Z = 1$.

Results. We fit an MRM-GP to the synthetic data, specifying $m_a = 1$ dimension across-region variables, $m_w = 1$ dimension within-region variable, and $Z = 2$ states. Figure 3(A) shows single-trial latent variable estimations that accurately reflect the latent dynamics and communications influenced by discrete states over two brain regions. State z_1 (depicted in blue) exhibits a periodic signal with a higher frequency and forward communication from brain region 1 to brain region 2. In contrast, state z_2 (shown in purple) displays an oscillatory pattern with a lower frequency and feedback communication from region 2 to region 1.

For a quantitative assessment of learned parameters, Figure 3(B) displays the estimated phase delays, frequencies, and length scales across different initializations, demonstrating close alignment with ground truths. Furthermore, Figure 3(C) illustrates the test log-likelihood for varying Z, revealing that both $Z = 2$ and $Z = 3$ provide similar and superior estimations compared to $Z = 1$ (see Appendix J for $Z = 3$ visualization). We also have synthetic experiments about parameter initialization and different parameter setting in Appendix F,G,H.

5.2. Local Field Potential Recordings

This section aims to explore interactions between the mouse's primary visual area (V1) and the visual anteromedial area (VISam) in the presence of an 8Hz drifting grating. We also aim to compare the performance and inference time cost with other multi-region methods, namely

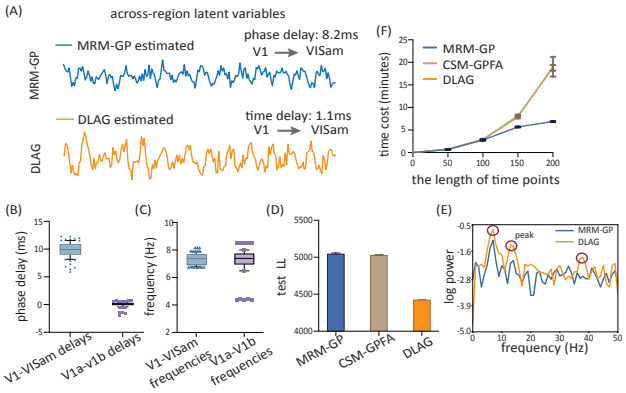

Figure 4. Applying MRM-GP to LFP recordings. **A**: Compare the across-region latent variables with DLAG. **B-C**: Visualize the learned phase delays and frequencies from V1-VISam and V1a-V1b. **D**: Compare the test LL with other multi-region methods: DLAG and CSM-GPFA. **E**: The power spectrum of across-region latent variables in **A**. The latent variable from DLAG exhibits three frequency peaks, while MRM-GP's latent variable has only one peak. **F**: Inference time comparison: MRM-GP has a linear time cost.

DLAG and CSM-GPFA.

Experimental setup. We conduct experiments using two sessions, each comprising eight orientation directions, resulting in 16 datasets. Each dataset includes 15 trials (10 as a training set and 5 as a testing set) of continuous-time local field potential recordings from approximately 20 neurons in V1 and approximately 25 in VISam. The initial sampling rate is 1000Hz, and we downsample it to 100Hz, resulting in 200 time points with 10ms bin size. We set hyperparameters $k = 2$, $R = 2$.

To determine the dimensionalities of across- and within-region latent variables, we adopt the approach outlined in (Gokcen et al., 2022). Initially, we apply Factor Analysis to identify the total number of latent variables required to elucidate the neural recordings for each region. A 5-fold cross-validation was employed to select the configuration yielding the highest test LL. Subsequently, given the selected total number of latent variables (M), we conduct a grid search for the dimensionalities of across- (m_a) and within-region (m_w), respectively. For each pair of (m_a, m_w), we run 5-fold cross-validation with MRM-GP and chose the setting with the highest test LL. Given this procedure, our final choice was $m_a = 1$ across-region variables and $m_w = 3$ within-region variables for both V1 and VISam. See Appendix B for the full comparison.

Results. We applied an MRM-GP to local field potential recordings with $Z = 1$ state. Figure 4(A) shows a comparison of single-trial across-region latent variables for one dataset (orientation 135°, session 721123822), and the within-region variables are in Appendix B. Both latent variables demonstrate an oscillatory structure, capturing characteristics of the external 8Hz drifting grating stimulus.

The MRM-GP's latent variable in this dataset is linked to a communication direction from V1 to VISam with an 8.2 ms phase delay, and the DLAG's latent variable shows a 1.1 ms time delay. Both delays fall within a single time bin (10ms) and are positive, suggesting consistent communication direction from V1 to VISam. The difference in their values arises because DLAG models time delay ($\delta_{pp'}$) in the kernel equation $K_{pp'}(\tau) = \exp(-\frac{1}{2\sigma^2}(\tau - \delta_{pp'})^2)$, which is independent of frequency and has a different interpretation from the phase delay ($\phi_{pp'}$). In this context, $\phi_{pp'}$ represents the delay in a specific frequency band. In contrast, $\delta_{pp'}$ signifies the delay for a latent variable with a mixture of multiple frequencies, as evidenced by its power spectrum with three frequency peaks (Figure 4(E)). Therefore, the divergence in values is acceptable if their directions align.

The left chart in Figure 4(B) illustrates the estimated phase delays across all 16 datasets. Each data point represents the phase delay for an individual run on a specific dataset. The findings suggest consistent communication from V1 to VISam across all datasets, with the phase delays clustered around 7.5Hz (the left chart in Figure 4(C)), which is consistent with external 8Hz stimulus.

To demonstrate that the MRM-GP itself is not the cause of delays, we first divide V1 randomly into two parts, V1a and V1b, each with channels of equal size. We then estimated the phase delays between them. The right chart in Figure 4(B) illustrates that across 16 datasets, all phase delays hover around zero within frequency bands around 7.5Hz (the right chart in Figure 4(C)). This suggests that the learned delays are a consequence of the data rather than the model.

Figure 4(D) shows that MRM-GP, a linear dynamics system approximation of CSM-GPFA, exhibits a similar test LL to CSM-GPFA. The higher test LL compared to DLAG suggests that the multi-region kernel (Eq. 1) outperforms DLAG's Squared Exponential kernel on these datasets. This is attributed to the former explicitly modeling frequencies through its kernel parameters and having a better frequency separation. Specifically, the across-region variable of the former has only one prominent frequency, whereas DLAG's across-region variable exhibits three peaks in Figure4(E), keeping consistent with the data spectrum in Appendix E.

Lastly, in Figure 4(F), we compare the inference time of MRM-GP, CSM-GPFA, and DLAG for 500 iterations. We achieved this by downsampling the recordings and creating four datasets with varying lengths of time points: 50, 100, 150, and 200. The results indicate that the time cost of MRM-GP increases linearly, whereas both CSM-GPFA and

DLAG exhibit cubic growth.

5.3. Neural Spike Trains

This section aims to evaluate MRM-GP's ability to identify switching states within the communications subspace while also discovering across-region communications using a distinct type of neural data.

Experimental setup. The simultaneous spike trains were obtained from the monkey's primary visual area (V1) and secondary visual cortex (V2) in the presence of a 6Hz moving grating. This dataset comprises four sessions, each featuring eight orientation directions, resulting in 32 datasets. Each dataset comprises 400 trials (64 time points with 20ms bin size for every trial), with 300 trials randomly selected as the training set and 100 trials as the testing set. In V1, there are approximately 90 neurons, while in V2, there are around 20 neurons. We set hyperparameters $k = 2$, $R = 2$.

Results. We fitted an MRM-GP to neural spikes trains with $m_a = 2$ dimensions across-region variables, $m_w = 2$ dimension within-region variable, and $Z = 2$ states. The configuration of m_a and m_w follows previous work (Gokcen et al., 2022) and adopts the same strategy mentioned in Section 5.2.

Figure 5(A) shows the across-region latent variables for one dataset (orientation 0°, session 106r001p26, ten trials are displayed, all variables are scaled by the variance explained in each region), and the within-region variables are in Appendix B. These latent variables indicate time-varying forward and feedback communications between V1 and V2. Different states exhibit distinct phase delays and frequencies. The first dimension of across-region variables (denoted as x_1^a) displays a periodic pattern caused by the external drifting grating stimulus, whereas the second dimension (denoted as x_2^a) exhibits a non-periodic signal with a single peak shortly after the stimulus onset.

Figure 5(B-C) presents the estimated phase delays and frequencies over multiple independent runs for 32 datasets. Each data point represents a dimension of across-region variables. Figure 5(B) corresponds to state z_1, indicating that most state z_1 dimensions exhibit across-region interactions within the 2Hz-8Hz range. Additionally, some dimensions display feedback communication with a large phase delay (>10ms) from V2 to V1, corresponding to state z_1 of x_2^a in Figure 5(A). However, there is variability across datasets for state z_1 with smaller phase delays (<10ms). Some show forward communication from V1 to V2, akin to state z_1 of x_1^a in Figure 5(A), while others indicate feedback communications from V2 to V1.

One explanation for this variability is that in certain datasets, x_1^a has a much weaker amplitude compared to x_2^a, making the weaker latent affected by the stronger one along with its delay. This leads to a feedback signal from V2 to V1 at state z_1. On the other hand, in some datasets (e.g., orientation 0°, session 106r001p26), x_1^a is not as weak, resulting in a forward signal from V1 to V2 at state z_1.

Figure 5(C) depicts the estimated phase delays and frequencies associated with state z_2. The findings indicate a clear separation of oscillatory communications into two frequencies. One involves 6Hz communications with small phase delays (referring to state z_2 of x_1^a in Figure 5(A)), while the other involves 1Hz communications with large phase delays (akin to state z_2 of x_2^a in Figure 5(A)).

The time-varying phase delays can be explained as follows: (1) For x_1^a, V1 triggers V2 to have oscillatory dynamics during state z_1, while in state z_2, V2 is already engaged, causing both regions to oscillate synchronously, resulting in a smaller phase delay than in state z_1. (2) For x_2^a, V2 consistently sends signals with a low frequency to V1, resulting in a larger phase delay due to the longer period as indicated by z_2. During state z_1, the stimulus onset triggers an intense signal from V2 to V1, leading to a smaller phase delay, which can be considered as an emergence of surprise or prediction error from V2 to V1 (Rao & Ballard, 1999).

Similar to Section 5.2, we also perform a control experiment by learning the phase delays between V1a and V1b. In Figure 5(D), the outcomes reveal zero-delay communications that are distributed across two frequencies (6Hz, 1Hz), suggesting that learned delays are a consequence of the data rather than the model.

Finally, we compare the test LL of MRM-GP, CSM-GPFA, and DLAG in Figure 5(E). The results indicate that MRM-GP with $Z = 2$ states achieves the highest LL, while MRM-GP with $Z = 1$ state exhibits a similar LL compared to CSM-GPFA, and both outperform DLAG. This suggests that (1) switching states exist in these datasets; (2) the multi-region kernel is more appropriate than the Squared Exponential kernel for modeling signals with sinusoidal structures.

6. Discussion

MRM-GP establishes the connection between a linear dynamics system (LDS) and a multi-output Gaussian Process (GP) explicitly modeling frequency-based communications and their directionality via phase delays within the latent space of neural data.

Connecting a complex-valued GP with an LDS is non-trivial. Although a complex-valued GP can be written as a multi-output real-valued GP (by twice the dimension), the resulted multi-output GP cannot be converted to an LDS by spectral factorization because the separability of the resulted multi-output kernel is not guaranteed.

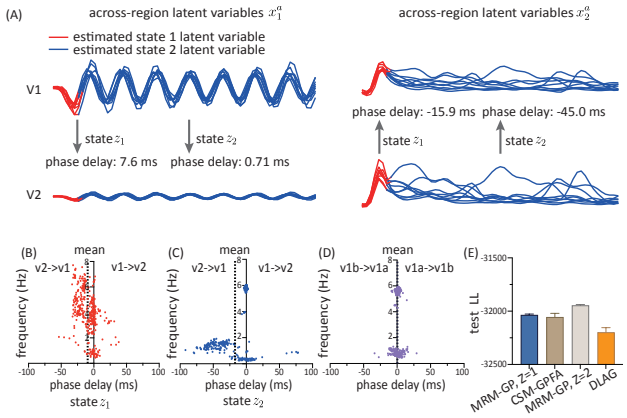

Figure 5. Applying MRM-GP to neural spike trains. **A**: Visualize across-region latent variables with state z_1 (in red) and state z_2 (in blue). The results suggest forward and feedback communications between V1 and V2, along with time-varying delays. **B-C**: Visualize estimated parameters through phase delays on the x-axis and frequencies on the y-axis. **D**: Depict phase delays and frequencies learned in the V1a-V1b control experiment, with their representation on the x-axis and y-axis, respectively. **E**: Compare the test LL with discrete states $Z = 1, 2$ and other multi-region methods: DLAG and CSM-GPFA.

Once the link is established, we can harness several advantages: (1) using the powerful representational capability of kernels, such as applying a multi-region kernel to model latent variables with periodic patterns; (2) achieving a linear computational cost; (3) incorporating time-varying frequencies and delays by introducing different discrete states.

We test MRM-GP using two distinct types of neural data. The findings showcase its capability to discover state-dependent latent communications across brain regions with a linear time inference cost.

Finally, the limitations of MRM-GP are twofold: (1) its reliance on separability for multi-output kernels, which restricts kernel selection options, and (2) its current model assumptions are unable to capture phenomena such as phase resetting and phase variability across different trials.

Acknowledgement

This work is supported by National Institutes of Health BRAIN initiative (1U01NS131810).

Impact Statement

The MRM-GP introduces an innovative and efficient method for investigating intricate interactions among brain regions. Its capacity to deliver an interpretable representation of multi-region neural data is poised to advance neuroscience, offering the potential for a more profound comprehension of brain function and disorders. This enhanced understanding of brain interactions has the prospect to drive advancements in neurotechnology, with potential benefits extending to fields such as brain-computer interfaces and personalized medicine.

References

Fox, E., Sudderth, E., Jordan, M., and Willsky, A. Nonparametric bayesian learning of switching linear dynamical systems. *Advances in neural information processing systems*, 21, 2008.

Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L., and Linderman, S. Recurrent switching dynamical systems models for multiple interacting neural populations. *Advances in neural information processing systems*, 33: 14867–14878, 2020.

Gokcen, E., Jasper, A. I., Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., and Yu, B. M. Disentangling the flow of signals between populations of neurons. *Nature Computational Science*, 2(8):512–525, 2022.

Gokcen, E., Jasper, A. I., Xu, A., Kohn, A., Machens, C. K., and Byron, M. Y. Uncovering motifs of concurrent signaling across multiple neuronal populations. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Grewal, M. S. and Andrews, A. P. *Kalman filtering: Theory and Practice with MATLAB*. John Wiley & Sons, 2014.

Harris, K. D. and Mrsic-Flogel, T. D. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013.

Hartikainen, J. and Särkkä, S. Kalman filtering and smoothing solutions to temporal gaussian process regression models. In *2010 IEEE international workshop on machine learning for signal processing*, pp. 379–384. IEEE, 2010.

Hom, R. A. and Johnson, C. R. Matrix analysis. *Cambridge University Express*, 455, 1985.

Hultman, R., Ulrich, K., Sachs, B. D., Blount, C., Carlson, D. E., Ndubuizu, N., Bagot, R. C., Parise, E. M., Vu, M.-A. T., Gallagher, N. M., et al. Brain-wide electrical spatiotemporal dynamics encode depression vulnerability. *Cell*, 173(1):166–180, 2018.

Kailath, T., Sayed, A. H., and Hassibi, B. *Linear estimation*. Number BOOK. Prentice Hall, 2000.

Kohn, A., Jasper, A. I., Semedo, J. D., Gokcen, E., Machens, C. K., and Byron, M. Y. Principles of corticocortical communication: proposed schemes and design considerations. *Trends in Neurosciences*, 43(9):725–737, 2020.

Miller, E. K., Lundqvist, M., and Bastos, A. M. Working memory 2.0. *Neuron*, 100(2):463–475, 2018.

Pei, F., Ye, J., Zoltowski, D., Wu, A., Chowdhury, R. H., Sohn, H., O'Doherty, J. E., Shenoy, K. V., Kaufman, M. T., Churchland, M., et al. Neural latents benchmark'21: evaluating latent variable models of neural population activity. *arXiv preprint arXiv:2109.04463*, 2021.

Rao, R. P. and Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.

Särkkä, S., Solin, A., and Hartikainen, J. Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering. *IEEE Signal Processing Magazine*, 30 (4):51–61, 2013.

Sayed, A. H. and Kailath, T. A survey of spectral factorization methods. *Numerical linear algebra with applications*, 8(6-7):467–496, 2001.

Semedo, J. D., Zandvakili, A., Machens, C. K., Byron, M. Y., and Kohn, A. Cortical areas interact through a communication subspace. *Neuron*, 102(1):249–259, 2019.

Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, 592(7852):86–92, 2021.

Solin, A. and Särkkä, S. Explicit link between periodic covariance functions and state space models. In *Artificial Intelligence and Statistics*, pp. 904–912. PMLR, 2014.

Solin, A. et al. Stochastic differential equation methods for spatio-temporal gaussian process regression. 2016.

Ulrich, K. R., Carlson, D. E., Dzirasa, K., and Carin, L. Gp kernels for cross-spectrum analysis. *Advances in neural information processing systems*, 28, 2015.

Wang, Y., Wu, Z., Li, C., and Wu, A. Extraction and recovery of spatio-temporal structure in latent dynamics alignment with diffusion model. *Advances in Neural Information Processing Systems*, 36, 2024.

Zandvakili, A. and Kohn, A. Simultaneous v1–v2 neuronal population recordings in anesthetized macaque monkeys. *CRCNS https://doi. org/10.6080/K0B27SHN*, 2019.

Zhu, H., Balsells-Rodas, C., and Li, Y. Markovian gaussian process variational autoencoders. In *International Conference on Machine Learning*, pp. 42938–42961. PMLR, 2023.

Zoltowski, D., Pillow, J., and Linderman, S. A general recurrent state space framework for modeling neural dynamics during decision-making. In *International Conference on Machine Learning*, pp. 11680–11691. PMLR, 2020.

A. Spectral Factorization.

Derivation for Eq. 6. Start with Eq. 3, taking Fourier transforms on both sides gives:

$$(i\omega)\mathbf{J}(i\omega)_m^{p,w} = \mathbf{F}_m^w \mathbf{J}(i\omega)_m^{p,w} + \mathbf{L}\mathbf{U}(i\omega). \tag{13}$$

Solving for $\mathbf{J}(i\omega)_m^{p,w}$ gives:

$$\mathbf{J}(i\omega)_m^{p,w} = ((i\omega - \mathbf{F}_m^w)\mathbf{I})^{-1}\mathbf{L}\mathbf{U}(i\omega). \tag{14}$$

Recall that $f(t)_m^{p,w}$ in Eq. 3 is a Complex Gaussian Process with single-output kernel $\mathbf{K}(\tau)^m = \sum_{r=1}^{R} a_r^{m\,2} \exp(-\frac{1}{2\sigma^{m\,2}}\tau^2 + i\eta^m(\tau))$ and its derivatives up to $(k-1)^{th}$. So, the spectral density matrix of this process $f(t)_m^{p,w}$ is:

$$\mathbf{S}_J(\omega) = \mathbb{E}[\mathbf{J}(i\omega)_m^{p,w} J(-i\omega)_m^{p,w\,\top}]. \tag{15}$$

Bring Eq. 14 into Eq. 15 gives:

$$\begin{aligned} \mathbf{S}_J(\omega) =& (\mathbf{F}_m^w - i\omega\mathbf{I})^{-1}\mathbf{L}\mathbb{E}[\mathbf{U}(i\omega)U(-i\omega)^\top]\mathbf{L}^\top(\mathbf{F}_m^w + i\omega\mathbf{I})^{-T}, \\ =& (\mathbf{F}_m^w - i\omega\mathbf{I})^{-1}\mathbf{L}v\mathbf{L}^\top(\mathbf{F}_m^w + i\omega\mathbf{I})^{-T}. \end{aligned} \tag{16}$$

Finally $S(\omega)$ in Eq. 6 is:

$$S(\omega) = \mathbf{G}\mathbf{S}_J(\omega)\mathbf{G}^\top. \tag{17}$$

Finding roots for $T(i\omega)$ in Eq. 8. Using the $T(i\omega)$'s coefficients $b_0, b_1, \ldots, b_{2k}$, we can create a companion matrix:

$$\mathbf{B} = \begin{bmatrix} 0 & 0 & \ldots & 0 & -\frac{b_{2k}}{b_0} \\ 1 & 0 & \ldots & 0 & -\frac{b_{2k-1}}{b_0} \\ \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & -\frac{b_1}{b_0} \end{bmatrix}, \tag{18}$$

where the eigenvalues of this matrix are the roots for $T(i\omega)$ (Hom & Johnson, 1985). Notably, the companion matrix is structured as a Hessenberg matrix, suggesting that its eigenvalues can be obtained through the QR algorithm with Givens rotation. This process has a time complexity of $\mathcal{O}(k)$ for each iteration.

B. DLAG latent variables for V1-V2 spike train data

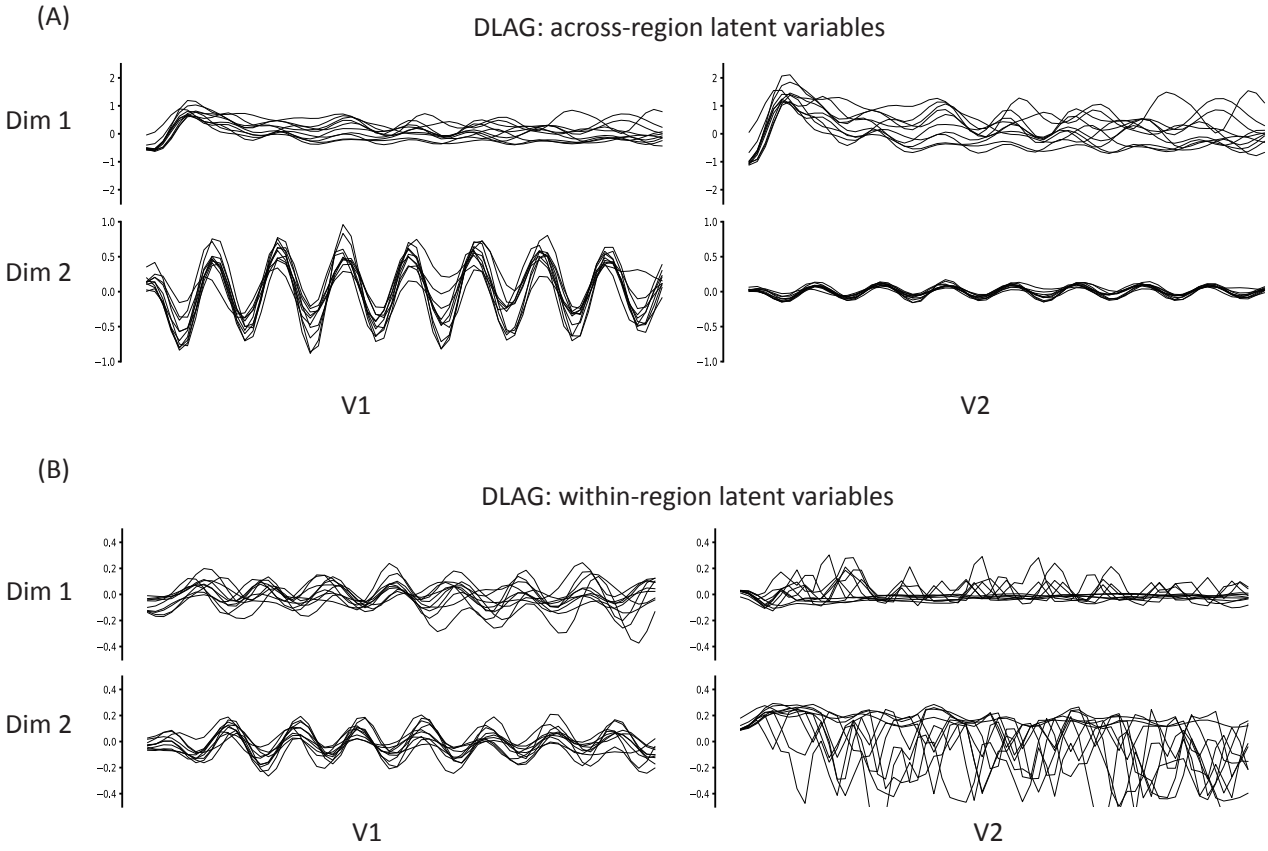

Figure 6. Across- and within-region latent variables for DLAG on V1-V2 spike train data. DLAG's two across-region latent variables both have some oscillations. However, there are fewer oscillations in the first dimension of our model's across-region variables (Figure 5), indicating a better disentanglement of frequencies.

C. Within-region latent variables for V1-V2 spike train data

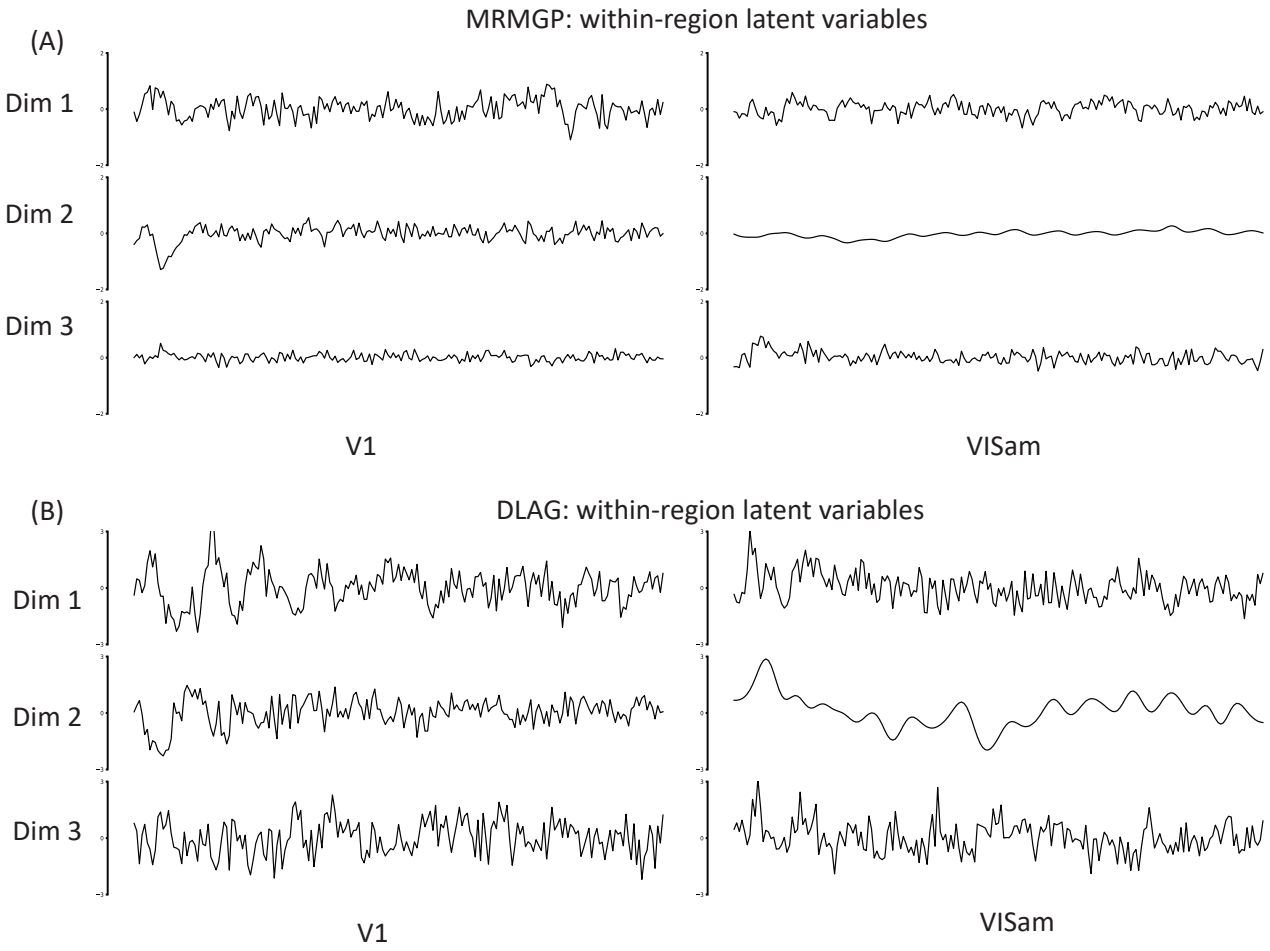

Figure 7. Within-region latent variables for MRM-GP and DLAG on local field potential recordings from V1 and VISam. MRMGP and DLAG exhibit similar within-region latent variables associated with high-frequency noise, indicating neural activities specific to a certain region.

D. Compare latent variables with different k

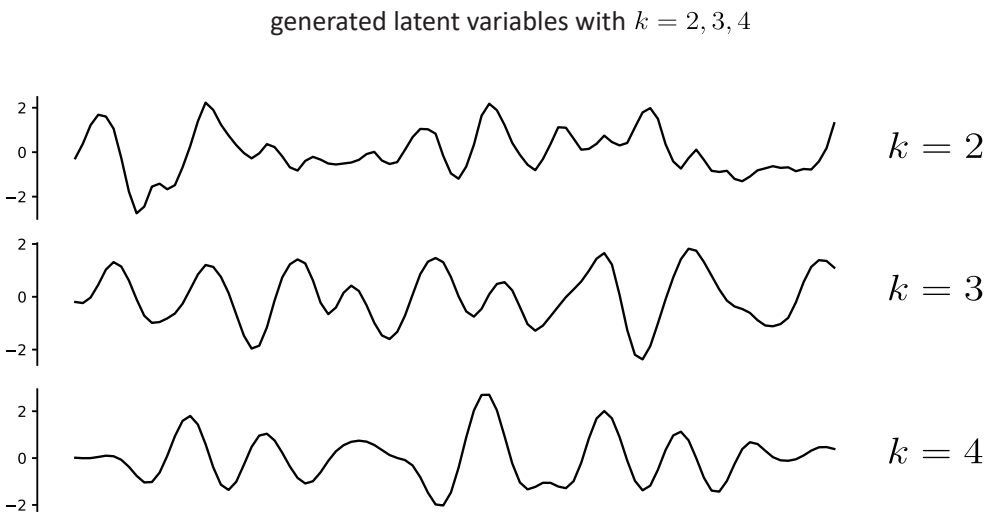

Figure 8. Compare the smoothness of latent variables with $k = 2, 3, 4$. We generate latent variables with the same frequency, length scale and phase delays, while changing the value of k. The results show that a larger k makes the latent variable smoother.

E. Spectrum of the LFP data

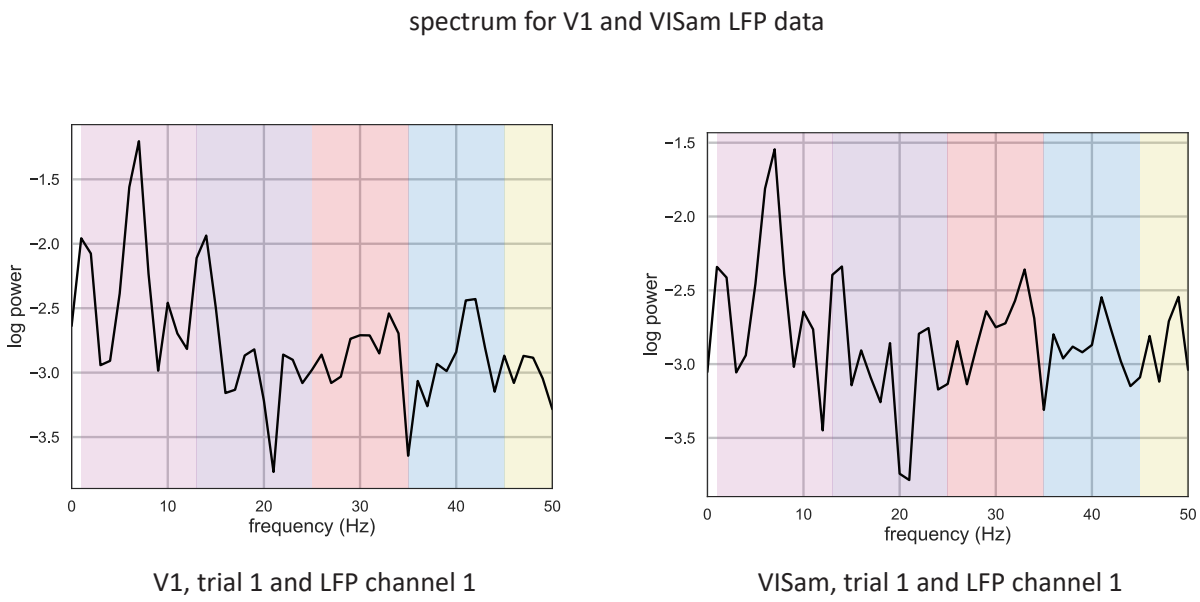

Figure 9. The spectrum for V1 and VISam LFP data with one trial and one channel shown. The multiple peaks in both V1 and VISam spectra indicate a mix of frequencies. In Figure 4, the spectrum of DLAG's across-region latent variable has three peaks, while our model's across-region latent variable has only one dominant peak, indicating a better frequency separation brought by the multi-region kernel in Eq.1.

F. Synthetic data with a wide range of frequencies and delays

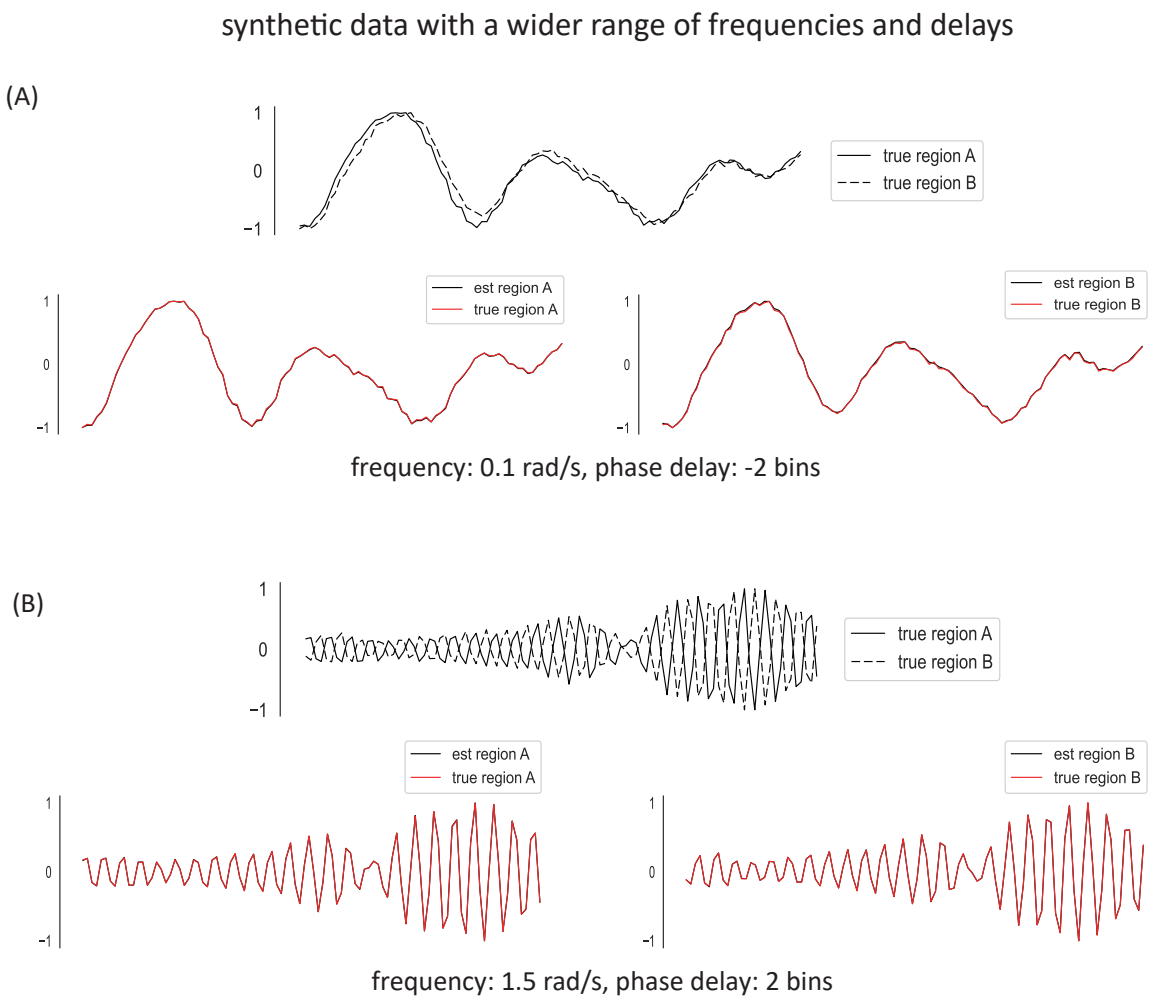

Figure 10. Synthetic data experiments with a wide range of frequencies and delays. We tested our model on a wide range of parameters, where the frequencies vary from 0.1 rad/s to 1.5 rad/s, and the delays vary from -2 to 2. Under all these cases, our model could perfectly recover the latent variables.

G. Synthetic data with different parameter initialization

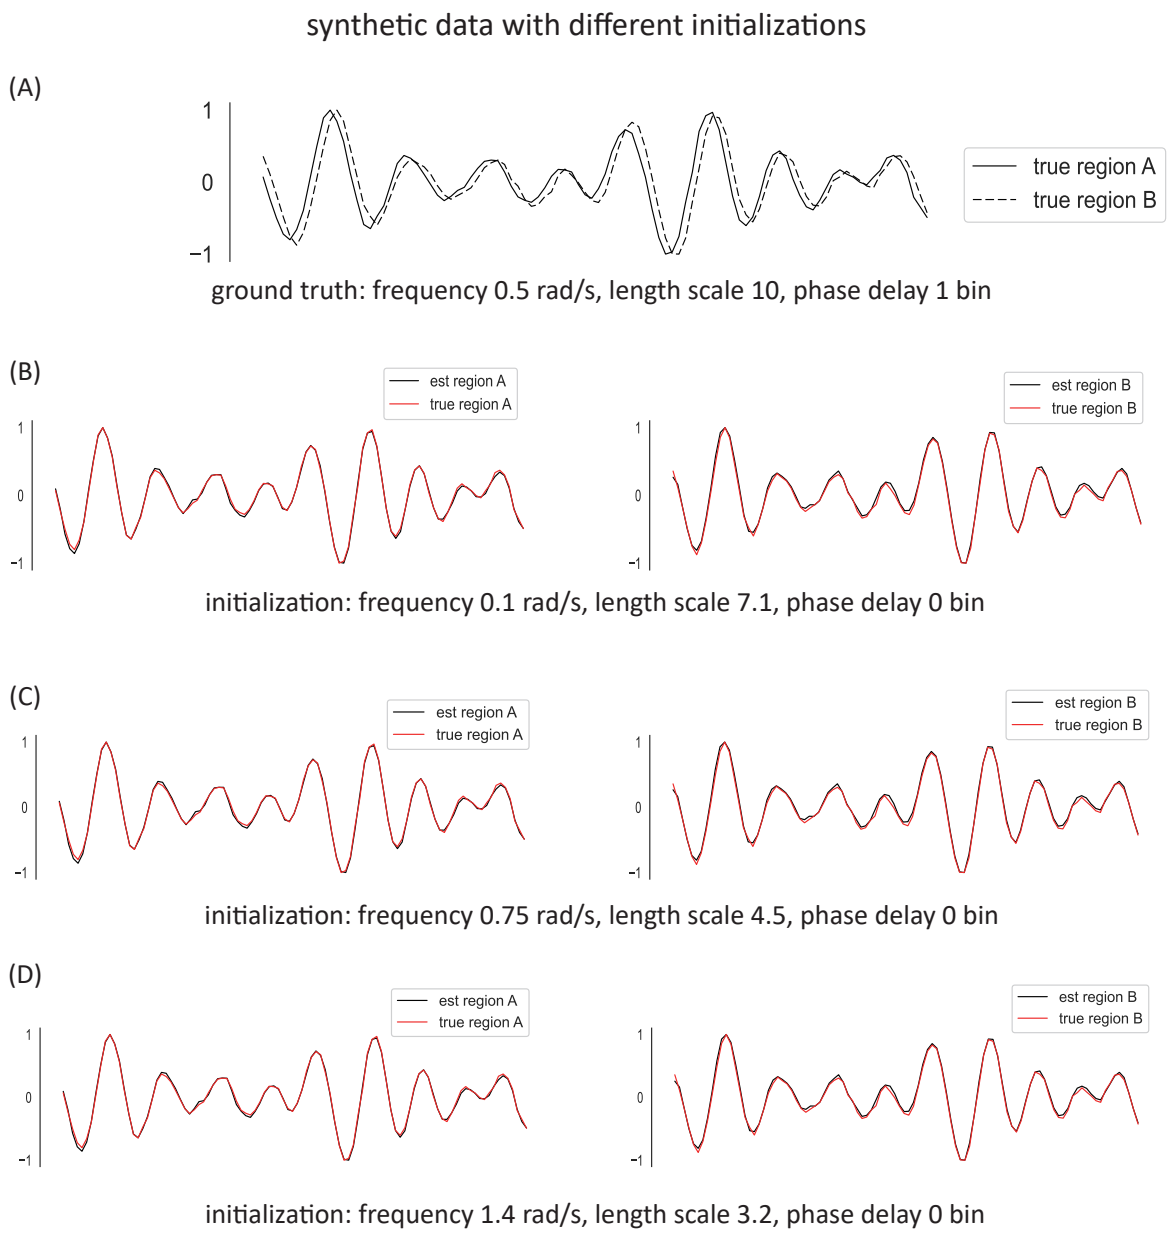

Figure 11. Synthetic data experiments with different parameter initialization. We tested our model with three initialization settings: from smoothness to unsmoothness. Under all these cases, our model could perfectly recover the latent variables, indicating that our model is not very sensitive to parameter initialization.

H. Synthetic data without pure oscillation

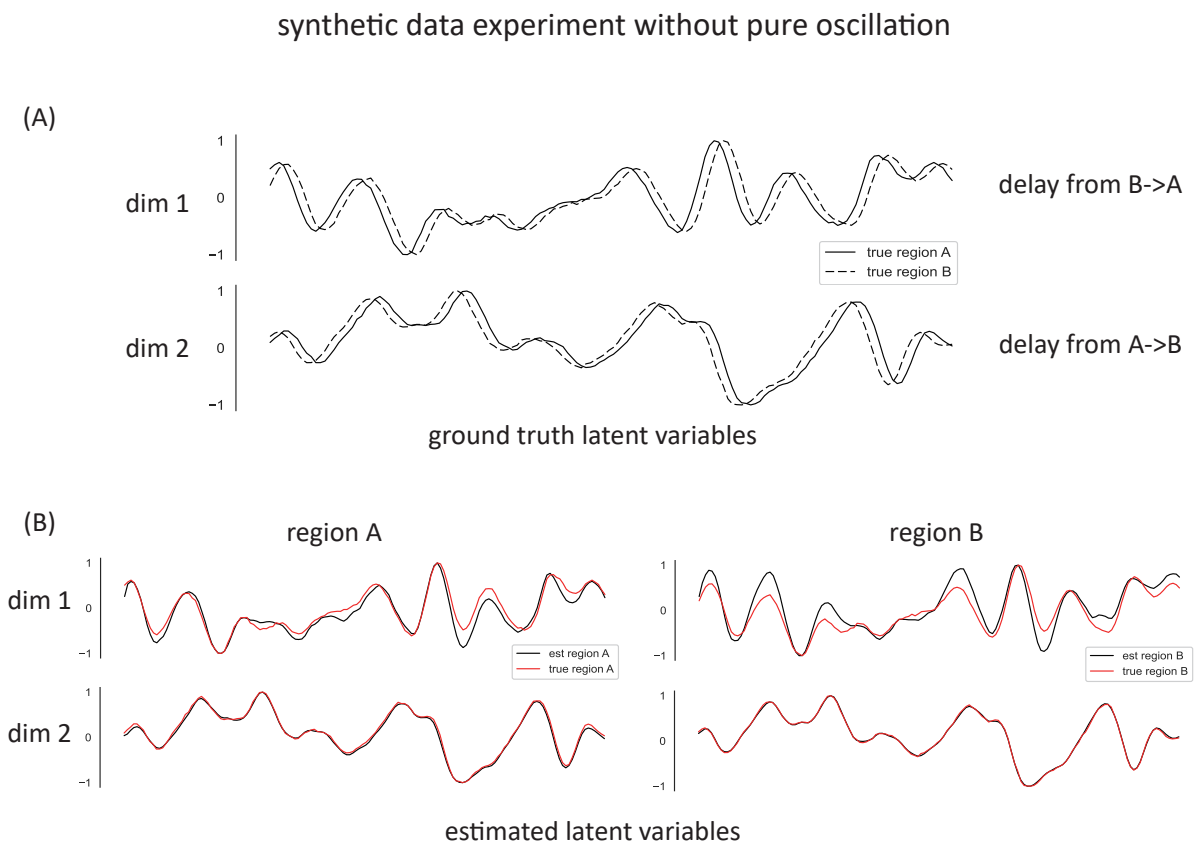

Figure 12. Synthetic data experiment without pure oscillation. We generated latent variables from a Gaussian Process with SE kernel. We added a manual shift to create delays, where one latent communication has a positive delay from Region A to Region B while the other has an opposite delay. The results show that our model can recover the true latent variables and correct delays between different regions. In summary, this demonstrates our model's ability to handle non-periodic data.

I. Cross-validating for the dimensionality of MRM-GP

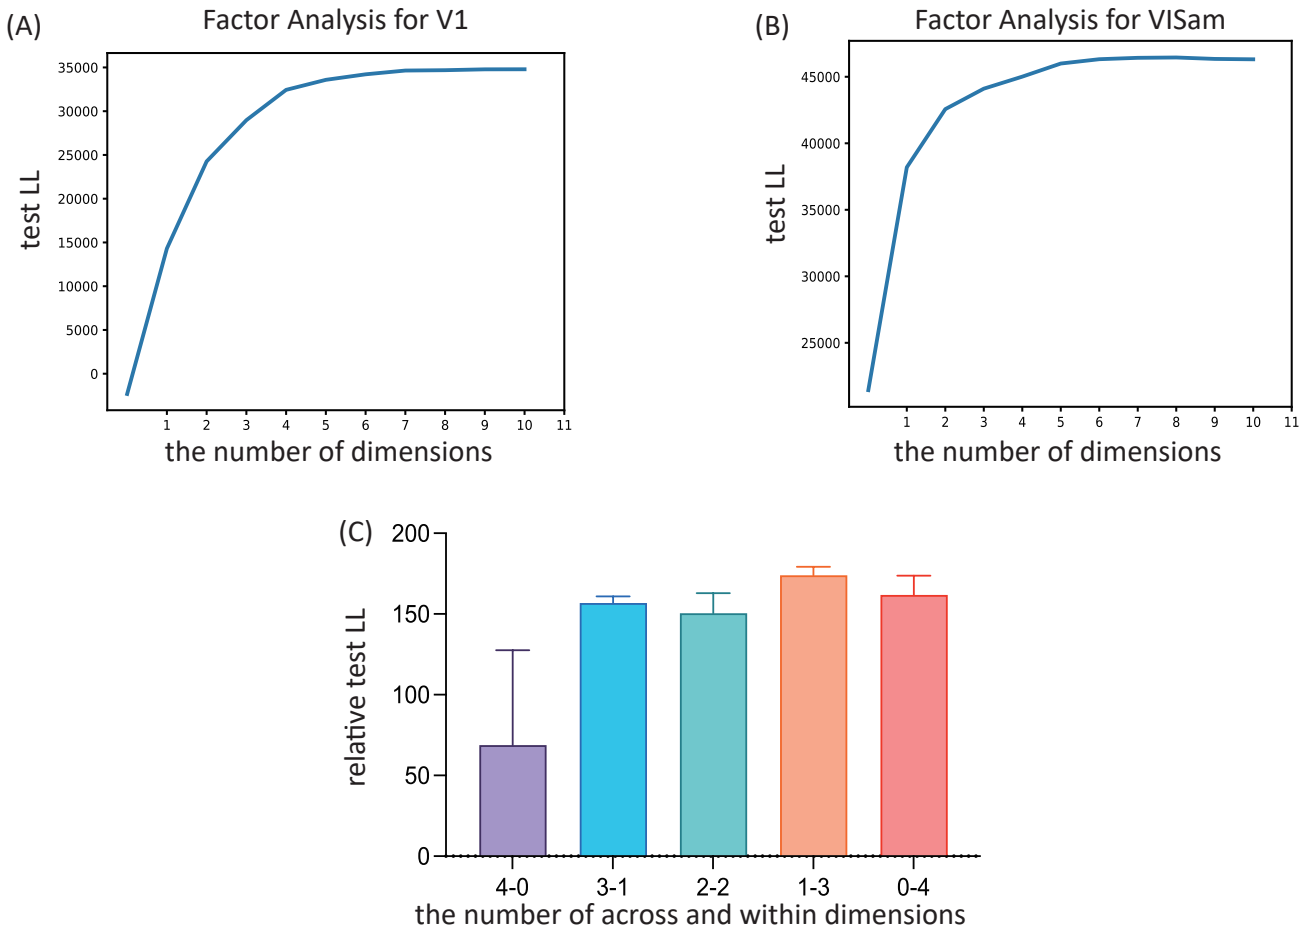

Figure 13. Determine the number of dimensions: (A-B) Factor Analysis is applied to V1 and VISam, revealing a gradual increase in test log-likelihood (LL) when the number of dimensions surpasses four (averaged over 5-fold cross-validation) . Considering both the inference time for all three models, we opt for four as the total dimensions for each region. (C) Conducting a grid search for across- and within-region variables using relative test LL as the metric—the difference between the model's test LL and Factor Analysis's test LL. The outcomes indicate that one across-region and three within-region variables yield the highest LL.

J. Applying MRM-GP to synthetic data with $Z = 3$

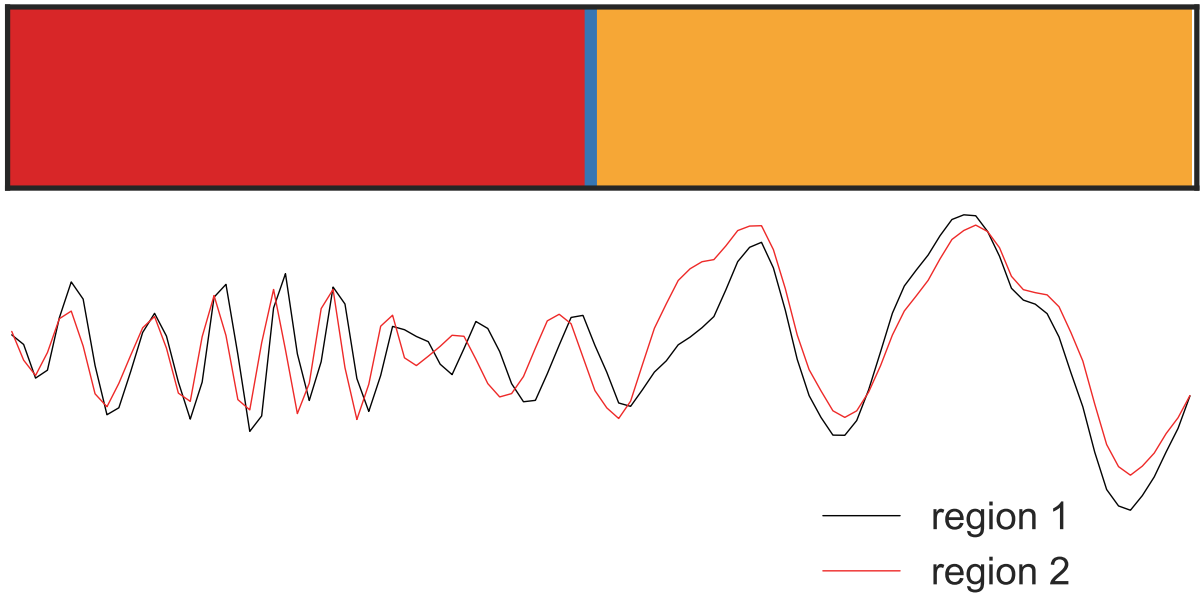

Figure 14. Applying MRM-GP to synthetic data with $Z = 3$. The result closely resembles Figure 3 when $Z = 2$, suggesting that the model can effectively learn the correct number of discrete states even specifying a large Z.