Gaussian Process Kernels for Cross-Spectrum Analysis

Kyle Ulrich [1] David Carlson [1] Kafui Dzirasa [2] Lawrence Carin [1]

[1]Department of Electrical and Computer Engineering, Duke University, Durham, NC

[2]Department of Psychiatry and Behavioral Sciences, Duke University, Durham, NC

INTRODUCTION

- Electrophysiological time-series data (e.g., LFPs) are quasi-periodic signals
  - The frequencies present in these time-series measurements are often highly correlated between regions
  - Phase synchrony patterns dependent on frequency band
- While multi-output Gaussian processes may be used to represent cross-amplitude spectra between channels, the lack of a cross-phase spectrum limits the utility of cross-spectrum analysis
- Our **cross-spectral mixture kernel** for multi-output Gaussian proceses is
  - Capable of representing the full cross-spectrum of multiple signals
  - Simple to interpret
  - Easily applied to neuroscience research for electrophysiological time-series

MULTI-OUTPUT GAUSSIAN PROCESSES

- **Data**: For input x_n, consider C output channel observations, $\boldsymbol{y}_n = [y_{n1}, \dots, y_{nC}]^T$
- An unobserved **latent function** $\boldsymbol{f}(x) = [f_1(x), \dots, f_C(x)]^T$ generates observations:

$$\boldsymbol{y}_n \sim \mathcal{N}(\boldsymbol{f}(x_n), \boldsymbol{H}^{-1}) \qquad \boldsymbol{H} = \text{diag}(\eta_1, \dots, \eta_C) \qquad (1)$$

- A **Gaussian process** prior on the latent function is formalized by

$$\boldsymbol{f}(x) \sim \mathcal{GP}(\boldsymbol{m}(x), \boldsymbol{K}(x, x')), \qquad (2)$$

  - The **mean function**: $\boldsymbol{m}(x) \in \mathbb{R}^C$ is often set to equal $\boldsymbol{0}$
  - The **covariance function**: $(\boldsymbol{K}(x, x'))_{c,c'} = k^{c,c'}(x, x') = \text{cov}(f_c(x), f_{c'}(x'))$
- The **Linear Model of Coregionalization** is given by

$$k^{c,c'}(x, x') = \sum_{q=1}^{Q} b_q(c, c') k_q(x, x'), \quad \text{or} \quad \boldsymbol{K}(x, x') = \sum_{q=1}^{Q} \boldsymbol{B}_q k_q(x, x'), \qquad (3)$$

  - **Input space kernel**: $k_q(x, x')$
  - **Output space kernel**: $b_q(c, c')$, **Coregionalization matrix**: $\boldsymbol{B}_q \in \mathbb{R}^{C \times C}$
- Evaluations $\boldsymbol{f} = [f_1(\boldsymbol{x}), \dots, f_C(\boldsymbol{x})]^T$ at locations $\boldsymbol{x} = [x_1, \dots, x_N]^T$ have a multivariate normal distribution $\mathcal{N}(\boldsymbol{f}; \boldsymbol{0}, \boldsymbol{K})$

$$\boldsymbol{K} = \begin{bmatrix} k^{1,1}(\boldsymbol{x}, \boldsymbol{x}) & \cdots & k^{1,C}(\boldsymbol{x}, \boldsymbol{x}) \\ \vdots & \ddots & \vdots \\ k^{C,1}(\boldsymbol{x}, \boldsymbol{x}) & \cdots & k^{C,C}(\boldsymbol{x}, \boldsymbol{x}) \end{bmatrix} = \sum_{q=1}^{Q} \boldsymbol{B}_q \otimes k_q(\boldsymbol{x}, \boldsymbol{x}), \qquad (4)$$

- The **marginal likelihood** is represented by

$$p(\boldsymbol{y}|\boldsymbol{x}) = \int p(\boldsymbol{y}|\boldsymbol{f}, \boldsymbol{x}) p(\boldsymbol{f}|\boldsymbol{x}) d\boldsymbol{f} = \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Gamma}), \qquad \boldsymbol{\Gamma} = \boldsymbol{K} + \boldsymbol{H}^{-1} \otimes \boldsymbol{I}_N, \qquad (5)$$

SPECTRAL KERNELS

- The **Spectral Gaussian** (SG) kernel:

$$k_{\text{SG}}(\tau; \theta) = \exp(-\frac{1}{2}\nu\tau^2)\cos(\mu\tau), \quad S_{\text{SG}}(\omega; \theta) = \frac{1}{2}[\mathcal{N}(\omega; -\mu, \nu) + \mathcal{N}(\omega; \mu, \nu)] \qquad (6)$$

- The **Spectral Mixture** (SM) kernel:

$$k_{\text{SM}}(\tau; \theta) = \sum_{q=1}^{Q} a_q k_{\text{SG}}(\tau; \theta_q), \qquad S_{\text{SM}}(\omega; \theta) = \sum_{q=1}^{Q} a_q S_{\text{SG}}(\omega; \theta_q), \qquad (7)$$

THE CROSS-SPECTRAL MIXTURE KERNEL

- Applying the **SM kernel** to the **LMC framework** yields:

$$\boldsymbol{K}_{\text{SM-LMC}}(\tau; \theta) = \sum_{q=1}^{Q} \boldsymbol{B}_q k_{\text{SG}}(\tau; \theta_q), \qquad (8)$$

  - The **coregionalization matrix** allows for the cross-spectrum to be defined
  - However, the **cross-phase spectrum** always equals zero
- We propose the **Cross-Spectral Mixture** (CSM) kernel:

$$k^{c,c'}_{\text{CSM}}(\tau; \theta) = \sum_{q=1}^{Q} \sqrt{a_{cq} a_{c'q}} \exp(-\frac{1}{2}\nu_q \tau^2) \cos(\mu_q(\tau + \phi_{c'q} - \phi_{cq})), \qquad (9)$$

  - The CSM kernel is an extension of the **SM-LMC kernel** with phase information
  - In the **LMC framework**, the CSM kernel is written in phasor notation:

$$\boldsymbol{K}_{\text{CSM}}(\tau; \theta) = \text{Re}\left\{\sum_{q=1}^{Q} \boldsymbol{B}_q \widetilde{k}_{\text{SG}}(\tau; \theta_q)\right\}, \qquad \boldsymbol{B}_q = \beta_q \beta_q^{\dagger},$$

$$\widetilde{k}_{\text{SG}}(\tau; \theta_q) = \exp(-\frac{1}{2}\nu_q \tau^2 + j\mu_q \tau), \qquad \beta_{cq} = \sqrt{a_{cq}} \exp(-j\psi_{cq}),$$

  - This defines a full **cross-amplitude spectrum** *and* **cross-phase spectrum**
- Samples drawn from the CSM and SM-LMC kernels:

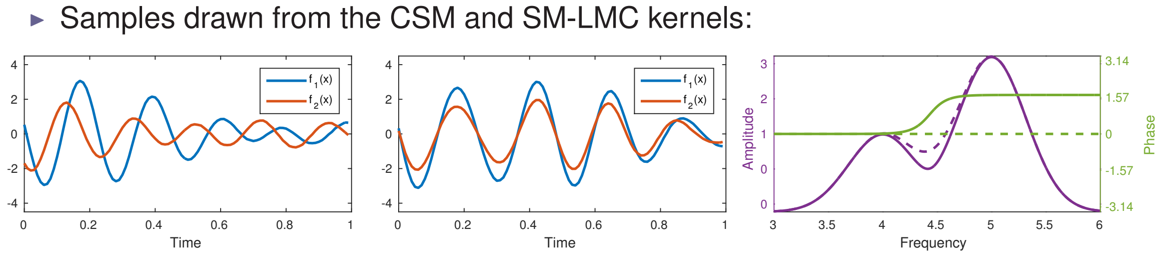

Left: Draw from the **CSM kernel** with corresponding cross-spectrum (solid, **right**)
Middle: Draw from the **SM-LMC kernel** with cross-spectrum (dashed, **right**)

APPLICATION

- Observations from C channels are segmented into W contiguous windows
- Data for window $w \in \{1, \dots, W\}$ at x_n^w are represented by $\boldsymbol{y}_n^w = [y_{n1}^w, \dots, y_{nC}^w]^T$
- Assign latent brain states according to a **Bayesian hidden Markov model**:

$$\zeta_1 \sim \text{Categorical}(\rho_0), \qquad \zeta_w \sim \text{Categorical}(\rho_{\zeta_{w-1}}) \;\; \forall w \geq 2,$$
$$\rho_0 \sim \text{Dirichlet}(\nu), \qquad \rho_\ell \sim \text{Dirichlet}(\nu),$$

- The **emission distribution** for each latent state is a Gaussian process

$$\boldsymbol{y}_n^w \sim \mathcal{N}(\boldsymbol{f}_w(x_n^w), \boldsymbol{H}_{\zeta_w}^{-1}), \qquad \boldsymbol{f}_w(x) \sim \mathcal{GP}(\boldsymbol{0}, \boldsymbol{K}(x, x'; \theta_{\zeta_w})), \qquad (10)$$

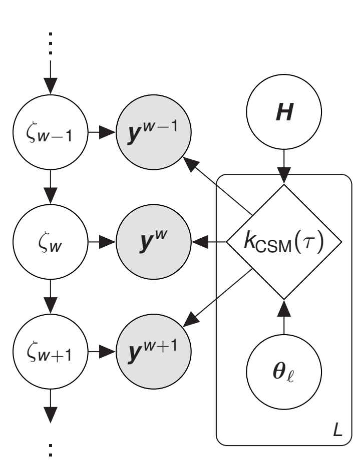

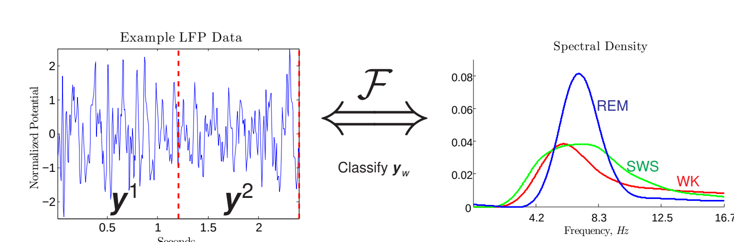

Left: Graphical model of HMM with GP emissions
Top: Example LFP data segmented into windows; each window is assigned a CSM kernel according to the cluster model

INFERENCE

- Kernel parameters: $\boldsymbol{\Theta} = \{\theta_\ell, \eta_\ell\}_{\ell=1}^L$
- Model variables: $\boldsymbol{\Omega} = \{\{\rho_\ell\}_{\ell=0}^L, \{\zeta_w\}_{w=1}^W\}$
- Model variables are inverted using **mean-field variational inference**
- Kernel parameters are chosen to maximize the **expected marginal likelihood**:

$$\mathcal{Q} = \sum_{w=1}^{W}\sum_{\ell=1}^{L} q(\zeta_w = \ell) \log \mathcal{N}(\boldsymbol{y}^w; \boldsymbol{0}, \boldsymbol{\Gamma}_\ell) \qquad (11)$$

  - $\boldsymbol{\Gamma}_\ell$ is the real component of $\widetilde{\boldsymbol{\Gamma}}_\ell = \sum_q \boldsymbol{B}_q^\ell \otimes \widetilde{k}_{\text{SG}}(\boldsymbol{x}, \boldsymbol{x}; \theta_\ell) + \boldsymbol{H}_\ell^{-1} \otimes \boldsymbol{I}_N$
  - This has naïve complexity $\mathcal{O}(N^3 C^3)$ due to the required inversion of $\boldsymbol{\Gamma}_\ell$
- Alternatively, consider the **Fourier transform** $\boldsymbol{z}^w = (\boldsymbol{I}_C \otimes \boldsymbol{U})^\dagger \boldsymbol{y}^w$
  - The equivalent marginal likelihood of the transformed data becomes:

$$\boldsymbol{z}^w \sim \mathcal{CN}(\boldsymbol{0}, 2\boldsymbol{S}_{\zeta_w}), \qquad \boldsymbol{S}_\ell = \delta^{-1} \sum_{q=1}^{Q} \boldsymbol{B}_q^\ell \otimes \boldsymbol{W}_q^\ell + \boldsymbol{H}_\ell^{-1} \otimes \boldsymbol{I}_N, \qquad (12)$$

  - $\boldsymbol{W}_q^\ell \approx \text{diag}([S_{\text{SG}}(\omega; \theta_{\ell q}), \boldsymbol{0}])$ is approximately diagonal

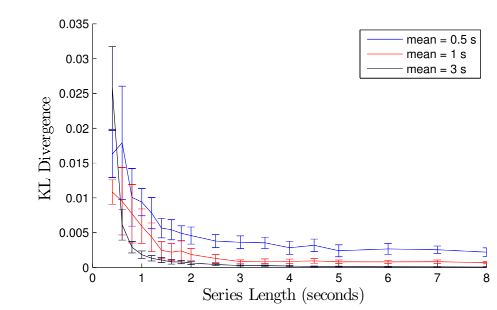

- The quality of this approximation depends on sequence length $N \cdot \delta$ and kernel parameters θ_q
- **Left**: A visualization of the **KL divergence** of the fitted marginal likelihood from the true marginal likelihood

RESULTS

Sleep data was collected as follows:

- *Three mice* were observed naturally transitioning through different levels of sleep
- *Twelve hours* of LFP data was recorded from *sixteen different brain regions*
- We present only **three hours** of sleep data from **four brain regions** of a **single mouse** for clarity

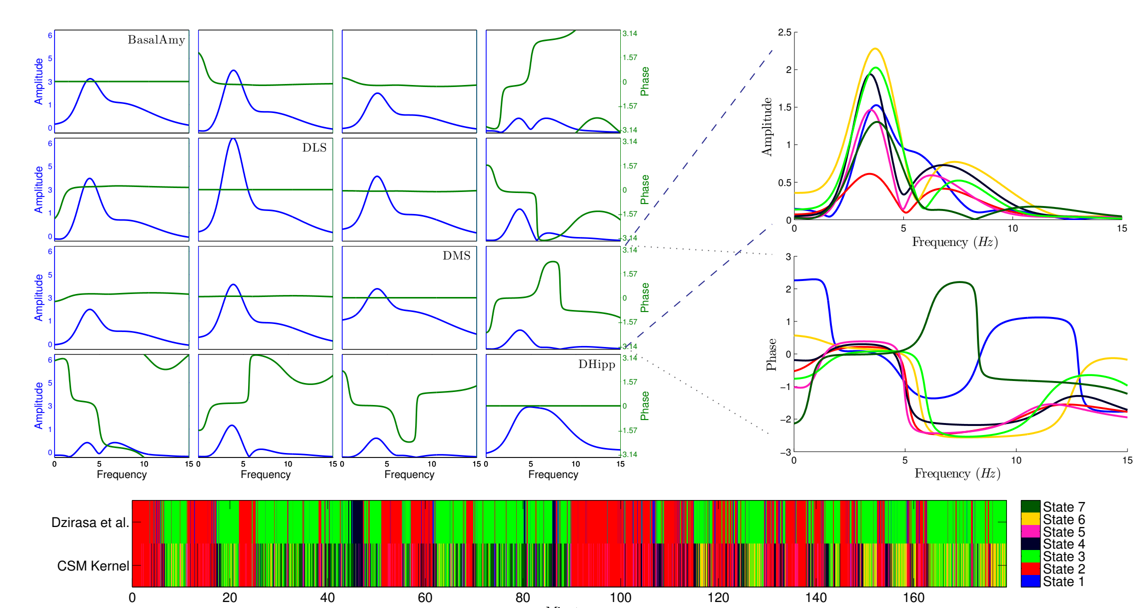

Top left: Full cross-spectrum of observations from the Basal Amygdala, Dorsomedial Striatum, Dorsolateral Striatum, and Hippocampus for a single state
Top right: All cross-amplitude and cross-phase spectra between DMS and DHipp
Bottom: Evolution of brain states according to HMM assignments

Presenter: Kyle Ulrich

Mail: kyle.ulrich@duke.edu