

MIT Open Access Articles

This is a supplemental file for an item in DSpace@MIT

Item title: A state-space model for
dynamic functional connectivity

Link back to the item: <https://hdl.handle.net/1721.1/126065>



A state-space model for dynamic functional connectivity

Sourish Chakravarty^{1,5,6}, Zachary D. Threlkeld⁷, Yelena G. Bodien⁶, Brian L. Edlow⁶, Emery N. Brown^{1–5}

Abstract—Dynamic functional connectivity (DFC) analysis involves measuring correlated neural activity over time across multiple brain regions. Significant regional correlations among neural signals, such as those obtained from resting-state functional magnetic resonance imaging (fMRI), may represent neural circuits associated with rest. The conventional approach of estimating the correlation dynamics as a sequence of static correlations from sliding time-windows has statistical limitations. To address this issue, we propose a multivariate stochastic volatility model for estimating DFC inspired by recent work in econometrics research. This model assumes a state-space framework where the correlation dynamics of a multivariate normal observation sequence is governed by a positive-definite matrix-variate latent process. Using this statistical model within a sequential Bayesian estimation framework, we use blood oxygenation level dependent activity from multiple brain regions to estimate posterior distributions on the correlation trajectory. We demonstrate the utility of this DFC estimation framework by analyzing its performance on simulated data, and by estimating correlation dynamics in resting state fMRI data from a patient with a disorder of consciousness (DoC). Our work advances the state-of-the-art in DFC analysis and its principled use in DoC biomarker exploration.

I. INTRODUCTION

Functional connectivity (FC) analysis can be broadly described as an analysis of “statistical dependencies among remote neuro-physiological events” [Fri11]. The FC measures are considered to be representative of long-range neural coordination across brain-regions. Conventionally, they are calculated as pair-wise correlation in neural signals, e.g. blood oxygen level dependent (BOLD) activation in functional magnetic resonance imaging (fMRI) data and or electrical activity in electroencephalography (EEG) data, across multiple brain regions. From a signal processing perspective, FC analysis is based upon estimation of correlation matrices

from multivariate time-series data. FC analysis is widely relevant in basic neuroscience [AM14] as well as translational clinical neuroscience research ([BCE17], [ECS⁺17]).

An emerging paradigm in FC research is that of dynamic functional connectivity (DFC) [TL15]. Conventionally, FC analysis is performed on an entire session of multi-site/multi-channel neural recordings to generate a correlation map. A key assumption of such an analysis is that the true correlation remains constant across the entire sequence of a recording session. This static FC (SFC) interpretation can be generalized to DFC by assuming that the true correlation map changes over time. Therefore, DFC analysis involves estimating a time-series of pairwise correlations from multivariate neural data. The DFC analysis yields information that is complementary to SFC analysis of the same dataset: SFC indicates the degree to which any two brain-regions are synchronous with each other on average, whereas DFC indicates how this degree of synchrony varies over time. Principled DFC analysis may aid translational research aimed at biomarker discovery for neurological disorders such as disorders of consciousness (DoC) [DTD⁺19].

Conventionally, DFC analysis generates dynamic estimates of correlations by assuming a sliding-window (SW) approach. In this approach, a sequence of instantaneous correlation measures is calculated by repeating the SFC analyses on a sequence of short data windows where the window-size is user-prescribed. A sub-optimal choice of window-size can lead to statistical challenges. For example, if the window size is too long, then the estimated correlation trajectory may be unable to capture faster dynamics. Conversely, if window size is too short then the estimates will be more sensitive to noise in the data. Therefore, it is preferable to use alternative DFC estimation techniques that can yield reliable estimates of correlation dynamics without requiring subjective windows. One such promising alternative has been proposed by Lindquist et al.[LXNC14] where the authors estimated correlation dynamics by fitting a dynamic conditional correlation model (DCC) to resting-state fMRI data. The DCC model originated in econometrics literature to estimate covariance trajectories in multivariate financial time-series data [ES01]. Motivated by this work, here we propose a novel approach based on state-space models (SSMs) to estimate DFC from neural data while avoiding data-windowing. In particular, we consider a specific class of SSMs known as Multivariate Stochastic Volatility (MVSF) models that also originated in econometrics research [CST18]. Such an MVSF model incorporates a matrix-variate latent state process to describe correlation dynamics in the observed neural signals. DFC statistics may then be derived by estimating the latent states and model

This work was partially supported by NIH Award P01-GM118629 (to E.N.B.), by funds from Massachusetts General Hospital (to E.N.B.), by funds from the Picower Institute for Learning and Memory (to E.N.B. and S.C.), by funds from the NIH/NINDS (DP2-HD101400, R21-NS109627, RF1-NS115268, K23-NS094538; to B.L.E.), by funds from James S. McDonnell Foundation (to B.L.E.), and by funds from the National Institute on Disability, Independent Living and Rehabilitation Research, Administration for Community Living (90DP0039, Spaulding- Harvard TBI Model System; to Y.G.B.).

¹Picower Institute for Learning and Memory, Massachusetts Institute of Technology (MIT), Cambridge, MA

²Department of Brain and Cognitive Sciences, MIT, Cambridge, MA

³Institute of Medical Engineering and Science, MIT, Cambridge, MA

⁴Harvard-MIT Division of Health Science and Technology

⁵Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital (MGH), Boston, MA

⁶Center for Neurotechnology and Neurorecovery, Department of Neurology, MGH, Boston, MA

⁷Dept. of Neurology and Neurological Sciences, Stanford University School of Medicine, Stanford, CA

parameters from the data. This SSM framework provides more flexibility in describing complex correlation dynamics [CG16], and in estimating instantaneous correlations even under scenarios of missing data and non-uniform sampling time-intervals of observations [SS82].

In Sec. II, we present our SSM, which is a modified version of the existing MVSV model [CST18], for estimating correlation dynamics in multivariate noisy neural data. Specifically, we do not consider certain modeling features of [CST18] such as a first-order vector autoregressive process to track the mean of observations, and additional latent states to track both gradual changes in channel-wise variance and temporally sharp regime transitions in mean and covariance dynamics of the signals. Except for minor differences in notational conventions and a major difference in the proposal distributions in Secs. II-F and II-G, the statistical model and estimation procedure used here are identical to [CST18] conditioned on the assumptions made in this work. In Sec. III, we numerically verify our model on synthetic data and demonstrate its utility in estimating time-varying correlations from BOLD-fMRI data from a DoC patient. Finally, in Sec. IV, we summarize the current work, and discuss potential next steps.

II. THEORY

A. State-Space Model (SSM)

Consider a sequence of multivariate neural data $\{y_k\}_{k=1}^K$. At the k -th instant, the observation vector, $y_k \in \mathbb{R}^m$, is assumed to be a realization of a multivariate normal distribution characterized by zero mean and covariance matrix Σ_k . This is represented as,

$$y_k \sim \mathcal{N}_m(0_m, \Sigma_k) \quad (1)$$

where, 0_m denotes a $m \times 1$ matrix of zeros. The covariance matrix, Σ_k , is decomposed as,

$$\Sigma_k = \Lambda_k \Omega_k \Lambda_k \quad (2)$$

where, Λ_k^2 is a diagonal matrix of variances and Ω_k denotes the current correlation matrix. The stochastic evolution of the sequence, $\{\Omega_k\}_{k=1}^K$, is assumed to be governed by a latent state sequence, $\{Q_k\}_{k=1}^K$ where Q_k is a real-valued positive-definite random matrix. The dynamics of this latent positive-definite process is described by,

$$Q_k^{-1} = \frac{1}{\nu} Q_{k-1}^{-d/2} E_k Q_{k-1}^{-d/2} \text{ such that } E_k \sim \mathcal{W}_m(\nu, I_m), \quad (3)$$

$$\Omega_k = \tilde{Q}_k^{-1} Q_k \tilde{Q}_k^{-1} \quad (4)$$

where, $\tilde{Q}_{k,ij} = Q_{k,ij}^{1/2} \delta_{ij}$, δ_{ij} denotes the Kronecker delta function, $\nu > m$, and $-1 \leq d \leq 1$. Here, we use $X \sim \mathcal{W}_m(\nu, S)$ to denote a real-valued, positive-definite, $m \times m$ random matrix X distributed according to a Wishart

probability density function (pdf) described as,

$$(2^{\frac{\nu m}{2}} \Gamma_m(\nu/2) |S|^{\frac{\nu}{2}})^{-1} |X|^{\frac{\nu-m-1}{2}} e^{\text{Tr}\left(-\frac{1}{2} S^{-1} X\right)}, \quad (5)$$

$$\Gamma_m(\nu/2) = \pi^{m(m-1)/4} \prod_{i=1}^m \Gamma((\nu - i + 1)/2). \quad (6)$$

and characterized by parameters S (a positive-definite matrix) and scalar $\nu (> m)$ [GN09]. In this work, we use $|\cdot|$ and $e^{\text{Tr}(\cdot)}$ to denote, respectively, determinant and exponential-of-trace operations on a square matrix. Note that Eq. (3) ensures that each point on the dynamic trajectory of Q_k will be positive-definite, and $Q_k^{-1} \sim \mathcal{W}_m(\nu, Q_{k-1}^{-d}/\nu)$ [GN09]. The matrix scaling operation in Eq. (4) preserves the unit-diagonal character of the resulting correlation matrices, Ω_k . Further simplifying our theoretical modeling in Eq. (3), we do not consider additional multiplicative coefficient matrices that can be used in relative scaling of past information (Q_{k-1}) [CST18, Eqs.8, 9], and assume $\Lambda_k = I_m$, where I_m denotes a $m \times m$ identity matrix.

B. Likelihood and model parameters

The total data log-likelihood can be written out as,

$$p(y_{1:K}, Q_{1:K} | \Theta) = p(Q_0) \prod_{k=1}^K p(y_k | Q_k) p(Q_k | Q_{k-1}, \Theta) \quad (7)$$

where, $p(X)$ denotes a pdf described on the appropriate space to which X belongs, and $\Theta = \{\nu, d\}$. We assume $Q_0 = I_m$ and set $p(Q_0) = 1$ [CST18].

C. Bayesian inference & Markov Chain Monte Carlo steps

For a given sequence of observations, $y_{1:K} \equiv \{y_k\}_{k=1}^K$, we employ a Bayesian inference framework to estimate posterior distributions on the DFC statistics, $Q_{1:K}^{-1} \equiv \{Q_k^{-1}\}_{k=1}^K$ and Θ . Using the model (Eqs. (1), (2), (3), (4)), the unit variance assumption, Bayes' rule and prior probability distributions $p(\nu)$ and $p(d)$, we can write the following posterior pdf,

$$p(Q_{1:K}^{-1}, \nu, d | y_{1:K}) \propto p(y_{1:K}, Q_{1:K}^{-1} | \nu, d) p(\nu) p(d), \quad (8)$$

where,

$$p(\nu) = \frac{\beta_\nu^{\alpha_\nu}}{\alpha_\nu} (\nu - m)^{\alpha_\nu - 1} \exp(-\beta_\nu(\nu - m)) \mathbb{I}_{(0, \infty)}(\nu - m), \quad (9)$$

$$p(d) = (1/2) \mathbb{I}_{[-1, 1]}(d), \quad (10)$$

where, (α_ν, β_ν) denote the user-prescribed hyper-parameters for the prior pdf on $(\nu - m)$ distributed according to a Gamma probability distribution denoted as $\text{Gamma}(\alpha_\nu, \beta_\nu)$. $\mathbb{I}_{(\mathcal{X})}(x)$ denotes an indicator function that takes value = 1 $\forall x \in \mathcal{X}$ and = 0, otherwise.

To sample from the joint posterior pdf (Eq. (8)) per [CST18], we use a Gibb's sampling approach [CG95] where in the j -th iteration we draw the j -th samples from the following conditional posterior pdf's,

$$Q_{k,(j)}^{-1} \sim p(Q_{k-1}^{-1} | Q_{k-1,(j)}^{-1}, Q_{k+1,(j-1)}^{-1}, \nu_{(j-1)}, d_{(j-1)}, y_{1:K}) \quad (11)$$

$$Q_{K,(j)}^{-1} \sim p(Q_{K,(j)}^{-1} | Q_{K-1,(j)}^{-1}, \nu_{(j-1)}, d_{(j-1)}, y_{1:K}) \quad (12)$$

$$\nu_{(j)} \sim p(\nu | Q_{1:K,(j)}^{-1}, d_{(j-1)}, y_{1:K}) = p(\nu | Q_{1:K,(j)}^{-1}, d_{(j-1)}) \quad (13)$$

$$d_{(j)} \sim p(d | Q_{1:K,(j)}^{-1}, \nu_{(j)}, y_{1:K}) = p(d | Q_{1:K,(j)}^{-1}, \nu_{(j)}) \quad (14)$$

where, the equalities in Eqs. (13) and (14) are due to the conditional independence assumption in the proposed SSM. To sample $x^{(j)} \sim p(x)$ (where, $p(x)$ may denote any of the $K + 2$ conditional pdf's in Eqs. (11)-(14)), given the last sample $x^{(j-1)}$, we perform the following general steps of MCMC algorithm [CG95]: we draw a candidate sample x^* from a relatively easy-to-sample *proposal* distribution whose pdf, conditioned on $x^{(j-1)}$, is denoted by $q(x^* | x^{(j-1)})$ and accept it as a new sample $x^{(j)}$ with an acceptance probability given by,

$$\min \left(1, \frac{g(x^*)q(x^{(j-1)} | x^*)}{g(x^{(j-1)})q(x^* | x^{(j-1)})} \right) \quad (15)$$

where, $p(x) \propto g(x)$. If we are unable to accept x^* , then we set $x^{(j)} = x^{(j-1)}$. In the following sections, Secs. II-D - II-G, we provide the details of our choice of $g(\cdot)$ and $q(\cdot | \cdot)$ that will be used in Eq. 15 for each of the conditional posterior pdfs, Eqs. (11)-(14). Note that with the expressions available to calculate $g(x^*)$ and $q(x^* | x^{(j-1)})$, we can also calculate $g(x^{(j-1)})$ and $q(x^{(j-1)} | x^*)$ by exchanging $x^{(j-1)}$ and x^* in the arguments.

D. Sampling Q_k^{-1} from conditional posterior pdf

The conditional posterior in Eq. (11) can be evaluated up to an unknown constant factor as,

$$p(y_k | Q_k^{-1}, \Theta) p(Q_{k+1}^{-1} | Q_k^{-1}, \Theta) p(Q_k^{-1} | Q_{k-1}^{-1}, \Theta) \propto g(Q_k^{-1}) \quad (16)$$

$$= |Q_k^{-1}|^{(\nu+1-m-1)/2} e^{\text{Tr} \left(-\frac{\nu}{2} S_k^{-1} Q_k^{-1} \right)} |Q_k^{-1}|^{(-d\nu)/2} \cdot |\tilde{Q}_k| e^{\text{Tr} \left(-\frac{1}{2} ((\tilde{Q}_k \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q}_k) Q_k^{-1} + \nu S_{k+1}^{-1} Q_{k+1}^{-1}) \right)} \quad (17)$$

$$= |Q_k^{-1}|^{(\nu+1-m-1)/2} \cdot e^{\text{Tr} \left(-\frac{1}{2} (\nu S_k^{-1} + \tilde{Q} \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q}) Q_k^{-1} \right)} |Q_k^{-1}|^{(-d\nu)/2} \cdot |\tilde{Q}_k| \frac{e^{\text{Tr} \left(-\frac{1}{2} ((\tilde{Q}_k \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q}_k) Q_k^{-1} + \nu S_{k+1}^{-1} Q_{k+1}^{-1}) \right)}}{e^{\text{Tr} \left(-\frac{1}{2} (\tilde{Q} \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q}) Q_k^{-1} \right)}} \quad (18)$$

where, $S_k \equiv Q_{k-1}^{-d}$. Note that Eq. (18) is obtained by multiplying and dividing Eq. (17) by $e^{\text{Tr} \left(-\frac{1}{2} (\tilde{Q} \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q}) Q_k^{-1} \right)}$ where $\tilde{Q} = (\tilde{Q}_{k-1} + \tilde{Q}_{k+1})/2$. Per [CST18], to sample a candidate realization, $Q_{k,*}^{-1}$ we use as proposal pdf, $q(Q_{k,*}^{-1} | Q_{k,(j-1)}^{-1})$, the pdf of $\mathcal{W}_m(\nu + 1, (\nu S_k^{-1} + \tilde{Q} \Lambda_k^{-1} y_k y_k^T \Lambda_k^{-1} \tilde{Q})^{-1})$.

E. Sampling Q_K^{-1} from conditional posterior pdf

The conditional posterior in Eq. (12) can be expressed upto an unknown constant as,

$$g(Q_K^{-1}) = |S_K^{-1}|^{\nu/2} |Q_K^{-1}|^{(\nu+1-m-1)/2} e^{\text{Tr} \left(-\frac{\nu}{2} S_K^{-1} Q_K^{-1} \right)} \cdot |\tilde{Q}_K| e^{\text{Tr} \left(-\frac{1}{2} (\tilde{Q}_K \Lambda_K^{-1} y_K y_K^T \Lambda_K^{-1} \tilde{Q}_K) \right)} Q_K^{-1} \quad (19)$$

. Per [CST18], to sample any candidate $Q_{K,*}$ we use as proposal pdf, $q(Q_{K,*} | Q_{K,(j-1)})$, the pdf corresponding to $\mathcal{W}_m(\nu + 1, S_K/\nu)$.

F. Sampling ν from conditional posterior pdf

The conditional posterior distribution in Eq. (13) can be written out up to a constant multiplying factor as $g(\nu)$, such that

$$\begin{aligned} \ln g(\nu) = & \mathbb{I}_{(m,\infty)} (\ln(\nu - m)^{\alpha_\nu - 1} - K \ln \Gamma_m(\nu/2) \\ & + (m\nu K/2) \ln \nu - (\nu/2)(2\beta_\nu + mK \ln 2 \\ & + \sum_{k=1}^K (\ln |S_k| - \ln |Q_k^{-1}| + \text{Tr}(S_k^{-1} Q_k^{-1}))) \end{aligned} \quad (20)$$

To generate a candidate sample ν^* we use as proposal pdf, $q(\nu^* | \nu^{(j-1)})$, the pdf corresponding to a shifted Gamma distribution denoted by $\text{Gamma}(\alpha_p, \beta_p)$. The characteristic parameters (α_p, β_p) are functions of a prescribed mode $\nu_M = \nu^{(j-1)}$ and user-prescribed constant variance ν_{var} , such that

$$\beta_p(\nu_M, \nu_{var}) = \frac{((\nu_M - m) + ((\nu_M - m)^2 + 4\nu_{var})^{1/2})}{2\nu_{var}} \quad (21)$$

$$\alpha_p(\nu_M, \nu_{var}) = 1 + (\nu_M - m)\beta_p(\nu_M) \quad (22)$$

Here, the proposal distribution, $q(\nu^* | \nu^{(j-1)})$, is designed such that the mode of the distribution is at last sample $\nu^{(j-1)}$, and the bounds on ν are respected. To calculate $q(\nu^{(j-1)} | \nu^*)$, we use a different shifted Gamma pdf whose parameters are calculated using Eqs. (21) and (22) with $\nu_M = \nu^*$.

G. Sampling d from conditional posterior pdf

The posterior pdf in Eq. (14) can be written out upto an unknown constant factor as,

$$\begin{aligned} g(d) = & \mathbb{I}_{[-1,1]} \exp \left(\sum_{k=1}^K (-d\nu/2) \ln |Q_{k-1}^{-1}| + \right. \\ & \left. (-\nu/2) \text{Tr}(Q_{k-1}^d Q_k^{-1}) \right) \end{aligned} \quad (23)$$

To sample a candidate d^* we use a proposal density $q(d^* | d^{(j-1)})$ which corresponds to $(1/2) \times$ the pdf of $\text{Beta}(a_p, b_p)$. For a given mean, $d_A = d^{(j-1)}$, constraint $b_p = 1/a_p$, and a user-prescribed constant threshold parameter a_f such that $1/a_f \leq a_p \leq a_f$ where $a_f > 1$, the characteristic parameter is given by,

$$\begin{aligned} a_p(d_A, a_f) = & \max \left(\min \left(\left(\frac{(1 + d_A)/2}{1 - (1 + d_A)/2} \right)^{1/2}, a_f \right), 1/a_f \right) \end{aligned} \quad (24)$$

The proposal distribution is designed such that its mean corresponds to the previous sample, while also having more probability mass around this mean and respecting the bounds on d . To determine $q(d^{(j-1)}|d^*)$, we use $Beta(a_p(d^*, a_f), 1/a_p(d^*, a_f))$.

III. RESULTS

A. Initialization of MCMC simulation

In all the DFC analyses performed here on synthetic and fMRI datasets, we consider bivariate observation sequences ($m = 2$), $K = 150$ and maximum MCMC iterations of $N_{MCMC} = 10000$. For the very first MCMC iteration, we initialize $Q_{k,(j=0)}^{-1} = I_m, \forall k = 1, \dots, K$. The hyperparameters for the prior on ν , we choose $\alpha_\nu = m + 2$, and $\beta_\nu = 1$, and initialize $\nu_{(j=0)} = m + (\alpha_\nu - 1)/\beta_\nu$. Also, we set $d_{(j=0)} = 0.5$, $\nu_{var} = 0.1$ and $a_f = 5$. With regard to post-processing the MCMC generated samples, we set the burn-in iterations for latent states and model parameters as 1000 and 4000, respectively. Also, we set the thinning interval for latent state and model parameters as 100 and 200, respectively. The numbers for burn-in and thinning interval are chosen based on visual inspection of the convergence in the respective Markov chains and the auto-correlation (after ignoring the samples from the burn-in iterations) in each MCMC chain.

B. DFC analyses of simulated data

Using the generative model (Eqs.(1), (3), (4), (2)) and setting $\nu = 5$, $d = 0.8$, $\Lambda_k = I_2 \forall k = 1, \dots, K$ we generate a single realization of bivariate observation $y_{1:K}$ (Fig. 1(a)). On this data, we run our MVSF model-based estimation procedure that is agnostic to the ground-truth model parameters. The ordered statistics from samples estimated from the posterior distribution of correlation (Fig. 1(b)), and the empirical histograms on ν and d (Fig. 1(c,d)) are reported. Note, in Fig. 1(b), the 2.5-th and 97.5-th percentiles include not only the true value, but also track the changes in true correlation trajectory. Comparing with the raw data, Fig. 1(a), we observe that where the estimated 2.5th percentile is above 0, the two time-series appear to be in synchrony (high positive correlation) with each other. On the other hand, when the estimated 97.5th percentile is below 0, the two signals are synchronous, but in the opposite sense (high negative correlation). Therefore, the posterior estimates generated by our DFC analyses enable us to infer whether the correlation is significant at any time instant. In Figs. 1(c) and 1(d), the posterior samples of the model parameters are indeed in the near-neighborhood of the ground truth. Also, note that the posterior estimates on ν and d use only $K = 150$ data-points. With longer sessions, these estimates can be expected to be more accurate.

C. DFC analyses of BOLD fMRI data from a patient

We demonstrate the utility of the DFC analyses pipeline, numerically demonstrated in the previous section, on 2 recording sessions of BOLD-fMRI data. These sessions were

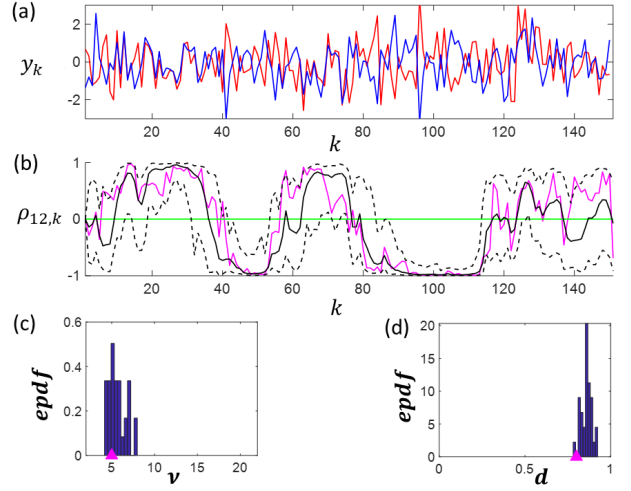


Fig. 1. MVSF-based estimation on simulated data. (a) Simulated observations, (b) True correlation (magenta) and ordered statistics on estimated posterior samples (2.5-th, 97.5-th percentiles indicated as dashed black line, while median indicated as solid black line), green line indicates MCMC initialization state. (c) and (d) present estimated empirical pdf's (epdf's) and true value (magenta triangle) on ν and d , respectively.

collected from a single DoC patient who underwent a resting-state fMRI scan (time to repeat = 2.4 seconds) while in post traumatic confusional state (PTCS; pre-recovery), and 6 months later after full recovery of consciousness (post-recovery). Details of the clinical trial, recording and data pre-processing procedures can be found in [TBR⁺18]¹. For the purpose of this proof-of-concept study, we use the denoised and artifact-rejected signal that is averaged within 2 regions-of-interest (ROI's) located in the medial pre-frontal cortex (MPFC) and in the posterior cingulate cortex (PCC). The MPFC and PCC are nodes of the default mode network, and in a healthy subject, these ROI's are known to have positively correlated BOLD activity [BCE17]. The time-series dataset from each of these 2 ROI's is standardized by shifting the mean to zero followed by dividing by the standard deviation, where both the mean and standard deviations are calculated across the entire session. The resulting standardized bivariate sequence constitutes the observations, $y_{1:K}$, to be analyzed.

For each session, using our DFC analysis framework we estimate correlation trajectories (Fig. 2) from the entire data sequence without application of a SW approach. The estimated correlation trajectory demonstrates how the degree of synchrony between the BOLD signals from MPFC and PCC gradually evolve over time. The inference of such smooth correlation trajectories is enabled by the MVSF model which incorporates the influence of neighboring time points on the current latent correlation state (Eq. (3)). Furthermore, this DFC analysis provides session-specific estimates of ν and d , which may be regarded as reduced-order descriptors of the correlation dynamics. For each of these sessions, we note that there are brief epochs where the 95% confidence interval trajectory of the estimated correlations does not

¹Denoising, artifact rejection and roi-based averaging steps are performed within CONN Matlab toolbox [WGNC12]

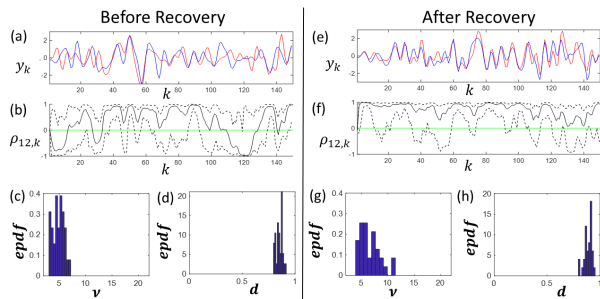


Fig. 2. DFC analyses of BOLD fMRI data from a PTCS patient, before and after recovery. (a, e) illustrate the standardized bivariate data from MPFC and PCC, (b, f) present the ordered statistics (2.5-th, and 97.5-th percentiles as dashed black lines, median as solid black line) calculated from the estimated samples of instantaneous correlation. Figure panels (c, g) and (d, h) present empirical pdf's (epdf's) on ν and d , respectively.

include 0. Therefore, we infer that for these 2 sessions of fMRI data, the correlation is indeed time-varying. From Fig. 2, we also observe that the correlation trajectory prior to recovery includes both positive and negative correlation values, but the trajectory after recovery remains mostly above zero. Thus, the MVSF model and the corresponding DFC statistics (latent states and model parameters), together provide a principled lower-dimensional objective vocabulary to describe the complex correlation dynamics in multivariate BOLD signals. In the absence of ground truth, the generative nature of the MVSF model allowed us to analyze the performance of the inference framework against synthetically generated ground truth (see Sec. III-B).

IV. CONCLUSION

In summary, we modified an MVSF model from the econometrics literature [CST18] to develop a novel DFC analysis approach for multivariate neural data. We used this approach to perform exploratory DFC analyses of bivariate resting-state BOLD fMRI data from a patient with DoC. We introduced principled modifications to the proposal densities within the MCMC simulations for estimating two key DFC model parameters. We numerically verified the estimation framework on synthetic data. We then deployed the numerically verified DFC analysis pipeline on two fMRI recordings from the patient, before and after recovery of consciousness. Our analysis of resting-state BOLD activity in this DoC patient indicates that the correlation between MPFC and PCC varied within the duration of the recording and correlation dynamics showed marked change when the patient recovered.

Our MVSF model-based DFC analysis pipeline does not require sliding windows, and it can estimate gradually evolving correlation trajectories as well as reduced-order statistical descriptors of correlation dynamics. Furthermore, the estimation framework results in probability distributions, learnt from both prior knowledge and available data, that quantify the uncertainty in both the estimated time-varying correlations and the model parameters. Such probability distributions can be used to compare correlations across

multiple time-points in the same session, and to infer dynamic connectivity graphs. A next step will be to extend our bivariate treatment to higher dimensions. Furthermore, this work suggests that exploring SSMs with matrix-variate latent processes for principled DFC analysis of neural data may aid DFC-based bio-marker discovery.

REFERENCES

- [AM14] Evan G Antzoulatos and Earl K Miller. Increases in functional connectivity between prefrontal cortex and striatum during category learning. *Neuron*, 83(1):216–225, 2014.
- [BCE17] Yelena G Bodien, Camille Chatelle, and Brian L Edlow. Functional networks in disorders of consciousness. In *Seminars in neurology*, volume 37, pages 485–502. Thieme Medical Publishers, 2017.
- [CG95] Siddhartha Chib and Edward Greenberg. Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335, 1995.
- [CG16] Joshua CC Chan and Angelia L Grant. Modeling energy price dynamics: GARCH versus stochastic volatility. *Energy Economics*, 54:182–189, 2016.
- [CST18] Roberto Casarin, Domenico Sartore, and Marco Tronzano. A Bayesian Markov-switching correlation model for contagion analysis on exchange rate markets. *Journal of Business & Economic Statistics*, 36(1):101–114, 2018.
- [DTD⁺19] Athena Demertzi, Enzo Tagliazucchi, Stanislas Dehaene, Gustavo Deco, Pablo Barttfeld, Federico Raimondo, Charlotte Martial, Davinia Fernández-Espejo, Benjamin Rohaut, HU Voss, et al. Human consciousness is supported by dynamic complex patterns of brain signal coordination. *Science advances*, 5(2):eaat7603, 2019.
- [ECS⁺17] Brian L Edlow, Camille Chatelle, Camille A. Spencer, Catherine J. Chu, Yelena G. Bodien, Kathryn L. O'Connor, Ronald E. Hirschberg, Leigh R. Hochberg, Joseph T. Giacino, Eric S. Rosenthal, and Ona Wu. Early detection of consciousness in patients with acute severe traumatic brain injury. *Brain*, 140(9):2399–2414, 07 2017.
- [ES01] Robert F Engle and Kevin Sheppard. Theoretical and empirical properties of dynamic conditional correlation multivariate garch. Technical report, National Bureau of Economic Research, 2001.
- [Fri11] Karl J Friston. Functional and effective connectivity: a review. *Brain connectivity*, 1(1):13–36, 2011.
- [GN09] Arjun K Gupta and Daya K Nagar. *Matrix variate distributions*. Chapman and Hall/CRC, 2009.
- [LXNC14] Martin A Lindquist, Yuting Xu, Mary Beth Nebel, and Brain S Caffo. Evaluating dynamic bivariate correlations in resting-state fmri: a comparison study and a new approach. *NeuroImage*, 101:531–546, 2014.
- [SS82] Robert H Shumway and David S Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of time series analysis*, 3(4):253–264, 1982.
- [TBR⁺18] Zachary D Threlkeld, Yelena G Bodien, Eric S Rosenthal, Joseph T Giacino, Alfonso Nieto-Castanon, Ona Wu, Susan Whitfield-Gabrieli, and Brian L Edlow. Functional networks reemerge during recovery of consciousness after acute severe traumatic brain injury. *Cortex*, 106:299–308, 2018.
- [TL15] Enzo Tagliazucchi and Helmut Laufs. Multimodal imaging of dynamic functional connectivity. *Frontiers in Neurology*, 6:10, 2015.
- [WGNC12] Susan Whitfield-Gabrieli and Alfonso Nieto-Castanon. Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain connectivity*, 2(3):125–141, 2012.