

Interactions Between Large-Scale Functional Brain Networks are Captured by Sparse Coupled HMMs

Thomas A. W. Bolton¹, *Student Member, IEEE*, Anjali Tarun, Virginie Sterpenich, Sophie Schwartz, and Dimitri Van De Ville, *Senior Member, IEEE*

Abstract—Functional magnetic resonance imaging (fMRI) provides a window on the human brain at work. Spontaneous brain activity measured during resting-state has already provided many insights into brain function. In particular, recent interest in dynamic interactions between brain regions has increased the need for more advanced modeling tools. Here, we deploy a recent fMRI deconvolution technique to express resting-state temporal fluctuations as a combination of large-scale functional network activity profiles. Then, building upon a novel sparse coupled hidden Markov model (SCHMM) framework, we parameterised their temporal evolution as a mix between intrinsic dynamics, and a restricted set of cross-network modulatory couplings extracted in data-driven manner. We demonstrate and validate the method on simulated data, for which we observed that the SCHMM could accurately estimate network dynamics, revealing more precise insights about direct network-to-network modulatory influences than with conventional correlational methods. On experimental resting-state fMRI data, we unraveled a set of reproducible cross-network couplings across two independent datasets. Our framework opens new perspectives for capturing complex temporal dynamics and their changes in health and disease.

Index Terms—Dynamic functional connectivity, total activation, innovation-driven co-activation patterns, sparse coupled hidden Markov model, ℓ_1 regularisation.

I. INTRODUCTION

SPONTANEOUS brain activity can be measured non-invasively in human volunteers using resting-state (RS) functional magnetic resonance imaging (fMRI). The study of RS functional connectivity (FC)—statistical interdependencies between brain regions’ activity traces—has

shown great potential in refining our understanding of human cognition [1] and in increasing our knowledge of the alterations caused by brain disease and disorder [2], [3]. This was in large part contributed by the discovery of a repertoire of large-scale functional networks [4]–[6], commonly termed resting-state networks (RSNs).

Conventional FC is measured as pairwise correlation between complete time courses of several minutes. Multiple time samples are thus required to generate a FC estimate, leading to an intrinsically assumed “stationarity” of functional relationships. This view has, however, recently been challenged by new methods that characterise time-dependent FC changes reflecting moment-to-moment reorganization of functional networks [7]. Therefore, *dynamic functional connectivity* (dFC) has emerged as a new research direction with important methodological developments.

In the majority of existing dFC works, changes in regional interactions are tracked through successive connectivity measures on overlapping temporal windows (*sliding window framework*). To extract an informative, limited subset of connectivity states, hard clustering [8]–[10] or subspace decomposition methods [11]–[13] are subsequently applied. Other attempts rely on frame-wise dFC analysis, without resorting to second-order connectivity measurements; amongst the main such efforts, temporal independent component analysis [14] and co-activation pattern analysis [15], [16] both stand out by their ability to reveal meaningful maps of brain activity. For more details on the wide landscape of available dFC tools, the reader is referred to [17] and [18].

An alternative approach is to incorporate knowledge about the hemodynamic response to better take into account neurological meaningful signal. For instance, Caballero Gaudes *et al.* [19] have proposed to retrieve event-related responses without timing information using regularisation strategies on deconvolved blood oxygenation level-dependent (BOLD) time courses. Along the same methodological line, deconvolution can also be performed based on a generalized total variation regularisation criterion, termed *total activation* (TA) [20]. For each brain voxel, this processing pipeline provides not only the activity-inducing (deconvolved) signal, but also the innovation (deconvolved and differentiated) signal. The latter can then be used to mark key frames of transient brain activity for temporal clustering to extract

Manuscript received June 19, 2017; accepted September 12, 2017. Date of publication September 21, 2017; date of current version December 29, 2017. This work was supported in part by the Swiss National Science Foundation under Grant 205321_163376, in part by the Center for Biomedical Imaging of the Geneva/Lausanne Universities and EPFL, in part by the Leenaards Foundation, in part by the Louis-Jeantet Foundation, and in part by the Bertarelli Foundation. (*Corresponding author: Thomas A. W. Bolton.*)

T. A. W. Bolton, A. Tarun, and D. Van De Ville are with the Center for Neuroprosthetics, Institute of Bioengineering, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland, and also with the Department of Radiology and Medical Informatics, University of Geneva, 1211 Geneva, Switzerland (e-mail: thomas.bolton@epfl.ch).

V. Sterpenich and S. Schwartz are with the Department of Neuroscience, University of Geneva, 1211 Geneva, Switzerland.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2017.2755369

innovation-driven co-activation patterns (iCAPs) [21], dynamically retrieved RSNs that can overlap in their spatial pattern of activity.

By mapping back iCAPs onto the voxelwise activity-inducing signals, a temporal expression profile can also be retrieved for each network. An outstanding question is then to probe the dynamics of those networks, which can be achieved through temporal modeling approaches. The most commonly applied tool in this setting has been the hidden Markov model (HMM). For instance, Ou *et al.* [22] could successfully parameterise the dynamics of functional brain states extracted by a sliding window framework and subsequent hard clustering, while Chiang *et al.* did so at the level of graph metrics [23]. In other works, functional brain states were estimated within the HMM pipeline, at the same time as their dynamics: to do so, each state was parameterised by a multivariate normal distribution, and the analysis focused either on its covariance matrix (interpreting a state as a *connectivity* pattern) [24], [25], or on its mean vector (regarding a state as an *activation* pattern) [26].

An important limitation of the above strategies is the assumption that at a given time, only one global brain state is expressed, whereas in reality, this whole-brain pattern reflects an overlap between the activity of multiple distinct functional networks [27]. With the TA/iCAPs framework, those separate building blocks are accurately extracted, and by parameterising their respective dynamics with HMMs, a clearer and more intuitive understanding of dFC is enabled. Here, we suggest a novel framework that can link those HMM descriptions together, so that possible dynamic cross-network relationships can also be characterised on top of intrinsic network dynamics. Several well-known observations in the literature based on stationary measures, such as the anti-correlation between the default mode network (DMN) and task-positive network (TPN) [28], or the triple network hypothesis in which the balance between those two networks is modulated by the salience network (SN) [29], are excellent candidates to be revisited in terms of dynamic interactions. Although some efforts towards this direction are emerging in the literature [30], they are hampered by the high computational load of the problem at hand, which has so far required to downscale the analysis to a limited subset of networks. We thus also propose to overcome this constraint by a sparsity-inducing regularisation term in our modeling framework, to enable only a limited set of cross-network couplings in data-driven manner.

In the following sections, we first briefly review the main features of the TA/iCAPs pipeline (II-A), and touch upon the key points in estimating network dynamics with standard HMMs (II-B). We then describe and justify all the main steps from our novel *sparse coupled hidden Markov model* (SCHMM) framework, including how cross-network couplings are introduced (II-C), how we ensure their sparsity at an optimal level of regularisation (II-D), and how coupling coefficients are thresholded to keep only the most significant ones (II-E). We then move to implementation details of the pipeline (II-F). Finally, the approach is validated on simulated data (II-G/III-A) and applied to real RS time courses (II-H/III-B).

II. METHODS

A. Total Activation and Innovation-Driven Co-Activation Patterns in Brief

Total activation (TA) is formulated as a regularised denoising problem; i.e., from the original BOLD signal matrix $\mathbf{Y} \in \mathbb{R}^{V \times T}$, where V is the number of voxels input to the algorithm and T the number of time points at hand, we look for the output $\mathbf{X} \in \mathbb{R}^{V \times T}$ such that

$$\tilde{\mathbf{X}} = \underset{\mathbf{X}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \mathcal{R}_T(\mathbf{X}) + \mathcal{R}_S(\mathbf{X}). \quad (1)$$

Let $\Delta_L = \Delta_D \Delta_{L_h}$ be the operator combining the deconvolution and differentiation operations (assuming a known hemodynamic response function). Then, $\Delta_L\{\mathbf{X}(v, \cdot)\} = \mathbf{U}_s(v, \cdot)$ is the *innovation signal*, peaking at the time points characterised by a signal change for voxel v , and the temporal regulariser is described as:

$$\mathcal{R}_T(\mathbf{X}) = \sum_{v=1}^V \lambda_T(v) \sum_{t=1}^T |U_s(v, t)|, \quad (2)$$

where $\lambda_T(v)$ is the temporal regularisation weight for voxel v .

Similarly, defining Δ_{Lap} as the 3D second-order difference operator and with $\lambda_S(t)$ the spatial regularisation weight at time point t , we have:

$$\mathcal{R}_S(\mathbf{X}) = \sum_{t=1}^T \lambda_S(t) \sum_{m=1}^M \sqrt{\sum_{v \in \mathcal{M}_m} \Delta_{Lap}\{X(v, t)\}^2}, \quad (3)$$

where we consider M regions as defined by a structural atlas [31], and \mathcal{M}_m is the set of voxels belonging to region m . This regularisation term imposes signal smoothness in each area from the atlas. For more details about implementing a discrete version of Δ_L and solving the TA problem, we refer to [32] and [20], respectively.

Following TA, the time points showing significant innovation across a sufficient fraction of brain voxels are extracted by using phase-randomized data for null distribution generation (see [21] for details). K-means clustering is applied to separate those frames into K distinct clusters. Let $I(t, k) = 1$ if the frame at time t is assigned to cluster k , and $I(t, k) = 0$ otherwise (i.e., either not retained as a significant innovation, or not assigned to cluster k); the set of iCAPs $\mathbf{C} \in \mathbb{R}^{V \times K}$ is then expressed as:

$$C(v, k) = \frac{\sum_{t=1}^T I(t, k) U_s(v, t)}{\sum_{t=1}^T I(t, k)}. \quad (4)$$

Finally, if $\mathbf{U}(v, \cdot) = \Delta_{L_h}\{\mathbf{X}(v, \cdot)\}$ is the deconvolved (activity-inducing) BOLD signal for voxel v , the activity time courses $\mathbf{T} \in \mathbb{R}^{K \times T}$ are obtained, for each time point t , as

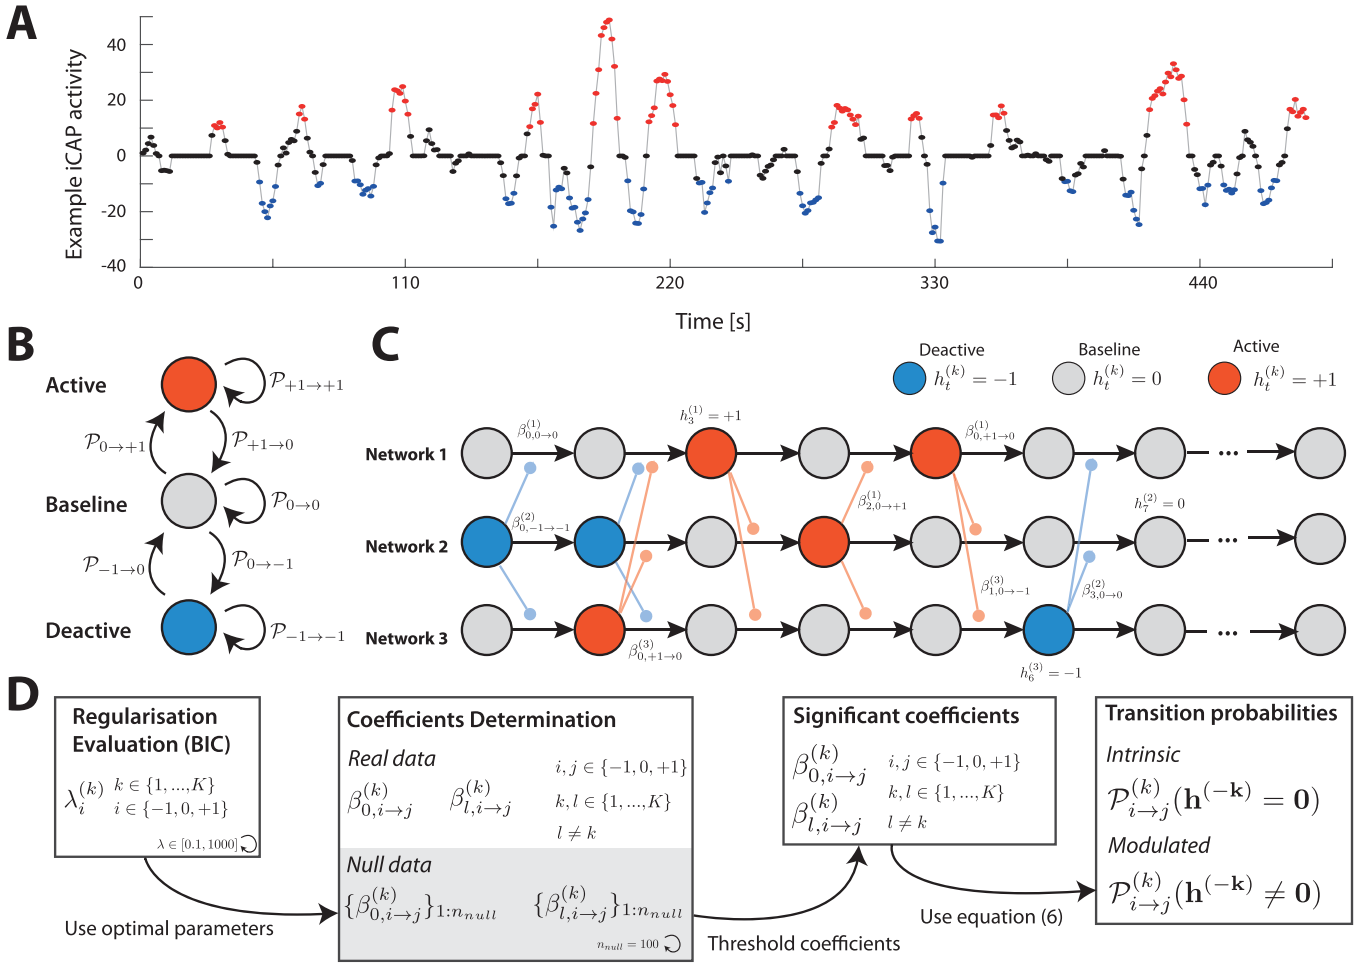


Fig. 1. Sparse coupled hidden Markov model framework. (A) Example iCAP activity time course (posterior DMN iCAP; see Figure 3A) for an indicative subject. Time points are colored according to the state to which they have been assigned by k-means clustering (see IV): deactive (blue), baseline (black) or active (red). (B) Three possible states of activity are hypothesised for each network: deactive (-1, blue), baseline (0, gray), and active (+1, red). From a time point to the next, networks have an intrinsic probability to transit across those states. (C) Example hidden state sequences for three networks, where some hidden states (h), intrinsic transition probability coefficients (β_0), and modulatory coefficients (β_l) are laid out (see II-C for details). (D) Global analytical pipeline of the SCHMM approach, where optimal regularisation parameters are first established (II-F), before the computation of modulatory coefficients on real and on null data (II-E). Significant modulatory coefficients are recovered, and converted into transition probabilities. Both intrinsic transition probabilities (when all other networks are at baseline activity level), and the ones under external modulatory influence (some other networks are (de)active), can be retrieved. K is the total number of analysed networks. BIC, Bayesian Information Criterion.

$T(k, t) = T_P(k, t) + T_N(k, t)$, with:

$$\begin{cases} \mathbf{T}_P(\cdot, t) = \operatorname{argmin}_{\mathbf{T}_P(\cdot, t)} \|\mathbf{U}(\cdot, t) - \mathbf{C}\mathbf{T}_P(\cdot, t)\|^2 \\ \text{s.t. } T_P(k, t) \in [0, +\infty[\\ \mathbf{T}_N(\cdot, t) = \operatorname{argmin}_{\mathbf{T}_N(\cdot, t)} \|\mathbf{U}(\cdot, t) - \mathbf{C}\mathbf{T}_N(\cdot, t)\|^2 \\ \text{s.t. } T_N(k, t) \in]-\infty, 0]. \end{cases} \quad (5)$$

B. Estimation of Network Dynamics With Parallel HMMs

In practice, iCAPs show a characteristic temporal profile involving excursions towards positive or negative levels of activity; Figure 1A illustrates this behaviour with an indicative, real iCAP activity time course. To capture these dynamics, let us denote by $h_t^{(k)}$ the activity state of iCAP k at time point t ; we assume that over time, this activity can switch between *deactive* ($h_t^{(k)} = -1$), *baseline* ($h_t^{(k)} = 0$), and *active* ($h_t^{(k)} = +1$) states, as depicted in Figure 1B. We denote

this set of possible activity states by $\mathcal{S} = \{-1, 0, +1\}$. Two assumptions fitting the data structure at hand are made at this stage: first, we enable only one discrete level of activation or deactivation; second, we consider a state diagram where a given network cannot directly transit from deactive to active state, or *vice versa*.

In a parallel hidden Markov model (PHMM) framework, each network k has its dynamics independently parameterised by its probability to start in state i , $\Pi_{k,i} = \mathcal{P}(h_1^{(k)} = i)$, and its probability to transit from state i to state j , $A_{i \rightarrow j}^{(k)} = \mathcal{P}(h_{t+1}^{(k)} = j | h_t^{(k)} = i)$. The observed values characterising each state i are also typically modeled by a normal distribution of mean $\mu_{k,i}$ and standard deviation $\sigma_{k,i}$.

C. Coupling Separate HMMs

In our framework, we hypothesise that transition probabilities across states for a given network evolve dynamically as

a function of the activity levels of the other networks (see Figure 1C). For instance, if one iCAP becomes active, it may increase the likelihood that one or several others enter a more active state as well, leading to a spatiotemporal sequence of brain activity akin to the ones put forward in previous dFC works [33].

The transition probability of network k from state i (at time t) to state j (at time $t + 1$), which we denote by $B_{t,i \rightarrow j}^{(k)}$, thus depends on two separate contributions: the intrinsic transition probability of network k itself, represented by the coefficient $\beta_{0,i \rightarrow j}^{(k)}$; and the modulatory influences of the other networks $l \neq k$, which we denote by $\beta_{l,i \rightarrow j}^{(k)}$. Following [34], we express the conditional probability to reach a given end state as a multinomial logistic regression:

$$B_{t,i \rightarrow j}^{(k)} = \mathcal{P}(h_{t+1}^{(k)} = j | h_t^{(k)} = i, \mathbf{h}_t^{(-k)}) = \frac{e^{\beta_{0,i \rightarrow j}^{(k)} + \sum_{l \neq k} \beta_{l,i \rightarrow j}^{(k)} h_t^{(l)}}}{\sum_{m \in \tilde{\mathcal{S}}} e^{\beta_{0,i \rightarrow m}^{(k)} + \sum_{l \neq k} \beta_{l,i \rightarrow m}^{(k)} h_t^{(l)}}}. \quad (6)$$

In this equation, $\mathbf{h}_t^{(-k)}$ refers to the activity level of all networks else than k , and the set $\tilde{\mathcal{S}}$ encompasses all the possible end states from the considered start state. If all networks else than k are in a baseline state of activity at time t ($h_t^{(l)} = 0$, for all $l \neq k$), then at this moment, only $\beta_{0,i \rightarrow j}^{(k)}$ contributes to the transition probability estimate (that is, network k evolves according to its intrinsic dynamics). If another network l is in the active state at time t , however, it can enhance/decrease (positive/negative $\beta_{l,i \rightarrow j}^{(k)}$) the probability of network k to transit from state i to state j from time t to time $t + 1$. If several other networks l are active, their modulatory influences sum up. In the case of deactive networks, the reasoning is the same, but the sign of the influence is flipped.

With this strategy, we thus assume fixed cross-network modulatory strengths (i.e., when a network modulates another, it always does so with the same magnitude), but those modulations are effective only at time points when the modulating networks are (de)active; this is what renders our approach dynamic.

D. Sparsity in Modulatory Influences

In practice, including all possible cross-network interactions in the model would amount to inferring $9K^2$ separate values, which becomes computationally demanding for a large number of networks, and also does not fit with our understanding of the brain, where only particular subsets of functionally related networks are expected to interact to generate the content of mind wandering [35].

For those reasons, we opt for a regularisation strategy where, for each network k and start state i , we constrain the set of incoming modulatory influences $\beta_{l,i \rightarrow j}^{(k)}$, $l \neq k$ to be sparse through ℓ_1 regularisation [36]. Formally, we thus impose:

$$\sum_{l \neq k} |\beta_{l,i \rightarrow j}^{(k)}| < \rho_{k,i}. \quad (7)$$

The choice of a *network-specific* regularisation level enables a more accurate representation of RS brain activity, where we hypothesise that some networks may receive a larger amount of modulating influences than others. For instance, the DMN has been associated to a wide array of brain functions [37], [38], and could thus be expected to be particularly coupled to other functional brain networks. Further, we also enabled *start state-specific* regularisation levels, because we wanted to include the possibility that modulating influences may not be equally potent on a deactive, an inactive, or an active network.

E. Thresholding of Coupling Coefficients

To ensure that modulatory coefficients are truly reflective of dynamic network interactions, we append another processing layer in which each $\beta_{l,i \rightarrow j}^{(k)}$ is compared to a distribution of values created under the null hypothesis of no such cross-talks. Only the coefficients that survive this thresholding step are included in the final model estimate (see Figure 1D, middle boxes).

To generate the null distributions, coefficients are recomputed at optimal regularisation levels (see II-F) on network time courses independently shifted by a random number of samples $n_s \in [1, T]$, in order to break down causality. This is done $n_{null} = 100$ times. For each coefficient, the 1st and 99th percentiles of the generated null distribution were chosen as thresholds, and only the coefficients lying outside of this interval were retained.

When the final set of coefficients has been obtained, Eq. (6) can be used to determine the transition probabilities of any network k both in the absence of any external modulation ($\mathbf{h}^{(-k)} = \mathbf{0}$), or under modulatory influences from the other networks ($\mathbf{h}^{(-k)} \neq \mathbf{0}$; see Figure 1D, rightmost box).

F. Implementation

To solve the SCHMM problem for network k and start state $i = 0$, for which there are three possible end states, we individually consider the set of coefficients related to each end state j by forming a partial quadratic approximation to the log-likelihood [34]. To retrieve $\beta_{0,i \rightarrow j}^{(k)}$ and $\beta_{l,i \rightarrow j}^{(k)}$ for $l \neq k$, we then need to minimize:

$$\frac{1}{2N} \sum_{t \in \mathcal{C}} [\omega_{t,j} (z_{t,j} - \beta_{0,i \rightarrow j}^{(k)} - \sum_{l \neq k} \beta_{l,i \rightarrow j}^{(k)} \tilde{h}_t^{(l)})^2] + \lambda_{k,i} \sum_{l \neq k} |\beta_{l,i \rightarrow j}^{(k)}|, \quad (8)$$

where $\mathcal{C} = \{t : h_t^{(k)} = i\}$ is the set of N selected data points (that is, the time points of interest when network k is in state i) and $\tilde{h}_t^{(l)}$ the estimated activity level of network l at time t . The coefficients of this regularised least square problem are given by:

$$\begin{cases} \omega_{t,j} = \tilde{B}_{t,i \rightarrow j}^{(k)} (1 - \tilde{B}_{t,i \rightarrow j}^{(k)}) \\ z_{t,j} = \tilde{\beta}_{0,i \rightarrow j}^{(k)} + \sum_{l \neq k} \tilde{\beta}_{l,i \rightarrow j}^{(k)} \tilde{h}_t^{(l)} + \frac{y_{t,j}^{(k)} - \tilde{B}_{t,i \rightarrow j}^{(k)}}{\omega_{t,j}}. \end{cases} \quad (9)$$

In the above, for the i to j state transition, $\tilde{\beta}_{0,i \rightarrow j}^{(k)}$ is the current estimate of the baseline regression coefficient of network k ,

$\tilde{\beta}_{l,i \rightarrow j}^{(k)}$ is the current estimate of the modulatory coefficient of network l on network k , and $\tilde{B}_{t,i \rightarrow j}^{(k)}$ is the current estimate of the transition probability of network k at time t . The term $y_{t,j}^{(k)} = \delta_{h_{t+1}^{(k)}, j}$ specifies, for network k , which data points from \mathcal{C} ($h_t^{(k)} = i$) are followed in the following time point by state j ($h_{t+1}^{(k)} = j$), denoting the transition of interest.

To solve this optimisation problem, we use PHMM outputs as initial transition probability coefficients ($\tilde{\beta}_{0,i \rightarrow j}^{(k)} = A_{i \rightarrow j}^{(k)}$) and initialise couplings at zero ($\tilde{\beta}_{l,i \rightarrow j}^{(k)} = 0$). To estimate the activity level of all networks and select the working set of data points \mathcal{C} for a given start state, we pick the most likely state at each time point according to the PHMM smoothed node marginals; i.e., $\mathcal{P}(h_t^{(k)} | \mathbf{T}(k, \cdot))$. Coefficients are iteratively updated across all three possible end states j , and recentered [34] as given by:

$$\begin{cases} \beta_{0,i \rightarrow j}^{(k)} \leftarrow \beta_{0,i \rightarrow j}^{(k)} - \bar{\beta}_{0,i}^{(k)} \\ \beta_{l,i \rightarrow j}^{(k)} \leftarrow \beta_{l,i \rightarrow j}^{(k)} - \max(\bar{\beta}_{l,i}^{(k)}, \hat{\beta}_{l,i}^{(k)}), \end{cases} \quad (10)$$

with $\bar{\beta}_{0,i}^{(k)} / \bar{\beta}_{l,i}^{(k)}$ the mean and $\hat{\beta}_{l,i}^{(k)}$ the median across end states, respectively.

Updates of the modulatory coefficients for a given network k are performed in random order, until convergence, through soft thresholding [39]. Let $R_{0,t,i \rightarrow j}^{(k)} = \sum_{l \neq k} \beta_{l,i \rightarrow j}^{(k)} \tilde{h}_t^{(l)}$ and

$R_{l,t,i \rightarrow j}^{(k)} = \sum_{m \neq k, l} \beta_{m,i \rightarrow j}^{(k)} \tilde{h}_t^{(m)}$; we then have:

$$\begin{cases} \beta_{0,i \rightarrow j}^{(k)} = \frac{\sum_{t \in \mathcal{C}} \omega_{t,j} (z_{t,j} - R_{0,t,i \rightarrow j}^{(k)})}{\sum_{t \in \mathcal{C}} \omega_{t,j}} \\ \beta_{l,i \rightarrow j}^{(k)} = \frac{\text{soft}(\sum_{t \in \mathcal{C}} \omega_{t,j} \tilde{h}_t^{(l)} (z_{t,j} - R_{l,t,i \rightarrow j}^{(k)}), \lambda_{k,i})}{\sum_{t \in \mathcal{C}} (\omega_{t,j} \tilde{h}_t^{(l)})^2}, \end{cases} \quad (11)$$

with $\text{soft}(x, \lambda) = \text{sign}(x)(|x| - \lambda)_+$ the soft thresholding operation applied on x with threshold λ .

Algorithmically speaking, the following is run for each start state i and network k :

```

Compute PHMM parameters
Initialise  $\beta_{0,i \rightarrow j}^{(k)}$  and  $\beta_{l,i \rightarrow j}^{(k)}$  for  $j \in \mathcal{S}$  and  $l \neq k$ 
Select working set  $\mathcal{C}$ 
Initialise end state  $j$ 
while  $\mathcal{L}^{(iter+1)} - \mathcal{L}^{(iter)} < \epsilon = 10^{-3}$  and  $n_{iter} < 200$  do
  Update  $\omega_{t,j}$  and  $z_{t,j} \forall j$ 
  Recenter  $\beta_{0,i \rightarrow j}^{(k)}$  and  $\beta_{l,i \rightarrow j}^{(k)}$  for  $l \neq k$ 
  Compute log-likelihood  $\mathcal{L}^{(iter)}$ 
  Compute  $\beta_{0,i \rightarrow j}^{(k)}$  and  $\beta_{l,i \rightarrow j}^{(k)}$  for  $l \neq k$ 
  Change end state
end while

```

In the simpler case of the start state $i = -1$ or $i = +1$, for which only two possible end states exist, end state iterative updates and recentering of coefficients are not needed.

To select optimal regularisation levels for each network/start state case, we perform grid search over the interval $\lambda_{k,i} \in [0.5, 1000]$, where we solve for modulatory coefficients as described above and select the scenario for which the Bayesian Information Criterion (BIC) [40] is minimal, that is, for which an optimum between model complexity (number of non-null coefficients) and data fitting quality is reached (see Figure 1D, leftmost box). As advised in [34], to speed up computations, warm restarts are used to initialise coupling coefficients, solving from larger to smaller $\lambda_{k,i}$ values.

G. Validation on Simulated Data

To validate our SCHMM pipeline, we generated 20 sets of simulated time courses (1000 data points per network in each set) for a system of three networks, with three possible activity states ($-1, 0, +1$) and normally distributed noise ($\sigma = 0.05$) added to each observation. We considered three different ground truth cases: (1) independent evolution of the networks, (2) modulation of network 2 onto network 1 to increase its overall activity, and (3) a similar modulation applied on both networks 1 and 3. Example activity time courses and ground truth transition probabilities for all networks and cases are presented in Figure 2, left column.

From the final modulatory coefficients retrieved by the SCHMM framework, we computed transition probabilities in the absence and in the presence of modulatory couplings using Eq. (6), and quantified the error made as the average absolute difference in probability with the ground truth across all possible transitions. We compared our SCHMM approach to the outcomes from the simpler PHMM scheme, where networks stay uncoupled (see II-B).

Alternatively, the set of modulatory coefficients can also be interpreted as a directed graph representation of network-to-network interactions. Coefficients are computed for all possible state transitions, but to simplify the analysis, we devised a summarising metric of *activity upregulation* describing the ability of a modulating network l , when turning active ($h_t^{(l)} = +1$), to entrain network k into a more active state itself.

For this purpose, we define the difference in transition probability between a case without any external influence by other networks, and one when only network l is active:

$$\Delta P_{l,i \rightarrow j}^{(k)} = \mathcal{P}(h_{t+1}^{(k)} = j | h_t^{(k)} = i, h_t^{(l)} = +1, \mathbf{h}_t^{(-k, -l)} = \mathbf{0}) - \mathcal{P}(h_{t+1}^{(k)} = j | h_t^{(k)} = i, \mathbf{h}_t^{(-k)} = \mathbf{0}). \quad (12)$$

If the value is positive, it means that for the considered i to j transition, there is an increased transition probability for network k upon activation of network l . We then define the activity upregulation of network l onto network k as:

$$K_{U,l}^{(k)} = \sum_{j>i} (\Delta P_{l,i \rightarrow j}^{(k)})_+ + \sum_{j<i} (-\Delta P_{l,i \rightarrow j}^{(k)})_+ \quad (13)$$

The value increases either if activity of network l makes transitions of network k towards more active states more likely (first term), or if it makes transitions towards lower activity states less likely (second term). One $K_{U,l}^{(k)}$ value can be seen as the edge of a directed graph from network l to network k .

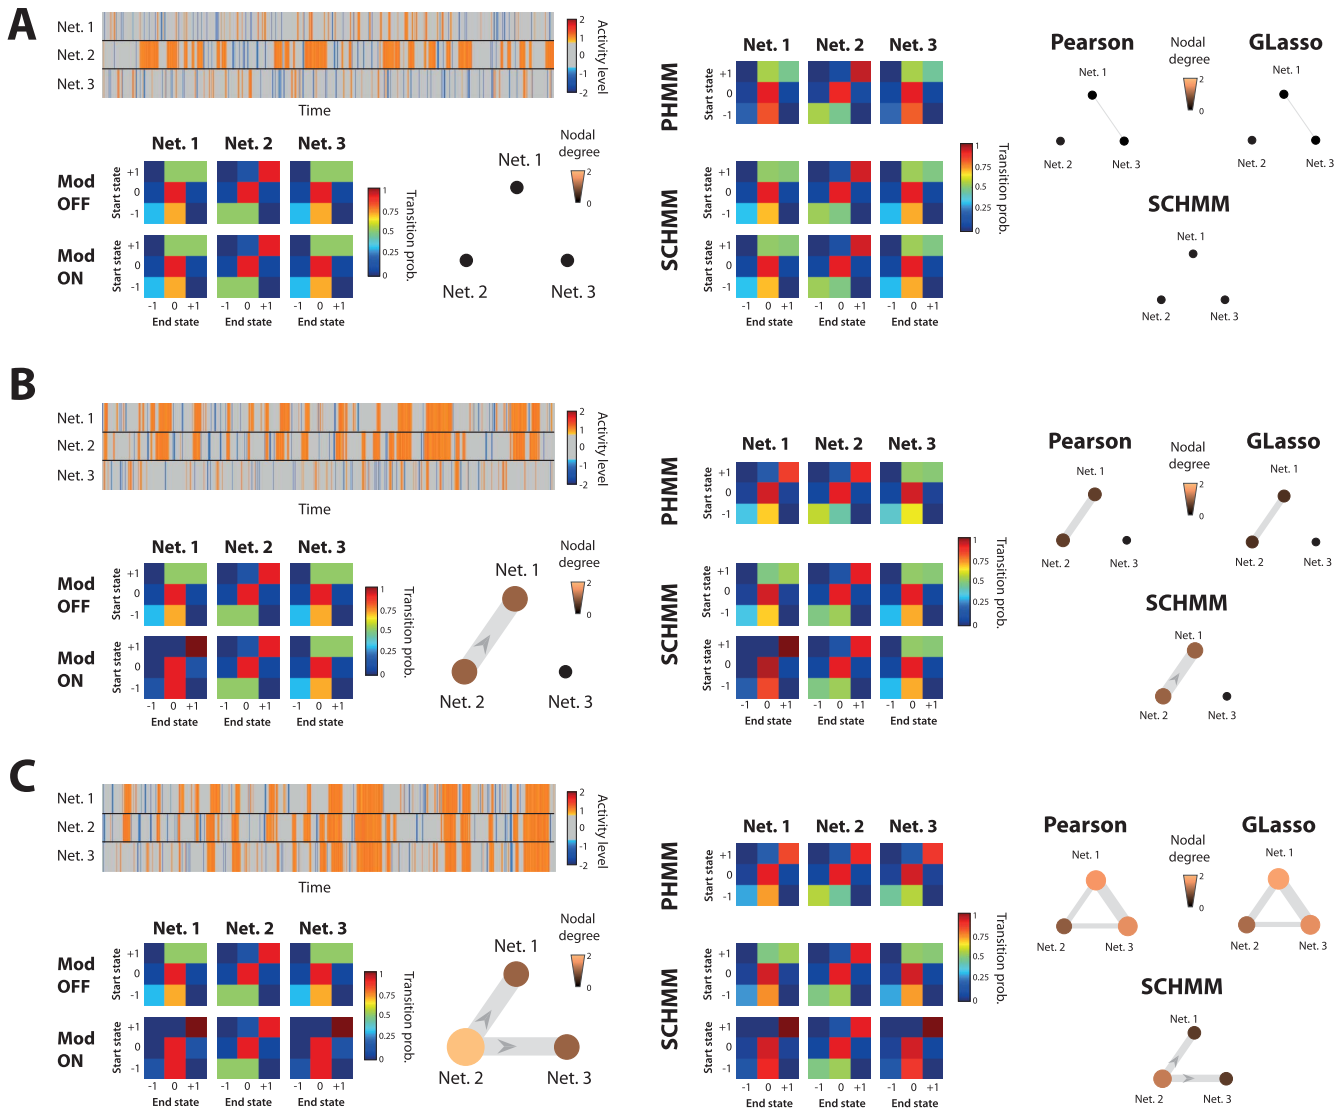


Fig. 2. Comparison of SCHMM results to other methods on simulated data. For three simulated cases (A, B and C), we display ground truth parameters of the system on the left hand side, including example time courses for the three simulated networks (top left plot), transition probability matrices describing the dynamics of the networks (bottom left), and a graph description of the modulations at play across networks (bottom right graph). In the middle panel, we present estimated transition probabilities with PHMMs (top row) or with the SCHMM approach, when cross-network modulations are enabled (bottom row) or not (middle row). In the right panel, we show the graph descriptions obtained using Pearson correlation coefficient (top left), the GLasso approach (top right), or the SCHMM approach (bottom). In the first scenario (A), networks evolve independently, and so the dynamics of the networks are stable over time. In the second scenario (B), network 1 has its dynamics altered when network 2 turns active (*Mod ON*). In the third scenario (C), networks 1 and 3 both have their dynamics similarly modified when network 2 turns active.

We compared the accuracy of this graph representation to the outcomes obtained with the more conventional Pearson correlation coefficient and graphical lasso (GLasso) [41] approaches. For the former case, coefficients were computed for each generated set of time courses, and thresholded using a null distribution approach similar to the SCHMM case (see II-E). For the latter case, where a sparse covariance matrix is obtained, grid search for an optimal regularisation level was performed, similarly to the SCHMM case (see II-F), prior to null data-based thresholding. In both settings, to obtain population-level measures that could be readily compared to the SCHMM activity upregulation metric, edge weight represented the fraction of subjects with a significant coefficient surviving past the thresholding process.

The graph representations of all three cases were compared to the ground truth using the average absolute difference in edge weight as the error measure. For each case, data was normalised so that the largest edge across the three examined simulated cases was set to 1. Because Pearson and GLasso outcomes are non-directional, we considered a symmetrical ground truth in those cases.

H. Application to Experimental fMRI Data

We wanted to determine whether reliable cross-network couplings could be retrieved on real RS data, and considered recordings from two independent datasets: the first (dataset DS1) was acquired on $n_{DS1} = 12$ healthy volunteers (38.4 ± 6 years old) with a Siemens 3T Trio TIM scanner,

TABLE I

EVOLUTION OF NETWORK DYNAMICS ESTIMATION ERROR ACROSS METHODS (PHMM VS SCHMM), MODULATION TYPE (ON VS OFF), EXAMINED SIMULATED CASES (A, B, C) AND NETWORKS (N1, N2, N3)

Method		PHMM		SCHMM	
Modulation		OFF	ON	OFF	ON
Case A	N1	0.0327	0.0327	0.0081	0.0081
	N2	0.0222	0.0222	0.0079	0.0079
	N3	0.0308	0.0308	0.0066	0.0066
Case B	N1	0.0855	0.1032	0.0166	0.0292
	N2	0.0187	0.0187	0.0066	0.0066
	N3	0.0245	0.0245	0.0045	0.0045
Case C	N1	0.0755	0.0925	0.0135	0.0151
	N2	0.016	0.016	0.0055	0.0055
	N3	0.1152	0.1226	0.0045	0.0234

using a 32-channel head coil and gradient-echo echo-planar imaging (TR/TE/FA = 1.1s/27ms/90°, matrix = 64 × 64, voxel size = 3.75 × 3.75 × 5.63mm³, 21 slices). We analysed the $T_{DS1} = 440$ data points (8.1min) from the iCAP time courses previously extracted from this dataset in [21], focusing on the restricted set of 13 functionally meaningful iCAPs discussed by the authors.

The second dataset (DS2) involved $n_{DS2} = 21$ healthy individuals (22 ± 2.3 years old) whose recordings were acquired with a Siemens 3T Trio TIM scanner, using a 12-channel head coil and gradient-echo echo-planar imaging (TR/TE/FA=2.1s/40ms/90°, matrix=128x84, voxel size=3.2 × 3.2 × 3.84mm³, 32 slices). We examined $T_{DS2} = 450$ functional volumes (15.8min), for which activity-inducing signals were computed by the TA framework [20]. Following normalisation to MNI space, DS1 iCAPs were back-projected onto those time courses to yield the analysed DS2 network activity profiles.

We compared the directional graphs resulting from our activity upregulation metric across datasets, and also to GLasso results. We assessed similarity in the set of retrieved couplings by the Jaccard index.

III. RESULTS

A. Validation on Simulated Data

Across the examined simulated examples, the estimation of modulatory coefficients for a particular network at optimal regularisation level always took less than a second for the deactive and active states, and varied from 1 to 5 minutes for the baseline state as a function of the assessed network and extent of regularisation. BIC grid search time was in the order of a minute for the active and deactive states, and climbed to around an hour for the baseline case. One iteration of the null data generation and computation process lasted for 6 to 7 minutes. Here and elsewhere, computations were run on an

Intel Xeon CPU E5 at 2.4GHz with 14 cores, 256GB RAM and Ubuntu 16.04.

Errors made in estimating the dynamics of simulated networks under different scenarios of modulation are displayed in Table I, for the PHMM case where networks are not coupled, and for our SCHMM framework where they are. In the three considered cases, when turning active, network 2 can influence both networks 1 and 3 (case C, Figure 2C), only network 1 (case B, Figure 2B), or no other network (case A, Figure 2A). For modulated networks, there are thus two different ground truth dynamics: the intrinsic one (*Mod OFF*), and the one upon modulation (*Mod ON*). It can be seen from error measurements that the SCHMM framework consistently outperformed the PHMM approach across networks, modulation cases and assessed scenarios.

In terms of graph representation, Pearson, GLasso and SCHMM approaches all successfully managed to retrieve the ground truth in case A (with respective errors of 0.0167, 0.0167 and 0) and case B (0.13, 0.16, 0). In the more complex case C, however, both the Pearson and GLasso graph estimates included an incorrect link between networks 1 and 3 (resulting in high errors of 0.683 and 0.63), whereas the SCHMM approach correctly retrieved the true graph structure (with a low error of 0.0002).

B. Application to Experimental fMRI Data

On two independently acquired RS datasets (DS1 and DS2), we next probed the existence of cross-network couplings across 13 iCAPs previously derived in [21] (see Figure 3A for spatial maps). The estimation of modulatory coefficients for a given network, at optimal regularisation level, always took less than a second for deactive and active start states, and around a minute in the baseline case. BIC grid search times were between 1 and 2 minutes for deactive and active start states, but a longer 2 to 3 hours for the baseline start state case. As for null data generation and computation, one complete iteration took from 5 to 15 minutes.

Of all possible $\beta_{1,i \rightarrow j}^{(k)}$ modulatory coefficients, 18.6%/21.79% (DS1/DS2), 58.97%/75.21% and 14.1%/17.31% survived the sparsity constraint for deactive, baseline and active start states, respectively. These values were reduced to 9.6%/14.1%, 30.77%/39.53% and 8.3%/12.18% following comparison to null data. The activity upregulation metric computed from those coefficients (see II-G) showed 29.49%/25% non-null coefficients, while with GLasso, 37.18%/33.3% of all possible couplings were retained.

Cross-network couplings found with the SCHMM and the GLasso approaches, for both examined datasets, are presented in Figure 3B. For DS1, almost half of the significant couplings found with the SCHMM framework matched GLasso-derived relationships ($J_{DS1} = 0.43$), and the same was observed for DS2 ($J_{DS2} = 0.5$). In particular, the MOT → AUD, pVIS → sVIS and DMN → ACC couplings were always amongst the strongest captured relationships.

Around half of the links captured with GLasso were shared across datasets ($J_{GLasso} = 0.49$). In the SCHMM case,

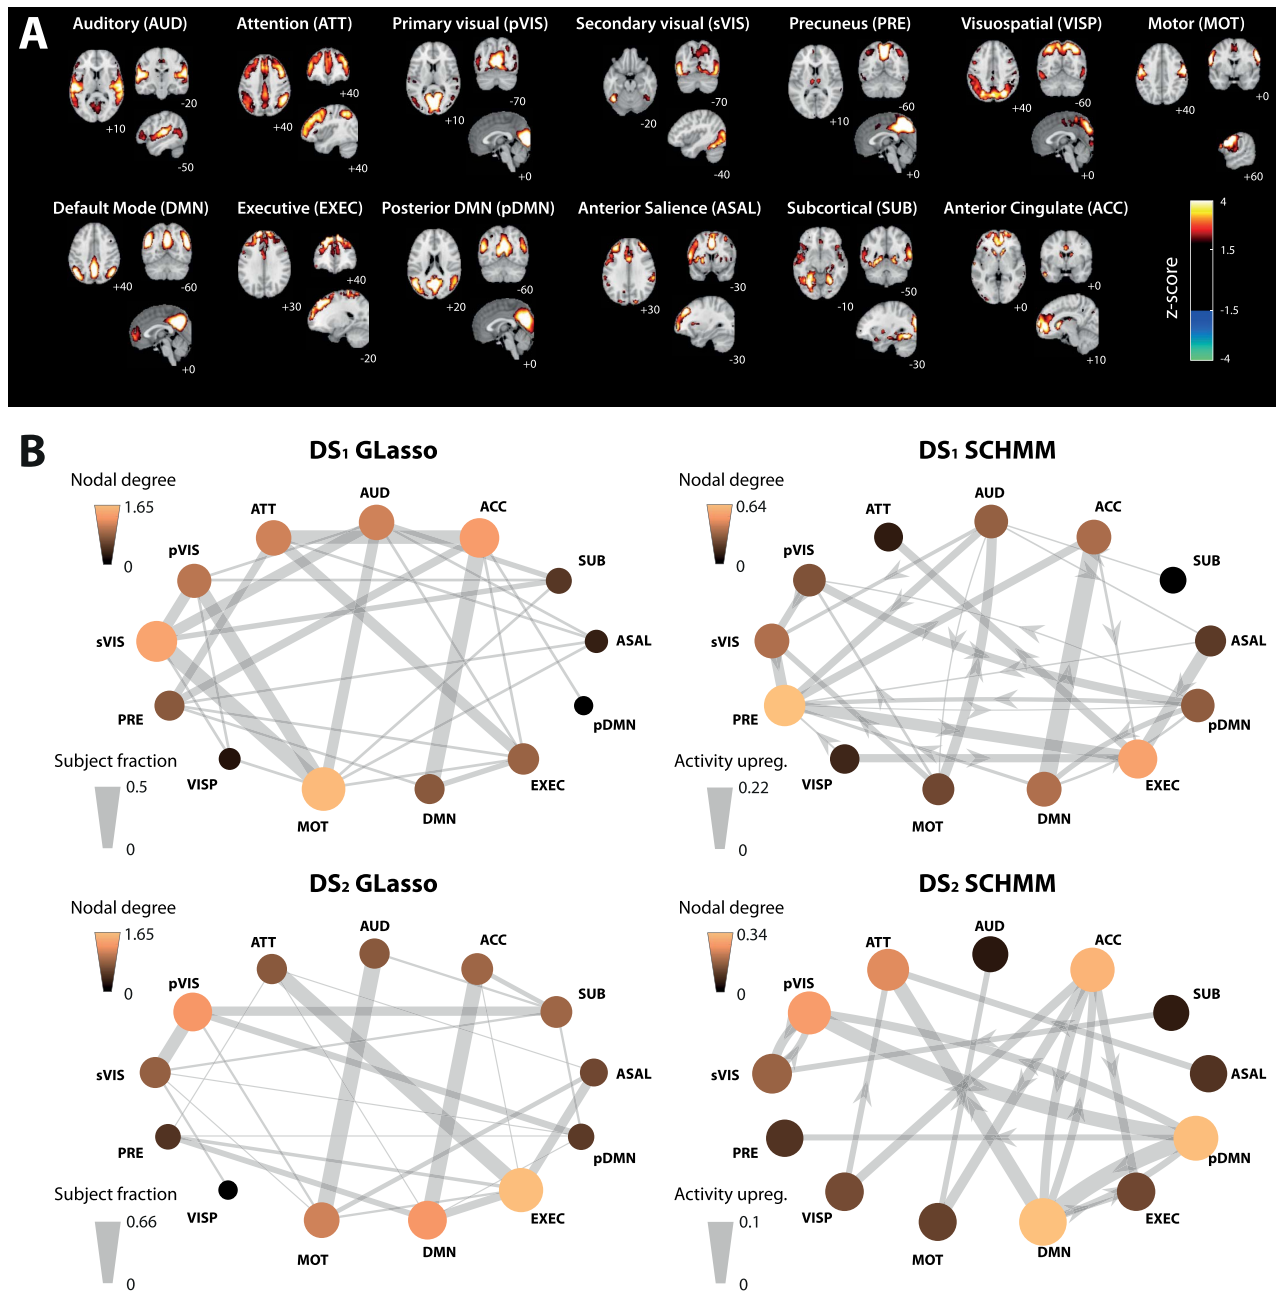


Fig. 3. Analysis of cross-network couplings on two independent datasets. (A) Spatially z-scored maps of the 13 examined networks, with denominations and abbreviations as in [21] and MNI coordinates shown in white caption below the brain slices. (B) For DS1 (top row) and DS2 (bottom row), graph representations of cross-network couplings as found with the GLasso (left column) and SCHMM (right column) approaches. In GLasso representations, edge weight stands for the fraction of subjects showing a significant relationship. In SCHMM displays, edge weight stands for activity upregulation values, and edges with a value lower than 0.01 (DS1) or 0.025 (DS2) are not displayed. The size and color coding of the nodes are proportional to their degree.

directional couplings were in agreement in almost a third of cases ($J_{SCHMM} = 0.28$), a lower value because of the additional directionality requirement. In addition, activity upregulation values were around two-fold lower in DS2. Interestingly, some of the strongest couplings consistently captured with the SCHMM across both datasets were not seen, or only barely detected (significant in only one subject) with the GLasso approach: this was the case of the pDMN→DMN, DMN→pDMN and pDMN→PRE links.

IV. DISCUSSION

The SCHMM framework successfully retrieved the ground truth dynamics of all examined networks across three simulated cases with an increasing amount of cross-network modulations, and constantly outperformed PHMMs in doing so. In our simulations including cross-network couplings, network 2 could modulate the others when turning active itself; this means that at some time points, the dynamics of the other networks were purely governed by their intrinsic

propensity to transit across activity states, while at others, their dynamics were altered. With a PHMM approach, those two types of moments are mixed in the estimated values, and thus cannot be disentangled. With the SCHMM, however, accurate transition probabilities can be retrieved both for intrinsic dynamics—setting modulatory influences to 0 in Eq. (6)—and for modulated ones.

On top of providing accurate estimates of network dynamics, the SCHMM approach could also successfully reproduce the ground truth graph structure (denoting directional modulations between networks) across all examined cases, using our summarising measure of activity upregulation. In the simpler cases with zero or one modulatory coupling, results were on par with Pearson and GLasso outcomes, but in the more elaborate case where network 2 was modulating the 2 others, the SCHMM was the only approach that correctly retrieved the true underlying graph structure, and did not mistakenly link networks 1 and 3. This is because, contrary to correlational approaches for which an edge reflects activity in two networks at the same time points (which may arise due to a third external source), the SCHMM considers whether the activity of a network at time t will drive a change in another from time t to time $t + 1$, somehow closer to effective connectivity tools [42].

Based on these results for simulated data, we could have expected a reduced amount of significant couplings on experimental fMRI data using the SCHMM, but this was not the case, possibly because when examining network relationships, indirect couplings may remain very limited, with direct network-to-network interactions dominating. Interestingly, around half of SCHMM cross-network couplings matched the ones retrieved with GLasso, and this held true on two independent datasets that we analysed. However, even if similar network-to-network relationships are retrieved, the SCHMM also recovers their directionality. For instance, the pVIS→sVIS link was more prominent than the sVIS→pVIS one in both datasets, which might indicate the dominant flow of visual information from low-level to high-level visual brain structures.

Some couplings were only detected by the SCHMM, and involved variants of the DMN, a system known to dissociate into separate subnetworks linked to different types of internal processes [38], [43]. The reason for the SCHMM sensitivity to those interplays may be that when the modulating network triggers enhanced activity in the modulated network, it also lowers in activity at the same time. This way, there is no temporally overlapping activity, and so no way for the GLasso to detect the relationship. Such spatiotemporal sequences in which a particular network (for instance, pDMN) progressively loses or gains some of its constituting nodes to change in spatial pattern (for instance, into DMN), perhaps because those nodes change their modular allegiance [44], [45], have already been resolved in RS recordings [33] without being further investigated.

Although insightful parallels could be drawn across our two examined datasets, the match in retrieved cross-network relationships remained partial, both for the GLasso and the SCHMM cases. Several factors may have contributed,

starting with the different ages of the studied populations, but we believe the main cause to be the different TRs of the acquisitions (1.1s for DS1, 2.1s for DS2), as SCHMM estimates, in particular, rely on frame-to-frame changes in activity. In accordance with this hypothesis, we noticed a roughly two-fold decrease of retrieved activity upregulation values in our DS2 dataset, possibly because rapid directional influences are more rarely observed in this setting.

Methodologically speaking, the use of sparsity-based strategies has already been suggested in past RS FC work, where sparsity was then either imposed at the level of functional connectivity matrices retrieved from a sliding window analysis (i.e., a limited amount of non-null connections was allowed in each matrix; see for example [8], [46], [47]), or at the level of global extracted functional connectivity brain states [24], [48]. With the present strategy, which is inspired from the bioinformatics field [49], we do not rely on any connectivity estimate, and we impose a restricted set of non-null modulatory coefficients onto a given network for a particular start state of activity. The power of the SCHMM approach is data-driven selection of relevant modulations, so that dimensionality of the problem remains affordable.

Finally, we note a few limitations and possible improvements of the current framework. First, to generate iCAP activity time courses, one could rely on an improved version of the TA approach that does not require the use of an atlas anymore, and instead imposes piecewise constant activity in space [50]. Then, to retrieve PHMM parameters, standard expectation maximisation (EM) [51] did not converge to a mixture solution that properly segregated activity states, and we thus resorted to a simpler approach where for each network k , state means $\mu_{k,i}$ were obtained as the centroids from a k -means clustering run on all time points $\mathbf{T}(k, \cdot)$, and standard deviations $\sigma_{k,i}$ were then computed on the time points assigned to each cluster (see Figure 1A). Only start and transition probabilities were iteratively updated within an EM scheme, keeping $\mu_{k,i}$ and $\sigma_{k,i}$ fixed.

Regarding the SCHMM framework itself, modulatory coefficients are so far of similar intensity, but opposite sign, when the modulating network is active or deactive, which may be an oversimplification. Also, parameters are assumed constant across the analysed subjects, which is a clear oversimplification knowing that individual fingerprinting can be reliably achieved on the basis of RS fMRI recordings [52]. In addition, there are alternatives to CHMM modeling, such as through fully-linked HMMs or dynamically multi-linked HMMs [53]; the comparison of those different approaches may consist in an interesting direction to follow. Finally, computational time is so far relatively high when one wishes to analyse systems made of more than a few networks; to improve in this regard, it could be interesting to consider a variant of the present model where only two different activity states are enabled, as the bulk of computational time is currently taken in solving for baseline start states of activity. This could then promote the application of the SCHMM framework to a dimensionally larger region-level setting.

V. CONCLUSION

In conclusion, our SCHMM framework showed promising potential to unravel directional cross-network couplings in fMRI data, including subtle interactions that could not be resolved with simpler correlational methods. We hope that our efforts shall pave the way towards more frequent analyses of brain network dynamics in future years.

ACKNOWLEDGMENT

The authors would like to thank Dr. F. I. Karahanoğlu for sharing part of the data analysed here.

REFERENCES

- [1] M. P. van den Heuvel and H. E. H. Pol, "Exploring the brain network: A review on resting-state fMRI functional connectivity," *Eur. Neuropsychopharmacol.*, vol. 20, no. 8, pp. 519–534, 2010.
- [2] M. Greicius, "Resting-state functional connectivity in neuropsychiatric disorders," *Current Opinion Neurol.*, vol. 21, no. 4, pp. 424–430, 2008.
- [3] M. D. Fox and M. Greicius, "Clinical applications of resting state functional connectivity," *Frontiers Syst. Neurosci.*, vol. 4, p. 19, Jun. 2010.
- [4] B. Biswal, F. Z. Yetkin, V. M. Haughton, and J. S. Hyde, "Functional connectivity in the motor cortex of resting human brain using echo-planar MRI," *Magn. Reson. Med.*, vol. 34, no. 4, pp. 537–541, 1995.
- [5] C. F. Beckmann, M. DeLuca, J. T. Devlin, and S. M. Smith, "Investigations into resting-state connectivity using independent component analysis," *Philos. Trans. Roy. Soc. London B, Biol. Sci.*, vol. 360, no. 1457, pp. 1001–1013, 2005.
- [6] J. S. Damoiseaux *et al.*, "Consistent resting-state networks across healthy subjects," *Proc. Nat. Acad. Sci. USA*, vol. 103, no. 37, pp. 13848–13853, 2006.
- [7] C. Chang and G. H. Glover, "Time-frequency dynamics of resting-state brain connectivity measured with fMRI," *NeuroImage*, vol. 50, no. 1, pp. 81–98, Mar. 2010.
- [8] E. A. Allen, E. Damaraju, S. M. Plis, E. B. Erhardt, T. Eichele, and V. D. Calhoun, "Tracking whole-brain connectivity dynamics in the resting state," *Cerebral Cortex*, vol. 24, no. 3, pp. 663–676, Mar. 2014.
- [9] E. Damaraju *et al.*, "Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia," *NeuroImage, Clin.*, vol. 5, pp. 298–308, Jul. 2014.
- [10] R. M. Hutchison and J. B. Morton, "Tracking the brain's functional coupling dynamics over development," *J. Neurosci.*, vol. 35, no. 17, pp. 6849–6859, Apr. 2015.
- [11] N. Leonardi *et al.*, "Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest," *NeuroImage*, vol. 83, pp. 937–950, Dec. 2013.
- [12] X. Li *et al.*, "Dynamic functional connectomics signatures for characterization and differentiation of PTSD patients," *Hum. Brain Mapping*, vol. 35, no. 4, pp. 1761–1778, Apr. 2014.
- [13] R. L. Miller *et al.*, "Higher dimensional meta-state analysis reveals reduced resting fMRI connectivity dynamism in schizophrenia patients," *PLoS ONE*, vol. 11, no. 3, p. e0149849, 2016.
- [14] S. M. Smith *et al.*, "Temporally-independent functional modes of spontaneous brain activity," *Proc. Nat. Acad. Sci. USA*, vol. 109, no. 8, pp. 3131–3136, 2012.
- [15] X. Liu and J. H. Duyn, "Time-varying functional network information extracted from brief instances of spontaneous brain activity," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 11, pp. 4392–4397, 2013.
- [16] X. Liu, C. Chang, and J. H. Duyn, "Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns," *Frontiers Syst. Neurosci.*, vol. 7, p. 101, Dec. 2013.
- [17] R. M. Hutchison *et al.*, "Dynamic functional connectivity: Promise, issues, and interpretations," *NeuroImage*, vol. 80, no. 4, pp. 360–378, Oct. 2013.
- [18] M. G. Preti, T. A. W. Bolton, and D. Van De Ville, "The dynamic functional connectome: State-of-the-art and perspectives," *NeuroImage*, to be published, doi: <https://doi.org/10.1016/j.neuroimage.2016.12.061>.
- [19] C. C. Gaudes, N. Petridou, S. T. Francis, I. L. Dryden, and P. A. Gowland, "Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses," *Hum. Brain Mapping*, vol. 34, no. 3, pp. 501–518, 2013.
- [20] F. I. Karahanoğlu, C. Caballero-Gaudes, F. Lazeyras, and D. Van De Ville, "Total activation: fMRI deconvolution through spatio-temporal regularization," *NeuroImage*, vol. 73, pp. 121–134, Jun. 2013.
- [21] F. I. Karahanoğlu and D. Van De Ville, "Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks," *Nature Commun.*, vol. 6, Jul. 2015, Art. no. 7751.
- [22] J. Ou *et al.*, "Characterizing and differentiating brain state dynamics via hidden Markov models," *Brain Topogr.*, vol. 28, no. 5, pp. 666–679, Sep. 2015.
- [23] S. Chiang *et al.*, "Time-dependence of graph theory metrics in functional connectivity analysis," *NeuroImage*, vol. 125, pp. 601–615, Jan. 2016.
- [24] H. Eavani, T. D. Satterthwaite, R. E. Gur, R. C. Gur, and C. Davatzikos, "Unsupervised learning of functional network dynamics in resting state fMRI," *Inf. Process. Med. Imag.*, vol. 23, pp. 426–437, Apr. 2013.
- [25] S. Ryali *et al.*, "Temporal dynamics and developmental maturation of salience, default and central-executive network interactions revealed by variational Bayes hidden Markov modeling," *PLoS Comput. Biol.*, vol. 12, no. 12, p. e1005138, 2016.
- [26] S. Chen, J. Langley, X. Chen, and X. Hu, "Spatiotemporal modeling of brain dynamics using resting-state functional magnetic resonance imaging with Gaussian hidden Markov model," *Brain Connectivity*, vol. 6, no. 4, pp. 326–334, 2016.
- [27] N. Leonardi, W. R. Shirer, M. D. Greicius, and D. Van De Ville, "Disentangling dynamic networks: Separated and joint expressions of functional connectivity patterns in time," *Hum. Brain Mapping*, vol. 35, no. 12, pp. 5984–5995, Dec. 2014.
- [28] M. D. Fox, A. Z. Snyder, J. L. Vincent, M. Corbetta, D. C. Van Essen, and M. E. Raichle, "The human brain is intrinsically organized into dynamic, anticorrelated functional networks," *Proc. Nat. Acad. Sci. USA*, vol. 102, no. 27, pp. 9673–9678, 2005.
- [29] V. Menon, "Large-scale brain networks and psychopathology: A unifying triple network model," *Trends Cognit. Sci.*, vol. 15, no. 10, pp. 483–506, 2011.
- [30] M. Sourty, L. Thoraval, D. Roquet, J.-P. Armspach, J. Foucher, and F. Blanc, "Identifying dynamic functional connectivity changes in dementia with Lewy bodies based on product hidden Markov models," *Frontiers Comput. Neurosci.*, vol. 10, p. 60, Jun. 2016.
- [31] N. Tzourio-Mazoyer *et al.*, "Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain," *NeuroImage*, vol. 15, no. 1, pp. 273–289, 2002.
- [32] F. I. Karahanoğlu, I. Bayram, and D. Van De Ville, "A signal processing approach to generalized 1-D total variation," *IEEE Trans. Signal Process.*, vol. 59, no. 11, pp. 5265–5274, Nov. 2011.
- [33] W. Majeed *et al.*, "Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans," *NeuroImage*, vol. 54, no. 2, pp. 1140–1150, Jan. 2011.
- [34] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *J. Stat. Softw.*, vol. 33, no. 1, pp. 1–22, 2010.
- [35] K. Christoff, Z. C. Irving, K. C. Fox, R. N. Spreng, and J. R. Andrews-Hanna, "Mind-wandering as spontaneous thought: A dynamic framework," *Nature Rev. Neurosci.*, vol. 17, pp. 718–731, Sep. 2016.
- [36] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Stat. Soc. B, Methodol.*, vol. 58, no. 1, pp. 267–288, 1996.
- [37] R. L. Buckner, J. R. Andrews-Hanna, and D. L. Schacter, "The brain's default network," *Ann. New York Acad. Sci.*, vol. 1124, no. 1, pp. 1–38, 2008.
- [38] J. R. Andrews-Hanna, J. S. Reidler, J. Sepulcre, R. Poulin, and R. L. Buckner, "Functional-anatomic fractionation of the brain's default network," *Neuron*, vol. 65, no. 4, pp. 550–562, 2010.
- [39] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani, "Pathwise coordinate optimization," *Ann. Appl. Stat.*, vol. 1, no. 2, pp. 302–332, 2007.
- [40] G. Schwarz, "Estimating the dimension of a model," *Ann. Stat.*, vol. 6, no. 2, pp. 461–464, 1978.
- [41] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, Jul. 2008.
- [42] K. J. Friston, "Functional and effective connectivity: A review," *Brain Connectivity*, vol. 1, no. 1, pp. 13–36, 2011.
- [43] J. R. Andrews-Hanna, J. Smallwood, and R. N. Spreng, "The default network and self-generated thought: Component processes, dynamic control, and clinical relevance," *Ann. New York Acad. Sci.*, vol. 1316, no. 1, pp. 29–52, 2014.

- [44] R. Baumgartner, G. Scarth, C. Teichtmeister, R. Somorjai, and E. Moser, "Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part I: Reproducibility," *J. Magn. Reson. Imag.*, vol. 7, no. 6, pp. 1094–1101, 1997.
- [45] R. F. Betzel, M. Fukushima, Y. He, X.-N. Zuo, and O. Sporns, "Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks," *NeuroImage*, vol. 127, pp. 287–297, Feb. 2016.
- [46] I. Cribben, R. Haraldsdottir, L. Y. Atlas, T. D. Wager, and M. A. Lindquist, "Dynamic connectivity regression: Determining state-related changes in brain connectivity," *NeuroImage*, vol. 61, no. 4, pp. 907–920, Jul. 2012.
- [47] C. Y. Wee, S. Yang, P. T. Yap, and D. Shen, "Sparse temporally dynamic resting-state functional connectivity networks for early MCI identification," *Brain Imag. Behav.*, vol. 10, no. 2, pp. 342–356, Jun. 2016.
- [48] H. Eavani, T. D. Satterthwaite, R. Filipovych, R. E. Gur, R. C. Gur, and C. Davatzikos, "Identifying sparse connectivity patterns in the brain using resting-state fMRI," *NeuroImage*, vol. 105, pp. 286–299, Jan. 2015.
- [49] H. Choi, D. Fermin, A. I. Nesvizhskii, D. Ghosh, and Z. S. Qin, "Sparsely correlated hidden Markov models with application to genome-wide location studies," *Bioinformatics*, vol. 29, no. 5, pp. 533–541, 2013.
- [50] Y. Farouj, F. I. Karahanoglu, and D. Van De Ville, "Regularized spatiotemporal deconvolution of fMRI data using gray-matter constrained total variation," in *Proc. IEEE 14th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2017, pp. 472–475.
- [51] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, Feb. 1989.
- [52] E. S. Finn *et al.*, "Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity," *Nature Neurosci.*, vol. 18, pp. 1664–1671, Oct. 2015.
- [53] L. Zhang, D. Samaras, N. Alia-Klein, N. Volkow, and R. Goldstein, "Modeling neuronal interactivity using dynamic Bayesian networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 18, 2006, p. 1593.