Contents lists available at ScienceDirect

# NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

# Time-resolved effective connectivity in task fMRI: Psychophysiological interactions of Co-Activation patterns



<sup>a</sup> Institute of Bioengineering, École Polytechnique Fédérale de Lausanne, Switzerland

<sup>b</sup> Department of Radiology and Medical Informatics, University of Geneva, Switzerland

<sup>c</sup> Department of Physics, University of Bristol, United Kingdom

<sup>d</sup> Department of Basic Neurosciences, University of Geneva, Switzerland

<sup>e</sup> Division of Development and Growth, Department of Pediatrics, University of Geneva, Switzerland

#### ARTICLE INFO

Keywords: PPI-CAPs Dynamic functional connectivity (dFC) Task fMRI Psychophysiological interaction (PPI) Co-activation patterns (CAPs) Framewise analysis

#### ABSTRACT

Investigating context-dependent modulations of Functional Connectivity (FC) with functional magnetic resonance imaging is crucial to reveal the neurological underpinnings of cognitive processing. Most current analysis methods hypothesise sustained FC within the duration of a task, but this assumption has been shown too limiting by recent imaging studies. While several methods have been proposed to study functional dynamics during rest, task-based studies are yet to fully disentangle network modulations.

Here, we propose a seed-based method to probe task-dependent modulations of brain activity by revealing Psychophysiological Interactions of Co-activation Patterns (PPI-CAPs). This point process-based approach temporally decomposes task-modulated connectivity into dynamic building blocks which cannot be captured by current methods, such as PPI or Dynamic Causal Modelling. Additionally, it identifies the occurrence of co-activation patterns at single frame resolution as opposed to window-based methods.

In a naturalistic setting where participants watched a TV program, we retrieved several patterns of coactivation with a posterior cingulate cortex seed whose occurrence rates and polarity varied depending on the context; on the seed activity; or on an interaction between the two. Moreover, our method exposed the consistency in effective connectivity patterns across subjects and time, allowing us to uncover links between PPI-CAPs and specific stimuli contained in the video.

Our study reveals that explicitly tracking connectivity pattern transients is paramount to advance our understanding of how different brain areas dynamically communicate when presented with a set of cues.

#### 1. Introduction

Since its introduction in the early 1990s (Ogawa et al., 1990), functional magnetic resonance imaging (fMRI) has played a growing role in advancing our knowledge of brain activity. Its technical developments have allowed the study of brain function at increasingly high spatial and, moreover, temporal resolution (Van Essen et al., 2013). Simultaneously, the progress in analysis techniques has exposed the joint importance of, on the one hand, how brain regions activate and, on the other hand, how their activity interacts to support complex cognitive processes. This view of the brain as a network of functionally linked regions has spawned the field of Functional Connectivity (FC) which, traditionally, uses Pearson's correlation to study temporal dependencies between separate brain regions over the duration of an entire resting-state fMRI run—typically in the order of several minutes (van den Heuvel and Hulshoff Pol, 2010).

The discovery that FC significantly fluctuates over time in restingstate fMRI recordings (Chang and Glover, 2010) first suggested that methods driven by averaging over long runs offer an incomplete picture of brain function, which led to a widespread effort to investigate *dynamic* Functional Connectivity (dFC; Hutchison et al., 2013; Calhoun et al.,

E-mail address: lorena.freitas@epfl.ch (L.G.A. Freitas).

https://doi.org/10.1016/j.neuroimage.2020.116635

Received 30 September 2019; Received in revised form 10 February 2020; Accepted 12 February 2020 Available online 25 February 2020

1053-8119/© 2020 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).





Abbreviations: AN, attention network; CAP, co-activation pattern; DMN, default mode network; FC, functional connectivity; fMRI, functional magnetic resonance imaging; FPN, frontoparietal network; HRF, haemodynamic response function; MNI, montreal neurological institute; PCC, posterior cingulate cortex; PPA, point process analysis; PPI, psychophysiological interaction; SD, standard deviation; SN, salience network; VSN, visuospatial network.

<sup>\*</sup> Corresponding author. Institute of Bioengineering, École Polytechnique Fédérale de Lausanne, Switzerland.

2014; Preti et al., 2016; Karahanoğ;lu & Van De Ville, 2017). Since then, several methodological developments have been proposed to capture this feature. The most common technique involves computing a metric characterising FC over gradually shifted temporal windows of data (*sliding window approach*; Leonardi et al., 2013; Allen et al., 2014). While this increases the temporal refinement of the analysis from minutes to several seconds, it cannot fully benefit from the current sub-second resolution of fMRI recordings because a minimal window size of at least 30 s is required to obtain reliable correlation estimates (Kucyi and Davis, 2014; Shen et al., 2016; Preti et al., 2016). Furthermore, shorter windows demand more stringent high-pass filtering of the original time courses to avoid spurious correlations due to aliasing, thus limiting the available information (Leonardi & Van De Ville, 2015; Zalesky and Breakspear, 2015).

Essentially, the sliding window approach keeps the idea of computing second-order statistics, but within each window. Dynamic Conditional Correlation (DCC; Lindquist et al., 2014), a method based on multivariate generalized autoregressive conditional heteroskedasticity models, overcomes some of the methodological issues inherent to traditional sliding window correlation by gradually refining the FC estimate with each new sample, thus requiring no *ad hoc* parameter settings. However, the associated gain of performance in capturing meaningful neuronal fluctuations has been inconsistent across studies (Choe et al., 2017; Damaraju et al., 2018). Other extensions include a more systematic use of the wavelet coherence transform (Rack-Gomer and Liu, 2012; Yaesoubi et al., 2015) as originally proposed in the exploratory results of Chang and Glover (2010).

Another category of analyses follows a framewise approach. Good examples of these are point process analysis (PPA)-based methods (Tagliazucchi et al., 2012), which select a subset of frames where a chosen seed is highly active and proceed from there. Liu et al. (2013) benefited from this to identify co-activation patterns (CAPs) by clustering the retained frames into groups of similar activation arrangements. Another way of performing an implicit selection of relevant information is through sparsity-driven detection of neuronal activation time points (through Sparse Paradigm Free Mapping; Caballero Gaudes et al., 2011; Petridou et al., 2013) or moments of transient activity (innovation-driven CAPs; Karahanoglu & Van De Ville, 2015). All of these approaches have been mainly used to study resting-state data.

Unsurprisingly, the dynamic nature of connectivity in the brain is also expressed in the presence of external stimulation or during task performance (Gonzalez-Castillo and Bandettini, 2018). The natural assumption is thus that certain brain areas interact differently during the course of a task experiment. Beta Series Correlations (BSC, Rissman et al., 2004) analysis, for example, reveals the absolute FC between brain regions under different stages of task performance. Another reasonable expectation is that different regions may change the way they interact with each other when performing different tasks. These can be studied using methods such as Effective Connectivity (EC) analyses to look into the influence one neural system has on another and how this relationship changes between task settings (which we will refer to as contexts or conditions in what follows). This can be done using methods as varied as regression models such as Psychophysiological Interaction (PPI) analysis (Friston et al., 1997) and its generalizations (McLaren et al., 2012); differential equation models such as dynamic causal modelling (Friston et al., 2003) and causal dynamic network modelling (Cao et al., 2019); structural equation modelling (Zhuang et al., 2008); and Granger causality (Wen et al., 2013). These approaches, however, do not explicitly reveal moment-to-moment interactions between experimental conditions and brain activity, but rather average over the duration of the task/rest epochs, which is probably too limiting as has been shown from high temporal resolution neuroimaging techniques (Ploner et al., 2009; Zhang et al., 2012). As a natural development from these, window-based approaches to compute time-varying networks have been used to capture changes in task-related functional connectivity (Di et al., 2015; Baczkowski et al., 2017; Ge et al., 2019), although this also comes with known

disadvantages as discussed above. Recent studies (e.g., Fransson et al., 2018) have steered from windowed correlations, but there is still a great need for novel techniques to study how tasks modulate moment-to-moment connectivity at high temporal resolution.

Here, we introduce Psychophysiological Interaction of Co-Activation Patterns (PPI-CAPs) as a novel seed-based approach to investigate timeresolved effective connectivity. Our aim was to create a method that reveals the modulation of moment-to-moment connectivity between brain regions under a specific context or during performance of a task. To illustrate our framework, we applied PPI-CAPs to an fMRI dataset where subjects were exposed to a naturalistic paradigm by watching a short episode of a TV program containing two types of scenes (conditions). Our method dissected the connectivity patterns elicited by subjects across time, and found several PPI-CAPs with at least one of three possible effects: 1) a seed effect, indicating that the pattern was highly correlated with the seed activity in general; 2) a condition effect, meaning that the pattern in question was significantly more elicited during one of the types of scenes than during the other one; and 3) an interaction effect, representing an interaction between condition and the relationship of that coactivation pattern with the seed. Additionally, we shed light on the consistency in effective connectivity patterns across subjects and time, allowing to uncover links between PPI-CAPs and specific movie cues. Overall, our approach contributes to the state-of-the-art by unraveling time-resolved, relevant information on brain dynamics during task performance that cannot be captured by other methods.

# 2. Methods

We first provide a global overview of the PPI-CAPs analysis pipeline (see Fig. 1). We start from fMRI data where an experimental modulation has been applied and the timing is known (Fig. 1A). Similar to the framework of conventional PPA and CAPs, we start by selecting frames at time points when a predefined is seed most highly (de-)activated (Fig. 1 B/C). A static analysis illustrates the relevance of the frames that have been selected (Fig. 1D). Then, we proceed to the dynamic (PPI-CAPs) analysis (Fig. 1E/F).

For the static analysis, we first multiply the selected frames by the sign of the seed and the centred modulating term (i.e., the contrast variable, as it encodes knowledge of the task paradigm) at the corresponding time points. Note that this sequence is equivalent to multiplying the original selected frames by the PPI variable, which corresponds to the multiplication of the sign of the seed and the task time courses. The average of these selected frames leads to a proxy of the conventional PPI results (Fig. 1D, bottom), which we refer to as the *static interaction map* (siMap). For a mathematical motivation of the link between conventional PPI and static PPI-CAP analyses, as well as a toy example that illustrates this relationship intuitively, we refer to A and Supplementary Fig. A1, respectively.

For the dynamic analysis, all the originally retained frames are subjected to K-means++ clustering (Arthur and Vassilvitskii, 2007) as described in section 2.3 (Fig. 1E). This step yields PPI-CAPs, their occurrence in time and the polarity of their constituting frames, which allows us to eventually identify meaningful statistics (i.e., effects of task, seed, or their interaction). The next sections further detail the different steps of the pipeline.

#### 2.1. Seed-based frame selection

The first step of the PPI-CAPs framework is to select a seed region according to prior knowledge about the task being studied or on an exploratory basis—a good discussion on how to choose a seed can be found in O'Reilly et al. (2012). The fMRI frames in which the seed reaches high magnitude values are then selected for further analysis. We refer to frames in which absolute seed activity is above the stipulated threshold as *suprathreshold frames*.



**Fig. 1. The PPI-CAPs analysis pipeline.** A) We begin with an experimental design containing different task blocks or contexts. Here, we used data acquired in a naturalistic paradigm in which subjects watched a film containing two main types of scenes: *fun* and *science*. B) For each subject, the z-scored signal from a selected seed is thresholded so that frames in which it is highly active or deactive (darkened in the subject time course) can be considered for further analysis. Orange and blue frames occur during *fun* and *science* scenes, respectively. C) Suprathreshold frames from all subjects are concatenated. From here on, we can proceed to a static or a dynamic analysis. D) Static analysis: suprathreshold frames from each condition are multiplied by the sign of the seed and task time courses (which corresponds to the PPI time course) and then averaged, yielding a proxy of the PPI analysis results: the static interaction map, or siMap, at either subject or group level. E) Dynamic analysis: suprathreshold frames from all subjects are clustered into a set of PPI-CAPs. Frame labels allow us to count how often a PPI-CAP occurs in each condition in its positive or negative polarity and compare these to the signs of each effect of interest. F) The frames' polarity after clustering tends to correlate with the sign of the effects it represents. By examining the confusion matrix for each effect (seed, task or PPI), we can determine if a PPI-CAP is strongly related to it (i.e., when the confusion matrix has no obvious pattern).

#### 2.2. Static analysis

In the context of CAP analysis (Tagliazucchi et al., 2012; Liu and Duyn, 2013), simple averaging of selected frames provides a good proxy for seed-based functional connectivity (see Appendix A), showing that the chosen subset of data contains relevant information. For large datasets, this observation holds for a wide range of thresholds retaining 15–90% of frames (Liu and Duyn, 2013, Fig. 1B), while for smaller datasets the threshold must aim at a trade-off between dramatically decreasing the amount of data used for analysis while keeping enough frames to reduce the effects of noise. Note that, besides improving the signal-to-noise ratio, including more frames in the analysis of smaller

datasets improves the chance that enough frames will be available for the dynamic analysis.

In a similar validation step, we perform an initial static analysis procedure that can be compared against conventional PPI. Specifically, suprathreshold frames were multiplied by a centred contrast variable—time points occurring during the first context had positive values and those belonging to the second one had negative values—and subsequently averaged (Fig. 1D, yellow inset). Here, we used the strategy described in Liu and Duyn (2013) to improve the signal-to-noise ratio of an fMRI frame, where a mask was created to cover voxels with the highest 10% and lowest 5% values, and all other voxels were set to zero. The spatial correlation between the resulting static interaction map

(siMap) and the PPI results is measured using spatial Pearson's correlation coefficient. This correlation was repeatedly calculated for siMaps including 5-100% of all frames both at subject and group levels, to test the robustness of the method to the choice of threshold. Since correlation values are range-constrained, when calculating their averages as described in section 3.1 and shown in Fig. 2C we first applied Fisher's z transform (the inverse hyperbolic tangent) to all Pearson's r values (Cox, 2008). After calculating the means, we then converted the results back to Pearson's r values by calculating the former's hyperbolic tangent. For the group maps, z-scored seed activations were calculated for all subjects, whose data were concatenated and frames selected as per individual subjects. The PPI analysis was performed using the standard procedure from SPM8 where three regressors are included in a general linear model design, relating to the time courses of the PPI, the seed, and the task, respectively. Note that the PPI regressor in SPM is derived by deconvolving the haemodynamic response function (HRF) from the seed time course, multiplying the latter by the task time course, and then reconvolving the final series with the HRF. The exact same task time course and seed region were used for both analyses.

## 2.3. Dynamic analysis

For a dynamic analysis of brain function during task, we identify the PPI-CAPs that may present a seed, task, or interaction effect (Fig. 1E). To this end, we applied K-means++ clustering on the suprathreshold frames of all subjects, using a *modulo*- $\pi$  cosine distance as a similarity measure, which we have named as such because it consists in the following: frames among which only the sign of the voxels was reversed were considered to be representations of the same pattern, with opposite polarity (e.g., if frame *a* showed an active prefrontal cortex and deactive occipital cortex, while frame *b* displayed a deactive prefrontal cortex and active occipital cortex, then, assuming  $\mathbf{F}_a$  and  $\mathbf{F}_b$  are the frames' voxel patterns,  $\mathbf{F}_a \approx -\mathbf{F}_b$ ).

Traditionally, the iterative algorithm to implement K-means consists of two steps: 1) assigning each data point to its closest centroid given the defined similarity metric d; 2) updating the centroids according to their assigned data points (e.g., by updating each centroid with the average of the normalized data points most recently assigned to it). At each iteration, the cluster label assigned to the  $i^{th}$  frame,  $L_i$ , is given by:

$$\mathbf{L}_{i} = \operatorname{argmin}(\operatorname{d}(\mathbf{C}_{\kappa}, \mathbf{F}_{i})), \tag{1}$$

where  $\kappa$  runs over the number of clusters,  $C_{\kappa}$  is the value of the  $\kappa^{th}$  cluster's centroid,  $F_i$  is the voxelwise activation pattern for the  $i^{th}$  frame, and d is the selected distance metric.

To add information regarding the frame polarity in PPI-CAPs, in our approach we set *d* to be the *modulo*- $\pi$  cosine distance (mpcos):

$$d(\mathbf{x}, \mathbf{y}) = \operatorname{mpcos}(\mathbf{x}, \mathbf{y}) = 1 - \left| \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\| \|\mathbf{y}\|} \right|,$$
(2)

where  $\mathbf{x} \cdot \mathbf{y} = \sum_{k} x_k y_k$  is the standard inner product and the norm is  $\mathbf{x} =$ 

 $\sqrt{\mathbf{x} \cdot \mathbf{x}}$ . The polarity  $P_i$  of frame  $F_i$  is thus equivalent to sign( $\mathbf{C}_{\kappa} \cdot \mathbf{F}_i$ ). From here on, we will describe frames for which  $P_i = 1$  as having a "positive polarity", and those for which  $P_i = -1$  as having a "negative polarity". Note how the metric in equation (2) compares to the standard cosine distance, given by:

$$\cos(\mathbf{x}, \mathbf{y}) = 1 - \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\| \|\mathbf{y}\|},$$
(3)

meaning that *modulo*- $\pi$  cosine implies no additional complexity as compared to the standard cosine distance metric. The value of each centroid  $C_{\kappa}$  is then updated as such:

$$\mathbf{C}_{\kappa} = \frac{1}{N_{\kappa}} \sum_{i=1}^{N_{\kappa}} P_i \frac{\mathbf{F}_i}{\|\mathbf{F}_i\|},\tag{4}$$

where  $N_{\kappa}$  is the number of frames in the  $\kappa^{th}$  cluster and  $||\mathbf{F}_i||$  is the magnitude of frame  $\mathbf{F}_i$ .

The final cluster centroids then form the PPI-CAPs. Each frame is thus annotated according to: 1) the time point it corresponds to; 2) the task or condition label which corresponds to that time; 3) the subject to whom it belongs; and 4) the polarity in which it occurred (positive or negative). This information can then be used to investigate differences in PPI-CAP occurrence across settings.

# 2.4. Significance assessment

If a PPI-CAP has a strong main or interaction effect, the polarity of the frames that constitute it will tend to correlate with the sign of that effect for the same time points. To visualise if this is the case, we can thus generate confusion matrices for each effect and PPI-CAP. When there is a strong correlation, the higher values of a confusion matrix will tend to load on one of its diagonals, and the relevance of this relationship can be measured by taking the matrix's determinant, which we will call the *det*-*index*. To test whether this value is significant, we can thus generate a null distribution by performing random permutations of the effect of interest's labels and re-calculate the det-index each time. Finally, we see where the real det-index stands in the distribution (Fig. 1F). For the results shown here, we performed 3000 random permutations for each test. We disclose the uncorrected p-values and indicate the significance level that should be used to correct for multiple corrections, controlling for the number of PPI-CAPs.



**Fig. 2.** The static interaction Map (siMap) using a subset of fMRI frames accords with the PPI contrast obtained from all the available data. A) Frames from time points when the considered seed is highly active are selected. A fraction of seed time course is shown for one subject. B) Retained frames are then multiplied by the PPI variable and averaged, yielding the static interaction map (siMap), as shown for the same subject with a yellow background. For comparison, the contrast map resulting from a PPI analysis using all frames is also shown with a red background. C) The siMap and PPI contrast map are reasonably correlated (r > 0.7) even when the former is computed using only 15% of the frames, and this similarity remains high as more frames are used, showing that the selection of relevant frames is robust to the choice of threshold. The dark blue curve represents the average correlation between the siMap and 1<sup>st</sup>-level PPI contrast map across subjects, and the light blue shading denotes the standard error. The correlation means were calculated by averaging Fisher's z transforms of the Pearson's r values, and transforming the result back to Pearson's r.

# 2.5. Choosing the number of PPI-CAPs

Effective Connectivity methods such as PPI provide a summary spatial map of task-specific seed relationship modulation. To disentangle which and when instantaneous patterns of activity support the summarized PPI findings, we must first determine the number of clusters into which to categorize the data. To this end, we employed Consensus Clustering (Monti et al., 2003). This approach applies K-means clustering on several subsamples of the data and calculates the *consensus matrix*  $\mathcal{M}$ . Each element  $\mathcal{M}(a, b)$  indicates the fraction of subsamples in which two frames *a* and *b* were both retained and clustered together. The optimal number of clusters can then be inferred by visual inspection of the ordered matrix  $\mathcal{M}$ , as well as of the cumulative distribution function (CDF) of  $\mathcal{M}$  for different values of *k*.

Additionally, for every k = 3, 4, ..., 8 we calculated the number of frames from each subject that contributed to each of the *k* PPI-CAPs. This helped us choose a *k* value for which the distribution of PPI-CAPs across subjects was roughly uniform. We applied consensus clustering for k = 3, 4, ..., 8 using 10 random subsamples for every *k*. Each subsample included 80% of the suprathreshold frames of all subjects, and K-means was computed for 50 random initializations for each subsample. To obtain the final clustering result, we applied K-means clustering with the optimum *k* on 100% of the suprathreshold frames and kept the best result from 50 random initializations, i.e., the one that minimised the total sum of modulo- $\pi$  cosine distances between frames and centroids.

# 2.6. Experimental data

To validate the method, we used fMRI data from 16 healthy subjects (mean age  $22.92 \pm 8.14$  years) watching a short TV program about the effects of sun exposure. The video alternated between two contexts: 1) images of several children playing by the beach (from here on described as fun); 2) scenes where scientific concepts were explained in a laboratory (we will call this context *science*). The movie can be watched online,<sup>1</sup> and a detailed description of the dataset can be found in Jochaut et al. (2015). One subject was excluded due to high motion for having >25% of frames with framewise displacement (Power et al., 2010) higher than 0.5 mm. Thus, 15 subjects were analysed (mean percentage of scrubbed frames = 2.2%, SD = 5.8%). 179 vol (Tim-Trio; Siemens, 40 transverse slices, voxel size =  $3 \text{ mm} \times 3 \text{ mm} \times 3 \text{ mm}$ ; repetition time = 2000 ms; echo time = 50 ms; field of view = 192) were available per subject, as well as an anatomical T1-weighted rapid acquisition gradient echo sequence (176 slices, voxel size = 1 mm  $\times$  1 mm  $\times$  1 mm, field of view = 256), acquired at the end of the scanning. All participants have given their written informed consent, which was approved by the local ethics committee (Biomedical Inserm protocol C08-39).

# 2.7. fMRI preprocessing

Functional images were preprocessed using SPM8 (Wellcome Department of Imaging Neuroscience, UK) where they were realigned to correct for head motion; coregistered with structural images; normalized in the Montreal Neurological Institute (MNI) stereotactic space; and spatially smoothed using a 6 mm full width at half maximum isotropic Gaussian kernel. In order to remove haemodynamic temporal blurring and better approximate neural activity, blood oxygenation level-dependent (BOLD) signals were deconvolved with the canonical haemodynamic response function from SPM8. This was done using an implementation of the Wiener filter from spm\_peb\_ppi.m, an SPM8 function that computes BOLD deconvolution in the context of PPI analyses. Voxels were then z-scored in time (Liu and Duyn, 2013).

# 2.8. Considerations for the application of PPI-CAPs

We applied PPI-CAPs on a movie-watching dataset to illustrate its potential to uncover task-related time-resolved effective connectivity in a realistic setting. To this end, the Posterior Cingulate Cortex (PCC) was selected as a seed region for this study due to its well documented connectivity arrangements (Lin et al., 2017) and description as a hub region (Andrews-Hanna et al., 2010). A previous study by our group on transient brain activity also revealed the PCC as the node with the most spatial overlap between networks (Karahanoglu & Van De Ville, 2015). We used SPM8's Check Orthogonality tool to understand how collinear the PCC activity was with our task paradigm, MATLAB's Skewness function to check that the activity was not skewed, and we verified that the magnitude of the seed time course did not correlate with the sign of the task (Supplementary Material A.2). To keep in line with previous work that inspired our method (Liu and Duyn, 2013), we report the thresholding step based on the percentage of data points kept for the dynamic analysis. As discussed later, PPI-CAPs is robust to a very wide range of seed activity thresholds for selecting frames to retain for analysis. For the results presented here, we use 60% of the available frames as a trade-off between optimising data usage from our dataset and obtaining clear brain patterns in the dynamic analysis, avoiding noise that is not averaged out due to a lack of frames in some PPI-CAPs.

# 3. Results

# 3.1. Static analysis

Single subject level. At a single subject level, the spatial pattern from the corresponding siMap, generated as the average of the 60% suprathreshold fMRI frames (when the PCC seed was highly active), was reasonably correlated with the resulting map from a 1<sup>st-</sup>level PPI Analysis (average spatial correlation:  $r = 0.76 \pm 0.28$ ). This was the threshold we chose in order to keep enough frames for the dynamic analysis, but the similarity was robust to the choice of threshold for frame selection across subjects (Fig. 2C): correlation was already high even when only 15% of suprathreshold frames were kept ( $r = 0.71 \pm 0.1$ ). To further illustrate that the correlation strength remains stable for varying cutoff values, we first calculated the mean subject-level Fisher's z-transformed correlations (see section 2.2) for each threshold that kept from 15 to 60% of original frames. Then, we computed the average of these means and transformed it back to Pearson's r, for which we obtained 0.76 with SD $\pm$ 0.04. The correlation gradually drops when fewer than 15% of the frames are used, which is expected as not enough frames are selected to average out the noise.

*Group level.* At the group level, the spatial pattern from the siMap obtained using 60% of frames also correlated (r = 0.85) with the pattern obtained from a 2<sup>nd</sup>-level PPI analysis (Supplementary Fig. A3), which showed a significant increase in effective connectivity between the PCC and the right V5 during *fun* scenes (height threshold T = 3.8, p < 0.001; right V5 MNI coordinates x = 39, y = -67, z = -2; cluster size = 185 voxels; p<sub>fwe-corr</sub> < 0.001).

Together, these results demonstrate that even a subset of fMRI frames, selected when the seed activity is highly active or strongly deactivated, contains relevant information about how its co-activation with other regions changes based on task context.

# 3.2. Dynamic analysis

*Choice of number of clusters.* An analysis of PPI-CAP occurrences for k = 3, 4, ..., 8 showed that for  $k \ge 6$ , some patterns never occurred in some subjects, while k = 4 and k = 5 were the cases with the most homogeneous distribution of pattern occurrence per subject (see Supplementary Fig. A4). Visual inspection of the consensus matrices showed that the most stable values for k (i.e., the values for which any two frames would most consistently be clustered together or separately) were k = 3

<sup>&</sup>lt;sup>1</sup> https://miplab.epfl.ch/index.php/miplife/research/supplement-asd-study.

and k = 4. Taking these observations together, we proceeded with the analysis generating 4 clusters, as this k value combined the beneficial features of: 1) yielding PPI-CAPs that are homogeneously distributed across subjects; 2) being highly stable (i.e., running the clustering several times would always produce similar results); and 3) representing a reasonable balance between variety and redundancy (Liu and Duyn, 2013).

Temporal decomposition of psychophysiological interactions into coactivation patterns. Our dynamic analysis revealed four recurring patterns of co-activation (Fig. 3A), all of which were significantly modulated by the seed, the context or an interaction between the two (PPI effect) after correcting for multiple comparisons (significance level  $\alpha$  0.05/4 = 0.0125). PPI-CAP1 includes nodes of the visuospatial (VSN) and attention (AN) networks correlated with the PCC and nodes of the fronto-parietal network (FPN) and salience networks (SN) anti-correlated with the PCC (p = 0.008). Additionally, the VSN and AN had a tendency to be more active, while the FPN and SN were more deactive, during science scenes as compared to fun ones (p = 0.019). PPI-CAP<sub>2</sub> combines the FPN correlated, plus the posterior insula and visual nodes anti-correlated, with the seed (p = 0.0009). PPI-CAP<sub>3</sub> corresponds to the default mode network (DMN) and was significantly correlated with the seed (p = 0.0009). It also appeared more often during fun scenes rather than science ones (p =0.008). PPI-CAP<sub>4</sub>, in turn, which contains the V5 and nodes of the VSN, showed not only a significant seed effect (p = 0.0009), but also a PPI effect (p = 0.0009). Supplementary Fig. A6 shows the exact times where each PPI-CAP appeared more consistently in their positive-or negative (meaning that the signs of the voxels should be flipped)-configuration across subjects, and Supplementary Fig. A5 shows their corresponding permutation histograms.

Consistency of PPI-CAPs across subjects and brain activity decoding. Our time-resolved method allowed us to investigate how consistent each PPI-CAP was across subjects and throughout the experiment. Fig. 4 illustrates this analysis for PPI-CAP<sub>4</sub>. By inspecting moments when the PPI-CAP appeared consistently for several subjects, we were able to identify specific video frames that elicited the co-activation pattern. For PPI-CAP<sub>4</sub>, moments of positive activation, that is, moments when the V5 and visuospatial network were more active, corresponded to scenes in which there was high motion (e.g., a group of children playing football), whereas a negative polarity of this pattern was related to moments of stillness (Fig. 4). Supplementary Figure A6 shows these results for all four PPI-CAPs: for PPI-CAP<sub>1</sub>, moments of positive polarity corresponded to scenes in which the main object (or character) was zoomed into and movements were slow, whereas a negative polarity (meaning a negative VSN and AN, with activated FPN and SN) was related to moments when some goal-oriented action was being performed. For example, at 1'38", while a group of children are sat at the beach, one of the girls is clearly reaching out for sand to build her castle. PPI-CAP2 seemed to be more strongly active during moments when there are lots of people on the scene, and more consistently negative (i.e., deactivated FPN and activated posterior insula) when the scene changes to only one person on screen, explaining concepts about the danger of sun exposure. PPI-CAP<sub>3</sub> appeared more with a positive polarity during zoomed-out scenes where many people were present and interacting, or when science concepts had been explained for a while in the laboratory scenario. These results illustrate PPI-CAPs' ability to link a pattern's occurrence to specific moments of the experimental paradigm.

# 4. Discussion

*PPI-CAPs: a tool to more accurately reveal task-based brain dynamics.* For decades, parametric statistical methods have been used for the analysis of task fMRI data (Friston et al., 1994; Eklund et al., 2016). Notably, the information revealed by the family of PPI approaches, where statistical analysis is performed on fMRI signal time courses to extract brain locations with context-dependent seed correlation, has greatly expanded our knowledge of brain function in health and disease in that time. For instance, Decety et al. (2008) investigated how healthy children experience empathy and moral reasoning when they view someone in pain.



Fig. 3. PPI-CAPs reveal patterns of coactivation that have a seed; task; or an interaction effect. A) We retrieved four PPI-CAPs from the Movie Watching dataset by clustering suprathreshold frames: PPI-CAP1 includes activated visuospatial and attention networks and deactive nodes of the frontoparietal network (FPN) and salience networks; PPI-CAP2 includes an activated FPN plus deactive posterior insula and visual nodes; PPI-CAP3 corresponds to nodes of the default mode network: and PPI-CAP<sub>4</sub> contains the V5 and visuospatial network. B) Confusion matrices depict how closely the polarity of the frames that make up each PPI-CAP relate to the sign of each effect. A clear diagonal (or anti-diagonal) pattern indicates a strong effect. All PPI-CAPs show a significant seed effect, PPI-CAP3 shows a significant task effect, and PPI-CAP4 shows a significant PPI effect, after Bonferroni correction for the number of PPI-CAPs. Raw p-values reported below the corresponding confusion matrices.



**Fig. 4. PPI-CAPs enables the time-resolved analysis of effective connectivity consistency across subjects.** For PPI-CAP<sub>4</sub>: (Top) The plot shows the percentage out of all 15 participants (y-axis) of subjects for whom this PPI-CAP was present at each fMRI frame (x-axis). Red and blue bars indicate the percentage for appearances in positive and negative configurations, respectively. Orange and blue background patches correspond to *fun* and *science* scenes, respectively. The histogram on the rightmost side of the top panel depicts the distribution of random consistency across subjects for each polarity, calculated by randomly permuting the time points on which the PPI-CAP appeared for each subject, then re-calculating the subject consistency across time. The dashed line represents the value of the 99% percentile from the random distribution, indicating that a subject consistency above this threshold is significant. Four sample time points of high consistency were highlighted. Numbers above the brain slices correspond to MNI coordinates. (Bottom) Video frames spanning a duration of 2 s, centred at each frame which corresponds to the timing of high consistency indicated in the top plot. For other PPI-CAPs see Supplementary Figure A6.

The connectivity observed in areas consistently engaged in moral behaviour and social interaction depended highly on intention and on whether the pain was self-inflicted or not, providing an empirical framework for studies of social cognition disorders in children. Steuwe et al. et al. (2015) showed that subcortical limbic and frontal loci become more connected to the locus coeruleus in female post-traumatic stress disorder patients when facing direct eye contact rather than averted gaze, potentially indicating an innate alarm system. More recently, a PPI analysis revealed that music intervention for preterm-born babies in neonatal intensive care units induces functional connectivity changes which suggest that music induces a more arousing and pleasant state (Lordier et al., 2018).

Meanwhile, point process analyses (Tagliazucchi et al., 2012) have proven to be a powerful tool for the study of resting-state brain data, by showing that large-scale brain activity could be condensed by solely analysing the time points when seed activity exceeds a given threshold, while nonetheless closely approximating seed-based correlation findings

#### (Liu et al., 2013).

In this work, we brought the advantages of these two methodologies together to expand the analysis of task-based recordings. Indeed, we showed that after modulation by the centred contrast variable, averaging as few as 15% fMRI frames with the strongest absolute seed activity yielded a spatial interaction pattern resembling the conventional PPI map calculated from the whole example dataset (Fig. 2C).

Beyond this initial equivalence as a first sanity check of the approach, we showed that functional brain connectivity across task conditions could be disentangled into a set of distinct building blocks, the PPI-CAPs, in analogy to the resting-state CAP methodology (Liu and Duyn, 2013; Liu et al., 2018). Whereas a traditional PPI map reflects, *on average*, the functional seed interplays that differ across task conditions, PPI-CAPs break this information down into separate seed co-activation patterns with their own spatiotemporal features. In addition, our approach also extends the information provided by traditional CAPs by capturing interaction effects between the task and the relationship of a

co-activation pattern with the seed.

Of note is the fact that the retrieved PPI-CAPs that do not show a context or PPI effect highlight shared functions between task settings, thus displaying common temporal expression features (i.e., similar occurrence levels across conditions). Statistically probing for a context effect enables, then, to distinguish these from co-activation patterns that are indeed task-modulated. This information cannot be obtained from previous PPI approaches. Further, extracting PPI-CAPs also enables the analyst to overcome the caveats arising from multiple comparisons that plague PPI analysis. This is because rather than mass univariate testing, K-means++ clustering (a multivariate, unsupervised technique) is performed to establish characteristic activation patterns, and only then followed by statistical testing on a much lower dimensional space.

Our methodological pipeline also yields polarity labels for each frame that contributes to a PPI-CAP, revealing moments among which a certain pattern has opposite signs-that is, regions that are highly activated in a frame labelled as "positive" will be deactivated in a "negative" frame of that PPI-CAP, and vice-versa. Since a given pattern may contain activated and deactivated voxels simultaneously, this also provides subregionspecific information on their relationship with the seed, task, or both, at any point in time. For instance, for a PPI-CAP whose frames' polarities correlate with the polarity of the seed (i.e., the confusion matrix with the seed is clearly diagonal), both jointly show their strongest signal values at the same time in all positively labelled frames. Accordingly, areas that appear in blue on that same PPI-CAP represent areas of seed-to-voxel anti-correlation when frames are labelled as positive, and vice-versa. Conversely, when a PPI-CAP has, for instance, a seed effect with the corresponding confusion matrix's pattern being anti-diagonal, we know that voxels depicted as red are, in fact, discordant with the seed at time points when a negative frame appears. This is why the interpretation of a PPI-CAP requires both the visualization of the co-activation map and the confusion matrix that shows how the polarity labels of its contributing frames relate to each effect's time course.

In the current version of our method the appearance of a PPI-CAP is defined only by averaging constituting frames, whose labels were assigned automatically during the clustering step. Since the first centroid initialization is made at random by K-means++ clustering, the polarity labels and, consequently, the pattern of the PPI-CAP they constitute, may be inverted in different runs. If preferred, this could be changed as a postprocessing step so that, if desired, the default pattern would be the one with the largest polarity. For example, if a PPI-CAP is made up mainly of negative frames, its flipped version could become the default (meaning that the voxels that currently appear in the pattern as deactive would be shown as active). In this case, the frame polarity labels would therefore be updated for their distribution to be correctly represented in the plots. Alternatively, polarity labels (as well as the pattern formed by the corresponding frames) could be arranged to default to the arrangement that favours a diagonal confusion matrix with the seed, to facilitate the interpretation of voxels in red or blue in relation to that region when there is a seed effect. Yet another option would be for the appearance of a PPI-CAP to default to the one that contains the most activated voxels (with frame labels being updated when necessary). Note, however, that none of these choices affects the interpretation of the results, as switching the labels will also invert the appearance of the PPI-CAP when it is updated as per Equation (4), so in the end the conclusions are unchanged; these suggestions are merely for visualization purposes.

*Neuroscientific relevance.* The above touches upon the striking complexity of functional brain activation during a task (Simony et al., 2016; Bolton et al., 2018b; Gonzalez-Castillo and Bandettini, 2018): in this, although standard PPI analysis already provides valuable insight into brain function at the cross-condition level (Kucyi et al., 2016), appropriately capturing truly occurring activity requires the deployment of better temporal resolution approaches as introduced here. Along this line of reasoning, none of the PPI-CAPs extracted in the present study strongly correlated with the results of a 2<sup>nd</sup>-level PPI analysis (rppI-CAP1 = 0.33; rppI-CAP2 = -0.07; rppI-CAP3 = 0.06; rppI-CAP4 = 0.31). The diversity

of the patterns revealed by PPI-CAPs in this work confirms the heterogeneity of PCC connectivity to large scale networks seen in previous studies (Liu and Duyn, 2013; Karahanoglu & Van De Ville, 2015), and their occurrence counts reveal how some of these relationships are modulated according to task. This suggests that the results from a conventional PPI analysis may yield a distorted picture of modulated activity as it likely never occurs as presented in the resulting map, and highlights PPI-CAPs' ability to reveal this effect more accurately.

The ability to characterise functional brain changes at the single frame level also offers the advantage to tie PPI-CAPs to specific subparts of the analysed task paradigm (see Spiers and Maguire, 2007 for a more general review on this analytical direction): in the results illustrated here, PPI-CAP<sub>4</sub> (visuospatial network) occurred upon strong movement during the displayed movie (Fig. 4), and all four patterns showed consistent, time-locked expression across subjects at specific time points (Supplementary Figure A6). The vivid and homogeneous recruitment of posterior motion processing areas actually squares well with previous findings considering the same dataset from another methodological angle (Bolton et al., 2018a).

Recently, the fMRI research community has striven to improve robustness and reproducibility by creating large-scale data acquisition and sharing initiatives (Van Horn and Toga, 2014; Van Essen et al., 2013; Poldrack et al., 2013), which brought along novel issues regarding data analysis (Smith and Nichols, 2018; Choudhury et al., 2014). This is because alongside the increase in data size, methods to analyse brain data have become increasingly complex. This combination may make some analyses simply computationally infeasible. It is thus an advantage for new analytical methods to obtain robust results even if running only on a portion of data, thereby circumventing computational cost issues that could make them impractical. From a researcher's point of view, even a linear decrease in computation time is significant, as a reduction of analysis time from weeks to several days can be decisive when trying to meet deadlines. This further highlights the importance of PPI-CAPs' efficient use of brain data, which is achieved by selecting a subset that already contains the relevant information needed to uncover our method's novel results.

Potential extensions of the PPI-CAPs approach. In the current methodology, PPI-CAPs are derived through a K-means clustering step, and are thus mutually exclusive in time. Previous work has already shown the merit of considering separate brain states to accurately describe taskbased data (Leonardi et al., 2014; Gonzalez-Castillo et al., 2015). However, a possible avenue to be explored would be to disentangle patterns that may overlap in time. An interesting option to achieve this would be to translate a recent total variation framework tailored to fMRI data (Karahanoglu & Van De Ville, 2015) to the task-based setting. As part of this new approach, the deconvolution method we currently apply would be replaced by Total Activation (TA) (Karahanoglu et al., 2013), which deconvolves BOLD signals and the haemodynamic response function using spatio-temporal regularisation to recover activity-inducing signals. These signals more closely reflect true neuronal activation than the indirect and noisy BOLD signals. Still based on Karahanoglu et al.'s work, by differentiating those, innovation signals are obtained-which reflect changes in activation intensity rather than pure amplitude. The development of our method would thus be to apply the frame selection and clustering steps on these signals to yield innovation-driven PPI-CAPs, which would represent spatial patterns of voxels whose signals transition simultaneously. Backprojecting these would then recover their time courses, thus revealing moments when different combinations of those patterns may overlap.

Another attractive avenue would be to consider the introduction of temporal relationships between successive time points, a feat that can be achieved both when considering sequential (Eavani et al., 2013; Vidaurre et al., 2017; Chen et al., 2016) or overlapping (Sourty et al., 2016; Bolton et al., 2018b) brain states. For instance, given that the present results revealed default mode network (PPI-CAP<sub>3</sub>), fronto-parietal network (PPI-CAP<sub>1</sub> and PPI-CAP<sub>2</sub>) and salience network (PPI-CAP<sub>2</sub>) contributions

during movie-watching, causal interplays across these networks could be assessed in accordance with the so called triple network model (Menon, 2011).

Aside from addressing temporal dynamics, an equally important issue is to optimally tackle the spatial dimension of the data. One extension could be the injection of a spatial prior in deriving PPI-CAPs (Zhuang et al., 2018). Another interesting aspect could be to study more closely the spatial variability of task-related functional activity patterns (Kiviniemi et al., 2011). A possible direction for this purpose could be to separately consider, for each PPI-CAP, the pools of frames linked to given task contexts, and carry out statistical comparisons at this level (Amico et al., 2014). Subtle spatial differences across task setting could then be revealed. However, this aspect also depends on the number of clusters used in the analysis: at larger values of K, different spatial patterns across contexts may rather be seen as context-specific PPI-CAPs. Here, given the relatively low amount of available data and the clearly optimal choice of K = 4 (Supplementary Figure A4), we did not pursue this side of the analyses.

Yet other extensions could be, on the one hand, to consider more sophisticated measures than PPI-CAP occurrences as features of interest (Chen and Glover, 2015) and, on the other hand, to broaden the analysis of PPI-CAPs to a meta-state perspective (Miller et al., 2016; Vidaurre et al., 2017), where a meta-state would symbolise a particular combination of expression and polarity of the investigated patterns.

Ultimately, the goal is of course the application of the developed tools to better understand both brain functions in healthy individuals, but also dysfunction in the case of neurological pathologies. Our approach enables us to easily address this last point, by adding a group factor to the employed nonparametric statistical assessment, enabling at the same time to gain insight into which features (seed, task, interaction) relate to the studied disease.

*Study considerations*. A natural limitation of this method is a consequence of the type of data: the spatiotemporal resolution can only be as good as that of the fMRI data used. The threshold for seed activation is the main free parameter in PPI-CAPs, but we showed that, as has been confirmed in other independent studies that followed a PPA approach (Liu et al., 2013; Tagliazucchi et al., 2012), the choice of this value does not affect the results within a very wide range of options. If the magnitude of the seed activity is not correlated with the sign of the task, then the resulting siMap will be proportional to the results of a PPI analysis, and its interpretation can follow the same guidelines as for PPI (see Friston et al., 1997 Fig. 5).

Recent work by Cole et al. (2019) has highlighted some challenges related to analysis of task-based functional connectivity. The authors find that simply describing the task paradigm as a dedicated regressor convolved with a canonical HRF, such as in PPI analysis, leaves a relatively large amount of false positives in the data. This is partly explained by the facts that the actual HRF shape varies across regions, and that task-related increases in activation will not necessarily always have the same amplitude, leaving residual activity in the data despite the use of a task regressor. In our work, we consider data that is deconvolved, a step for which we assume a canonical HRF shape. Given the impact of HRF variability on task-based analyses, and the advantage offered by approaches in which the HRF can be modeled individually across regions, future work should enable the use of similar strategies in a deconvolution setting. A second point made by the authors relates to the possible impact of differences in task-evoked activation amplitude across epochs. In our case, while we consider a naturalistic paradigm, we delineate fun and science task sub-blocks. There is thus the risk that during one type of block, task-based activation takes varying amplitudes, which would not be accounted for in the modelling of the task. Note that we do not explicitly rely on amplitude information in deriving and interpreting PPI-CAPs: rather, we examine how much the expression of a PPI-CAP (the polarity that it takes) across frames is in line with that of the task paradigm, the seed paradigm, or the PPI. A whole-brain pattern seen in a PPI-CAP thus may jointly represent those three effects. In any case, future work should keep the points above in mind.

# 5. Conclusion

We presented a novel analysis that temporally decomposes taskmodulated functional connectivity into dynamic building blocks which cannot be captured by static methods such as PPI analysis. We demonstrated that the PPI-CAPs approach successfully identifies dynamic taskdependent patterns using only a subset of the available data, which will lead to a linear decrease in computation time for large datasets proportional to the reduction in data size. Moreover, we illustrated how our method can be used to analyse brain activity at a resolution as low as the scanner's repetition time. Finally, we indicated how our method can expand other existing techniques and proposed new avenues for future research. Taken together, these show that our approach provides a more accurate picture of brain activity during task performance.

# Funding

This research was supported by the Swiss National Science Foundation (SNF): Switzerland [grant number 320030B\_182832]. It was also supported in part by the Bertarelli Foundation: Swizerland (to TB and DVDV); the Center for Biomedical Imaging (CIBM): Switerland; and the National Agency for Research: Switzerland (tempofront grant number 04701 to ALG).

## Appendix A. Link between conventional PPI and the stationary PPI-CAP

Within the PPA/CAPs framework, it is well known that averaging the frames where the seed exceeds a "well chosen" threshold  $\varphi$ , yields a proxy for the conventional seed connectivity map (Tagliazucchi et al., 2012; Liu and Duyn, 2013). This observation can be explained in the following way: let us assume that the activity time course *S*(*t*) of a seed voxel is a realization over time, t = 1, ..., N, of a random variable *S* for which E[S] = 0 and  $E[S^2] = 1$ . To construct conventional seed connectivity, a general linear model (GLM) is put forward to explain the time course *V*(*t*) of a target voxel as:

$$V(t) = \beta S(t) + \varepsilon(t),$$

(A.1)

where  $\beta$  is the parameter weight (i.e., functional connectivity) and  $\varepsilon(t)$  is the residual and assumed to be a realization of the noise with  $E[\varepsilon] = 0$  and  $E[\varepsilon^2] = \sigma^2$ . In addition, we assume the noise to be independent from the seed voxel activity; i.e.,  $E[S \cdot \varepsilon] = 0$ . Multiplying both sides of Eq. (A.1) with S(t) and taking the expectation leads to:

$$E[V \cdot S] = \beta E[S^2] + E[S \cdot \varepsilon], \tag{A.2}$$

which simplifies into  $\beta = E[V \cdot S]$ .

To obtain the stationary map in the CAPs approach, we want the expectation value of the target voxel for frames where the seed exceeds a threshold

L.G.A. Freitas et al.

$$I_{CAP+} = E[V:S > \varphi]$$
(A.3)

$$= E[\beta S + \varepsilon : S > \varphi] \tag{A.4}$$

$$=\beta E[S:S>\varphi] + E[\varepsilon:S>\varphi],$$
(A.5)

where in the first step we have substituted the previous GLM of Eq. (A.1) and ":" implies conditioning. Due to the independence of noise and seed voxel, and  $E[\varepsilon] = 0$ , the second term of Eq. (A.5) will vanish, and we obtain:

$$I_{\text{CAP+}} = \beta E[S:S > \varphi], \tag{A.6}$$

which means that the expectation of suprathreshold frames in the PPA method will be proportional to the result of a conventional correlation map, with the constant of proportionality given by  $E[S|S > \varphi]$ , which only depends on the seed, and not the target voxel. A similar relationship can be derived when the seed activity is below the negative threshold:

$$I_{\text{CAP}-} = E[V: S < -\varphi] \tag{A.7}$$

$$= E[\beta S + \varepsilon : S < -\phi] \tag{A.8}$$

$$=\beta E[S:S<-\varphi]+E[\varepsilon:S<-\varphi]$$
(A.9)

$$=\beta E[S:S<-\varphi] \tag{A.10}$$

In the case of a two-sided threshold, using the identity sign(A)A = |A|, we can conclude that:

$I_{CAP} = E[\operatorname{sign}(S)V :  S  > \varphi]$	(A.11)
$= E[\beta S \operatorname{sign}(S) + \varepsilon \operatorname{sign}(S) :  S  > \varphi]$	(A.12)
$=\beta E[S \operatorname{sign}(S): S  > \varphi] + E[\varepsilon \operatorname{sign}(S): S  > \varphi]$	(A.13)
$=\beta E[ S  :  S  > \varphi].$	(A.14)

Based on this observation, we can generalize to PPI for which the GLM becomes:

$$V(t) = \beta_S S(t) + \beta_T T(t) + \beta_{PPI} P(t) + \varepsilon(t), \tag{A.15}$$

where T(t) is the time course of the task and P(t) is the "interaction term" given by P(t) = S(t)T(t). The stationary map of the PPI-CAPs analysis then averages values of the target voxel multiplied with the sign of the interaction term, where the selection is based on the criterion of the seed exceeding the threshold:

$$\begin{split} I_{\text{PPI-CAP}} &= E[\text{sign}(ST)V : |S| > \varphi] \\ &= E[\text{sign}(ST)(\beta_{S}S + \beta_{T}T + \beta_{PPI}P + \varepsilon) : |S| > \varphi] \end{split}$$

 $= \beta_{S} E[\operatorname{sign}(ST)S: |S| > \varphi] + \beta_{T} E[\operatorname{sign}(ST)T: |S| > \varphi] + \beta_{PPI} E[\operatorname{sign}(ST)P: |S| > \varphi].$ 

Using the identities sign(A)A = |A| and sign(AB)A = sign(B)|A|, this becomes:

$$\begin{split} I_{\text{PPI-CAP}} &= \beta_S E[\text{sign}(T)|S| : |S| > \varphi] + \beta_T E[\text{sign}(S)|T| : |S| > \varphi] + \beta_{PPI} E[|ST| : |S| > \varphi] \\ &\approx \beta_{PPI} E[|ST| : |S| > \varphi]. \end{split}$$

The last approximation can be made when the contributions of the first two terms are negligible. For the first term, we assume the absolute value of *S* to be independent of the task (that weights the expectation operator) and *T* has equal number of positive and negative values, in which case the first term will be 0. For the second term, if the task is just a change of sign, then its absolute value will be fixed and uncorrelated to the sign of the seed. Thus the second term is proportional to  $E[sign(S) : |S| > \varphi]$ , which is necessarily bound between  $\pm 1$  and will be exactly 0 if S is symmetric.

Therefore, what remains is the final term and we find that the result of a PPI analysis will be proportional to the SiMap. This brief derivation also shows that the stationary PPI-CAP can be contaminated by "leakage" from the seed and/or task contributions if these assumptions are not fulfilled.

# Appendix B. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2020.116635.

#### L.G.A. Freitas et al.

#### References

- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. Cerebr. Cortex 24, 663–676. https://doi.org/10.1093/cercor/bhs352.
- Book of Gomez, F., Di Perri, C., Vanhaudenhuyse, A., Lesenfants, D., Boveroux, P., Bonhomme, V., Brichant, J.F., Marinazzo, D., Laureys, S., 2014. Posterior cingulate cortex-related co-activation patterns: a resting state fMRI study in propofol-induced loss of consciousness. PloS One 9, 1–9. https://doi.org/10.1371/ journal.pone.0100012.
- Andrews-Hanna, J.R., Reidler, J.S., Sepulcre, J., Poulin, R., Buckner, R.L., 2010. Functional-anatomic fractionation of the brain's default network. Neuron 65, 550–562. https://doi.org/10.1016/j.neuron.2010.02.005.
- Arthur, D., Vassilvitskii, S., 2007. K-Means ++ : the advantages of careful seeding. Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms 1027–1035.
- Baczkowski, B.M., Johnstone, T., Walter, H., Erk, S., Veer, I.M., 2017. Sliding-window analysis tracks fluctuations in amygdala functional connectivity associated with physiological arousal and vigilance during fear conditioning. Neuroimage 153, 168–178. https://doi.org/10.1016/j.neuroimage.2017.03.022.
- Bolton, T.A., Jochaut, D., Giraud, A.L., Van De Ville, D., 2018a. Brain dynamics in ASD during movie-watching show idiosyncratic functional integration and segregation. Hum. Brain Mapp. 39, 2391–2404. https://doi.org/10.1002/hbm.24009.
- Bolton, T.A.W., Tarun, A., Sterpenich, V., Schwartz, S., Van De Ville, D., 2018b. Interactions between large-scale functional brain networks are captured by sparse coupled HMMs. IEEE Trans. Med. Imag. 37, 230–240. https://doi.org/10.1109/ TMI.2017.2755369.
- Caballero Gaudes, C., Petridou, N., Francis, S.T., Dryden, I.L., Gowland, P.A., 2011. Paradigm free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. Hum. Brain Mapp. https://doi.org/10.1002/hbm.21452 n/a–n/a.
- Calhoun, V.D., Miller, R., Pearlson, G., Adal, T., 2014. The chronnectome : time-varying connectivity networks as the next frontier in fMRI data discovery, 84, 262–274. https://doi.org/10.1016/j.neuron.2014.10.015.
- Cao, X., Sandstede, B., Luo, X., 2019. A Functional Data Method for Causal Dynamic Network Modeling of Task-Related fMRI, vol. 13, pp. 1–19. https://doi.org/10.3389/ fnins.2019.00127.
- Chang, C., Glover, G.H., 2010. Time–frequency dynamics of resting-state brain connectivity measured with fMRI. Neuroimage 50, 81–98. https://doi.org/10.1016/ j.neuroimage.2009.12.011.
- Chen, J.E., Glover, G.H., 2015. Functional Magnetic Resonance Imaging Methods 2, 289–313. https://doi.org/10.1007/s11065-015-9294-9.
- Chen, T., Cai, W., Ryali, S., Supekar, K., Menon, V., 2016. Distinct global brain dynamics and spatiotemporal organization of the salience network. PLoS Biol. 14, 1–21. https://doi.org/10.1371/journal.pbio.1002469.
- Choe, A.S., Nebel, M.B., Barber, A.D., Cohen, J.R., Xu, Y., Pekar, J.J., Caffo, B., Lindquist, M.A., 2017. Comparing test-retest reliability of dynamic functional connectivity methods. Neuroimage 158, 155–175. https://doi.org/10.1016/ j.neuroimage.2017.07.005.
- Choudhury, S., Fishman, J.R., McGowan, M.L., Juengst, E.T., 2014. Big data, open science and the brain: lessons learned from genomics. Front. Hum. Neurosci. 8, 1–10. https:// doi.org/10.3389/fnhum.2014.00239.
- Cole, M.W., Ito, T., Schultz, D., Mill, R., Chen, R., Cocuzza, C., 2019. Task activations produce spurious but systematic in fl ation of task functional connectivity estimates. Neuroimage 189, 1–18. https://doi.org/10.1016/j.neuroimage.2018.12.054.
- Cox, N.J., 2008. Speaking Stata: correlation with confidence, or Fisher's z revisited. STATA J. 8, 413–439. https://doi.org/10.1177/1536867x0800800307.
- Damaraju, E., Tagliazucchi, E., Laufs, H., Calhoun, V.D., 2018. Connectivity Dynamics from Wakefulness to Sleep. bioRxiv preprint. https://doi.org/10.1101/380741.
- Decety, J., Michalska, K.J., Akitsuki, Y., 2008. Who caused the pain? An fMRI investigation of empathy and intentionality in children. Neuropsychologia 46, 2607–2614.
- Di, X., Fu, Z., Chan, S.C., Hung, Y.S., Biswal, B.B., 2015. Task-related Functional Connectivity Dynamics in a Block-Designed Visual Experiment, vol. 9, pp. 1–11. https://doi.org/10.3389/fnhum.2015.00543.
- Eavani, H., Satterthwaite, T.D., Gur, R.E., Gur, R.C., Davatzikos, C., 2013. Unsupervised Learning of Functional Network Dynamics in Resting State fMRI. Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 7917, pp. 426–437. https://doi.org/10.1007/ 978-3-642-38868-2\_36. LNCS.
- Eklund, A., Nichols, T.E., Knutsson, H., 2016. Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. In: Proceedings of the National Academy of Sciences, p. 201602413.
- Fransson, P., Schiffler, B.C., Thompson, W.H., 2018. Brain network segregation and integration during an epoch-related working memory fMRI experiment. Neuroimage 178, 147–161. https://doi.org/10.1016/j.neuroimage.2018.05.040.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.-P., Frith, C.D., Frackowiak, R.S., 1994. Statistical parametric maps in functional imaging: a general linear approach. Hum. Brain Mapp. 2, 189–210.
- Friston, K., Buechel, C., Fink, G., Morris, J., Rolls, E., Dolan, R., 1997. Psychophysiological and modulatory interactions in neuroimaging. Neuroimage 6, 218–229. https:// doi.org/10.1006/nimg.1997.0291.
- Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19, 1273–1302. https://doi.org/10.1016/S1053-8119(03)00202-7.
- Ge, B., Wang, H., Wang, P., Tian, Y., Zhang, X., Liu, T., 2019. Discovering and Characterizing Dynamic Functional Brain Networks in Task FMRI.

- Gonzalez-Castillo, J., Bandettini, P.A., 2018. Task-based dynamic functional connectivity: recent findings and open questions. Neuroimage 180, 526–533. https://doi.org/ 10.1016/j.neuroimage.2017.08.006.
- Gonzalez-Castillo, J., Hoy, C.W., Handwerker, D.A., Robinson, M.E., Buchanan, L.C., Saad, Z.S., Bandettini, P.A., 2015. Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. Proc. Natl. Acad. Sci. Unit. States Am. 112, 8762–8767. https://doi.org/10.1073/pnas.1501242112.arXiv arXiv: 1408.1149.
- Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 2013. Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage 80, 360–378. https://doi.org/10.1016/ j.neuroimage.2013.05.079.arXiv:NIHMS150003.
- Jochaut, D., Lehongre, K., Saitovitch, A., Devauchelle, A.-D., Olasagasti, I., Chabane, N., Zilbovicius, M., Giraud, A.-L., 2015. Atypical coordination of cortical oscillations in response to speech in autism. Front. Hum. Neurosci. 9, 1–12. https://doi.org/ 10.3389/fnhum.2015.00171.
- Karahanoglu, F.I., Van De Ville, D., 2015. Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. Nat. Commun. 6, 7751. https://doi.org/10.1038/ncomms8751.
- Karahanoğlu, F.I., Van De Ville, D., 2017. Dynamics of large-scale fMRI networks: deconstruct brain activity to build better models of brain function. Current Opinion in Biomedical Engineering 28–36. https://doi.org/10.1016/j.cobme.2017.09.008.
- Karahanoglu, F.I., Caballero-Gaudes, C., Lazeyras, F., Van De Ville, D., 2013. Total activation: FMRI deconvolution through spatio-temporal regularization. Neuroimage 73, 121–134. https://doi.org/10.1016/j.neuroimage.2013.01.067.
- Kiviniemi, V., Vire, T., Remes, J., Elseoud, A.A., Starck, T., Tervonen, O., Nikkinen, J., 2011. A sliding time-window ICA reveals spatial variability of the default mode network in time. Brain Connect. 1, 339–347. https://doi.org/10.1089/ brain.2011.0036.
- Kucyi, A., Davis, K.D., 2014. Dynamic functional connectivity of the default mode network tracks daydreaming. Neuroimage 100, 471–480. https://doi.org/10.1016/ j.neuroimage.2014.06.044.
- Kucyi, A., Hove, M.J., Esterman, M., Hutchison, R.M., Valera, E.M., 2016. Dynamic brain network correlates of spontaneous fluctuations in attention. Cerebr. Cortex bhw029. https://doi.org/10.1093/cercor/bhw029.
- Leonardi, N., Van De Ville, D., 2015. On spurious and real fluctuations of dynamic functional connectivity during rest. Neuroimage 104, 430–436. https://doi.org/ 10.1016/j.neuroimage.2014.09.007.
- Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-m., Schluep, M., Vuilleumier, P., Ville, D.V.D., 2013. Principal components of functional connectivity : a new approach to study dynamic brain connectivity during rest. Neuroimage 83, 937–950. https://doi.org/10.1016/j.neuroimage.2013.07.019.
- Leonardi, N., Shirer, W.R., Greicius, M.D., Van De Ville, D., 2014. Disentangling dynamic networks: separated and joint expressions of functional connectivity patterns in time. Hum. Brain Mapp. 35, 5984–5995. https://doi.org/10.1002/hbm.22599.
- Lin, P., Yang, Y., Gao, J., Pisapia, N.D., Ge, S., Wang, X., Zuo, C.S., Levitt, J.J., Niu, C., 2017. Dynamic Default Mode Network across Different Brain States. Nature Publishing Group, pp. 1–13. https://doi.org/10.1038/srep46088.
- Lindquist, M.A., Xu, Y., Nebel, M.B., Caffo, B.S., 2014. Evaluating dynamic bivariate correlations in resting-state fMRI: a comparison study and a new approach. Neuroimage 101, 531–546. https://doi.org/10.1016/ i.neuroimage.2014.06.052.arXiv:NIHMS150003.
- Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proc. Natl. Acad. Sci. Unit. States Am. 110, 4392–4397. https://doi.org/10.1073/pnas.1216856110.arXiv arXiv: 1408.1149.
- Liu, X., Chang, C., Duyn, J.H., 2013. Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. Front. Syst. Neurosci. 7, 1–11. https://doi.org/ 10.3389/fnsys.2013.00101.
- Liu, X., Zhang, N., Chang, C., Duyn, J.H., 2018. Co-activation patterns in resting-state fMRI signals. Neuroimage 180, 485–494. https://doi.org/10.1016/ j.neuroimage.2018.01.041.
- Lordier, L., Loukas, S., Grouiller, F., Vollenweider, A., Vasung, L., Meskaldij, D., Lejeune, F., Pittet, M.P., Borradori-Tolsa, C., Lazeyras, F., et al., 2018. Music processing in preterm and full-term newborns: a psychophysiological interaction (PPI) approach in neonatal fMRI. Neuroimage 185, 857–864. https://doi.org/ 10.1016/j.neuroimage.2018.03.078.
- McLaren, D.G., Ries, M.L., Xu, G., Johnson, S.C., 2012. A generalized form of contextdependent psychophysiological interactions (gPPI): a comparison to standard approaches. Neuroimage 61, 1277–1286. https://doi.org/10.1016/ j.neuroimage.2012.03.068.arXiv:NIHMS150003.
- Menon, V., 2011. Large-scale brain networks and psychopathology: a unifying triple network model. Trends Cognit. Sci. 15, 483–506. https://doi.org/10.1016/ j.tics.2011.08.003 arXiv:NIHMS150003.
- Miller, R.L., Yaesoubi, M., Turner, J.A., Mathalon, D., Preda, A., Pearlson, G., Adali, T., Calhoun, V.D., 2016. Higher dimensional meta-state analysis reveals reduced resting fMRI connectivity dynamis in schizophrenia patients. PloS One 11, e0149849. https://doi.org/10.1371/journal.pone.0149849.
- Monti, S., Tamayo, P., Mesirov, J., Golub, T., 2003. Consensus clustering : a resamplingbased method for class discovery and visualization of gene expression microarray data. Mach. Learn. 52, 91–118.
- O'Reilly, J.X., Woolrich, M.W., Behrens, T.E.J., Smith, S.M., Johansen-Berg, H., 2012. Tools of the trade: psychophysiological interactions and functional connectivity. Soc. Cognit. Affect Neurosci. 7, 604–609. https://doi.org/10.1093/scan/nss055.

#### L.G.A. Freitas et al.

Ogawa, S., Lee, T.M., Kay, A.R., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. Unit. States Am. 87, 9868–9872.

- Petridou, N., Gaudes, C.C., Dryden, I.L., Francis, S.T., Gowland, P.A., 2013. Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. Hum. Brain Mapp. 34, 1319–1329. https://doi.org/10.1002/ hbm.21513.
- Ploner, M., Schoffelen, J.M., Schnitzler, A., Gross, J., 2009. Functional integration within the human pain system as revealed by Granger causality. Hum. Brain Mapp. 30, 4025–4032. https://doi.org/10.1002/hbm.20826.
- Poldrack, R.A., Barch, D.M., Mitchell, J.P., Wager, T.D., Wagner, A.D., Devlin, J.T., Cumba, C., Koyejo, O., Milham, M.P., 2013. Toward open sharing of task-based fMRI data: the OpenfMRI project. Front. Neuroinf. 7, 1–12. https://doi.org/10.3389/ fninf.2013.00012.
- Power, J.D., Fair, D.A., Schlaggar, B.L., Petersen, S.E., 2010. The development of human functional brain networks. Neuron 67, 735–748. https://doi.org/10.1016/ i.neuron.2010.08.017.
- Preti, M.G., Bolton, T., Van De Ville, D., 2016. The dynamic functional connectome : stateof-the-art and perspectives. Neuroimage. https://doi.org/10.1016/ j.neuroimage.2016.12.061.
- Rack-Gomer, A.L., Liu, T.T., 2012. Caffeine increases the temporal variability of restingstate BOLD connectivity in the motor cortex. Neuroimage 59, 2994–3002.
- Rissman, J., Gazzaley, A., Esposito, M.D., 2004. Measuring Functional Connectivity during Distinct Stages of a Cognitive Task, vol. 23, pp. 752–763. https://doi.org/ 10.1016/j.neuroimage.2004.06.035.
- Shen, H., Li, Z., Qin, J., Liu, Q., Wang, L., Zeng, L.L., Li, H., Hu, D., 2016. Changes in functional connectivity dynamics associated with vigilance network in taxi drivers. Neuroimage 124, 367–378. https://doi.org/10.1016/j.neuroimage.2015.09.010.
- Simony, E., Honey, C.J., Chen, J., Lositsky, O., Yeshurun, Y., Wiesel, A., Hasson, U., 2016. Dynamic reconfiguration of the default mode network during narrative comprehension. Nat. Commun. 7, 12141. https://doi.org/10.1038/ncomms12141.
- Smith, S.M., Nichols, T.E., 2018. Statistical challenges in "big data" human neuroimaging. Neuron 97, 263–268. https://doi.org/10.1016/j.neuron.2017.12.018.
- Sourty, M., Thoraval, L., Roquet, D., Armspach, J.-P., Foucher, J., Blanc, F., 2016. Identifying dynamic functional connectivity changes in dementia with lewy bodies based on product hidden markov models. Front. Comput. Neurosci. 10, 1–11. https:// doi.org/10.3389/fncom.2016.00060.
- Spiers, H.J., Maguire, E.A., 2007. A navigational guidance system in the human brain. Hippocampus 626, 618–626. https://doi.org/10.1002/hipo.

- Steuwe, C., Daniels, J.K., Frewen, P.A., Densmore, M., Theberge, J., Lanius, R.A., 2015. Effect of direct eye contact in women with PTSD related to interpersonal trauma: psychophysiological interaction analysis of connectivity of an innate alarm system. Psychiatr. Res. Neuroimaging 232, 162–167.
- Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D.R., 2012. Criticality in Large-Scale Brain fMRI Dynamics Unveiled by a Novel Point Process Analysis, vol. 3, pp. 1–12. https://doi.org/10.3389/fphys.2012.00015.
- van den Heuvel, M.P., Hulshoff Pol, H.E., 2010. Exploring the brain network: a review on resting-state fMRI functional connectivity. Eur. Neuropsychopharmacol 20, 519–534. https://doi.org/10.1016/j.euroneuro.2010.03.008.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., 2013. The Wu-minn human connectome project: an overview. Neuroimage 80, 62–79. https://doi.org/10.1016/j.neuroimage.2013.05.041.arXiv:NIHMS150003.
- Van Horn, J.D., Toga, A.W., 2014. Human neuroimaging as a "big data" science. Brain Imaging and Behavior 8, 323–331. https://doi.org/10.1007/s11682-013-9255-y.
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. Proc. Natl. Acad. Sci. Unit. States Am. 114, 201705120. https://doi.org/10.1073/pnas.1705120114.arXiv arXiv:1408.1149.
- Wen, X., Rangarajan, G., Ding, M., 2013. Is granger causality a viable technique for analyzing fMRI data? PloS One 8. https://doi.org/10.1371/journal.pone.0067428.
- Yaesoubi, M., Allen, E.A., Miller, R.L., Calhoun, V.D., 2015. Dynamic coherence analysis of resting fMRI data to jointly capture state-based phase , frequency , and timedomain information. Neuroimage 120, 133–142. https://doi.org/10.1016/ j.neuroimage.2015.07.002.
- Zalesky, A., Breakspear, M., 2015. Towards a statistical test for functional connectivity dynamics. Neuroimage 114, 466–470. https://doi.org/10.1016/ j.neuroimage.2015.03.047.
- Zhang, Z.G., Hu, L., Hung, Y.S., Mouraux, A., Iannetti, G.D., 2012. Gamma-band oscillations in the primary somatosensory cortex–A direct and obligatory correlate of subjective pain intensity. J. Neurosci. 32, 7429–7438. https://doi.org/10.1523/ JNEUROSCI.5877-11.2012.
- Zhuang, J., Peltier, S., He, S., LaConte, S., Hu, X., 2008. Mapping the connectivity with structural equation modeling in an fMRI study of shape-from-motion task. Neuroimage 42, 799–806. https://doi.org/10.1016/j.neuroimage.2008.05.036.
- Zhuang, X., Walsh, R.R., Sreenivasan, K., Yang, Z., Mishra, V., Cordes, D., 2018. Incorporating spatial constraint in co-activation pattern analysis to explore the dynamics of resting-state networks: an application to Parkinson's disease. Neuroimage 172, 64–84. https://doi.org/10.1016/j.neuroimage.2018.01.019.