



HAL
open science

TbCAPs: A toolbox for co-activation pattern analysis

Thomas A.W. Bolton, Constantin Tuleasca, Diana Wotruba, Gwladys Rey, Herberto Dhanis, Baptiste Gauthier, Farnaz Delavari, Elenor Morgenroth, Julian Gaviria, Eva Blondiaux, et al.

► **To cite this version:**

Thomas A.W. Bolton, Constantin Tuleasca, Diana Wotruba, Gwladys Rey, Herberto Dhanis, et al.. TbCAPs: A toolbox for co-activation pattern analysis. *NeuroImage*, 2020, 211, pp.116621. 10.1016/j.neuroimage.2020.116621 . hal-02538773

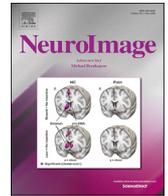
HAL Id: hal-02538773

<https://hal.sorbonne-universite.fr/hal-02538773>

Submitted on 9 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



TbCAPs: A toolbox for co-activation pattern analysis

Thomas A.W. Bolton^{a,b,*}, Constantin Tuleasca^{c,d,e,f}, Diana Wotruba^g, Gwladys Rey^h,
Herberto Dhanisⁱ, Baptiste Gauthierⁱ, Farnaz Delavari^j, Elenor Morgenroth^a, Julian Gaviria^h,
Eva Blondiauxⁱ, Lukasz Smigielskiⁱ, Dimitri Van De Ville^{a,b}

^a Institute of Bioengineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

^b Department of Radiology and Medical Informatics, University of Geneva (UNIGE), Geneva, Switzerland

^c Centre Hospitalier Universitaire Bicêtre, Service de Neurochirurgie, Paris, France

^d Sorbonne Université, Faculté de Médecine, Paris, France

^e Lausanne University (UNIL), Faculty of Biology and Medicine (FBM), Lausanne, Switzerland

^f Centre Hospitalier Universitaire Vaudois (CHUV), Neurosurgery Service and Gamma Knife Center, Lausanne, Switzerland

^g Program for Sustainable Development of Mental Health, University of Zurich, Zurich, Switzerland

^h Department of Neuroscience, University of Geneva (UNIGE), Geneva, Switzerland

ⁱ Laboratory of Cognitive Neuroscience, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

^j Developmental Imaging and Psychopathology Laboratory, Office Médico-Pédagogique, Department of Psychiatry, University of Geneva (UNIGE), Geneva, Switzerland

ARTICLE INFO

Keywords:

Dynamic functional connectivity
Frame-wise analysis
Co-activation pattern analysis
Task-positive network
Attention
Continuous performance
Open source software

ABSTRACT

Functional magnetic resonance imaging provides rich spatio-temporal data of human brain activity during task and rest. Many recent efforts have focussed on characterising dynamics of brain activity. One notable instance is co-activation pattern (CAP) analysis, a frame-wise analytical approach that disentangles the different functional brain networks interacting with a user-defined seed region. While promising applications in various clinical settings have been demonstrated, there is not yet any centralised, publicly accessible resource to facilitate the deployment of the technique.

Here, we release a working version of TbCAPs, a new toolbox for CAP analysis, which includes all steps of the analytical pipeline, introduces new methodological developments that build on already existing concepts, and enables a facilitated inspection of CAPs and resulting metrics of brain dynamics. The toolbox is available on a public academic repository at https://c4science.ch/source/CAP_Toolbox.git.

In addition, to illustrate the feasibility and usefulness of our pipeline, we describe an application to the study of human cognition. CAPs are constructed from resting-state fMRI using as seed the right dorsolateral prefrontal cortex, and, in a separate sample, we successfully predict a behavioural measure of continuous attentional performance from the metrics of CAP dynamics ($R = 0.59$).

1. Introduction

Functional magnetic resonance imaging (fMRI) has enabled to track temporal changes in activity levels at the whole-brain scale by means of the blood oxygenation level-dependent (BOLD) contrast, a proxy for neural activation (Logothetis et al., 2001). In addition to more traditional task-based studies in which BOLD changes are mapped to a paradigm of interest (Friston et al., 1994), the characterisation of statistical interdependence between remote brain locations—termed *functional connectivity* (Friston, 1994)—in the resting-state, and the concomitant definition of large-scale *resting-state networks* (RSNs), has been a popular endeavour

(Biswal et al., 1995; Fox et al., 2005; Damoiseaux et al., 2006; Power et al., 2011), with great benefits for the understanding of cognition and disease (van den Heuvel and Hulshoff Pol, 2010; Greicius, 2008; Fox and Greicius, 2010).

Over the past years, it has become increasingly appreciated that cross-regional relationships do not remain static over the course of a full scanning session (Chang and Glover, 2010): instead, a given region rearranges its interactions along time, in ways that have been addressed with very diverse analytical tools (see Hutchison et al. (2013); Preti et al. (2017) for exhaustive reviews of the *dynamic functional connectivity* field).

In one family of approaches that has been developed, it is assumed

* Corresponding author. Institute of Bioengineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

E-mail address: thomas.bolton@epfl.ch (T.A.W. Bolton).

<https://doi.org/10.1016/j.neuroimage.2020.116621>

Received 11 October 2019; Received in revised form 16 January 2020; Accepted 6 February 2020

Available online 10 February 2020

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

that only few salient time points contain the information of interest that shapes whole-brain correlational relationships; selecting only these frames, by means of a seed-based thresholding process, already enables to derive accurate RSN maps, even if as little as 10% of data points are retained (Tagliazucchi et al., 2012). The analysis then moves from a second-order correlation-based characterisation to a first-order activation viewpoint, and reduces computational load, a desirable feat in light of the numerous large-scale acquisition initiatives embraced by the fMRI community (Van Essen et al., 2013; Nooner et al., 2012; Holmes et al., 2015).

Building on this point process analysis concept, and inspired by the dynamic viewpoint on resting-state brain function, Liu and Duyn (2013) hypothesised that at different moments in time, the seed region of interest would display distinct interactions with the rest of the brain. A k-means clustering step was thus appended to frame selection, so that fMRI volumes with a large enough seed activity would be partitioned into a limited set of *co-activation patterns* (CAPs).

Since then, co-activation pattern analysis has started to gain momentum as a potent tool to reveal functional brain dynamics subtleties: analyses taking the posterior cingulate cortex (PCC) as a seed revealed alterations of spatial intensity level and occurrence in specific CAPs (Amico et al., 2014; Di Perri et al., 2018), while in adolescent depression, Kaiser et al. (2019) showed that the time spent in a specific frontoinsular-default network CAP positively correlated with symptoms severity. In other work, the renormalisation of CAP occurrences in patients with essential tremor following surgical intervention could be tracked (Tuleasca et al., 2019).

In parallel to clinical applications, the technical details of the approach have also been addressed, in terms of retaining activation *versus* deactivation time points (Di and Biswal, 2015), extending it to the whole brain (Liu et al., 2013), designing novel metrics of interest (Chen et al., 2015), or constraining the extent of spatial overlap across CAPs (Zhuang et al., 2018). For more details, the reader is pointed at the recent review of Liu et al. (2018).

Here, we wish to further foster the development of CAP analysis by releasing a dedicated toolbox, which enables to easily navigate through the steps of the analytical pipeline through a graphical user interface, and also offers additional technical developments regarding frame selection and metrics computation. While the mathematical underpinnings of CAP analysis are relatively straightforward, we hope that providing such a resource will encourage practitioners to embrace the method, and that it will become easier to compare CAP analyses based on subtle, but sometimes important, differences in the processing pipeline. Through this resource, we also aim at preventing the variability in analytical results that may otherwise arise due to implementation differences alone (Bowring et al., 2019).

In addition, to exemplify the use of our toolbox, we describe an application of CAP analysis in the yet unaddressed setting of predicting cognitive skills: in a battery of healthy individuals, we show that continuous performance in a visual attention and vigilance task correlates with the expression profile of task-positive network (TPN) CAPs.

2. Materials and methods

2.1. Co-activation pattern analysis theory

Let us consider the data matrix $\mathbf{X}_s \in \mathbb{R}^{V \times T}$ for subject s , where V is the number of voxels to consider in the analysis and T the number of time points. Each voxel-wise time course is temporally z-scored, so that $\mu_i = \frac{\sum_{t=1}^T X_s(i,t)}{T} = 0$ and $\sigma_i = \sqrt{\frac{\sum_{t=1}^T (X_s(i,t) - \mu_i)^2}{T-1}} = 1$, for all $i = 1, 2, \dots, V$.

Co-activation pattern analysis requires the definition of a seed region, whose interactions with the rest of the brain will be probed. Formally, a set of voxels \mathcal{S} that one wishes to consider is specified, and a time point t of the seed activation time course is then given by:

$$S_{\text{seed}}(t) = \frac{\sum_{i \in \mathcal{S}} X_s(i,t)}{|\mathcal{S}|}, \quad \text{for all } t \in 1, 2, \dots, T. \quad (1)$$

Only time points when the seed time course takes sufficiently extreme values (denoting significant seed (de)activation) are considered. Let the activation threshold be T , we then construct the subject-specific set of time points \mathcal{T}_s that satisfies $S_{\text{seed}}(t) > T$ (if we wish to consider solely activation moments) or $S_{\text{seed}}(t) < -T$ (if we are interested in seed deactivation time points).

In this work, in addition to the above standard CAP methodology, we propose an extension in which more than one seed region can be considered: for each seed j and subject s , a set of time points $\mathcal{T}_{s,j}$ is derived. Assuming J separate seeds, one can then consider the time points when all seed time courses jointly take extreme values:

$$\mathcal{T}_{s,\text{Intersection}} = \bigcap_{j=1}^J \mathcal{T}_{s,j}. \quad (2)$$

Alternatively, one may instead be interested in the moments when at least one of the seed regions becomes strongly (de)active:

$$\mathcal{T}_{s,\text{Union}} = \bigcup_{j=1}^J \mathcal{T}_{s,j}. \quad (3)$$

Finally, other additional criteria can be incorporated at the time point selection step: for instance, given the deleterious impact that head motion exerts on BOLD signals even following standard preprocessing (Power et al., 2012; Van Dijk et al., 2012; Satterthwaite et al., 2012), it may be desirable to only retain the frames for which framewise displacement does not exceed a threshold M .

After having selected the frames to keep for each subject, the next step is the population-level clustering of data points into CAPs. K-means clustering is used for this purpose, to optimise:

$$\underset{\mathcal{C}}{\text{argmin}} \sum_{k=1}^K \sum_{s=1}^S \sum_{t \in \mathcal{T}_s \cap \mathcal{C}_k} \text{dist}(\mathbf{X}_s(\cdot, t), \mathbf{c}_k), \quad (4)$$

where K is the number of co-activation patterns to derive, $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_K\}$ summarises the hard assignment of the frames to each CAP, and \mathbf{c}_k is the spatial map for co-activation pattern k . The dist function depends on the type of distance to use in the algorithm. In addition, since k-means clustering is an iterative process with no guaranteed convergence towards the global optimum, the algorithm is run n_{rep} times.

In several previous works using CAPs, it was also suggested to solve Equation (4) after setting to 0 the voxel intensity values that, for each frame of interest, would not be part of the largest P_P or P_N percents—for positive-valued and negative-valued voxels, respectively (Liu and Duyn, 2013; Liu et al., 2013).

Table 1 summarises the different parameters that are defined for CAP analysis, and also highlights the default values that we used in this work.

2.2. Metrics characterising CAP dynamics

Once all retained frames have been assigned to CAP representatives, it becomes possible to construct, for each subject, an empirical transition probability matrix \mathbf{A}_s that summarises the likelihood to transit from a given CAP at time t to another at time $t + 1$. Another available piece of information regards the likelihood to transit from and back to the baseline state (when the seed was not significantly (de)active). Further, if separate subject populations are used in computing CAPs and deriving associated metrics (as in our example application below), there are also occurrences of entries into an extra state associated to frames that could not be matched to any CAP with sufficient certainty.

An indicative example of averaged transition probability matrix across subjects is displayed in Fig. 1A (left column). Individual elements of the transition probability matrix may be considered as such (Chen et al., 2015), which would amount to a total of K^2 values per subject. To meaningfully lower the amount of features of interest, we propose to

Table 1

Parameters to define for co-activation pattern analysis. The seed region was extracted from a TPN independent component map derived in Shirer et al. (2012). Investigated cluster number values $\{K\}$ were determined through consensus clustering, a subsampling-based assessment of clustering robustness (Monti et al., 2003), and the final choice for the analyses was defined based on an exploratory assessment of the brain/behaviour correlation significance.

Parameter	Description	Default value
J	Number of seeds to use	1
\mathcal{S}	Voxel set to use as seed	Right dorsolateral prefrontal cortex
Polarity	Sign of the seed excursions to consider	Activation
Seed combination	Whether all or at least one seed should be (de)active to retain a time point	n.a.
T	Threshold for frame selection	1.5
M	Threshold of framewise displacement above which to scrub	0.3 mm
K	Number of clusters to use	16
n_{rep}	Number of replicates of k-means	50
P_p	Percentage of positive-valued voxels to keep in each frame for clustering	100
P_N	Percentage of negative-valued voxels to keep in each frame for clustering	100
dist	Distance measure used for clustering	corr

rather view the available information as a directional graph representation, from which a series of summarising metrics can be derived (Rubinov and Sporns, 2010). First, by sampling the diagonal elements of the matrix, we obtain a measure of *resilience* for each CAP: that is, the likelihood to remain in the same configuration from time t to $t + 1$. Second, after having set the diagonal elements of the matrix to 0, we can define the *in-degree* k_{in} (how likely a CAP is visited from any other), the *out-degree* k_{out} (how likely a CAP is exited towards any other), and the *betweenness centrality*—how important a CAP is regarding the shortest paths between other pairs (Freeman, 1979). In total, the feature space has thus been reduced from K^2 to $4 \cdot K$. This alternative viewpoint is exemplified in Fig. 1A (right column).

In several works, *counts* or *occurrences* (that is, how many times a given CAP is expressed) were used as metrics of interest (Di Perri et al., 2018; Kaiser et al., 2019; Tuleasca et al., 2019). We verified that our suggested metrics also include the information rendered by the counts: as seen in Fig. 1B, the average correlation across CAPs between counts and in-degree, out-degree or resilience exceeded 0.8 (respectively $\rho = 0.83 \pm 0.11$, $\rho = 0.85 \pm 0.08$ and $\rho = 0.81 \pm 0.07$). From pair-wise comparisons between our four metrics, it can also be seen that in-degree and out-degree are strongly correlated ($\rho = 0.87 \pm 0.1$), while resilience and betweenness centrality capture separate information given their more moderate correlations (for resilience: $\rho = 0.45 \pm 0.11$, $\rho = 0.5 \pm 0.11$ and $\rho = 0.3 \pm 0.19$ with in-degree, out-degree and betweenness centrality, respectively; for betweenness centrality: $\rho = 0.59 \pm 0.13$ and $\rho = 0.59 \pm 0.13$ with in-degree and out-degree, respectively). Despite their overall similarity, we decided to retain both in-degree and out-degree as they still yielded different values in specific CAP cases.

In addition to the above metrics that summarise the transitory behaviour across different CAPs, an interesting complement is the assessment of which CAPs are entered from the baseline state of seed activity, as well as of which CAPs are the ones expressed just before a return to baseline activity. With this additional information, a total of $6 \cdot K$ features of interest are available per subject (as 6 metrics are computed for each of K different CAPs). These are the summarising measures that we use in our example application.

2.3. TbCAPs: implementation

We implemented the CAPs processing pipeline as a toolbox in Matlab version 2017a (The MathWorks, Natick, USA). This software is freely accessible at https://c4science.ch/source/CAP_Toolbox.git. It contains a

graphical user interface to facilitate the different steps of the pipeline. In addition, we also provide a scripted version of a typical analysis pipeline for power-users. An illustrative display of the graphical user interface at the end of a typical analysis is provided in Fig. 2. Next, we concisely describe the steps to be performed by the user, and the available options at each stage of the analysis. For more details on all existing functionalities, alternative example applications of the toolbox in clinical settings, and more specific suggestions regarding data preparation and quality control based on our past experience, the reader is pointed at the TbCAPs User Manual that accompanies this work as Supplementary Material.

2.3.1. Data loading

Prior to CAP analysis, the data at hand should have already underwent standard resting-state fMRI preprocessing steps, such as realignment, co-registration, regression of covariates of no interest, and filtering (we advise to only high-pass filter the data). All the volumes to analyse should have been warped to MNI space (as co-activation patterns will be derived from the whole population data). Particular care should be taken, during preprocessing, to attenuate physiology-related artefacts as much as possible, as they may otherwise exert pervasive effects on the BOLD signal (see Caballero-Gaudes and Reynolds (2017) for a recent review). Selected example strategies include the application of independent component analysis-based denoising approaches (Griffanti et al., 2014; Pruim et al., 2015), or the use of a set of regressors reflecting physiological variables (Glover et al., 2000).

Before loading the data to analyse in the toolbox, the user should first define how to mask it (that is, which voxels should be part of the analysis, excluding for instance out-of-brain ones); to do so, a popup window enables to choose between several mask options, after which the **A1. Set mask** button should be pressed. The user is then prompted to select any directory containing part of the functional data to analyse: this will enable to convert the chosen mask into the resolution of the functional data. The prefix specifying the data of interest (e.g., “sw” for an SPM preprocessing with warping and smoothing) should be provided through a dedicated text box.

In a second step, clicking the **A2. Load data** button prompts the user to select all the directories containing the functional data to analyse as part of a given group. We assume here that the data is arranged in a BIDS format (Gorgolewski et al., 2016), which implies in particular that different series of functional volumes to analyse should be located in different directories. Following loading, the data is z-scored by the toolbox, making it fully ready for CAP analysis. In addition, a text file summarising the results from the realignment step (that is, containing the 6 motion time courses) should be present in each directory, so that a framewise displacement time course can be constructed and enable subsequent scrubbing of corrupted frames; if such a file is not available, null motion is assumed and the analysis continues nonetheless.

2.3.2. Spatio-temporal selection

Following data loading, the user is prompted to select one or more seeds to use in the analysis (**B. Select seed file(s)** button): all seed files should be entered at once, each as an MNI space NIFTI volume, which does not need to be at the same spatial resolution as the functional data (this is automatically handled by the toolbox). The brain areas covered by the seed(s) can be inspected in brain slice representations, which can be navigated through by means of dedicated sliders. For now, we allow up to three separate seeds to be entered for the analyses. In addition, the interested user can also plot the average seed-based correlation map across subjects associated to the first selected seed.

The next step is to select which types of events should be retained (activation versus deactivation), and if more than one seed was selected, whether all seeds should show an extreme event at once to select a time point (**Intersection** option), or if a frame should be kept as long as at least one does so (**Union** option). The user is also prompted to determine the threshold T to use in selecting frames (or alternatively, a percentage P of most (de)active frames to retain), and the threshold M above which

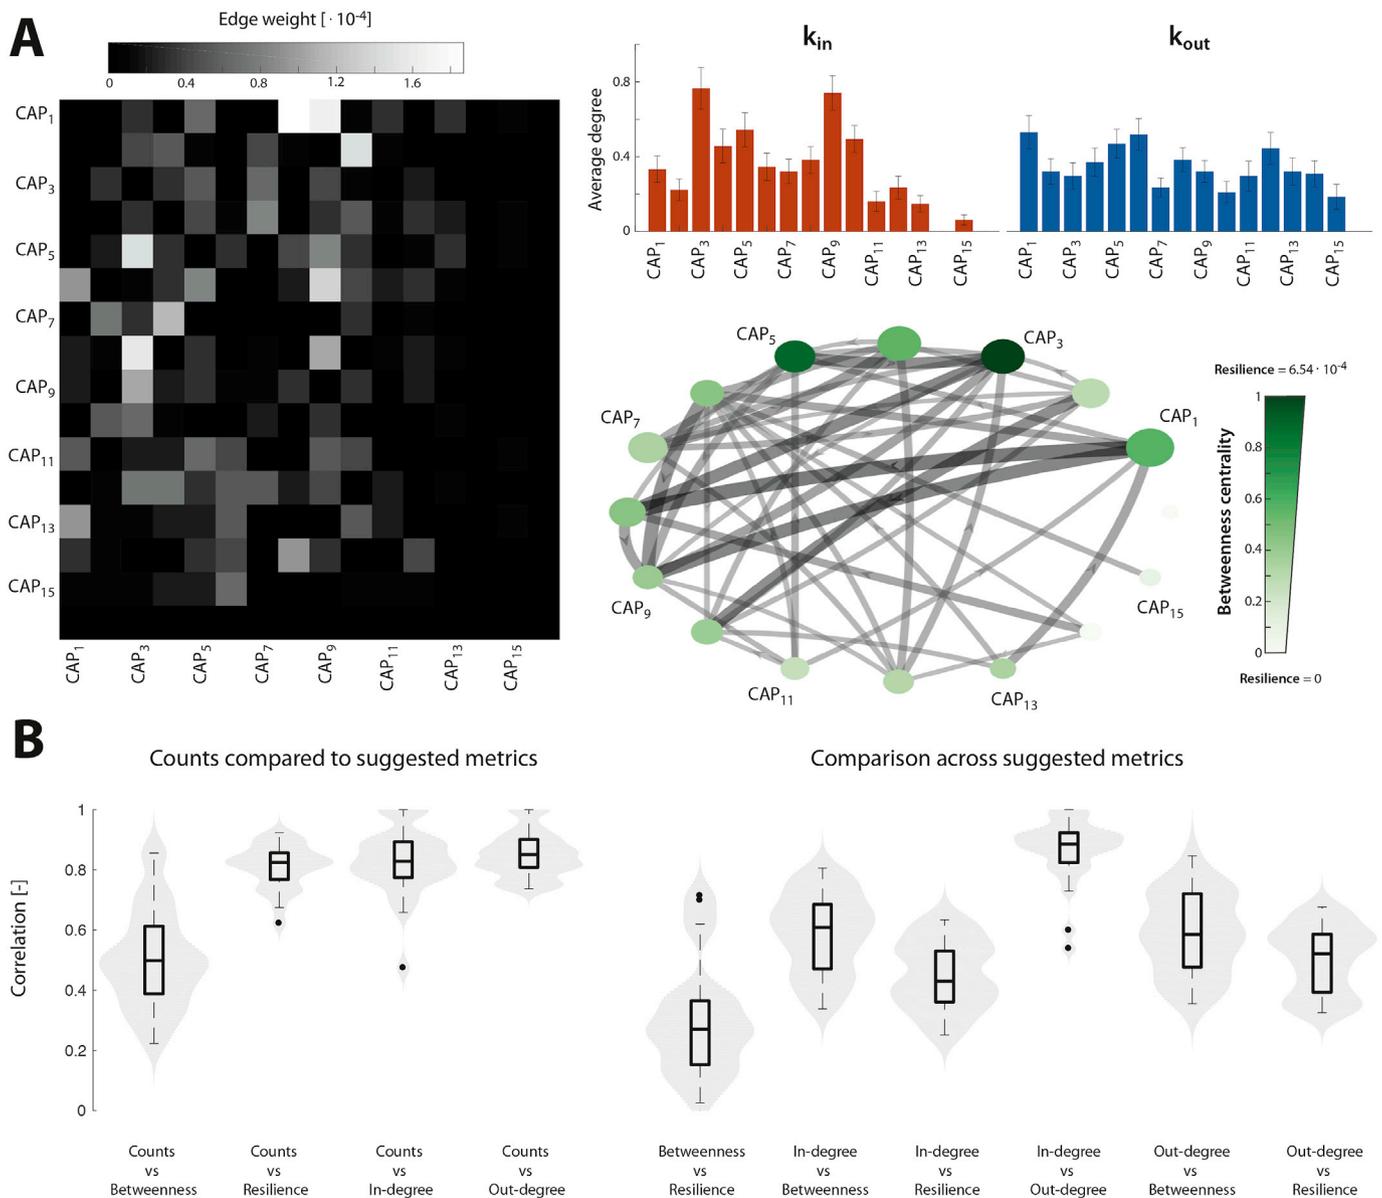


Fig. 1. Generation of CAP metrics. (A) Transitions across CAPs can be viewed in terms of individual transition probabilities for a total of K^2 features (left), or alternatively, a directional graph representation can be constructed (right) to extract in-degree (red bars), out-degree (blue bars), betweenness centrality (green colour coding of the nodes in the bottom plot) and resilience (size of the nodes) information, for a reduced total of $4 \cdot K$ features. In this example, which considers $K = 16$ CAPs, the feature space would thus be lowered from 256 to 64 (four-fold reduction). Error bars denote standard error of the mean, and the displayed transition probabilities, betweenness centrality values and resilience values are averages across subjects. (B) The extracted metrics contain equivalent information compared to the more traditional occurrences, and each such metric characterises partly different aspects of CAP dynamics, as seen by moderate correlations. Each box plot/violin plot representation depicts correlation values across $K = 16$ CAPs. Illustrations from this figure are generated from the data presented in our example application, with $A_p = 0\%$ (i.e., very lenient frame assignment), and without assigning scrubbed volumes.

frames will be deemed excessively corrupted by head motion, and scrubbed out.

At the end of this process, clicking on the **Select time points** button performs the frame selection process, and summarises the percentage of kept volumes across subjects in a violin plot representation. Note that in lieu of a seed-based analysis, we also implemented an alternative seed-free option, following Liu et al. (2013), where frames are retained regardless of any seed. To run this option, the **Seed-free analysis** button should be clicked instead of seed loading.

2.3.3. Generation of co-activation patterns

Regarding the subsequent generation of CAPs, if the optimal number of clusters K to select is not known *a priori*, we offer the possibility to run

consensus clustering (Monti et al., 2003), where clustering is run many times from $K = 2$ to a user-specified K_{max} using a subsample of the data (the percentage of data points to use is specified by P_{CC} , and the number of iterations by N). A good clustering solution is one for which across folds, two frames are either always clustered together, or never clustered together (but not an intermediate case). We quantify this by the Percentage of Ambiguously Clustered pairs, or PAC (Senbabaoglu et al., 2014), and display the stability measure $1 - PAC$. Extended details on consensus clustering can be found in the TbCAPs User Manual.

Following the definition of how many CAPs should be extracted (parameter K in the interface), k-means clustering can be run by clicking the **Cluster** button. The first 5 CAPs with most occurrences across the subject population are displayed, and can be visually inspected and

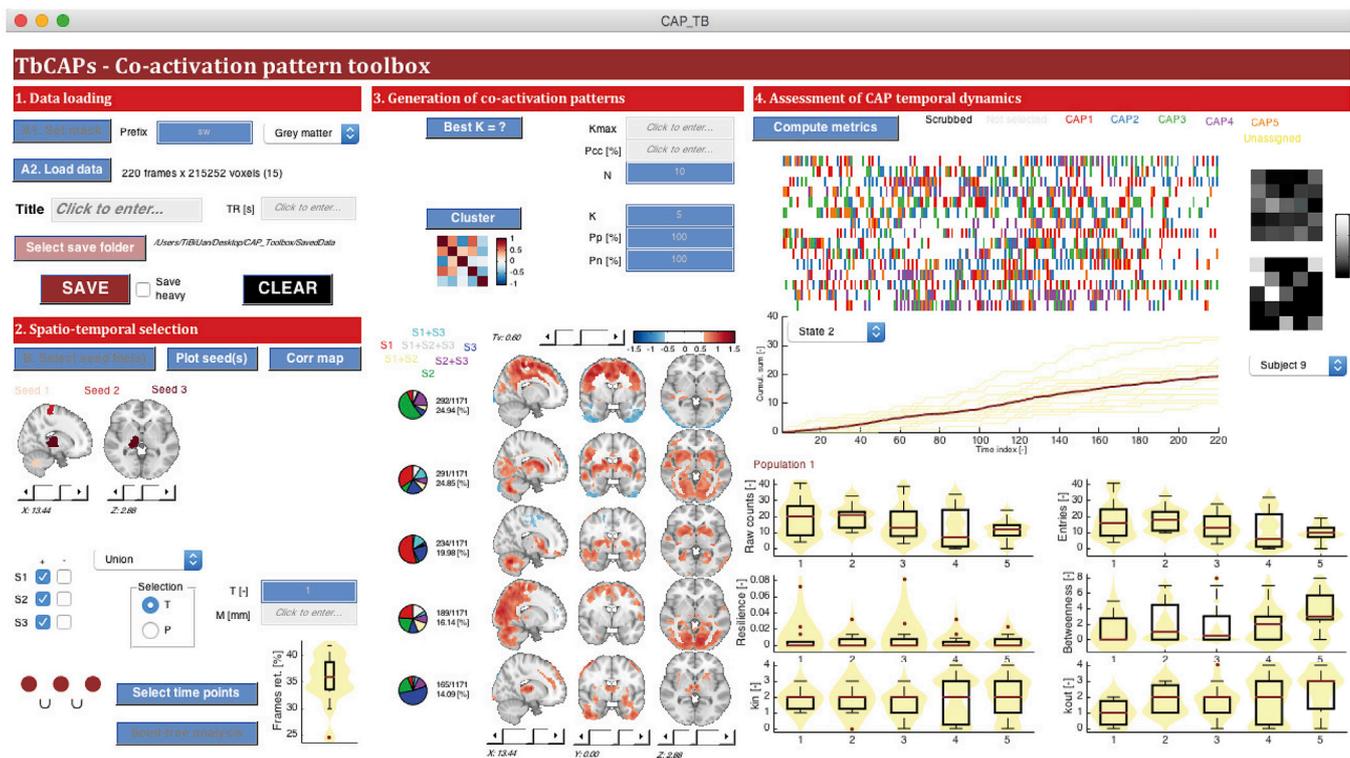


Fig. 2. Illustration of the TbCAPs graphical user interface. The typical output of a CAP analysis is displayed for a set of healthy volunteers, with three seeds chosen within a functional circuitry associated to essential tremor. This data was explicitly analysed in [Tuleasca et al. \(2019\)](#).

navigated through by dedicated sliders. As a complement, the matrix of spatial similarity between CAPs is also provided, and if using the **Union** option in a multi-seed analysis, the user is shown pie charts summarising, for each CAP, what fraction of frames was selected in a given seed combination configuration.

2.3.4. Computation of metrics

Finally, upon clicking the **Compute metrics** button, displays of CAP expression time courses and cumulative CAP expression along time appear on screen. The latter can be adjusted to selectively view the cumulative counts of a given CAP, across subjects and as a population average. Transition probability matrices can also be inspected in terms of average entries at the population level, or for each subject, with the option to select the one to display the data for.

In addition, 6 violin plot representations summarising the distribution of computed metrics across subjects for each CAP are also provided: they reflect (1) raw counts, (2) number of entries in a given CAP, (3) resilience, (4) betweenness centrality, (5) in-degree and (6) out-degree.

2.3.5. Analysis of multiple subject populations

In some settings, the user may wish to compare different subject populations (for instance, healthy controls and a clinical group): this can be done by sequentially loading up to 4 different populations at the start of the analysis. CAPs will be derived from the first population, and there is then the need to assign the frames from the other populations to the CAPs by a matching process. In doing so, the spatial correlation between a frame to assign and the CAP to which it is most similar is compared to the distribution of spatial correlations of the frames from population 1 that belong to the CAP in question: if the A_P^{th} percentile of this distribution is exceeded, assignment is performed; else, the frame is left unassigned and belongs to an extra $(K+1)^{\text{th}}$ cluster. Additional details regarding multi-population analyses are provided in the TbCAPs User Manual.

2.4. Application to experimental fMRI data

2.4.1. Functional data preprocessing

As a proof of feasibility and application of TbCAPs, we considered a sample of 181 subjects from the *Human Connectome Project* ([Van Essen et al., 2013](#)), aged between 26 and 35 years old. This sample originated from a slightly larger, randomly selected set of subjects that had at least one fully exploitable resting-state scanning session on which to apply the method, and less than 5% of recorded behavioural entries that were missing; a few subjects from this original set were discarded due to errors in preprocessing. Details regarding acquisition parameters can be found elsewhere ([Smith et al., 2013](#)), but briefly, the data was acquired at a TR of 0.72 s over 15 min (for a total of 1200 fMRI volumes), with a spatial resolution following initial preprocessing steps of $2 \text{ mm} \times 2 \text{ mm} \times 2 \text{ mm}$.

We started from the *minimally preprocessed* resting-state data (first session, LR acquisition direction). The first 10 samples of the data were discarded. We then performed linear detrending, and regressed out low-frequency components of the discrete cosine transform basis with a cutoff frequency at 0.01 Hz. Due to collinearity with this basis, we did not regress out average white matter or cerebrospinal fluid time courses. We also chose not to regress motion parameter time courses, as motion is handled within the co-activation pattern pipeline by scrubbing, and because recent evidence points at the fact that motion regression schemes may not always be beneficial in the context of brain/behaviour analyses ([Bolton et al., 2020](#)). As for global signal regression, given the lack of a clear consensus ([Murphy and Fox, 2017](#)), we preferred to leave the data as untouched as possible and did not include it.

Following the regression step, the data was scrubbed at a framewise displacement threshold of 0.3 mm, and excised volumes were estimated with cubic spline interpolation. Although scrubbing is performed within TbCAPs, we reasoned that if we wished to try and assign scrubbed frames to CAPs in our additional analyses regarding head motion, it would make more sense to have previously corrected these volumes to the best of our abilities.

Then, individual fMRI volumes were smoothed at a full width at half maximum value of 5 mm, and in order to make the analyses computationally more affordable, spatial resolution was downsampled at $3 \text{ mm} \times 3 \text{ mm} \times 3 \text{ mm}$. Eventually, z-scoring was handled within TbCAPs as a final preprocessing step prior to CAP analysis.

2.4.2. Selection of seed and behaviour of interest

As a behaviour of interest to study, we selected the Short Penn Continuous Performance Test (SCPT), which quantifies continuous sustained attention (Gur et al., 2010). In more details, participants see vertical and horizontal red lines flash on screen, and from block to block, must respond either when the lines form a number, or a letter. The lines are displayed for 300 ms, followed by a 700 ms inter-trial interval.

We started from raw behavioural entries provided by the HCP, for 951 different subjects. There are 8 available SCPT measures: amount of true positives, false positives, true negatives or false negatives, median response time for true positive responses, sensitivity, specificity and longest run of non-responses. In order to reduce this information into one summary measure while filling in missing behavioural entries, we performed probabilistic PCA (Bishop, 1999). The output composite score positively correlated with true positives, true negatives, sensitivity and specificity ($\rho = 0.24, 1.00, 0.25$ and 1.00 , respectively), thus summarising overall task performance. We z-scored this output measure across subjects, in order to quantify performance with respect to the overall population. We then extracted the behavioural data related to the 181 subjects considered in this work.

To study sustained attention, we focussed on a right dorsolateral prefrontal cortex seed from the task-positive network, which we extracted from the associated independent component map provided by Shirer et al. (2012). Our hypothesis was that the expression of different TPN configurations would relate to sustained attentional performance.

2.4.3. Co-activation pattern analysis details

We resorted to a threshold $T = 1.5$ to select active frames, and performed scrubbing with a framewise displacement threshold $M = 0.3 \text{ mm}$.

To avoid double dipping (Kriegeskorte et al., 2009), CAPs were extracted from a randomly selected subset of 100 subjects, while we performed correlations with behaviour for the remaining 81 only. To determine the optimal number of clusters, we used consensus clustering (Monti et al., 2003). We then ran k-means $n_{\text{rep}} = 50$ separate times, keeping the best solution. We included all voxels in the analyses ($P_P = P_N = 100\%$), and used spatial correlation as our distance measure; given two similarly-sized vectors \mathbf{x} and \mathbf{y} , this thus yields $\text{dist}(\mathbf{x}, \mathbf{y}) = 1 - \text{corr}(\mathbf{x}, \mathbf{y})$.

Following the extraction of the CAPs on our 100 *training subjects*, we determined which CAP was expressed at each retained fMRI volume of the other 81 subjects. To do so, we used the aforementioned assignment process with A_P ranging from 0 to 100%.

2.4.4. Assessment of brain/behaviour relationship

As imaging metrics of interest, we considered in-degree, out-degree, betweenness centrality and resilience for each CAP, and also included the amount of excursions from the baseline state, and the amount of excursions back to the baseline state. Thus, we generated a total of $6 \cdot K$ imaging features per subject.

After having obtained the behavioral scores $\mathbf{b} \in \mathbb{R}^{81 \times 1}$ and metrics $\mathbf{M} \in \mathbb{R}^{81 \times 6 \cdot K}$ for our population of subjects, we used Partial Least Squares (PLS) analysis (McIntosh and Lobaugh, 2004; Krishnan et al., 2011) to probe the existence of a brain/behaviour relationship.

Briefly, consider a matrix of behavioural features $\mathbf{B} \in \mathbb{R}^{S \times n_B}$ and a matrix of imaging metrics $\mathbf{M} \in \mathbb{R}^{S \times n_M}$. Assuming that $n_B < n_M$, and using the singular value decomposition, the covariance between these two sets is given by:

$$\mathbf{R} = \mathbf{M}^T \mathbf{B} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T = \sum_{i=1}^{n_B} \sigma_i \mathbf{u}_i \mathbf{v}_i^T, \quad (5)$$

where each column in \mathbf{U} and \mathbf{V} contains the weights (so called *salience*s) that respectively multiply imaging and behavioural markers to yield a maximised covariance between both sets. The associated singular value σ_i is proportional to the fraction of covariance explained by the component at hand.

In our case, since $n_B = 1$ (we only consider one behavioural measure), only one covariance component is retrieved, which implies $v_1 = 1$. The interesting information lies in \mathbf{u}_1 : positive-valued saliences highlight metrics that are larger in subjects who show a greater cognitive ability, and negative-valued saliences are associated to metrics that, when larger, impede attentional performance.

Prior to running the algorithm, each of the $6 \cdot K$ features was z-scored across subjects. In order to assess significance, we reran PLS 1000 times after having randomly shuffled the subject entries in one of the two matrices; to non-parametrically derive a p-value, the singular value of the actual covariance component was compared to the null distribution constructed from this permutation process.

Further, to establish the stability of the salience weights, we reran PLS 1000 times using a randomly selected subsample of 80% of the data, and computed a bootstrap score for each salience weight as its mean across folds divided by its standard deviation.

2.4.5. Influence of head motion on quantified metrics

While scrubbing enables to minimise the deleterious impacts of motion on the analysis and compute clean CAPs, discarding frames also has the potential to distort transition probability estimates. For example, a succession of three frames in the same state (which would amount to a higher resilience for the CAP in question) would not be captured if the middle frame is scrubbed out.

To verify that our findings were minimally sensitive to this effect, we ran another series of analyses in which we also performed the aforementioned assignment process on scrubbed frames, with a similar A_P range as for assigning test subject frames. This way, frames strongly distorted by head motion still do not enter the analysis, but more mildly affected fMRI volumes can be matched to their CAP. We verified the reproducibility of our findings upon this additional analytical step.

3. Results

Consensus clustering results are displayed in Fig. 3A for K values ranging from 10 to 40. Positive peaks highlight good candidate values (see figure legend for details); such values are present for diverse numbers of clusters (more notably at $K = 16, 22, 32$). While a lower number of clusters yields a reduced feature space and more interpretable outcomes, CAPs may not be segmented finely enough to resolve insightful dynamic properties regarding cognition. Our strategy was thus to first perform an exploratory assessment, in which we evaluated the significance of the brain/behaviour correlation across a set of candidate K values ($K_{\text{opt}} = \{14, 16, 22, 32\}$) and assignment thresholds (forcing the assignment of all frames, or using $T_P = [0 : 5 : 100]$), to select the relevant parameters to proceed forward with.

The results of this exploratory assessment are displayed in Fig. 3B, when scrubbed frames are discarded (left panel) or also assigned to the CAPs at threshold A_P (right panel). Both settings yield very similar significance values, which is good evidence that remaining head motion effects only have a minor influence on the analyses. As the assignment threshold increases (that is, less and less frames are assigned because the criterion becomes more and more stringent), significance generally decreases. A smooth spot can be observed for $K = \{16, 22\}$ and $A_P < 15\%$, which indicates that this granularity is optimal in the context of behavioural prediction. We selected $K = 16$ and $A_P = 0\%$ as values for more detailed subsequent analyses.

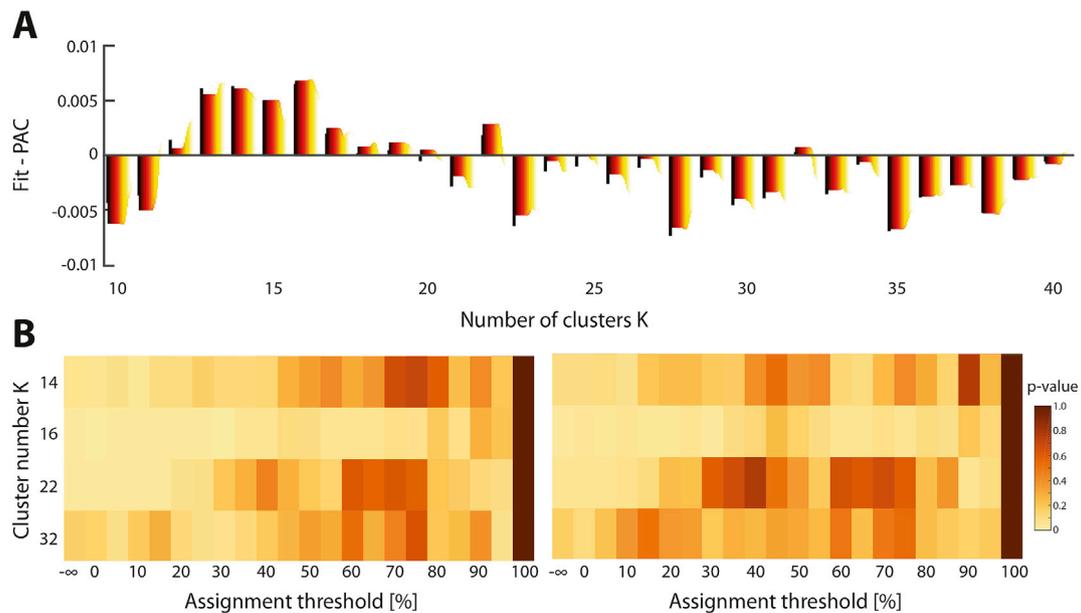


Fig. 3. Parameter selection. (A) Consensus clustering results for a range of K values from 10 to 40. Outputs from the algorithm are percentages of ambiguously clustered pairs (PAC); since this measure requires the definition of an interval of consensus values within which clustering is deemed *ambiguous*, we investigated a range of values as colour-coded from black to yellow. For each value, we fitted a decreasing exponential to capture the overall tendency of PAC values going down with larger cluster number. The y-axis of the plot depicts the difference between this fitted value and the actual value: thus, a positive difference means that the considered cluster number is more satisfying than what would be predicted in terms of the overall behaviour. (B) Across candidate cluster number values selected from consensus clustering, and assignment threshold values, significance of the relationship between CAP metrics and attentional abilities, as quantified by the p-value obtained upon PLS analysis. The left panel depicts the results for which scrubbed frames are not considered at all, while in the right panel, scrubbed frames were also assigned to the CAPs. The infinity symbol is used to depict a case in which assignment is done for all frames, even if a frame is not sufficiently close to any CAP when comparing its spatial correlation to the associated correlation distribution.

CAPs are displayed, for this chosen parameter set, in Fig. 4A, while their involvement in driving the brain/behaviour relationship, as quantified by salience weights across our range of investigated metrics, is depicted in Fig. 4B. The correlation between actual and predicted attentional performance was strongly significant ($R = 0.587$, $p < 0.001$; Fig. 4C). The associated covariance component found by PLS analysis was significant at $p = 0.003$. Note that this relationship is derived from only subjects that were not used to construct the CAPs.

CAP₁ depicts co-activation of a range of resting-state networks, including the auditory, somatomotor, visual and salience ones. Attentional performance was better in the subjects that transitioned more frequently from the baseline state of seed activity to this CAP. CAP₁ was also more often the entry point towards other CAPs in high performance subjects, as seen from a strongly positive out-degree salience weight.

Good subjects in terms of continuous performance also more often entered CAP₂ and CAP₇ from other states (as seen from positive in-degree salience weights), and these same 2 CAPs were also more influential in the transitory behaviour of CAP dynamics (since betweenness centrality salience weights also showed large positive values). In both CAPs, the seed region co-activates with a restricted set of areas including the right inferior parietal cortex (for both), the posterior cingulate cortex and medial prefrontal cortex (for CAP₂), and the right anterior prefrontal cortex (for CAP₇).

CAP₃ and CAP₅ were associated to good attentional abilities from the viewpoint of several metrics, which emphasises the importance of their expression: for CAP₃, it involved resilience, in-degree and out-degree, while for CAP₅, it was return to baseline, resilience, in-degree and betweenness centrality. Both CAPs include strong co-activation with the right inferior parietal cortex, and for CAP₃, also with the left cerebellum lobule VI and a subpart of the occipital cortex.

CAP₄ and CAP₁₄ were the only states whose expression was detrimental to attentional performance, in terms of betweenness centrality for the former, and of return to baseline for the latter. CAP₄ displayed bilateral right superior cortex and anterior prefrontal cortex co-activation

with the seed, while for CAP₁₄, involved areas were the fusiform gyrus, parahippocampal cortex, and a diffuse right lateralised spot covering parts of the auditory, secondary somatosensory and posterior insular cortices.

The majority of the other CAPs that did not show any link to attentional performance involved co-activations with regions that were not part of the attentional networks: for instance, CAP₆ includes the anterior cingulate cortex and anterior insula; CAP₈ contains the anterior cingulate, visual and right somatosensory cortices; CAP₉ showcases primary visual and auditory cortices; CAP₁₀ shows the angular gyrus and part of the precuneus; and CAP₁₁ and CAP₁₅ mostly highlight ventral medial prefrontal cortex signal.

4. Discussion

In this work, we have introduced TbCAPs, a toolbox for co-activation pattern analysis, which provides practitioners with an intuitive dedicated graphical user interface as well as a powerful scripting equivalent. It provides an easy control over all key analytical parameters of the technique, novel methodological additions for augmented analyses, and facilitated visualisation of the resulting CAPs and associated metrics. Although we have focussed on the usefulness of CAP analysis in the resting-state setting, we also remark that nothing precludes the use of the technique in task-based investigations.

As most CAP studies to date have revealed the potential of the approach in clinical settings (Amico et al., 2014; Di Perri et al., 2018; Kaiser et al., 2019; Tuleasca et al., 2019), we sought to demonstrate the relevance of the technique in another context; *i.e.*, rather than considering a *classification* problem in which two or more distinct subject populations are separated, we considered a *regression* task in which we attempted to explain attentional abilities within a more homogeneous population in a continuous vigilance task by means of CAPs dynamics.

We observed that of all the extracted CAPs showing coupling with the right dorsolateral prefrontal seed, the large majority either did not

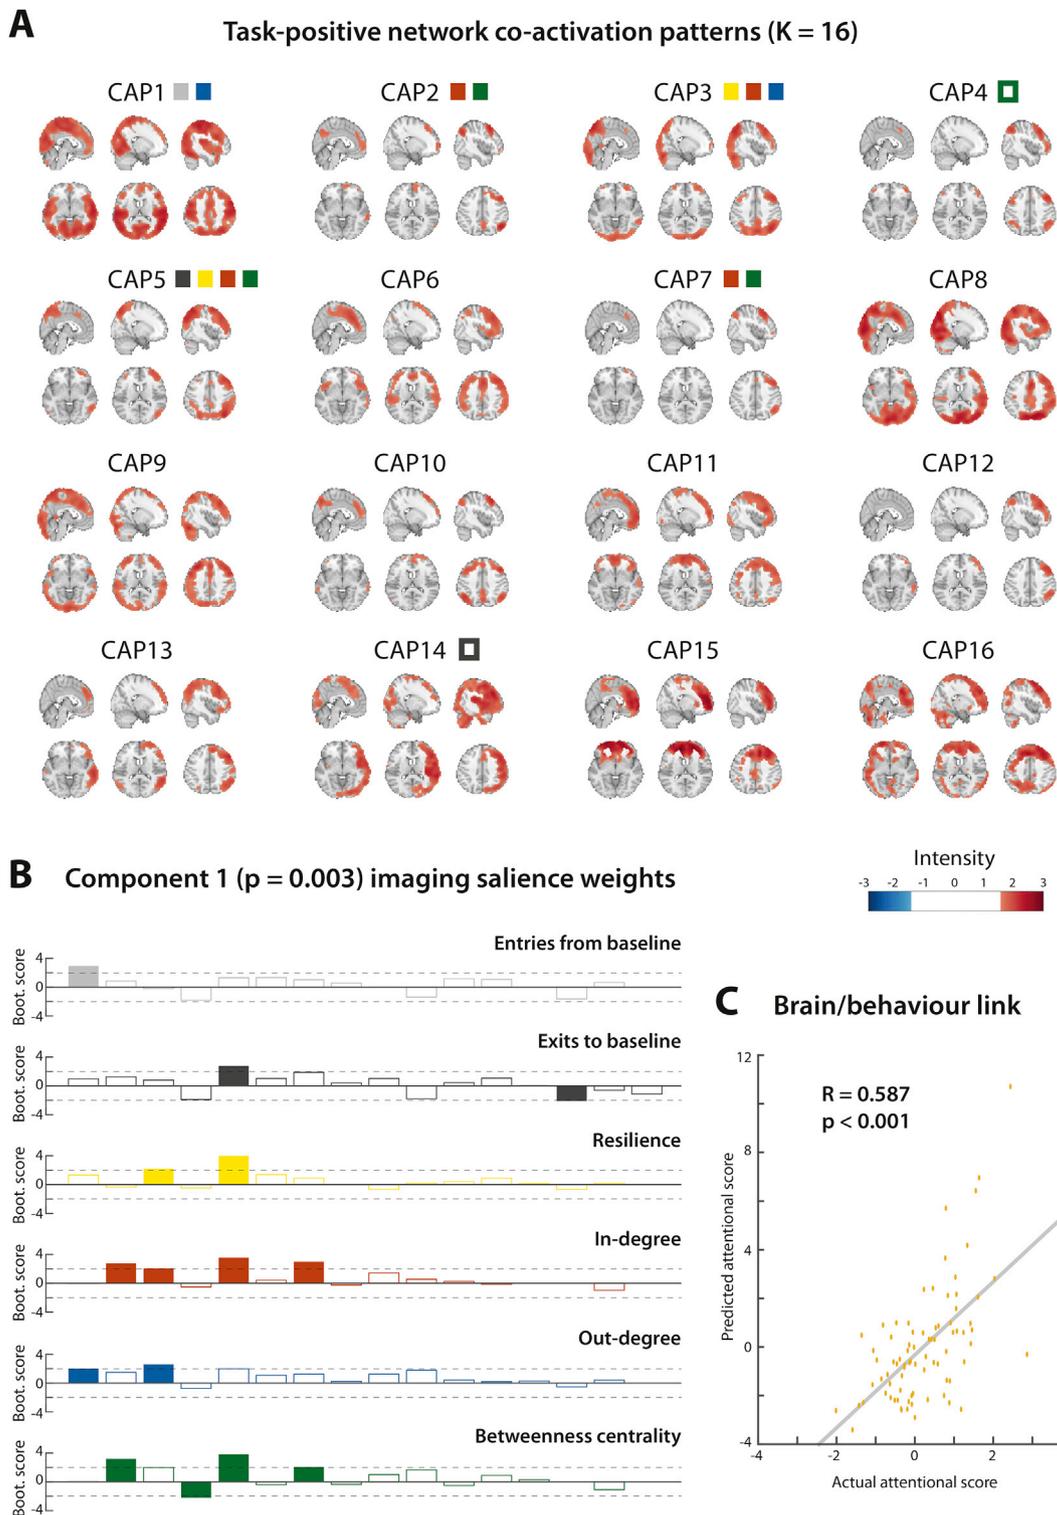


Fig. 4. Co-activation patterns, and relationship to attentional abilities. (A) The 16 obtained CAPs are plotted, with coloured rectangles symbolising the CAPs that are particularly important (bootstrap score of the associated salience weight larger than 2) for the brain/behaviour relationship as quantified from a given metric (grey: entries from baseline state, black: exits to baseline state, yellow: resilience, red: in-degree, blue: out-degree, green: betweenness centrality). Filled rectangles highlight beneficial CAPs (positive bootstrap score), while hollow rectangles depict detrimental CAPs (negative bootstrap score). (B) Salience weights across all 16 CAPs and the 6 investigated CAP metrics. The threshold bootstrap score past which a weight is considered significant is highlighted by a horizontal dashed line. Empty bars denote non-significant weights, while filled bars represent significant ones. (C) Actual attentional performance score (x-axis) versus predicted values with PLS (y-axis).

appear to be involved in attentional abilities, or showed positive salience weights indicating a positive impact of their expression. This is not so surprising, given that our analysis was focussed on a region of the attention network in the first place. The common feature of beneficial CAPs appeared to be the coupling of an array of other regions previously pinpointed in continuous performance tasks, including the inferior parietal cortex, cerebellum lobule VI or occipital cortex (Häger et al., 1998; Ogg et al., 2008; Tana et al., 2010). At the same time, these beneficial CAPs also barely involved coupling of other functionally distinct networks.

The one CAP for which the above reasoning does not hold is CAP₁: despite the involvement of a very diverse set of regions, it was also retrieved as beneficial for attentional performance. More precisely, contrarily to most of the others, the expression of this CAP appears to be essential at the start of a seed activation sequence: indeed, salience weights were large specifically for the entries from baseline and out-degree metrics. In other words, there is first a transition from baseline to this CAP, followed by the exit of that configuration to reach more spatially well-defined states. This involvement of short-lived periods of extensive cross-network interactions in mediating some aspects of human cognition has recently started to be appreciated as an insightful functional brain mechanism (Betzel et al., 2016; Fukushima et al., 2018).

The fact that all the probed metrics significantly contributed to explain attention is good evidence in favour of the temporal complexity of functional brain dynamics: instead of an instantaneous characterisation or a one-frame expression of a telling functional state, what truly matters is a complex mix between how activation starts (captured by the from-baseline and to-baseline metrics), how transitions occur across distinct functional states (as seen from in-degree, out-degree and betweenness centrality), and how lasting a given state is (as quantified by resilience). Our characterisation relates to the broad family of temporal modelling approaches, of which notable examples include the use of graph-theoretical analysis for energy landscape (Kang et al., 2019), or hidden Markov models—HMMs (Vidaurre et al., 2017; Bolton et al., 2017).

A future actual use of HMMs in CAP analysis would make it possible to not only estimate the transitory behaviour across CAPs, but also the parameters governing the expression of the voxelwise patterns of BOLD signal. In addition, this full characterisation of the system would enable the generation of new data, going beyond the mere computation of empirical estimates as done now. However, HMM-based approaches require extensive amounts of data to converge properly, which hinders the exploration of subject-specific properties with typical data amounts (Bolton et al., 2017). It is to bypass this issue that we instead set, in our current approach, to derive composite metrics that incorporate the information from several CAPs at once. Going back to the above example where we considered $K = 16$ CAPs, we could thus lower our amount of imaging features from $K^2 = 256$ (all of which should be estimated with an HMM) to $6 \cdot K = 96$, and this feature extraction strategy will become more and more beneficial as the number of examined CAPs increases.

Prediction of continuous performance abilities from resting-state fMRI recordings has been shown possible in previous functional connectivity work relying on second-order correlational measures across brain regions (Rosenberg et al., 2016). More recently, this characterisation has been pushed to the dynamic level by Fong et al. (2019), who showed that prediction can also be successfully achieved when *temporal variability*, which quantifies fluctuations in functional connectivity over the course of a scanning session, is used as a metric of interest. Our prediction accuracy is on par with the one achieved in this whole-brain analysis, despite focussing on one seed region. Interestingly, the authors described the fact that lowered temporal variability was beneficial for better attentional performance, and that many of the most important features for prediction involved an executive control brain network: this is fully consistent with our results, in which increased resilience (which can be expected to yield lowered temporal variability) of CAPs featuring executive control areas is beneficial.

In comparison to clinical applications of CAP analysis, in which between 3 and 8 CAPs are typically considered, a finer granularity was required in the present work (see Fig. 3). This is not surprising given that the regression problem at hand here is more challenging than a classification task, as we need to predict a value within a continuum. Furthermore, the functional underpinnings of inter-individual differences in cognitive abilities are likely more subtle than when comparing subjects across consciousness or disease severity levels. In fact, if a too low number of CAPs is extracted, patterns with a different cognitive relevance are averaged together as a single cluster, which impedes prediction; conversely, if a too large number of CAPs is extracted, meaningful configurations become further segmented, and statistical power is lost due to the smaller amount of frames constituting each CAP.

A limitation of our work, and of any standard CAP analysis, is that computations are performed on BOLD time courses that have not been freed from hemodynamic effects (*i.e.*, that have not undergone deconvolution with a hemodynamic response function—HRF—estimate). The parameters of the HRF vary across brain regions (Handwerker et al., 2004), and such differences can confound functional connectivity estimates (Rangaprakash et al., 2018). However, we believe that such impacts remain minimal in the present analysis, since the areas located in most CAPs are also consistently found in the literature on attentional performance. Furthermore, we do not focus on the spatial patterns of the CAPs, but instead, on the transitory dynamics between them. Since the HRF also varies across subjects (Aguirre et al., 1998), we cannot rule out that our association to behaviour was partly influenced by such effects, but given the fact that our sample of subjects was distributed over a narrow age range of 10 years—a leading factor in HRF variability (D'Esposito et al., 1999), we consider this an unlikely scenario.

While the example application introduced here involved fast TR (0.72 s) data acquired in a multi-band setting, it should be remarked that the findings of CAP analyses may partly vary as a function of the employed acquisition type. Since the HRF acts as a bottleneck factor, setting a limit below which functional dynamics cannot be resolved more finely anymore even at faster acquisition paces, the influence of the TR *per se* may remain limited. However, another associated problematic is the differential influence of physiological rhythms on the functional data (Chen et al., 2019). This even extends to the motion time courses typically used in data preprocessing—including scrubbing as performed within TbCAPs, since additional physiology-driven components are observed at faster TRs (Power et al., 2019). All in all, we thus wish to emphasise the importance of freeing the data from such physiological impacts as well as possible using the available resources for this purpose (Glover et al., 2000; Griffanti et al., 2014; Pruim et al., 2015).

In future work, it will be interesting to examine clinical or cognitive research hypotheses at the broader focus level of more than one seed region. As alluded to above, this is already feasible with our current toolbox version, and may enable to better bridge the gap between seed-based and whole-brain analyses. We also foresee additional technical developments in the near future, such as the possibility to extract *co-activation sequences* (that is, series of successive fMRI volumes) rather than CAPs. Finally, we would like to encourage the motivated readers to help us in further improving our publicly accessible toolbox, so that it can become an even more multimodal package integrated with other widely used existing software.

CRediT authorship contribution statement

Thomas A.W. Bolton: Formal analysis, Writing - original draft, Conceptualization, Software, Writing - review & editing. **Constantin Tuleasca:** Data curation, Validation. **Diana Wotruba:** Validation. **Gwladys Rey:** Validation. **Herberto Dhanis:** Validation. **Baptiste Gauthier:** Validation. **Farnaz Delavari:** Validation. **Elenor Morgenroth:** Validation, Writing - review & editing. **Julian Gaviria:** Validation. **Eva Blondiaux:** Validation. **Lukasz Smigielski:** Validation. **Dimitri Van De Ville:** Conceptualization, Supervision, Writing - review & editing.

Acknowledgments

Constantin Tuleasca gratefully acknowledges receipt of a ‘Young Researcher in Clinical Research Grant’ (Jeune Chercheur en Recherche Clinique) from the University of Lausanne (UNIL), Faculty of Biology and Medicine (FBM), and the Lausanne University Hospital (CHUV). In addition, the authors would like to thank Raphaël Liégeois for his rereading of the manuscript, as well as the anonymous Reviewers who commented on this work, as they significantly contributed to improve the quality of the manuscript and of the toolbox by their insightful suggestions.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.neuroimage.2020.116621>.

References

- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998. The variability of human, BOLD hemodynamic responses. *Neuroimage* 8 (4), 360–369.
- Amico, E., Gomez, F., Di Perri, C., Vanhauzenhuyse, 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 (6), 1–9.
- Betzel, R.F., Fukushima, M., He, Y., Zuo, X., Sporns, O., 2016. Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. *Neuroimage* 127, 287–297.
- Bishop, C.M., 1999. *Bayesian PCA*. Curran Associates, pp. 382–388.
- Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34 (4), 537–541.
- Bolton, T.A.W., Tarun, A., Sterpenich, V., Schwartz, S., Van De Ville, D., 2017. Interactions between large-scale functional brain networks are captured by sparse coupled HMMs. *IEEE Trans. Med. Imag.* 37 (1), 230–240.
- Bolton, T.A.W., Kebets, V., Glerean, E., Zöllner, D., Li, J., Yeo, B.T.T., Caballero-Gaudes, C., Van De Ville, D., 2020. Agito ergo sum: Correlates of spatio-temporal motion characteristics during fMRI. *Neuroimage* 209, 116433.
- Bowring, A., Maumet, C., Nichols, T.E., 2019. Exploring the impact of analysis software on task fMRI results. *Hum. Brain Mapp.* 40 (11), 3362–3384.
- Caballero-Gaudes, C., Reynolds, R.C., 2017. Methods for cleaning the BOLD fMRI signal. *Neuroimage* 154, 128–149.
- Chang, C., Glover, G.H., 2010. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage* 50 (1), 81–98.
- Chen, J.E., Chang, C., Greicius, M.D., Glover, G.H., 2015. Introducing co-activation pattern metrics to quantify spontaneous brain network dynamics. *Neuroimage* 111, 476–488.
- Chen, J.E., Polimeni, J.R., Bollmann, S., Glover, G.H., 2019. On the analysis of rapidly sampled fMRI data. *Neuroimage* 188, 807–820.
- Damoiseaux, J.S., Rombouts, S.A.R., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. *Proc. Natl. Acad. Sci.* 103 (37), 13848–13853.
- D'Esposito, M., Zarahn, E., Aguirre, G.K., Rypma, B., 1999. The effect of normal aging on the coupling of neural activity to the BOLD hemodynamic response. *Neuroimage* 10 (1), 6–14.
- Di, X., Biswal, B.B., Jan, 2015. Dynamic brain functional connectivity modulated by resting-state networks. *Brain Struct. Funct.* 220 (1), 37–46.
- Di Perri, C., Amico, E., Heine, L., Annen, J., Martial, C., Larroque, S.K., Soddu, A., Marinazzo, D., Laureys, S., 2018. Multifaceted brain networks reconfiguration in disorders of consciousness uncovered by co-activation patterns. *Hum. Brain Mapp.* 39 (1), 89–103.
- Fong, A.H.C., Yoo, K., Rosenberg, M.D., Zhang, S., Li, C.R., Scheinost, D., Constable, R.T., Chun, M.M., 2019. Dynamic functional connectivity during task performance and predicts individual differences in attention across studies. *Neuroimage* 188, 14–25.
- Fox, M.D., Greicius, M., 2010. Clinical applications of resting state functional connectivity. *Front. Syst. Neurosci.* 4, 19.
- Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci.* 102 (27), 9673–9678.
- Freeman, L.C., 1979. Centrality in social networks conceptual clarification in Hawaii nets conferences. *Soc. Network.* 1 (3), 215–239.
- Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* 2 (1), 56–78.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J., Frith, C.D., Frackowiak, R.S.J., 1994. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2 (4), 189–210.
- Fukushima, M., Betzel, R.F., He, Y., de Reus, M.A., van den Heuvel, M.P., Zuo, X., Sporns, O., 2018. Fluctuations between high- and low-modularity topology in time-resolved functional connectivity. *Neuroimage* 180, 406–416.
- Glover, G.H., Li, T., Ress, D., 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn. Reson. Med.* 44 (1), 162–167.
- Gorgolewski, K.J., Auer, T., Calhoun, V.D., Craddock, R.C., Das, S., Duff, E.P., Flandin, G., Ghosh, S.S., Glatard, T., Halchenko, Y.O., et al., 2016. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data* 3, 160044.
- Greicius, M., 2008. Resting-state functional connectivity in neuropsychiatric disorders. *Curr. Opin. Neurol.* 21 (4), 424–430.
- Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., et al., 2014. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* 95, 232–247.
- Gur, R.C., Richard, J., Huggett, P., Calkins, M.E., Macy, L., Bilker, W.B., Brensinger, C., Gur, R.E., 2010. A cognitive neuroscience-based computerized battery for efficient measurement of individual differences: standardization and initial construct validation. *J. Neurosci. Methods* 187 (2), 254–262.
- Häger, F., Volz, H., Gaser, C., Mentzel, H., Kaiser, W.A., Sauer, H., 1998. Challenging the anterior attentional system with a continuous performance task: a functional magnetic resonance imaging approach. *Eur. Arch. Psychiatr. Clin. Neurosci.* 248 (4), 161–170.
- Handwerker, D.A., Ollinger, J.M., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* 21 (4), 1639–1651.
- Holmes, A.J., Hollinshead, M.O., O'Keefe, T.M., Petrov, V.I., Fariello, G.R., Wald, L.L., Fischl, B., Rosen, B.R., Mair, R.W., Roffman, J.L., et al., 2015. Brain Genomics Superstruct Project initial data release with structural, functional, and behavioral measures. *Scientific Data* 2, 150031.
- 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.
- Kaiser, R.H., Kang, M.S., Lew, Y., Van Der Feen, J., Aguirre, B., Clegg, R., Goer, F., Esposito, E., Auerbach, R.P., Hutchison, R.M., et al., 2019. Abnormal fronto-insular-default network dynamics in adolescent depression and rumination: a preliminary resting-state co-activation pattern analysis. *Neuropsychopharmacology* 44, 1604–1612.
- Kang, J., Pae, C., Park, H., 2019. Graph-theoretical analysis for energy landscape reveals the organization of state transitions in the resting-state human cerebral cortex. *PLoS One* 14 (9), 0222161.
- Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S.F., Baker, C.I., 2009. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* 12 (5), 535–540.
- Krishnan, A., Williams, L.J., McIntosh, A.R., Abdi, H., 2011. Partial least squares (PLS) methods for neuroimaging: a tutorial and review. *Neuroimage* 56 (2), 455–475.
- 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.
- Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc. Natl. Acad. Sci.* 110 (11), 4392–4397.
- Liu, X., Zhang, N., Chang, C., Duyn, J.H., 2018. Co-activation patterns in resting-state fMRI signals. *Neuroimage* 180, 485–494.
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412 (6843), 150–157.
- McIntosh, A.R., Lobaugh, N.J., 2004. Partial least squares analysis of neuroimaging data: applications and advances. *Neuroimage* 23, 250–263.
- Monti, S., Tamayo, P., Mesirov, J., Golub, T., 2003. Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. *Mach. Learn.* 52 (1), 91–118.
- Murphy, K., Fox, M.D., 2017. Towards a consensus regarding global signal regression for resting state functional connectivity MRI. *Neuroimage* 154, 169–173.
- Nooner, K.B., Colcombe, S., Tobe, R., Mennes, M., Benedict, M., Moreno, A., Panek, L., Brown, S., Zavitz, S., Li, Q., et al., 2012. The NKI-Rockland sample: a model for accelerating the pace of discovery science in psychiatry. *Front. Neurosci.* 6, 152.
- Ogg, R.J., Zou, P., Allen, D.N., Hutchins, S.B., Dutkiewicz, R.M., Mulhern, R.K., 2008. Neural correlates of a clinical continuous performance test. *Magn. Reson. Imag.* 26 (4), 504–512.
- Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59 (3), 2142–2154.
- Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., et al., 2011. Functional network organization of the human brain. *Neuron* 72 (4), 665–678.
- Power, J.D., Lynch, C.J., Silver, B.M., Dubin, M.J., Martin, A., Jones, R.M., 2019. Distinctions among real and apparent respiratory motions in human fMRI data. *Neuroimage* 201, 116041.
- Preti, M.G., Bolton, T.A.W., Van De Ville, D., 2017. The dynamic functional connectome: state-of-the-art and perspectives. *Neuroimage* 160, 41–54.
- Pruim, R.H., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J.K., Beckmann, C.F., 2015. ICA-AROMA: a robust ICA-based strategy for removing motion artifacts from fMRI data. *Neuroimage* 112, 267–277.
- Rangaprakash, D., Wu, G., Marinazzo, D., Hu, X., Deshpande, G., 2018. Hemodynamic response function (HRF) variability confounds resting-state fMRI functional connectivity. *Magn. Reson. Med.* 80 (4), 1697–1713.

- Rosenberg, M.D., Finn, E.S., Scheinost, D., Papademetris, X., Shen, X., Constable, R.T., Chun, M.M., 2016. A neuromarker of sustained attention from whole-brain functional connectivity. *Nat. Neurosci.* 19 (1), 165.
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52 (3), 1059–1069.
- Satterthwaite, T.D., Wolf, D.H., Loughhead, J., Ruparel, K., Elliott, M.A., Hakonarson, H., Gur, R.C., Gur, R.E., 2012. Impact of in-scanner head motion on multiple measures of functional connectivity : relevance for studies of neurodevelopment in youth. *Neuroimage* 60 (1), 623–632.
- Senbabaoglu, Y., Michailidis, G., Li, J.Z., 2014. Critical limitations of consensus clustering in class discovery. *Sci. Rep.* 4, 6207.
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebr. Cortex* 22 (1), 158–165.
- Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., et al., 2013. Resting-state fMRI in the human connectome project. *Neuroimage* 80, 144–168.
- Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D.R., 2012. Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. *Front. Physiol.* 3, 1–12.
- Tana, M.G., Montin, E., Cerutti, S., Bianchi, A.M., 2010. Exploring cortical attentional system by using fMRI during a continuous performance test. *Comput. Intell. Neurosci.* 2010, 329213.
- Tuleasca, C., Bolton, T.A.W., Régis, J., Najdenovska, E., Witjas, T., Girard, N., Delaire, F., Vincent, M., Faouzi, M., Thiran, J., et al., 2019. Normalization of aberrant pretherapeutic dynamic functional connectivity of extrastriate visual system in patients who underwent thalamotomy with stereotactic radiosurgery for essential tremor: a resting-state functional MRI study. *J. Neurosurg.* 1, 1–10.
- 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 (8), 519–534.
- Van Dijk, K.R., Sabuncu, M.R., Buckner, R.L., 2012. The influence of head motion on intrinsic functional connectivity MRI. *Neuroimage* 59 (1), 431–438.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., 2013. The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79.
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci.* 114 (48), 201705120.
- 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.