Reverse inference via connectivity fingerprinting: task sensitivity, task specificity, and task potency

Authors:

Roselyne Chauvin[1,2], Maarten Mennes[2], Jan Buitelaar[1,2], Christian Beckmann[1,2,3]

Institutions:

1Radboud University Medical Center, Department of Cognitive Neuroscience, Nijmegen, The Netherlands,

2Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, The Netherlands

3 FMRIB

Keywords: task, rest, functional connectivity

Number of words: 7125

Number of Figures: 9

Abstract

In recent years, several large-scale neuroimaging efforts have been launched in an attempt to tackle a potential lack of power in the context of small effect sizes. However, within these large-scale efforts different cognitive tasks are typically treated independently, while the availability of such large databases provides much greater potential for integration. Additionally, the field of functional connectivity seems to face redundant results involving a reproducible baseline architecture of networks connectivity and would need to go beyond that representation to understand how cognition emerge from variation of that structure. In this paper, we propose a framework that aim to provide a new feature to study the functional connectivity related to cognitive involvement in tasks by removing the common architecture of the functional connectivity, i.e. taking the resting state as a baseline functional structure. As such we can propose a feature of connectivity modulation by a task: task potency. Supplementary, our framework aims to bring all tasks connectivity to a comparative level to enable to study all tasks of a cohort together. We demonstrate the use of such a feature by comparing three tasks of the NeuroIMAGE cohort (ref) and describing the behaviour of the task potency by looking at commonality and differences between tasks. We then moved to related the task potency as a new feature for the reverse inference investigation.

1. Introduction

Advances in functional brain imaging have provided tremendous insight into the neural correlates of cognition by relating behavioural descriptions to local changes in brain activity or metabolism. Typical experimental studies investigate specific cognitive functions (Geyer 2010, Zilles 2011), and thereby inform about the *sensitivity* of a response within a brain area to the experimental manipulation of interest (Aguirre 2002, Harley 2004). Yet, single neuroimaging studies do not allow making inferences about whether an observed area exclusively responds to cognitive function A or whether it is also sensitive to manipulation of function B. As such, they cannot inform about the *specificity* of a brain area for the tested cognitive function.

To be informative about specificity rather than mere sensitivity and thus allow reverse inference (Poldrack2006, Aguirre), study participants would need to be probed for various cognitive functions across a broad repertoire of domains. However, such multi-paradigm investigations are technically and logistically challenging and therefore remain rare. Accordingly, to indirectly infer specific behavioural relevance for the neural responses they observed authors typically resort to previously reported results via literature meta-analysis or alternative initiatives that implement Bayesian analysis or machine learning techniques to validate study results (Hutzler 2014, Poldrack 2011, Schwartz 2013, Varoquaux 2014). However, such literature-based techniques suffer from typical biases associated with the publication process including article imprecision, 'File Drawer' issues, and p-value hunting (Poldrack 2006).

Here we describe an innovative framework for disentangling sensitivity and specificity. We do this in the context of emerging large cohort functional activation studies that involve multiple experimental fMRI designs allowing within-subject comparisons. Examples include the Human Connectome Project (REF) and the UK Biobank aimed at deciphering the complex relationship between brain functions, cognition, and the functional and structural human connectome within a normal cohort. Similar projects are translating these efforts to the clinical domain (e.g., NeuroIMAGE, PNC, ABCD). We capitalize on the increased statistical sensitivity to effect size differences that within-subject designs offer and

introduce the concept of *task potency*, which characterises response magnitude in relation to experimental manipulations.

Our approach focuses on task sensitivity, task specificity, and task potency at the level of mesoscopic functional connectivity within the brain. Functional connectivity can successfully differentiate between mental states (Shirer W. 2011), corroborating the idea that mental states can be defined based on specific connectivity profiles, similar to a "brain fingerprint" (YEO 2011, NN paper). Accordingly, even more than local brain activity, distributed patterns of brain activity might be the best reflection of mental states and behavioural performance. As an example, a region's task involvement is related to its whole brain connectivity pattern, where a region's connectivity strength with the task-positive network effectively determines its task-related activity (Mennes et al., 2010). In another example, memory-learning performance could be accurately predicted based on whole-brain connectivity, where connections between regions not involved in the task were equally important predictors (Yamashita et al. 2015). This evidence suggests that both general as well as task-specific neural processes support task performance. Corroborating the idea of a functional baseline is the observation that task-related networks can be delineated using resting state data (Smith et al. 2009). Such baseline network structure supports the idea that specific cognitive states are produced by a baseline of common network activity supplemented by specific modulations across the brain's functional architecture, rather than being orchestrated by independent activity in single regions. An important corollary is that cognitive function emerges by embedding unique regional activity within the context of larger network processes, as described in the 'massive redeployment hypothesis' (Anderson 2007, Dan Lloyd 2000).

We aim to disentangle general and task-specific processes by indexing the presence or absence of significant connectivity under different tasks. We do this by quantifying the associated *potency*, i.e., amplitude, of a functional connectivity modulation in relation to experimental manipulations. We hypothesize to observe that general neural processes will occur across multiple tasks but might be small in potency, while specific neural processes will be related to a single manipulation. Of note, a specific cognitive process could be of

importance for solving several tasks. In this manuscript, we consider three tasks, yet the framework is readily extensible to tasks not included here.

We aim to link connectivity changes to regional task specialisation and involvement at the network level by defining functional connectivity fingerprints using a top-down, functionally defined, hierarchical atlas that divides 11 larger networks into 184 smaller regions (Van Oort et al. Arxiv). After robustly defining edges, i.e., connections, that are relevant for a task, our approach allows differentiating experimental designs in their specificity of modulating these connections (i.e., is this change in brain connectivity specifically associated with task A or task B?). In cases of shared sensitivity where an edge is modulated by the presence of multiple tasks we will be able to determine which task is the more potent one. We do this by indexing the deviation in connectivity in each task from a common baseline derived from resting state fMRI data. Using effect-sizes on population distributions of these connectivity deviations, we can quantify the potency of selected functional connections. As such, the amplitude of modulation becomes a feature that defines task specificity as a new tool to do reverse inference and to understand the relative effect of one task within a larger battery of tasks.

2. Methods

2.1. Participants

We use MRI data from the NeuroIMAGE sample (N total > 800 participants; see von Rhein et al., 2014). In the current analyses, we included data from healthy control participants only (initial N=385) who each performed at least one of the following tasks during fMRI scanning: working memory (Spatial WM, REF), reward processing (MID, REF), or response inhibition (Stop Signal Task, REF). In addition, to the task-based MRI scans each participant completed a task-free resting state fMRI scan. fMRI acquisition parameters are shown in table 1. All participants also completed an anatomical scan for registration purposes (T1-weigthed MPRAGE, TR=2730 ms, TE=2.95 ms, T1=1000ms, flip angle=7, matrix size=256x256, FOV=256mm, 176 slices with 1mm isotropic voxels).

fMRI scans exhibiting limited brain coverage or excessive head motion were excluded from further processing. Limited brain coverage was defined as having less than 97% overlap with the MNI152 standard brain after image registration. Applying this criterion excluded 47 subjects (details in table 1). In addition, we excluded from each task those participants who were among top 5% in terms of head motion as quantified by RMS-FD, the root mean square of the frame-wise displacement computed using MCFLIRT (Jenkinson et al, 2002). Applying these criteria resulted in the inclusion of data from 371 healthy controls, comprising 218 resting state acquisitions, 224 stop signal task acquisitions, 236 MID acquisitions, and 257 Spatial WM acquisitions. Participants ranged in age between 8.6 and 30.5; mean=17.4; sd=3.7; 46.1% were male.

2.2. fMRI Preprocessing

All fMRI acquisitions were processed using tools from FSL 5.0.6. (FSL; http://www.fmrib.ox.ac.uk/fsl; Smith et al, 2004; Woolrich et al., 2009; Jenkinson et al, 2012). We employed the following pipeline: removal of the first volumes to allow magnetization equilibration (see table 1), head movement correction by volume-realignment to the middle volume using MCFLIRT, global 4D mean intensity normalization, spatial filtering with a 6mm FWHM Gaussian kernel. We then denoised all preprocessed data for

motion-related artefacts. We used ICA-AROMA to detect motion-related artefacts in single-subject data and subsequently regressed these artefacts out of the data using fsl_regfilt (Pruim 2015a, Pruim 2015b). Finally, we regressed out mean signals from CSF and white matter, and applied a 0.01Hz temporal high-pass filter (Gaussian-weighted least square straight line fit to the data).

For each participant, all acquisitions were registered to its high-resolution T1 image using Boundary-Based Registration (BBR) available in FSL FLIRT (REFS). All high-resolution T1 images were registered to MNI152 space using 12-dof linear registration available in FLIRT and further refined using non-linear registration available in FSL FNIRT (REFS). Transformations were not applied. Instead we used the inverse of the obtained transformations to bring a hierarchical atlas of brain regions to the participant's native space (see below). All analyses were performed in each participant's native space.

Table 1: fMRI acquisitions parameters

	Resting state	Stop Signal Task	MID	Spatial WM
Image acquisition parameters				
General parameters	TE=40 ms, FOV=224mm, 37 axial slices, flip angle=80, matrix size=64x64, in-plane resolution=3.5mm, slice thickness/gap=3.0mm/0.5mm			
N volumes	>260	86 * 4 blocks	>300	107 * 4 blocks
TR in ms: min-max [mean (std)]	1860-1960[1959.5(6.75)]	2150-2340[2338.3(17.87)]	2150-2340[2339.9(12.93)]	2340
N first volumes rejected*	5	4	5	3
Participant characteristics				
N initial	302	239	256	266
N rejected for limited brain coverage	73	4	8	6
N rejected for head motion	11	11	12	13
N used in final analyses	X	X	X	X
RMS-FD min-max	0.026 - 1.930	0.025 - 0.947	0.027 - 0.976	0.029 - 1.504
RMS-FD mean (std)	0.171 (0.224)	0.093 (0.098)	0.145 (0.145)	0.149 (0.191)
Age min - max	8.6 - 30.5	8.6 - 27	9.1 - 30.1	8.6 - 30.1
Age mean (std)	17 (3.5)	17.4 (3.6)	17.6 (3.7)	17.4 (3.6)
% male	45.8%	45.5%	44.5%	47%
Participants who also completed a resting state scan	X	111	123	144

* The number of initial volumes removed from further analyses varied to ensure comparability with earlier studies that used these data. Note that this variation will have very limited impact on the current analyses.

2.3. Connectome atlas

For each functional imaging scan we defined connectivity matrices using regions defined in a hierarchical whole-brain functional atlas (Van Oort et al. Arxiv). This atlas contains 185 non-overlapping regions and was defined through Instantaneous Correlation Parcelation (ICP) as applied to resting state fMRI data of 100 participants of the Human Connectome Project (HCP; REF). In short, ICP aims to parcel larger regions into subregions based on signal homogeneity, where the optimal number of subregions is determined based on split-half reproducibility at the cohort level.

Figure 1 illustrates the hierarchical brain atlas, where areas were grouped in 11 higher-level networks: 9 resting state networks (visual1, visual2, motor, right attention, left attention, auditory, default mode network (DMN), fronto-temporal and cingulum), and 2 anatomical structures (subcortical areas, and cerebellum). These higher-level networks respectively contained 19, 12, 22, 22, 18, 8, 18, 13, 7, 23, and 23 subregions.

Connectivity matrices were calculated in each participant's native space. To this end, we transformed the atlas to each participant's native space using the inverse of the anatomical to MNI152 non-linear warp, and the inverse of the linear transformation of the functional image to the participant's high resolution anatomical image. Voxel-membership in brain parcels was established on the basis of majority overlap. Areas that were on average across our population over 50% outside of the brain were rejected from further analyses. This resulted in the rejection of one area in brainstem. For consistency, we removed the 5 others brainstem areas. As a result, we used 179 areas to compute connectivity matrices, as explained in section 2.4.

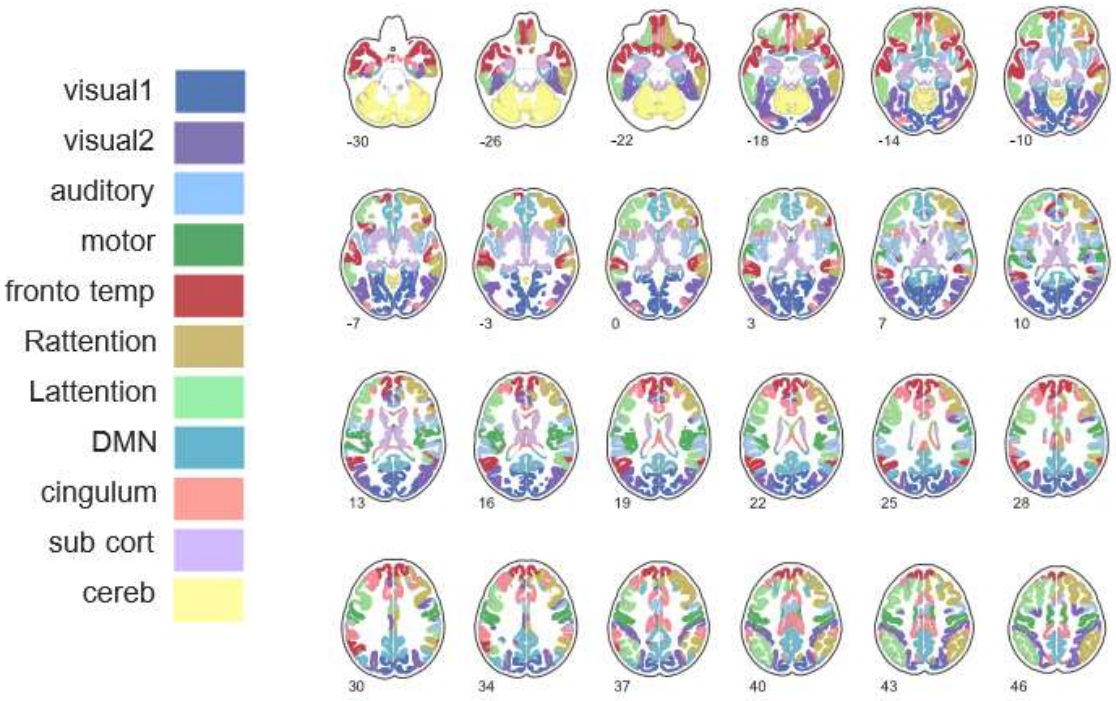

Figure 1 : ICP atlas with 179 areas represented in their corresponding overarching networks.

2.4. Connectivity calculation

For each participant and each task (rest, Spatial WM, MID, Stop) we calculated 179x179 connectivity matrices, by cross-correlating the time series of all regions in the atlas. We obtained each region's time series through multivariate spatial regression, using all 179 regions as regressors and each task's preprocessed time series as dependent variable. The resulting regional time series were demeaned. For the Spatial WM and Stop task we temporally concatenated the time series of individual runs. Using these time series, we calculated 179x179 partial correlation matrices through inverting covariance matrices estimated by the Ledoit-Wolf normalization algorithm (Ledoit and Wolf, 2004) as implemented in nilearn (http://nilearn.github.io/). Finally, all pair-wise correlations were Fisher r-to-Z transformed.

To allow comparing connectivity values between acquisitions, we normalized the connectivity values within each matrix to fit a Gaussian distribution. Yet, we were cautious not to affect the tails of the connectivity distributions as these represent the most

interesting connectivity values. Therefore, we modelled the obtained connectivity values per task using a Gaussian-gamma mixture model (ref) and used the main Gaussian, i.e., the one fitting the body of the distribution, to normalize our connectivity values. In practice, we applied the mixture modelling to the upper triangle values of each connectivity matrix and subsequently normalized the connectivity values by subtracting the mean and dividing by the standard deviation of the obtained Gaussian model. As a result, the values within the normalized, z-transformed partial correlation matrices are readily comparable across participants and tasks.

To allow interpretation of the task-based connectivity matrices in terms of their deviation from the resting state baseline connectivity we standardized each participant's task-based connectivity matrix. Specifically, we standardized each individual-level pair-wise correlation by subtracting the mean and dividing by the standard deviation obtained for that pair-wise correlation during rest across all participants that completed a resting state scan. As such, each task-based pair-wise correlation or edge quantifies how connectivity for that edge differed from that edge's connectivity during the resting state. After standardization we obtain for each participant an individual connectivity matrix for each of its task acquisitions. We refer to these matrices as 'task potency' matrices. For each task, we finally create group-level task potency matrices by averaging across participant matrices.

2.5. Task-based Fingerprints

To focus on those connections that characterize a task's functional fingerprint we selected for each task those edges that showed a relevant deviation from rest. To this end, we converted the group-level task potency matrices to Z-statistic matrices by subtracting the mean and dividing by the standard deviation calculated for each task matrix. For each task we selected those edges where $|Z| >= 2.3$. We refer to those edges as task-based fingerprints. The task-based fingerprints form the basis to define task sensitivity, task specificity, and task potency of each edge in the fingerprint.

2.6. Network-based summary metrics

Our hierarchical atlas defines 179 areas as a subdivision of 11 large-scale networks. As such, we can average across all edges within each network to summarize sensitivity, specificity, and potency at the network level. Additionally, we can differentiate edges that link areas within a network (within-network edges) from edges that link areas between two different networks (between-network edges). Importantly, as our networks are derived from networks characterized in the literature as task-related (Smith et al. 2009), we can interpret our connectivity-based fingerprints in the context of specific cognitive networks. By comparing the within- and between-network connections we assessed whether a specific task involved specific networks or resulted in an overall, diffuse potentiation of connectivity. In practice, to derive network-level scores, we calculated the percentage of selected edges included in each network. This was done for each entry in the 11x11 network connectivity matrix, and allowed quantifying the selection of edges at the within- (diagonal matrix entries) and between-network (off-diagonal matrix entries) level.

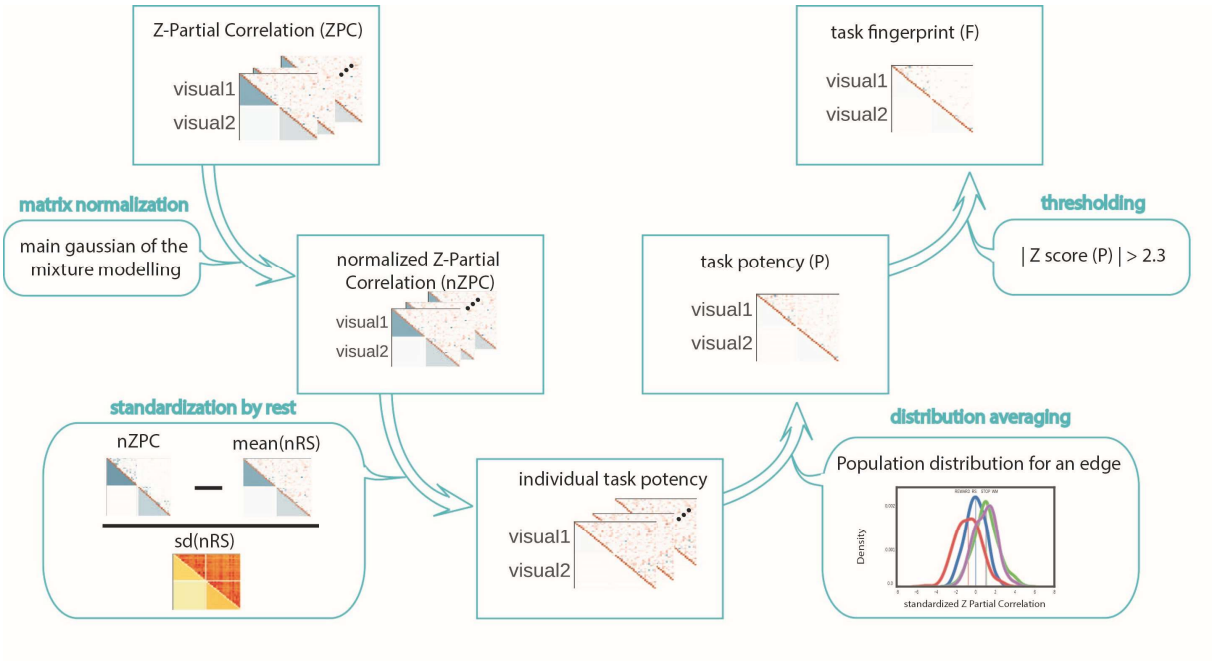

Figure 2 framework to obtain connectivity-based task fingerprints. The framework starts at the participant-level with obtaining a partial correlation matrix (Fisher-Z transformed), which is normalized, and subsequently standardized by population-based resting state connectivity. A group task fingerprint can be obtained by averaging the individual task potency matrices and thresholding based on the z-score of the group potency.

2.7. Sensitivity, Specificity, and Task Potency

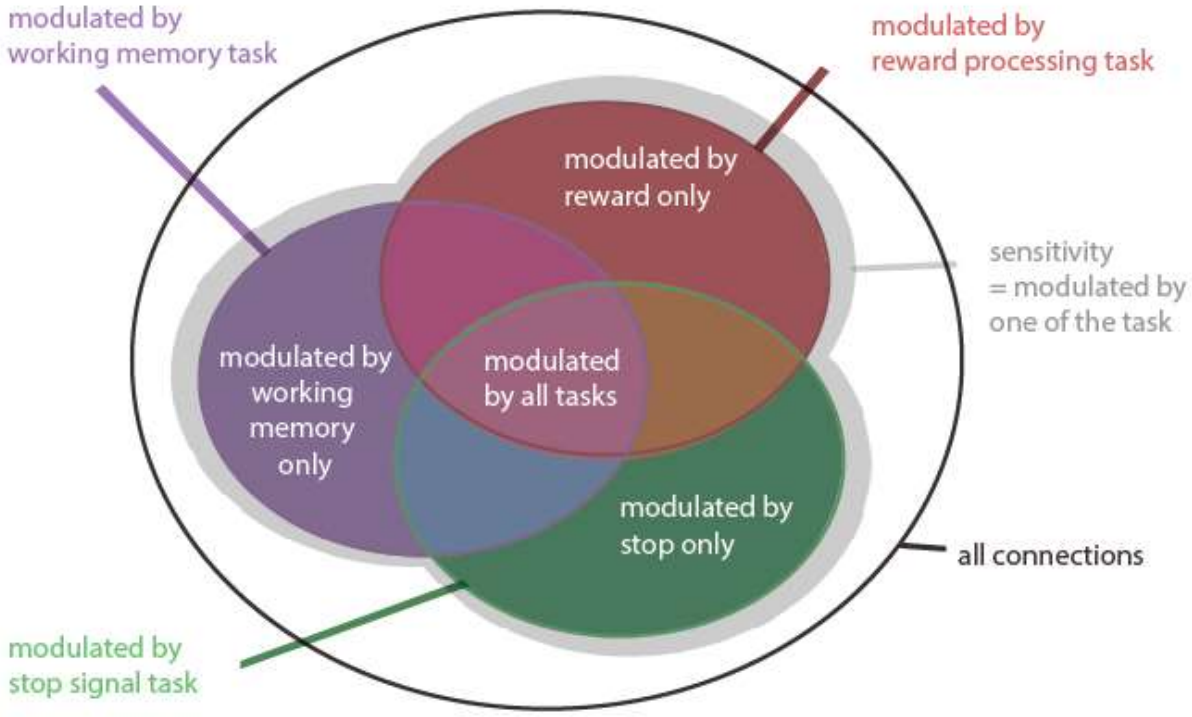

Figure 3 Overview of sensitivity and specificity as applicable to each edge within a task potency matrix.

Figure 3 illustrates how we can characterize each edge within the task potency matrix in terms of its sensitivity and specificity to the different tasks included in our investigation. An edge was regarded ***sensitive*** to task modulation when selected (i.e., above the |Z|>=2.3 threshold) in at least one of the tasks. We computed an overall sensitivity matrix that indexed for each edge the number of tasks that it was sensitive to (ranging from 0 to 3). An edge can be sensitive to modulation by several tasks, yet with a different level of potency in one task compared to another. This is not considered when assessing sensitivity.

Task ***specific*** edges were those edges that were selected for one task only. Note that specificity is determined by the collection of available tasks. The bigger the difference in cognitive constructs between tasks, the easier it is to define specificity of involved edges.

Within the edges showing common sensitivity we aimed to differentiate tasks based on their ***potency*** to enlist a certain edge. To this end, we assessed whether one task was significantly more potent than others by comparing the distribution of task potency values

for that edge across participants, between tasks. This comparison relies on the normalization of all matrices prior to the standardization by the resting state connectivity (see Figure XX) to remove possible differences due to e.g., different length of an individual task experiment or possible effects related to the complexity of the task design. We quantified the difference in potency between sensitive tasks by computing Cohen's *d* between the obtained distributions. We required a minimum difference of 0.3 in Cohen's *d* in order to decide which task specifically potentiated an edge. If all three tasks similarly potentiated an edge it was labelled as "common". In contrast, if two tasks similarly potentiated an edge (i.e., the difference in Cohen's *d* between them was <0.3), the edge was labelled as "undefined", as it is neither specific to one task, nor common to all. The same method was used to label task preference in the 11-network framework. In this case, the input were edges from the 11x11 matrices for each participant corresponding to the sum of potency across edges for each network as selected from the group-level matrices. Finally, we investigated task preference at the level of areas (i.e., columns in our connectivity matrices). To this end, we summed potency for group-level selected edges selected across an area's potential 179 connections.

As such, we assessed task potency at the level of individual edges, across edges and areas involved in larger networks, and across edges involving a specific area. This will help illustrating various potency profiles; for instance, a task can strongly potentiate 1-2 specific edges or, can result in general potentiation of connections involving specific areas, or be confined to within-network connections.

2.8. Reproducibility of the edge selection procedure

Every single participant has associated within-subject differences relative to the cohort-derived network. In order to assess reproducibility of our group-level fingerprint pattern, we defined individual task fingerprints applying the same selection procedure as above, but applied to the individual task potency matrices (i.e., select those edges with a $|Z|>=2.3$ in the individual task fingerprints). This enables us to quantify subject-specific variations in the edge selection and thereby permits quantification of reproducibility across participants. We indexed the number of times an edge was selected across participants. This proportion is interpreted as the reproducibility of an edge's potency.

3. Results

3.1. Task-Based Fingerprints

Overall, the connectivity structure derived from each task was very similar to the connectivity structure derived from the resting state scan. Figure 4 illustrates this for the reward task (see supplement for the other tasks). To evaluate connectivity sensitivity to each task we created task-based fingerprints by standardizing the task connectivity by the resting state connectivity, resulting in a matrix quantifying each edge's functional potency, as illustrated in Figure 4c. This task-based fingerprint served as the basis to identify sensitive edges and assess their task specificity and potency.

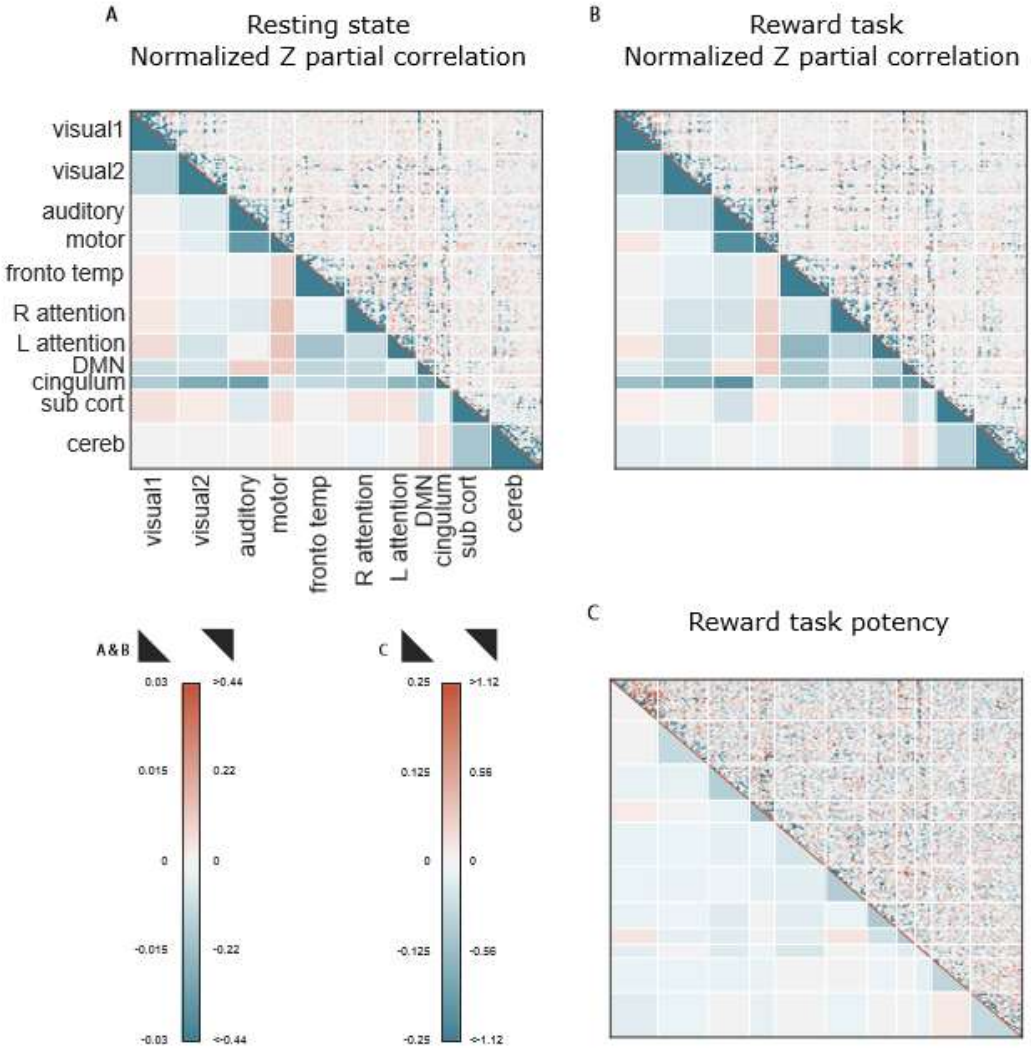

Figure 4 from left to right average normalized resting state ZPC, average normalized Reward task ZPC and average reward task potency. Upper triangle displays the 179x179 connectivity fingerprints; lower triangle displays average summary per network.

3.2. Task Sensitivity

Using each task's functional fingerprint, we determined which edges were sensitive to modulation by that task by selecting those edges that showed a z-score above 2.3 in the average potency across participants. Figure 5 illustrates the characteristics of the selected edges across the three different tasks.

We observed a higher prevalence of sensitive edges for within-network connections compared to between-network connections. As expected, especially the visual1 and motor networks showed a particularly high percentage of within-network sensitive edges. In contrast, the number of sensitive between-network connections was considerably lower, with on average only 3.55% of edges selected across the 11 networks versus 29.5% of within network connections. We highlight the result for the cingulum network which included twice as many sensitive between-network connections compared to the other networks.

When comparing the percentage of sensitive edges between tasks we observed that overall all tasks yielded similar percentages of selected edges. Yet, we observed two exceptions: the within-network DMN and cingulum connections, where respectively the Stop and Working-memory tasks did not yield any sensitive within-network edges.

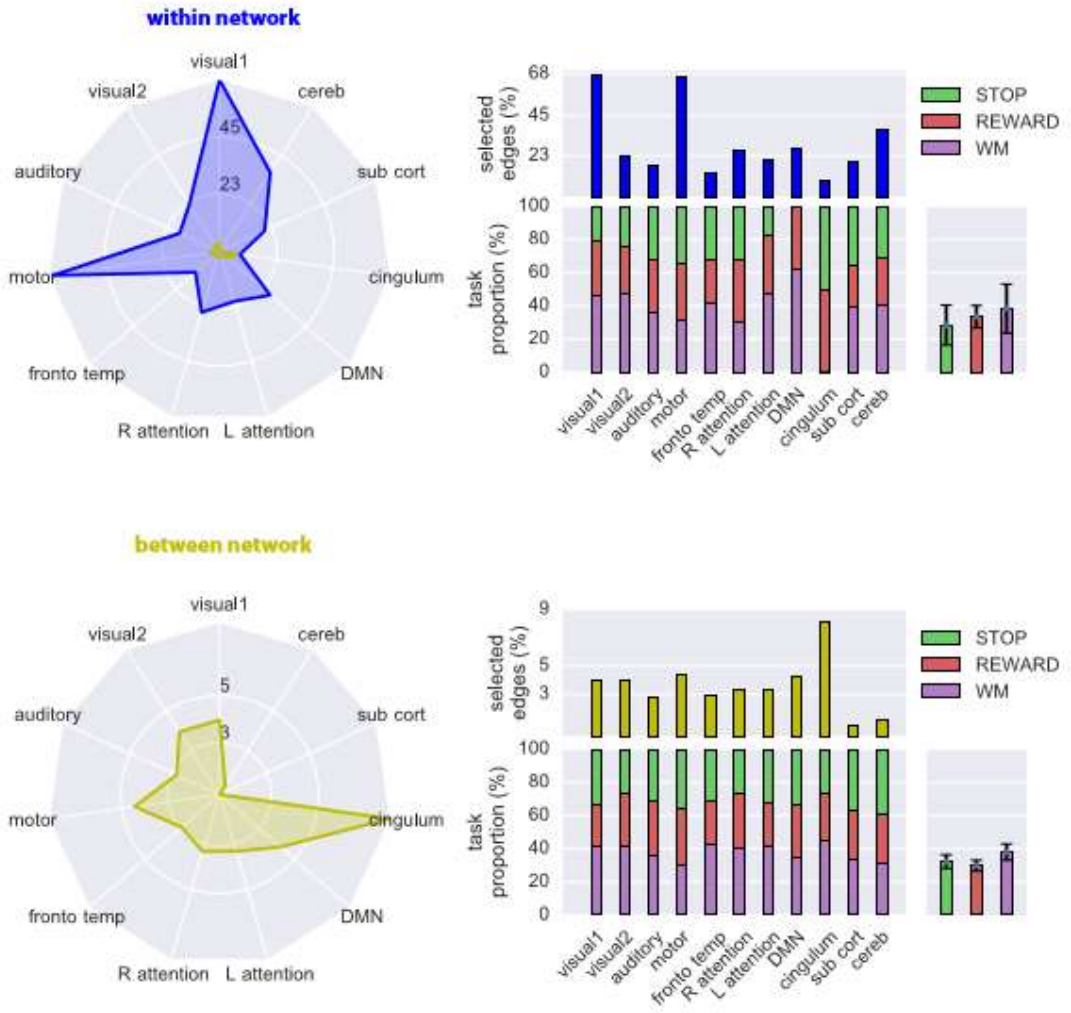

Figure 5 radar plot of the percentage of edges showing connectivity changes (sensitivity) for each of the 11 networks. We observed a larger percentage of sensitivity for within-network connections (top), compared to between-network connections (bottom). Bar plots on the right illustrate edge sensitivity for each task.

Overall, we observed greater sensitivity to task modulation in within- compared to between-network connections. This result suggests that during task performance there is increased change in communication within a network, rather than connecting regions across the brain. This result supports the hypothesis that the brain strongly segregates information at the level of individual networks, while more lightly integrating information between networks (Deco 2015). In the context of reverse inference, these results support the idea

that task activity builds on the resting state architecture, preferentially modulating connections within well-connected networks. However, as indicated before, going beyond overall statements requires assessing the task-specificity of each sensitive edge.

3.3. Task Specificity

To disentangle overlapping connectivity modulation by the different tasks (i.e., reverse inference), we define the task specific edges by dividing the obtained collection of sensitive edges into those that were modulated by one task only (i.e. specific to a certain task) versus those that were sensitive to modulation by several tasks (figure X methods), and in particular edges modulated by all tasks.

Figure 6 illustrates how the sensitive edges modulated by one task only or by all tasks were represented in the 11 networks. Across both the within- and between-network connections we observed that on average 29.65% of the sensitive edges were specific to a certain task, versus 13.1% that were modulated by all tasks. Note that this also means that 57.25% of sensitive edges was modulated by more than one, yet not all, tasks.

On average, relative to the number of sensitive edges, the number of specific within- and between-network edges was fairly similar. The average ratio of specific and common within-network connections was 1.75(+/- 1.24) and was 3.2(+/- 1.42) for between-network connections. Importantly, we did not detect specific within-network connections for the cingulum network.

However, when assessing the specificity for each particular task, the within- and the between-network connections yielded a different pattern. While the between-network connections were almost equally distributed between tasks, we observed greater variability in task specificity for the within-network connections, with some networks showing a notable task-specific response. For reward, the left attention and motor networks showed higher specificity of within-network modulations, while the fronto-temporal and default mode network did not exhibit any reward-specific within-network modulation. For the stop signal task, the right attention, fronto-temporal and motor networks exhibited most specific connectivity modulation. Similar to the reward task, we observed no specific connectivity

modulation within the DMN. For the working memory task, we observed exclusive specificity of within-DMN connectivity modulations. In contrast, the working memory task did not modulate specific connections within the motor and right attention networks.

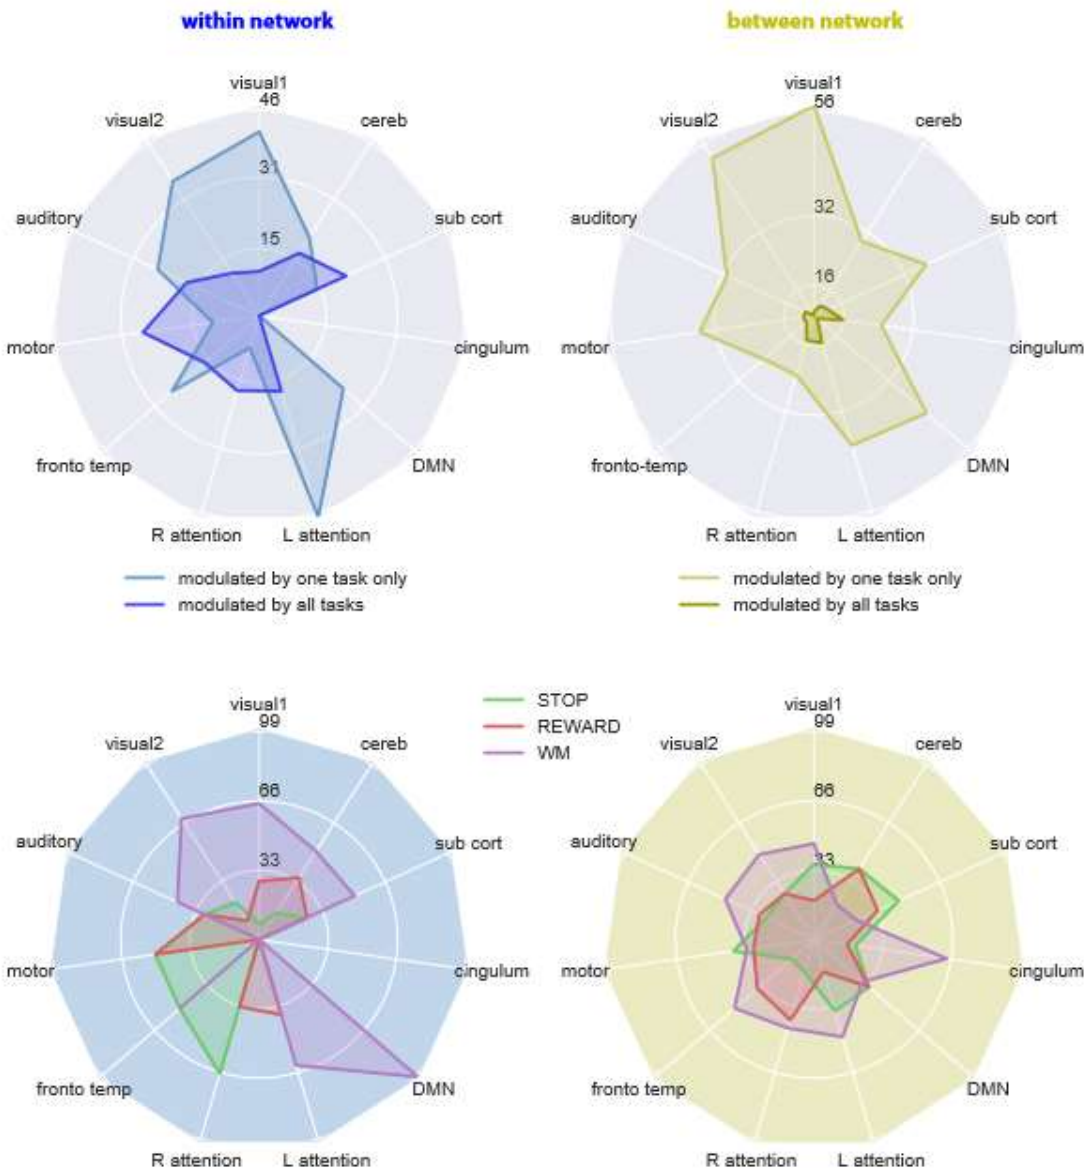

Figure 6 percentage of sensitive edges for each task that end up to be specific to this task, e.g., percentage of edges selected in the group potency of a task that are not present in another task group potency selection. Percentage corresponding to the within- vs between-network connectivity is differentiated per network. As an example: ~38% of the sensitive edges within the visual1 network were modulated by one task only. Of those 38% sensitive edges, about 60% was modulated only during WM performance. In contrast, ~10% of the sensitive edges in the visual1 network were modulated during performance of all tasks.

3.4. Reproducibility of the selection across individual fingerprints

The result above shows that different tasks exhibit a specific pattern of network potentiation, which can be accessed by comparing a cohort of tasks. However, a large proportion of sensitive edges are represented in all tasks. To estimate whether this method can help to evaluate individual task fingerprints in a reproducible manner, we studied the detection rate of edges sensitive to one or all tasks across each individual's task fingerprints.

Specifically, we investigated whether the group selection was reproducible at the individual level and in particular how well the task-specific connections where represented at the individual level. We defined the individual fingerprint by selecting edges that showed a Z score higher than 2.3 in the individual-level task fingerprints. We computed the sum of selected edges across the population for each task. We observed a set of edges with high selectivity across participants for each task: 2.5% of edges within the union of individual mask were selected by minimum 15% of the population and in each of the three tasks. As shown in Figure 8, those edges mainly linked homotopic areas of each hemisphere, including bilateral motor areas, cerebellum, attention networks, visual 1 areas, and bilateral putamen (see figure 8). 88,8% of this set of edges were also selected in the group-level selection, and 45.8% of this set of edges were modulated by all tasks (see figure 7).

In contrast to the set of highly selected edges, we note that the individual selection showed important variability: 88.5% of edges were selected in less than 10% of the individual fingerprints (figure 7 black dashed line). This variability is interesting as we can investigate parameters influencing the connectivity response to a task. The most consistently selected edges between participants involved connections sensitive to all tasks as most of the most reproducibly selected edges belong to edge modulated by all tasks. Whereas, highest inter-individual selection variability regarded task-specific edges as edges selected in only one tasks at the group level show a lower individual selection reproducibility than edges selected in all tasks at the group level (figure 7).

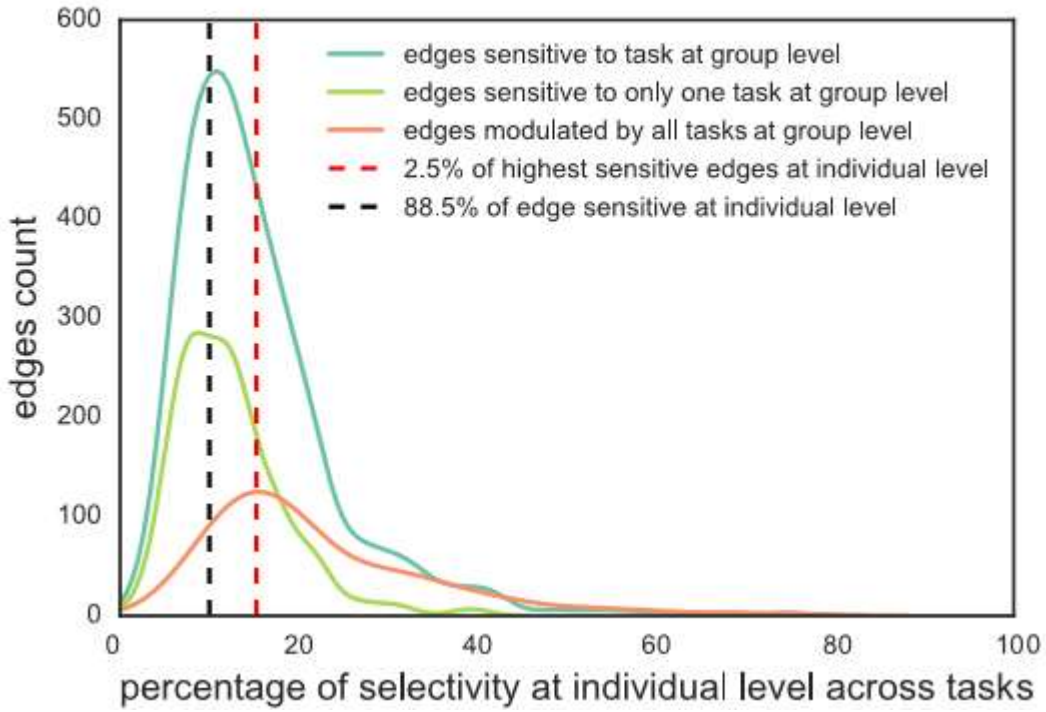

figure 7 distribution of selectivity of edges showing any sensitivity at the individual level, and corresponding occurrence of edges selected at the group level and occurrence of edges modulated by all task at the group level

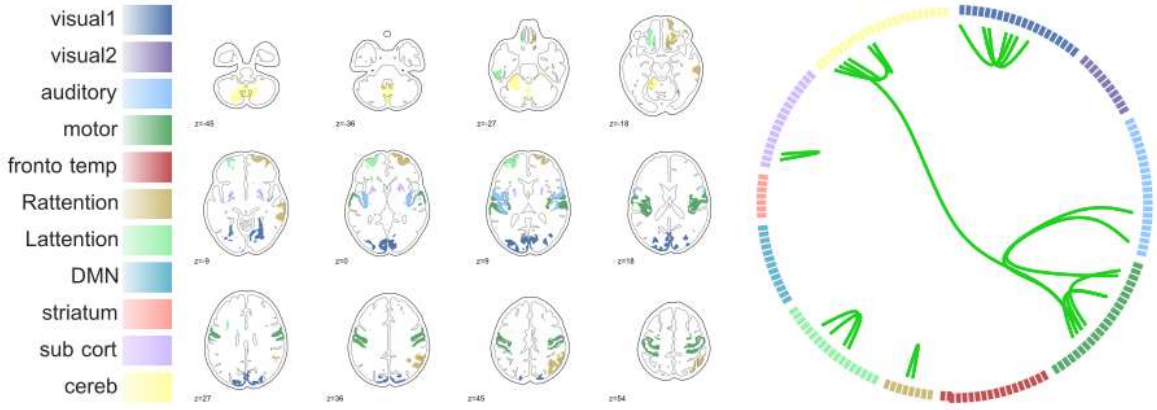

figure 8 edge selected over 20% of subjects across tasks condition (right brain slices) and areas corresponding to them in the 4x3 brain slices, the color of area and sphere correspond to the networks where they belong to (legend on the left)

3.5. The most potent task

Even if each task specific modulation of connectivity reflects network enlistment by a specific experiment, this information is only accessible in comparison to other tasks. Additionally, to define edges as tasks specific brings limitations as soon as the number of task increases or when tasks share cognitive processes differently involved in the experimental design. Therefore, we propose to move from the binary concept of specificity and sensitivity to a measure of connectivity potency to enlist a connection, a network, or an area of a task.

Using the potency as a representation of the strength of enlistment of connections in a task, we can characterise which task potentiated an edge, an area, or a network most strongly (if at all). To this end, we assessed whether a task showed a significantly higher modulation of a given feature (i.e. an edge, an area, or a network) relative to other tasks by computing the Cohen's *d* for comparing the population distributions of task potency of a feature between the different tasks. The task with the highest potency and a significant cohen's *d* difference with other tasks was selected as most potent task. For edges, we investigated the potency distribution of each edge separately (fig 9, upper triangle of the matrix). For networks (fig 9, lower triangle of the matrix) and areas (fig 9 brain slices), we looked at the sum of potency across selected edges within a task for a network or involving a given area.

Sensitive connections were mostly potentiated by a specific task, instead of showing common potentiation by all tasks. 51,3% of sensitive edges were modulated by one task only vs 29,2% of edges modulated by all tasks. In edges that were modulated by several tasks, 15,7% was modulated more by one task compared to the others: 69,2% by the working memory task, 20,1% by the reward task, and 10,6% by the stop signal task.

At the area and network level, the working memory task offered strong potentiation across the brain. The reward task principally potentiated the right attention networks and areas as well as areas from the reward circuit (anterior cingulate cortex, prefrontal area,

thalamus). As for the stop task, connections involving the visual 1 network and cerebellum are mainly potentiated.

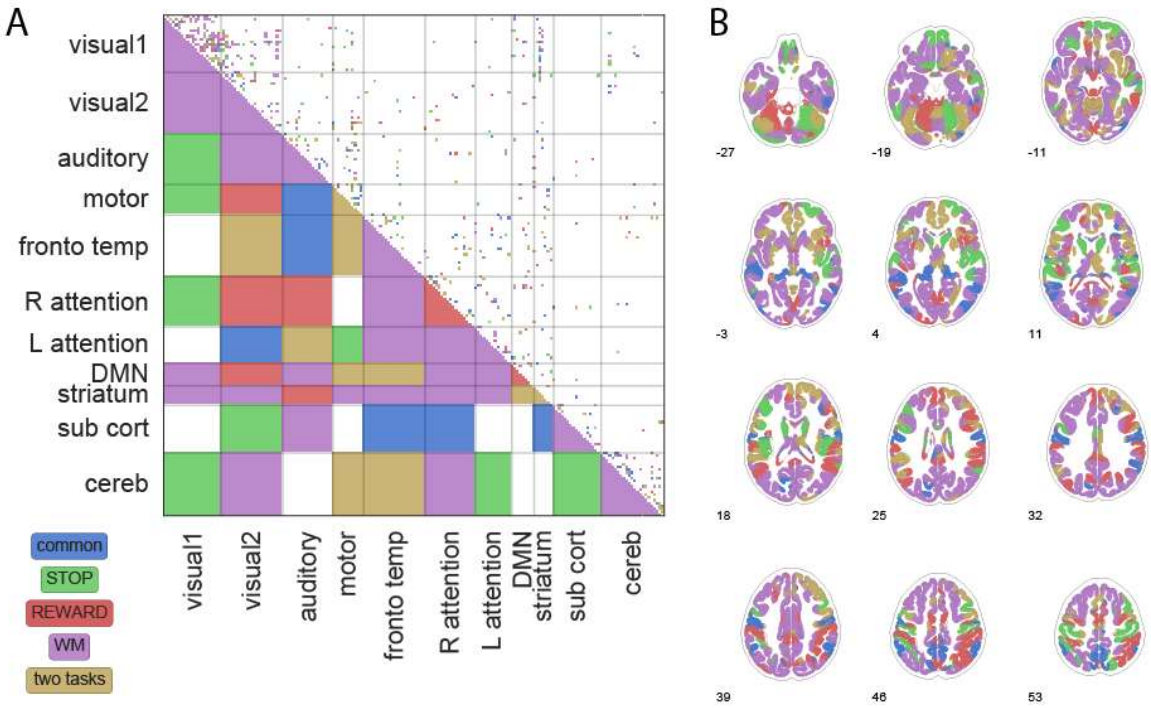

Figure 9 The most potent task illustrated for areas (left), edges (right, upper triangle), and networks (right (lower triangle). The brain slices on the right illustrate for each area which task on average potentiated edges involving this area the most. The same is shown for individual edges and at the network-level in the matrix on the right. Next to being most strongly potentiated by one specific task, we observed potency that was common to all tasks (i.e., no differentiation between tasks), or significant for some but not all tasks (undefined).

4. Discussion

Describing the task-related fingerprint at the network level

Studies showed that portion of the task functional connectivity is different from the resting state functional connectivity (Cole 2014, geerligs 2015), showing task-specificity (shirer 2012). However, so far, studies did not assess involvement of networks neither characterise the relevance of the location of edges modulations for tasks.

In this study, sensitive edges to task modulation were mainly located between areas within the same resting state network, with a particular involvement of motor, visual 1, and the

cerebellum networks across task. As all of our tasks involved the perception of stimuli on a screen as well as a motor response, this result is expected, however it also suggests that lower cognitive functions (as perception and action) request more connectivity modulation. This larger connectivity modulation can highlight a more straightforward and automatic response to incoming stimuli and a standardize motor activity. After all, visual and motor area show a highly constrained organization highly evolutionally conserved reflecting in a lower variability between individual (Mueller 2014) but higher flexibility within individual (Laumann 2015).

A second result in line with the literature of functional connectivity is the prevalence of task modulation within resting state networks (cole 2014). We also can say now that task modulates more the within-network connectivity than between-network one. One would say it modulates more the integration than the segregation of information.

Comparing results from tasks of a cohort

The comparison of three tasks that assess distinct cognitive domains (memory, inhibition, and reward processing) allowed us to distinguish connections that were specific to each task, as well as those common to all tasks. A large portion of connections was shared between tasks even if these tasks belong to distant cognitive domain. However, the observed task specificity is consistent with existing literature. For example, the high specificity of DMN connectivity in the working memory task is in accordance with the idea that areas within the DMN are involved in working memory [REFS]. However, the specificity of a task is defined relatively to our set of tasks from the same cohort (here 3). As such, this definition cannot be extended for reverse inference purposes yet. However most potent task to initiate connectivity response can be defined with the task potency knowledge according to ones' hypotheses. This output is relevant to optimize experimental designs.

Investigating reverse inference

The modulation of an edge during a task is related to the task processing. We can hypothesis that if tasks are cognitively related, the task-modulation connectivity from the baseline, i.e. task potency, will behave as on a continuum. However, if edges are involved in

the realisation of two independent cognitive function, the task potency may present two distinct values, one for each cognitive process. Led by this idea, we investigated the potency of task to involve edge differently. We compute a Cohen's d between task distribution of value for a given edge sensible to more than one task. Some edges demonstrate to be able to present different states, i.e. they present distinguishable modulation values in different tasks. This result supports the proposal to integrate dimensionality to define building blocks of cognition.

We outlined a novel framework for establishing reverse inference. Using a large dataset including multiple tasks, we propose to identify which task is most potent in modulating an edge. Task potency is defined by comparing connectivity extracted from a task acquisition to baseline connectivity extracted from a resting state scan. Our approach to compare each task against a resting baseline is different from current implementations that try to quantify how specific a given task-activation pattern is for a cognitive function by mining available literature to resolve the reverse inference problem.

Individual variability in task potency

We looked at the edges modulated by the task at the individual level. We observe a great variability in the selectivity of connections, except for a set of edges that mainly belonged to edges modulated by all tasks at the group level selection. Theses edges are linking mainly bilateral connections. These interhemispheric connections already show a strong resting state and underlying structural connectivity support (Hermundstad - 2013). Now we can also say that they are strongly modulated, in a reliable way in all tasks. We aim to move away from this high reproducible connectivity to capture individual differences linked to task performance.

Using the task potency to predict individual differences

Functional connectivity of tasks has been used to categorize tasks (shirer 2012), to predict performance or IQ (Schultz 2016) or tasks activity (tavor 2016). The task-functional connectivity seems a good candidate to understand individual variability in performance or individual characteristics. However, a better understanding of how functional connectivity

modulation lead to task performance is needed to extend the use of the task-functional connectivity prediction. Understanding the task-related modulation would enable modelling it to predict functional connectivity from individual that deviate from the norm, e.g. pathological response or alternative strategy use to respond to a task. In contrast to the method used in Saad et al to use the individual resting state connectivity to predict task connectivity, we aim to remove this strong individual fingerprint to access the task specific, reproducible characteristics. Once this task connectivity modulation characterises, we would be able to understand why variability in the resting state can predict task connectivity.

In this study, we investigate a standardization of task connectivity by considering each edge distribution of value in the population, we used population baseline. One limitation of this current implementation of this method is that we consider all resting state as equal. When studying differences between groups or variability between individual, we aim to consider a non-uniform population, having non-equal connectivity baseline. We can define group baseline only if we are leading group comparison. However, it will enhance differences between group due to a preloaded group information. A group standardization will interact with data-driven analysis like machine learning. To avoid constraining individual connectivity to fit into a group one, we can consider individual baseline (i.e. regressing an individual's resting state from his task-connectivity). Analysis has been realized and reproducibility of results can be found in the supplementary materials. The main difference between the two methods is that this other method enhance individual differences. Indeed, variances for each edge potency in the population are significantly higher in every task by subtracting the resting state connectivity compare to standardizing by a population distribution of resting connectivity. This other method can be of great interest for clinical application, as it magnifies differences in the sample and because we do not expect clinical subgroup to share the same resting state characteristics that control subjects.

To further validate the use of the task-potency, we would benefit from linking the task modulation from task performance. Knowing that competition between DMN and task-positive networks is related to performance in a task (Kelly 2008), we could hypothesis that the task-potency would maximise this effect. Similarly, Cohen et al (2014) showed that optimal integration between cingolu-opercular and fronto-parietal improves behavioural

accuracy and cognitive control. With the task-potency, we would expect to observe this modulation without a priori hypothesis on location of connectivity modulation. The interest of our framework is to provide an alternative way of computing a networks-based task-fingerprint to enhance single subject differences and connectivity modulation by task. This connectivity modulation enhancement facilitates the study of neural theory behind cognitive function.

References:

Aguirre, G.K., Detre, J.A., Zarahn, E., and Alsop, D.C. (2002). Experimental Design and the Relative Sensitivity of BOLD and Perfusion fMRI. NeuroImage *15*, 488–500.

D'Esposito, M., Ballard, D., Aguirre, G.K., and Zarahn, E. (1998). Human Prefrontal Cortex Is Not Specific for Working Memory: A Functional MRI Study. NeuroImage *8*, 274–282.

Hutzler, F. (2014). Reverse inference is not a fallacy per se: Cognitive processes can be inferred from functional imaging data. NeuroImage *84*, 1061–1069.

Poldrack, R.A. (2006). Can cognitive processes be inferred from neuroimaging data? Trends in Cognitive Sciences *10*, 59–63.

Poldrack, R.A. (2011). Inferring mental states from neuroimaging data: From reverse inference to large-scale decoding. Neuron *72*, 692–697.

Schwartz, Y., Thirion, B., and Varoquaux, G. (2013). Mapping cognitive ontologies to and from the brain. arXiv:1311.3859 [cs, Q-Bio, Stat].

Varoquaux, G., and Thirion, B. (2014). How machine learning is shaping cognitive neuroimaging. GigaScience *3*, 28.

Aguirre, G.K., Feinberg, F.T.E., (eds, M.J.F., and Neurology, B. Functional Imaging in Behavioral Neurology and Cognitive Neuropsychology.

Cohen, J.R., Gallen, C.L., Jacobs, E.G., Lee, T.G., and D'Esposito, M. (2014). Quantifying the Reconfiguration of Intrinsic Networks during Working Memory. PLoS ONE *9*, e106636.

Cole, M.W., Reynolds, J.R., Power, J.D., Repovs, G., Anticevic, A., and Braver, T.S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. Nat Neurosci *16*, 1348–1355.

Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., and Petersen, S.E. (2014). Intrinsic and Task-Evoked Network Architectures of the Human Brain. Neuron *83*, 238–251.

Hermundstad, A.M., Brown, K.S., Bassett, D.S., Aminoff, E.M., Frithsen, A., Johnson, A., Tipper, C.M., Miller, M.B., Grafton, S.T., and Carlson, J.M. (2014). Structurally-Constrained Relationships between Cognitive States in the Human Brain. PLoS Comput Biol *10*.

Rissman, J., Gazzaley, A., and D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. NeuroImage *23*, 752–763.

Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., and Greicius, M.D. (2012). Decoding Subject-Driven Cognitive States with Whole-Brain Connectivity Patterns. Cereb Cortex *22*, 158–165.

Mennes, M., Kelly, C., Colcombe, S., Castellanos, F.X., and Milham, M.P. (2013). The Extrinsic and Intrinsic Functional Architectures of the Human Brain Are Not Equivalent. Cereb Cortex *23*, 223–229.

Yamashita, M., Kawato, M., and Imamizu, H. (2015). Predicting learning plateau of working memory from whole-brain intrinsic network connectivity patterns. Sci. Rep. *5*.

Cole, D.M., Smith, S.M., and Beckmann, C.F. (2010). Advances and Pitfalls in the Analysis and Interpretation of Resting-State FMRI Data. Front Syst Neurosci *4*.

Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol *106*, 1125–1165.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci U S A *106*, 13040–13045.

Harley, T.A. (2004). Does Cognitive Neuropsychology have a future? Cogn Neuropsychol *21*, 3–16.

Geyer, S., Weiss, M., Reimann, K., Lohmann, G., and Turner, R. (2011). Microstructural Parcellation of the Human Cerebral Cortex – From Brodmann's Post-Mortem Map to in vivo Mapping with High-Field Magnetic Resonance Imaging. Front Hum Neurosci *5*.

Zilles, K., and Amunts, K. (2010). Centenary of Brodmann's map — conception and fate. Nature Reviews Neuroscience *11*, 139–145.

A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices", Ledoit and Wolf, Journal of Multivariate Analysis, Volume 88, Issue 2, February 2004, pages 365-411.

Hermundstad, A.M., Bassett, D.S., Brown, K.S., Aminoff, E.M., Clewett, D., Freeman, S., Frithsen, A., Johnson, A., Tipper, C.M., Miller, M.B., et al. (2013). Structural foundations of resting-state and task-based functional connectivity in the human brain. PNAS *110*, 6169–6174.

Mueller, S., Wang, D., Fox, M.D., Yeo, B.T.T., Sepulcre, J., Sabuncu, M.R., Shafee, R., Lu, J., and Liu, H. (2013). Individual Variability in Functional Connectivity Architecture of the Human Brain. Neuron *77*, 586–595.

Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.-Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U.F., et al. (2015). Functional System and Areal Organization of a Highly Sampled Individual Human Brain. Neuron *87*, 657–670.