

LARGE SCALE META-ANALYTIC CARTOGRAPHY OF HUMAN FRONTAL
CORTEX

by

ALEJANDRO DE LA VEGA

B.A., Pomona College, 2009

M.A., University of Colorado, 2012

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Philosophy
Department of Psychology and Neuroscience
2016

This thesis entitled:
Large-scale meta-analytic cartography of human frontal cortex
written by Alejandro de la Vega
has been approved for the Department of Psychology and Neuroscience

Marie Banich

Tor Wager

McKell Carter

Scott Vrieze

Rafael Frongillo

Date_____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

ABSTRACT

De La Vega, Alejandro Isaac (Ph.D., Psychology and Neuroscience)

Large-scale meta-analytic cartography of human frontal cortex

Thesis directed by Professor Marie T. Banich

The field of human brain mapping has made immense progress in recent years by making tens of thousand associations between the brain and psychological states using functional magnetic resonance imaging (fMRI). However, there is a growing appreciation of the limited ability to determine the specificity between brain-cognition mappings in individual studies. Without surveying a diverse range of psychological states, it is difficult to know if a brain region is preferentially recruited by a given state, or a more domain-general process that underlies it. In a related issue, several recent efforts have attempted to find the fundamental computational units of the brain by using statistical learning techniques to identify discrete regions on the basis of properties that constrain information processing, such as connectivity. However, it's not clear how well these brain atlases describe the high-level functional organization of the brain.

In this dissertation, I apply relatively unbiased data-driven methods to a database of nearly 12,000 fMRI studies to comprehensively map psychological states to discrete regions in human frontal cortex— a complex, high-level association area of the brain. On the basis of activation patterns across studies, I identify functionally distinct whole-brain networks composed of spatially contiguous subregions. While each network exhibits distinct functional associations, subregions

within each network show much more similar, yet dissociable profiles. In contrast with strong localizationist accounts, we find distributed associations between psychological states and brain anatomy, suggesting moderate functional selectivity in many parts of frontal cortex.

In the last section, I quantitatively assess various approaches for clustering the brain into discrete regions by comparing novel meta-analytic atlases to existing brain atlases from other brain modalities. Across a variety of metrics, I find evidence that meta-analytic atlases are robust and may provide a better account of the task-dependent organization of the brain than atlases from other brain modalities. I conclude by discussing future approaches for using large-scale meta-analysis to better understand how the brain gives rise to psychological function.

To my grandmother, Noemi

ACKNOWLEDGMENTS

The inherently isolated nature of doctoral training belies the vast support network without which this thesis would not have been possible. This network ranges from people intimately familiar with the content of this thesis— and the trials and tribulations associated with its production— to friends who were blissfully unaware and otherwise enriched my life and indirectly contributed to my success.

I am especially grateful to my advisor, Marie Banich, for encouraging my intellectual curiosity and allowing me to get into trouble, while simultaneously reminding me of the tangible impacts that research can have actual individuals. I appreciate that she understood not just as a graduate student, but as a (sometimes) stubborn human being, and extended me patience that I was able to repay. I am also very grateful to Tal Yarkoni, whose influence steered the course of my academic career towards neuroinformatics. He encouraged a balance between intellectual curiosity and critical skepticism that scientifically inspired me, and trusted my abilities and potential at critical times. I am also thankful to Andrew Reineberg, and other members of my lab, for the frequent— sometimes hourly— conversations that shaped this work.

Finally, but certainly not least, I am grateful to my family for their encouragement and enduring love and support. Nadia, you've been there for the best and worst moments and believed in me always. Adriana, you never allowed me to forget why my work was important and have been encouraging me the literal day one. Francisco, thank you for showing me what an exemplary scientist is.

CONTENTS

CHAPTER

I.	Introduction	1
II.	Large-scale meta-analysis of human medial frontal cortex reveals tripartite functional organization	9
III.	Large-scale meta-analysis of human lateral frontal cortex	37
IV.	Cross-modal evaluation of whole-brain atlases	84

BIBLIOGRAPHY.....	112
-------------------	-----

APPENDIX

I.	Full associations between LFC regions and topics	129
----	--	-----

CHAPTER I

Introduction

A fundamental goal of cognitive neuroscience is to precisely map the computational processes performed by anatomically discrete regions in the human brain. Although a precise ‘cartography’ is not sufficient for understanding brain function, it allows researchers to formulate mechanistic theories of information processing across the brain (Friston, 2002). Functional cartography, or human brain mapping, began by making associations between behavioral changes in response to focal brain lesions and has enjoyed great success in revealing specific patterns of functional specialization throughout the brain. Further progress was also made by systematically mapping the electrical response of neurons in animals using invasive electrophysiological methods. For example, in a hallmark study, Hubel and Wiesel mapped the structural and functional architecture of cat primary visual cortex, discovering orientation selective neurons in V1 (Hubel & Wiesel, 1962).

The advent of functional magnetic resonance imaging (fMRI) (Kwong et al., 1992) enabled an explosion of human brain mapping by allowing researchers to measure the whole brain’s response to a relatively unconstrained range of psychological phenomena. In the decades since, tens of thousands of studies have correlated individual activation foci to carefully controlled psychological states to understand the localization of specific psychological states. Moreover, structural and functional connectivity imaging methods have precisely characterized the anatomical and functional connectivity between brain regions, revealing complex whole-brain networks underlying human behavior.

The importance of large-scale approaches in cognitive neuroscience

Despite the enormous amount of fMRI studies that have been conducted, there are several roadblocks that prevent a comprehensive understanding of functional-anatomical mappings. A significant limitation inherent to fMRI is a relatively low signal to noise ratio (SNR). As a consequence, the majority of published fMRI studies are vastly underpowered and report a large number of false-positives and inflated effect sizes (Button et al., 2013; Wager, Lindquist, & Kaplan, 2007; Yarkoni, 2009). In fact, the average fMRI study has only around 20% power to detect a medium sized effect (Yarkoni, 2009). Moreover, traditional fMRI analysis techniques conduct what are called ‘mass univariate’ analyses at the smallest unit in imaging: the voxel. After correcting for multiple comparisons, the spatial maps resulting from these techniques comprise a small subset of the true underlying brain signal correlated with the psychological state of interest.

A critical analysis technique that helps overcome some of these shortcomings is quantitative meta-analysis (Wager et al., 2007). In meta-analysis, individual peak coordinates are extracted from multiple studies that purportedly engaged participants in similar psychological states. These peaks quantitatively combined to determine the regions significantly associated with the psychological phenomena of interest. For example, Shackman et al., 2011 used a meta-analysis to map the anatomical overlap between pain, negative affect, and cognitive control. Shackman found evidence that these three processes engaged an overlapping section of anterior midcingulate cortex (aMCC) and in conjunction with anatomical and electrophysiological evidence argued that pain and negative affect signal the need to adaptively change motor plans to avoid future negative outcomes.




Figure 1.1 Negative affect, pain and cognitive control activate a common region within the aMcc. The map depicts the results of a meta-analysis of 380 activation foci derived from 192 experiments and involving more than 3,000 participants. The uppermost panel shows the spatially normalized foci for each domain. The next panel shows thresholded activation likelihood estimate (ALE) maps for each domain considered in isolation. The two lowest panels depict the region of overlap across the three domains. Reproduced from Shackman et al. 2011

Although meta-analyses allow for fine-grained testing of functional-anatomical hypotheses, without surveying a wide range of unrelated psychological states, it is possible to fall prey to what has been dubbed the ‘reverse inference’ problem (Poldrack, 2006). Traditional fMRI studies are designed to infer the probability of brain activity given the psychological states induced in the study— or what is known as ‘forward inference’ [$P(\text{activity} \mid \text{state})$]. In contrast, true ‘reverse inference’ requires determining which psychological states are probable given a pattern of brain activity [$P(\text{state} \mid \text{activity})$]. However, to conduct a proper reverse inference, it is necessary to survey a wide range of unrelated psychological states to determine the *specificity* between activity in a given brain region and a

psychological state. This is particularly problematic as the base rate of activation varies widely across the brain (Figure 1.2). Certain regions, like anterior midcingulate cortex are active in such a large proportion of studies that it is very difficult to determine if specific psychological states (such as pain, negative affect and cognitive control) *preferentially* recruit this region.




Figure 1.2. Frequency of activation across the brain. For each voxel across the brain, I display the proportion of studies in which it's active. Regions critical for goal-directed cognition, such as medial and dorsolateral prefrontal cortex, exhibit high rates of activation, in turn making it difficult to determine which states preferentially recruit them. Lighter colors represent greater activation rates.

Fortunately, there has been a recent growth in the development of large-scale meta-analysis frameworks, such as Neurosynth (Yarkoni, Poldrack, Nichols, Van Essen, & Wager, 2011) and BrainMap (Laird, Lancaster, & Fox, 2005), which

allow researchers to more formally formulate ‘reverse inferences’. This is particularly true of the Neurosynth framework, as it was specifically designed to scale as the literature grows by automatically extracting activation coordinates and semantic meta-data of fMRI studies. As of 2016, Neurosynth includes over 11,000 fMRI studies, encompassing a diverse range psychological manipulations. By widely surveying across the psychological literature, large-scale meta-analysis allows researchers to quantitatively determine the specificity of brain-behavior relationships.

To demonstrate the importance of appropriately modeling reverse inference, we recreated Shackman’s (2011) meta-analysis using forward inference— akin to a traditional meta-analysis— using Neurosynth. Similar to Shackman (2011), we find overlap between pain, negative affect and cognitive control in aMCC (Figure 1.3a; overlap shown in white). However, we also find a very similar pattern of overlap when we perform a forward inference analysis of three theoretically unrelated psychological states: ‘social cognition’, ‘vision’ and ‘long term memory’ (Figure 1.3b). This demonstration highlights the importance of large-scale meta-analyses that appropriately quantify preferential psychological recruitment across the brain.




Figure 1.3. a) Forward inference meta-analysis of pain, negative affect and cognitive control, showing distinct overlap in anterior midcingulate cortex (aMCC). b) Forward inference meta-analysis of three theoretically unrelated constructs

(social cognition, vision and long term memory) shows similarly striking overlap in aMCC. Overlap is indicated in white.

Finding the right brain units

A related problem facing cognitive neuroscience is determining how the brain's complex anatomy is spatially organized into units that give rise to psychological function. It is of great interest to define functionally dissociable units as this can facilitate the formulation of theories linking brain to behavior (Poldrack & Yarkoni, 2016). Much progress has been made on this front by using data-mining techniques on a variety of brain data from various modalities. For example, a particularly popular strategy has been to apply unsupervised learning algorithms to connectivity data that describes either the anatomical connections (Beckmann, Johansen-Berg, & Rushworth, 2009; Johansen-Berg et al., 2004) or temporal correlations in fMRI signal (i.e. functional connectivity; Craddock, James, Holtzheimer, Hu, & Mayberg, 2012; Shen, Tokoglu, Papademetris, & Constable, 2013; Yeo et al., 2011). These methods have greatly informed our understanding of the organization of the brain, revealing large-scale brain networks that were not previously widely appreciated (Figure 1.4).




Figure 1.4. Seven-whole brain networks estimated from intrinsic functional connectivity in resting state fMRI. Reproduced from Yeo et al., (2011).

However, a shortcoming of these brain-centric methods is that they are generally void of functional data linking brain units to distinct psychological states. As such, these methods cannot directly speak to the psychological function of these brain regions. Moreover, it is not clear if the units derived from these various brain measures are necessarily the units that best explain the *functional differences* observed during behavioral performance in task-related fMRI.

Dissertation overview

In the present dissertation, I seek to advance large-scale meta-analytic techniques by making a link between psychological function and anatomical brain units in three investigations. In the first two investigation, I use unsupervised data-driven techniques to identify spatially distinct regions in medial and lateral frontal cortex (Chapter 2 and 3, respectively) on the basis of co-activation patterns across a wide variety of fMRI studies. I then use classification techniques to decode the psychological states that best predict activity for each region, revealing theoretically informative brain-behavior mappings. I chose to study frontal cortex as the topology of psychological states is less well understood in higher-level association cortex. Moreover, as a consequence of frontal cortex being centrally involved in a wide variety of behaviors, the base rate of activation in certain frontal regions is very high and particularly vulnerable to the ‘reverse inference’ problem. These two studies provide comprehensive and relatively unbiased functional-anatomical mappings of human frontal cortex.

In the final chapter, I evaluate the quality of various strategies for meta-analytic parcellation and compare the utility of these meta-analytic brain atlases to those derived from other brain modalities. In an effort to objectively choose ‘the

right brain units', I use classification to evaluate how well brain atlases from different modalities are able to predict psychological states. In this study, I find evidence that meta-analytically derived atlases may provide a more accurate and useful representation of the underlying functional-anatomical organization of the human brain than those derived from various other modalities. I conclude this dissertation with brief concluding remarks summarizing the contributions of the present studies.

CHAPTER 2

Large-scale meta-analysis of human medial frontal cortex reveals tripartite functional organization

The medial frontal cortex (MFC) is purported to play a key role in a number of psychological processes, including motor function, cognitive control, emotion, pain and social cognition. However, the precise correspondence of psychological states onto discrete medial frontal anatomy remains elusive. Several recent attempts to define distinct functional sub-regions of MFC have been based on morphology (Palomero-Gallagher, Zilles, Schleicher, & Vogt, 2013; Vogt, 2016) in-vivo structural connectivity (Beckmann et al., 2009; Johansen-Berg et al., 2004; Neubert, Mars, Thomas, Sallet, & Rushworth, 2014; Sallet et al., 2013) and functional connectivity (Andrews Hanna, Reidler, Sepulcre, Poulin, & Buckner, 2010). Although such studies map key properties which constrain information processing in MFC, it's unclear if these boundaries correspond to patterns of brain activity observed during behavioral performance (Amunts & Zilles, 2015; Eickhoff et al., 2007; Mattar, Cole, Thompson-Schill, & Bassett, 2015). Moreover, as these methods do not measure the brain's response to various psychological challenges, they cannot directly identify the (potentially separable) functional associates of MFC sub-regions.

To this end, task-based functional MRI (fMRI) has suggested that distinct foci of MFC activation may be associated with specific psychological manipulations. For

example, the supplementary motor area (SMA) and pre-SMA have been associated with the planning and initiation of movements (Kennerley, Sakai, & Rushworth, 2004; Leek & Johnston, 2009; Roland, Larsen, Lassen, & Skinhøj, 1980), while midcingulate cortex (MCC) has been implicated in various aspects of cognitive control (Botvinick, Nystrom, Fissell, Carter, & Cohen, 1999; Brown & Braver, 2005; Holroyd et al., 2004; Milham et al., 2001; Shenhav, Botvinick, & Cohen, 2013), fear (Etkin, Egner, & Kalisch, 2011; Milad et al., 2007; B. A. Vogt & Vogt, 2003), and pain processing (Rolls et al., 2003; B. A. Vogt, 2016; Wager et al., 2013). Further anterior, medial prefrontal cortex (mPFC) and the rostral anterior cingulate cortex (rACC) have been associated with affective processes, including emotion (Etkin et al., 2011; K. A. Lindquist, Wager, Kober, Bliss-Moreau, & Barrett, 2012), autonomic function (Critchley et al., 2003), and valuation (Hare, Camerer, & Rangel, 2009), as well as internally oriented processes, such as mentalizing (Baumgartner, Götze, & Fehr, 2012) and autobiographical memory (Spreng & Grady, 2010).

Despite the large number of neuroimaging studies, there have been few large-scale efforts to comprehensively map the full range of psychological functions onto medial frontal anatomy. Most meta-analyses are restricted to a subset of empirical findings relevant to candidate cognitive states hypothesized to be important (e.g. negative affect, pain, cognitive control; Shackman et al., 2011) or a specific anatomical region of interest (e.g., Palomero-Gallagher et al., 2015). This relatively narrow scope limits the ability to address the specificity of activation of psychological states across the MFC more broadly. That is, without considering a

wide representative range of psychological states, it is difficult to determine whether particular psychological processes preferentially recruit specific subdivisions of MFC. This limitation, widely known as the reverse inference problem (Poldrack, 2006), is particularly acute for portions of MFC which commonly activate in a large proportion of fMRI studies, raising questions about whether these regions are selectively involved in specific mental functions (Nelson, Dosenbach, Cohen, Wheeler, Schlaggar, & Petersen, 2010a; Yarkoni et al., 2011).

Here we address these issues by creating a comprehensive mapping between psychological states and MFC anatomy using Neurosynth (Yarkoni et al., 2011), a framework for large-scale fMRI meta-analysis composed of nearly 10,000 studies. We first clustered MFC voxels into functionally separable regions at several spatial scales based on their co-activation across studies with the rest of the brain (Kober et al., 2008; Robinson, Laird, Glahn, Lovallo, & Fox, 2010; S. M. Smith et al., 2009; Toro, Fox, & Paus, 2008). In contrast to cytoarchitectonic and connectivity based parcellations, the present analysis identified clusters with distinct signatures of activation across a wide range of psychological manipulations. This procedure revealed three zones along the rostro-caudal axis that further fractionated into nine sub-regions. We then characterized each cluster's functional profiles using multivariate classification, revealing broad functional shifts between the three zones, and subtler variations between their corresponding sub-regions. Collectively, our results provide a comprehensive functional map of the human MFC using relatively unbiased data-driven methods.

Materials & Methods




Figure 2.1. Methods overview. A) Whole brain co-activation of MFC voxels was calculated and k-means clustering was applied resulting in spatially distinct clusters. B) For each cluster, thresholded whole-brain co-activation maps were generated. C) We generated functional preference profiles for each cluster by determining which psychological topics best predicted their activation.

Database. We analyzed version 0.4 of the Neurosynth database, (Yarkoni et al., 2011), a repository of 9,721 fMRI studies and over 350,000 activation peaks that span the full range of the published literature. The studies included human subjects of either sex. Each observation contains the peak activations for all contrasts reported in a study's table as well as the frequency of all of the words in the article abstract. A heuristic but relatively accurate approach is used to detect and convert reported coordinates to the standard MNI space (see: Yarkoni et al., 2011). As such, all activations and subsequent analyses are in MNI152 coordinate

space. The scikit-learn Python package (Pedregosa et al., 2011) was used for all machine learning analyses. Analyses were performed using the core Neurosynth python tools (<https://github.com/neurosynth/neurosynth>); code and data to replicate these analyses on any given brain region at any desired spatial granularity are available

as a set of IPython Notebooks (<https://github.com/adelavega/neurosynth-mfc>).

Co-activation-based clustering. We clustered individual voxels inside of a MFC mask based on their meta-analytic co-activation with voxels in the rest of the brain (Figure 2.1A). First, we defined a MFC mask excluding voxels further than 10mm from the midline of the brain, posterior to the central sulcus ($Y < -22\text{mm}$) and ventral to vmPFC ($Z < -32\text{mm}$). Next, we removed voxels with low grey matter signal by excluding voxels with either fewer than 30% probability of grey matter cortex according to the Harvard-Oxford anatomical atlas, or very low activation rates in the database (less than 80 studies per voxel). In general, Neurosynth’s activation mask (derived from the standard MNI152 template distributed with FSL) corresponded highly with probabilistic locations of cerebral cortex, with the exception of portions of precentral gyrus and far ventromedial prefrontal cortex— which showed low activation although they were more than 50% likely to be in cerebral cortex.

Next, we calculated the co-activation of each MFC voxel with the rest of the brain by correlating the target voxel’s activation pattern across studies with the rest of the brain. Activation in each voxel is represented as a binary vector of length

9,721 (the number of studies). A value of 1 indicated that the voxel fell within 10 mm of an activation focus reported in a particular study, and a value of 0 indicated that it did not. Because correlating the activation of every MFC voxel with every other voxel in the brain would result in a very large matrix (15,259 MFC voxels x 228,453 whole-brain voxels) that would be computationally costly to cluster, we reduced the dimensionality of the whole brain to 100 components using principal components analysis (PCA; the precise choice of number of components does not materially affect the reported results). Next, we computed the Pearson correlation distance between every voxel in the MFC mask with each whole-brain PCA component. We applied k-means clustering to this matrix (15,259 MFC voxels x 100 whole-brain PCA components) to group the MFC voxels into 2-15 clusters. K-means was used for clustering as this algorithm is computationally efficient, widely used, and shows reasonably high goodness-of-fit characteristics (Thirion, Varoquaux, Dohmatob, & Poline, 2014). We used the k-means++ initialization procedure, ran the algorithm 10 times on different centroid seeds and selected the output of these consecutive runs with the lowest inertia to avoid local minima.

Since the optimality of a given clustering depends in large part on investigators' goals, the preferred level of analysis, and the nature and dimensionality of the available data, identifying the 'correct' number of clusters is arguably an intractable problem (Eickhoff, Thirion, Varoquaux, & Bzdok, 2015; Poldrack & Yarkoni, 2016; Varoquaux & Thirion, 2014). However, in the interest of pragmatism, we attempted to objectively select the number of clusters using the

silhouette score, a measure of within-cluster cohesion. The silhouette coefficient was defined as $(b - a) / \max(a, b)$, where a is the mean intra-cluster distance and b is the distance between a sample and the nearest cluster of which the sample is not a part. Solutions that minimized the average distance between voxels within each cluster received a greater score. To estimate the uncertainty around silhouette scores, we used a permutation procedure previously employed by our group (Wager, Davidson, Hughes, Lindquist, & Ochsner, 2008).

To understand the anatomical correspondence of the resulting clusters, we calculated the probability of voxels in each cluster of occurring in probabilistic regions from the Harvard-Oxford atlas (H-O). We refer to H-O's Juxtapositional Lobule Cortex as Supplementary Motor Area (SMA) for consistency. We also compared the location of clusters to regions from cytoarchitectonic atlases of medial motor areas (Picard & Strick, 1996), mid-cingulate cortex (B. A. Vogt, 2016) and vmPFC (S. Mackey & Petrides, 2014). To be precise, sub-regions in the nine-cluster solution were given alphanumeric labels in addition to descriptive names.

Co-activation profiles. Next, we analyzed the differences in whole brain co-activation between the resulting clusters (Figure 2.1B). To highlight differences between clusters, we contrasted related sets of clusters. For the three-cluster solution, we contrasted the co-activation of each cluster (e.g. 'posterior zone') with the other two clusters (e.g. 'middle' and 'anterior' zones). For the nine-cluster solution, we contrasted the co-activation of each cluster (e.g. 'SMA') with spatially adjacent clusters that fell within the same zone of the three-cluster solution (e.g.

‘pre-SMA’). To do so, we performed a meta-analytic contrast between studies that activated a given cluster and studies that activated control clusters. The resulting images identify voxels with a greater probability of co-activating with the cluster of interest than with control clusters. For example, voxels in grey in the first panel of Figure 2.3B indicate voxels that are active more frequently in studies in which SMA [P1] is active than in studies in which pre-SMA [P2] is active. We calculated p-values for each voxel using a two-way chi-square test between the two sets of studies and thresholded the co-activation images using the False Discovery Rate ($q < 0.01$). The resulting images were binarized for display purposes and visualized using the NiLearn library for Python.

Topic modeling. Although term-based meta-analysis maps in Neurosynth closely resemble the results of manual meta-analyses of the same concepts, there is a high degree of redundancy between terms (e.g. ‘episodes’ and ‘episodic’), as well as potential ambiguity as to the meaning of an individual word out of context (e.g. ‘memory’ can indicate working memory or episodic memory). To remedy this problem, we employed a reduced semantic representation of the latent conceptual structure underlying the neuroimaging literature: a set of 60 topics derived using latent dirichlet allocation (LDA) topic-modeling (Blei, Ng, & Jordan, 2003). This procedure was identical to that used in a previous study (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012b), except for the use of a smaller number of topics and a much larger version of the Neurosynth database. The generative topic model derives 60 independent topics from the co-occurrence across

studies of all words in the abstracts fMRI studies in the database. Each topic loads onto individual words to a varying extent, facilitating the interpretation of topics; for example, a working memory topic loads highest on the words 'memory, WM, load', while an episodic memory topic loads on 'memory, retrieval, events'. Note that both topics highly load on the word “memory”, but the meaning of this word is disambiguated because it is contextualized by other words that strongly load onto that topic. Out of the 60 topics, we excluded 25 topics representing non-psychological phenomena— such as the nature of the subject population (e.g. gender, special populations) and methods (e.g., words such as “images”, “voxels”)— resulting in 35 psychological topics. See Table 1 for a list of topics most associated with MFC.

Topic definitions

Topic name	Highest loading words
gaze	eye gaze movements eyes visual saccades saccade target
decision-making	decision choice risk decisions choices uncertainty outcomes risky
episodic	memory events imagery autobiographical retrieval episodic memories future
motor	motor movement movements sensorimotor primary finger control imagery
social	social empathy moral person judgments mentalizing mental theory
reward	reward anticipation monetary responses rewards motivation motivational loss
switching	cues target trials cue switching stimulus targets preparation
conflict	conflict interference control incongruent trials stroop congruent cognitive
inhibition	inhibition control inhibitory stop motor trials nogo cognitive
fear	fear anxiety threat responses conditioning cs extinction autonomic
WM	memory performance cognitive wm tasks verbal load executive
pain	pain painful stimulation somatosensory intensity noxious heat nociceptive

Table 2.1. Topics most strongly associated with MFC regions used in Figure 2.4. Ten strongest loading words for each topic are listed, in descending order of association strength.

Meta-analytic functional preference profiles. We generated functional preference profiles by determining which psychological topics best predicted each MFC cluster’s activity across fMRI studies (Figure 2.1C). First, we selected two sets

of studies: studies that activated a given cluster—defined as activating at least 5% of voxels in the cluster— and studies that did not— defined as activating no voxels in the cluster. For each cluster, we trained a naive Bayes classifier to discriminate these two sets of studies based on psychological topics. We chose naive Bayes because (i) we have previously had success applying this algorithm to Neurosynth data (Yarkoni et al., 2011); (ii) these algorithms perform well on many types of data (Androutsopoulos, Koutsias, & Chandrinos, 2000); (iii) they require almost no tuning of parameters to achieve a high level of performance; and (iv) they produce highly interpretable solutions, in contrast to many other machine learning approaches (e.g., support vector machines or decision tree forests).

We trained models to predict whether or not fMRI studies activated each cluster, given the semantic content of the studies. In other words, if we know which psychological topics are mentioned in a study how well can we predict whether the study activates a specific region? We used 4-fold cross-validation for testing and calculated the mean score across all folds as the final measure of performance. We scored our models using the area under the curve of the receiver operating characteristic (AUC-ROC) –a summary metric of classification performance that takes into account both sensitivity and specificity. AUC-ROC was chosen because this measure is not detrimentally affected by unbalanced data (Jeni, Cohn, & la Torre, 2013), which was important because each region varied in the ratio of studies that activated it to the studies that did not.

To generate functional preference profiles, we extracted from the naive Bayes models the log odds-ratio (LOR) of a topic being present in active studies versus inactive studies. The LOR was defined as the log of the ratio between the probability of a given topic in active studies and the probability of the topic in inactive studies, for each region individually. LOR values above 0 indicate that a psychological topic is predictive of activation of a given region. To determine the statistical significance of these associations, we permuted the class labels and extracted the LOR for each topic 1000 times. This resulted in a null distribution of LOR for each topic and each cluster. Using this null distribution, we calculated p-values for each pairwise relationship between psychological concepts and regions, and reported associations significant at the $p < 0.001$ threshold. Finally, to determine if certain topics showed greater preference for one cluster versus another, we conducted exploratory, post-hoc comparisons by determining if the 95% confidence intervals (CI) of the LOR of a specific topic for a one region overlapped with the 95% CI of the same topic for another region. We generated CIs using bootstrapping, sampling with replacement and recalculating log-odds ratios for each region 1000 times.




Figure 2.2. Co-activation-based clustering of MFC results. A) Mid-sagittal view at three levels at granularity: three broad zones, nine and twelve sub-regions. Clusters in nine sub-region solution are given both descriptive and alphanumeric names for reference. SMA: supplementary motor area. pre-SMA: pre-supplementary motor area; MCC: midcingulate cortex. pgACC: pre-genual anterior cingulate cortex; dmPFC: dorsal medial PFC; vmPFC: ventromedial PFC. B) Axial view of nine sub-regions. C) Silhouette scores of real (green) and permuted (blue) clustering solutions. Clustering was performed on permuted data 1000 times for each k to compute a null distribution (p -values for all clusters $< .001$). Silhouette scores reached local maxima at 3 regions and plateaued after 9.

Results

Functionally separable regions of medial frontal cortex. We

identified spatially dissociable regions on the basis of shared co-activation profiles with the rest of the brain (Chang, Yarkoni, Khaw, & Sanfey, 2013; Kober et al., 2008; S. M. Smith et al., 2009; Toro et al., 2008), an approach that exploits the likelihood of a voxel co-activating with another voxel across studies in the meta-analytic database (Figure 2.2). Because structure-to-function mappings can be identified at multiple spatial scales, we iteratively extracted 2- through 15-cluster

solutions and assessed their validity using the silhouette score—a commonly used measure of inter-cluster coherence. Permutation analyses indicated that the null hypothesis of random clustering could be rejected for all solutions, with silhouette scores reaching local maxima at 3 clusters (Figure 2.2C). The plateauing of silhouette scores suggests that there is little objective basis for selecting one solution over another past around 9 clusters (Thirion et al., 2014). We have therefore opted to focus on the 3-cluster and 9- cluster solutions because they provide greater theoretical parsimony than more fine-grained solutions.

At the coarsest level, MFC divided into three broad bilateral clusters organized along the rostral-caudal axis. The nine-cluster solution revealed additional fine-grained topographical organization, with each of the three major zones fractionating into 2-4 smaller regions (84% of all voxels within each zone overlapped with its putative sub-regions). We henceforth refer to the clusters from the 3-cluster solution as “zones” to differentiate them from clusters in the 9-cluster solution, which we refer to as “sub-regions”.

To better understand the anatomical location of our clusters, we compared them to previously defined sub-regions from the Harvard-Oxford (H-O) probabilistic structural atlas and well-known cytoarchitectonic studies. Although we did not necessarily expect our clusters to conform precisely to morphologically derived regions, we nonetheless observed moderate correspondence— suggesting morphological properties constrain, but not determine function. Within the posterior zone, we identified two clusters (Figure 2.2A; SMA [P1] & pre-SMA[P2])

with a high probability of occurring in SMA according to H-O. The two clusters were approximately delineated by the vertical commissure anterior (VCA), consistent with cytoarchitectonic delineations (Picard and Strick, 1996). However, SMA [P1] spanned multiple cytoarchitectonic areas— extending ventrally to include portions of Picard & Strick’s cingulate zones— suggesting these morphologically distinct areas co-activate similarly across tasks.


In the middle zone, we identified four clusters consistent with midcingulate cortex (MCC). In particular, two anterior and two posterior clusters delineated from each other a few millimeters anterior to the VCA, consistent with Vogt’s definition of anterior and posterior midcingulate cortex (B. A. Vogt, 2016). The two dorsal clusters (pdMCC [M1] & adMCC [M2]) showed a high probability of falling within H-O’s paracingulate gyrus, whereas the two ventral clusters (pvMCC [M3] & avMCC [M4]) showed a high probability of falling in the cingulate gyrus proper. Unlike some cytoarchitectonic studies, we did not identify any regions exclusively located in the cingulate sulcus, such as the rostral cingulate zone.

In the anterior zone, the most dorsal cluster (dmPFC [A1]) included medial aspects of H-O’s frontal pole and superior frontal gyrus, and was entirely outside of the anterior cingulate gyrus. Ventrally, we identified a second cluster (pgACC [A2]) which was primarily located within pregenual aspects of the anterior cingulate gyrus, but also included pregenual portions of paracingulate gyrus. Finally, the most ventral cluster (vmPFC [A3]) encompassed both pregenual aspects of the ACC

and medial OFC, similar to the vmPFC area of interest used in cytoarchitectonic studies (Mackey and Petrides, 2014).

Next, to provide direct insight into the functions of the clusters we identified, we applied two approaches. First, we determined which other brain regions co-activate with each cluster, in order to reveal their functional networks. Second, we used semantic data from Neurosynth to determine which psychological states predict the activation of each cluster.

A) Functional zones



B) Sub-regions




Figure 2.3. Meta-analytic co-activation contrasts for (A) three zones and B) nine sub-regions. Colored voxels indicate significantly greater co-activation with the seed region of the same color (at right) than control regions in the same row. The three zones showed distinct co-activation patterns, while sub-regions within each zone showed fine-grained co-activation differences. Images are presented using

neurological convention and were whole-brain corrected using a false discovery rate (FDR) of $q = 0.01$. Major subcortical structures are labeled; Thal: thalamus, Hipp: hippocampus; Amyg: amygdala; DS: dorsal striatum; VS: ventral striatum.

Meta-analytic co-activation profiles. We directly contrasted co-activation patterns of the three functional zones— i.e., we sought to identify voxels that co-activated to a stronger degree with each zone than with the other two (Figure 2.3A). The posterior zone showed greater bilateral co-activation with primary motor cortex (PMC) and superior parietal cortex (SPC), anterior cerebellum, and posterior insula (pIns) as well subcortical regions such as the thalamus and dorsal striatum (DS)—a co-activation pattern consistent with motoric function. The middle zone co-activated with anterior aspects of the thalamus as well as regions in the frontoparietal control network such as dorsolateral prefrontal cortex (DLPFC), anterior insula (aIns) and SPC. Finally, the anterior zone showed a qualitatively different pattern, co-activating to a greater extent with default network regions such as angular gyrus, hippocampus and posterior cingulate cortex (PCC) (Andrews-Hanna, 2012). The anterior zone also showed greater co-activation with subcortical regions important for affect— the amygdala and ventral striatum (VS).

To understand the differences in co-activation found within each zone, we directly contrasted the co-activation patterns of each zone's sub-regions (Figure 2.3B). In the posterior zone, SMA [P1] showed greater co-activation with somatosensory cortices and pIns while pre-SMA [P2] showed greater co-activation with posterior DLPFC, including the inferior frontal junction (IFJ), as well as

aIns— regions associated with goal-directed cognition (Chang et al., 2013; Nelson, Dosenbach, Cohen, Wheeler, Schlaggar, & Petersen, 2010b). Within the middle zone, we found that all four sub-regions strongly co-activated with various aspects of the insula. However, pvMCC [M3] was more strongly co-activated with pIns, SII and the brain stem—important regions for pain processing (B. A. Vogt, 2005; Wager et al., 2013). In contrast, avMCC [M4] co-activated more strongly with ventral aIns and lateral OFC—regions previously associated with reward-driven learning (Stalnaker, Cooch, & Schoenbaum, 2015). In contrast, both dorsal MCC [M1 & M2] clusters were more strongly associated with dorsal aIns and frontoparietal control regions (e.g DLPFC, SPC). However, adMCC [M2]’s co-activation extended anteriorly into the frontal pole, whereas pdMCC [M1] more strongly co-activated with motor cortices. Subcortically, pvMCC [M3] showed greater co-activation with the thalamus and dorsal striatum while avMCC showed greater co-activation with the left amygdala. However, daMCC [M2] also showed robust co-activation with portions of thalamus and dorsal striatum.

Within the anterior zone, pgACC [A2] did not show many co-activation differences from its neighbors. Surprisingly, both dmPFC [A1] and vmPFC [V3] showed greater co-activation with PCC – a key default network region. In addition, dmPFC [A1] robustly co-activated with portions of the so-called ‘mentalizing’ network, such as the tempo-parietal junction (TPJ) (R. M. Carter & Huettel, 2013) and the superior temporal sulcus (STS) (Zilbovicius et al., 2006), as well as lateral PFC, including inferior and middle frontal gyri. Finally, vmPFC [A3] showed strong

co-activation with subcortical regions, including VS and the amygdala, extending into the hippocampus. As a whole, these co-activation patterns demonstrate that the regions we identified are involved with distinct functional networks, and suggest that there are likely broad functional differences across MFC zones, accompanied by fine-grained differences within each sub-region.




Figure 2.4. Functional preference profiles of MFC clusters. Each cluster was profiled to determine which psychological concepts best predicted its activation. Top) Each of the three functional zones we identified showed distinct functional profiles with broad shifts across cognitive domains Bottom) Within each zone, sub-regions showed fine-grained shifts in function. Strength of association is measured in log odds-ratio (LOR), and permutation-based significance ($p < 0.001$) is indicated next to each psychological concept by color-coded dots corresponding to each region. Meta-analytic functional preference profiles

Next, we used a data-driven approach that surveyed a broad range of psychological states to determine if MFC clusters are differentially recruited by

psychological states. For each cluster, we trained a multivariate classifier to predict which studies activated the cluster using a set of 35 psychological topics derived by applying a standard topic modeling approach to the abstracts of articles in the database (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012a) (Table 1). From the resulting fitted classifiers, we calculated a measure of how strongly each topic indicated that a study activated a given cluster (measured as the log odds-ratio [LOR] of the probability of a each topic in studies that activated a given cluster to the probability of the same topic in studies that did not activate the cluster). LOR values over 0 indicate that the presence of that topic in a study predicts activity in a given region. We restricted interpretation to significant associations ($p < 0.001$) and additionally report 95% confidence intervals of LORs whenever we comparatively discuss sets of regions. As the latter comparisons are post-hoc and exploratory, caution in interpretation is warranted.

Although the following results demonstrate relatively high loadings between specific topics and regions (e.g. ‘motor’ and SMA), classification using all 35 topics yielded much better performance (mean AUC-ROC: 0.61) than when using only the most predictive topic of each region (mean AUC-ROC: 0.54). The relatively poor performance when using only one topic suggests low selectivity between psychological states and any one region.

Across the three broad MFC zones, we observed distinct functional patterns, consistent with their divergent patterns of functional co-activation (Figure 2.3). The posterior zone was primarily involved with motor function (including gaze),

consistent with its co-activation with motor regions. The middle zone was primarily associated with various facets of cognitive control, but was also implicated in negative affect—pain and fear— as well as decision-making. Consistent with its distinct pattern of co-activation, the anterior zone showed a robust shift away from goal-directed cognition and was strongly associated with affective processes such as reward, fear and decision-making, as well as internally oriented processes such as episodic memory and social processing.

Inspection at a finer spatial scale revealed that sub-regions within each zone showed more subtle patterns of psychological function, similar to the fine-grained variations in co-activation previously observed for each sub-region. In the posterior zone (Figure 2.4, bottom left), activity in both clusters was similarly predicted by motor function and switching. However, exploratory post-hoc tests suggested that SMA [P1] was more strongly associated with pain, while pre-SMA [P2] was more strongly associated with working memory (WM) (95% CI LOR. ‘pain’: SMA [0.6, 1.1], pre-SMA [-0.1, 0.4]; ‘WM’, SMA [-0.2, 0.1], pre-SMA [0.2, 0.4]).

In the middle zone (Figure 2.4, bottom middle), activity in all four sub-regions was significantly predicted by aspects of cognitive control (i.e. conflict and inhibition) and pain. However, post-hoc exploratory tests indicated dorsal MCC (M1 & M2) was more strongly associated with WM than ventral MCC (M3 & M4) (95% CI LOR. ‘pdMCC [0.5, 0.8], adMCC [0.4, 0.6], pvMCC [0, 0.15], avMCC [0, 0.3]) whereas ventral MCC showed a stronger association with affect (95% CI LOR. ‘fear’: pdMCC [-0.1, 0.4], adMCC [-0.4, 0.2], pvMCC [0.7, 1.2], avMCC [0.4, 0.9]; ‘reward:

pdMCC [-0.4, 0.1], adMCC [-0.4, 0.1], pvMCC [0.3, 0.7], avMCC [0.3, 0.8]; ‘pain’: pdMCC [0.3, 0.8], adMCC [0.2, 0.7], pvMCC [1.1, 1.5], avMCC [0.6, 1.1]). Finally, both anterior clusters showed a greater association with decision-making than their posterior counterparts (95% CI LOR. pdMCC [-0.2, 0.3], adMCC [0.3, 0.8], pvMCC [-0.2, 0.4], avMCC [0.6, 1.1])

In the anterior zone (Figure 2.4, bottom right), activity across all three sub-regions was significantly predicted by episodic memory and social processing; however, the association with social processing was maximal for dmPFC [A3] (95% CI LOR. dmPFC [1.3, 1.7], pgACC [0.7, 1], vmPFC [0.6, 1]). In contrast, the reverse was true for reward and decision-making; we observed a gradient such that the association with reward and fear was greatest going ventrally, reaching a maximum in vmPFC (95% CI LOR. ‘reward’: dmPFC [-0.4, 0.3], pgACC [0.5, 1], vmPFC [1.2, 1.7]; ‘fear’: dmPFC [-0.4, 0.3], pgACC [0.2, 0.7], vmPFC [0.8, 1.3]).

Discussion

In the current study, we identified and functionally characterized regions of the medial frontal cortex by applying a data-driven approach to a large-scale database of ~10,000 fMRI studies. We defined regions on the basis of differences in co-activation patterns across a wide variety of psychological manipulations— a more direct measure of function than morphology or connectivity. We identified three broad zones arranged along the rostral-caudal axis that further fractionated into 2-4 sub-regions. Finally, we used multivariate classification analyses to identify the psychological topics most strongly predictive of activity in each region, revealing

broad shifts in function between the three broad zones and more fine-grained differences between sub-regions within each zone. In the following sections, we discuss theoretical implications for each zone as well as future challenges.

Posterior zone. Posterior MFC spanned various regions previously associated with motoric function—such as SMA, pre-SMA, and motor cingulate zones. This zone further fractioned into a posterior and anterior cluster similarly to cytoarchitectonic (Vorobiev, Govoni, Rizzolatti, Matelli, & Luppino, 1998) and connectivity parcellations (Johansen-Berg et al., 2004; Kim et al., 2010). As a whole, posterior MFC was primarily associated with motor function and co-activated with key motor regions such as primary motor cortex and thalamus. However, SMA [P1] showed a greater association with pain processing and greater co-activation with key pain regions such as SII and thalamus, suggesting this region may be important for initiating movements in response to pain. In contrast, pre-SMA [P2] showed a stronger association with cognitive control and co-activated with regions important for goal-directed cognition (e.g. DLPFC, aIns). These results are generally consistent with a large line of work suggesting that pre-SMA is responsible for more complex motor actions that presumably require cognitive control (Picard & Strick, 1996).

Middle zone. The middle MFC zone spanned portions of the cingulate and paracingulate gyri consistent with existing definitions of midcingulate cortex (MCC) (Vogt, 2016). In contrast to claims of pain-selectivity in MCC (Lieberman & Eisenberger, 2015), all four middle sub-regions were associated with pain and

cognitive control. This finding is broadly consistent with adaptive control hypotheses, which postulates that MCC integrates negative affective signals with cognitive control in order to optimize actions in the face of action-outcome uncertainty (Cavanagh & Shackman, 2015; Shackman et al., 2011). However, the present results additionally suggest functional differences between sub-regions of MCC. Notably, both dorsal MCC clusters were more strongly associated with WM—and showed greater co-activation with other cognitive control regions— while ventral MCC was more strongly associated with affect and co-activated more strongly with subcortical regions, such as amygdala and striatum. Importantly, ventral MCC was associated not only with negative affect and pain, but also reward. Thus, the present results suggest that ventral aspects of MCC may incorporate low-level affective signals into cognitive control, whereas dorsal MCC may be more important for aspects of cognitive motor control that require working-memory or resolving interference. Finally, we also observed that both anterior MCC clusters were more strongly associated with decision-making than posterior clusters, consistent with theories that incorporate reward-driven decision-making processes into the optimization of cognitive control (Alexander & Brown, 2011; Brown & Braver, 2005).

Anterior zone. Anterior MFC exhibited a distinct functional profile with strong associations with affect, decision-making, social cognition, and episodic memory, accompanied by co-activation with the default network. Yet, our results suggest that anterior MFC zone is not a unitary area, and fractionated into

functionally differentiable subregions. DmPFC [A1] was most strongly associated with social processing, consistent with studies linking dmPFC to social perception and self-referential thought (Mitchell, Banaji, & Macrae, 2005) and consistent with its robust co-activation with TPJ— a region hypothesized to be important for mentalizing (Baumgartner et al., 2012; Denny, Kober, Wager, & Ochsner, 2012). pgACC [A2] showed a less specific functional pattern, showing moderate associations with both affective processes and decision-making, perhaps consistent with descriptions of a default network ‘hub’ region in mPFC (Andrews Hanna et al., 2010; van den Heuvel & Sporns, 2013). Finally, vmPFC [A3] was primarily associated with affective processes, such as reward and fear, consistent with its robust sub-cortical co-activation. Although some have characterized vmPFC as a ‘valuation’ system (Lebreton, Jorge, Michel, Thirion, & Pessiglione, 2009), our results suggest that vmPFC is equally important for other affective processes, such as fear. Thus, vmPFC may play a more general role of incorporating sub-cortical affective signals into cortex, while more dorsal regions contextualize this affective information (Roy, Shohamy, & Wager, 2012).

Future challenges. While the present results provide valuable insights into the functional neuroanatomy of MFC, a number of important challenges remain for future research. Although the present analyses revealed distinct functional profiles for each region in MFC, it is notable that no region was selectively activated by a single psychological concept. This functional diversity is evident in that at least two

distinct topics were significantly associated with each cluster and our classifier's poor ability to predict activation using only the single most strongly associated topic for each region. These results suggest a complex many-to-many mapping between brain regions and cognitive processes— in contrast to recent claims of functional selectivity in MFC (Lieberman and Eisenberger, 2015; c.f., Wager et al, in press). This heterogeneity is consistent with an enormous wealth of electrophysiological data demonstrating that virtually all areas of association cortex contain distinct, but overlapping, neuron populations with heterogeneous functional profiles (Kvitsiani et al., 2013; Shidara & Richmond, 2002; Sikes, Vogt, & Vogt, 2008).

Although the present results provide a comprehensive snapshot of MFC function, many have argued that brain regions dynamically assume different roles (Shackman, Fox, & Seminowicz, 2015) and modulate their connectivity as a function of task demands (Cole, Bassett, Power, Braver, & Petersen, 2014; Mattar et al., 2015). Moreover, MCC is likely to be among the most heterogeneous brain regions (Anderson, Kinnison, & Pessoa, 2013) as evidenced by its very high activation rate (Nelson, Dosenbach, Cohen, Wheeler, Schlaggar, & Petersen, 2010c; Yarkoni et al., 2011). Thus, because the functional co-activation profiles presented here represent averages across tasks, they may mask task-dependent co-activation structure. For example, it's possible that ventral MCC co-activates more strongly with the amygdala during 'fear', but co-activates with posterior insula during 'pain'. An interesting avenue of future research will be to precisely characterize how co-

activation and functional patterns of MFC change as a function of context through large-scale meta-analysis.

Moreover, although our parcellation was moderately consistent with boundaries based on cytoarchitecture and connectivity (e.g. the distinction between SMA and pre-SMA), we observed several discrepancies. For example, we did not identify separate cingulate motor zones (Picard & Strick, 1996), suggesting morphologically distinct regions can co-activate similarly to support high-level psychological function (e.g. ‘motor function’). Systematic modeling of the relationship between anatomy and task evoked activation– similarly to existing models linking resting state and anatomical connectivity (Goñi et al., 2014)– are needed to better understand the nature of such discrepancies.

The present report also provides the ability to generate hypotheses that can be more carefully tested in future studies using the candidate psychological functions discussed here. For example, our result suggests that ventral MCC had a higher association with affect than dorsal MCC. However, given the wide inter-subject variability in paracingulate anatomy (Paus et al., 1996) it would be prudent to explore this suggestion in a single sample with subject-level anatomical registration. This hypothesis might also be explored by large-scale meta-analyses that combine functional and anatomical data to more precisely localize activity to detailed anatomical variation. Moreover, the present findings can be improve the development of future multivariate classifiers by providing better prior information

as to the regions that may specifically predict psychological states (e.g. Wager et al., 2013).

Finally, there are several limitations of Neurosynth that can be addressed in future research. First, the topic model we employ is data-derived from the semantic content of papers. Although these topics provide a substantial improvement over term based meta-analysis (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012a), these topics are still based purely on the frequency with which terms appear in the abstracts of articles and are not able to capture more complex semantic structures. The adoption of a standardized ontology of psychological concepts and tasks, such as the cognitive atlas (Poldrack et al., 2011), will greatly improve the ability of future meta-analyses to discriminate more fine-grained theories. Second, the quality of activation data in Neurosynth is inherently limited due to its automatically generated nature. Although previous validation analyses have shown that these limitations are unlikely to contribute systematic biases (Yarkoni et al., 2011), coordinate based meta-analyses are generally limited in comparison to their image-based counterparts (Salimi-Khorshidi, Smith, Keltner, Wager, & Nichols, 2009). Sharing of statistical images in databases such as NeuroVault (Gorgolewski et al., 2015) will greatly improve the fidelity of future meta-analyses.

Conclusion. In the present study, we provide a comprehensive functional map of the human medial frontal cortex using unbiased data-driven methods. Although the anatomy of this area has been extensively studied, the present study

more directly identified putative sub-regions with distinct functional profiles across a wide-variety of psychological states. The present results can serve as a foundation to generate and test more fine-grained hypotheses in future studies.

CHAPTER 3

Large-scale meta-analysis of human lateral frontal cortex

Flexible behavior is the hallmark of human and nonhuman primates. Such flexible behavior enables the navigation of complex, rapidly changing environments, the pursuit of goals in the face of various obstacles, planning for hypothetical future events and the communication of complex ideas with others using language. Decades of research have implicated the lateral frontal cortex (LFC) as critical for high-level behavior distal from low-level sensory representations (Goldman & Rakic, 1996; Miller et al. 2011).

However, the mapping between fundamental psychological processes and discrete lateral frontal anatomy remains actively debated. In recent years, much progress has been made in understanding the LFC's organization by mapping various properties that constrain information processing. For example, more precise and detailed maps of the microstructural properties of LFC have suggested potentially dissociable regions in lateral frontal cortex (Figure 3.1A) (Petrides, 2005). In addition, detailed study of anatomical connectivity has revealed potentially dissociable regions in human LFC with distinct primate analogues (Neubert et al., 2014; Orr, Smolker, & Banich, 2015; Sallet et al., 2013). Although such methods have helped carefully characterize important properties of LFC, it is unclear to what extent the boundaries derived using these methods correspond to

the functional organization observed during behavioral performance (Amunts & Zilles, 2015; Eickhoff et al., 2007; Mattar et al., 2015).

An alternative approach that has helped map the functional correlates of distinct behavioral phenotypes is the quantitative meta-analysis of functional MRI (fMRI) studies. Such meta-analyses help overcome the low power in most individual fMRI studies and have resulted in more precise spatial maps of key processes based in LFC, such as working-memory (Nee et al., 2013; Wager & Smith, 2003), switching (Derrfuss, Brass, Neumann, & Cramon, 2005a), language (Binder, Desai, Graves, & Conant, 2009; Turkeltaub, Eden, Jones, & Zeffiro, 2002), mentalizing (Gilbert et al., 2006) and self-referential processing (Northoff et al., 2006). However, due to the effort required to compile meta-analyses and because most researchers are interested in a particular psychological domain, most meta-analyses are typically focused on a particular sub-region of LFC or a subset of domain-specific empirical findings.




Figure 3.1. A) Cyto-architectonic parcellation of human lateral frontal cortex based on Petrides and Pandya (1996). B) Control networks of the human brain derived using graph theory in rs-fc MRI based on Power et al., (2012). Frontoparietal network (yellow) and cingulo-opercular network (purple). C) Base rate of activation for voxels across the brain in Neurosynth. Note that key regions of the LFC, such and mid-LPFC, are active in a high percentage of studies.

The narrow scope of most existing meta-analyses necessarily limits the extent of their impact for two reasons. First, it is critical to interpret the association of a region to psychological states in the broader context of nearby anatomical neighbors and regions in the same network to be able to better differentiate subtle functional differences between regions. Second, due to a limitation known as the reverse inference problem (Poldrack, 2006), without contrasting studies that activate a region of interest to a diverse range of studies that do not, it is difficult to determine if a psychological function preferentially recruits a region, or if this association is due to domain-general functions underlying many psychological states. This limitation is particularly acute in regions of the brain active frequently across a broad range of tasks (Figure 3.1C). Hence, by the very nature of the LFC being involved in a broad range tasks due to its critical role in flexible behavior, subregions of this area are particularly difficult to associate with specific mental operations (Nelson, Dosenbach, Cohen, Wheeler, Schlaggar, & Petersen, 2010c; Yarkoni et al., 2011).

Here we address these issues by creating a comprehensive mapping between psychological states and LFC using Neurosynth (Yarkoni et al., 2011), a framework for large-scale fMRI meta-analysis composed of nearly 11,500 studies. First, we used a data-driven method that exploits the observation that functionally related regions frequently co-activate together across studies to defined functional subregions of LFC (Robinson et al., 2010; S. M. Smith et al., 2009; Toro et al., 2008). In recognition that brain-wide functional networks likely supersede LFC as

organizational units, we identified regions within LFC at various spatial scales, including a broad network-level solution and a finer-grained regionally-specific scale. In order to avoid restricting ourselves to an unnecessarily small subset of regions and to avoid defining biased or arbitrary boundaries, we clustered the whole cortex and selected regions that had a substantial number of voxels in LFC without using an explicit mask to mask voxels within LFC. We then characterized the functional profile of each resulting region using multivariate classification, explicitly contrasting studies that activated each region to those that did not, revealing dissociable psychological profiles that were relatively similarly within each network. Collectively, we provide a comprehensive and unbiased functional-anatomical mapping of LFC using the largest meta-analytic database presently available.

Materials & Methods




Figure 3.2. Methods overview. A) Co-activation across studies with the rest of

the brain was calculated for every cortical voxel and whole-brain clustering results were obtained using Ward hierarchical clustering. We chose three spatial scales to focus on using the silhouette method and selected clusters in LFC from the whole-brain clustering solutions. B) We contrasted the whole-brain co-activation of clusters, grouping clusters for comparison that grouped together at coarser spatial scales (i.e. clusters that were in the same network) C) We generated functional preference profiles for each cluster by determining which psychological topics best predicted their activation across studies in the database.

Dataset. We analyzed version 0.6 of the Neurosynth database (Yarkoni et al., 2011), a repository of 11,406 fMRI studies and over 410,000 activation peaks that span the full range of the published neuroimaging literature. Each observation contains the peak activations for all contrasts reported in a study's table as well as the frequency of all of the words in the article abstract. A heuristic but relatively accurate approach is used to detect and convert reported coordinates to the standard MNI space (see: Yarkoni et al., 2011). As such, all activations and subsequent analyses are in MNI152 coordinate space. The scikit-learn Python package (Pedregosa et al., 2011) was used for all machine learning analyses. Analyses were performed using the core Neurosynth python tools (<https://github.com/neurosynth/neurosynth>).

Lateral frontal cortex mask. To select clusters from whole-brain clustering solutions in lateral frontal cortex, we defined an LFC anatomical mask. Crucially, we only used this mask to select clusters that fell within this mask, and not to exclude individual voxels. First, we included voxels with a greater than 30% chance of falling in the frontal lobes according to the Montreal Neurological Institute structural probabilistic atlas (Collins, Holmes, Peters, & Evans, 1995; Mazziotta et al., 2001) and excluded medial voxels within 14mm of the midline. To focus on lateral frontal cortex, we excluded voxels that were exclusively located on the

orbital surface— ensuring to include lateral orbitofrontal voxels— by removing voxels in the superior and medial orbital gyri according to the AAL atlas and voxels with a greater than 30% probability of falling in ‘Frontal Operculum Cortex’ in the Harvard-Oxford atlas. Finally, we also excluded far ventral voxels of OFC ($Z < -14\text{mm}$) that were not excluded using anatomical atlases.

Co-activation clustering. Next, we clustered individual grey-matter cortical voxels across the whole brain based on their meta-analytic co-activation with the whole brain across studies in the database (Figure 3.2A). In order to avoid potentially biased or arbitrary cluster boundaries, we clustered the whole cortex and selected clusters for further analysis that fell within an anatomically defined LFC mask. Critically, we did not mask out voxels that were slightly outside of our mask— we either included or excluded entire clusters. This was particularly important for clusters near the edge of our LFC mask— as functional boundaries may not conform to anatomical boundaries— and at coarse clustering solutions— given the well-established finding that at least 4-5 whole-brain networks include voxels in lateral frontal cortex (Power et al., 2011; Yeo et al., 2011). For whole-cortex clustering, we excluded voxels with less than 30% probability of falling in grey matter according to the Harvard-Oxford anatomical atlas and those with very low activation in the database (less than 100 studies per voxel). In general, Neurosynth’s activation mask (derived from the standard MNI152 template distributed with FSL) corresponded highly with probabilistic locations of cerebral

cortex, with the exception of portions of dorsal precentral gyrus— which showed low activation although it was more than 50% likely to be in cerebral cortex.

We calculated the co-activation between each cortical voxel and every other voxel in the brain (including sub-cortex) by correlating the target voxel's activation pattern across studies with the rest of the brain. Activation in each voxel is represented as a binary vector of length 11,406 (the number of studies). A value of 1 indicated that the voxel fell within 10 mm of an activation focus reported in a particular study, and a value of 0 indicated that it did not. Because correlating the activation of every cortical voxel with every other voxel in the brain would result in a very large matrix (112,358 cortical voxels x 171,534 whole-brain voxels) that would be very computationally costly to cluster so as to identify distinct LFC regions. Hence, we reduced the dimensionality of the whole brain to 100 components using principal components analysis (PCA; the precise choice of number of components does not materially affect the reported results). Next, we computed the Pearson correlation distance between every voxel in the MFC mask with each whole-brain PCA component, resulting in a matrix that described the frequency with which each cortical voxel co-activated with the rest of the brain.

As an additional pre-processing step, we standardized each cortical voxel's co-activation with other brain voxels to ensure clustering would be driven by relative differences in whole brain co-activation and not the overall activation rate of each voxel. That is, if two voxels co-activated with similar voxels across the brain, we should consider them to be relatively similar even if one of those voxels activates

more frequently (and thus has slightly stronger correlations with all voxels). This adjustment was particularly important as preliminary analyses indicated that regions with very high rates of activation (e.g. pre-SMA/mid-cingulate cortex) more readily clustered into multiple clusters with few voxels, reflecting base rates in activation, although differences in their functional associations were minimal. Indeed, preliminary analyses confirmed that standardizing the co-activation matrix alleviated this concern. At $k = 70$, the mean activation rate of each cluster showed no correlation with voxel size when Z-scoring was used ($r=0.05$), as compared to when the raw co-activation matrix was used ($r = -0.65$) at $k = 70$ (Figure 3.3). Additionally, the range of cluster sizes was compressed, resulting in more evenly sized clusters. Cluster sizes ranged from 352 to 4546 voxels using the raw activation, compared to a range of 560 to 2862 voxels using standardized co-activation. See Chapter 4 for a more in depth investigation of the implications of this preprocessing strategy.




Figure 3.3. Relationship between base rate and cluster size. Standardizing the co-activation matrix prior to clustering reduced the relationship between

average base rate of a cluster and the size in voxels.

We applied hierarchical clustering with Ward’s linkage to the normalized co-activation matrix, resulting in a whole-brain linkage matrix. Ward’s clustering was selected as this algorithm is recommended as the best compromise between accuracy (e.g., fit to data) and reproducibility for clustering fMRI data (Thirion et al., 2014). However, this clustering algorithm is seldom used for whole-brain clustering because the computational time increases cubically [$\Theta(N^3)$] as a function of samples. We employed the fastcluster algorithm (Müllner, 2013)—a package of libraries that enable efficient hierarchical clustering [$\Theta(N^2)$ —to achieve whole-brain clustering.

Since the optimality of a given clustering depends in large part on investigators’ goals, the preferred level of analysis, and the nature and dimensionality of the available data, identifying the ‘correct’ number of clusters is arguably an intractable problem (Eickhoff et al., 2015; Poldrack & Yarkoni, 2016; Varoquaux & Thirion, 2014). However, in order to attempt to objectively guide the choice of choice of number of clusters to further analyze, we selected viable solutions using the silhouette score— a measure of within-cluster cohesion. Crucially, as we were specifically interested in the fit of the clustering to lateral frontal cortex, we only calculated the silhouette score with respect to voxels within our lateral frontal cortex mask. The silhouette coefficient was defined as $(b - a) / \max(a, b)$, where a is the mean intra-cluster distance and b is the distance between a sample and the nearest cluster of which the sample is not a part. Solutions that minimized the

average distance between voxels within each cluster received a greater score. Once having selected three k solution sizes to focus on ($k = 5, 33$ and 70 whole-brain clusters), we extracted only those clusters from these solutions that had a substantial percentage of voxels in our LFC mask. We varied the percentage of voxels within our LFC mask required to include a region across granularities with the objective maximizing coverage in LFC without including extraneous clusters with little presence in LFC. We arrived at 10% of voxels in a cluster within LFC at $k=5$, 25% of voxels at $k=33$ and 50% of voxels at $k=70$.

To understand the anatomical correspondence of the resulting clusters, we consulted a variety of anatomical and cytoarchitectonic atlases. To locate each cluster anatomically, we used the probabilistic Harvard-Oxford atlas (H-O) that is packaged with FSL. We also visually compared the location of our clusters to the Petrides' (2005) and Jülich micro-anatomical atlases included in FSL (Eickhoff et al., 2007; Mazziotta et al., 2001). Regions were assigned names in accordance to Brodmann areas (BA) whenever clusters were sufficient small to correspond to a single area (e.g. 'area 9/46v'). Clusters were given functional names when they spanned multiple cytoarchitectonic areas (e.g. IFJ) or multiple clusters spanned a single cytoarchitectonic area (e.g. PMd & PMv). Note that although names were assigned to ease the discussion of these regions, we do not make strong claims of correspondence between functionally and anatomically defined regions, as we observed several discrepancies throughout LFC.

Co-activation profiles. Next, we analyzed the differences in whole brain co-activation between the resulting clusters (Figure 3.2B) in order to understand the patterns of co-activation that differentiates these clusters. To highlight differences between clusters, we contrasted related sets of clusters. Related clusters were defined as clusters within each network, as defined by each section in our results section (e.g. ‘posterior LFC’ clusters). The organization of these clusters was informed by the hierarchical structure provided by the clustering dendrogram. For example, in Figure 3.10b, we contrast the co-activation of three clusters in the default network. To do so, we performed a meta-analytic contrast between studies that activated a given cluster, and studies that activated control clusters. The resulting images identify voxels with a greater probability of co-activating with the cluster of interest than with control clusters. For example, voxels in blue in Figure 3.10b indicate voxels that are active more frequently in studies in which ‘area 9’ is active than in studies in which other default network clusters are active. We calculated p-values for each voxel using a two-way chi-square test between the two sets of studies and thresholded the co-activation images using the False Discovery Rate ($q < 0.01$). In Figure 3.7b, the more liberal threshold of $q < 0.05$ was used as too few voxels were significantly different at $q < 0.01$. The resulting images were binarized for display purposes and visualized using the NiLearn library for Python.

Topic modeling. Although term-based meta-analysis maps in Neurosynth closely resemble the results of manual meta-analyses of the same concepts, there is a high degree of redundancy between terms (e.g. ‘episodes’ and ‘episodic’), as well as

potential ambiguity as to the meaning of an individual word out of context (e.g. ‘memory’ can indicate working memory or episodic memory). To remedy this problem, we employed a reduced semantic representation of the latent conceptual structure underlying the neuroimaging literature: a set of 60 topics derived using latent dirichlet allocation (LDA) topic-modeling (Blei et al., 2003). This procedure was identical to that used in a previous study (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012b), except for the use of a smaller number of topics and a much larger version of the Neurosynth database. The generative topic model derives 60 independent topics from the co-occurrence of all words in the abstracts of fMRI studies in the database. Each topic loads onto individual words to a varying extent, facilitating the interpretation of topics; for example, a working memory topic loads highest on the words “memory, WM, load”, while an episodic memory topic loads on “memory, retrieval, events:”. Note that both topics highly load on the word “memory”, but the meaning of this word is disambiguated because it is contextualized by other words that strongly load onto that topic. Although the set of topics included 25 topics representing non-psychological phenomena—such as the nature of the subject population (e.g. gender, special populations) and methods (e.g., words such as “images”, “voxels”)—these topics were not explicitly excluded as they were rarely the strongest loading topics for any region. For all of our results, we focus on a set of 16 topics that strongly loaded onto lateral frontal cortex clusters (Table 3.1). These topics were obtained by determining the two strongest loading topics for each region.

Topic name	Top words
action	action actions motor goal mirror planning imitation execution
attention	attention attentional visual spatial search location orienting target
conflict	conflict interference incongruent stroop congruent selection competition color
emotion	emotional emotion regulation affective pictures emotions arousal affect
gaze	eye gaze eyes movements saccades target saccade visual
inhibition	inhibition inhibitory stop motor sustained nogo transient suppression
memory	memory retrieval encoding recognition episodic items recall words
mentalizing	social empathy moral person judgments mentalizing mental mind
motor	motor movement movements sensorimotor finger somatosensory sensory force
novelty	target targets novelty oddball distractor distractors deception mismatch
pain	pain stimulation somatosensory painful intensity sensory chronic noxious
reward	reward sleep anticipation monetary rewards motivation incentive loss
semantics	semantic words word lexical verbs abstract meaning verb
speech	speech auditory sounds sound perception voice acoustic listening
switching	switching rule executive switch rules flexibility shifting aggression
WM	memory working wm load verbal maintenance delay encoding

Table 3.1. Topics most strongly associated with lateral frontal regions. Eight strongest loading words for each topic are listed, in descending order of association strength.

Meta-analytic functional preference profiles. We generated functional preference profiles by determining which psychological topics best predicted each cluster’s activity across fMRI studies (Figure 3.2C). First, we selected two sets of studies: studies that activated a given cluster– defined as activating at least 5% of voxels in the cluster– and studies that did not– defined as activating no voxels in the cluster. For each cluster, we trained a naive Bayes classifier to discriminate these two sets of studies based the loading of psychological topics onto individual studies. We chose naive Bayes because (i) we have previously had success applying this algorithm to Neurosynth data (Yarkoni et al., 2011); (ii) these algorithms perform well on many types of data (Androutsopoulos et al., 2000), (iii) they require almost no tuning of parameters to achieve a high level of performance; and (iv) they

produce highly interpretable solutions, in contrast to many other machine learning approaches (e.g., support vector machines or decision tree forests).

We trained models to predict whether or not fMRI studies activated each cluster, given the semantic content of the studies. In other words, if we know which psychological topics are mentioned in a study how well can we predict whether the study activates a specific region? We used 4-fold cross-validation for testing and calculated the mean score across all folds as the final measure of performance. We scored our models using the area under the curve of the receiver operating characteristic (AUC-ROC)— a summary metric of classification performance that takes into account both sensitivity and specificity. AUC-ROC was chosen because this measure is not detrimentally affected by unbalanced data (Jeni et al., 2013), which was important because each region varied in the ratio of studies that activated it to the studies that did not.

To generate functional preference profiles, we extracted from the naive Bayes models the log odds-ratio (LOR) of a topic being present in active studies versus inactive studies. The LOR was defined, for each region, as the log of the ratio between the probability of a given topic in active studies and the probability of the topic in inactive studies, for each region. LOR values above 0 indicate that a psychological topic is predictive of activation of a given region. To determine the statistical significance of these associations, we permuted the class labels and extracted the LOR for each topic 1000 times. This resulted in a null distribution of LOR for each topic and each cluster. Using this null distribution, we calculated p-

values for each pairwise relationship between psychological concepts and regions, and reported associations significant after controlling for multiple comparisons using False Discovery Rate with $q < 0.01$. Finally, to determine if certain topics showed greater preference for one cluster versus another, we conducted exploratory, post-hoc comparisons by determining if the 95% confidence intervals (CI) of the LOR of a specific topic for a one region overlapped with the 95% CI of the same topic in another region. We generated CIs using bootstrapping, sampling with replacement and recalculating log-odds ratios for each region 1000 times. A full reference figure of the loading between topic and regions, including CIs, is available in Appendix I. The ordering of the labels around the polar plot was determined using hierarchical clustering with average linkage, resulting in an order that concisely conveyed the functional differences between LFC's sub-regions.

Results




Figure 3.4. Silhouette scores, a measure of intra-cluster cohesion, for lateral frontal

cortex from 4-100 whole-brain clusters. We chose to focus on three levels of granularity— 5, 33, and 70 whole-brain clusters— highlighted in red. Although silhouette score did not strongly indicate a single optimal solution, solutions from 7-32 clusters received relatively low scores, and thus were avoided.

Hierarchical clustering of lateral frontal cortex. We identified spatially dissociable regions on the basis of shared co-activation profiles with the rest of the brain (Chang et al., 2013; S. M. Smith et al., 2009; Toro et al., 2008), an approach that exploits the likelihood of a voxel co-activating with another voxel across studies in the meta-analytic database. To avoid defining arbitrary boundaries for regions in the lateral frontal cortex, we clustered the entire cortex and selected clusters that had a significant number of voxels within an anatomically defined LFC mask— allowing clusters to span beyond LFC and excluding clusters that were primarily outside of LFC.

In order to map structure to function across various spatial scales, we extracted 4– to 100– flat whole-brain clusters and guided the selection of three scales for further analysis by evaluating the inter-cluster coherence within LFC using the silhouette score (Figure 3.4). Silhouette scores began moderately high from 4-6 clusters, before dipping from 6-32 whole-brain clusters and rising consistently again after 33 clusters. This pattern was consistent with evidence suggesting there are around six distributed whole-brain ‘networks’ (Yeo et al., 2011) and previous observations that the accuracy of clustering increases monotonically with the number of clusters (Craddock et al., 2012; Thirion et al., 2014).

Since silhouette scores did not suggest a strong preference for a single dominant solution, we focused on three well-spaced levels of granularity– avoiding the trough between 6-32 clusters. At the low end, we chose 5 clusters– as this scale had the highest silhouette score of network-level solutions– and an intermediary solution of 33 clusters, as this scale showed a substantial increase in coherence than 32 clusters. At a finer scale, we chose to focus on the 70- cluster solution, as this solution resulted in 15 spatially contiguous clusters in LFC clusters that largely separated from distal brain regions in parietal cortex. At coarser scales, clusters were not spatially contiguous and resembled “networks”. Importantly, since our goal was to identify spatially contiguous regions in LFC, we primarily used the two coarser solutions to examine the hierarchical organization of the finer grained regions, and organize subsequent analyses accordingly.




Figure 3.5. Whole-cortex co-activation based hierarchical clustering reveals 4 networks in lateral cluster that fractionate into constituent sub-regions. From a full cortical parcellation, we selected clusters in lateral frontal cortex at three spatial scales. (Left) From five whole-cortex clusters, we identified four clusters with voxels in lateral frontal cortex resembling large-scale whole-brain networks: task-positive control network (red), default network (purple), somatosensory-motor network (green) and the ventral attention network (blue). At 33 whole-brain clusters, these networks fractionated into 10 LFC clusters and at 70 clusters two clusters in the fronto-parietal network further fractionated into 3-4 clusters, resulting in a total of 15 clusters in lateral frontal clusters. At 70 clusters, all clusters were spatially contiguous and a majority of their voxels fell in our LFC mask, hence our focus on this granularity. Lighter colors indicate voxels outside of LFC in the 5 and 33 cluster solutions.

In the five cluster whole-cortex solution (Figure 3.5), we identified four broad networks that showed moderate correspondence to previously described large-scale networks (Power et al., 2011; Yeo et al., 2011). Although the functional networks we identified were not isomorphic with resting-state networks— in part because our measure of fit suggested choosing a coarser solution— these results are consistent

with the view that large-scale functional networks supersede the anatomically defined area of lateral frontal cortex as organizational units.

Spanning around half of LFC, primarily in prefrontal cortex, we identified a cluster consistent with previous descriptions of the fronto-parietal control network (Yeo et al., 2021; dice coefficient (d) = 0.56), which also spanned medial-frontal and anterior insular aspects of the ventral attention network (d = 0.21). Also in prefrontal cortex, we identified a cluster closely (d = 0.62) matching previous extensive previous descriptions of the default network (Andrews-Hanna, 2012). In posterior LFC, we identified two clusters primarily situated in primary motor cortex (PMC) within LFC. The more dorsal of the two moderately overlapped with Yeo's somatosensory-motor network (dice coefficient (d) = 0.36), encompassing dorsal primary motor and somatosensory cortices and the supplementary motor area (SMA), while also extending slightly more posterior into lateral aspects of Yeo's dorsal attention network (d =0.31). Immediately ventral, we identified a second network with moderate overlap to Yeo's somatosensory-motor network (d =0.44) that also spanned lateral aspects of what is referred to as the ventral attention or cingulo-opercular network (d =0.34).

Each of these networks further fractionated into 1-9 clusters in the $k = 70$ solution that were almost entirely located in LFC (Figure 3.6). To better understand the organization and function of each of these clusters— for each network separately— we describe their anatomical and functional correspondence. Because the two posterior networks resulted in only three clusters at $k = 70$, we have

grouped them in subsequent analyses as ‘posterior LFC’. To provide direct insight into the functions of the clusters we identified, we applied two approaches. First, we determined which other brain regions co-activate with each cluster across studies, revealing distinct whole-brain functional networks for each cluster. Second, we used semantic data from Neurosynth to determine which psychological states predict the activation of each cluster, resulting in a unique meta-analytic functional preference profile for each.



Figure 3.6. Final set of fifteen LFC clusters derived from a $k = 70$ whole-brain co-activation based clustering.

Posterior LFC

Anatomical correspondence. In the far posterior aspects of lateral frontal cortex, we identified two networks, which we refer to as dorsal and ventral somatosensory-motor networks (Figure 3.7a). In the dorsal network, we identified two clusters that were located almost entirely in LFC: dorsal and ventral lateral premotor cortex— PMd and PMv, respectively. Both of these areas were located in the dorsal half of BA 6 (Eickhoff et al., 2007), although PMd was located slightly

more anterior; PMv slightly encroached into primary motor cortex as a result of its slightly more posterior location. Notably, PMd also included a small number of voxels outside of LFC in the right primary somatosensory cortex. Investigation at finer-grained levels of granularity indicated these voxels remained grouped even past 100 whole-cortex clusters, suggesting the co-activation of these regions is strongly coupled. At the coarser solution of $k=33$, PMd grouped with the medial supplementary motor area (SMA), suggesting these regions perform relatively similar roles in motor function.

Immediately ventral to these two regions was the only lateral frontal cluster associated with the ventral somatosensory network. This center of this cluster was located in the fundus of the central sulcus, and extended onto ventral primary motor cortex (PMC) and ventral primary somatosensory cortices (SCX); as such, we refer to this cluster as “PMC/SCX_v”. The lack of a clean boundary between clusters within and outside LFC, and across distinct cytoarchitectonic areas suggests that anatomical boundaries do not necessarily reflect task-dependent functional boundaries, at least for these sensori-motor regions.




Figure 3.7. Meta-analysis of posterior LFC. a) Individual clusters projected onto an inflated surface. PMd and area PMv belonged to a dorsal somatosensory-motor network and area PM/SCZv to a ventral somatosensory-motor network b) Differences in co-activation. Colored voxels activated more frequently in the seed cluster of the same color was also active. c) Functional preference profiles for each cluster, revealing distinct psychological signatures. Strength of association is measured in log odds-ratio (LOR), and permutation-based significance ($q < 0.05$) is indicated next to each topic by color-coded dots corresponding to each cluster.

Meta-analytic co-activation profiles. Next, we examined differences in co-activation with the rest of the brain across fMRI studies, in order to better understand the functional differences between these regions (Figure 3.7b). To do so, we directly contrasted co-activation patterns of the three clusters— i.e., we sought to identify voxels across the brain that co-activated to a stronger degree with each cluster than with the other two (note that each cluster trivially co-activates with

itself, as studies that activate a given cluster necessarily show robust activity within that cluster). PMd showed greater co-activation across parietal cortex, extending from the inter-parietal sulcus (IPS) into the superior parietal lobule (SPL), and mid-DLPFC— regions implicated in executive function and goal directed cognition. PMv, on the other hand, showed greater co-activation with ventrolateral prefrontal cortex (i.e. IFG pars orbitals) and pre-SMA. Although co-activation cannot directly speak to the functional role of these regions, these results suggest dorsal and ventral premotor cortex perform dissociable roles. Finally, PMC/SCXv showed a somewhat more distinct pattern, exhibiting greater co-activation with the posterior insula (pIns) and secondary somatosensory cortex (SII), posterior MCC and anterior medial prefrontal cortex (mPFC). This more distinct pattern is consistent with PMC/SCXv's grouping into a different network from the two premotor clusters.

Meta-analytic functional preference profiles. Next, we used a data-driven approach that surveyed a broad range of fMRI studies to determine which psychological states differentially recruited each LFC cluster (Figure 3.7c). For each cluster, we trained a multivariate classifier to predict if studies activated the cluster using a set of 60 psychological topics derived by applying a standard topic modeling approach to the abstracts of articles in the database (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012a). From the resulting fitted classifiers, we calculated a measure of how strongly each topic indicated that a study activated each cluster (measured as the log odds-ratio [LOR] of the probability of each topic in

studies that activated a given cluster to the probability of the same topic in studies that did not activate the cluster). Values over 0 indicate that the presence of that topic in a study predicts activity in a given region. We report the results of 16 psychological topics that loaded strongly onto LFC regions (Table 3.1) and restrict interpretation to significant associations using False Discovery Rate (FDR; $q < 0.01$). In addition, whenever we comparatively discuss sets of regions, we determined significance if the 95% confidence interval (CI) of a given topic did not overlap between two regions. As the latter comparisons are post-hoc and exploratory, caution in interpretation is warranted. A full reference figure of 95% CI for all regions is reported in Appendix I.

All three sub-regions showed relatively similar functional signatures, justifying their grouping. However, PMC/SCXv showed the most distinct signature, consistent with the relatively different co-activation pattern exhibited by this region. Notably, all three clusters were significantly associated with motor function (e.g. ‘motor’ & ‘gaze’), although this relationship was stronger for the two pre-motor clusters. However, only pre-motor clusters were associated with higher-level motor planning (i.e. ‘action’) and working-memory– suggesting these clusters are involved in relatively high-level motoric function. Moreover, PMd was significantly associated with ‘conflict’ and ‘attention’, consistent with its stronger co-activation with regions implicated in attention control such as lateral parietal cortex and DLPFC. In contrast, ‘PMC/SCZv’ was significantly associated with language topics (i.e. ‘semantics’ and ‘speech’), consistent with its relative proximity to the primary

auditory cortex. Moreover, this cluster was strongly associated with ‘pain’, consistent with its co-activation with pIns and SII– key pain processing regions (Rolls et al., 2003; Wager et al., 2013). In sum, although clusters in posterior lateral frontal cortex showed relatively similar functional profiles– focused primarily on motor function– the present results suggest a dorsal-ventral gradient of function, with more dorsal regions being more involved in attentional control, and more ventral regions with language and pain processing.

Fronto-parietal network

Anatomical correspondence. The majority of lateral frontal cortex belonged to a fronto-parietal whole-brain network that additionally spanned the lateral parietal cortex (LPC), anterior insula (aI), pre-SMA, mid-cingulate cortex (MCC), and the precuneus outside of LFC. This network resembled previous descriptions of the fronto-parietal network in addition to including some regions sometimes ascribed to cingulo-opercular (Power et al., 2011) or ventral attention networks (Yeo et al., 2011), such as aMCC. This network fractionated into nine LFC clusters in the 70 whole-brain solution (Figure 3.8). Four of these clusters grouped at $k=33$ into a single ‘mid’ lateral prefrontal cluster, and three grouped into a ‘rostral’ lateral prefrontal cluster. Two additional clusters did not group until later in the clustering process, but we have organized them into a ‘caudal’ group due to their spatial proximity in caudal LPFC.

In caudal LPFC, we identified two distinct dorsal clusters (Figure 3.8a). The most posterior cluster was located anterior to the premotor cortex— near BA 6 and 8— extending from lateral superior frontal gyrus, into middle frontal gyrus dorsal to the intermediate frontal sulcus (i.e. area 9/46d). This cluster, which we refer to as area 6/8, was co-located with the frontal eye fields (FEF)— a region important for volitional eye saccades (Paus, 1996). Lying anterior and ventral to ‘area 6/8’, we identified a cluster that spanned a small area of caudal area 9/46 (cluster 9/46c), extending across the intermediate frontal sulcus into dorsal and ventral cytoarchitectonic sub-divisions. Although this cluster extended well into mid-DLPFC, as far as area 9/46v, it was notable that this cluster did not group with other mid-LPFC clusters in whole-brain clustering, suggesting this cluster may exhibit a dissociable functional signature.

Anterior and ventral to caudal LPFC, we identified four clusters that grouped together into a single cluster at 33 clusters of granularity (Figure 3.8b). These four clusters spanned most of what many refer to as ‘mid’ lateral prefrontal clusters. Most dorsally, we identified a cluster that spanned the majority of area 9/46v ventral to the intermediate frontal sulcus, extending well into the fundus of the inferior frontal sulcus. Although this cluster was primarily lateralized to the left hemisphere, it did include a small number of voxels in right 9/46v. As such, we refer to this cluster as ‘9/46v’. In the right hemisphere, we identified a cluster that spanned the entirety of BA45. This cluster, which we refer to as right IFG (IFG [R]), extended dorsally into the inferior frontal sulcus near area 9/46v. Notably, the

contralateral analogue of this cluster was not a part of the fronto-parietal network. This hemispheric asymmetry is consistent with the observation that right IFG activation is consistently observed during goal-directed cognition, and hence groups with regions in the fronto-parietal network, whereas left IFG is more consistently involved in language processing. Anterior to both of these clusters, we identified a bilateral cluster located in the caudal end of the inferior frontal sulcus, spanning precentral, inferior frontal and middle frontal gyri. The cluster was mostly buried in the fundus of the sulci and is consistent with previous reports, and co-activation based parcellations, of an IFJ region (e.g. MNI coordinates 48, 4, 33; Brass, Derrfuss, Forstmann, & Cramon, 2005; Muhle-Karbe, Derrfuss, Lynn, Neubert, Fox, Brass, & Eickhoff, 2015a). Finally, ventral to this cluster in the right hemisphere, we identified a fourth cluster (cluster 44 [R]), which was located in the posterior end of IFG, spanning BA44 and abutting BA6.

At the most anterior portion of LFC, we identified three bilateral clusters which grouped together in the $k = 33$ solution into a single fronto-polar cluster (FP) (Figure 3.8c). These three clusters spanned the entirety of lateral BA10 and fractionated along a ventral-dorsal axis, consistent with a recent DTI parcellation of frontal pole (Orr et al., 2015). Notably, none of these clusters extended into medial aspects of BA10 or ventrally into orbitofrontal cortex, consistent with recent cytoarchitectonic evidence of a lateral-medial distinction in frontal pole (Bludau et al., 2014). The most dorsal of these three clusters extended well into rostral portions of BA 9/46d, bilaterally. As such, we refer to this cluster as 9/46dr. The

next two clusters separated along a dorsal/ventral axis in BA 10; as such, as refer to these clusters as '10d' and 10v', respectively.




Figure 3.8. Anatomical location and meta-analytic contrast of lateral frontal clusters of the fronto-parietal network. Left) a) Two clusters located in caudal frontal cortex. b) Four clusters located in mid-lateral pre-frontal cortex, which grouped together into single cluster at 33 whole-cortex clusters of

granularity. c) Three clusters located in rostrolateral pre-frontal cortex, which grouped together into a single cluster at $k = 33$. Clusters were assigned labels corresponding to cytoarchitectonic areas whenever possible. In cases where the region spanned many cytoarchitectonic areas, broader anatomical (e.g. inferior frontal junction [IFJ]) (inferior frontal junction [IFJ]) labels were assigned. Right) Colored voxels indicate significantly greater co-activation with the seed region of the same color than other lateral frontal regions in the fronto-parietal network. Images are presented using neurological convention and were whole-brain corrected using a false discovery rate (FDR) of $q = 0.00001$ to prevent excessive overlap.

Meta-analytic co-activation profiles. Next, we contrasted the whole-brain co-activation of the LFC clusters that fell within the fronto-parietal network, which revealed distinct patterns for each region (Figure 3.8, right panel). We observed a striking pattern of co-activation differences, such that the majority of clusters co-activated with distinct sub-portions of other cortical association areas. Across parietal cortex, each LFC cluster co-activated most strongly with distinct areas across a gradient extending from tempo-parietal junction (TPJ) to the lateral parieto-occipital sulcus. For example, clusters ‘9/46c’ and all fronto-polar clusters showed greater co-activation with parietal cortex ventral to the intraparietal sulcus. In contrast, area ‘6/8 and all four ‘mid’ LPFC clusters showed greater co-activation with the intraparietal sulcus itself and areas slightly dorsal to it. This gradient of co-activation across LPC is consistent observations from rsfc-fMRI suggesting association cortex is composed of parallel interdigitated networks that are highly integrative in nature (Yeo et al., 2011).

Similarly, all clusters except right IFG co-activated most strongly with distinct portions of pre-SMA and MCC. Generally, more anterior clusters co-

activated more strongly with more anterior portions of pre-SMA/MCC. For instance, cluster '10d' co-activated most strongly with a portion of mid-cingulate cortex that extended into perigenual ACC. In contrast, cluster '44 [R]' co-activated with a more posterior portion in the supplementary motor area (SMA). Given the importance of MFC for motoric and executive function, and the strong coupling of these regions at a network, these results suggest that distinct areas of lateral frontal cortex work in concert with distinct medial areas in support of goal-directed cognition.

In addition to the differences in co-activation across parietal and frontal cortex, we observed strong differences in co-activation with the insula. Cluster 44 [R] showed the most distinct pattern, exhibiting strong co-activation with posterior insula (pI), an important region for pain and sensorimotor processing (Chang et al., 2013). In contrast, the remaining clusters showed strong co-activation with different portions of anterior insula (aI). For instance, IFJ co-activated most strongly with dorsoanterior Insula, a sub-region implicated in goal-directed cognition. In contrast, areas 10d and 10v generally showed greater co-activation with ventroanterior insula, which has been implicated in chemo-sensory processing.

In sum, the primary differences in co-activation across fronto-parietal regions in lateral frontal cortex were within highly active 'task-positive' association areas, perhaps due to a requirement for different types of multi-modal information depending on the type of cognitive control that needs to be exerted by each region. This finding is consistent with the hypothesis that association cortex across the brain is composed of parallel inter-digitated networks with high levels of cross talk


between each other (Yeo et al., 2011). The exceptions to this pattern were area 9/46c and the two clusters in BA10, which showed stronger co-activation areas of the default network, including vmPFC and PCC. This pattern is consistent with the hypothesis that frontal pole may serve as a relay between the task-positive fronto-parietal network and the internally-oriented default network (Burgess, Dumontheil, & Gilbert, 2007).

Meta-analytic functional preference profiles. Having observed distinct patterns of co-activation between fronto-parietal LFC clusters, we probed the semantic data in Neurosynth to determine if psychological states differentially recruited each cluster (Figure 3.9). Consistent with distributed role for the fronto-parietal network in goal-directed cognition, all nine clusters were significantly associated with working-memory, all clusters except 10d and 10v were associated with conflict, and seven clusters were associated with switching. The present results are inconsistent with focal anatomical locations for high-level executive processes and instead suggest these processes likely rely on distributed firing across fronto-parietal network to support goal-directed cognition in the face of interference and conflict (Nee & Brown, 2012).


Caudal fronto-parietal LFC. Despite the overall functional similarities between these regions across core aspects of cognitive control, each cluster exhibited distinguishing functional characteristics. Consistent with its co-location with the frontal eye fields, ‘6/8’ was the only cluster significantly associated with saccadic eye movements (i.e ‘gaze’) in the fronto-parietal network, and was also associated with

‘attention’. This pattern suggests that the area ‘6/8’ may be important for directing attention to relevant stimuli in the external environment to support downstream information processing. However, ‘6/8’ was also significantly associated with ‘action’– a topic representing motor planning– as well as a ‘working-memory’ topic. This latter result is notable as a recent lesion study suggests that FEF may play a causal role in working memory (Mackey, Devinsky, Doyle, Meager, & Curtis, 2016). As such, these present results suggest the region containing the FEF is not merely involved in saccadic eye movements, but plays an important role in higher-level cognition.

a) Caudal



b) Mid



c) Rostral




Figure 3.9. Meta-analytic functional preference profiles for lateral frontal regions in the fronto-parietal network.

Each cluster was profiled to determine which psychological concepts best predicted its activation. Each of the three functional groups we identified showed distinct functional profiles, although appreciable variation was observed for each individual cluster. Strength of association is measured in log odds-ratio (LOR), and permutation-based significance corrected using false discovery rate (FDR) of $q = 0.01$ is indicated next to each psychological concept by color-coded dots corresponding to each region.

Cluster 9/46 showed the least distinctive functional signature, showing no significant associations outside of core EF processes. Given that this cluster did not join with other fronto-parietal regions until much later in the clustering process suggests this region may be involved in psychological states not well characterized by our topic model, or is involved in a domain-general process that supersedes these more specific processes.

Mid fronto-parietal LFC. In mid-LPFC, cluster 9/46v and IFJ showed similar functional profiles, exhibiting robust associations with various executive functions (e.g. ‘wm’, ‘conflict’, ‘switching’) in addition to semantics. Cluster 9/46v showed a particularly strong association with executive control processes, as the relationship between this region and ‘conflict’ was significantly greater than all other fronto-parietal clusters except IFJ. These results are consistent with a hypothesized role for mid-DLPFC as the seat of high-level executive processes, although the cluster we identified is more ventral— extending into inferior frontal sulcus— than has been suggested previously (Petrides, 2005). However, the association of these regions to ‘semantic’ processing suggests that language and executive function are not mutually exclusive processes, consistent with the hypothesis that language relies on core executive function processes. This view is in contrast with models in which left ventrolateral PFC is mainly related to language function.

These results are also consistent with the hypothesis that IFJ is involved in switching (Brass et al., 2005; Derrfuss, Brass, Neumann, & Cramon, 2005b) and is

underappreciated in its contributions to cognitive control. However, many other clusters were similarly strongly associated with switching, suggesting IFJ is not likely to be focally responsible for this phenomenon. However, IFJ was also characterized by its significant association with low and high level motor function (i.e. ‘motor’, ‘action’)– an association shared only by 44 [R] in the fronto-parietal network. As the only region strongly associated with both executive processes and motor function, IFJ is well positioned as a unique mediator between high-level plans and task-sets and low level motoric innervation. This view is consistent with the hypothesis that IFJ is important for integrating motor representations with high-level abstract aspects of cognitive control (De Baene, Albers, & Brass, 2012). The potential ubiquity of such a process across domains may explain its extremely high rate of activation across a wide range of studies. In contrast, cluster 44 [R]– with its much lower associations with executive functions and a significant association with ‘pain’– may be important for introducing negative affective signals that may require an immediate change in plans into such a process. Notably, a similar role has been attributed to anterior mid-cingulate cortex (Cavanagh & Shackman, 2015; Shackman et al., 2011), but present models may overlook area 44’s contributions to this process.

Finally, rIFG, showed a relatively distinct functional signature to other mid LPFC clusters. This cluster was more weakly associated with conflict, working memory and switching– processes not typically attributed to ventrolateral PFC. In contrast, rIFG showed a significant association with ‘inhibition’– consistent with an

extensive literature on the role of right inferior frontal gyrus in inhibitory processes (Aron, Robbins, & Poldrack, 2004; Depue, Orr, Smolker, Naaz, & Banich, 2016; Munakata et al., 2011). rIFG was also strongly associated with ‘emotion’, consistent with the hypothesis that this region is crucial for effective emotion regulation and reappraisal (Frank, Dewitt, & Hudgens-Haney, 2014; Opiella et al., 2015; Wager et al., 2008). However, the relationship between ‘inhibition’ and rIFG was not particularly strong or significantly greater than with other regions, suggesting rIFG may play a more general role that is not well encapsulated by the present topics.

Rostral fronto-parietal LFC. The three most rostral regions of the fronto-parietal network located within LFC showed relatively similarly functional profiles, consistent with their similar pattern of co-activation. Like other clusters in the FPN, these fronto-polar clusters showed robust— although not particularly strong— associations with various executive processes. Notably, both clusters ‘9/46dr’ and ‘10d’ showed a robust association with ‘inhibition’, while cluster ‘10d’ was also significantly associated with ‘novelty’. This pattern was remarkably similar to that shown by ‘rIFG’, suggesting inhibitory control is not the sole provenance of that area. However, these regions were not associated with ‘emotion’— in contrast to rIFG. This pattern is potentially consistent with hierarchical models of control in LPFC, which postulate that more rostral regions represent more abstract goals (Badre & D’Esposito, 2009; Botvinick, 2008). This is particularly plausible in light of the lack of association between these regions and any low-level processes such as motor function or affect. However, it’s not clear the present pattern of results

suggests that these rostral areas are more abstract in nature than mid-DLPFC (e.g. cluster 9/46v), unless ‘novelty’ detection is construed as a more abstract process than conflict processing.

Finally, the most ventral fronto-polar region, cluster ‘10v’, showed a more distinct pattern, with weaker associations with all executive processes. In contrast, this cluster was significantly associated with ‘reward’ (at a lower threshold, $q < 0.05$), consistent with its location near orbitofrontal cluster and its co-activation with vmPFC. These results are consistent with existing hypotheses that suggest that the ventral frontal pole is particularly important for relaying information that represents the value of stimuli to effectively guide goal-directed behavior (Orr et al., 2015).

Default network

Anatomical correspondence. The final network we identified in lateral frontal cortex showed a strong resemblance to previous descriptions of the ‘default network’ (Andrews-Hanna, 2012; Power et al., 2011; Yeo et al., 2011), spanning anterior mPFC, and PCC outside of LFC. Within LFC, we identified three distinct clusters (Figure 8a). The first two clusters were positioned adjacent to each other in ventrolateral prefrontal cortex. The larger of the two spanned the entire left inferior frontal gyrus (IFG [L]) while the third cluster was positioned immediately anterior to it in lateral orbitofrontal cortex and IFG orbitalis in the right hemisphere. The latter region is consistent with cytoarchitectonic area 47/12. Most dorsally, we

identified a third cluster consistent within BA9 (Petrides, 2005) extending from superior frontal gyrus to dorsal middle frontal gyrus across the superior frontal sulcus. The grouping of IFG and BA9 with the default network instead of fronto-parietal is highly consistent with various parcellations based on rs-fMRI (Power et al., 2011; Yeo et al., 2011). BA9, in particular, has long been noted for its lack of input from lateral and medial parietal cortex (Petrides & Pandya 1984, 1999; Cavada & Goldman-Rakic 1989; Andersen et al. 1990). Thus, despite the proximity of area 9 to area 9/46v, the results of the clustering suggest these regions will exhibit a distinct functional profile from fronto-parietal clusters.

Meta-analytic co-activation profiles. Next, we contrasted the whole-brain co-activation of fronto-parietal LFC clusters, revealing distinct patterns for each region (Figure 8b). Left IFG showed the most distinct pattern, co-activating with portions of the fronto-parietal network such as mid-DLPFC and pre-SMA. This pattern is consistent with the fact that this cluster's contralateral homologue grouped with the fronto-parietal network. Hence, although this cluster's connectivity was similar enough to other default network regions to form a part of this network, this region may not be fully dissociable from the fronto-parietal network. Left IFG also showed stronger co-activation with middle temporal gyrus, consistent with a possible role in language. In contrast, cluster 47/12, which is located anatomically near left IFG, showed strong co-activation with orbitofrontal cortex, vmPFC, PCC and anterior temporal lobe— key regions of the default network. Similarly, area 9 showed robust co-activation with the rest of the default

network, firmly placing this region within the default network and not fronto-parietal network.




Figure 3.10. Lateral frontal regions of the default network

a) Individual clusters projected onto an inflated surface. b) Differences in co-activation between the three regions. Colored voxels activated more frequently in studies in the seed cluster of the same color was also active. c) Functional preference profiles for each cluster, revealing distinct psychological signatures for each sub-region. Strength of association is measured in log odds-ratio (LOR), and permutation-based significance is indicated next to each topic by color-coded dots corresponding to each region.

Meta-analytic co-activation profiles. Having observed distinct patterns of co-activation between default LFC clusters, we probed the semantic data in Neurosynth to determine if psychological states differentially recruited each cluster (Figure 3.10c). Consistent with left IFG's co-activation with other fronto-parietal

regions, this cluster was significantly associated with various executive functions, including ‘conflict’, ‘WM’ and ‘switching’. This finding further highlights the distributed nature of core executive processes across frontal regions beyond those that group with the fronto-parietal network. However, it is notable left IFG was not associated with inhibition, consistent with the strong lateralization of this process onto right IFG. Finally, consistent with this region’s overlap with Broca’s area and co-activation with the middle temporal gyrus, left IFG was significantly associated with language topics. Notably, left IFG was the only lateral prefrontal region associated with ‘speech’. However, we did not find an association between left IFG and ‘motor’ or ‘action’, despite the long held belief that Broca’s area is important for motor function in language. The present functional profile of left IFG is consistent with recent electrophysiological data suggesting Broca’s area is involved in the generation of novel speech motor plans, but not mere motor function (Flinker et al., 2015). Left IFG was also notable for its robust association with ‘semantic’ function—more so than any other region. This pattern is consistent with the hypothesis that left IFG is a critical higher-level region in a broader ‘semantic’ system in the brain (Binder et al., 2009).

The two other lateral frontal clusters of the default network showed very distinct functional profiles, as unlike the other prefrontal clusters, they showed no association with any executive processes. This fact is particularly notable for area 9, given its anatomical proximity to mid-DLPFC and the ‘area 6/8’. Instead, both regions showed robust associations with ‘emotion’ and ‘mentalizing’—consistent

with their placement into the default network. These findings are consistent with the hypothesis that these regions, as part of the dorsal medial subsystem of the default network play a critical role in mentalizing and conceptual processing (Andrews Hanna, Smallwood, & Spreng, 2014b; Spreng & Grady, 2010; Spreng, Sepulcre, Turner, Stevens, & Schacter, 2013).

Finally, it was notable that all three clusters that reside within the default network were associated with ‘memory’ function. This is consistent with a long line of evidence supporting the role of these regions in autobiographical, integrally oriented cognition (Andrews Hanna, Saxe, & Yarkoni, 2014a). Moreover, the left IFG is purported to play a key role in controlled memory retrieval (Badre & Wagner, 2007; Snyder, Banich, & Munakata, 2011)– a hypothesis supported by the joint association between executive processes and memory in this region. However, it is also notable that memory was associated with many other clusters in the fronto-parietal networks (i.e. 9/46v, 10d, 10v, 9/46c and right IFG). As such, memory is likely to be widely distributed across lateral frontal cortex and the distinct role played by these regions may require more fine-grained modeling of memory sub-processes.

Discussion

We applied data-driven methods to the largest meta-analytic database available to produce a systematic mapping between discrete lateral frontal anatomy and psychological function. By taking a broad scope both functionally and

anatomically, we provide a comprehensive view of the psychological states that predict activity across this complex and heterogeneous area. The present results are inline with extensive work suggesting multiple whole-brain networks extend into LFC and support a distinct range of psychological functions. Within each network, we identified spatially contiguous subregions that exhibited relatively similar, but dissociable psychological profiles. However, in contrast to strict localizationism, we find that specific psychological states are distributed throughout LFC, consistent with the view that association cortex is composed of parallel integrative circuits (Friston, 2002; Yeo et al., 2011), rather than highly specialized and isolated computational units.

Distributed lateral-frontal regions support goal-directed cognition. A striking pattern evident in our results is the extent to which core executive functions required for externally oriented, goal-directed cognition were distributed throughout (and in some cases beyond) lateral-frontal sub-regions of the fronto-parietal network. This is in contrast to models that hypothesize that specific executive processes are supported by discrete computational units. For example, sustained activity in LPFC during working memory tasks has been hypothesized to reflect the active maintenance of a representation in domain-specific buffers (Baddeley, 2003). However, in the present study we find that working memory preferentially recruits activity across a wide range of regions extending from posterior LFC (i.e. dorsal premotor cortex) to the lateral frontal pole. This is consistent with a more recent view that suggests sustained activity across these

regions reflects domain-general processes which are required to flexibly guide behavior in support of the task goals (Curtis & Lee, 2010; Postle, 2016; Riggall & Postle, 2012). Similarly, updating task representations when switching task sets has been hypothesized to preferentially recruit specific LFC regions such as the inferior frontal junction (Brass et al., 2005; Derrfuss, Brass, Cramon, Lohmann, & Amunts, 2009; Derrfuss, Brass, Neumann, & Cramon, 2005b; Muhle-Karbe, Derrfuss, Lynn, Neubert, Fox, Brass, & Eickhoff, 2015b). However, we find that ‘switching’ is similarly predictive of activity across all LFC fronto-parietal sub-regions. As such, the present findings are more consistent with the view that goal-oriented cognition is supported by distributed ‘controllers’ that rely on highly distributed information processing (Power & Petersen, 2013).

Although the present results are consistent with the importance of network level dynamics, we identified complex multivariate patterns for each sub-region that support dissociable roles. For instance, although many individual regions were associated with core executive functions, only IFJ showed additionally robust associations with high and low level motor function. These results are consistent with the recent appreciation of IFJ as a core cognitive control region (Brass et al., 2005; Muhle-Karbe, Derrfuss, Lynn, Neubert, Fox, Brass, & Eickhoff, 2015b) and suggest this region may be particularly important for biasing motoric representations in support of high level goals. Put differently, this region may be important for resolving response conflict, a role typically ascribed to midcingulate cortex (Botvinick et al., 1999; C. S. Carter et al., 1998). In contrast, area 9/46v in

mid-DLPFC was the region most strongly recruited by core executive processes, but showed no associations with ‘lower-level’ processes such as attention and motor function, suggesting this region is more important for the biasing of abstract representations in more domain-specific regions of posterior cortex (Badre, 2008; Banich, 2009).

We also found a distinct pattern of functional associations in rostral regions of the fronto-parietal network. Although regions of the lateral frontal pole remained significantly associated with core aspects of executive function such as switching and working memory, these regions showed markedly weaker associations with ‘conflict’ and showed strong associations with ‘inhibition’ and— in the case of area 10d— ‘novelty’ detection. These results are consistent with the ‘gateway hypothesis’ (Burgess et al., 2007) which suggests that area 10 is important for re-directing current processing to novel environmental and internal cues. This theory is compatible with recent theories that suggest that context monitoring (Chatham et al., 2012) or attentional capture (Sharp et al., 2010) are important pre-requisites of response inhibition and extends these results by suggesting that such a function may not be the sole provenance of right IFG.

Distinct functional signature of lateral-frontal default network. We observed three sub-regions that are associations with the so-called default network (Andrews-Hanna, 2012) in lateral frontal cortex, consistent with extensive characterization of this network using rsfc-fMRI (Buckner, Andrews Hanna, & Schacter, 2008; Power et al., 2011; Yeo et al., 2011). Notably, despite the close

spatial proximity of these regions to fronto-parietal regions robustly associated with executive function, these regions showed very distinct co-activation and psychological profiles. Particularly in the case of area 9, which lies immediately dorsal to area 9/46, we found no association with executive function topics, suggesting the relatively low functional-anatomical selectivity we observed within the fronto-parietal network was not due to poor spatial resolution in our approach. Instead, areas 9 and 47/12 were preferentially recruited by internally oriented processes such as ‘mentalizing’, ‘emotion’ and ‘memory’. This pattern is consistent with these regions hypothesized role as part of the dorso-medial subsystem of the default network in self-generated conceptual processing (Andrews Hanna et al., 2014b).

However, we observed a unique pattern in the left IFG that suggests this region may play an intermediary role between the default and fronto-parietal networks. Although this region grouped with the default network in coarser clustering solutions, and is present in this network in other well-validated atlases (Power et al., 2011; Yeo et al., 2011), we observed that this region showed significant, although moderate, associations with core executive function topics. Additionally, this region showed a very robust association with semantics as well as speech—consistent with its co-location in Broca’s area. This intermediate pattern suggests language production requires aspects supported by both networks. Speech may require both goal oriented motor control—supported by the fronto-parietal network— and access to personally relevant semantic information— supported by the

default network (Binder et al., 2009; Binder & Desai, 2011). The present results are additionally consistent with the recent hypothesis that Broca's area is important for higher-level aspects of language production, such as choosing the appropriate words, rather than low-level motor function which is likely executed in other regions, such as pre-SMA and SMA (Flinker et al., 2015).

Future challenges and limitations. A difficult challenge in cognitive neuroscience is developing the appropriate psychological constructs that distinguish activity in related brain regions. Appropriately modeling the differences between nuanced psychological concepts is particularly difficult for large-scale meta-analyses, as there is no established ontology of psychological constructs, unlike in fields such as genetics (Botstein, Cherry, Ashburner, & Ball, 2000). In the present study, we used a data-driven set of topics derived from the abstracts of fMRI papers to represent major psychological phenomena. Although these topics are a major improvement on more simple term based features, due to their data-driven nature they are likely to misrepresent psychological dimensions which are hypothesized to be important for differentiating regions. For example, in our set of 60 topics, only a single topic represented long term memory function, and likely combined memory retrieval and autobiographical memory processes. Although the Neurosynth framework allows researchers to develop custom meta-analyses that can be used to test apriori predictions, the myriad of combinations in which studies can be combined is not conducive to establishing the optimal differentiating dimensions of psychological function between regions.

The classification based approach we employed is a step in the direction of quantifying the extent to which a given set of psychological features explains variability in brain activity. A promising future direction is to use classification based approaches to find the psychological dimensions that best differentiate patterns in activity between related regions, such as regions within a network. In combination with the adoption of standardized cognitive ontologies, such as the Cognitive Atlas (Poldrack et al., 2011; Poldrack & Yarkoni, 2016), such large-scale approaches should help the development of novel theories of functional brain organization. Moreover, given the limited quality of the summarized coordinate based data in Neurosynth (Salimi-Khorshidi et al., 2009) the widespread sharing of richer statistical images in databases such as NeuroVault (Gorgolewski et al., 2015) will greatly improve the fidelity of future meta-analyses.

Conclusion. In the present study, we used relatively unbiased data-driven methods to comprehensively psychological states to individual regions in lateral frontal cortex. These regions were found to belong to large-scale whole-brain networks and generally shared functional properties with other regions in the same network. Moreover, we found that various specific psychological processes which have been previously hypothesized to map onto specific brain regions were widely distributed throughout lateral frontal cortex. However, we identified dissociable functional signature for each sub-region, suggesting that lateral frontal cortex supports a wide variety of psychological state through a mixture of network-level dynamics and moderate degree of functional specialization.

CHAPTER 4

Cross-modal evaluation of whole-brain atlases

Dividing the brain into non-overlapping spatially contiguous regions is of much interest to the scientific community for both theoretical and pragmatic reasons. From a theoretical standpoint, it has been hypothesized that discrete regions perform selective computational roles, such as the recognition of faces (Kanwisher, McDermott, & Chun, 1997), detection of motion (Martinez-Trujillo et al., 2005) and recognition of fear (Adolphs, Tranel, Damasio, & Damasio, 1995) among others. Although more advanced analysis techniques suggest that such representations are likely to be coded in a much more distributed fashion (Haxby et al., 2001; Wager et al., 2013), it is nonetheless theoretically useful as simplifying assumption to describe the brain as a series of cortical areas that differ along various properties that modulate their information processing abilities— such as structure, connectivity and functional associations (Eickhoff et al., 2015).

From a pragmatic standpoint, in the analysis of functional MRI data— especially when the researcher has strong apriori predictions and power is low— it is useful to reduce the dimensionality of the brain using a region of interest (ROI) approach. In such an approach, instead of conducting a whole-brain analysis across thousands of voxels, researchers use independently defined ROIs to extract BOLD

signal from brain regions hypothesized to play a role in the task at hand (Poldrack, 2007). ROIs are typically selected either from previous studies that targeted similar psychological processes, or from one of the many existing brain atlases that provide comprehensive sets of regional boundaries. These brain atlases are typically constructed by grouping together regions with similarities in micro-anatomical structure (e.g. cyto-, receptor-, and myelo- architecture; (Amunts & Zilles, 2015; Mazziotta et al., 2001; Vogt, 2009), macroanatomical structure (e.g. gyrification; (Desikan et al., 2006; Eickhoff et al., 2007), anatomical connectivity (e.g. probabilistic tractography; (Beckmann et al., 2009; Johansen-Berg et al., 2004; Neubert et al., 2014; Sallet et al., 2013), functional connectivity (e.g. Gordon et al., 2015; Power et al., 2011; Shen et al., 2013; Yeo et al., 2011) or meta-analytic co-activation (Eickhoff et al., 2015).

However, choosing among these atlases is difficult for a variety of reasons. First, brain atlases often have inconsistent regional definitions vis a vis each other (Bohland, Bokil, Allen, & Mitra, 2009). Second, it is not well established how well relatively static properties of the brain— such as anatomy— or functional connectivity measured at rest correspond to the pattern of brain activity observed during behavioral performance. As such, the choice of an atlas for selecting ROIs, or interpreting the functional significance of whole-brain fMRI analysis is typically unprincipled and requires making somewhat arbitrary and flexible choices.

Improper ROI choice, however, can potentially have detrimental effects on the functional significance of the subsequent results. For example, if one were

interested in building a classifier to predict if subjects had observed motion in a visual paradigm, a useful ROI to include in the analysis would region MT— an area known to respond to motion (Dubner & Zeki, 1971; Maunsell & Van Essen, 1983). However, if this ROI was improperly specified and included voxels from regions not sensitive to motion, such as area V3, the classifier’s ability to predict subjects’ experienced motion would suffer.

This is not merely a pragmatic concern. The proper definition of brain regions is critical for theory development. For example, area 9 in lateral prefrontal cortex is hypothesized to belong to the default network and play a role in internal mentation, while the spatially nearby area 9/46d plays a much different role in externally oriented, goal-directed cognition. Properly specifying the boundaries between these regions is critical for understanding their functional role. Improperly specifying area 9 as a ‘task control’ region in a study could lead to misleading theoretical conclusions.

For these various reasons, an important but relatively unstudied question is the quantitative evaluation of existing brain atlases. In the only such study to our knowledge, Thirion et al., 2014 evaluated the impact of algorithm choice on the accuracy and reproducibility of fMRI parcellations. Using simulated data with pre-specified numbers of “true clusters”, Thirion and colleagues demonstrated the difficulty of recovering the true number of underlying parcels, underscoring the difficulty in choosing an optimal number of clusters. Next, using task contrasts from fMRI data in real subjects, they compared the performance of k-means, Ward

hierarchical and spectral clustering. They found that although k-means clustering produced the most accurate representation of the underlying data, Ward hierarchical clustering produced more reliable parcellations. As such, they concluded that Ward hierarchical clustering provides a reasonable compromise between accuracy and reproducibility for clustering fMRI data.

However, several important questions remain. First, it is not known if the same recommendations Thirion et al., 2014 put forth apply to summarized coordinate based meta-analytic data— such as the data in the Neurosynth database. Although they used task-related fMRI to develop and test their parcellations, the range of activation patterns represented in their data was necessarily limited, as they studied a total set of 19 task contrasts across two studies. As such, a primary goal of the present study is to evaluate the external validity of various parcellations using a more diverse range of psychological paradigms using large-scale meta-analytic data. Second, although Thirion and colleagues demonstrated that functionally derived parcellations better fit task-fMRI data than anatomical atlases, the majority of existing atlases are derived from fMRI at rest. In the present study we sought to test if parcellations derived from functional data across various tasks demands better represent the underlying functional organization of the brain than those defined at rest. We did so in a variety of ways.

First, we evaluated the choice of algorithm and preprocessing strategy by quantifying the within-cluster cohesion and reproducibility across various spatial scales. Second, we present a novel method for evaluating parcellations that assesses

their ability to predict a diverse range of psychological states, under the assumption that atlases with more functionally homogenous regions will perform better. Using this method, we compare the performance of meta-analytic co-activation parcellation to five well known anatomical and resting state atlases across various spatial scales. We conclude by making recommendations for future researchers and propose that regional boundaries derived from large-scale meta-analysis better reflect the underlying functional-anatomical structure of the brain than atlases from other modalities.

Materials and Methods

Dataset. We analyzed version 0.6 of the Neurosynth database (Yarkoni et al., 2011), a repository of 11,406 fMRI studies and over 410,000 activation peaks that span the full range of the published literature. Each observation contains the peak activations for all contrasts reported in a study’s table as well as the frequency of all of the words in the article abstract. A heuristic but relatively accurate approach is used to detect and convert reported coordinates to the standard MNI152 coordinate space (see Yarkoni et al., 2011).

Meta-analytic co-activation clustering. The goal of meta-analytic co-activation is to group together functionally similar voxels of the brain under the assumption that regions with similar function will appear in the same studies, i.e co-activate (Eickhoff et al., 2015). As such, a typical approach in meta-analytic clustering is to create an array that represents the distance between voxels as the

extent to which each voxel in the brain co-activates across studies with every other voxel in the brain. This array is generated by computing the pairwise distance between a vectorized representation of each voxel's activation across studies in the database— where a value of 1 indicates the voxel was active in a given study, a value of 0 indicates it was not, excluding voxels with very low signal (i.e. active in less than 100 studies). However, as this procedure would result in a large, computationally intractable matrix (151,527 by 151,527 voxels), we use principal components analysis (PCA) to reduce the dimensionality of one axis to 100 principal components. Next, we calculate the pairwise Pearson correlation between each principal component and every voxel in the brain, resulting in a matrix C^{vp} (here v are voxels to be clustered and p are whole-brain principal components), which represents the frequency with which each voxel co-activates with the rest of the brain. This matrix C^{vp} is then entered into the clustering algorithm in order to cluster voxels in the brain to a given spatial granularity.

However, at this step, an additional standardization pre-processing step is potentially necessary to ensure neuroscientifically interpretable results. C^{vp} is purported to represent the difference in co-activation between every voxel in a manner akin to the resting state functional connectivity. That is, if two voxels show a high degree of co-activation with the same set of voxels (e.g. the fronto-parietal network) and low co-activation with another set (e.g. default network), these voxels are thought to show similar differences to the rest of the brain and should be grouped together. However, we have anecdotally observed that clustering

algorithms are sensitive to the base rate of activation of voxels. Voxels that are frequently active across studies, such as the insula, tend to form much smaller clusters than regions with a lower rate of activation. This artifact likely arises because more frequently active voxels show a greater range of differences with the rest of the brain. However, since from a theoretical standpoint we are interested in the relative difference in the pattern co-activation between voxels, standardizing the co-activation matrix may result in more evenly sized and interpretable clusters. An additional goal of this study is to test if standardizing the co-activation matrix C^{vp} within voxels will correct for this artifact without introducing other unwanted costs.

Clustering algorithms.

K-means. K-means clustering is among the most popular clustering algorithms across a variety of domains. K-means groups voxels by minimizing inertia (i.e. within-cluster sum-of-squares) between observations (e.g. voxels) and the centroid of their corresponding cluster. A potential drawback of this algorithm is that it assumes that clusters are convex and isotropic, and performs poorly when the underlying structure is not spherical in Euclidean space. K-means clustering requires one to specify the number of clusters and recompute at each level of granularity desired.

Ward's hierarchical clustering. Hierarchical agglomerative clustering is a less popular, but plausible alternative clustering method recommended by Thirion et al., 2014. Hierarchical clustering groups observations using a bottom-up

approach, starting by grouping the two most similar observations, followed by the next most similar pair of observations (or clusters) and so on. Hierarchical clustering produces a dendrogram, which represents the full linkage tree between clusters across levels of granularity. As such, this algorithm does not require one to specify the number of clusters apriori and can provide qualitatively useful information as to the organization of the brain into networks and regions at various spatial scales. Here, we use Ward’s linkage criterion, which minimizes the total within-cluster variance similarly to k-means clustering. Notably, Ward’s method is seldom used for whole-brain clustering because the computational time increases cubically [$\Theta(N^3)$] as a function of samples. We employed the fastcluster algorithm (Müllner, 2013)—a package of libraries that enable efficient hierarchical clustering [$\Theta(N^2)$ —to achieve whole-brain clustering.

Evaluation of parcellations.

Within-cluster cohesion. We assessed the overall quality of clustering solutions using the silhouette coefficient: a measure of how similar each sample is to the cluster it was assigned. The silhouette coefficient is defined as:

$$s = \frac{b - a}{\max(a, b)}$$

where a is the mean intra-cluster distance and b is the distance between a sample and the nearest cluster of which the sample is not a part. Solutions that minimized the average distance between voxels within each cluster received a greater score.

Reproducibility. Given the goal of gleaning neuroscientifically useful information from parcellations—especially given that the grouping of voxels can be interpreted as a plausible computational unit in the brain— it is desirable to choose a solution with high reproducibility (Eickhoff et al., 2015; Thirion et al., 2014). We assessed the reproducibility of the resulting parcellations across the previously outlined various strategy choices by measuring the consistency between pairs of bootstrapped clustering solutions (as in Thirion et al., 2014). We generated bootstrapped solutions by resampling studies from the database with replacement fifty times and applying our previously outlined clustering procedure repeatedly. We then computed the reproducibility between each pair of bootstrapped parcellations using the adjusted Rand index (ARI)— a measure of the similarity of two vectors. Importantly, ARI is impervious to the specific labels assigned to each cluster and is adjusted for chance, allowing for apples-to-apples comparisons across k. ARI scores range from -1 to 1, with 1 indicating a perfect match (and -1 a perfect systematic mismatch) and 0 indicating chance. ARI is defined as:

$$ARI = \frac{\sum_{ij} \binom{n_{ij}}{2} - [\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}] / \binom{n}{2}}{\frac{1}{2}[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2}] - [\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}] / \binom{n}{2}}$$

where n_{ij} , a_i , b_j are pairs of matching observation between the two vectors being compared.

Psychological topic prediction. It follows that clusters that better represent the computational units that generate BOLD signal should serve as better predictors of psychological states as a function of the activation observed in a

given study. Put differently, the selection of a preferred clustering solution can be thought of as a feature engineering problem, in which the goal is to select brain features that best predict psychological states. As such, we assessed the validity of our various clustering solutions, as well as existing parcellations from various brain modalities, on the basis of their ability to predict the presence of psychological states across studies in Neurosynth. Note that in contrast to our previous predictive models, we trained the present models to predict psychological topics using activity within our clusters, rather than predicting activity using a combination of topics. This was done in part to emulate predictive analyses on raw fMRI data, which often use the brain to predict psychological outcomes.

We used a standard set of 100 topics derived using latent dirichlet allocation (LDA) topic-modeling (Blei et al., 2003) that represented the latent conceptual structure underlying the neuroimaging literature (Poldrack, Mumford, Schonberg, Kalar, Barman, & Yarkoni, 2012b). For each topic individually, we selected two sets of studies: studies that loaded onto the topic above a threshold and studies that did not meet this threshold. We selected a topic loading threshold of 0.001 as this resulted on an average of 16% of studies loading onto each topic – a sufficient number for training of classifier but not so large that a heterogeneous set of non-representative studies would load for each. Although we only present results with this threshold, results did not qualitatively vary with different thresholds.

For each cluster, we trained a naive Bayes classifier to discriminate these two sets of studies based on the mean activity within each cluster in a given whole-brain

parcellation map. We chose naive Bayes because (i) we have previously had success applying this algorithm to Neurosynth data (Yarkoni et al., 2011); (ii) these algorithms perform well on many types of data (Androutsopoulos et al., 2000), (iii) they require almost no tuning of parameters to achieve a high level of performance; and (iv) they produce highly interpretable solutions, in contrast to many other machine learning approaches (e.g., support vector machines or decision tree forests).

We used 4-fold cross-validation for assessing the performance of our classifier across all topics. Crucially, we nested the parcellation of the brain into our cross-validation, to avoid optimistically biasing the performance of our classification. In other words, we ensure to use training data to generate clusters, and test the performance of these clusters of new, unseen data. We generated clusters using the 3/4 of our studies that were selected for training, fit our naïve Bayes model to the training data, and tested the performance on the unseen 1/4 of remaining studies. We scored our models using the area under the curve of the receiver operating characteristic (AUC-ROC) – a summary metric of classification performance that takes into account both sensitivity and specificity. AUC-ROC was chosen because this measure is not detrimentally affected by unbalanced data (Jeni et al., 2013), which was important because each topic varied in the ratio of studies that loaded onto it.

Cross modal assessment. We compared the predictive performance of our meta-analytically defined whole-brain parcellations to those derived from other brain modalities. In all of the following comparisons, we ensured our parcellation

closely matched the comparison atlas by re-running our clustering using the same whole-brain masked used in each atlas. First, we compared our parcellation to two widely used anatomical atlases: the AAL and Harvard-Oxford (HO) probabilistic atlas. For the HO atlas, we threshold each cluster at 25% probability and combined the resulting regions to generate flat cluster maps.

We also compared our parcellation to three publically available whole-brain resting-state functional connectivity fMRI (rsfc-fMRI) atlases. The first atlas was produced using a spatially constrained spectral clustering algorithm and computed across a wide range of spatial scales (Craddock et al., 2012), allowing us to compare our parcels across a wide range of granularity. The second atlas is a well-known atlas of large-scale networks across the brain generated by applying k-means clustering to a dataset of 1000 subjects (Yeo et al., 2011). The third atlas by Gordon et al., 2015, was produced by applying boundary mapping on a large-scale high quality resting state dataset (Gordon et al., 2015). This atlas has been previously established to exhibit higher within cluster signal homogeneity than other popular atlases, suggesting this atlas may serve as a rigorous comparison.

Results


Co-activation clustering strategy selection. First, we evaluated various possible strategy choices in meta-analytic co-activation clustering, including algorithm choice. Across the whole-brain, we extracted clusters across various levels of granularity (i.e. 20-180 regions) using k-means and Ward hierarchical clustering. For each algorithm, we used the original co-activation matrix as well as a

standardized co-activation matrix to attempt to reduce potential artifacts introduced by the wide variability in base rate across the brain. In Figure 4.1, we display the cluster solutions across the four combination of strategies at $k = 100$ clusters. In all of the resulting solutions, most clusters encompassed well-known anatomical structures and functional regions. At low granularity, we identified whole-brain distributed functional networks such as the default, fronto-parietal and visual networks. At higher levels of granularity— despite no explicit spatial constraint— we typically observed spatially contiguous clusters that corresponded reasonably well to discrete anatomical structures, such as the amygdala, and cortical functional regions such as the visual cortex. However, k-means produced smoother, more spherical clusters than Ward hierarchical clustering. The smoothness of clusters was not affected by standardizing the co-activation matrix, resulting in qualitatively similar clusters within each algorithm.


Next, we examined if standardizing the co-activation matrix prior to clustering reduced the unwanted relationship between cluster size and activation rate. In Figure 4.2, we display the correlation between cluster size (number of voxels) and the mean base rate of activation of each voxel in the cluster. Across both algorithms, clusters with a higher rate of activation were much smaller in size when using the original co-activation distance matrix. Although this relationship was greatest at coarse scales, peaking at $k = 20$, it remained substantial even at a finer-grained levels of resolution. Standardizing the co-activation matrix greatly reduced this relationship— and even reversed it at very coarse scales. At finer-grained scales (i.e. 60-180 clusters), this relationship decreased substantially. For example, at $k=100$, the correlation between base rate and cluster size decreased from $r = -.7$ to r

= -.49 for k-means and $r = -.5$ to $r = -.33$ for Ward. Moreover, the range of cluster sizes was substantially compressed; for instance using k-means, cluster sizes ranged from 443 to 3918 voxels using the original co-activation at $k = 100$, but this range decreased to 948 to 3277 using standardized co-activation. These results suggest standardizing co-activation is an important pre-processing step to assure similarly sized clusters.


Ward - Original co-activation



Ward - Scaled co-activation



K-means - Original co-activation



K-means - Scaled co-activation




Figure 4.1. Clustering solutions for k-means and Ward clustering with original and scaled co-activation at $k = 100$ clusters. All clustering strategies resulted in spatially contiguous clusters that encompassed known functional and anatomical regions.




Figure 4.2. Relationship between cluster size (in voxels) and mean activation rate across clusters. Using the raw, original co-activation distance matrix, clusters with high activation rate were much smaller in size. This artifact was reduced across k-means and Ward clustering by standardizing the co-activation data prior to clustering.

Within-cluster cohesion. A fundamental way to measure the quality of a clustering solution is by the within-cluster cohesion of the resulting clusters. High quality solutions will maximize the similarity of the samples that compose each

cluster across. In other words, better clustering solutions result in brain clusters composed of voxels with more similar whole-brain co-activation. Here, we use the silhouette score to measure within-cluster cohesion across the various combinations of clustering strategies (Figure 4.3). Across all strategies, silhouette scores increased as the spatial granularity of the clustering increased, plateauing to some extent after 80-100 regions. Notably, there were no well-defined local maxima, suggesting choosing a spatial granularity is non-trivial and should be done in the context of other measures or qualitative, theory-driven judgments.




Figure 4.3. Within-cluster cohesion, measure by the silhouette coefficient, across various levels of granularity and clustering strategies in co-activation based clustering. K-means clustering exhibits greater silhouette scores, suggesting clusters in those solutions have greater co-activation homogeneity.

We observed substantial differences in within-cluster cohesion based on the clustering strategy employed. K-means clustering exhibited substantially greater within-cluster cohesion than Ward hierarchical clustering at all spatial scales, indicating k-means clusters have greater co-activation homogeneity. Across both algorithms, standardized co-activation matrices resulted in greater silhouette coefficients, suggesting that this important preprocessing step results in higher quality clustering. This is consistent with recommendations from the machine learning literature that normalization of the input data prior to clustering is often beneficial.

Reproducibility. We examined the spatial reliability of clustering by computing the similarity between pairs of bootstrapped solutions for each clustering strategy across spatial scales (Figure 4.4a). Across all clustering parameter choices, clustering solutions were much more reliable than chance, suggesting meta-analytic parcels can be interpreted with some degree of confidence. Moreover, our reliability was within the same range of those observed with raw fMRI data (Thirion et al., 2014). However, in contrast to Thirion et al., 2014, k-means exhibited greater reliability than Ward clustering. This finding is consistent with the qualitatively smooth nature of k-means clusters and suggests that k-means may be an appropriate algorithm if reliability is of particular importance.

We also observed that reliability peaked near 100 regions for k-means and 40 regions for Ward, as opposed to near 200 regions in Thirion et al., 2014. Although reliability remained high at 180 regions, this discrepancy may suggest that

summarized meta-analytic data is inherently lower in resolution than raw fMRI data at the subject level. In addition, across both algorithms, standardized co-activation resulted in less reliable clusters. A possible reason for this may be that standardizing introduces an additional data-dependent step to the clustering (because the mean of each cluster will be different for each re-sample), making the clustering solution more susceptible to over-fitting. However, this decrement was lower, and essentially negligible in Ward clustering compared to k-means.




Figure 4.4. Evaluation of co-activation based clustering strategies across spatial granularities. A) Reliability of clustering as measured by the adjusted Rand index (ARI). K-means clustering produced more reliable and consistent clustering solutions across bootstrapped samples. B) Performance in predicting psychological topics across studies in the database. All solutions performed similarly and predicted psychological topics above chance, and performed better at finer scaled of granularity.

Psychological topic prediction. An alternative way to evaluate the validity of the various clustering solutions is in their ability to predict the psychological concepts in the Neurosynth database. The key idea underlying this approach is the assumption that regions in the brain represent functionally

homogenous computational units. As such, activity in regions underlying specific psychological process should predict the presence of that psychological state, to some extent. For example, electrophysiological data has indicated that neurons in extrastriate area MT are tuned to the speed and direction of moving visual stimuli (Dubner & Zeki, 1971; Maunsell & Van Essen, 1983). Thus, activity in this area should predict with some degree of accuracy if a subject is observing motion. If a single cluster encompassed area MT, it follows that this cluster would be a strong predictor of a psychological topic in Neurosynth representing ‘motion’. However, if this cluster were poorly formed and only partially encompassed MT in addition to voxels that do not respond to motion, the predictive power of such a cluster would drop. Whole-brain clustering solutions that respect functional divisions in the brain should more accurately predict activity across a wide variety of psychological topic than atlases that do not respect these boundaries.

To evaluate our clustering solutions, we used regions as features in a naïve Bayes model to predict Neurosynth semantic topics that represent discrete psychological states and evaluated the performance using the receiver operating characteristic area under the curve (ROC-AUC). Critically, we evaluated performance using cross validation, and included the clustering step in the cross validation loop, generating clusters from training data not used in testing, to avoid an optimistic bias in performance. Summing across all topics for each parcellation map, Figure 4.4b displays classification performance as a function of spatial granularity for each parcellation strategy. Performance was better than chance across all algorithms and increased as a function of number of regions. This finding is consistent with previous evidence that functional specificity increases as regions

become smaller (Poldrack, 2006), suggesting larger regions average distinct functional activation patterns.

K-means clustering performed marginally better than Ward hierarchical clustering, suggesting k-means clustering better represents functional boundaries in the brain. However, this difference was very small— hierarchical clustering is not prohibitively low in performance. Moreover, standardizing the co-activation matrix also had a small effect on performance— in k-means clustering it resulted in poorer performance whereas it improved performance in Ward clustering. It bears repeating that the small differences shown here suggest that other factors, such as reliability or interpretability of the solutions, may be more important for choosing the appropriate co-activation parcellation strategy.

Comparison to clustering from other modalities. Having characterized the performance of meta-analytic co-activation clustering across various parcellation strategies, we sought to compare the performance of our functionally defined clusters to brain atlases from brain modalities. We compared our clustering solutions to two anatomical atlases: AAL and Harvard-Oxford, in addition to three well-validated rsfc-fMRI atlases generated with different algorithms: Yeo et al., (2011), Craddock et al., (2012), and Gordon et al., (2015). In order to ensure a fair, apples-to-apples comparison, we generated matching co-activation based parcellations for each atlas, using the exact same mask. In order to simplify comparisons, we only compare these brain atlases to co-activation based clustering with the k-means algorithm.




Figure 4.5. Psychological topic prediction performance across brain atlases from various modalities. We compared the performance of predicting psychological topics using co-activation based parcels as features to the performance using resting-state and anatomical atlases. Left) Co-activated based parcellation consistently outperformed a resting-state atlas generated using spectral clustering (Craddock, 2013). Right) Co-activation based clustering outperformed anatomical atlases (AAL & Harvard-Oxford) and two well-validated resting state atlases at different scales (Yeo, et al., 2011; Gordon, 2015).

Co-activation based clustering was better able to predict psychological topics compared to all other atlases we used for comparison (Figure 4.5). Across a wide range of spatial granularities, co-activation based clustering achieved greater classification performance (mean roc-auc: 0.549) compared to the Craddock (2013) atlas (mean roc-acu: 0.543). Notably, the difference between the two atlases in performance was similar across all level of granularity. Co-activation based clustering also outperformed both the Harvard-Oxford and AAL anatomical atlases, although the difference between Harvard-Oxford and co-activation based clustering was the smallest across all comparisons. Finally, co-activation based clustering also outperformed both the Yeo and Gordon resting state atlases. The Yeo atlas provided

a particularly fair comparison as this atlas was generated using the same algorithm (k-means) as the co-activation parcellation here. On the other hand, the Gordon atlas was a rigorous test as it was specifically tested for high resting state signal homogeneity, besting many competing atlases. The Gordon atlas was additionally relatively high in spatial resolution (334 regions), suggesting co-activation based parcellation outperforms resting state atlases even when clusters are relatively small.

Discussion

We assessed the performance of co-activation based parcellation across various spatial granularities using three distinct metrics: within-cluster homogeneity, bootstrapped reproducibility and ability to predict psychological topics across studies in Neurosynth. We used these metrics to compare the performance of various clustering strategies employed for co-activation based clustering, suggesting performance varies based on the strategy employed. Finally, we compared co-activation based clustering to atlases from other brain modalities on the basis of their ability to classify the presence of psychological topics across studies in Neurosynth, finding that co-activation based clustering outperforms them across spatial scales. Below we discuss in more depth the implications of our results.

Co-activation parcellation strategies and tradeoffs. We compared two clustering algorithms, k-means and Ward hierarchical clustering, and evaluated the impact of standardizing the co-activation matrix prior to clustering. However, to

properly understand the implications of these results, it is important to consider that there is no clear ‘winning’ strategy, as every choice introduces various tradeoffs. The appropriate clustering solution and granularity chosen will vary depending on the individual researcher’s goals. For example, standardizing the co-activation matrix resulted in more evenly sized clusters that were less impacted by differences in activation rate across the brain. Moreover, standardized co-activation resulted in greater within-cluster cohesion of the resulting clusters, suggesting this strategy results in cohesive and interpretable clusters. However, standardization resulted in less reproducible clusters, suggesting caution must be taken in interpreting the precise boundaries of regions. Moreover, standardized co-activation led to a slight cost in classification performance when using k-means clustering. As such, if the researcher’s goal is to choose ROIs that will best predict a given psychological state (e.g. predicting pain given brain activity), it may be prudent to use k-means clustering with the unaltered co-activation matrix.

Similarly, choosing the appropriate algorithm for clustering depends on the researcher’s priorities. K-means clustering exhibited smoother clusters and outperformed Ward hierarchical clustering with respect to within-cluster cohesion and reproducibility. As such, k-means clustering is a reasonable choice for many applications. However, hierarchical clustering may provide greater interpretability, as the relationship between clusters at various levels of granularity can be visualized using a dendrogram. As such, Ward clustering may be preferable when trying to understand the large-scale organization of a spatially large area of the

brain. As long as one does not make strong, deterministic arguments about the hierarchical organization of the resulting clusters, and simply uses them to guide the exploration of functional-anatomical brain organization, hierarchical clustering remains a viable strategy. However, if one is less concerned about interpretability, perhaps because one is interested in a smaller area of the brain that is easier to qualitatively understand, or the goal is to maximize the accuracy of a multivariate classifier to predict psychological states, k-means may be a better choice.

Comparison to atlases from other brain modalities. Meta-analytic co-activation parcellation consistently outperformed atlases from other modalities in the ability to predict psychological topics across studies in Neurosynth. Taken at face value, this finding suggests that brain atlases derived from meta-analytic fMRI data best capture the functional organization of brain activity (Eickhoff et al., 2015). This was relatively unsurprising when compared to anatomical atlases, as it is generally well accepted that functional boundaries often supersede gross anatomical boundaries, such as gyri and sulci. For example, in Chapter 3, we found that voxels from primary motor and sensory cortices grouped together into a single cluster despite being in anatomically distinct regions.

In contrast, rsfc-fMRI derived atlases are unconstrained by anatomy and are hypothesized to reflect functional organization of the brain. However, these more functional atlases were similarly outperformed by co-activation based parcellations, despite the well documented similarity between resting state and meta-analytic co-activation networks (S. M. Smith et al., 2009). One reason why rsfc-fMRI atlases

performed more poorly than meta-analytic parcellation is that— by definition— resting state signal is recorded when the subjects are resting without task in the scanner. As such, the organization of the brain when it is unchallenged, or perhaps in the specific state of mind wandering, may be different than the activity pattern observed across a wide range of psychological states.

Consistent with this hypothesis, recent efforts have demonstrated that functional connectivity is constrained by anatomical connections (Goñi et al., 2014) and dynamically changes when measured during task performance (Cole et al., 2014; Mattar et al., 2015). As such, connectivity signal at rest may more strongly reflect static anatomical differences than meta-analytic co-activation, suggesting that the parcels derived from such data less directly reflect the functional, task-dependent organization of the brain. Computational modeling of functional connectivity as a function of task demands is a promising future direction to better understand the nature of this signal (Mattar et al., 2015). However, this approach is inherently limited by the functional diversity (or lack thereof) that can be measured in a single population. In contrast, meta-analytic based parcellation surveys a broad range of psychological states and as such may serve as the best estimate of functional organization presently available.

Pragmatic advantages of meta-analytic co-activation atlases. The present results suggest that meta-analytically derived parcellations may provide several pragmatic advantages. Often, fMRI researchers are interested in defining a priori ROI's to extract brain signal in an effort to reduce the number of statistical

comparisons made in a study. The present results suggest that meta-analytic co-activation parcellation may provide a better source of a priori ROIs than atlases from other modalities that less directly measure functional brain signal. As co-activation based parcellations are derived from a very diverse set of behavioral manipulations, the resulting parcels are likely to be useful when researchers do not have strong a priori predictions as to the predicted pattern of brain activity. If researchers have strong a priori beliefs— for example, if researchers are inducing pain and wish to measure pain related signal— they may be better served by deriving topics specific ROIs using reverse inference maps in Neurosynth or from a targeted manually constructed meta-analysis. Conversely, the rigorously tested resting-state atlases, such as the recently released Gordon (2015) atlas, are likely to provide better a priori ROIs for rsfc-fMRI studies, in which the interest is to extract homogenous signal for each ROI at rest.

Limitations and future challenges. There are several limitations to our approach that may limit the generalizability of our results. In particular, the Neurosynth database is composed of summarized statistics (i.e. activation peaks) that are reported in tables of fMRI studies. Moreover, although unlikely, it is possible that the automated heuristic conversion of Talairach to MNI coordinates may contribute systematic bias to the database. As such, a possible reason why meta-analytic co-activation parcellation outperforms atlases from other modalities may be due to some degree of over fit to idiosyncratic peculiarities of Neurosynth. To truly demonstrate meta-analytic atlases more accurately reflect functional

boundaries, it will be interesting to compare the performance of various atlases in their ability to predict psychological states in raw fMRI data across individual subjects. If co-activation based parcels are better able to predict the intensity across specific psychological states across people, such as pain or working-memory load, it would provide strong evidence that meta-analytic atlases better reflect the underlying functional-anatomical organization of the brain.

It is also important to recognize that the definition of a computational unit in the brain will vary depending on the level of analysis and will greatly benefit from converging evidence across modalities. For example, although co-activation based parcellation grouped together portions of the primary motor and somatosensory cortices into a single cluster, it is well known from decades of electrophysiological studies that these areas differ in the specific computational role they play. Primary somatosensory cortex receives and process afferent information while primary motor cortex, through efferent cortico-spinal projects, actuates movements. At this level of analysis, it is self-evident these two regions play distinct computational roles. However, it is also informative that these regions strongly co-activate and thus are grouped into a single cluster at coarser scales, effectively forming unit at a higher level of analysis. Precise computational modeling of the relationship between different modalities across levels of abstraction is necessary to understand dynamic information processing in the brain.

Conclusion. We assessed various strategies for meta-analytic co-activation clustering using metrics that measure distinct qualities that researchers may want

to maximize. Moreover, we have demonstrated that meta-analytic parcellations may more accurately reflect the brain's underlying functional-anatomical organization than brain atlases from other modalities and suggest future directions to objectively evaluate the plethora of existing organizational schemes of the brain.

Summary. Across three studies in this dissertation, I have used large-scale meta-analysis to study the functional-anatomical organization of the brain. In two of these studies, I applied relatively unbiased data-driven methods to the largest existing database of neuroimaging studies to comprehensively map psychological states to discrete anatomical units in frontal cortex. The results of these studies extend upon the existing literature by providing a reference of the most robust associations between psychological semantic meta-data and frontal cortex anatomy. I have outlined various findings that suggest novel brain-cognition associations that can be tested more focused follow up studies. However, I have also tried to argue the importance of large-scale neuro-informatics in the cognitive neuroscience ecosystem by quantitatively demonstrating that such approaches can lead to unique insights and potentially result in better formal representations of functional-anatomical organization in the human brain.

BIBLIOGRAPHY

- Adolphs, R., Tranel, D., Damasio, H., & Damasio, A. R. (1995). Fear and the human amygdala. *Journal of Neuroscience*, 15(9), 5879–5891.
- Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, 14(10), 1338–1344.
<http://doi.org/10.1038/nn.2921>
- Amunts, K., & Zilles, K. (2015). Architectonic Mapping of the Human Brain beyond Brodmann. *Neuron*, 88(6), 1086–1107.
<http://doi.org/10.1016/j.neuron.2015.12.001>
- Anderson, M. L., Kinnison, J., & Pessoa, L. (2013). Describing functional diversity of brain regions and brain networks. *NeuroImage*, 73(C), 50–58.
<http://doi.org/10.1016/j.neuroimage.2013.01.071>
- Andrews Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., & Buckner, R. L. (2010). Functional-Anatomic Fractionation of the Brain's Default Network. *Neuron*, 65(4), 550–562. <http://doi.org/10.1016/j.neuron.2010.02.005>
- Andrews Hanna, J. R., Saxe, R., & Yarkoni, T. (2014a). Contributions of episodic retrieval and mentalizing to autobiographical thought: Evidence from functional neuroimaging, resting-state connectivity, and fMRI meta-analyses. *NeuroImage*, 91(C), 324–335. <http://doi.org/10.1016/j.neuroimage.2014.01.032>
- Andrews Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014b). The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, 1316(1), 29–52.
<http://doi.org/10.1111/nyas.12360>
- Andrews-Hanna, J. R. (2012). The Brain's Default Network and Its Adaptive Role in Internal Mentation. *The Neuroscientist*, 18(3), 251–270.
<http://doi.org/10.1177/1073858411403316>
- Androutsopoulos, I., Koutsias, J., & Chandrinou, K. V. (2000). An evaluation of naive bayesian anti-spam filtering. *Proceedings of the Workshop on Machine Learning in the New Information Age*, 9–17.
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8(4), 170–177.
<http://doi.org/10.1016/j.tics.2004.02.010>

- Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Reviews Neuroscience*, 4(10), 829–839. <http://doi.org/10.1038/nrn1201>
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5), 193–200. <http://doi.org/10.1016/j.tics.2008.02.004>
- Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews Neuroscience*, 10(9), 659–669. <http://doi.org/10.1038/nrn2667>
- Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, 45(13), 2883–2901. <http://doi.org/10.1016/j.neuropsychologia.2007.06.015>
- Banich, M. T. (2009). Executive Function The Search for an Integrated Account. *Current Directions in Psychological Science*, 18(2), 89–94. <http://doi.org/10.1111/j.1467-8721.2009.01615.x>
- Baumgartner, T., Götte, L., Gügler, R., & Fehr, E. (2012). The mentalizing network orchestrates the impact of parochial altruism on social norm enforcement. *Human Brain Mapping*, 33(6), 1452–1469. <http://doi.org/10.1002/hbm.21298>
- Beckmann, M., Johansen-Berg, H., & Rushworth, M. F. S. (2009). Connectivity-Based Parcellation of Human Cingulate Cortex and Its Relation to Functional Specialization. *Journal of Neuroscience*, 29(4), 1175–1190. <http://doi.org/10.1523/JNEUROSCI.3328-08.2009>
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15(11), 527–536. <http://doi.org/10.1016/j.tics.2011.10.001>
- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex (New York, N.Y. : 1991)*, 19(12), 2767–2796. <http://doi.org/10.1093/cercor/bhp055>
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3, 993–1022.
- Bludau, S., Eickhoff, S. B., Mohlberg, H., Caspers, S., Laird, A. R., Fox, P. T., et al. (2014). Cytoarchitecture, probability maps and functions of the human frontal pole. *NeuroImage*, 93, 260–275.
- Bohland, J. W., Bokil, H., Allen, C. B., & Mitra, P. P. (2009). The Brain Atlas Concordance Problem: Quantitative Comparison of Anatomical Parcellations. *PLoS ONE*, 4(9), e7200. <http://doi.org/10.1371/journal.pone.0007200>
- Botstein, D., Cherry, J. M., Ashburner, M., & Ball, C. A. (2000). Gene Ontology: tool

for the unification of biology. *Nat Genet.*

- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12(5), 201–208.
<http://doi.org/10.1016/j.tics.2008.02.009>
- Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, 402(6758), 179–181. <http://doi.org/10.1038/46035>
- Brass, M., Derrfuss, J., Forstmann, B., & Cramon, von, D. Y. (2005). The role of the inferior frontal junction area in cognitive control. *Trends in Cognitive Sciences*, 9(7), 314–316. <http://doi.org/10.1016/j.tics.2005.05.001>
- Brown, J. W., & Braver, T. S. (2005). Learned Predictions of Error Likelihood in the Anterior Cingulate Cortex. *Science*, 307(5712), 1118–1121.
<http://doi.org/10.1126/science.1105783>
- Buckner, R. L., Andrews Hanna, J. R., & Schacter, D. L. (2008). The Brain's Default Network. *Annals of the New York Academy of Sciences*, 1124(1), 1–38.
<http://doi.org/10.1196/annals.1440.011>
- Burgess, P. W., Dumontheil, I., & Gilbert, S. J. (2007). The gateway hypothesis of rostral prefrontal cortex (area 10) function. *Trends in Cognitive Sciences*, 11(7), 290–298. <http://doi.org/10.1016/j.tics.2007.05.004>
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365–376.
<http://doi.org/10.1038/nrn3475>
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, 280(5364), 747–749.
<http://doi.org/10.1126/science.280.5364.747>
- Carter, R. M., & Huettel, S. A. (2013). A nexus model of the temporal–parietal junction. *Trends in Cognitive Sciences*, 17(7), 328–336.
<http://doi.org/10.1016/j.tics.2013.05.007>
- Cavanagh, J. F., & Shackman, A. J. (2015). Frontal midline theta reflects anxiety and cognitive control: Meta-analytic evidence. *Journal of Physiology-Paris*, 109(1-3), 3–15. <http://doi.org/10.1016/j.jphysparis.2014.04.003>
- Chang, L. J., Yarkoni, T., Khaw, M. W., & Sanfey, A. G. (2013). Decoding the Role of the Insula in Human Cognition: Functional Parcellation and Large-Scale Reverse Inference. *Cerebral Cortex*, 23(3), 739–749.
<http://doi.org/10.1093/cercor/bhs065>

- Chatham, C. H., Claus, E. D., Kim, A., Curran, T., Banich, M. T., & Munakata, Y. (2012). Cognitive Control Reflects Context Monitoring, Not Motoric Stopping, in Response Inhibition. *PLoS ONE*, 7(2), e31546. <http://doi.org/10.1371/journal.pone.0031546>
- Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S., & Petersen, S. E. (2014). Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron*, 83(1), 238–251. <http://doi.org/10.1016/j.neuron.2014.05.014>
- Collins, D. L., Holmes, C. J., Peters, T. M., & Evans, A. C. (1995). Automatic 3-D model-based neuroanatomical segmentation. *Human Brain Mapping*, 3(3), 190–208. <http://doi.org/10.1002/hbm.460030304>
- Craddock, R. C., James, G. A., Holtzheimer, P. E., Hu, X. P., & Mayberg, H. S. (2012). A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*, 33(8), 1914–1928. <http://doi.org/10.1002/hbm.21333>
- Critchley, H. D., Mathias, C. J., Josephs, O., O'Doherty, J., Zanini, S., Dewar, B. K., et al. (2003). Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain*, 126(10), 2139–2152. <http://doi.org/10.1093/brain/awg216>
- Curtis, C. E., & Lee, D. (2010). Beyond working memory: the role of persistent activity in decision making. *Trends in Cognitive Sciences*, 14(5), 216–222. <http://doi.org/10.1016/j.tics.2010.03.006>
- De Baene, W., Albers, A. M., & Brass, M. (2012). The what and how components of cognitive control. *NeuroImage*, 63(1), 203–211. <http://doi.org/10.1016/j.neuroimage.2012.06.050>
- Denny, B. T., Kober, H., Wager, T. D., & Ochsner, K. N. (2012). A meta-analysis of functional neuroimaging studies of self- and other judgments reveals a spatial gradient for mentalizing in medial prefrontal cortex. *Journal of Cognitive Neuroscience*, 24(8), 1742–1752. http://doi.org/10.1162/jocn_a_00233
- Depue, B. E., Orr, J. M., Smolker, H. R., Naaz, F., & Banich, M. T. (2016). The Organization of Right Prefrontal Networks Reveals Common Mechanisms of Inhibitory Regulation Across Cognitive, Emotional, and Motor Processes. *Cerebral Cortex (New York, N.Y. : 1991)*, 26(4), 1634–1646. <http://doi.org/10.1093/cercor/bhu324>
- Derrfuss, J., Brass, M., Cramon, von, D. Y., Lohmann, G., & Amunts, K. (2009). Neural activations at the junction of the inferior frontal sulcus and the inferior precentral sulcus: Interindividual variability, reliability, and association with sulcal morphology. *Human Brain Mapping*, 30(1), 299–311. <http://doi.org/10.1002/hbm.20501>

- Derrfuss, J., Brass, M., Neumann, J., & Cramon, von, D. Y. (2005a). Involvement of the inferior frontal junction in cognitive control: Meta-analyses of switching and Stroop studies. *Human Brain Mapping*, 25(1), 22–34. <http://doi.org/10.1002/hbm.20127>
- Derrfuss, J., Brass, M., Neumann, J., & Cramon, von, D. Y. (2005b). Involvement of the inferior frontal junction in cognitive control: Meta-analyses of switching and Stroop studies. *Human Brain Mapping*, 25(1), 22–34. <http://doi.org/10.1002/hbm.20127>
- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., et al. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3), 968–980. <http://doi.org/10.1016/j.neuroimage.2006.01.021>
- Dubner, R., & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, 35(2), 528–532. [http://doi.org/10.1016/0006-8993\(71\)90494-X](http://doi.org/10.1016/0006-8993(71)90494-X)
- Eickhoff, S. B., Paus, T., Caspers, S., Grosbras, M.-H., Evans, A. C., Zilles, K., & Amunts, K. (2007). Assignment of functional activations to probabilistic cytoarchitectonic areas revisited. *NeuroImage*, 36(3), 511–521. <http://doi.org/10.1016/j.neuroimage.2007.03.060>
- Eickhoff, S. B., Thirion, B., Varoquaux, G., & Bzdok, D. (2015). Connectivity-based parcellation: Critique and implications. *Human Brain Mapping*, 36(12), 4771–4792. <http://doi.org/10.1002/hbm.22933>
- Etkin, A., Egner, T., & Kalisch, R. (2011). Emotional processing in anterior cingulate and medial prefrontal cortex. *Trends in Cognitive Sciences*, 15(2), 85–93. <http://doi.org/10.1016/j.tics.2010.11.004>
- Flinker, A., Korzeniewska, A., Shestyuk, A. Y., Franaszczuk, P. J., Dronkers, N. F., Knight, R. T., & Crone, N. E. (2015). Redefining the role of Broca's area in speech. *Proceedings of the National Academy of Sciences of the United States of America*, 112(9), 2871–2875. <http://doi.org/10.1073/pnas.1414491112>
- Frank, D. W., Dewitt, M., & Hudgens-Haney, M. (2014). Emotion regulation: quantitative meta-analysis of functional activation and deactivation. *Neuroscience &* <http://doi.org/10.1016/j.neubiorev.2014.06.010>
- Friston, K. (2002). Beyond phrenology: what can neuroimaging tell us about distributed circuitry? *Annual Review of Neuroscience*, 25, 221–250. <http://doi.org/10.1146/annurev.neuro.25.112701.142846>
- Gilbert, S. J., Spengler, S., Simons, J. S., Steele, J. D., Lawrie, S. M., Frith, C. D., & Burgess, P. W. (2006). Functional specialization within rostral prefrontal cortex

- (area 10): a meta-analysis. *Journal of Cognitive Neuroscience*, 18(6), 932–948. <http://doi.org/10.1162/jocn.2006.18.6.932>
- Goldman-Rakic, P. S. (1996). The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 351(1346), 1445–1453. <http://doi.org/10.1098/rstb.1996.0129>
- Goñi, J., van den Heuvel, M. P., Avena-Koenigsberger, A., Velez de Mendizabal, N., Betzel, R. F., Griffa, A., et al. (2014). Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences of the United States of America*, 111(2), 833–838. <http://doi.org/10.1073/pnas.1315529111>
- Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., & Petersen, S. E. (2015). Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, 26(1), 288–303. <http://doi.org/10.1093/cercor/bhu239>
- Gorgolewski, K. J., Varoquaux, G., Rivera, G., Schwarz, Y., Ghosh, S. S., Maumet, C., et al. (2015). NeuroVault.org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain. *Frontiers in Neuroinformatics*, 9(17), 8. <http://doi.org/10.3389/fninf.2015.00008>
- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System. *Science*, 324(5927), 646–648. <http://doi.org/10.1126/science.1168450>
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539), 2425–2430. <http://doi.org/10.1126/science.1063736>
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G. H., & Cohen, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience*, 7(5), 497–498. <http://doi.org/10.1038/nn1238>
- Jeni, L. A., Cohn, J. F., & la Torre, De, F. (2013). Facing Imbalanced Data Recommendations for the Use of Performance Metrics. *International Conference on Affective Computing and Intelligent Interaction and Workshops : [Proceedings]. ACII (Conference), 2013*, 245–251. <http://doi.org/10.1109/ACII.2013.47>
- Johansen-Berg, H., Behrens, T. E. J., Robson, M. D., Drobnjak, I., Rushworth, M. F. S., Brady, J. M., et al. (2004). Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proceedings of the*

- National Academy of Sciences of the United States of America*, 101(36), 13335–13340. <http://doi.org/10.1073/pnas.0403743101>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17(11), 4302–4311.
- Kennerley, S. W., Sakai, K., & Rushworth, M. F. S. (2004). Organization of action sequences and the role of the pre-SMA. *Journal of Neurophysiology*, 91(2), 978–993. <http://doi.org/10.1152/jn.00651.2003>
- Kim, J.-H., Lee, J.-M., Jo, H. J., Kim, S. H., Lee, J. H., Kim, S. T., et al. (2010). Defining functional SMA and pre-SMA subregions in human MFC using resting state fMRI: Functional connectivity-based parcellation method. *NeuroImage*, 49(3), 2375–2386. <http://doi.org/10.1016/j.neuroimage.2009.10.016>
- Kober, H., Barrett, L. F., Joseph, J., Bliss-Moreau, E., Lindquist, K., & Wager, T. D. (2008). Functional grouping and cortical–subcortical interactions in emotion: A meta-analysis of neuroimaging studies. *NeuroImage*, 42(2), 998–1031. <http://doi.org/10.1016/j.neuroimage.2008.03.059>
- Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J. Z., & Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature*, 498(7454), 363–366. <http://doi.org/10.1038/nature12176>
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., et al. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, 89(12), 5675–5679. <http://doi.org/10.1073/pnas.89.12.5675>
- Laird, A. R., Lancaster, J. L., & Fox, P. T. (2005). BrainMap: the social evolution of a human brain mapping database. *Neuroinformatics*, 3(1), 65–78. <http://doi.org/10.1385/NI:3:1:065>
- Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An Automatic Valuation System in the Human Brain: Evidence from Functional Neuroimaging. *Neuron*, 64(3), 431–439. <http://doi.org/10.1016/j.neuron.2009.09.040>
- Leek, E. C., & Johnston, S. J. (2009). Functional specialization in the supplementary motor complex. *Nature Reviews Neuroscience*, 10(1), 78–author reply 78. <http://doi.org/10.1038/nrn2478-c1>
- Lieberman, M. D., & Eisenberger, N. I. (2015). The dorsal anterior cingulate cortex is selective for pain: Results from large-scale reverse inference. *Proceedings of the National Academy of Sciences of the United States of America*, 112(49),

15250–15255. <http://doi.org/10.1073/pnas.1515083112>

- Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E., & Barrett, L. F. (2012). The brain basis of emotion: A meta-analytic review. *Behavioral and Brain Sciences*, 35(03), 121–143. <http://doi.org/10.1017/S0140525X11000446>
- Mackey, S., & Petrides, M. (2014). Architecture and morphology of the human ventromedial prefrontal cortex. *European Journal of Neuroscience*, 40(5), 2777–2796. <http://doi.org/10.1111/ejn.12654>
- Mackey, W. E., Devinsky, O., Doyle, W. K., Meager, M. R., & Curtis, C. E. (2016). Human Dorsolateral Prefrontal Cortex Is Not Necessary for Spatial Working Memory. *J Neurosci*, 36(10), 2847–2856. <http://doi.org/10.1523/JNEUROSCI.3618-15.2016>
- Martinez-Trujillo, J. C., Tsotsos, J. K., Simine, E., Pomplun, M., Wildes, R., Treue, S., et al. (2005). Selectivity for speed gradients in human area MT/V5. *Neuroreport*, 16(5), 435–438.
- Mattar, M. G., Cole, M. W., Thompson-Schill, S. L., & Bassett, D. S. (2015). A Functional Cartography of Cognitive Systems. *PLoS Computational Biology*, 11(12), e1004533. <http://doi.org/10.1371/journal.pcbi.1004533>
- Maunsell, J. H., & Van Essen, D. C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *Journal of Neurophysiology*, 49(5), 1148–1167. <http://doi.org/10.1152/jn.00847.2015>
- Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., et al. (2001). A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 356(1412), 1293–1322. <http://doi.org/10.1098/rstb.2001.0915>
- Milad, M. R., Quirk, G. J., Pitman, R. K., Orr, S. P., Fischl, B., & Rauch, S. L. (2007). A Role for the Human Dorsal Anterior Cingulate Cortex in Fear Expression. *Biological Psychiatry*, 62(10), 1191–1194. <http://doi.org/10.1016/j.biopsych.2007.04.032>
- Milham, M. P., Banich, M. T., Webb, A., Barad, V., Cohen, N. J., Wszalek, T., & Kramer, A. F. (2001). The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. *Cognitive Brain Research*, 12(3), 467–473. [http://doi.org/10.1016/S0926-6410\(01\)00076-3](http://doi.org/10.1016/S0926-6410(01)00076-3)
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167–202. <http://doi.org/10.1146/annurev.neuro.24.1.167>

- Mitchell, J. P., Banaji, M. R., & Macrae, C. N. (2005). The link between social cognition and self-referential thought in the medial prefrontal cortex. *Journal of Cognitive Neuroscience*, 17(8), 1306–1315.
<http://doi.org/10.1162/0898929055002418>
- Muhle-Karbe, P. S., Derrfuss, J., Lynn, M. T., Neubert, F. X., Fox, P. T., Brass, M., & Eickhoff, S. B. (2015a). Co-Activation-Based Parcellation of the Lateral Prefrontal Cortex Delineates the Inferior Frontal Junction Area. *Cerebral Cortex*, 1–17. <http://doi.org/10.1093/cercor/bhv073>
- Muhle-Karbe, P. S., Derrfuss, J., Lynn, M. T., Neubert, F. X., Fox, P. T., Brass, M., & Eickhoff, S. B. (2015b). Co-Activation-Based Parcellation of the Lateral Prefrontal Cortex Delineates the Inferior Frontal Junction Area. *Cerebral Cortex*, bhv073–17. <http://doi.org/10.1093/cercor/bhv073>
- Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences*, 15(10), 453–459.
- Müllner, D. (2013). fastcluster: Fast Hierarchical, Agglomerative Clustering Routines for Rand Python. *Journal of Statistical Software*, 53(9).
<http://doi.org/10.18637/jss.v053.i09>
- Nee, D. E., & Brown, J. W. (2012). Rostral–caudal gradients of abstraction revealed by multi-variate pattern analysis of working memory. *NeuroImage*, 63(3), 1285–1294. <http://doi.org/10.1016/j.neuroimage.2012.08.034>
- Nee, D. E., Brown, J. W., Askren, M. K., Berman, M. G., Demiralp, E., Krawitz, A., & Jonides, J. (2013). A Meta-analysis of Executive Components of Working Memory. *Cerebral Cortex*, 23(2), 264–282. <http://doi.org/10.1093/cercor/bhs007>
- Nelson, S. M., Dosenbach, N. U. F., Cohen, A. L., Wheeler, M. E., Schlaggar, B. L., & Petersen, S. E. (2010a). Role of the anterior insula in task-level control and focal attention. *Brain Structure and Function*, 214(5-6), 669–680.
<http://doi.org/10.1007/s00429-010-0260-2>
- Nelson, S. M., Dosenbach, N. U. F., Cohen, A. L., Wheeler, M. E., Schlaggar, B. L., & Petersen, S. E. (2010b). Role of the anterior insula in task-level control and focal attention. *Brain Structure and Function*, 214(5-6), 669–680.
<http://doi.org/10.1007/s00429-010-0260-2>
- Nelson, S. M., Dosenbach, N. U. F., Cohen, A. L., Wheeler, M. E., Schlaggar, B. L., & Petersen, S. E. (2010c). Role of the anterior insula in task-level control and focal attention. *Brain Structure and Function*, 214(5-6), 669–680.
<http://doi.org/10.1007/s00429-010-0260-2>
- Neubert, F.-X., Mars, R. B., Thomas, A. G., Sallet, J., & Rushworth, M. F. S. (2014).

- Comparison of Human Ventral Frontal Cortex Areas for Cognitive Control and Language with Areas in Monkey Frontal Cortex. *Neuron*, 81(3), 700–713. <http://doi.org/10.1016/j.neuron.2013.11.012>
- Northoff, G., Heinzl, A., de Greck, M., Bermpohl, F., Dobrowolny, H., & Panksepp, J. (2006). Self-referential processing in our brain--a meta-analysis of imaging studies on the self. *NeuroImage*, 31(1), 440–457. <http://doi.org/10.1016/j.neuroimage.2005.12.002>
- Opialla, S., Lutz, J., Scherpiet, S., Hittmeyer, A., Jäncke, L., Rufer, M., et al. (2015). Neural circuits of emotion regulation: a comparison of mindfulness-based and cognitive reappraisal strategies. *European Archives of Psychiatry and Clinical Neuroscience*, 265(1), 45–55. <http://doi.org/10.1007/s00406-014-0510-z>
- Orr, J. M., Smolker, H. R., & Banich, M. T. (2015). Organization of the Human Frontal Pole Revealed by Large-Scale DTI-Based Connectivity: Implications for Control of Behavior. *PLoS ONE*, 10(5), e0124797. <http://doi.org/10.1371/journal.pone.0124797>
- Palomero-Gallagher, N., Eickhoff, S. B., Hoffstaedter, F., Schleicher, A., Mohlberg, H., Vogt, B. A., et al. (2015). Functional organization of human subgenual cortical areas: Relationship between architectonical segregation and connectional heterogeneity. *NeuroImage*, 115(C), 177–190. <http://doi.org/10.1016/j.neuroimage.2015.04.053>
- Palomero-Gallagher, N., Zilles, K., Schleicher, A., & Vogt, B. A. (2013). Cyto- and receptor architecture of area 32 in human and macaque brains. *The Journal of Comparative Neurology*, 521(14), 3272–3286. <http://doi.org/10.1002/cne.23346>
- Paus, T. (1996). Location and function of the human frontal eye-field: A selective review. *Neuropsychologia*, 34(6), 475–483. [http://doi.org/10.1016/0028-3932\(95\)00134-4](http://doi.org/10.1016/0028-3932(95)00134-4)
- Paus, T., Tomaiuolo, F., Otaky, N., MacDonald, D., Petrides, M., Atlas, J., et al. (1996). Human cingulate and paracingulate sulci: pattern, variability, asymmetry, and probabilistic map. *Cerebral Cortex*, 6(2), 207–214. <http://doi.org/10.1093/cercor/6.2.207>
- Pedregosa, F., Varoquaux, G., Gamfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(Oct), 2825–2830.
- Petrides, M. (2005). Lateral prefrontal cortex: architectonic and functional organization. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1456), 781–795. <http://doi.org/10.1098/rstb.2005.1631>
- Picard, N., & Strick, P. L. (1996). Motor Areas of the Medial Wall: A Review of Their

- Location and Functional Activation. *Cerebral Cortex*, 6(3), 342–353.
<http://doi.org/10.1093/cercor/6.3.342>
- Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Sciences*, 10(2), 59–63.
<http://doi.org/10.1016/j.tics.2005.12.004>
- Poldrack, R. A. (2007). Region of interest analysis for fMRI. *Social Cognitive and Affective Neuroscience*, 2(1), 67–70. <http://doi.org/10.1093/scan/nsm006>
- Poldrack, R. A., & Yarkoni, T. (2016). From Brain Maps to Cognitive Ontologies: Informatics and the Search for Mental Structure. *Annual Review of Psychology*, 67(1), 587–612. <http://doi.org/10.1146/annurev-psych-122414-033729>
- Poldrack, R. A., Kittur, A., Kalar, D., Miller, E., Seppa, C., Gil, Y., et al. (2011). The cognitive atlas: toward a knowledge foundation for cognitive neuroscience. *Frontiers in Neuroinformatics*, 5, 17. <http://doi.org/10.3389/fninf.2011.00017>
- Poldrack, R. A., Mumford, J. A., Schonberg, T., Kalar, D., Barman, B., & Yarkoni, T. (2012a). Discovering Relations Between Mind, Brain, and Mental Disorders Using Topic Mapping. *PLoS Computational Biology*, 8(10), e1002707. <http://doi.org/10.1371/journal.pcbi.1002707.s002>
- Poldrack, R. A., Mumford, J. A., Schonberg, T., Kalar, D., Barman, B., & Yarkoni, T. (2012b). Discovering relations between mind, brain, and mental disorders using topic mapping. *PLoS Computational Biology*, 8(10), e1002707. <http://doi.org/10.1371/journal.pcbi.1002707>
- Postle, B. R. (2016). Working memory functions of the prefrontal cortex. In M. Watanabe (Ed.), (Vol. Prefrontal Cortex as an Executive, Emotional and Social Brain., pp. 1–14). Springer.
- Power, J. D., & Petersen, S. E. (2013). Control-related systems in the human brain. *Current Opinion in Neurobiology*, 23(2), 223–228. <http://doi.org/10.1016/j.conb.2012.12.009>
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., et al. (2011). Functional Network Organization of the Human Brain. *Neuron*, 72(4), 665–678. <http://doi.org/10.1016/j.neuron.2011.09.006>
- Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *J Neurosci*, 32(38), 12990–12998. <http://doi.org/10.1523/JNEUROSCI.1892-12.2012>
- Robinson, J. L., Laird, A. R., Glahn, D. C., Lovallo, W. R., & Fox, P. T. (2010). Metaanalytic connectivity modeling: Delineating the functional connectivity of the human amygdala. *Human Brain Mapping*, 31(2), 173–184.

<http://doi.org/10.1002/hbm.20854>

- Roland, P. E., Larsen, B., Lassen, N. A., & Skinhøj, E. (1980). Supplementary motor area and other cortical areas in organization of voluntary movements in man. *Journal of Neurophysiology*, 43(1), 118–136.
- Rolls, E. T., O'Doherty, J., Kringelbach, M. L., Francis, S., Bowtell, R., & McGlone, F. (2003). Representations of Pleasant and Painful Touch in the Human Orbitofrontal and Cingulate Cortices. *Cerebral Cortex*, 13(3), 308–317. <http://doi.org/10.1093/cercor/13.3.308>
- Roy, M., Shohamy, D., & Wager, T. D. (2012). Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends in Cognitive Sciences*, 16(3), 147–156. <http://doi.org/10.1016/j.tics.2012.01.005>
- Salimi-Khorshidi, G., Smith, S. M., Keltner, J. R., Wager, T. D., & Nichols, T. E. (2009). Meta-analysis of neuroimaging data: A comparison of image-based and coordinate-based pooling of studies. *NeuroImage*, 45(3), 810–823. <http://doi.org/10.1016/j.neuroimage.2008.12.039>
- Sallet, J., Mars, R. B., Noonan, M. P., Neubert, F. X., Jbabdi, S., O'Reilly, J. X., et al. (2013). The Organization of Dorsal Frontal Cortex in Humans and Macaques. *Journal of Neuroscience*, 33(30), 12255–12274. <http://doi.org/10.1523/JNEUROSCI.5108-12.2013>
- Shackman, A. J., Fox, A. S., & Seminowicz, D. A. (2015). The cognitive-emotional brain: Opportunities and challenges for understanding neuropsychiatric disorders. *Behavioral and Brain Sciences*, 38, e86. <http://doi.org/10.1017/S0140525X14001010>
- Shackman, A. J., Salomons, T. V., Slagter, H. A., Fox, A. S., Winter, J. J., & Davidson, R. J. (2011). The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nature Reviews Neuroscience*, 12(3), 154–167. <http://doi.org/10.1038/nrn2994>
- Sharp, D. J., Bonnelle, V., De Boissezon, X., Beckmann, C. F., James, S. G., Patel, M. C., & Mehta, M. A. (2010). Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proceedings of the National Academy of Sciences*, 107(13), 6106–6111. <http://doi.org/10.1073/pnas.1000175107>
- Shen, X., Tokoglu, F., Papademetris, X., & Constable, R. T. (2013). Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *NeuroImage*, 82(C), 403–415. <http://doi.org/10.1016/j.neuroimage.2013.05.081>
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*,

- 79(2), 217–240. <http://doi.org/10.1016/j.neuron.2013.07.007>
- Shidara, M., & Richmond, B. J. (2002). Anterior Cingulate: Single Neuronal Signals Related to Degree of Reward Expectancy. *Science*, 296(5573), 1709–1711. <http://doi.org/10.1126/science.1069504>
- Sikes, R. W., Vogt, L. J., & Vogt, B. A. (2008). Distribution and properties of visceral nociceptive neurons in rabbit cingulate cortex. *Pain*, 135(1), 160–174. <http://doi.org/10.1016/j.pain.2007.09.024>
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences of the United States of America*, 106(31), 13040–13045. <http://doi.org/10.1073/pnas.0905267106>
- Snyder, H. R., Banich, M. T., & Munakata, Y. (2011). Choosing Our Words: Retrieval and Selection Processes Recruit Shared Neural Substrates in Left Ventrolateral Prefrontal Cortex. *Dx.Doi.org*, 23(11), 3470–3482. http://doi.org/10.1162/jocn_a_00023
- Spreng, R. N., & Grady, C. L. (2010). Patterns of brain activity supporting autobiographical memory, prospection, and theory of mind, and their relationship to the default mode network. *Journal of Cognitive Neuroscience*, 22(6), 1112–1123. <http://doi.org/10.1162/jocn.2009.21282>
- Spreng, R. N., Sepulcre, J., Turner, G. R., Stevens, W. D., & Schacter, D. L. (2013). Intrinsic architecture underlying the relations among the default, dorsal attention, and frontoparietal control networks of the human brain. *Journal of Cognitive Neuroscience*, 25(1), 74–86. http://doi.org/10.1162/jocn_a_00281
- Stalnaker, T. A., Cooch, N. K., & Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. *Nature Neuroscience*, 18(5), 620–627. <http://doi.org/10.1038/nn.3982>
- Thirion, B., Varoquaux, G., Dohmatob, E., & Poline, J.-B. (2014). Which fMRI clustering gives good brain parcellations? *Frontiers in Neuroscience*, 8(171), 167. <http://doi.org/10.3389/fnins.2014.00167>
- Toro, R., Fox, P. T., & Paus, T. (2008). Functional Coactivation Map of the Human Brain. *Cerebral Cortex*, 18(11), 2553–2559. <http://doi.org/10.1093/cercor/bhn014>
- Turkeltaub, P. E., Eden, G. F., Jones, K. M., & Zeffiro, T. A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *NeuroImage*, 16(3 Pt 1), 765–780. <http://doi.org/10.1006/nimg.2002.1131>
- van den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in Cognitive Sciences*, 17(12), 683–696. <http://doi.org/10.1016/j.tics.2013.09.012>

- Varoquaux, G., & Thirion, B. (2014). How machine learning is shaping cognitive neuroimaging. *GigaScience*, 3(1), 28. <http://doi.org/10.1186/2047-217X-3-28>
- Vogt, B. (2009). *Cingulate Neurobiology and Disease*. Oxford University Press.
- Vogt, B. A. (2005). Pain and emotion interactions in subregions of the cingulate gyrus. *Nature Reviews Neuroscience*, 6(7), 533–544. <http://doi.org/10.1038/nrn1704>
- Vogt, B. A. (2016). Midcingulate cortex: Structure, connections, homologies, functions and diseases. *Journal of Chemical Neuroanatomy*. <http://doi.org/10.1016/j.jchemneu.2016.01.010>
- Vogt, B. A., & Vogt, L. (2003). Cytology of human dorsal midcingulate and supplementary motor cortices. *Journal of Chemical Neuroanatomy*, 26(4), 301–309. <http://doi.org/10.1016/j.jchemneu.2003.09.004>
- Vorobiev, V., Govoni, P., Rizzolatti, G., Matelli, M., & Luppino, G. (1998). Parcellation of human mesial area 6: cytoarchitectonic evidence for three separate areas. *European Journal of Neuroscience*, 10(6), 2199–2203. <http://doi.org/10.1046/j.1460-9568.1998.00236.x>
- Wager, T. D., & Smith, E. E. (2003). Neuroimaging studies of working memory: a meta-analysis. *Cognitive, Affective, & Behavioral Neuroscience*, 3(4), 255–274. <http://doi.org/10.3758/CABN.3.4.255>
- Wager, T. D., Atlas, L. Y., Lindquist, M. A., Roy, M., Woo, C.-W., & Kross, E. (2013). An fMRI-Based Neurologic Signature of Physical Pain. *New England Journal of Medicine*, 368(15), 1388–1397. <http://doi.org/10.1056/NEJMoa1204471>
- Wager, T. D., Davidson, M. L., Hughes, B. L., Lindquist, M. A., & Ochsner, K. N. (2008). Prefrontal-Subcortical Pathways Mediating Successful Emotion Regulation. *Neuron*, 59(6), 1037–1050. <http://doi.org/10.1016/j.neuron.2008.09.006>
- Wager, T. D., Lindquist, M., & Kaplan, L. (2007). Meta-analysis of functional neuroimaging data: current and future directions. *Social Cognitive and Affective Neuroscience*, 2(2), 150–158. <http://doi.org/10.1093/scan/nsm015>
- Yarkoni, T. (2009). Big Correlations in Little Studies: Inflated fMRI Correlations Reflect Low Statistical Power—Commentary on Vul et al. (2009). *Perspectives on Psychological Science*, 4(3), 294–298. <http://doi.org/10.1111/j.1745-6924.2009.01127.x>
- Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, 8(8), 665–670. <http://doi.org/10.1038/nmeth.1635>

- Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., et al. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165. <http://doi.org/10.1152/jn.00338.2011>
- Zilbovicius, M., Meresse, I., Chabane, N., Brunelle, F., Samson, Y., & Boddaert, N. (2006). Autism, the superior temporal sulcus and social perception. *Trends in Neurosciences*, *29*(7), 359–366. <http://doi.org/10.1016/j.tins.2006.06.004>

APPENDIX I

Full associations between LFC regions and topics

