Abstract

# Individualized and Task-Specific Functional Brain Mapping

Mehraveh Salehi

2019

Understanding the human brain, with its remarkable ability to control higher thought, behavior, and memory, remains one of the greatest intellectual challenges in all of science. Recent advances in neuroimaging techniques, such as fMRI, have made it possible to measure neural activity in the human brain. However, fMRI data are not only massive in size but also spatially and temporally complex. In my thesis, I develop advanced data summarization techniques to study brain function and its link to behavior. While traditional fMRI studies have typically collapsed data across individuals and cognitive states, here, I develop personalized algorithms that can preserve individual and state differences. I establish that individual and state differences in brain maps (also known as parcellations) are reliable and functionally relevant, such that features extracted from these parcellations can successfully predict behavioral phenotypes of individuals as well as their brains' cognitive states. In the first empirical chapter (Chapter 2), I propose a network-based predictive model which successfully predicts attentional abilities for novel individuals and generalizes across attention tasks and behavioral measures. In Chapter 3, I develop a submodular-based algorithm for individualized network-level brain functional parcellation and demonstrate that the brain networks are different across individuals, and show reliable sex-specific patterns which can be utilized to predict individual sex. In Chapter 4, I demonstrate that the functional networks are not only different across individuals, but also reliably vary across cognitive states, such that they form a predictive signature of the underlying cognitive state. In Chapter 5, I propose a new submodular-based algorithm for individualized node-level functional parcellation and show that compared to the group-level parcellation, the proposed individualized parcellations exhibit a stronger predictive power in predicting IQ, highlighting its ability in preserving behaviorally relevant information. In Chapter 6, I demonstrate that the individual functional parcellations are unique and reliable, such that they can act as an identifying 'fingerprint'. In the final empirical chapter (Chapter 7), I challenge one of the longstanding assumptions in the field by demonstrating that the individualized node-level functional parcellations are not fixed, but they reliably reconfigure with brain's cognitive state. I demonstrate that a measure of node configuration as coarse as node size can significantly predict the task condition under which the data were acquired, as well as the within-condition task performance. The work presented here provides a theoretical and empirical foundation for future research seeking data-driven approaches to understanding brain function, and its relationship to behavior.

# Individualized and Task-Specific Functional Brain Mapping

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

in Candidacy for the Degree of

Doctor of Philosophy

by

Mehraveh Salehi

Dissertation Directors: Robert Todd Constable and Amin Karbasi

December 2019

## Table of Contents

## List of Figures and Tables

Chapter 4: Individualized functional networks reconfigure with cognitive state ........ 122

Chapter 5: A Submodular Approach to Create Individualized Parcellations of the Human Brain ........ 179

Chapter 6: Individual parcellations are unique and reliable, forming an identifying fingerprint ........ 193

Chapter 7: There is no single functional atlas even for a single individual: Parcellation of the human brain is state dependent ........ 213

# Acknowledgments

The first day I joined Yale, I could not dream about researching on the most fundamental and intriguing subject (and object) in the world: the human brain. Yet, it happened. And now, I cannot think of any better way to spend the past five (actually three) years of my life. I am incredibly grateful to the many people who helped me grow through this journey and be who I am and where I am now. This journey would have never happened without you, seriously!

First and foremost, my advisor, Todd Constable, you are the reason for where I am now. You have always surprised me with your endless support. You allowed me to freely explore opportunities, from taking on different research projects, to initiating collaborations, to taking MBA courses, to doing internships, to managing my work-life balance. You have been a true exemplar of leadership for me, the first *exemplar* that any submodular function would select! I have learned a ton from you, way more than what they teach in MBA courses.

Amin Karbasi, I cannot thank you enough for being there for me from the most vulnerable times of failure to the most exhilarating times of success. From my first day of Ph.D. to the last day, you have been the person to consult with whenever any challenges arose or any lingering questions I had about professional life. You taught me math and life. You made this whole game a *Nash equilibrium*. Thank you for being such an amazing influence in my life!

Dustin Scheinost, you taught me adjustment. You were the person who had the answer to ALL of my questions, from the bug in line 122 of the code to the story that best fits the paper. You taught me how to move beyond deep inside of an equation to address the high-level problem. Without you, I would not have gained the skills to summarize my thesis (in 3 minutes). After all, if there is one thing I do know, that is *data summarization*.

Xenios Papademetris, you were the beginning of this exciting journey. If it were not because of you, I would have not been exposed to this sublime world of neuroscience. You have been a role model of genuine kindness to me. You thought me “what I don’t know is far more relevant than what I do know.” Your sharp and insightful comments continue to encourage me.

Leandros Tassiulas, you trusted me when trust was what I needed most. You helped me find my scientist identity. Thank you for your tremendous support and encouragement in research.

Jeff Bilmes, you thought me there are no *diminishing returns* to growth and learning. Your enthusiasm and curiosity about life inspire me. I look forward to learning more and more from you.

My amazing lab mates and friends at the MRRC: Daniel Barron, I do not know how to thank you for your *stable* support and friendship, and your *flexible* discussions of life, science, and entrepreneurship. I feel so lucky to have you in my life. I am also very grateful for all your help with editing (including this acknowledgement!). Xilin Shen, your past and current work have built the foundation for this thesis. Thank you for being so generous with your time and knowledge. Abby Greene, having you beside me, both literally and figuratively, has been a continuous source of motivation. Thank you for all your help and patience with data collection and processing. Emily Finn, you thought me more than you think, I have always admired you for your determination and energy in life. Monica Rosenberg, you are an example of passion, positivity, and collaboration when it comes to science. Thank you for all you have thought me. Eve Lake, I feel indebted to individualized parcellations for initiating a beautiful friendship between us. You mean a lot to me! Dave O'Connor, thank you for cheering me up with your continuing encouragement, and of course your higher-order jokes! Corey Horien, thank you for spreading authenticity and charm. You have been the social glue of the lab, for sure. Stephanie Noble, I still remember the excitement I had when figured we are in the same lab. You are the kindest lab mate one can have. John Onofrey, thank you for always encouraging me in my entrepreneurial interests. I really hope our paths cross in the future. And many more friends whose appearance ignited new energy to move forward: Kartiga Selva, Siyuan Gao, Sarah Yip, Sophie Holmes, Emre Kiziltug, Kangjoo Lee, Yanitza Rodriguez, Zach Saltzman, and many more. I have enjoyed every single minute of my time with you, from random chats during lunch, to deep conversations on drinks. Keep in touch, and I cannot wait to see where your next adventures take you.

My great lab mates at YINS: Chris Harshaw, Ehsan Kazemi, Marko Mitrovic, and Lin Chen. You have helped me save the world with submodularity! I cannot wait to see your future accomplishments.

All the members of MRRC and YINS, thank you for creating such a friendly environment to work in. I am especially grateful to Jitendra Bhawnani, Jagriti Arora, Lesley Nadeau-Foti, Tricia Ginnetti (from MRRC), Emily Hau and Kim Kuzina (form YINS), for celebrating every little success with me along the way.

The amazing team at the Office of Career Strategy and the Graduate Writing Lab at the Poorvu Center, in particular, Hyun Ja Shin, Elena Kallestinova, Brian Frenette; because of you, I reached a new understanding and appreciation of my thesis. Thank you for all your support for 3MT.

My friends from school: Ghazal Roostaei, Mahsa Ghasemi, Roksana Ehsani, Sahar Farahbakhsh, Shirin Goshtashbpour, Shiva Bahrami, and many more. You have shaped my dream of *changing the world,* which I still carry from those glorious days.

My friends from undergraduate: Mohammad Mahmoudi, Farnaz Ghaedipour, Kamyar Azizzadeh, Parsa Mirdehghan, Mehdi Shams, Mohsen Ghannadzadeh, Sarah Mostafavi, Saba Medghalchi, Amir Mohammad Masoudi, Rasoul Shafieipour, Armin Ellessar, Farhad Heydarian, Alireza Alizadeh, and many more. An essential part of my identity is intertwined with you, and the great times we spent together. I cannot wait to see you all again soon!

Razieh Armin and Vahid Abedini, it was only one year, but you made here feel like home after so many years! Thank you!

Hosna Sheikholeslami, Yahya Paya, Ali Abdi, Mani Mirzaei, and Mana Vahid, thank you for making the first years being away from home less bitter.

Neelima Sharma, you are so special to me. Thank you for all the moments of joy and sadness we shared. I cannot imagine having spent the past 5 years without you. If nothing else, having found you is a big success of this past years for me.

Behzad Borhan, your arrival to my life coincides with all the fantastic things, and I believe this is not a coincidence. You are a big inspiration for me. You make me the complete me I want to be.

Finally, my family, you have been the meaning behind all this. My grandparents, my parents: Roya Karimi and Davoud Salehi, and my brother: Parsa Salehi, if I were to write about you, and all the great things you have done for me, the entire thesis would have not sufficed, let alone this acknowledgement. You were the home I needed most all these past 5 years. You were where I brought all my anxiety and took all the love and cherish. You are the reason I believe in myself. I love you.

## Contributors

*The following authors contributed to Chapter 2:*
Dustin Scheinost, Monica D. Rosenberg, Emily S. Finn, Marvin M. Chun, R. Todd Constable

*The following authors contributed to Chapter 3:*
Amin Karbasi, Xilin Shen, Dustin Scheinost, R. Todd Constable

*The following authors contributed to Chapter 4:*
Amin Karbasi, Daniel S. Barron, Dustin Scheinost, R. Todd Constable

*The following authors contributed to Chapter 5:*
Amin Karbasi, Dustin Scheinost, R. Todd Constable

*The following authors contributed to Chapter 6:*
Amin Karbasi, Dustin Scheinost, R. Todd Constable

*The following authors contributed to Chapter 7:*
Abigail S. Greene, Amin Karbasi, Xilin Shen, Dustin Scheinost, R. Todd Constable

# Chapter 1: Introduction

The human brain is a complex network, consisting of functionally interconnected regions whose coordinated effort gives rise to different functions. Understanding what these regions are, how they interact, and how this interaction forms a wide range of behavior has long been an essential question for human neuroscience.

Neuroimaging techniques have provided a unique opportunity to tackle this question in a data-driven way. Advances in neuroimaging techniques have allowed us to approximately measure the neural activity in the brain. Among all other techniques, functional Magnetic Resonance Imaging (fMRI) has received increased attention due to its unparalleled ability in providing a non-invasive whole-brain coverage of awake behaving humans. fMRI measures blood oxygen-level dependent (BOLD) signal as a proxy for neural activity. Although BOLD signals are far from a direct measure of the neural activity, studies have shown that the two are coupled (Logothetis, et al. 2001). The spatial resolution in fMRI is primarily determined by the volume of the smallest imaging unit, i.e., voxel. Current fMRI pulse sequences can offer spatial resolution of up to 1mm$^3$ and temporal resolution of 1 volume every second or less.

The fMRI data are massive in size and complex in their spatio temporal patterns. In the past few years, with advances in the neuroimaging techniques and the increased attention on the publicly available data sets, the field has observed a proliferation of neuroimaging data, understanding which requires advanced data analysis techniques. Understanding these data, with its all complexity is the first step towards understanding the brains.

Submodularity is at the core of computational methods developed for big data analysis, due to its unique properties which allow for efficient data summarization algorithms. Submodularity, in brief, is a property of set functions (i.e., functions that are defined on a set of items) which refers to a natural diminishing returns condition, that is, the marginal value of adding an item to a set decreases as the set we possess increases. Because of this property, submodular functions can be efficiently optimized using scalable algorithms. Submodular functions are important, not just because they provide a unique opportunity for efficient optimization algorithms, but because they are a ubiquitous property of many real-world problems, thus making it a very natural assumption in many data analysis problems. Submodular function optimization algorithms have proven useful in a wide range of applications including viral marketing (Kempe, et al. 2003), outbreak detection (Leskovec, et al. 2007), sensor placement (Krause and Guestrin 2012), feature selection (Das, et al. 2012), and natural language processing (Lin and Bilmes 2011). Given this past success, submodular functions and submodular optimization form a great candidate for brain data analysis.

For my doctoral thesis, I have developed submodularity-based optimization algorithms to intelligently summarize the brain data into informative neural signals. Unlike the traditional population-based brain studies which focus at group-level inferences, the proposed algorithms are able to capture meaningful individual differences, and state differences within an individual. I demonstrate that individual differences in brain functional organization are meaningful and predictive of individual differences in behavior. The present work builds the theoretical and empirical foundation for future data-driven research seeking to characterize the brain function and its relationship to behavior.

## 1. Neuroscience background

Human neuroscience research has long been interested in delineating the functional subunits in the brain and elucidate the functional role of these brain regions, through a technique known as *human brain parcellation*. These regions are often referred to as parcels, nodes, cortical areas, or regions of interest (RoI), in the community. Another line of research, known as *functional connectivity*, is interested in elucidating the interaction between these regions, by modeling the brain as a network of spatially distributed but functionally interconnected regions. In what follows I will provide an overview of these two well-established fields, and explain how the present work contributes to the advances of both.

Perhaps one of the earliest attempts to human brain parcellation dates back to 1909 when Korbinian Brodmann defined 52 regions based on cytoarchitecture (see Figure 1) (Brodmann 1909). Since then, there have been numerous attempts to identifying functionally meaningful regions based on different brain properties such as anatomy (Desikan, et al. 2006; Tzourio-Mazoyer, et al. 2002), microarchitecture (i.e., cyto-, myelo-, and chemo-architectonic features) (Henssen, et al. 2016; Nieuwenhuys, et al. 2015b; Zilles and Palomero-Gallagher 2001), and more recently function (Craddock, et al. 2012; Thirion, et al. 2006; Thomas Yeo, et al. 2011a; Toga, et al. 2012). Depending on the particular property considered, different brain atlases arise. Even within a particular property, different techniques and algorithms would lead to different brain atlases, and there has not yet been a consensus among neuroscientists on what is the best atlas for the human brain. Although no atlas has emerged as the dominant one to date, there is an assumption that such an atlas exists and the field has been trying to identify that by increasing the amount of data, using different modalities, and advancing the parcellation techniques. Identifying this atlas is a critical step towards understanding the brain. It forms the very initial step for many subsequent analyses, such as, functional connectivity studies and functional network parcellations (Sporns 2011).

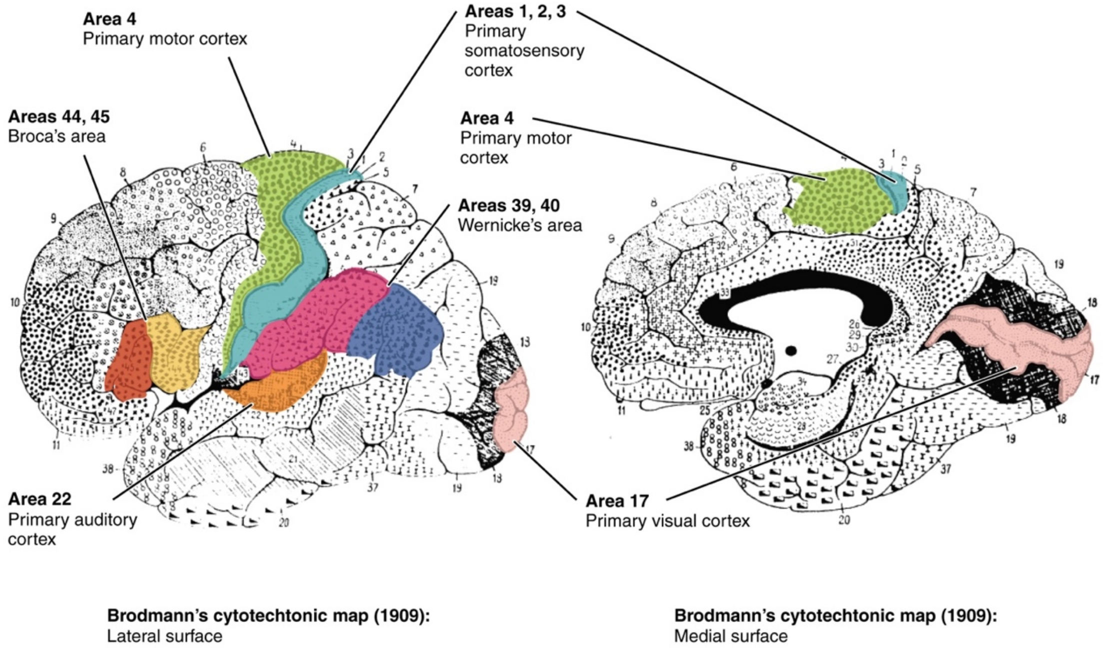

**Figure 1. Brodmann's cytotechtonic map.** Brodmann's areas, with regions marked according to their function. Reproduced from (OpenStax).

Functional connectivity research attempts to understand the interaction between brain regions by modeling the brain as a network. Nodes in this network are typically identified by a whole-brain atlas—for example, using a functional parcellation algorithm. Edge strengths are computed by measuring the synchrony of activity across these brain nodes. Synchrony can be defined by any statistical measure of similarity, with Pearson correlation being the most widely used measure in the field (Sporns 2011). Figure 2 shows an illustration of the general pipeline for identifying the brain regions and constructing the functional connectivity matrix based on them, a well-established pipeline which has been widely employed in the community.

Beyond the local connectivity patterns captured by edge strengths, global coordination patterns of brain network have also been of great interest to the field. Ample evidence suggests that the human brain has a modular functional organization (Sporns and Betzel 2016), consisting of multiple spatially-distributed modules, which are referred to as networks or systems in the community (Figure 2d). Here, the term "network" is equivalent to the concept of "subnetwork" or "community" in graph theory. Previous work has consistently identified between 5 to 20 functional networks using resting-state fMRI data and a wide range of network-level parcellation algorithms (Beckmann, et al. 2005; Damoiseaux, et al. 2006; Dosenbach, et al. 2007; Fox, et al. 2006; Lee, et al. 2012; Meunier, et al. 2009; Power, et al. 2011; Smith, et al. 2009; Thomas Yeo, et al. 2011a). The spatial organization of these networks is shown to be in alignment with task activation patterns (Smith, et al. 2009). Functional networks support a wide range of cognitive functions (Dosenbach, et al. 2007; Laird, et al. 2011) and their alterations have been linked to a range of clinical disorders (Bush 2011; Stern, et al. 2012; Zhu, et al.

2012). See (Smitha, et al. 2017) for more review and discussion on the topic of network-level parcellation.

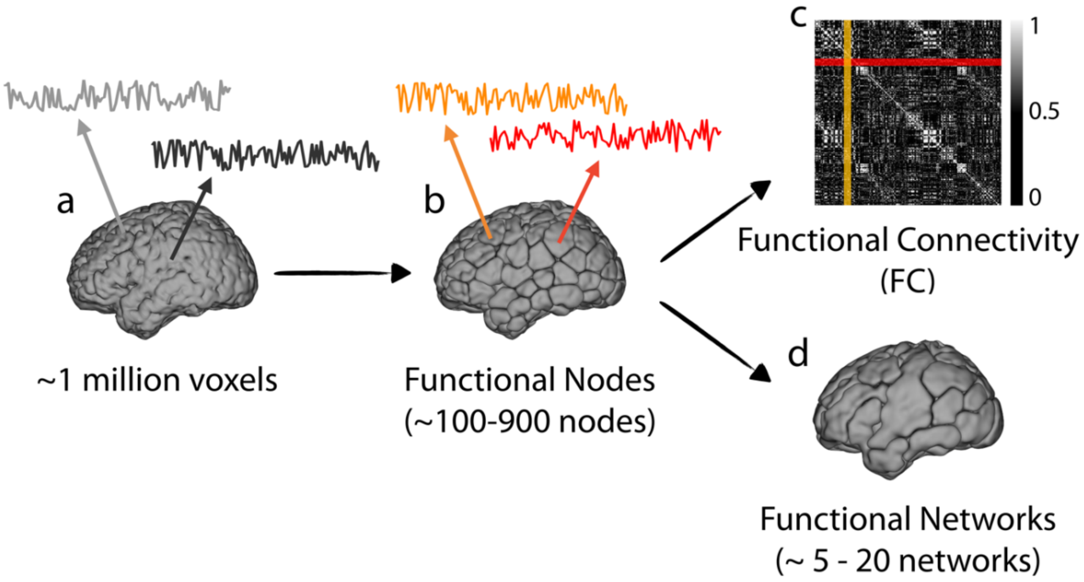

**Figure 2. Functional parcellation and connectivity pipeline.** a) Human brain voxel-level data, consisting of b) A node-level functional atlas, defined based on a functional parcellation algorithm which groups voxels into nodes according to the synchrony of their timeseries. The temporal signal (timeseries) of a node is calculated by taking the average over all timeseries of voxels in that nodes. c) Functional connectivity matrix ($n \times n$), defined by computing similarity of timeseries across all possible pairs of nodes ($n$). Element $(i, j)$ in this matrix represents the similarity between node $i$ and node $j$'s timeseries, where similarity is typically defined by Pearson correlation. d) Delineation of functional networks according to the temporal synchrony of the nodes.

Traditionally, fMRI studies have taken a group-level direction by collapsing data from all individuals in a population to draw inferences about the general blueprints of the brain function. Although group-level models provide insight into the neural mechanisms of the brain, they typically ignore the large amount of inter-individual and inter-state variability in brain function. Considering the individual variability is particularly important for clinical practices where there is considerable heterogeneity in neural and behavioral phenotypes of individuals. Preserving and appreciating the individual differences is an important step towards discovering functional *neuromarkers* of psychiatric disorders—a potential application of fMRI which many have been hoping for.

As such, recent fMRI studies have started moving towards individualized neuroscience with emphasis on behavior prediction. There have been numerous efforts in the past decade to characterize individual differences in connectivity patterns and link them to differences in behavior. A recent work form our group has shown that individual connectivity profiles are unique and reliable, acting as an identifying fingerprint (Finn, et al. 2015c). These connectivity patterns also predict individual differences in a wide range of cognitive traits such as intelligence (Finn, et al. 2015c), memory (Weber, et al. 2017), and attention (Rosenberg, et al. 2016d), or clinical symptoms such as Alzheimer's Disease (Lin, et al. 2018), schizophrenia (Du, et al. 2016), autism (Price, et al. 2014), and depression (Barron, et al. 2018). This great success in predicting individual differences from connectivity patterns highlights the potential for fMRI-based neuromarker discovery and paves the way for more personalized approaches in clinical practices.

Yet, all past work in functional connectivity have relied on group-defined atlases to assess connections across nodes and extract network properties. However, similar inter-individual variability is expected at the node-level, that is, when identifying brain regions (or nodes) at the very first step of the connectivity analysis pipeline. To what extent spatial topography of nodes vary across individuals, and how this variability is related to variations in behavior remains an interesting open question. A recent work has shown that, in many cases, inter-individual variability in nodes' spatial configuration can be interpreted as changes in functional connectivity (Bijsterbosch, et al. 2018). It is critical to isolate the effect of node configuration from connectivity changes, and also to increase the homogeneity of the nodes identified by functional parcellation approaches. If nodes are defined by collapsing data over the entire population, they are likely to include a mixture of non-homogenous timeseries which is not representative of any of the individual nodes, and thus may lead to erroneous outcomes. The same challenges hold when identifying functional networks or systems at the population level. Individualized approaches to parcellation are therefore of great interest.

Compared to population-level, individualized parcellation is typically more challenging as it requires extraction of meaningful information from a pool of confound factors, using less data, lower signal to noise ratio. Therefore, to be practically useful, an individualized parcellation technique must fulfill certain criteria, including (1) having high reproducibility within individuals; (2) having enough sensitivity to functional differences across individuals; (3) preserving correspondences across individuals to facilitate comparative studies; (4) being computationally efficien to allow recalculation of the atlas for every individual; and ideally, (5) extracting functionally relevant features which are predictive of behavior and generalize across populations.

Despite these challenges, a number of successful individualized parcellation approaches have been introduced over the past decade, using techniques such as k-means (Flandin, et al. 2002; Kahnt, et al. 2012), hierarchical clustering (Arslan and Rueckert 2015; Blumensath, et al. 2013; Meunier, et al. 2010; Moreno-Dominguez, et al. 2014), and spectral clustering (Chen, et al. 2013; Craddock, et al. 2012; Shen, et al. 2013; Van Den Heuvel, et al. 2008). Despite the initial success, the current parcellation methods fall short in meeting at least two of the above-mentioned requirements, namely the computational efficiency and the predictive power.

More importantly, all previous functional brain parcellations have focused on resting-state fMRI data, acquired from individuals at rest while they are not performing any explicit task. As such, these approaches are fundamentally limited in their ability to capture variability across different tasks and cognitive states. Yet we know the brain functional organization is not fixed, rather it dynamically reconfigures across cognitive states (Braun, et al. 2015; Cole, et al. 2014; Mennes, et al. 2013), possibly in adaptation with the underlying cognitive processes' demands (Gonzalez-Castillo, et al. 2015; Niendam, et al. 2012). However, all the commonly adopted parcellations to date, whether at the individual or group level, have defined a single functional parcellation with the underlying assumption that nodes and networks are homogenous in function and invariant in size, shape or position regardless of brain state.

How the spatial topography of nodes and networks vary across cognitive states, to what extent such variations are functionally meaningful and reflective of the underlying cognitive processes, and whether the state-dependent parcellations could lead to more insights about the relationship between brain and behavior are all questions that I attempt to answer in my thesis.

## 2. Summary of following chapters

This thesis includes six empirical chapters (Chapters 2-7), all about developing data summarization algorithms which allow for a better understanding of the brain at the individual and state level. Five of these chapters (Chapters 2-5 and 7) are already published or are currently in press (Salehi, et al. 2018a; Salehi, et al. 2018b; Salehi, et al. 2017a; Salehi, et al. 2017b). Chapter 6 describes preliminary results from an ongoing project. Below is a brief description of each empirical chapter.

While there is significant evidence that patterns of brain functional connectivity reconfigure with brain state, it is not fully clear how these reconfigurations relate to behavior and attention. In chapter 2, I develop a generalized, interpretable, and predictive model of attention based on brain connectivity reconfiguration. I show that models based on connectivity reconfiguration predict individual differences in attention, both within- and across four independent data sets with distinct attention-demanding tasks and behavioral measures: working memory task (*n*-back), continuous performance task (gradCPT), and the Attention Network Task (ANT). A post hoc analysis of predictive features reveals that individuals with greater flexibility in their connectivity strength during rest, but higher stability during task, exhibit better attention performance. While this chapter investigates the inter-individual and inter-state variability in functional connectivity, it starts from a fixed functional atlas both across individuals and states. The observed connectivity reconfiguration provokes the question of whether the underlying functional atlas reconfigures across individuals and even within an individual, across cognitive states. In the subsequent chapters, I attempt to answer this question at node- and network-level.

In Chapter 3, I introduce a parcellation technique to delineate functional networks both at the individual- and group-level. The proposed technique deploys the notion of *submodularity* to jointly parcellate the cerebral cortex while establishing an inclusive correspondence between the individualized functional networks. Using this parcellation technique, I successfully establish a cross-validated predictive model that predicts individuals' sex, solely based on the parcellation schemes (i.e. the node-to-network assignment vectors). The successful sex predictions illustrate that individualized parcellation of functional networks can reveal subgroups in a population and suggests that the use of a global network parcellation may overlook fundamental differences in network organization. This is particularly important for studies comparing patients versus controls or even patient subgroups. This individualized approach to the study of brain’s functional organization has many implications in cognitive and clinical neuroscience.

In chapter 4, I extend the analysis of chapter 3 to explore the possibility of functional networks changing across cognitive states, even when considered within an individual. There is extensive evidence that functional organization of the human brain varies dynamically as the brain switches between task demands, or cognitive states (Braun, et al. 2015; Gonzalez-Castillo, et al. 2015; Mennes, et al. 2013; Niendam, et al. 2012; Shine, et al. 2016a; Shirer, et al. 2012). However, the network structure of the brain has been considered static. In this chapter, I formulate an individualized and state-specific network parcellation pipeline, using fMRI data obtained across multiple cognitive states (task-evoked and rest conditions) and across multiple subjects. By assessing the parcellation reconfigurations, I demonstrate that the brain's functional networks are not spatially fixed, but that many nodes change their network assignments as a function of cognitive state. I show these reconfigurations are highly robust and reliable such that they can be used to predict cognitive states with up to 97% accuracy.

Chapter 5 extends the network-level parcellation analysis to delineation of functional nodes at the individual-level. All previous chapters were premised on the implicit assumption that functional nodes are fixed and shared across individuals and states. In Chapter 5, I challenge this assumption by developing an individualized parcellation algorithm to identify functional nodes specific to each individual. The algorithm proposed in earlier chapters requires computing pairwise distances across all elements of the ground set. However, computing the pairwise distances across millions of voxels is computationally expensive, and thus not scalable to the voxel-level analysis. In this chapter, I address the computational challenge by introducing a distributed approach which leverages previously developed group-level parcellation as the initial point, and morphs it to account for the individual's functional organization. I then evaluate the strength of the proposed individualized parcellation in capturing informative individual specific features by showing that it improves the accuracy of a predictive model that predicts IQ using functional connectivity.

In Chapter 6, I present preliminary results suggesting that the individual variability in functional parcellations are substantial, and reproducible across sessions. Using fMRI data from multiple individuals and sessions, I show that individuals can be identified with up to 99% accuracy solely based on their functional parcellation pattern, suggesting that the individualized functional parcellations are robust and reliable, acting as an identifying "fingerprint".

Yet, perhaps contrary to what is traditionally assumed, individualized functional parcellations are not fixed, but they reliably reconfigure by task. In Chapter 7, the final chapter, I provide strong evidence that functional parcellation of the human brain does not yield a fixed atlas of node definitions, rather, nodes reconfigure substantially and in a meaningful manner, according to brain state. Neuroimaging studies have long sought to define a single functional atlas that "best" represents the brain, assuming that such an atlas exists, at least at the individual level. Using fMRI data from a highly sampled subject as well as two independent replication data sets, I challenge this assumption, showing, for the first time, that node boundaries reconfigure across states, while being consistent for a given state. These reconfigurations are robust and reliable both within a subject across different sessions, and across different subjects, such that a topographic measure of node size can significantly predict the task condition and the within-condition task performance.

These findings have a wide range of implications for human brain mapping and, more generally, systems neuroscience. Most functional connectivity studies to date have used a fixed functional atlas. The present results, however, suggests that a single functional atlas is not applicable to all states. The state-dependent changes in connectivity could be attributable, at least in part, to reconfiguration of the underlying parcellation, and this should be considered as part of the interpretation of the results. I hope the proposed individualized state-specific parcellation methods will spark future research into custom state-dependent atlases, which may in turn enhance our understanding of the brain's state-dependent functional organization.

While much work is still left to be done, and many improvements yet to be brought about, this body of research hopes to provide a data-driven methodology to facilitate a more comprehensive account of human brain at the individual and state level. It is my hope that the work presented here serves as a theoretical and empirical foundation for future research into individual and state differences in brain function, behavior, and disease.

# Chapter 2: Brain network stability and flexibility form a generalizable predictive model of sustained attention

## Abstract

While there is significant evidence that patterns of brain functional connectivity reconfigure with brain state, it is not fully clear how these reconfigurations relate to behavior and attention. Using a higher-order measure of connectivity, we develop a generalized, interpretable, and predictive model of attention based on brain network reconfiguration—that is, the degree to which networks exhibit stable or flexible functional connectivity profiles. We show that models based on network reconfiguration predict individual differences in attention, both within- and across four independent data sets with distinct attention-demanding tasks and behavioral measures: working memory task (*n*-back), continuous performance task (gradCPT), and the Attention Network Task (ANT). A post hoc analysis of predictive features reveals that individuals with greater flexibility in their network strength during rest, but higher stability during task, exhibit better attentional abilities.

# 1. Introduction

While all of us have experienced paying attention in our life, the term “attention” is notoriously difficult to define, perhaps due to its complex and non-unitary notion (Chun, et al. 2011; Lavie, et al. 2004; Pashler 1998; Rosenberg, et al. 2017a). Previous studies have developed different classifications and taxonomies of attention (Chun, et al. 2011). An influential model of attention divides attentional mechanisms into two classes of top-down (goal-directed), and bottom-up (stimulus-driven) approaches (Corbetta and Shulman 2002). Another well-established model of attention divides the attention into three sub-systems according to the neural processes they involve: alerting, orienting, and executive-control (Fan, et al. 2005; Posner and Petersen 1990; Thiel, et al. 2004). These attentional mechanisms are common to everyone; however, their efficiency differs across individuals. Therefore, while developing such concrete taxonomies of attention holds great potential in revealing fundamentals of attention system, their performance in predicting inter-individual variability is not clear.

With advancements in neuroimaging techniques (such as fMRI) and the proliferation of neuroimaging data, there has been increased attention on data-driven approaches to studying attention. These models traditionally take a group-level approach, describing features of neural signals that are, on average, related to attention across subjects (Fan, et al. 2005). Although group-level models can provide insight into neural mechanisms of attention, they do not necessarily capture inter-individual variation in attentional abilities. More recently, individual-level based analyses have received increased attention and studies have begun to develop data-driven models to describe (Fortenbaugh, et al. 2015; Shine, et al. 2016b) and predict (Madhyastha, et al. 2015; Rosenberg, et al. 2016d;

Rosenberg, et al. 2017b; Yoo, et al. 2018) individual differences in attentional abilities. Our recently developed model, connectome based predictive modeling (CPM) (Shen, et al. 2017), has shown success in predicting a range of behavioral tests (Beaty, et al. 2018; Finn, et al. 2015b; Hsu, et al. 2018; Rosenberg, et al. 2016d). Through these models, it has become evident that functional connectivity patterns estimated during rest and task performance can provide a reliable neuromarker of cognitive and attention function.

We recently developed and validated a technique, connectome-based predictive modeling (CPM) (Shen, et al. 2017), that uses individual functional connectomes to predict cognitive and attention function(Finn, et al. 2015a; Rosenberg, et al. 2016a; Rosenberg, et al. 2016b). CPM models have shown success in predicting individual differences in a range of attention measures (Rosenberg, et al. 2017a), have generalized to predict attention deficit hyperactivity disorder (ADHD) symptoms (Rosenberg, et al. 2016a) and pharmacologically induced changes in attention function (Rosenberg, et al. 2016b), and have contributed to our understanding of the networks underlying attention, both during task-performance and resting-state. Through these models, attention has been characterized as a network property of the brain, where networks are statically defined within each functional condition.

Despite significant evidence on stability of functional networks across sessions and states (Cole, et al. 2014; Finn, et al. 2017; Gratton, et al. 2018; Laumann, et al. 2016; Noble, et al. 2017a), researchers have long recognized that the human brain displays network-level reconfiguration during resting-state (Calhoun, et al. 2014; Zalesky, et al. 2014), task performance (Braun, et al. 2015; Gonzalez-Castillo, et al. 2015; Shine, et al. 2016a), and in transition from resting-state to task (Mennes, et al. 2013; Niendam, et al.

2012; Shirer, et al. 2012). To what extent such changes are meaningful and can contribute to our understanding of behavior, remains an open question. Specifically, there is an implicit assumption that connectivity patterns need to be fundamentally stable in order to serve as a reliable neuromarker for behavioral variables and clinical symptoms (Gratton, et al. 2018). This is in line with characterization of attention as a network property of the brain. However, the state-evoked network-level reconfigurations are also likely to be important in understanding the neural mechanism underlying attention.

Here we build upon our previous findings and extend them to show that attention is, in fact, a property of the brain network *reconfiguration*. We develop a cross-validated predictive model whicih uses network reconfigruation across multiple rest and task states to predict individual differences in attention, both within- and across four independent data sets with distinct attention-demanding tasks and behavioral measures. The tasks include the *n*-back working memory task (which was used in two of the data sets), the gradual-onset continuous performance task (gradCPT), and the Attention Network Task (ANT). Further investigation of the predictive model reveals that individuals with greater flexibility in their network strength during resting-state, but higher stability during task states, exhibit better attentional abilities. The differential role of network stability and flexibility in characterization of attention is an interesting observation, which is further corroborated and replicated in four independent data sets with distinct participants, tasks, behavioral measures, and scan sites.

## 2. Results

### 2.1 Higher-order measure of network reconfiguration

These analyses were performed using fMRI data from four independent studies with individuals performing separate attention tasks. Each data set included two runs of resting-state and two or more runs of task-based fMRI (variable across data sets). Details of the tasks have been published elsewhere (Barch, et al. 2013; Rosenberg, et al. 2015b; Rosenberg, et al. 2018; Rosenberg, et al. 2016d) and are briefly described below.

Three data sets were collected at Yale University. The first consisted of 27 subjects performing *n*-back working memory task (*n*-back task (Rosenberg, et al. 2015b)) for three runs, the second consisted of 18 subjects performing the gradual-onset continuous performance task (gradCPT (Esterman, et al. 2012; Rosenberg, et al. 2016d; Rosenberg, et al. 2013)) for three runs, and the third consisted of 38 subjects performing the attention network task (ANT (Rosenberg, et al. 2017b)) for five runs. The fourth data set was from the Human Connectome Project (HCP (Van Essen, et al. 2013a)) 900 Subjects release (S900) and consists of 717 subjects performing the *n*-back working memory task for two runs, one with left-right (LR) and the other with right-left (RL) phase encodings. Performance was assessed using sensitivity (d′) for the *n*-back and gradCPT data sets, the coefficient of variation of correct-trial reaction time (RT CV) for the ANT data set, and the median correct-trial reaction time (Median RT) for the HCP data set. Only a limited number of subjects overlapped across Yale data sets which were excluded as appropriate (see Methods for more details).

Functional connectivity matrices were assessed using a functional brain atlas (Shen, et al. 2013) consisting of 268 nodes covering the whole brain. The Pearson correlation coefficients between the time courses of each possible pair of nodes were calculated and normalized using Fisher’s z-transformation to construct 268×268 symmetrical connectivity matrices, where each element represents a functional connection, or edge, between two nodes. Network strength vectors were calculated from the connectivity matrices by summing over all the edge values incident to a node and taking the absolute value of the sum, resulting in a 268×1 vector. This was performed for each subject and each run separately (Figure 1, steps 1 and 2).

Next, for every subject, the Pearson correlation coefficients between strength vectors from task runs (denoted as Task Correlation or TC) and rest runs (denoted as Rest Correlation or RC) were calculated and normalized using Fisher's z-transformation (Figure 1, step 3). TC and RC matrices measure the similarity (or difference) of network strength across runs of task performance and rest, respectively. They summarize the network strength reconfiguration across runs, such that higher values of TC or RC indicate a higher functional stability in brain connection and lower values reflect a higher functional flexibility in the network. We refer to TC and RC as higher-order measures of network reconfiguration, as they are derived from two successive correlations and thus contain information from the higher moments. Figure 1 provides a visual illustration of the entire pipeline, which was repeated for every subject in all four data sets.

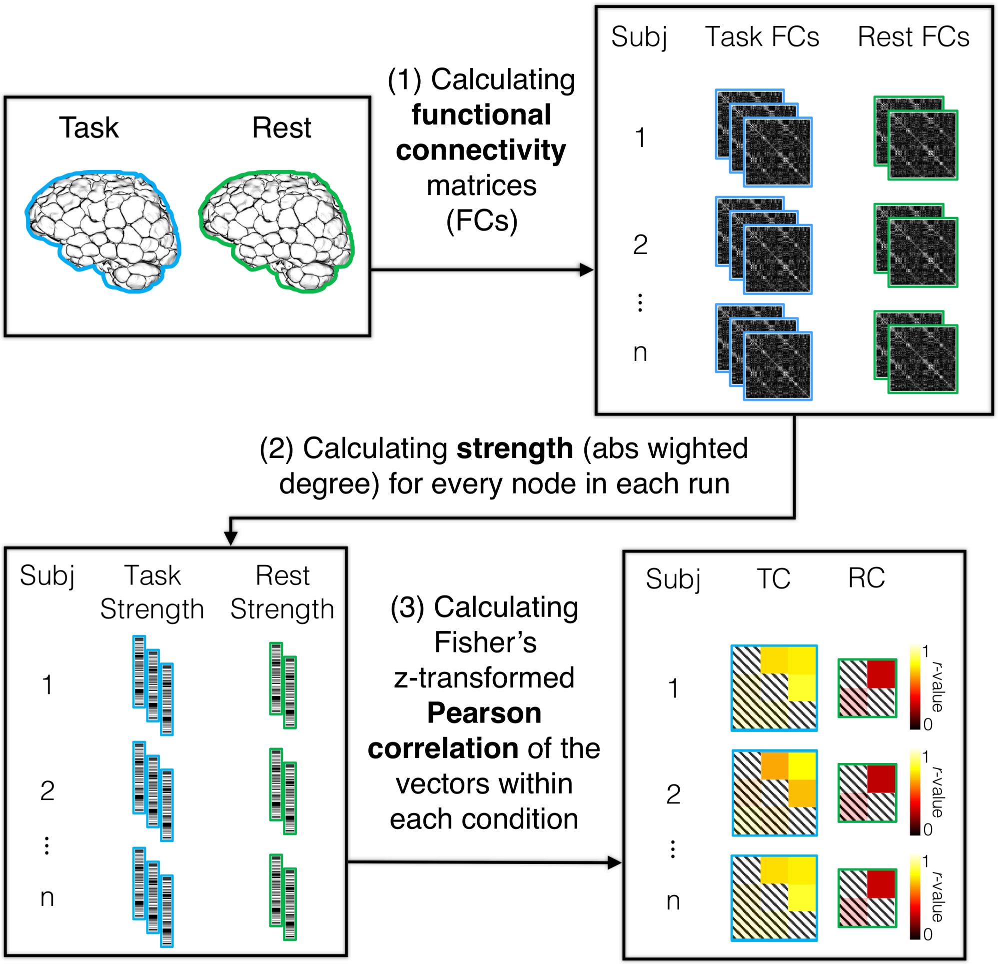

**Figure 1. The higher-order predictive modeling pipeline.** For each functional run, the functional connectivity matrix is calculated, and the corresponding strength vector is derived. Next, the correlation between the strength vectors are calculated for each of the task and rest conditions, yielding two higher-order matrices: TC (Task Correlation) and RC (Rest Correlation).

Figure 2 displays TC (blue border) and RC (green border) matrices averaged over all subjects in each data set. We observed that TC values are on average significantly higher than RC values ($p<0.0001$, two-sided Mann-Whitney U test, replicated for all four data sets), suggesting that there is significantly higher stability in network strength between task runs relative to the rest. This pattern was observed in all four data sets, indicating the generality of this observation across multiple attention tasks.

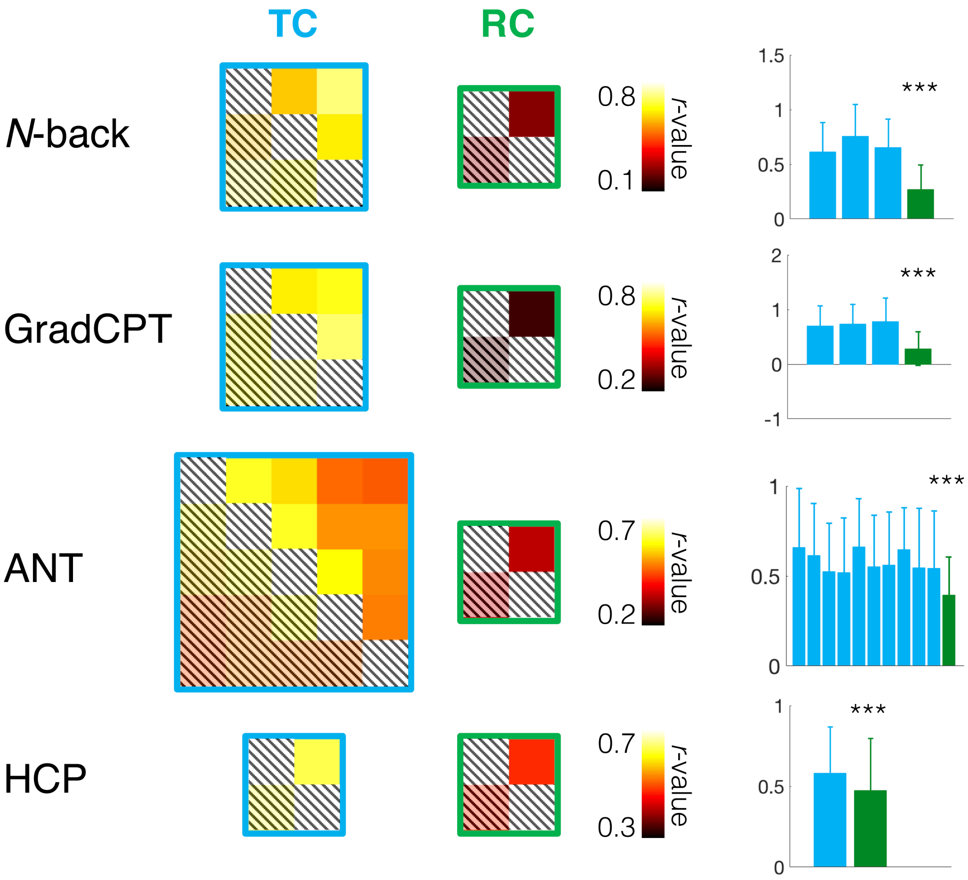

**Figure 2. A visualization of the higher order measures (TC and RC).** TC matrix was calculated for every subject as the pairwise Pearson correlation of the strength vectors derived from different runs of the task performance. Similarly, RC matrix was derived by correlating the strength vectors across different runs of resting state. All Pearson *r*-values were normalized using Fisher z-transformation. The average of these matrices across subjects is depicted on the left. The distribution (across subjects) of each element of the TC and RC is displayed on the right. A two-sided Mann-Whitney U test was performed to test whether $TC_{avg}$ values were significantly higher than $RC_{avg}$ values; *** indicates $p<0.001$.

### 2.2 Generalized predictive model of attention from network reconfiguration

We next investigated individual differences in network strength reconfiguration. Motivated by the pattern observed at the population level (i.e., TC>RC), we hypothesized that the differences between TC and RC may predict individual differences in attentional abilities (here operationalized as task performance). To test our hypothesis, we developed a linear regression model to predict task performance using TC and RC as features, and tested this model using both internal (within-data set) and external (cross-data set) validations. Since different data sets had different TC and RC matrix sizes, we used the average of TC elements ($TC_{avg}$) and RC elements ($RC_{avg}$) as predictive features, and designed a linear model as follows:

$$y = \beta_0 + \beta_{TC} \times TC_{avg} + \beta_{RC} \times RC_{avg}, \qquad (1)$$

where y represents the task performance measure, $\beta_0$, $\beta_{TC}$, $\beta_{RC}$ are the model coefficients which are learned from the training set, and $TC_{avg}$ and $RC_{avg}$ are the higher order connectivity measures calculated for every subject.

### 2.3 Internal validation: predicting attention within a data set

To determine whether TC and RC predict task performance in novel individuals, we performed a leave-one-out cross-validation procedure. For each data set, we trained our linear model over n-1 subjects and tested to predict task performance for the left-out subject. The model's predictive power was assessed using Pearson correlation between the predicted behavioral scores and the observed ones (Figure 3, diagonal plots) and the significance of the correlations were assessed by permutation testing (see Methods). The model significantly predicted task performance for novel subjects in all four data sets

(Figure 3: *n*-back, $r=0.26$, $p<0.05$; gradCPT, $r=0.75$, $p<0.01$, ANT, $r=0.18$, $p<0.05$; HCP, $r=0.12$, $p<0.01$). It suggests that models defined on network strength reconfiguration, that is, TC and RC, can significantly predict attention and generalize across subjects, tasks, and attention measures.

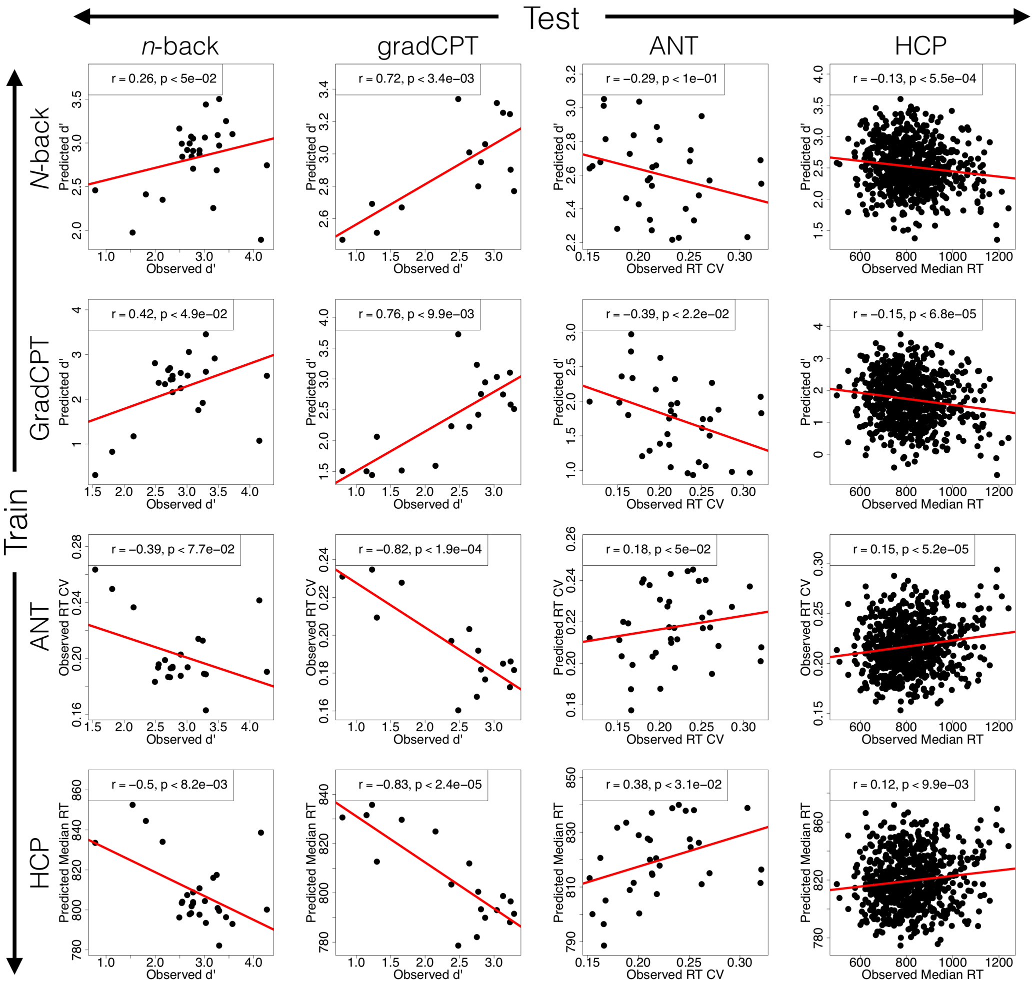

**Figure 3. The predictive analysis using a linear model with TCavg and RCavg as features**. The predicted behavioral scores versus observed ones are displayed for within-data set (diagonal elements) and cross-data set (off-diagonal elements) predictions. Within data set predictions were performed using leave-one-out cross-validation analysis. At every step a linear model was trained on n-1 subjects and tested on the left-out subject. Pearson correlation between the predicted and observed behavioral scores is reported along with the significance computed by a permutation test. Cross-data set predictions were performed by training a linear model on participants of one data set and predicting the performance for the other set. Of note, when the model is trained (tested) on *n*-back or gradCPT (behavior = d') and tested (trained) on ANT (behavior = RT CV) or HCP (behavior = Median RT), the trained and tested behaviors are inherently inversely correlated with each other. Therefore, the negative correlation between the observed and predicted scores is expected.

## 2.4 External validation: predicting attention across data sets

As a more powerful indicator of model generalizability, we tested whether models defined on network strength reconfiguration can predict attentional abilities across independent data sets. For every pair of data sets, we trained our linear model (defined in Eq. 1) on subjects of one data set (training set) and tested to predict task performance for the novel subjects in another data set (testing set). The model's predictive power was assessed using Pearson correlation between the predicted behavioral scores and the observed ones. Note that in the cross-data set prediction setting, we predict how the subjects in the testing set would hypothetically perform if they were to perform the task in the training set. As such, the generalizability of the model is inherently limited by the difference in the tasks and the performance measures. Nonetheless, all sixteen pairwise models, with the exception of [ANT, *n*-back] pairs, yielded significant predictions ($p<0.05$; Figure 3, off-diagonal plots).

Models were successful when trained and tested on data sets with matching task performance measures, that is [n-back, gradCPT] pairs with sensitivity (d') and [ANT, HCP] pairs with reaction time statistics (RT CV and Median RT). Specifically, when trained on subjects in the *n*-back data set, a model significantly predicted d' values of unseen subjects in the gradCPT data set ($r=0.72$, $p<0.003$). A separate model trained on the gradCPT data set significantly predicted the d' scores for the novel subjects in the *n*-back data set ($r=0.42$, $p<0.05$). Similarly, when trained on subjects in the ANT data set, a model significantly predicted RT CV for subjects in the HCP data set ($r=0.15$, $p<0.00005$). A separate model trained on HCP data set predicted Median RT for subjects in the ANT data set ($r=0.38$, $p<0.03$).

Interestingly, predictive models also generalized between data sets with non-matching performance measures. Specifically, when trained on subjects in the *n*-back data set, a model significantly predicted d' values for unseen subjects in the HCP data set ($r=-0.13$, $p<0.0006$). A separate model trained on the HCP data set significantly predicted the Median RT scores for the novel subjects in the *n*-back data set ($r=-0.5$, $p<0.008$). Similarly, when trained on subjects in the gradCPT data set, a model significantly predicted d' for novel subjects in the ANT data set ($r=-0.39$, $p<0.02$) and HCP data set ($r=-0.15$, $p<0.00007$). The reverse models were also successful and significantly predicted the RT CV ($r=-0.82$, $p<0.0002$) and Median RT ($r=-0.83$, $p<0.00002$) for subjects in the gradCPT data set. The negative correlation between predicted and observed scores is expected given that higher sensitivity values (d') correspond to better attentional performance whereas higher reaction time statistics (RT CV and Median RT) correspond to worse attention. Note that the ceiling model performance is determined by both the reliability of the attention tasks and the theoretical relationship between d' and RT statistics, which is unlikely to be a perfect inverse correlation.

When trained on Yale data sets ($n=18$, 27, and 39) and tested on HCP data set ($n=717$), the training sample was smaller than the testing sample. This could result in high variance in the parameter estimation, which is a barrier to model performance and generalizability. Nonetheless, all three linear models significantly predicted behavior for HCP subjects.

That the models built on TC and RC significantly predicted attention across four independent data sets with independent subjects, different attention tasks, different performance measures, and distinct acquisition sites is a remarkable finding, highlighting the significance of the information captured by these higher order measures (TC and RC).

## 2.5 Interpreting the predictive features, the differential role of task and rest

One advantage of the proposed linear model based on TC and RC is that it offers an immediate interpretation of the underlying brain mechanism. TC and RC both measure the underlying functional network reconfiguration across runs. Motivated by this, we sought to investigate how TC and RC are related to attention by analyzing their coefficients in the linear model (i.e., TC coefficient ($\beta_{TC}$) and RC coefficient ($\beta_{RC}$) in Eq. (1)

Figure 4 displays TC coefficients (blue) and RC coefficients (green) for the linear models trained and tested on all pairwise combination of the data sets. When trained on $n$-back and gradCPT data sets with sensitivity (d') as the observed task performance, the model learns positive coefficients for TC ($\beta_{TC} > 0$) and negative coefficients for RC ($\beta_{RC} < 0$), suggesting that d' is significantly positively related to TC (p<0.03, one-tailed t-test) and significantly negatively related to RC (p<0.03, one-tailed t-test). When trained on ANT and HCP data sets with RT statistics as task performance, the model learns negative coefficients for TC ($\beta_{TC} < 0$) and positive coefficients for RC ($\beta_{RC} > 0$), suggesting that RT statistics are significantly negatively related to TC (p<0.02, one-tailed t-test) and positively related to RC. The p-values are reported after ruling out the possibility of multicollinearity (see Table S1 for details). Given that larger d' and smaller RT statistics both indicate better attentional abilities, our findings suggest that higher TC and lower RC predict better attention performance.

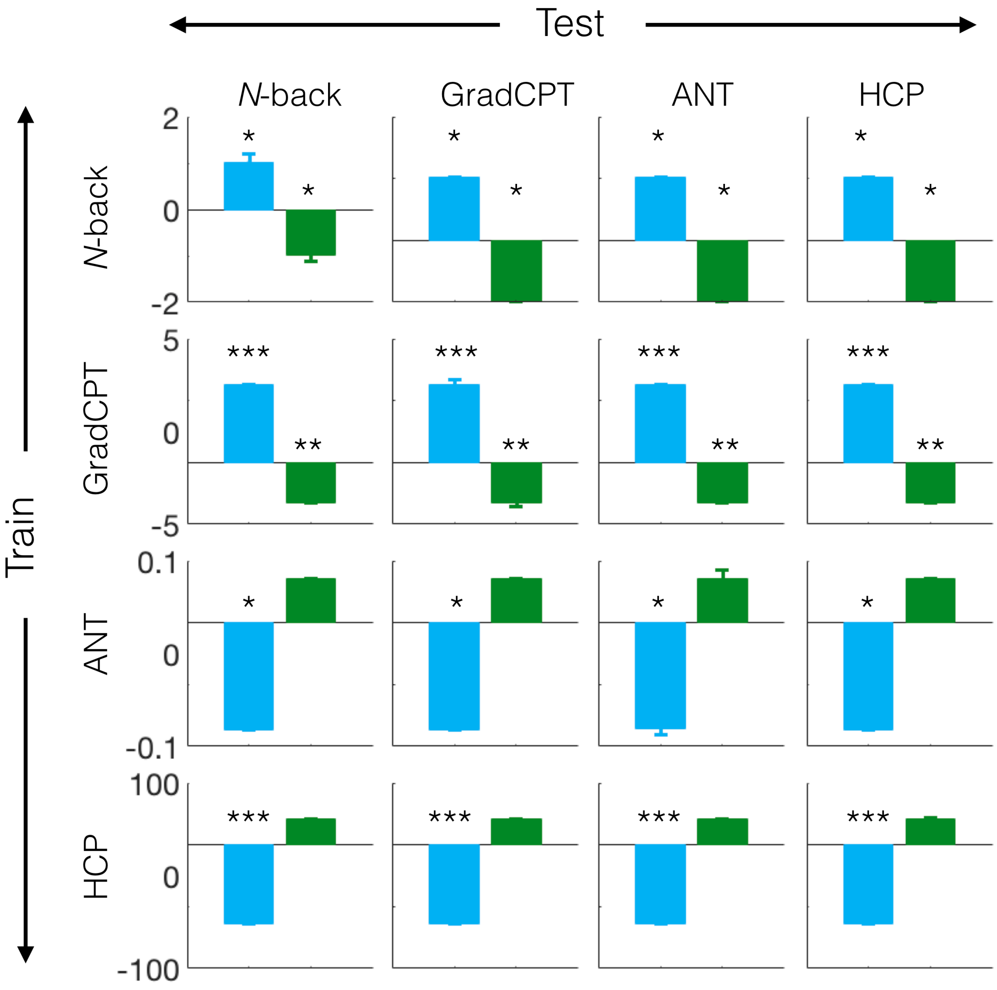

**Figure 4. The linear model coefficients for the within-data set and cross-data set predictive analysis.** TC coefficients are positively associated with attentional abilities (whether measured by higher sensitivity [n-back and gradCPT] or lower reaction time [ANT and HCP]). Conversely, RC coefficients are negatively associated with attentional performance. The significance of the relationship between each predictor and the output is assessed by one-tailed t-tests (using the Pr(>|t|) component in R's linear model). See Table S1 for the statistical significance and multicollinearity analysis. *** = p<1e-03, ** = p<1e-02, and * = p<5e-02.

Rooted in the network strength reconfigurations, higher TC indicates a higher stability (or consistency) in network organization during task runs. Similarly, lower RC reflects higher flexibility (or variability) in the network rewiring during rest. Together, these findings demonstrate the differential role that task and rest play in attentional abilities. Replicated across four independent data sets with varied attentional challenges, our findings suggest that better attentional abilities are associated with higher *stability during task* runs and higher *flexibility during rest* runs.

We also studied whether the linear model based on TC or RC values alone (i.e., a one-dimensional linear model) could predict attention across data sets. We observed that whereas the model trained on RC element alone could not predict attention, the model trained on TC element alone was successful (Figure S1), yet the predictions were less significant than or comparable to the models trained on both TC and RC. This was the case despite the lower dimensionality of the former (which could reduce variance and enhance performance), suggesting that RC can contribute unique information about individual differences in attention.

## 3. Discussion

We proposed a generalized, interpretable and predictive model of attention based on brain network reconfiguration—that is, the degree to which networks exhibit stable or flexible functional connectivity profiles. The proposed model predicted individual differences in attention abilities within and across four independent data sets with different attention tasks (the gradual-onset continuous performance task, the Attention Network Task, a working memory task collected at Yale, and a different working memory task collected by the HCP Consortium), different measures of attention (sensitivity vs. reaction time statistics), and different data acquisition sites (Yale University vs. Washington University). Further investigation of these models indicated higher stability in network strength during task runs and higher flexibility during rest runs are together predictive of better attentional abilities.

### 3.1 A data-driven and predictive characterization of attention

Previous work has shown great success in defining taxonomies of attention common across individuals (Chun, et al. 2011; Corbetta and Shulman 2002; Fan, et al. 2005; Posner and Petersen 1990; Thiel, et al. 2004). While they hold great potential in describing the fundamentals of attention mechanism in the brain, they do not address individual variability in attentional abilities. More recently, studies have increased focus on individual variability and have developed descriptive (i.e., correlation-based) models of attention from brain functional connectivity (Braun, et al. 2015; Fortenbaugh, et al. 2015; Shine, et al. 2016a; Shine, et al. 2016b). These studies have been successful in explaining the causal mechanisms of attention; however, it is not clear if they can predict behavior for novel subjects—that is, to accurately forecast behaviors that have not yet been observed—thus, their use for clinical and translational practices is debatable. Although modern neuroscience and psychology is increasingly focusing on predictions, rather than explanation, the distinction between the two is rarely appreciated (Yarkoni and Westfall 2017).

More recently, studies have started taking pains to develop more cross-validated predictive models that generalize to novel subjects (Doyle, et al. 2015; Haynes 2015;

Woo, et al. 2017). Even within these predictive models, the majority adopt internal validation strategies such as leave-one-subject-out cross-validation (Ekman, et al. 2012), as opposed to external validation using independent samples. Nonetheless, external validation is necessary for testing whether models of symptoms, traits, and behavior generalize beyond a single training data, a necessary precursor for practically useful models (Yarkoni and Westfall 2017). Only recently have cognitive neuroscientists attempted to validate predictive models in independent datasets; the work by Rosenberg et. al (Rosenberg, et al. 2016d) was the very first study to demonstrate the use of brain networks for external prediction of attention. Here, we propose a model that predicts attention across four independent datasets with different attention tasks, behavioral measures, data acquisition sites and imaging parameters. In doing so, we meet the highest standard of generalizability, that is, external validation across entirely independent samples (Riley, et al. 2016).

### 3.2 Characterizing attention through network reconfiguration

Previous work in our lab has shown that attention is a network property of the brain (Rosenberg, et al. 2016d). Using a novel technique, connectome-based predictive modeling (CPM), we had previously identified a set of predictive edges in the brain that can predict sustained attention in a cross-validated setting (Finn, et al. 2015c; Rosenberg, et al. 2016d; Shen, et al. 2017). This work builds on our previous findings and extends it to show that attention is, in fact, a property of the brain network *reconfiguration*. Our work is the first to demonstrate that measures of whole-brain network reconfiguration predict attention across independent samples. By focusing on the similarities of connections across runs, rather than the connections themselves, we provide a non-stationary model to investigate attention.

Given the dynamic notion of attention mechanisms in the brain, its reflection in brain network reconfiguration is expected. By leveraging network changes across runs of task (TC) and rest (RC), we capture the factors of attention that are common across a wide range of attention tasks. Our findings demonstrate the differential role that task and rest runs play in attentional abilities; whereas stability in network strength during task performance predicts better attention (Figure 4, TC coefficients positively related to task performance), stability during rest actually predicts *worse* attentional abilities (Figure 4, RC coefficients negatively related to task performance). This observation was replicated across all within- and cross-data set predictions, indicating the generalizability of this finding across attention tasks and attention measures.

### 3.3 Network stability during task and flexibility during rest predict better attention

Attentional mechanisms are complex, requiring the brain to select, modulate, and sustain the focus on information that are most relevant to behavior. Given the similarity of the challenges performed across task runs, it is reasonable to expect similar patterns of information selection and modulation to be leveraged for the emergence of the most relevant behavior. By exploring the coefficients of our linear model built on higher-order measures, we observe that task-correlation (TC) has a positive contribution to the task performance (Figure 5). That the higher similarity across task runs indicates a better task performance could demonstrate an optimal use of attentional resources by an engaged and attentive brain. This observation highlights the contribution of *neural stability* during task runs in sustaining attention and vigilance for a goal-directed behavior. This phenomenon was also observed at the population level with TC values being significantly higher than RC (Figure 2).

Resting state, on the other hand, imposes insufficient constraints on cognitive processes in the brain (Stoffers, et al. 2015). In the absence of overt stimuli and external task demands, resting-state cognitive processes tend to be internally directed, similar to what is experienced during mind-wandering (Gruberger, et al. 2011; Smallwood and Schooler 2006). Given the unconstrained notion of resting states, such processes are more likely to vary across different runs. Through the same exploration of linear model coefficients, we observed that rest-correlation (RC) negatively contributes to the task performance, that is, lower similarities across rest runs provoke a better task performance. These observations are in line with recent studies investigating the association of *neural flexibility* with cognitive flexibility (Braun, et al. 2015), and learning (Bassett, et al. 2011; Reddy, et al. 2018). Consistent with previous work, we demonstrated that higher network flexibility during rest, defined as variation in network strength, provokes heightened attention.

Recent work has characterized these network-level reconfiguration in the context of flexibility, and has shown individual differences in network flexibility can predict differences in a wide range of cognitive processes including memory (Akers, et al. 2014), learning (Bassett, et al. 2011), and cognitive control. However, Our work builds on top of previous findings and demonstrates, for the first time, that the two features together form a predictive and generalizable model of individual differences in attention.

### 3.4 Whole-brain characterization of attention

The proposed predictive features—TC and RC—are derived from the whole brain. They capture the correlation between *strength* vectors which have elements from the whole brain, and the *strength* measure itself is a whole-brain network measure. Characterizing attention by whole-brain network measures is in line with recent findings (Castellanos and Proal 2012; Rosenberg, et al. 2016d) suggesting that attention, like other higher-order cognitive abilities, is likely encoded in distributed neural systems involving networks of many regions (Finn, et al. 2015c). While the traditional function-location mappings (Fan, et al. 2005) provide a systematic approach to study attention, restricting our measurements to a specific region of the brain may not provide a holistic view of the complex attentional mechanisms ranging from information gathering to selection of relevant information, modulation of the selected information to sustaining the attention over time information (Cohen, et al. 2015; Desimone and Duncan 1995). Models that incorporate information from the entire brain, like the one developed here, may better account for the concerted and inter-connected processes that give rise to attention (Rosenberg, et al. 2017a; Rosenberg, et al. 2016d), and therefore could better capture an “overall” ability of subjects to sustain attention independent of a specific task or a specific measure.

### 3.5 Additional consideration

We estimated network reconfiguration by considering the entire run, estimated across different runs, in this work conditions with acquisitions spanning a series of continuous performance, event-related, and blocked tasks. Future work can extend this work to characterize attention by investigating the network reconfigurations over considerably shorter periods of time than these minutes-long task intervals, particularly given the growing literature on the dynamic nature of functional brain organization (Cohen 2018).

To conclude, we proposed a generalized, interpretable, and predictive model of attention based on brain network reconfiguration. The proposed model summarizes the functional connectivity patterns that remain stable—and those that change—across different runs of task (TC) and rest (RC). Our findings suggest that more stability and consistency during task runs and higher flexibility during rest runs contribute to better attention. Together, these findings demonstrate the differential role of functional network stability and flexibility in attention. The specificity and generalizability of the current approach is promising, suggesting implications beyond this work. Future work may seek to extend this approach to predict a wider range of cognitive and clinical traits.

## 4. Materials and Methods

Four independent data sets were used in this work. Three data sets were acquired at Yale University (Rosenberg, et al. 2015b; Rosenberg, et al. 2018; Rosenberg, et al. 2016d), and the fourth data set was from Human Connectome Project 900 Subject Release (S900) (Barch, et al. 2013; Van Essen, et al. 2013b). Details of the tasks have been published elsewhere (Barch, et al. 2013; Rosenberg, et al. 2015b; Rosenberg, et al. 2018; Rosenberg, et al. 2016d) and are described below.

### 4.1 Yale data sets

#### 4.1.1 Participants and processing

All data were acquired with participants' written consent and in compliance with procedures approved by the Yale University Human Subjects Committee. Three cohorts of subjects were used in this study (*n*-back, gradCPT, and ANT data sets), which are described in details below:

### 4.1.2 The *n*-back data set

*Participants.* Twenty-eight participants (16 females, ages 18–33 years, mean age = 25 years) were recruited from Yale University and the surrounding community. Previous work described a subset of this data set (Rosenberg, et al. 2015b); after publication six additional participants were collected with identical procedures and scan parameters. Due to excessive head motion, defined *a priori* as >2 mm translation or >3° rotation during a single run, we excluded one subject and applied the rest of our analysis to the remaining twenty-seven subjects.

*Paradigm, stimuli, and procedures.* As described previously (Rosenberg, et al. 2015a), scan sessions began with a high-resolution anatomical scan, followed by a 6-min resting-state fMRI scan, three 8.4-min *n*-back task runs, another 6-min resting scan, and a face/scene localizer scan (not analyzed here). Task stimuli consisted of face images overlaid on irrelevant distractor scenes. Participants were instructed to respond via button press to novel faces and withhold response to rare face repeats. Each task run included three experimental blocks: low load (1-back), working memory load (2-back), and perceptual load (1-back with degraded face stimuli). Performance was measured with sensitivity (*d'*), or hit rate relative to false alarm rate. *d'* was calculated separately for each task condition (1-back, 2-back, and degraded 1-back tasks). Here, we used *d'* on the 1-

back task as our primary performance measure. We focused on 1-back performance because the 2-back and degraded 1-back conditions were significantly more difficult than the gradCPT and ANT. Thus, 1-back *d′* was the most comparable to the behavioral measures in our other two data sets.

### 4.1.3 The gradCPT data set

*Participants.* Twenty-five subjects (13 females, ages 18–32 years, mean age = 22.7 years) from Yale University and the surrounding community performed a sustained attention task, the gradual-onset continuous performance task (gradCPT; (Esterman, et al. 2012; Rosenberg, et al. 2013)) during fMRI. Here we excluded seven subjects with at least one missing fMRI run and analyzed the remaining eighteen subjects.

*Paradigm, stimuli, and procedures.* As described in detail in Rosenberg et al. (Rosenberg, et al. 2016d), fMRI scan sessions included a high-resolution anatomical scan, a 6-min resting-state fMRI scan, three 13:44-min gradCPT runs, and another 6-min resting scan. During gradCPT runs, participants saw a central scene stimulus gradually transition from one image to the next at a constant rate. They were instructed to respond to each city scene but not to rare target mountains. Again, performance was assessed with sensitivity (*d′*).

### 4.1.4 The ANT data set

*Participants.* Forty-four subjects (29 females, ages 18–37 years, mean age = 23.9 years) performed an Attention Network Task (ANT) (Fan, et al. 2005; Fan, et al. 2002). The data set is described in detail in Rosenberg et al. (Rosenberg, et al. 2017b). We excluded five subjects for whom at least one of the first five fMRI task runs were missing. We further excluded one subject due to an unexpected interruption in the scanner. We then applied the rest of our analysis to the remaining thirty-eight subjects.

*Paradigm, stimuli, and procedures.* ANT scan sessions began with a high-resolution anatomical scan, followed by two consecutive 6-min resting-state runs and six 7:05-min task runs. Since eleven subjects did not have data for the sixth run, it was excluded from the analysis. The ANT requires participants to determine whether a central arrow points left or right, ignoring surrounding distractor arrows. On each trial, participants see a spatially informative cue, a spatially uninformative center cue, or no cue. On congruent trials, the central arrow points in the same direction as the four flanker arrows. On incongruent trials, the central arrow points in the opposite direction of the flankers. By comparing reaction times (RTs) on different trial types, the ANT measures three components of attention: alerting (mean RT on no-cue trials – mean RT on center-cue trials), orienting (mean RT on center-cue trials – mean RT on spatial-cue trials), and executive control (mean RT on incongruent trials – mean RT on congruent trials). Here we assessed *overall* ANT performance with RT coefficient of variation (RT CV), that is, the standard deviation divided by the mean of correct-trial RT. RT CV, a robust measure of attention (Adólfsdóttir, et al. 2008; Wojtowicz, et al. 2013), provides a summary measure of ANT performance by measuring the variability of responses across all trial types—that is, how sensitive participants are to different cue and target types (Rosenberg, et al. 2017b). In this task, RT CV was a better measure of overall performance than response accuracy, which was typically near ceiling.

### 4.1.5 Imaging parameters and Preprocessing

All cohorts were acquired with the same imaging parameters. FMRI data acquisition was performed on a 3T Siemens Trio TIM system equipped with a 32-channel head coil at the Yale Magnetic Resonance Research Center. Functional runs were acquired using a multiband gradient echo-planar imaging (EPI) sequence with the following parameters: repetition time (TR) = 1000 ms, echo time (TE) = 30 ms, flip angle = 62°, acquisition matrix = 84 × 84, in-plane resolution = 2.5 $mm^2$, 51 axial-oblique slices parallel to the ac–pc line, slice thickness = 2.5, multiband 3, acceleration factor = 2. MPRAGE parameters were as follows: TR = 2530 ms, TE = 2.77, flip angle = 7°, acquisition matrix = 256 × 256, in-plane resolution = 1.0 $mm^2$, slice thickness = 1.0 mm, 176 sagittal slices.

Data were analyzed using BioImage Suite (Joshi, et al. 2011) and custom scripts in Matlab (Mathworks). Motion correction was performed using SPM (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/). Linear and quadratic drift, mean signal from cerebrospinal fluid, white matter, and gray matter and a 24-parameter motion model (6 motion parameters, 6 temporal derivatives, and their squares) were also regressed from the data. Finally, data were temporally smoothed with a zero mean unit variance Gaussian filter. Images were warped to common space using a series of linear and nonlinear transformations as previously described (Shen, et al. 2013).

## 4.2 Human Connectome Project (HCP) Data set

### 4.2.1 Participants and processing

The HCP data set (Van Essen, et al. 2013b) includes data from 897 healthy individuals (S900) scanned during nine functional conditions (seven tasks and two rest). For details of the data acquisition parameters see Uğurbil et al. (Uğurbil, et al. 2013) and Smith et al (Smith, et al. 2013b).

Calculating TC required at least two runs of the same task. We used data from two runs of Working Memory (WM) task, with left-right (LR) and right-left (RL) phase-encodings. WM (*n*-back tasks*)* was selected as it was the closest task to the ones performed at Yale data sets. To calculate RC, we used the two phase-encodings (LR and RL) of the resting-state data collected on day 2 (REST2). All the analyses were repeated with resting-state runs on day 1 (REST1) and the same results held.

*Participants.* Analyses were restricted to n=717 subjects (392 females; age = 22–36+) for whom (i) data were available for all nine functional conditions (with left-right (LR) and right-left (RL) phase encoding), and (ii) working-memory task performance measures were available. Among the two reported HCP WM task performance measures—accuracy and median reaction time—we used the average of the median reaction time across all trials (Median RT), because the response accuracy was near ceiling for most subjects.

*Paradigm, stimuli, and procedures.* Functional MRI scans were acquired during two different days: Day 1, part of which was used here, included two runs (LR and RL) of the working memory (WM) task (5:01 min per run), incentive processing (gambling) task (3:12 min), motor task (3:34 min), and rest (14:33 min). The details of task design have been previously described (Barch, et al. 2013; Van Essen, et al. 2013b). Here, we provide a brief description of WM task and an overview of the relevant aspects. In this task, participants performed a visual *n*-back task, with blocked 0-back and 2-back conditions using four stimulus categories (faces, places, tools, body parts). Each run consisted of 8 task blocks (10 trials each), with each stimulus category used twice, and 4 fixation blocks. Each block started with a 2.5 s cue indicating the task type (0-back versus 2-back) and the target (for 0-back).

### 4.2.2 Imaging parameters and Preprocessing

The HCP minimal preprocessing pipeline was employed (Glasser, et al. 2013), which includes artifact removal, motion correction and registration to MNI space. Further preprocessing steps were performed using BioImage Suite (Joshi, et al. 2011) and included standard preprocessing procedures (Finn, et al. 2015b) including regressing 24 motion parameters, regressing the mean time courses of the white matter and cerebrospinal fluid as well as the global signal, removing the linear trend, and low pass filtering.

## 4.3 Functional network construction

Functional connectivity matrices were assessed using a functional brain atlas (Shen, et al. 2013) consisting of 268 nodes covering the whole brain; this atlas was defined using resting-state data from a separate population of healthy subjects (Finn, et al. 2015b). The Pearson correlation coefficients between the time courses of each possible pair of nodes were calculated and normalized using Fisher's z transformation to construct 268×268 symmetrical connectivity matrices. This was done for each subject and for each run separately.

### 4.4 Higher-order network measures: task correlation (TC) and rest correlation (RC)

Network strength vectors were calculated from the connectivity matrices by summing over all the edge values incident to a node and taking the absolute value of the sum, yielding a 268×1 strength vector. Each element $i$ in the strength vector indicates the connectivity strength between node $i$ and the rest of the brain. Let us denote the matrix including the strength vectors computed during all task runs $V_{Task}$ and those computed during rest runs $V_{Rest}$. Higher order measures of connectivity were computed by taking the Pearson correlation between every pair of strength vector in $V_{Task}$ (denoted as Task Correlation or TC; see Eq. 2) and every pair of strength vector in $V_{Rest}$ (denoted as Rest Correlation or RC; see Eq. 3).

$$TC = corr\ (V_{Task}) = corr\left(\begin{bmatrix} d_{1,1} & \cdots & d_{1,T} \\ \vdots & \ddots & \vdots \\ d_{268,1} & \cdots & d_{268,T} \end{bmatrix}\right), \tag{2}$$

$$RC = corr\ (V_{Rest}) = corr\left(\begin{bmatrix} d_{1,1} & \cdots & d_{1,R} \\ \vdots & \ddots & \vdots \\ d_{268,1} & \cdots & d_{268,R} \end{bmatrix}\right), \tag{3}$$

Where T is the total number of task runs and R is the total number of rest runs. For the purpose of predictive modeling, we took the average of the elements in TC and RC.

### 4.5 Predictive modeling pipeline

We used linear regression as our learning algorithm to predict behavioral scores. To make the predictive model generalizable across different data sets with different number of task and rest runs, we considered the average of TC elements ($TC_{avg}$) and RC elements ($RC_{avg}$) as our predictive features, resulting in a 2-dimensional model as shown in Eq. 4.

$$y = \beta_0 + \beta_{TC} \times TC_{avg} + \beta_{RC} \times RC_{avg}, \tag{4}$$

where y is the behavioral score, $\beta_0$ is the bias, and $\beta_{TC}$ and $\beta_{RC}$ are the model's coefficients. All coefficients (including bias) are learned in a fully cross-validated setting both within and across different data sets.

### 4.6 Within-data set predictive modeling

We performed a leave-one-out cross-validation analysis on each data set. At each step, we left data from one individual out of our training set. We trained our model with the remaining n-1 individuals, where n is the size of the data set. Then, we tested the result for the left-out subject. We repeated this for every individual until we obtained predicted behavioral scores for all n subjects.

P-values for the leave-one-out cross-validation were calculated using non-parametric permutation test. This is because the traditional r-to-p conversions assume the degree of freedom is equal to n-2, where n is the number of subjects for each data set. However, analyses in the leave-one-out folds are not independent, so the number of degrees of freedom is overestimated. To perform permutation test, we randomly permuted the behavioral scores 100 times and each time ran the permuted scores through the predictive pipeline and calculated the correlation between the predicted values and the permuted scores.

### 4.7 Cross-data set predictive modeling

External generalizability of the model was examined by employing cross-data set predictive analyses, i.e., external validation. This was performed by training a linear model on participants of one data set and testing it to predict the behavioral scores for the participants of another data set. For Yale data, some of the participants overlapped across data sets. To avoid non-independency of samples, we removed the overlapping subjects from the test set and left them in the training set.

*N-back data set.* A linear model was trained on the subjects of *n*-back data set (n=27) using 1-back d' values as output. The model was then used to predict d' values for the subjects in the gradCPT data set (n=14, after removing 4 overlapping subjects), ANT data set (n=33, after removing 5 overlapping subjects), and HCP data set (n=717). Note that for the participants of ANT and HCP we did not have the true d' values, instead, we plotted the predicted d' values versus the observed RT CV (for ANT) and Median RT (for HCP), both of which are proxies for attention performance (Rosenberg, et al. 2017b) and are inherently inversely correlated with d' values. That is, higher d' values and lower statistics of reaction time are both indicative of better attentional abilities.

*GradCPT data set.* A linear model was trained on the subjects of gradCPT data set (n=18) using d' values as output. The model was then used to predict d' values for the subjects in the *n*-back data set (n=23, after removing 4 overlapping subjects), ANT data set (n=35, after removing 3 overlapping subjects), and HCP data set (n=717).

*ANT data set.* A linear model was trained on the subjects of ANT data set (n=38) using the coefficient of variation of the correct-trial reaction time (RT CV) as output. The model was then used to predict RT CV values for the subjects in the *n*-back data set (n=22, after removing 5 overlapping subjects), gradCPT data set (n=15, after removing 3 overlapping subjects), and HCP data set (n=717).

*HCP data set.* A linear model was trained on the subjects of HCP data set (n=717) using the median correct-trial reaction time (Median RT) as output. The model was then used to predict Median RT values for the subjects in the *n*-back data set (n=27), gradCPT data set (n=18), and ANT data set (n=38).

Additionally, it is important to note that the order of rest and task runs were different across the four data sets. That is, for *n*-back and gradCPT data sets rest runs were separated by task performance, whereas for ANT and HCP data sets rest runs were performed successively before task performance. Nonetheless, the predictive models generalized across all four data sets.

### 4.8 Analysis of coefficients in the predictive model

To safely interpret the coefficients of the predictive model, we first ruled out the possibility of multicollinearity. To this end, we calculated variance of inflation factor (VIF), which measures how much the variance of an estimated regression coefficient increases because of collinearity. In practice, there is typically a small amount of collinearity among the predictors. It is generally believed that a VIF value that exceeds 5 or 10 indicates a problematic amount of multicollinearity. The significance of the relationship between each predictor and the output was assessed by one-tailed t-tests (using the Pr(>|t|) component in R's linear model).

## 5. Code availability

The 268-node functional parcellation is available online on the BioImage Suite NITRC page (https://www.nitrc.org/frs/?group_id=51). MATLAB and R scripts were written to perform the analyses described; this code can be found at https://github.com/YaleMRRC/TC_RC

## 6. Supplemental Materials

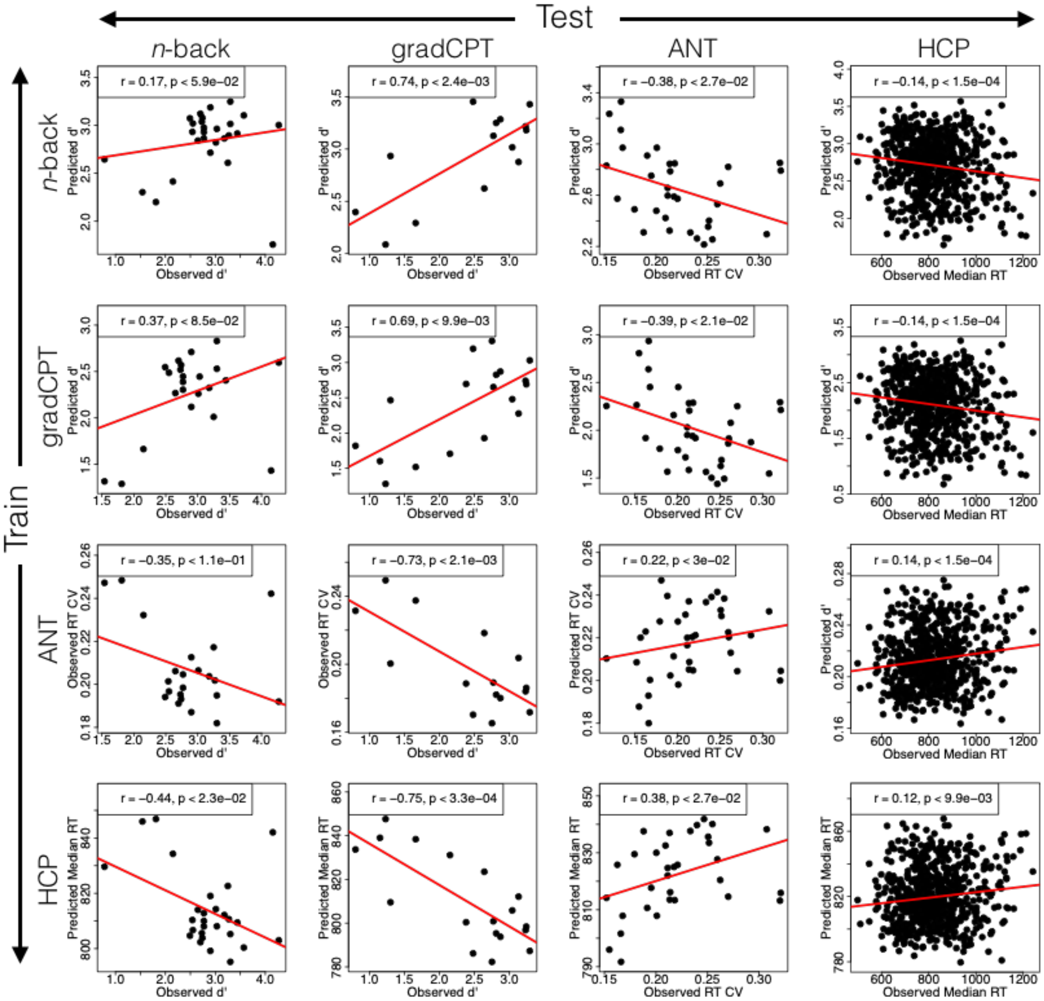

**Figure S1. Within and cross-data set predictive analysis using only TCavg as feature in the linear model.** The predicted behavioral measures versus observed ones are displayed for within-data set (diagonal elements) and cross-data set (off-diagonal elements) predictions. Comparing to the 2-dimensional linear model on both $\mathrm{TC}_{\mathrm{avg}}$ and $\mathrm{RC}_{\mathrm{avg}}$ (Figure 3 in the main manuscript), this mode model performs less significantly in most 2cases, despite its lower dimension (which causes lower variance error and lower chance of overfitting). This observation highlights the unique and meaningful information that RC could add to the model.

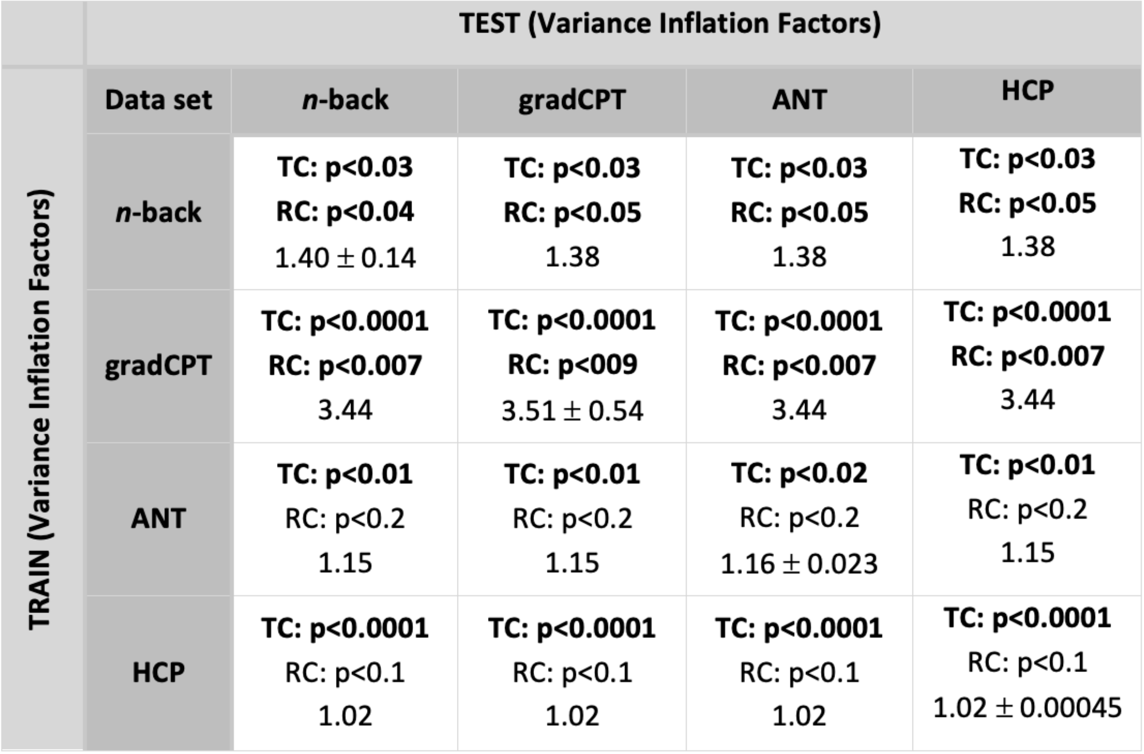

| | | **TEST (Variance Inflation Factors)** | | | |
|---|---|---|---|---|---|
| **TRAIN (Variance Inflation Factors)** | **Data set** | ***n*-back** | **gradCPT** | **ANT** | **HCP** |
| | ***n*-back** | **TC: p<0.03**<br>**RC: p<0.04**<br>1.40 ± 0.14 | **TC: p<0.03**<br>**RC: p<0.05**<br>1.38 | **TC: p<0.03**<br>**RC: p<0.05**<br>1.38 | **TC: p<0.03**<br>**RC: p<0.05**<br>1.38 |
| | **gradCPT** | **TC: p<0.0001**<br>**RC: p<0.007**<br>3.44 | **TC: p<0.0001**<br>**RC: p<009**<br>3.51 ± 0.54 | **TC: p<0.0001**<br>**RC: p<0.007**<br>3.44 | **TC: p<0.0001**<br>**RC: p<0.007**<br>3.44 |
| | **ANT** | **TC: p<0.01**<br>RC: p<0.2<br>1.15 | **TC: p<0.01**<br>RC: p<0.2<br>1.15 | **TC: p<0.02**<br>RC: p<0.2<br>1.16 ± 0.023 | **TC: p<0.01**<br>RC: p<0.2<br>1.15 |
| | **HCP** | **TC: p<0.0001**<br>RC: p<0.1<br>1.02 | **TC: p<0.0001**<br>RC: p<0.1<br>1.02 | **TC: p<0.0001**<br>RC: p<0.1<br>1.02 | **TC: p<0.0001**<br>RC: p<0.1<br>1.02 ± 0.00045 |

**Table S1. Multicollinearity analysis using variance of inflation factor (VIF).** To safely interpret the coefficients of the predictive model, we first ruled out the possibility of multicollinearity. To this end, we calculated variance of inflation factor (VIF), which measures how much the variance of an estimated regression coefficient increases because of collinearity. VIF values are reported for each pair of training and testing data sets. For within-data set predictions, the mean and standard deviation across all leave-one-out models are reported. The smallest possible value for VIF is 1, which indicates the complete absence of collinearity. Typically in practice, there is a small amount of collinearity among the predictors. As a rule of thumb, a VIF value that exceeds 5 or 10 indicates a problematic amount of multicollinearity. The significance of the relationship between each predictor and the output is assessed by one-tailed t-tests (using the Pr(>|t|) component in R's linear model). The significant values are displayed in bold.

# Chapter 3: An Exemplar-based Approach to Individualized Parcellation Reveals the Need for Sex Specific Functional Networks

## Abstract

Recent work with functional connectivity data has led to significant progress in understanding the functional organization of the brain. While the majority of the literature has focused on group-level parcellation approaches, there is ample evidence that the brain varies in both structure and function across individuals. In this work, we introduce a parcellation technique that incorporates delineation of functional networks both at the individual- and group-level. The proposed technique deploys the notion of "submodularity" to jointly parcellate the cerebral cortex while establishing an inclusive correspondence between the individualized functional networks. Using this parcellation technique, we successfully established a cross-validated predictive model that predicts individuals' sex, solely based on the parcellation schemes (i.e. the node-to-network assignment vectors). The sex prediction finding illustrates that individual parcellation of functional networks can reveal subgroups in a population and suggests that the use of a global network parcellation may overlook fundamental differences in network organization. This is a particularly important point to consider in studies comparing patients versus controls for example or even patient subgroups. Network organization may differ between individuals and global configurations should not be assumed. This

approach to the individualized study of functional organization in the brain has many implications for both neuroscience and clinical applications.

## 1. Introduction

The human brain is functionally segregated into multiple spatially-distributed networks, and how best to divide or parcellate the brain into these networks is a fundamental question for neuroscience (Power, Cohen et al. 2011, Yeo, Krienen et al. 2011, Yang, Fan et al. 2016). Resting-state functional magnetic resonance imaging (fMRI) studies have consistently identified a number of brain networks that replicate across different datasets (Power, Cohen et al. 2011, Yeo, Krienen et al. 2011) and overlap with task activation patterns (Smith, Fox et al. 2009). The spatial organization of these networks is thought to support a wide range of cognitive functions (Dosenbach, Fair et al. 2007, Laird, Fox et al. 2011), and such networks have been shown to be altered in clinical disorders (Bush 2011, Stern, Fitzgerald et al. 2012, Zhu, Wang et al. 2012).

The majority of previous work on parcellating the brain into networks has been focused on group-level analyses (Power, Cohen et al. 2011, Yeo, Krienen et al. 2011, Shen, Tokoglu et al. 2013, Gordon, Laumann et al. 2014) with the aim of defining a set of networks that generalizes over all individuals. Group-level analysis is typically accomplished by collapsing data from individuals, either by averaging the subject's connectivity matrices (Power, Cohen et al. 2011, Yeo, Krienen et al. 2011) or by concatenating time courses from each subject, as in the case of Independent Component Analysis (ICA) (Beckmann, DeLuca et al. 2005, Smith, Fox et al. 2009). As a result, these approaches do not preserve information regarding inter-individual variability.

Nevertheless, emerging studies have highlighted the importance of inter-individual variability in functional connectivity in contributing to individual differences in behavior and cognition (Van Horn, Grafton et al. 2008, Baldassarre, Lewis et al. 2012, Mueller, Wang et al. 2013, Zilles and Amunts 2013, Calluso, Tosoni et al. 2015, Finn, Shen et al. 2015, Smith, Nichols et al. 2015, Finn and Constable 2016, Rosenberg, Finn et al. 2016). Such inter-individual variability in functional connectivity is likely to be expressed at the network level and thus should be revealed by functional parcellation schemes.

Accordingly, individual-level parcellation of the brain into networks has recently received increased attention. To enable functional network parcellation at the individual-level, one plausible approach is to apply a back-projection from the group-level parcellation. This approach has been prevalent in ICA studies; and techniques such as principle component analysis (PCA) back-projections (Calhoun, Adali et al. 2001) and GLM dual regression approaches (Beckmann, Mackay et al. 2009) have been developed. However, studies have reported notable limitations for ICA approaches at the individual-level (Zuo, Kelly et al. 2010), including shortcomings to address inter-subject variation, limitations in scaling to higher dimensions (i.e. finer grained parcellations), and high sensitivity to artifacts such as motion, scanner noise, and physiological noise (McKeown, Jung et al. 1998, McKeown, Hansen et al. 2003, Cole, Smith et al. 2010). To reduce the impact of ICA limitations in addressing inter-subject variability, extensions of this method such as independent vector analysis (IVA) have been proposed (Lee, Lee et al. 2008, Michael, Anderson et al. 2014). While promising, IVA is highly sensitive to each individual data and suffers from excessive computational burden and memory requirements (Michael, Anderson et al. 2014). More recently, studies have used functional connectivity, as derived from BOLD fMRI, to establish individualized networks (Eickhoff, Thirion et al. 2015), using techniques such as *k*-means (Flandin, Kherif et al. 2002, Kahnt, Chang et al. 2012), hierarchical clustering (Bellec, Rosa-Neto et al. 2010, Meunier, Lambiotte et al. 2010, Moreno-Dominguez, Anwander et al. 2014, Arslan and Rueckert 2015), spectral clustering (Thirion, Flandin et al. 2006, Van Den Heuvel, Mandl et al. 2008, Craddock, James et al. 2012, Chen, Li et al. 2013, Shen, Tokoglu et al. 2013), and boundary mapping (Cohen, Fair et al. 2008, Gordon, Laumann et al. 2014, Laumann, Gordon et al. 2015). Although many of these approaches are promising, none of them provide a unified framework that incorporates joint individual- and group-level functional network parcellations with a comprehensive correspondence across the identified networks.

Wang et al. parcellated resting-state fMRI data into a number of coherent networks using an iterative parcellation approach that requires an initial group-level parcellation as a reference (Wang, Buckner et al. 2015). Their approach requires this initiation step and thus cannot be used when there is no representative group-level parcellation. Similarly, Shen et al. provided a joint individual- and group-level parcellation approach through optimization of a rotation function derived from individualized functional connectivity (Shen, Tokoglu et al. 2013). This approach, however, requires the same dataset for the group- and individual-level parcellations, and thus does not provide a generalizable parcellation scheme that can be used across datasets.

Here we develop a comprehensive parcellation framework that overcomes the above concerns through a three-step flexible pipeline. The proposed method exploits “exemplar-based clustering” that seeks to summarize the massive amount of data using a relatively small number of representative exemplars (Dueck and Frey 2007, Badanidiyuru, Mirzasoleiman et al. 2014, Mirzasoleiman, Karbasi et al. 2016). Using “exemplars” provides a flexible one-to-one mapping of the functional networks across subjects, easing localization of inter-individual variability over the cortex. Moreover, an intuitive notion of diminishing returns, known as "submodularity", is utilized to provide an efficient optimization algorithm with provable bounds (Nemhauser, Wolsey et al. 1978). Unlike many other individual-level parcellations that are initiated from a group-level parcellation scheme to derive the corresponding functional networks for individuals (Zuo, Kelly et al. 2010, Gordon, Laumann et al. 2015, Wang, Buckner et al. 2015, Gordon, Laumann et al. 2016), our method moves a step forward by initiating from the individual data. We show that this approach has a higher sensitivity to individual variations and thus provides the basis for more powerful inferences. We evaluate our parcellation approach using clustering validation measures of stability and reproducibility. Finally, we compare our method with the two individual-level parcellations mentioned above – Shen. et al. (Shen et al., 2013) and Wang et al. (Wang et al., 2015) – in two different aspects: (1) internal clustering evaluation, and (2) sensitivity to inter-individual variations (i.e. predictive power). Of note, although there exists potentially interesting individual variability in functional organization both at the node- and network-levels, the focus of this initial work is to delineate the network-level organization. It should be noted, however, that the approach described here can be applied to voxel-level data in order to define a node-level functional atlas.

## 2. Theory

### 2.1 Overview

Exemplar-based clustering algorithms summarize massive datasets through the selection of a relatively small set of representative exemplars. Our proposed algorithm seeks to select K exemplar regions (representing our networks) across the cerebral cortex. The clustering algorithm then assigns each of the nodes to one of the exemplars, i.e., one of the networks.

Most techniques for identifying exemplars define an objective function that measures the "representativeness" of each set of exemplars with regard to the full dataset. Often, these objective functions satisfy an intuitive notion of diminishing returns called submodularity (Nemhauser, Wolsey et al. 1978): for instance, if given two sets of exemplars $S_1$ and $S_2$ with $S_1 \subseteq S_2$, adding a new element to $S_1$ is more beneficial than adding it to the super set, $S_2$, as the new element can potentially add more information to $S_1$ rather than $S_2$. When using this concept of submodularity, the problem of finding K exemplars can be reduced to maximizing a non-negative monotone submodular set function subject to a cardinality constraint (i.e., a bound on the number K of elements that can be selected) (Krause and Golovin 2012, Mirzasoleiman, Karbasi et al. 2016). Simple greedy algorithms can efficiently maximize these objective functions (Nemhauser, Wolsey et al. 1978). See (Mirzasoleiman, CH et al. 2016) for recent developments of submodular maximization methods.

In the following, we formally define our submodular function following the work of Krause et al. (Krause and Golovin 2012), and define the greedy algorithm, and exemplar-based clustering. We subsequently present our algorithm and the details of our implementation.

## 2.2 Submodular functions

Submodularity is a property of set functions, i.e., functions $f\colon 2^V \rightarrow R$ that assign each subset $S \subseteq V$ a value $f(S)$. Here, $V$ is a finite set, commonly called the ground set, and $S$ is a finite subset of $V$. The definition of submodularity relies on a notion of discrete derivative, also called the marginal gain. An important subclass of submodular functions (used in the proposed algorithm) are those which are monotone.

**Definition 1.1 (Discrete derivative)** For a set function $f\colon 2^V \rightarrow R$, $S \subseteq V$, and $e \in V$, let $\Delta f(e|S) := f(S \cup \{e\}) - f(S)$ be the discrete derivative of $f$ at $S$ with respect to $e$.

**Definition 1.2 (Submodularity)** A function $f\colon 2^V \rightarrow R$ is submodular if for every $A \subseteq B \subseteq V$ and $e \in V \backslash B$ it holds that $\Delta(e|A) \geq \Delta(e|B)$. Meaning that adding an element $e$ to a set $A$ increases the utility at more than (or at least equal to) adding it to $A$'s superset, $B$, suggesting a natural diminishing returns.

**Definition 1.3 (Monotonicity)** A function $f : 2^V \rightarrow R$ is monotone if for every $A \subseteq B \subseteq V$, $f(A) \leq f(B)$. Equivalently, function $f$ is monotone if and only if all its discrete derivatives are nonnegative, i.e., for every $A \subseteq V$ and $e \in V$ it holds that $\Delta(e|A) \geq 0$.

### 2.3 The greedy algorithm for optimization of the submodular function

In general, maximizing a non-negative monotone submodular function subject to a cardinality constraint, i.e.,

$$\max_{S \subseteq V} f(S) \;\; s.t. \;\; |S| \leq K, \tag{1}$$

is NP-hard (Feige 1998). However, a seminal result of (Nemhauser, Wolsey et al. 1978) proves that a simple greedy algorithm provides the best approximation (≈ 63%) to the optimal solution. In practice, this approximation is significantly closer to the optimal solution (see Supplementary Materials and Figure S5 for an empirical evaluation). The greedy algorithm starts with an empty set $S_0 = 0$, and at each iteration $i$, it selects and adds the element $\{e_i^*\} \in V$ such that the marginal gain is maximized, i.e.,

$$e_i^* = \arg\max_{\mathrm{e} \in \mathrm{V}} \Delta f(e|S_{i-1}) \coloneqq \arg\max_{\mathrm{e} \in \mathrm{V}} f(S_{i-1} \cup \{e\}) - f(S_{i-1}), \tag{2}$$

$$S_i = S_{i-1} \cup \{e_i^*\}. \tag{3}$$

The algorithm continues until the cardinality constraint is reached, i.e., until $|S| = K$.

### 2.4 Exemplar-based clustering

Exemplar-based clustering provides an approach to summarize the data by introducing a set of $K$ exemplars that best represents the full dataset. A classic way of identifying such exemplars is solving the $k$-medoids problem, by minimizing the sum of pairwise distances between the elements of the dataset and the exemplars (see Friedman et al. (Friedman, Hastie et al. 2001) for more details on $k$-medoid problems). Specifically, assume we are given a dissimilarity function $d: V \times V \to R$, where $d$ encodes the dissimilarities between the elements of the ground set $V$. The $k$-medoid problem minimizes the following loss function:

$$L(S) = \frac{1}{|V|} \sum_{v \in V} \min_{e \in S} d(v, e). \tag{4}$$

$L(S)$ measures how much information we lose if we represent all the data points in each cluster, with its corresponding exemplar.

By introducing an appropriate auxiliary element $v_0$, we can turn $L$ into a monotone submodular function, so that the minimization of (4) is equivalent to the maximization of the following monotone submodular function (5), and can be efficiently solved by the greedy algorithm:

$$f(S) = L(v_0) - L(S \cup v_0). \tag{5}$$

Technically, any vector $v_0$ satisfying the following condition can be used as an auxiliary exemplar:

$$\max_{v' \in V} d(v, v') \leq d(v, v_0), \ \ \forall v \in V \backslash S. \tag{6}$$

This condition implies that the distance between the auxiliary element and all of the data points must be greater than the pairwise distances between the data points.

Note that in contrast to the classical clustering algorithms (such as $k$-means), the exemplar-based clustering is very general in that it does not require the distance function $d$ to be symmetric nor to obey triangle inequality. All it requires for $d$ is nonnegativity. Here we used the squared Euclidean distance as the dissimilarity function:

$$d(x, x') = \|x - x'\|^2. \tag{7}$$

Herein, we utilize the submodularity of our utility function further to implement an accelerated version of the greedy algorithm, called lazy greedy (Minoux 1978).

### 2.5 Cortical parcellation algorithm

In this section, we deploy the aforementioned algorithm to parcellate the cerebellar cortex into $K$ functional networks. For each individual $j \in \{1, \dots, J\}$, we have a matrix $V^j_{N \times T}$, where $N$ denotes the number of regions in the brain and $T$ represents the number of time points. Each region $n \in \{1, \dots, N\}$ of the brain, forms a vector in a $T$-dimensional space, denoted as $v^j_n$. We aim to find $K$ exemplar labels $S = \{e_1, e_2, \dots, e_K\}$ whose corresponding exemplar set for each individual $j$, i.e., $S^j = \left\{v^j_{e_1}, v^j_{e_2}, \dots, v^j_{e_K}\right\} \subseteq V^j$, maximizes a desired utility function. In order to jointly consider the information of each individual and the group, we define a natural objective utility function as follows:

$$F(S) = \sum_{j=1}^{J} f^{j}(S^{j}), \tag{8}$$

where $S$ is the exemplar label set and $S^j$ is the set including the corresponding exemplar vectors in individual $j$. In addition, $f^j(S^j)$ is the utility function of individual $j$ defined according to equations (2-5), and $J$ is the total number of subjects. Note that submodularity is preserved under non-negative linear combination and thus $F(S)$ remains a non-negative monotone submodular function that can similarly be optimized by the greedy algorithm. Also note that $f^j$ is a function that is locally defined for each individual $j$, meaning that it takes the label set $S = \{e_i\} \in \{1, \dots, N\}$ of regions and considers the corresponding vectors in each individual. The algorithm finally selects the $K$ exemplar labels for which the corresponding exemplar vectors in each individual minimize the sum of loss functions over all individuals. After $K$ exemplars are obtained for each individual, the algorithm assigns each region $n$ in individual $j$ (i.e. vector $v_n^j$) to the closest exemplar, i.e.,

$$Exemplar\left(v_n^j\right) = \arg\min_{\mathrm{e_i} \in S} (v_n^j, v_{e_i}^j). \tag{9}$$

Thus, the brain is parcellated into $K$ networks each represented by an exemplar. In order to obtain the group-level parcellation, we employ a majority vote algorithm over all subjects. In other words, region $n$ is assigned to network $k$ if the majority of individuals vote for this assignment.

Overall, the proposed algorithm operates in three steps. First, the exemplar-search step finds the global exemplars over all subjects. Second, the individual-clustering step parcellates each individual's brain by greedily maximizing a utility function, defined according to the group data. Third, the group-clustering step takes the majority vote of all individual clusters. The pseudocode in Figure 1 shows each step in more detail.

**Exemplar-based cortical parcellation algorithm – individual and group approach**

**Input**

- BOLD signals for all subjects $V = \{V^1, \dots, V^J\}$, where $V^j_{N\times T} \in V$ contains $N$ normalized and zero-centered time-courses in a $T$-dimensional space.
- $K$: Number of networks.

**Algorithm:**

```
1) Exemplar-search
Initialize: $S_0 \leftarrow \emptyset$
Initialize: $v_0 \leftarrow [3,0,0,\dots,0]^T$
for $i = 1\ to\ K$ do
    $e_i^* = \arg\max_{e_i} \Delta F(v_{e_i} | S_{i-1}) := \arg\max_{e_i} \sum_{j=1}^{J} f^j(S_{i-1}^j \cup \{v_{e_i}^j\}) - f^j(S_{i-1}^j)$,
    where:
            $f^j(S_i^j) = L^j(v_0) - L^j(S_i^j \cup v_0)$,
            $L^j(S_i^j) = \frac{1}{|V^j|} \sum_{n=1}^{N} \min_{e \in S_i} d(v_n^j, v_e^j)$,
            $d(v_n^j, v_e^j) = \|v_n^j - v_e^j\|^2$
    $S_i = S_{i-1} \cup \{e_i^*\}$
Return $S_K$

2) Individual-clustering
for $j = 1\ to\ J$ do
    $S^j = V^j_{S_K} = \{v^j_{e_1}, v^j_{e_2}, \dots, v^j_{e_K}\}$
    for $n = 1\ to\ N$ do
            $Exemplar(v_n^j) = \arg\min_{e_i \in S}(v_n^j, v_{e_i}^j) \longrightarrow Cluster(v_n^j) = \arg\min_{i \in \{1,\dots,K\}}(v_n^j, v_{e_i}^j)$
Return: $Cluster(v_n^j)\ \forall n \in N, j \in J$
It returns the individual-level parcellations.

3) Group-clustering (Majority voting)
for $n = 1\ to\ N$ do
    $Cluster(v_n) = \arg\max_{i} \sum_{j=1}^{J} \mathbf{1}_{\{Cluster(v_n^j) = i\}}$
    where $\mathbf{1}$ represents the indicator function, i.e.:
            $\mathbf{1}_{\{Cluster(v_n^j) = i\}} = \begin{cases} 1, & \text{if } Cluster(v_n^j) = i \\ 0, & \text{if } Cluster(v_n^j) \neq i \end{cases}$
Return: $Cluster(v_n)$
It returns the group-level parcellation.
```

**Figure 1. Pseudocode explaining the three steps of exemplar-based parcellation.** In the first step (exemplar-search), K=2, …, 25 exemplars are derived for each individual with a group constraint, i.e. by greedily optimizing a nonnegative monotone submodular function defined as the summation of the utility function over individuals. In the second step (individual-clustering), for each single individual, every node in the cortical area is assigned to its closest exemplar, where closeness is defined using a squared Euclidean distance function. Finally, in the third step (group-clustering), the group-level parcellation is derived by majority voting over all individual-level parcellations (i.e. the node-to-network assignment vectors).

One significant advantage of this algorithm is that there is a straightforward mapping between the parcellation of each individual to each other and to the group, as each network is represented by a global exemplar. Thus, we do not require another algorithm to retrieve the correspondences. This facilitates direct comparison between individuals and with the group.

## 3. Material and methods

### 3.1 Participants and processing

Data were obtained from the 900 subject release dataset in Human Connectome Project (HCP) (Van Essen, Smith et al. 2013). Analysis was limited to 825 subjects for which the complete scan data were available for each of the two resting states: REST1 and REST2. For details of scan parameters, see Uğurbil et al. (Uğurbil, Xu et al. 2013) and Smith et al. (Smith, Beckmann et al. 2013). Starting with the minimally preprocessed HCP data (Glasser, Sotiropoulos et al. 2013), further preprocessing steps were performed using BioImage Suite (Joshi, Scheinost et al. 2011) and included regressing 12 motion parameters (Movement_Regressors_dt.txt), regressing the mean time courses of the white matter and cerebrospinal fluid as well as the global signal, removing the linear trend, and low pass filtering (as previously described in (Finn, Shen et al. 2015)). We employed a functional brain atlas (Shen, Tokoglu et al. 2013) consisting of 188 nodes covering the cortex of the brain. This atlas was defined on a separate population of healthy subjects (Finn, Shen et al. 2015).

### 3.2 Functional distance matrix

Time courses from the two resting-state conditions (REST1 and REST2) and the two functional runs with opposing phase-encoding directions (left-right, "LR", and right-left, "RL") were concatenated and further used to generate a ground set consisting of $N$ vectors in $T$-dimensional space. All the data points were normalized to a unit norm sphere centered at origin, and a point with the norm greater than two was used as the auxiliary exemplar (Eq. 6 in the Theory Section 2.4). For each individual, the pairwise squared Euclidean distances between the data points were calculated, and a matrix of size $188 \times 188$ was obtained. Next, the greedy algorithm was employed to find the best $K = \{2, 3, \dots, 25\}$ exemplars according to the algorithm described above.

### 3.3 Stability and convergence

We examined the stability and convergence behavior of our group-level parcellation. Stability was defined as the robustness of the output to slight perturbations to the input (Von Luxburg 2010), which was examined here in terms of the variations both in the group size and the selection of subsets of individuals from the larger group. Convergence was examined through the rate at which the output parcellation converged to the final result as the input merged to span the entire dataset. We started with twenty-five subjects and incremented the number of subjects used in the network parcellation in steps of twenty-five. At each step $t$, we employed the exemplar-search algorithm (part 1 in Algorithm) over the set of $25 \times t$ subjects, which we refer to as training set herein. Using the exemplars derived from the training set, we applied the individual-clustering algorithm (part 2 in Algorithm) over both the training set (i.e. $25 \times t$ subjects) and the entire dataset (i.e. 825 subjects), obtaining two sets of individual-level parcellations. Next, we employed majority voting (part 3 in Algorithm) over both the training set and the entire dataset, and then calculated the Hamming distances (Hamming 1950) between the group-level parcellation derived from the training set (which here is called the perturbed parcellation) and the full parcellation derived from the entire dataset (Figure 3).

### 3.4 Reproducibility of group-level parcellation

We investigated whether our proposed method was generalizable across different sets of subjects using two different pipelines (Figure 4B). In the first pipeline, the dataset consisting of 800 subjects was split into two equal size subsamples, and the three-step parcellation algorithm – including exemplar-search, individual-clustering, and group-clustering – was applied on each half independently (Figure 4B, Right). Finally, the overlap between the two group-level parcellation schemes was calculated using the Dice coefficient (Dice 1945). The process was repeated 100 times for different permutations of subjects (Figure 4A, blue error bars). In the second pipeline, the data set was similarly split into two equal size subsets, but this time, a training-testing strategy was utilized. We employed the first part of the algorithm (i.e. exemplar-search) over group 1 (referred to as the training set). Next, using the exemplars derived from the training set, we ran the rest of the algorithm (i.e. individual-clustering and group-clustering) over the individuals in group 2 (referred to as the testing set) as well as group 1 (the training set), obtaining two group-level parcellation schemes (Figure 4B, Left). The overlaps between the two parcellations were computed using the Dice coefficient. The same process was repeated 100 times for different permutations of subjects (Figure 4A, orange error bars). Note that in the second pipeline, there is a direct one-to-one mapping between the two parcellation schemes, through their common exemplars. In other words, all the regions with the same exemplar labels, are assigned to the same clusters (networks). This straightforward mapping across functional networks at individual-level and group-level is a unique advantage of the exemplar-based clustering. It also provides a cross-validation approach to the parcellation schemes through training-testing settings.

### 3.5 Reproducibility of individual-level parcellations across rest sessions

To investigate the reproducibility of parcellations at the individual-level, we repeated the parcellation analyses, this time taking into account the data from REST1 and REST2 scan sessions separately. As with the group-level reproducibility analysis, we employed two separate pipelines: First, we computed the individualized parcellations for each rest session, using the previously computed exemplars that were derived from the joint consideration of the two sessions. The advantage of this approach is that it preserves the correspondences between the resulting networks for each individual across the two sessions (referred to as the global exemplars). In the second pipeline, we recalculated the exemplars for the two rest sessions independently, and used them to parcellate the individuals within each session. We refer to these as the local exemplars. For each pipeline, we calculated the Dice coefficients between the parcellation results of every individual in the two rest sessions: REST1 and REST2 (Figure S2).

### 3.6 Mapping highly variable regions

For each node in the cortex, we investigated the number of individuals that voted for the appointed network in the group-level parcellation; a measure labeled as F1 (or the frequency of the 1st mode). This measure captures the consistency of the node-to-network assignments across individuals, and thus , the inverse of F1 (1/F1) could be an indicator of the inter-individual variability. Another metric of interest in the literature is the frequency of the 2nd mode (known as F2), i.e., the number of occurrences for the second most frequent network assignment. To further address the confidence of the node-to-network-assignments across all individuals, the ratio between F1 and F2 (i.e. F1:F2 ratio) was calculated. Similarly, to underscore the variability of regions, the inverse ratio (i.e. F2:F1 ratio or F2/F1) was considered. A high value for the inverse F1 and the F2:F1 ratio reflects greater variability in the network assignment. For each node, the two inter-individual variability measures (1/F1, F2/F1) were calculated and summed up over the number of networks ranging from K=2 to K=25, then the resulting numbers were scaled to the range (0, 100) (Figure 5).

### 3.7 Sex-prediction

To illustrate that this individualized parcellation approach provides meaningful information, we next demonstrated a data-driven predictive model based on parcellation (i.e. node-to-network assignments) to predict the sex for each individual. We used a gradient boosting machine (GBM) with 100 estimators (also known as decision trees) and 0.05 learning rate (see the code for more details on parameters) in a ten-fold cross-validated setting (Friedman, Hastie et al. 2001). Each time, we fed the predictive model with node-to-network-assignment vectors for the individuals in the training set as features, and their corresponding sex as output. We predicted sex for the unseen fold of data across a varying number of networks from K=2 to K=25. The reported accuracies are the mean and standard deviations across all ten folds (Figure 6A, blue error bars). To confirm that our prediction results were highly significant, we applied a nonparametric permutation testing by generating a null distribution via randomly shuffling the outputs (i.e. sex) 100 times and running the generated vectors through our predictive model (Figure 6A, orange error bars).

Of note, since we initially defined the network parcellations across all individuals and then used the same individuals for the sex prediction, these were not two independent samples. It is unlikely that this dependency has confounded the results, for two main reasons: first, the parcellation step was employed agnostic to the individuals' sex. That is, the same parcellation algorithm was employed on both male and female subjects, with no prior knowledge on their sex. Second, the employed predictive model (GBM) is a non-parametric model with no sensitivity to the dependency of samples. Nevertheless, we tested for the potential biases by employing the parcellation and the prediction steps on two independent subsets. In one analysis, we split the entire population into two equal-size sets (each with 400 subjects) and employed the training-testing framework described earlier (see the second pipeline in Method Section 3.4 and Figure 4B [Left]), i.e. defined the exemplars on the training set and used those exemplars to parcellate individuals in both training and testing sets. We next conducted our predictive analysis by training on one set and testing on the other. The accuracies remained significant (Figure S3) despite the smaller size of the training set. In another analysis, we employed both the parcellation and predictions in a 10-fold cross-validated setting. That is, we divided the entire population into 10 folds. At every step, the exemplars were calculated from the 9 training folds and used to parcellate the entire population. A GBM model was trained on the 9 training folds and tested to predict the sex for the one left-out testing fold. The entire procedure was repeated until each fold was left out once. The prediction accuracies remained significantly higher than chance (Figure S4).

A benefit of using gradient boosting machines is that after the decision trees are constructed, it is relatively straightforward to retrieve the importance of each feature. Importance is explicitly calculated as the number of times that each feature was used to make key decisions in the single decision tree, i.e. decisions that improve the performance measure. The feature importance is weighted by the number of observations within each decision tree, and then averaged across all of the trees within the model. As our GBM model was fit with 188 features indicating the network assignment of each node, we simply derived the importance of each node in sex identification by assessing the corresponding importance attribute. We further scaled the importance scores to the range (0, 100) as shown in Figure 7.

### 3.8 Comparison with other approaches

We compared our proposed exemplar-based parcellation algorithm with two well-established individual-level parcellation methods: (1) our earlier rotation-based individual-level parcellation (Shen, Tokoglu et al. 2013) and (2) Wang's iterative scaling individual-level parcellation (Wang, Buckner et al. 2015). Shen's method had two free parameters, α which tunes the smoothing kernel's standard deviation, and λ, which adjusts the level of similarity between individuals and the group. We set α=0.2 and let λ take values in the range (0.1-0.6), with smaller numbers representing lower similarity.

For the sake of clarity, we only report the results for the two ends of the interval (λ=0.1 and λ=0.6); similar results were found for other values of λ. For some specific number of networks (e.g. K=17, 18) Shen's algorithm terminated at some lower K values, and these were then used instead of the input K. Similarly, we derived the individualized parcellations for Wang's method starting with their K=7 and K=17 group-level parcellation schemes (Thomas Yeo, 2011). For their averaging step, we deployed different weighting schemes, but the results were highly similar across different weighting schemes for the clustering evaluation measures and for sex-prediction analysis. Thus, we present results that used standard averaging.

To quantify the results of the comparison, we used two independent frameworks. In the first step, two clustering validation techniques were applied and in the second step, the sensitivity of these methods to inter-individual variability was examined through comparisons of the predictive power in a sex discrimination analysis.

We utilized two internal clustering validation measures – the Dunn Index (Dunn 1973) and the Davies-Bouldin Index (Davies and Bouldin 1979) – that are commonly reported in the literature (Halkidi, Batistakis et al. 2001, Ghosh, De et al. 2007, Saitta, Raphael et al. 2007, Ziegler, König et al. 2010, Fichtinger, Martel et al. 2011). Similar to all other internal clustering validations, the Dunn and the DB indices utilize the clustered data itself to measure compactness and cluster separation. The Dunn index identifies to what extent the clustering scheme is successful in maximizing the inter-cluster distance while minimizing the intra-cluster distance. For K clusters, the Dunn index is defined as the ratio between the minimal inter-cluster distance to the maximal intra-cluster distance, according to Eq. 10:

$$Dunn_K = \min_{1<i,j<K} \left\{ \frac{\min\limits_{x \in C_i, y \in C_j} d(x,y)}{\max\limits_{1<k<K} \max\limits_{x,y \in C_k} d(x,y)} \right\}, \tag{10}$$

where $d(x, y)$ is the Euclidean distance between the two vectors $x$ and $y$. Therefore, for a given assignment of clusters, a higher Dunn index indicates better clustering. We computed the Dunn index for each individual-level parcellation derived from the three different methods, for the number of clusters (networks) varying from K=2 to K=25. The Davies-Bouldin index (DB) measures the average similarity between each cluster and its most similar one, and is defined according to Eq. 11:

$$DB_k = \frac{1}{K} \sum_{i=1}^{K} \max_{1<j<K, j \neq i} \left\{ \frac{\left(\frac{1}{n_i} \sum_{x \in C_i} d(x,c_i)^2\right)^{\frac{1}{2}} + \left(\frac{1}{n_j} \sum_{x \in C_j} d(x,c_j)^2\right)^{\frac{1}{2}}}{d(c_i,c_j)} \right\}, \tag{11}$$

with $n_i$ the number of points and $c_i$ the centroid of cluster $C_i$. Since the objective is to obtain clusters with minimum intra-cluster and maximum inter-cluster distances, small values for DB are desired. Similarly, the DB indices were calculated using the three individual-level parcellations (described above) for the number of clusters (networks) ranging from K=2 to K=25.

Finally, we assessed the predictive power of our proposed model in comparison with the two other approaches. We employed a sex-prediction analysis described previously, this time using the individual-level parcellations resulting from Shen's and Wang's algorithm. We calculated the accuracies for the number of networks varying in the range K=2 to K=25.

We note here that there is a subtle change in the accuracy results each time the algorithm is executed. This is due to the randomness of ten-fold-cross-validation and also the initial state of the GBM. In the comparison of different methods, we fixed all these parameters and thus the result is for the same initial state and the same assignments of data points to the folds.

### 3.9 Implementation

The parcellation code was written in MATLAB. Clustering was performed on a workstation with 64 GB of RAM, and a 3.4 GHz Intel Xeon processor with 24 cores. Run time for our proposed method with K=25 was 442.15 second for the exemplar-search, 1.95 second for the individual-clustering, and 0.22 second for the group-clustering. Predictive analysis code was written in Python using scikit-learn library (Pedregosa, Varoquaux et al. 2011)

## 4. Results

### 4.1 Visualization of parcellations as a function of the number of networks

One advantage of using the greedy algorithm to solve the optimization of our submodular function is that it provides a hierarchy of nested clusters (through defining the new exemplars while maintaining the older ones) and hence enables an illustrative visualization for different granularities/resolutions as the number of networks is gradually incremented from K=2 to K=25. At K=2, the brain is divided into two subnetworks that are associated with default mode network (DMN) – which is known as the task-negative network – and the rest of the brain, which attributes to the task-positive network. At K=11, many canonical networks (including the DMN, frontoparietal network (FPN), and sensorimotor network (SMN)) are observable (Figure 2). For K>11 the changes are subtle and more difficult to observe (Figure S1).

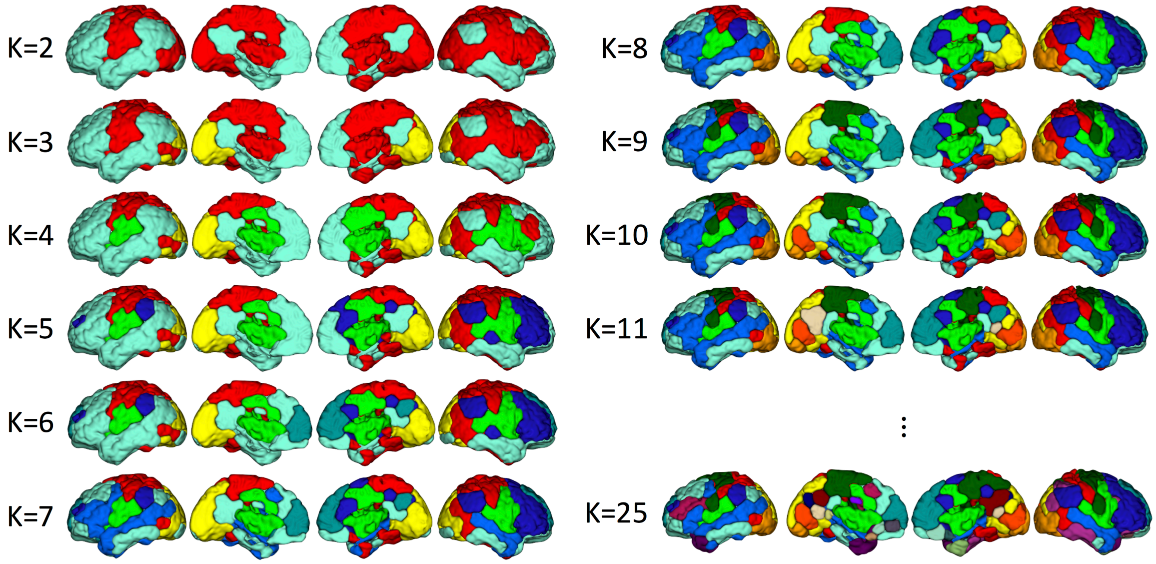

**Figure 2. The group-level parcellation schemes for the number of networks ranging from K=2 to K=25.** At K=2, the brain is roughly divided into the default mode network (DMN) and task-positive network. As K is increased, the greedy algorithm discovers new exemplars while preserving the former ones, and hence parcellates the brain in a hierarchical setting. For example, at K=3, the visual network is separated from the DMN and task-positive network. When K is increased to K=11, many canonical networks (including the DMN, frontoparietal network (FPN), and sensorimotor network (SMN)) are observable. K=25 was the finest resolution parcellation derived here. For K>11 the changes are subtle and more difficult to observe (Figure S1).

### 4.2 Stability and convergence of group-level parcellation as a function of group size

For all numbers of networks, increasing the number of subjects in the training set (on average) decreases the distance between the perturbed parcellation, created with a subsample of subjects, and the final parcellation, created with all subjects (Figure 3). The decrease in the error bars indicates that the distance between the perturbed parcellations resulting from random selection (of the same number) of subjects is also decaying. These findings suggest the algorithm converges to the final solution as the input expands to the entire set. Furthermore, for any number of networks, the average distance between the perturbed and the final parcellation is relatively small: when only using 25 subjects, the perturbed parcellation exhibited an average of 16% difference (i.e. 84% overlap) with the final parcellation. These findings suggest the stability of the algorithm to perturbations to the size of the input and to the selection of the subjects. That the exemplars derived from a relatively small portion of dataset produce parcellations highly similar to the final parcellation scheme (with 84% overlap on average) is a promising result with non-trivial implications for cross-dataset validations.

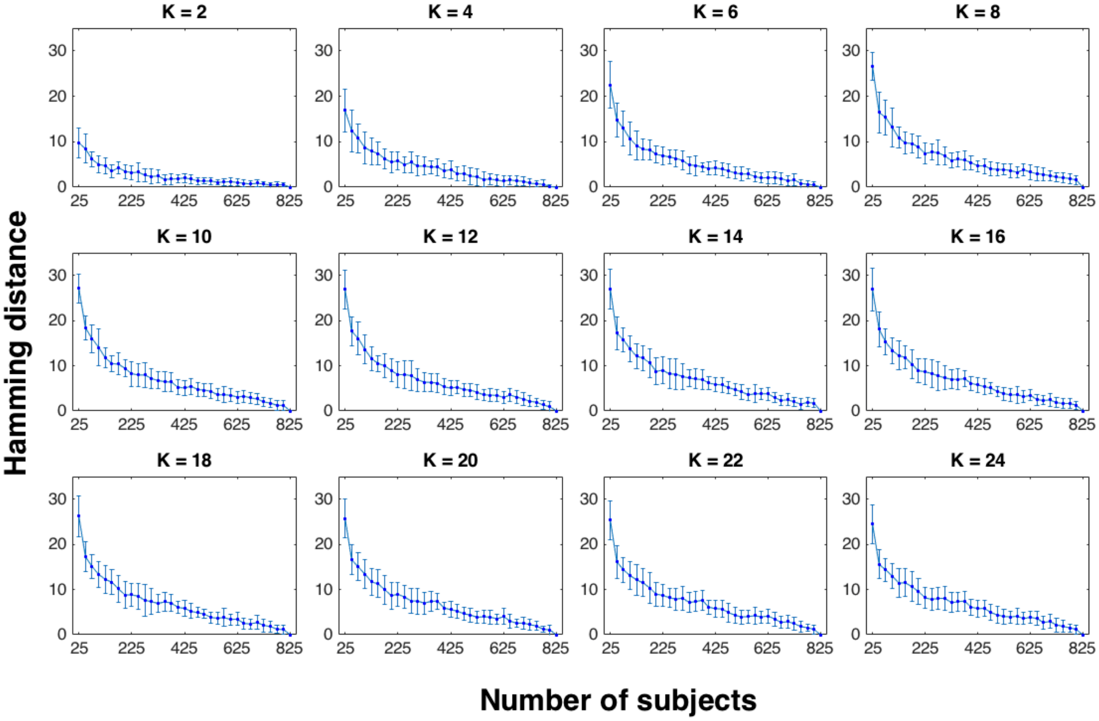

**Figure 3. Stability and convergence of the group-level parcellation algorithm as a functional of group size and individual selection.** For each number of networks taking even values in the range K=2, …, 25, the Hamming distance between the two parcellation schemes is displayed: 1) the group-level parcellation derived from the training set (that is a portion of the full dataset) and 2) the group-level parcellation derived by considering the entire dataset (with 825 subjects). On the x-axis, the number of subjects in the training set is displayed. On the y-axis the Hamming distance (i.e. the number of network differences in the node-to-network assignment vectors) is displayed. Error bars correspond to the variations resulting from 100 permutations for the selection of subjects for the training set. The model is stable to the variation in the group size as the average difference between the perturbed parcellation using a subset of the subjects and the final parcellation using all the subjects is bounded and less than 30 (16% of the full vector with 188 nodes). The model converges to the final solution with a general decaying rate both in the average distance between the perturbed and the final parcellations, and in the error bar lengths. Error bars are proxy of the distances between perturbed parcellations using the same number of subjects selected from the entire dataset over 100 permutations.

### 4.3 Reproducibility of group-level parcellations

Using two non-overlapping subsets, the Dice coefficients between the parcellation results of group 1 and group 2 are depicted in Figure 4A (blue error bars). For all number of networks, there is on average more than 70% overlap between the two parcellations with the overlap generally greater than 80%. Using a training-testing replication method, the Dice coefficients between the training and testing group's parcellation schemes are depicted in Figure 4A (orange error bars). On average, the two parcellations have approximately 96% overlapping occurrences. As anticipated, the Dice coefficients for the second pipeline are significantly higher than the first, in part due to having common exemplars.

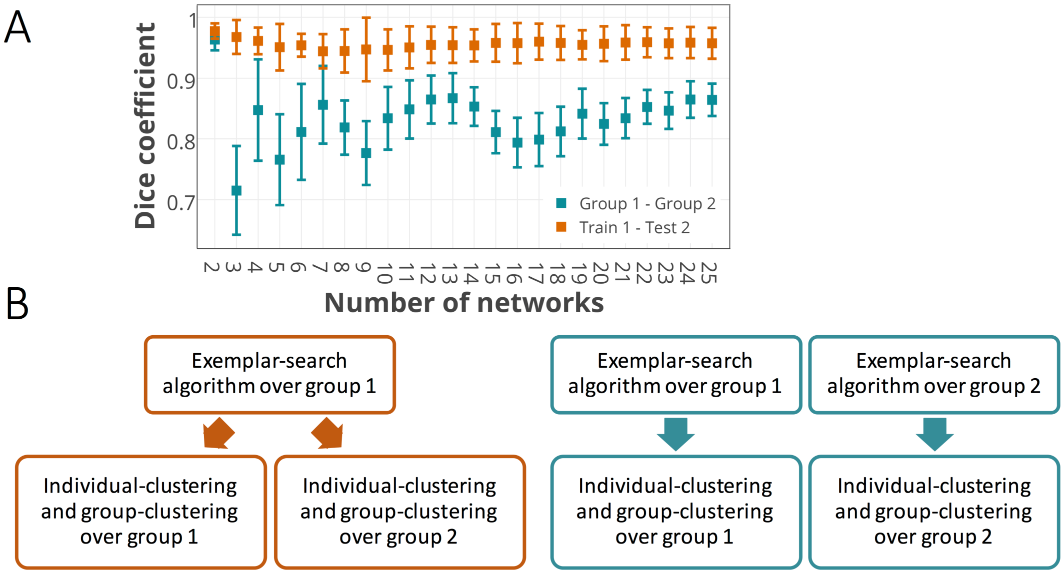

**Figure 4. Reproducibility of the group-level parcellation measured by the Dice coefficient.** A) Dice coefficients between the group-level parcellation of two equal-size sets (with 400 subjects). The reproducibility is examined by two different pipelines shown in part B. The colors match between the error bars (part A) and the diagrams (part B). The blue error bars represent the Dice coefficient between the parcellations derived by running the entire three-step algorithm over each subset (group 1 and group 2) separately, as displayed by the right (blue) diagram in part B. The orange error bars show the Dice coefficient between the two group-level parcellations with the same exemplars (derived from group 1). Due to having a setting similar to training-testing validation, group 1 is called train 1 here and group 2 is called test 2. It corresponds to the left (orange) diagram in part B. B) The two pipelines for addressing the reproducibility of the group-level parcellation algorithm. The Dice coefficient between the parcellation outcomes of the left diagram is depicted in orange, and the corresponding measure for the parcellation of the right diagram is depicted in blue in part A.

### 4.4 Reproducibility of individual-level parcellations across rest sessions

The Dice coefficients between each individual's parcellations across the two rest sessions are depicted in Figure S2. The orange bars correspond to the analysis with global exemplars (across the two sessions). The blue bars display the comparison result using local session-specific exemplars. There is on average 72% overlap between the parcellation results across the two sessions, using the global exemplars. This number decreases to 63% when employing the local exemplars. We note that reliability of individual parcellations across different sessions is subject to various factors including system noise, physiological noise, and intrinsic cognitive processes (Krüger and Glover 2001, Bennett and Miller 2010). Thus, the reliability of the parcellation results could be confounded by factors other than the specific parcellation algorithm employed, and hence warrants further investigation.

### 4.5 Inter-individual variability of individual-level parcellations

Figure 5 displays the sorted distribution of inter-individual variability (in node-to-network assignments) across nodes, using two measure of variability: 1/F1 (Figure 5A) and F2/F1 (Figure 5C). It suggests that there are regions with relatively high values for both measures summed across all numbers of networks. These regions, that follow relatively similar patterns for 1/F1 and F2/F1 across all numbers of networks, display high variation, and lower consensus, in their network assignments between the individual- and the group-level parcellation. These regions are predominantly localized in higher-order association cortices in the frontal, parietal and temporal lobes (Figure 5B, D). In particular, the frontoparietal network, default mode network, and anterior cingulate cortex display high 1/F1 and F2/F1 scores. On the contrary, primary-sensory regions, including the visual network, sensorimotor network, and medial temporal lobe display relatively lower 1/F1 and F2/F1 values. These latter regions demonstrate a higher consistency between the individualized and the group-level parcellation.

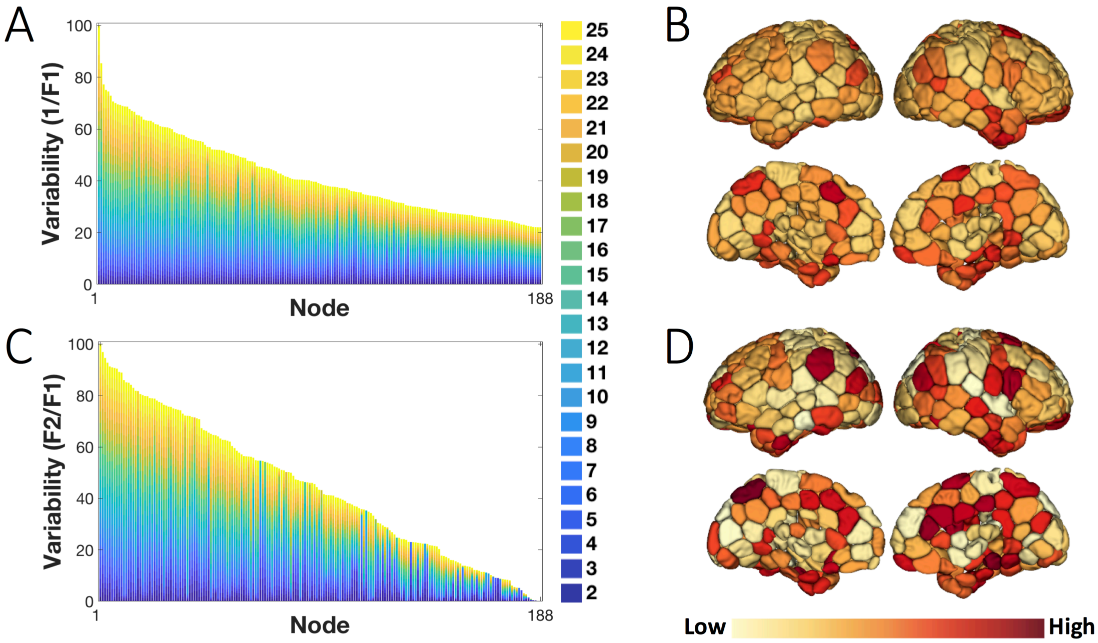

**Figure 5. Inter-individual variability measured by the 1st and 2nd votes in the majority voting.** A) The inverse F1 is displayed for all the cortical nodes in the brain, sorted from high to low. For all numbers of networks (K=2, …, 25), inverse F1 measures are collapsed, scaled, and depicted in a barplot. As F1 measures the number of individuals who voted for the group-vote node-to-network assignment, the inverse F1 is a measure of variability between individuals and the group, with a higher measure indicating higher variability and lower confidence. B) The inverse F1 depicted on the brain after summing over all numbers of networks. C) The ratio between the second (F2) and the first (F1) vote for the node-to-network assignments is displayed for all cortical nodes in the brain, sorted from high to low. Similarly, the F2:F1 ratio is a measure of variability across individuals, as a high F1 and a low F2 corresponds to a confident network assignment reproduced across individuals. The barplot displays the corresponding measure for all numbers of networks (K = 2, …, 25) stacked on top of each other and scaled to the range (0,100). D) F2:F1 ratio depicted on the brain after summing over all numbers of networks. The higher-order association areas in the frontal, parietal and temporal lobes display higher inverse F1 and F2:F1 ratio values compared to primary-sensory areas.

### 4.6 Sex-predictions

Figure 6A displays the sex prediction accuracies for a range of network numbers (K=2, …, 25) using gradient boosting machine (GBM) as a classifier. The accuracies are reported as the mean and standard deviation across all folds (blue bars). The accuracies for the null model are also depicted (orange bars). We observe that the model predicts sex for an unseen individual with the average accuracies ranging from 61% for K=2 to 70% for K=22, with the maximum of 75% for K=22. These reported accuracies are significantly higher than random accuracies (p-value<1e-10), for all numbers of networks even as low as K=2, suggesting meaningful information is stored in the individualized parcellations.

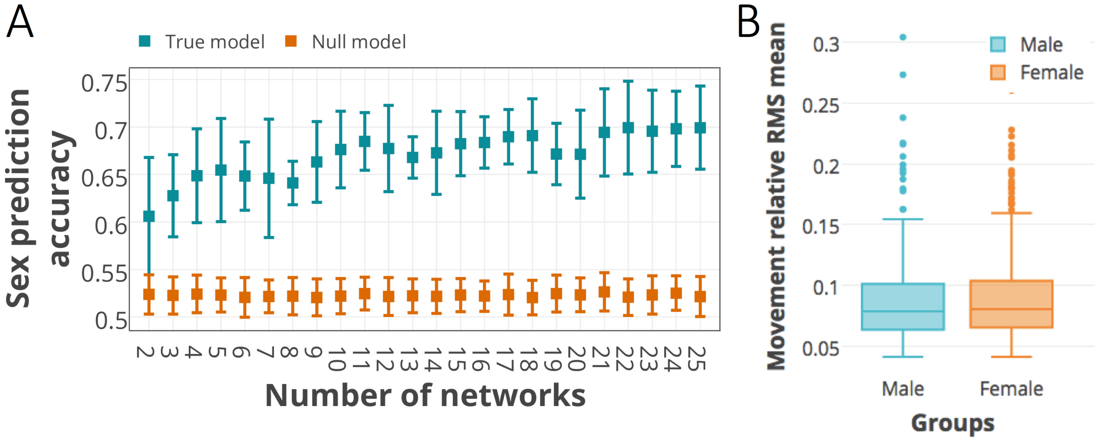

**Figure 6. Sex prediction accuracies using parcellation schemes as features, for the numbers of networks from K=2 to K=25.** A) The sex prediction accuracies for a 10-fold cross-validation using gradient boosting machine (GBM) as the classifier. The classifier is fed with the node-to-network assignment vectors (with 188 elements) as features and a binary output (male vs. female) is predicted for an unseen fold of subjects. The mean and standard deviation across all folds are depicted in blue error bars. To determine the significance of our predictive model, the accuracies derived from the null distributions are also depicted in orange error bars. B) 2-tailed t-test comparison of the head motion between the two sex groups. There are no significant differences in head motion between female (N=458, m=8.9e-02, sd=3.41e-02) and male (N=367, m=8.8e-02, sd=3.55e-02) subjects (two-tailed t-test: t(823) = 0.47, p = 0.64).

We also tested for differences in head motion between the two sex groups (Figure 6B), as motion could be a confound for our predictive analysis. We calculated the average frame-to-frame displacement from the Movement_RelativeRMS.txt for each run and averaged over the 4 runs. Using two-tailed t-tests, there were no significant differences in head motion between female (N=458, m=8.9e-02, sd=3.41e-02) and male (N=367, m=8.8e-02, sd=3.55e-02) subjects (two-tailed t-test: t(823) = 0.47, p = 0.64) (Figure 6B).

To illustrate which regions were the most different between females and males, we utilized the "feature importance" attribute from gradient boosting machine classifier. Figure 7A illustrates the sorted distribution of the importance scores for all the features used for classification, that is a vector of 188 cortical regions. We observe that regions in the anterior and posterior cingulate cortex, precuneus, superior parietal lobule, superior frontal gyrus, parahippocampal gyrus and inferior temporal gyrus (including anterior temporal pole) show relatively high importance scores (Figure 7B). These regions, predominantly located in the default mode network (DMN) and the frontoparietal network (FPN), have been consistently associated with sex differences in the literature (Biswal, Mennes et al. 2010, Scheinost, Finn et al. 2015).

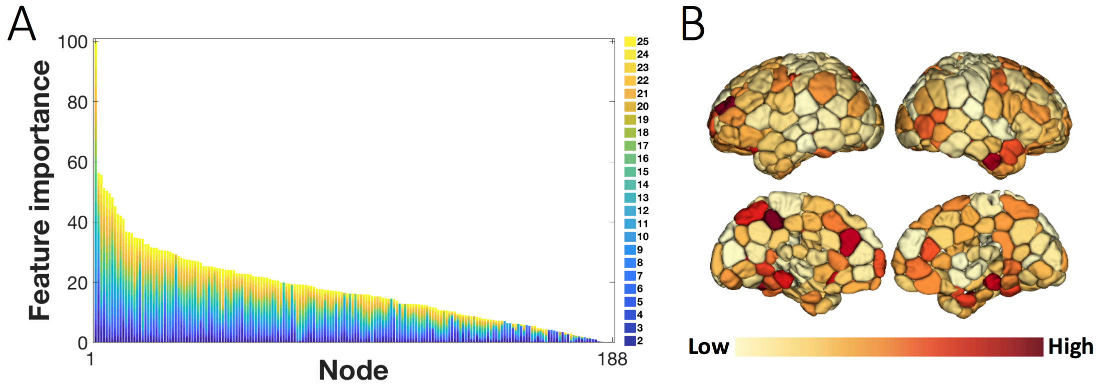

**Figure 7. Node importance in the sex-discrimination predictive model.** A) The sorted distribution of node importances in discriminating sex based on the parcellation schemes. The importance is derived from the "feature importance" attribute of the GBM sex classifier and scaled to the range (0,100). B) The feature importance measures depicted on the brain after summing up over all numbers of networks. Regions in the anterior and posterior cingulate cortex, precuneus, superior parietal lobule, superior frontal gyrus, parahippocampal gyrus and inferior temporal gyrus have relatively high importance scores.

### 4.7 Comparison with other methods

Figure 8 displays the clustering evaluation results from three methods for varying number of networks, from K=2 to K=25 (Dunn Index: Figure 8A; DB index: Figure 8B). Figure 8A (Left) reports the Dunn index for the exemplar-based and Shen parcellation for even Ks. Figure 8A (Right) compares the same measure among all three parcellation approaches (exemplar-based, Shen, and Wang) for K=7, 17. Higher values of Dunn index indicate a better clustering algorithm, with larger intra-cluster and smaller inter-cluster similarities. Figure 8B (Left) depicts the DB index for the exemplar-based and Shen's approach for even Ks. Figure 8B (Right) displays the DB index for all three methods for K=7 and K=17. By definition, lower values for the DB index indicate a better clustering algorithm. These results suggest that our proposed exemplar-based algorithm is able to cohesively parcellate the brain for each individual, specifically for larger values of K.

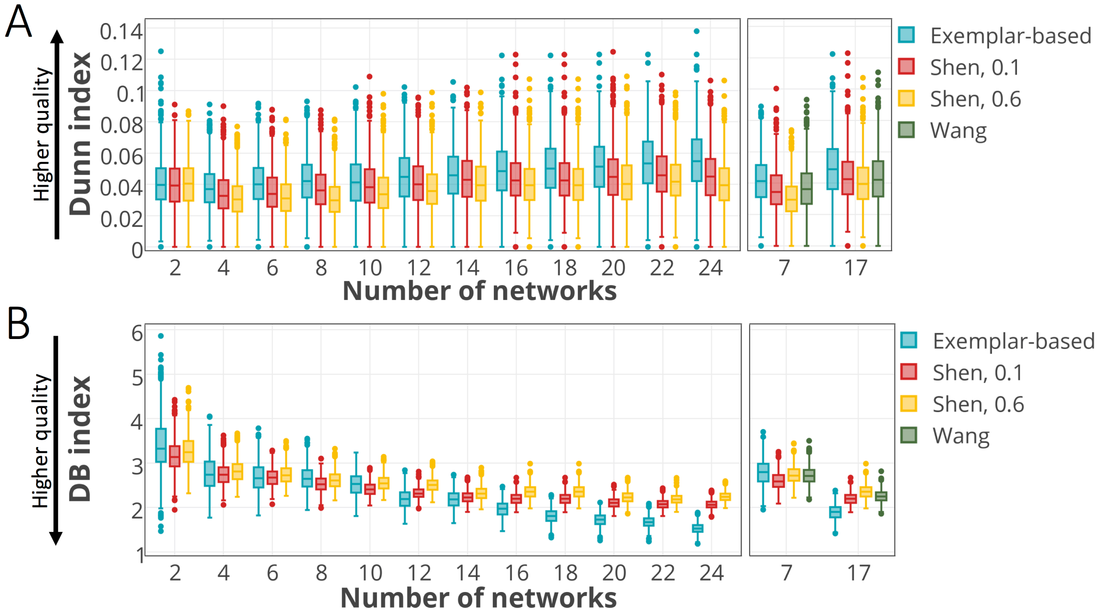

**Figure 8. Comparison of clustering evaluation measures (the Dunn and the Davies-Bouldin (DB) indices) across the three methods.** A) The comparison of Dunn index between the exemplar-based method and Shen's approach for even values of K=2, …, 24 (Left), and the comparison over all three methods for K=7 and K=17 (Right). A higher Dunn index represents higher clustering quality with more compactness within clusters and more separation between clusters. B) The comparison of DB index between the exemplar-based method and Shen's approach for even values of K=2, …, 24 (Left) and the comparison over all three methods for K=7 and K=17 (Right). A lower DB index indicates a higher clustering quality.

In the second step of comparison, we seek to address the model's predictive power in a sex discrimination analysis, using a GBM classifier. Figure 9 (Left) displays the classification accuracies (the mean and standard deviation across all folds) for exemplar-based parcellation and Shen's method, with the number of networks ranging from K=2 to K=24, only taking even values. Figure 9 (Right) compares the three methods for K=7 and K=17.

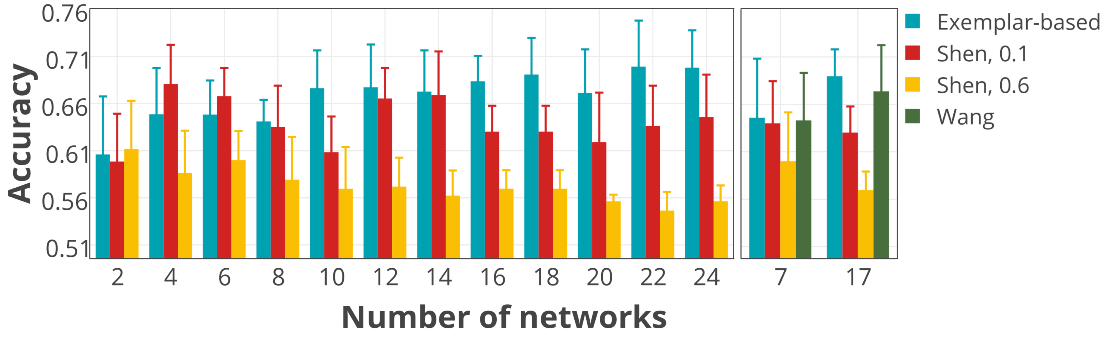

**Figure 9. Comparison of the models' ability to preserve the inter-individual variability as measured by sex-prediction accuracies.** The individual-level parcellation schemes derived from each model are separately fed to the GBM classifier. The classification accuracies (the mean and standard deviation across all folds) for exemplar-based parcellation and Shen's method are displayed with the numbers of networks ranging from K=2 to K=24, taking even values (Left). The classification accuracies for all three methods are displayed for K=7 and K=17 (Right).

# 5. Discussion

A novel algorithm has been introduced here that utilizes submodular optimization to parcellate the cerebral cortex into functional networks at both the group- and the individual-levels. At the group-level, the proposed algorithm has favorable stability, convergence, and replicability properties. At the individual-level, regions of high variability in parcellations overlaps with known regions of high inter-individual variability in functional connectivity and parcellation. We showed that our algorithm performs well on internal clustering validation measures and more importantly it eliminates the cross-subject correspondence problem for a group when parcellating individuals. Finally, using only the individual differences in network parcellation vectors, we built a predictive model using a ten-fold cross-validated framework that predicts sex for the left out subjects with greater than 70% accuracy. This finding that network definitions are sex specific suggests that network studies need to take sex into account and that the same network should not be applied to the population as a whole. These prediction results show the benefit of individual-level parcellation for extracting additional information that would otherwise be missed by simply using a full group-level parcellation.

## 5.1 Exemplar-based clustering for individual network-level parcellation

Exemplar-based clustering algorithms have been successfully applied in a wide variety of data-mining applications. Exemplar-based approaches are conceptually similar to clustering methods such as *k*-means where we aim to find a set of representative points that best fit the data as a whole. Although *k*-means algorithms yield satisfactory results for problems with a small number of clusters, they generally suffer from sensitivity to the initialization (also called seeding). As the *k*-means cost function is highly non-convex, the commonly used iterative algorithms converge to local optima depending on the initialization. One key difference between the exemplar-based methods and *k*-means is that the former restricts the selection of the representative points to the actual observed data points. By doing so, instead of minimizing a continuous loss function, we maximize a discrete submodular function for which the classical greedy algorithm provides the best approximation to the optimal solution. Note that in general there are exponentially many possibilities. However, submodularity allows us to find a near-optimal solution in linear time (Mirzasoleiman, Badanidiyuru et al. 2015). In fact, exemplar-based clustering is empirically more robust to noise and outliers than *k*-means methods or its close variants such as Wang's iterative brain parcellation (Wang, Buckner et al. 2015). There are other variations of *k*-means that include soft assignment of nodes to clusters, such as fuzzy *c*-means (FCM) (Bezdek 2013). Similarly, the proposed exemplar-based approach could be extended to incorporate probabilistic assignment of nodes to networks, where the probability of assigning a node to a network is proportional to the inverse distance of the node to the corresponding exemplar. Finally, the greedy algorithm smoothly splits the old networks, similar to hierarchical clustering methods. This is in contrast to our earlier work (Shen, Tokoglu et al. 2013) where for each value of K, a different network is proposed without considering correspondences.

## 5.2 Comparison of algorithms for individual-level networks

We compared our exemplar-based parcellation with two other algorithms for delineating individual-level networks: Shen's rotation based algorithm (Shen, Tokoglu et al. 2013) and Wang's iterative scaling algorithm (Wang, Buckner et al. 2015). These two methods take contrasting approaches from each other to define individual-level networks, leading to different strengths and limitations. Shen's method assumes that the entire dataset is accessible to jointly create new group- and individual-level networks. This method is well-suited for studies where the network structure of the current population is not applicable for preexisting parcellations and a new parcellation must be created. However, this method may give different group-level networks for each study, and thus a lack of one-to-one correspondence between studies. Alternatively, Wang's method assumes a group-level parcellation can be modified to fit an individual's networks with limited changes to the gross topology of the group-level networks. This algorithm is well-suited for many studies where a preexisting group-level parcellation is a reasonable assumption. However, if an individual's networks differ from the group-level parcellation such as in the case of brain tumors or other pathology (Ghumman, Fortin et al. 2016), it is not clear how well this algorithm will perform. Given these limitations, neither approach can generalize to multiple applications. In contrast, our exemplar-based parcellation algorithm can be used to accomplish either of these purposes. We show our algorithm's ability to find exemplars and parcellate individual-level networks in the main analysis (Figure 2 and 4) and to find individual-level networks given a set of exemplars in the split-half analysis (Figure 3 and 4). In this sense, our algorithm generalizes these contrasting approaches.

### 5.3 The need for individualized networks

Our finding that individual-level parcellations can predict sex demonstrates a problem of group-level parcellations. As the sex prediction relies only on network organization (not the connectivity based on these networks, as reported in Satterthwaite et al. (Satterthwaite, Wolf et al. 2015)), these results show that important information can be missed with group-level parcellations. If a basic characteristic such as sex in a cohort of healthy controls of a similar age, it is reasonable to assume that other characteristics linked to connectivity such as age (Hampson, Tokoglu et al. 2012), cognition (Finn, Shen et al. 2015, Smith, Nichols et al. 2015, Rosenberg, Finn et al. 2016), and neuropsychiatric diagnosis (Fornito and Harrison 2012) could also show different individual-level networks. Overall, this finding suggests the need for individual-level parcellation algorithms, like our approach, to address individual differences, while maintaining a one-to-one correspondence of networks across subjects.

### 5.4 Localizing inter-individual variability

Our findings suggest that the greatest inter-individual variability in network organization is located in limbic, parietal, and prefrontal regions (Figure 5). These findings are consistent with the previous studies that have examined inter-individual variability in connectivity (Mueller, Wang et al. 2013, Miranda-Dominguez, Mills et al. 2014, Finn, Shen et al. 2015, Mejia, Nebel et al. 2016), and parcellations (Gordon, Laumann et al. 2015, Laumann, Gordon et al. 2015, Wang, Buckner et al. 2015). Accumulating evidence suggests that the neural systems subserving higher-order association cortices display more inter-individual variability in their connectivity profiles than those in sensorimotor regions (Frost and Goebel 2012, Mueller, Wang et al. 2013). These regions further match with maps of evolutionary cortical expansion (Zilles, Armstrong et al. 1988) and long-range integration and regional segregation (Sepulcre, Liu et al. 2010), whose reflection on parcellation is expected.

### 5.5 Sex differences

Recent neuroimaging studies have reported sex differences in functional connectivity (Kilpatrick, Zald et al. 2006, Biswal, Mennes et al. 2010, Scheinost, Finn et al. 2015, Zhang, Cahill et al. 2016). We observed that regions in the anterior and posterior cingulate cortex, precuneus, superior parietal lobule, superior frontal gyrus, parahippocampal gyrus and inferior temporal gyrus exhibited relatively high importance scores. These regions, predominantly located in DMN and FPN, have been reported to display sex differences (Biswal, Mennes et al. 2010, Scheinost, Finn et al. 2015). These regions have also been consistently identified as functional hubs in the brain (Zuo, Ehmke et al. 2012, van den Heuvel and Sporns 2013), showing a high density of connections. When taken together, these observations suggest that functional hubs exhibit different network organization in males and females, consistent with previous studies (Tomasi and Volkow 2012).

Note, however, that the main focus of the presented analysis was not to demonstrate sex differences in the functional organization of the brain. This could have been achieved using more informative features, such as functional connectivity matrices with information regarding all edges. Nor was it to distinguish between the two sex groups. Instead, the sex prediction was used to demonstrate that group effects can lead to different network definitions and thus patient versus control or patient group comparisons should not assume that the use of global network definitions is appropriate. As an aside, it is also interesting to note that given the minimal information stored in the functional node-to-network assignment vectors, composed of all integer values (1, 2, … K), it is impressive that such group effects can be detected.

### 5.6 Strengths and limitations

This study has several strengths. Unlike many other parcellation algorithms (Beckmann, DeLuca et al. 2005, Power, Cohen et al. 2011, Thomas Yeo, Krienen et al. 2011, Wang, Buckner et al. 2015, Gordon, Laumann et al. 2016), our proposed approach does not depend on thresholds, or the selection of hyper parameters. Our method provides a one-to-one mapping across subjects and no additional algorithms are needed to map network correspondences. However, there are several limitations that should be noted. Individual-level networks could be influenced by individual differences in physiological noise (Rogers, Morgan et al. 2007) and head motion (Van Dijk, Sabuncu et al. 2012). As males and females did not show differences in motion, our network differences as a function of sex are unlikely due to motion.

In this work, the starting point was a 188-node functional atlas. It would be quite reasonable to begin, instead, at the voxel-level as individual node definitions may differ between the sexes whereas our starting point assumes they are the same. Starting at the node level reduces the computational burden because the process of defining nodes already provides a large dimensionality reduction step. Moreover, as the atlas was derived from an independent dataset, this can also reduce the chance of overfitting the data of interest. On the other hand, it may suffer from the propagation of registration noise and misalignment of the preexisting atlas. The approach described above, however, is applicable at the voxel-level and this can therefore be used for node definitions at the individual-level while maintaining cross-subject correspondences.

Even though we started from a node-level atlas, the number of features used by the predictive model (d=188) was relatively high comparing to the number of samples (n=825). This may lead to a larger variance in the model and thus make it harder to generalize over novel subjects. We deliberately did not reduce the number of features, in order to achieve higher accuracies, as we sought to employ a fully transparent and data-driven analysis of the feature space with all the nodes included as features. In this regard, we also did not force any prior knowledge on the importance of nodes. Nevertheless, the achieved accuracies are comparable (or higher for some Ks) to the previous models that have used full functional connectivity data (Satterthwaite, Wolf et al. 2015).

## 6. Conclusion

In conclusion, we present a novel algorithm to parcellate individual-level networks using exemplar-based clustering with submodularity optimization. The algorithm compares favorably with existing algorithms when parcellating nodes into individual-level networks while maintaining cross-subject correspondences. Using networks defined at the individual-level, we demonstrated that brain network organization differs between the sexes as indicated by our ability to predict sex with 70% accuracy. The sex prediction finding illustrates that individual parcellation of functional networks can reveal subgroups in a population and suggests that the use of a global network parcellation may overlook fundamental differences in network organization in subgroups. This is a particularly important point to consider in studies comparing patients versus controls for example or even patient subgroups. Network organization may differ between individuals and global configurations should not be assumed.

## 7. Supplemental Materials

While the theoretical lower bound for the greedy algorithm is $\approx 63\%$ of the optimal solution, in practice, it provides a solution significantly closer to the optimal one. To demonstrate this, we employed an empirical evaluation. We calculated a data-driven upper bound for the optimal solution, following the equations below:

$$f(S^*) \leq f(S^* \cup S_k) \tag{1}$$

$$= f(S_k) + \sum_{j=1}^{k} \Delta\left(v_j^* | S_k \cup \{v_1^*, v_2^*, \dots, v_{j-1}^*\}\right) \tag{2}$$

$$\leq f(S_k) + \sum_{v \in S^*} f(v|S_k) \tag{3}$$

$$\leq f(S_k) + \sum_{v \in V \setminus S_k} f(v|S_k), \tag{4}$$

where $S^*$ is the optimal exemplar set, and $S_k$ is the exemplar set calculated by the greedy algorithm. Equation 5 provides an upper bound for the optimum solution, $f(S^*)$. Figure S5 compares the upper bound with the utility function derived from the greedy algorithm. This demonstrates that the greedy algorithm is significantly closer to the upper bound than to the 63% of the optimal solution, as provided in theory.

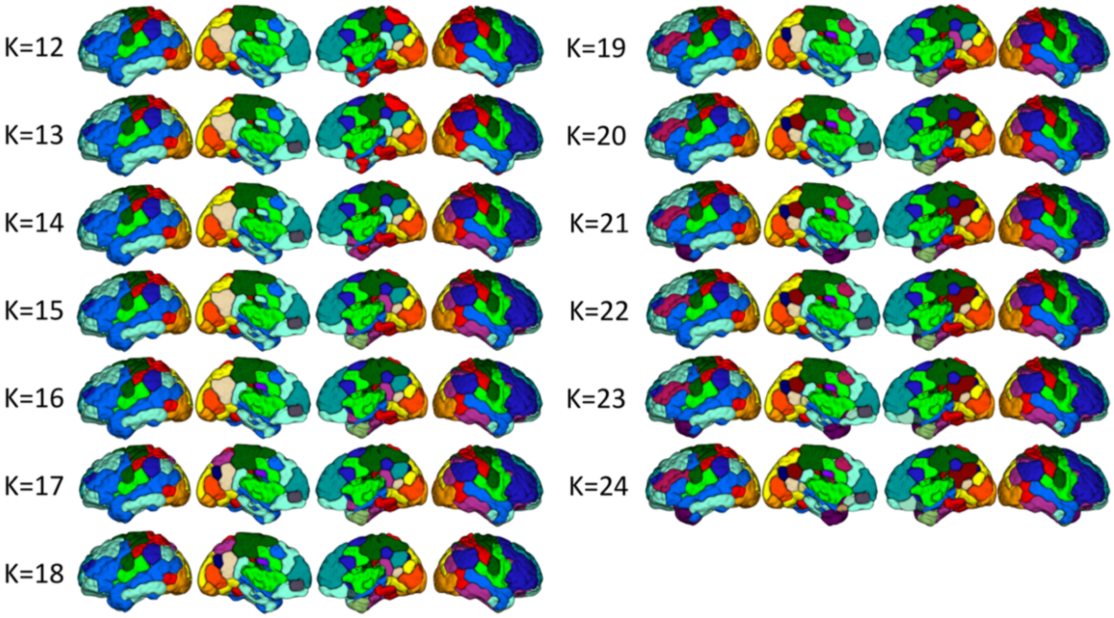

**Figure S1. The group-level parcellation schemes for the numbers of networks ranging from K = 12 to K=24.** As the number of networks increases, the larger networks tend to split and slightly rearrange to accommodate new ones, similar to a hierarchical nested scheme.

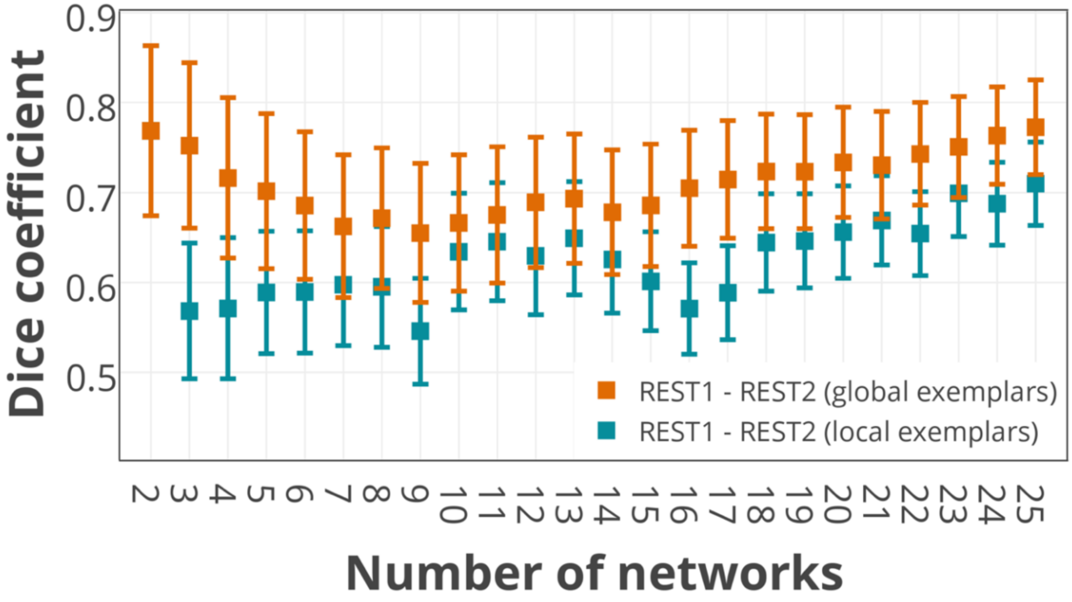

**Figure S2. Reproducibility of the individual-level parcellations across the two rest sessions (REST1 and REST2) measured by the Dice coefficient.** Dice coefficients between the individualized parcellations derived from the two rest sessions. The reproducibility is examined by two different pipelines. The orange error bars display the Dice coefficient between the individualized parcellations derived from the same global exemplars (derived from concatenation of REST1 and REST2, as explained in Method Section 3.2). The blue error bars represent the Dice coefficient between the parcellations derived by running the entire three-step algorithm within each rest session separately. The error bars display the standard deviation across subjects.

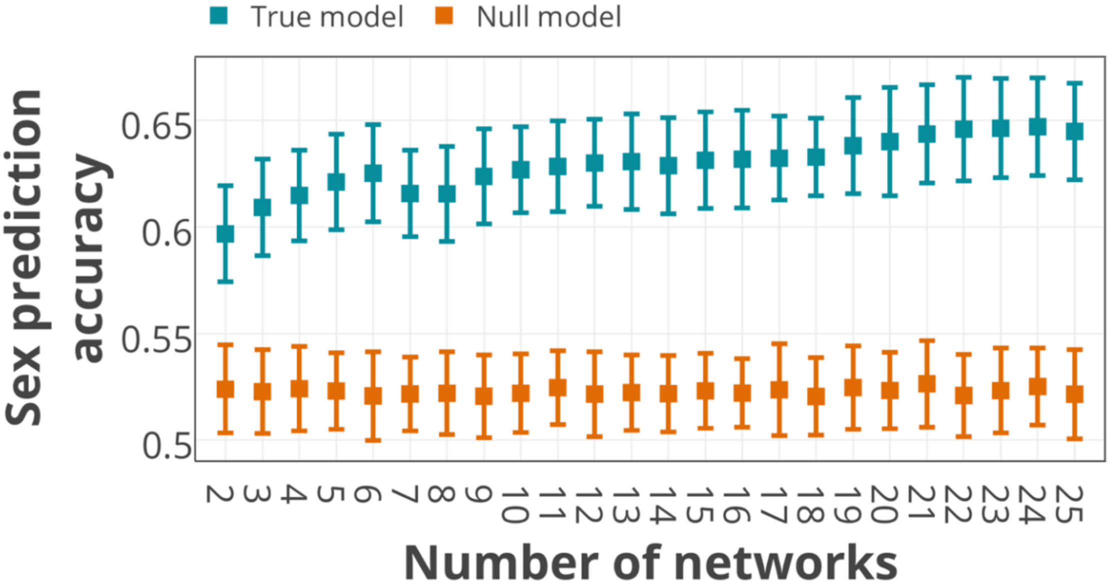

**Figure S3. Sex prediction accuracies using parcellation schemes derived from two non-overlapping subsets of 400 subjects.** Two sets of individualized parcellations are derived using two non-overlapping subsets (training set and testing set, each with 400 subjects). Exemplars are derived from the training set and used to parcellate the entire dataset (comprised of both training and testing sets). A GBM classifier is trained on the training set using the parcellation schemes (i.e. node-to-network assignment vectors) as features. Next, it is tested on the testing set. The entire process is repeated 100 times and the prediction accuracies for K=2, …, 25 are reported (as the means and standard deviations over 100 runs).

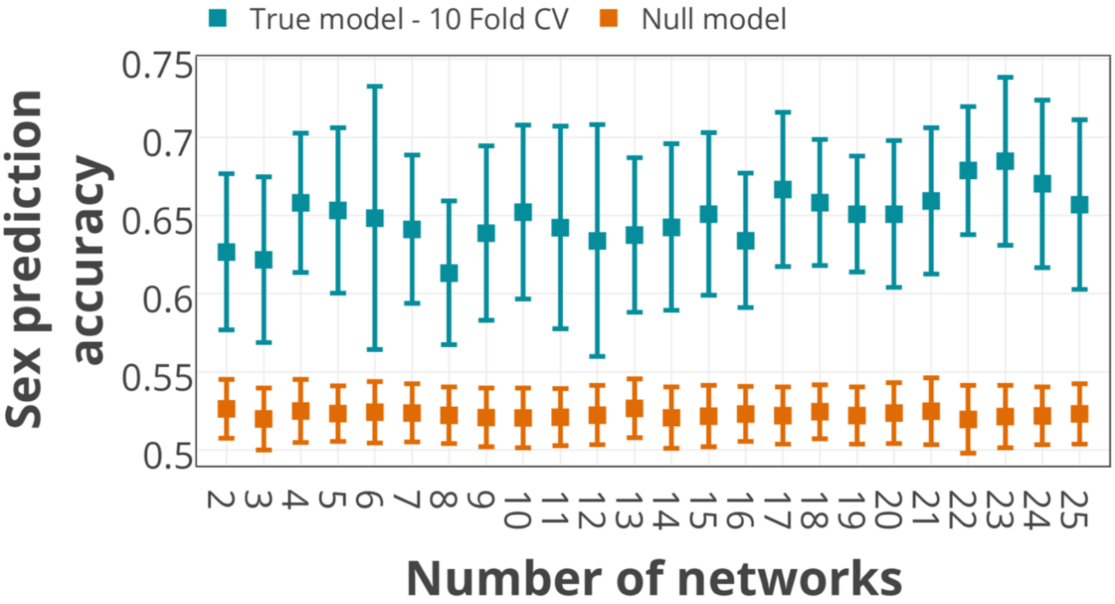

**Figure S4. Sex prediction accuracies using parcellation schemes, derived from a 10-fold cross-validation.** The sex prediction accuracies are reported for a 10-fold cross-validated parcellation scheme using gradient boosting machine (GBM) as the classifier. The entire population is divided into 10 folds. At every step, exemplars were calculated from 9 folds (the training set) and used to parcellate all the subjects. Next, a GBM classifier is trained on the 9 training folds and tested to predict the sex for the one left-out testing fold. The GBM classifier is fed with the node-to-network assignment vectors (with 188 elements) as features and a binary output (male vs. female) is predicted for an unseen fold of subjects. The entire procedure is repeated until each fold is left out once. The mean and standard deviation across all folds are depicted in blue error bars. To determine the significance of predictions, the accuracies derived from the null distributions are also depicted in orange error bars.

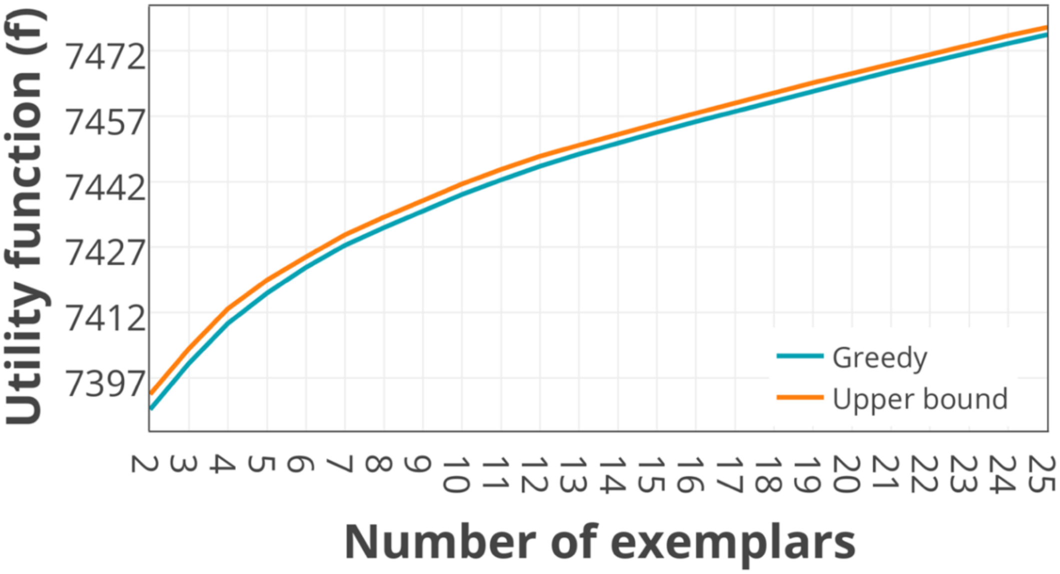

**Figure S5. Comparison of the greedy algorithm and the data-driven upper bound for the optimal solution.** While in theory the greedy algorithm provides a ≈ 63% approximation of the optimal solution, this is only a lower bound on the greedy algorithm's performance. Here, we empirically show that the greedy solution (blue line) is considerably close to the optimal solution's upper bound (orange line), calculated from the data.

# Chapter 4: Individualized functional networks reconfigure with cognitive state

## Abstract

There is extensive evidence that functional organization of the human brain varies dynamically as the brain switches between task demands, or cognitive states. This functional organization also varies across subjects, even when engaged in similar tasks. To date, the network organization of the brain has been considered static. In this work we use fMRI data obtained across multiple cognitive states (task-evoked and rest conditions) and across multiple subjects, to measure state- and subject-specific functional network parcellation (the assignment of nodes to networks). Our parcellation approach provides a measure of how node-to-network assignment (NNA) changes across states and across subjects. We demonstrate that the brain's functional networks are not spatially fixed, but that many nodes change their network membership as a function of cognitive state. Such reconfigurations are highly robust and reliable to the extent that they can be used to predict cognitive state with up to 97% accuracy.

## 1. Introduction

The human brain is organized into functional networks that reflect the coordinated effort of individual subunits (or nodes) to execute specific functions (Finn et al., 2015; Power et al., 2011; Smith et al., 2009; Thomas Yeo et al., 2011). Previous studies have identified 5-20 networks during the resting-state that putatively represent the "intrinsic" functional organization of the brain (Meunier et al., 2009; Power et al., 2011; Smith et al., 2012; Thomas Yeo et al., 2011). These networks have been associated with a wide range of cognitive tasks (Dosenbach et al., 2007; Grayson et al., 2014; Laird et al., 2011) and alterations in the spatial organization of these networks have been linked to clinical disorders (Greicius et al., 2004; Stern et al., 2012; van Eimeren et al., 2009). Increasingly, there is evidence that the functional organization of the brain is fluid and reconfigures with task demands or specific cognitive states (Cole et al., 2014; Krienen et al., 2014; Mennes et al., 2012). Understanding network reconfiguration across different task states is a key step towards understanding the functional organization of the brain.

Recent studies have made significant progress in defining state-evoked changes in functional connectivity elicited by task performance (Cole et al., 2010; Cole et al., 2014; Cole et al., 2016; Cole et al., 2013; Davison et al., 2015; Di et al., 2013; Krienen et al., 2014; Shine et al., 2016; Telesford et al., 2016). However, none of these studies explicitly examined the possibility that the spatial topography of functional networks reconfigure, with most studies explicitly assuming that the networks remain spatially unchanged across tasks (Cole et al., 2014; Cole et al., 2016; Gratton et al., 2018; Telesford et al., 2016). Such analyses have been primarily restricted to investigating differences in connectivity between networks, and not whether networks spatially reconfigure across tasks. Changes in functional network spatial topography, however, could be driven by changes in functional connectivity across different tasks. Another line of research takes a more abstract perspective by computing global network measures (such as modularity and participation coefficient) and comparing these measures across different states (Bassett et al., 2011; Cohen and D'Esposito, 2016; Di et al., 2013; Kinnison et al., 2012; Mohr et al., 2016; Shine et al., 2016). While such approaches address the modular reconfiguration of the brain as a whole, they do not quantify how or whether the same networks change across states. Such studies also typically define a new set of networks (ranging between 3 to 7) for every state, whose correspondence to the putative resting-state networks is unclear (Bassett et al., 2011; Di et al., 2013; Kinnison et al., 2012; Mattar et al., 2015a). Finally, many studies do not directly examine cross-subject variations in network reorganization (Krienen et al., 2014), or consider cross-subject and cross-state changes as similar notions for defining reorganization (Mattar et al., 2015b). However, the functional organization of the brain varies significantly across subjects, yielding large individual differences in network definitions (Gordon et al., 2017a; Gordon et al., 2017b). In this work, we account for both cross-subject and cross-state variability in functional network organization demonstrating the dynamic flexibility of this organization.

A new approach is introduced to dynamically map node membership in functional networks across task-evoked and resting states, and across individuals. We apply our recently developed exemplar-based parcellation method (Salehi et al., 2018) to assign a set of pre-defined nodes to individualized, state-specific functional networks, while preserving precise network correspondences across subjects and states. We demonstrate that resting-state networks reconfigure as nodes change their allegiance in specific manners according to the task condition or cognitive state. We demonstrate that the observed state-evoked reconfiguration of networks is reliable, through the use of a cross-validated predictive model that decodes the cognitive state of unseen subjects based on their node-to-network assignments (NNA). Elucidating network configurations across a number of tasks reveals three classes of nodes based on the stability of these node-to-network assignments: 1) steady nodes that exhibit the same network assignment across states and subjects; 2) flexible nodes that change their network assignment across states, but are consistent across subjects, and 3) transient nodes that change their network assignment across both subjects and states. To give behavioral context for these node classes, we use a large-scale meta-analysis of task activation studies from BrainMap to assign behavioral domains and paradigm classes to each node. Together, our findings provide a comprehensive view of how large-scale functional networks (as putatively defined during resting-state) are modulated by task performance. If networks are to be defined via functional clustering of nodes, then it is essential to consider that such definitions are fluid and cognitive-state dependent.

## 2. Materials and Methods

### 2.1 Participants and processing

Functional MRI data from the Human Connectome Project (HCP) (Van Essen et al., 2013) were analyzed. We limited the analysis to 718 subjects for which data were available for all nine functional conditions (hereafter, “states”): MOTOR, GAMBLING, WORKING MEMORY (WM), EMOTIONAL, LANGUAGE, RELATIONAL, SOCIAL, REST1, and REST2. For details of scan parameters, see (Uğurbil et al., 2013) and (Smith et al., 2013). Starting with the minimally preprocessed HCP data (Glasser et al., 2013), further preprocessing steps were performed using BioImage Suite (Joshi et al., 2011) and included regressing 12 motion parameters (Movement_Regressors_dt.txt), regressing the mean time courses of the white matter and cerebro-spinal fluid as well as the global signal, removing the linear trend, and low pass filtering (as previously described in (Finn et al., 2015)).

### 2.2 Functional distance and functional connectivity matrices

Functional matrices were assessed using a functional brain atlas (Shen et al., 2013) consisting of 268 nodes covering the whole brain; this atlas was defined using resting-state data from a separate population of healthy subjects (Finn et al., 2015). For every subject in each functional condition, time courses for the two functional runs with opposing phase-encoding directions (left-right, “LR”, and right-left, “RL”) were concatenated and further used to generate a ground set $V$ consisting of $N = 268$ vectors in $T$-dimensional space, where $T$ indicates the length of scan session. To construct the parcellations, we first normalized all the data points to a unit norm sphere centered at origin. We then calculated the pairwise squared Euclidean distances between the data points, yielding a distance matrix of size 268 × 268 for each subject and each state. We also constructed functional connectivity matrices to visualize the brain connectivity patterns in different states. To this end, the Pearson correlation coefficients between the time courses of each possible pair of nodes were calculated and normalized using Fisher's z-transformation. Each element of the functional connectivity matrix represents a functional connection, or edge, between two nodes.

### 2.3 Individualized and state-specific functional network parcellation

We used exemplar-based parcellation (Salehi et al., 2018) to assign nodes into functional networks in an individualized manner for each state and subject, such that every subject acquired an individualized node-to-network assignment (NNA) for each state. This parcellation method provides an approach to summarize data by introducing a set of $K$ exemplars that best represents the full data.

We followed the same methodology explained in (Salehi et al., 2018), by attempting to find $K$ exemplar labels $S = \{e_1, e_2, \dots, e_K\}$ with $e_i \in \{1, \dots, N = 268\}$, whose corresponding exemplar vectors for each individual $j$ in each state $m$, i.e., $S^{m,j} = \{v_{e_1}^{m,j}, v_{e_2}^{m,j}, \dots, v_{e_K}^{m,j}\} \subseteq V^{m,j}$, maximizes a desired utility function, as follows:

$$F(S) = \sum_{m=1}^{M} \sum_{j=1}^{J} f^{m,j}(S^{m,j}), \tag{1}$$

where $S$ is the exemplar label set and $S^{m,j}$ is the set including the corresponding exemplar vectors in subject $j \in \{1, \dots, J = 718\}$ at state $m \in \{1, \dots, M = 9\}$. Here,

$f^{m,j}: 2^{V^{m,j}} \to \mathbb{R}$ is the utility function of subject $j$ at state $m$. This utility function is derived from the following loss function $L(S)$, which measures how much information we lose if we summarize the entire ground set to the exemplar set $S$ by representing each data point with its closest exemplar:

$$L^{m,j}(S) = \frac{1}{|V^{m,j}|} \sum_{n=1}^{N} \min_{e \in S} d(v_e^{m,j}, v_n^{m,j}), \tag{2}$$

where $V^{m,j}$ is the ground set of size 268 and dimension $T$ for subject $j$ at state $m$, and $v_e^{m,j} \in S^{m,j} \subseteq V^{m,j}$ is the node exemplar vector corresponding to the exemplar label $e \in S$.

While minimization of (2) subject to the cardinality constraint would output the desired exemplar set, this minimization is NP-hard. We transform this minimization into maximization of a monotone, non-negative, submodular function (see (Fujishige, 2005) for a formal definition and for properties of submodular functions) by introducing an auxiliary exemplar, $v_0$, as follows (Krause and Gomes, 2010):

$$f^{m,j}\left(S^{m,j}\right) = L^{m,j}(v_0) - L^{m,j}\left(S^{m,j} \cup v_0\right), \tag{3}$$

Maximization of (3) subject to the cardinality constraint is tractable using general greedy algorithm, which is known to provide an efficient 1-1/e ≈ 0.63 approximation to the optimal solution (Nemhauser et al., 1978). The auxiliary exemplar can be an arbitrary vector whose distance to every data point is greater than the pair-wise distances between data points. Here we used $v_0 = [3,0, \dots ,0]$ as the auxiliary exemplar and employed an accelerated version of the greedy algorithm, known as Lazy Greedy (Minoux, 1978).

After $K$ exemplar labels were identified, the corresponding exemplar vectors in every subject and state were used to parcellate the brain. We assigned each node $n$ in subject $j$

and state $m$ (i.e. vector $v_n^{m,j}$) to the closest exemplar (which represents the corresponding community or network), i.e.:

$$\text{Network}\left(v_n^{m,j}\right) = \arg\min_{\mathrm{e_i} \in S} d(v_n^{m,j}, v_{e_i}^{m,j}). \tag{4}$$

Thus, the brain was parcellated into $K$ networks each represented by an exemplar.

To generate the state-specific population-level parcellations, we employed the winner takes all algorithm over all subjects in each state. That is, region $n$ was assigned to network $k$ at state $m$ if the majority of subjects in that state voted for this assignment. As the number of networks ($K$) is an arbitrary parameter, we repeated the analysis for $K = 2 - 50$ networks. Results presented in the main text are for $K = 12$ networks, as it was the minimum number of networks after which the assignments stabilized with significantly less changes in the NNAs (see Supplementary Materials, Figure S3).

### 2.4 Functional network reconfigurations across states

We quantified the state-evoked network reconfigurations by computing the pair-wise Hamming distances (Hamming, 1950) between every pair of state-specific population-level parcellation vectors. This resulted in a distance matrix $D_{9\times9}$, where every element $d_{m_1,m_2}$ represents the number of nodes that changed their assignment from state $m_1$ to state $m_2$. To quantify each individual network reconfigurations, we considered REST1 as the benchmark for comparison. For every network $k$, we computed the normalized Hamming distance between the NNA vector limited to the nodes that are assigned to network $k$ during REST1, and the vector containing the same nodes' network assignment during every other state.

### 2.5 Decoding cognitive states by node-to-network assignments (NNA) as features

To demonstrate that the state-evoked reconfigurations are robust across subjects and specific to each state, we built a cross-validated predictive model that decodes the cognitive state of each individual brain based solely on the NNAs. We employed two separate predictive pipelines: in the first pipeline, we employed a two-class (binary) classification on every pair of states. For every pair of state $m_1$ and $m_2$, we combined the corresponding populations resulting in a data set with $718 \times 2$ subejcts. In this case, the chance accuracy was equal to 50%. In the second pipeline, we employed an eight-class classification over the entire data set (excluding REST2). REST2 was excluded to eliminate the reduncdancy of resting-state session and balance the probability space evenly across sessions. For both pipelines, we developed a predictive model using gradient boosting machine [GBM (Friedman, 2001)] with 100 estimators (or decision trees) and 0.05 learning rate. We randomly divided the entire population into a training set and a testing set. We calculated exemplars using the training set and used those to parcellate the entire population. Next, we trained a GBM on the NNAs of the training set and predicted the state for the unseen subjects in the testing set. To evaluate the significance of the results, we employed non-parametric permutation test: we randomly permuted the output vector (here the cognitive states) 1000 times and each time ran the permuted values through the same predictive pipeline and calculated the accuracy.

Note that we restricted the exemplar identification to the training set to assure the independency of the training and the testing set throughout the pipeline. However, we repeated the entire analysis using the initial parcellation schemes in a 10-fold cross-validated setting (Figure S4). The predictive pipeline was implemented using Python's Scikit-learn library (Pedregosa et al., 2011).

### 2.6 Classifying nodes based on NNA reconfigurations

To quantify the state-evoked network reconfigurations, we employed a statistical ensemble analysis. For every state, we randomly divided the entire population into two equal-size sets, and employed the population-level parcellation on each set independently using the winner-takes-all strategy over the corresponding subjects. We repeated this 50 times, to derive a distribution of 100 state-specific population-level parcellations. Using this distribution, we quantified which nodes changed their NNAs across states and subjects. To this end, we calculated the entropy (a metric of uncertainty from information theory (Borda, 2011)) of the NNA histograms. We used two measures of entropy: 1) $\mathrm{Entropy}_{\mathrm{cross-subject}}$ to measure the variation across multiple groups of subjects within the same state (see Eq. 5-6), and 2) $\mathrm{Entropy}_{\mathrm{cross-state}}$ to measure the variation across multiple states for the same group of subjects (see Eq. 7-8). A high entropy in either of these measures means high variation in NNAs.

$$\mathrm{Entropy}_{\mathrm{cross-state}}(p,m) = -\sum_{i=1}^{100}\sum_{k=1}^{K} \Pr_{\mathrm{cross-state}}\left(p_{i,m} \in k\right) \times \log \Pr_{\mathrm{cross-state}}\left(p_{i,m} \in k\right), \tag{5}$$

$$\Pr_{\mathrm{cross-state}}(p_{i,m} \in k) = \frac{1}{M}\sum_{m=1}^{M} \mathbf{1}(p_{i,m} \in k), \tag{6}$$

$$\mathrm{Entropy}_{\mathrm{cross-subj}}(p,m) = -\sum_{m=1}^{M}\sum_{k=1}^{K} \Pr_{\mathrm{cross-subj}}\left(p_{i,m} \in k\right) \times \log \Pr_{\mathrm{cross-subj}}\left(p_{i,m} \in k\right), \tag{7}$$

$$\Pr_{\text{cross-subj}}(p_{i,m} \in k) = \frac{1}{100}\sum_{i=1}^{100} \mathbf{1}(p_{i,m} \in k), \tag{8}$$

where $\mathbf{1}$ represents the indicator function:

$$\mathbf{1}(p_{i,m} \in k) = \begin{cases} 1, & \text{if } p_{im} \in k \\ 0, & \text{if } p_{im} \notin k \end{cases} \tag{9}$$

Here, $\Pr_{\text{cross-state}}(p_{i,m} \in k)$ indicates the probability of node $p$ in iteration $i \in \{1, \ldots, 100\}$ and state $m \in \{1, \ldots, M\}$ to be assigned to network $k \in \{1, \ldots, K\}$, averaged over all states. Similarly, $\Pr_{\text{cross-subj}}(p_{i,m} \in k)$ indicates this probability averaged over all iterations. In this regard, every node falls into a two-dimensional space with cross-state and cross-subject entropies as its axes. We then grouped the nodes into three entropy classes, based on their position in this space: 1) *steady* nodes, those with zero cross-state and cross-subject entropies, 2) *flexible* nodes, those with positive (non-zero) cross-state entropies, and relatively low cross-subject entropies, and 3) *transient* nodes, regions with positive cross-state entropies and relatively high cross-subject entropies. To define a threshold to separate low and high cross-subject entropies, we took the average of this measure after excluding steady nodes ($\tau_{\text{mean}} = 29.47$). Eq. 10 formally defines the three entropy classes:

$$\text{Node class} = \begin{cases} \text{Steady,} & \text{if } \mathrm{E}_{\text{cross-state}} = \mathrm{E}_{\text{cross-subj}} = 0 \\ \text{Flexible,} & \text{if } \mathrm{E}_{\text{cross-state}} > 0 \text{ and } \mathrm{E}_{\text{cross-subj}} \leq \tau_{\text{mean}}, \\ \text{Transient,} & \text{if } \mathrm{E}_{\text{cross-state}} > 0 \text{ and } \mathrm{E}_{\text{cross-subj}} > \tau_{\text{mean}} \end{cases} \tag{2}$$

where E stands for Entropy.

We also performed a subject-based entropy analysis, where instead of performing a statistical ensemble analysis to create a pool of 100 population-level parcellations, we considered every single individual-level parcellation, resulting in a pool of size 718 (and thus replaced 100 with 718 in the equations above.) The subsequent analyses are based on the population-based entropies, and the subject-based entropies are included in the Supplementary Materials (Figure S5).

### 2.7 Functional characterization of entropy classes

To provide behavioral context for these entropy classes, we conducted a meta-analysis using the behavioral domains and paradigm classes reported in the BrainMap Functional Database (Fox and Lancaster, 2002). We examined the association between BrainMap's 55 behavioral domains and 108 paradigm classes (reflecting 10,467 task-activation experiments from 27,820 subjects) within the node entropy classes (steady, flexible, and transient), both at the node-level (with 268 nodes) and at the network-level (with 12 networks). The details are referenced in Supplementary Materials.

### 2.8 Graph visualization of the functional network reconfiguration across states

To provide a complementary analysis on how different networks interact with each other, we employed a graph visualization approach. While the delineation of functional networks represents the integration and segregation of different regions, functional connectivity analysis—using graph visualization—provides additional information on connections both within- and between-networks. This analysis further differentiates the notion of *community* and *connection,* as motivated by previous studies (Bassett et al., 2015; Mattar et al., 2015b). To this end, we constructed a population-level connectivity matrix for each state by taking the average of the individualized connectivity matrices. For the purpose of visualization, we summarized the average matrix to the top 10% edges and removed all self-loops. The resulting matrix for each functional condition was visualized as a graph using R package networkd3 [https://CRAN.R-project.org/package=networkD3]. These force-directed graphs aim to visually organize networks such that the energy of the graph as a whole is minimized. This is accomplished by assigning both repulsive and attractive forces to each pair of nodes such that the nodes with stronger interconnections are displayed closer to each other and the ones with weaker connections are shown as more distant.

### 2.9 Effects of head motion on the state-evoked network reconfigurations

Head motion is a known confound of connectivity analyses (Van Dijk et al., 2012). We ruled out the possibility that the network differences are derived by systematic differences in head motion across functional conditions. We employed a pairwise Wilcoxon signed-rank test on the mean frame-to-frame displacement (Movement_RelativeRMS.txt) across functional conditions and corrected for multiple comparisons using Bonferroni correction (see Figure S9).

## 3. Results

### 3.1 Functional network reconfigurations across states

Figure 1a displays state-specific population-level networks for 12 networks (K = 12) across nine functional conditions. While the overall network structure remained largely intact, specific nodes changed their network assignment in different task states. For example, the posterior areas of the caudate nucleus were associated with the ventral attention network during rest, but their network assignments changed to the language or the sensorimotor network during task performance (Figure 1a). To quantify network reconfiguration across states, we calculated the Hamming distance (Hamming, 1950)—or the number of nodes that change their NNAs—for each pair of states (Figure 1b). The two resting-states (REST1 and REST2) had the most similar functional organization with only a 5.6% (15/268) difference in their NNAs. In contrast, the smallest reconfiguration between any pair of states was over three times larger: WM and GAMBLING, 19% (51/268). GAMBLING and SOCIAL demonstrated the highest level of network reconfiguration (45.5% difference). SOCIAL was the most distinct with 41.9% network reconfiguration on average.

To quantify the level of state-evoked reconfiguration for each network we computed the ratio of the nodes that change their network assignments from REST1 to every other state (Figure 1c). Among all eight states, REST2 exhibited the maximum similarity to REST1, with 7 of 12 networks exhibiting no reconfiguration. Networks including subcortical/cerebellar nodes exhibited the maximum reconfiguration, while the visual networks (visual I and II), and the default mode network (core) exhibited the least reconfiguration from REST1 to all other states. The sensorimotor network displayed the maximum reconfiguration during MOTOR task, and the language network displayed the maximum reconfiguration during LANGUAGE task. Similarly, a majority of the higher-order association networks—including the frontoparietal network, the dorsal attention network, and the cingulo-opercular network—showed maximum reconfiguration during the SOCIAL cognition task (Figure 1a). These observations suggest that the task relevant nodes organize into specific networks in order to execute a particular task. These fuzzy network profiles are displayed in detail in Figure S1.

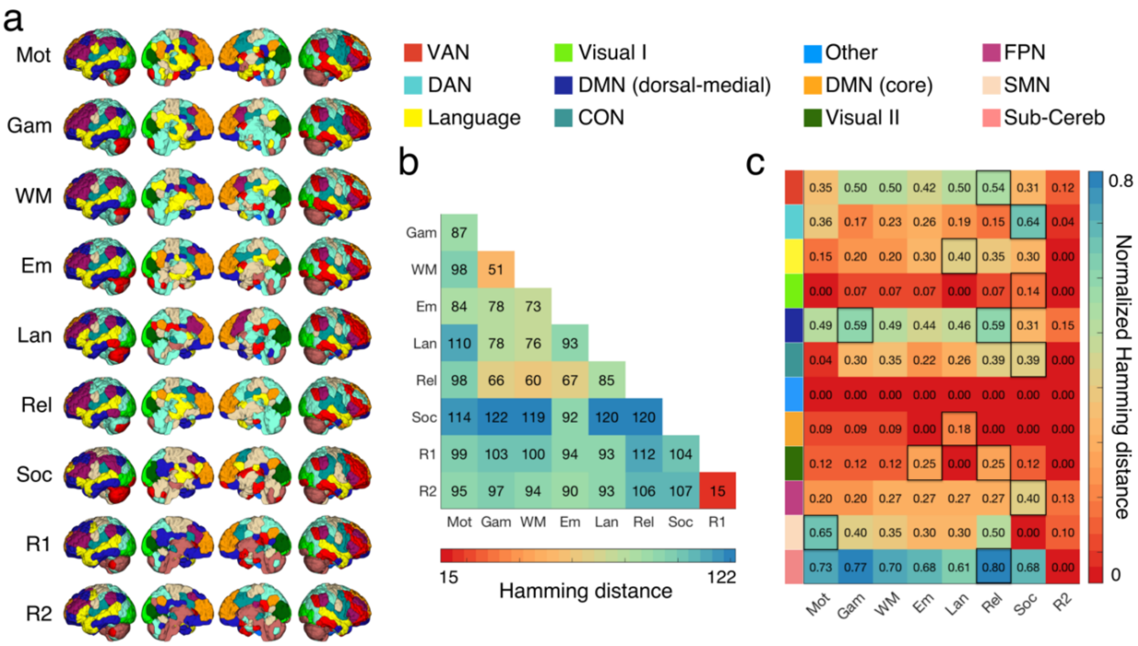

**Figure 1. Functional network reconfigurations across states.** a) Each row shows a state-specific population-level parcellation, calculated using the winner-takes-all strategy for each functional condition (Mot: MOTOR, Gam: GAMBLING, WM: WORKING MEMORY, Em: EMOTION, Lan: LANGUAGE, Rel: RELATIONAL, Soc: SOCIAL, R1: REST1, and R2: REST2). b) Hamming distance between every pair of state-specific population-level parcellation schemes is visualized. Hamming distance is calculated as the number of different NNAs between the two states, with values ranging from 0 (perfect matching) to the number of nodes (here, 268). As expected, the two resting states have the most similar parcellations (15/268 = 5.6% different NNAs). SOCIAL and GAMBLING are the least similar parcellations (122/268 = 45/5% different NNAs). SOCIAL task is the most distinct with 41.9% (=112.25/268) different NNAs on average. c) The ratio of the nodes that changed their NNAs from REST1 to every other state is calculated for every network separately. For every network, the state during which the maximum reconfiguration is observed is highlighted in black. VAN: ventral attention network, DAN: dorsal attention network, Visual I: primary visual, DMN: default mode network, CON: cingulo-opercular network, Visual II: secondary visual, FPN: frontoparietal network, SMN: sensorimotor network, Sub-Cereb: subcortical/cerebellum. See Figure S1 for a probabilistic illustration of these network definitions.

### 3.2 Node-to-Network Assignments (NNA) decode cognitive states

Figure 2 illustrates the accuracy of the predictive models that predict the cognitive state of unseen subjects, based on their NNAs. For the two-class (binary) classification (Figure 2a), the random accuracy was 50%. The minimum prediction accuracy was associated with the REST1-REST2 pair (accuracy=55%) consistent with the fact that these two runs reflect similar cognitive states. We successfully decoded all other pairs of states with accuracies considerably higher than random, with minimum, maximum, and average accuracy of 84%, 97%, and 94%, averaged over all $K$s from 2-50 (Figure 2a). Similarly, the 8-class classification accuracies were also significantly higher than random with minimum, maximum, and average accuracy of 66%, 87%, and 78%, averaged over all numbers of networks: $K$s from 2-50 (Figure 2b). That we could significantly predict cognitive state based solely on NNAs suggest that functional networks reorganize due to changes in state in a robust, reliable, and predictive manner across subjects.

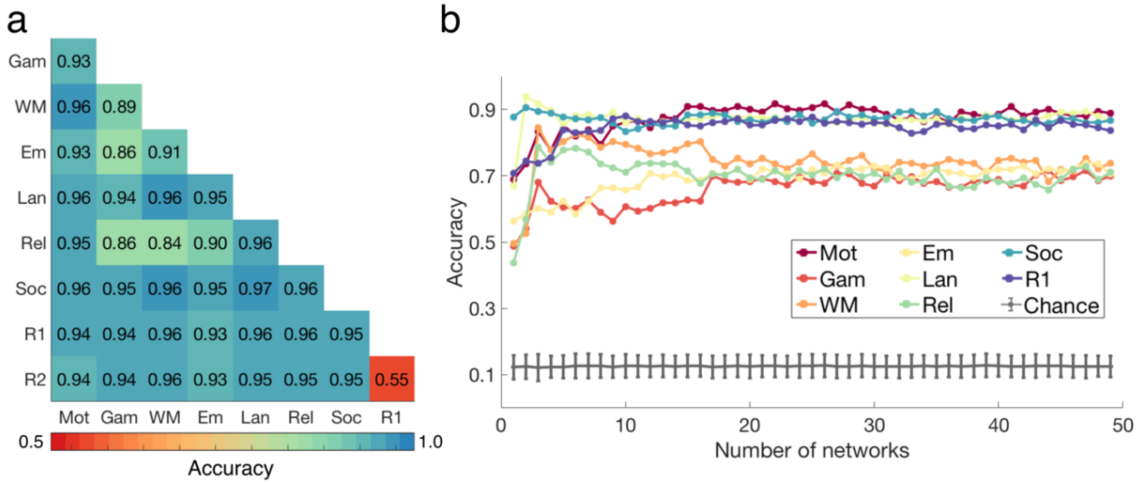

**Figure 2. Decoding cognitive states by node-to-network assignments (NNA) as features.** The subject pool was split into training and testing sets of equal-size. The exemplars (nodes with fixed NNAs both across states and subjects, representing each network) were computed using the training set and were used to parcellate the entire population. A GBM was trained on the training set using the NNA vectors as features. It was then used to predict the cognitive state of the novel subjects in the test set. a) The binary classification accuracies averaged over number of networks ranging from 2 to 50 are displayed. b) The 8-class classification accuracies are displayed for all number of networks ($K = 2 - 50$). The gray error bars represent the chance accuracy, computed as the mean and standard deviation of 1000 permutations. For $K$s larger than 17 the accuracies tend to stabilize with marginal variation. As expected, the minimum accuracies are observed for $K = 2$, since the entire vector only consists of two numbers: 1 or 2, depending on whether the node is assigned to network #1 or network #2. GAMBLING, RELATIONAL, EMOTION, and WM display lower prediction accuracies than MOTOR, SOCIAL, LANGUAGE, and REST. See Figure S4 for a 10-fold cross-validated predictive model with parcellation schemes derived from the entire population (not just the training set as shown above).

We also ruled out the possibility that the state-specific network reconfigurations are derived by the differences in the number of time points across different task conditions. To this end, we trimmed all functional runs to equal the length of the shortest task, (SOCIAL 176 time points), and repeated the analysis by generating the individualized and state-specific networks and re-running the steps in the predictive model. Similar networks reconfigurations and prediction accuracies were observed (Figure S10), suggesting that the differences in network configurations across states are not driven by differences in the amount of data, but reflect the underlying cognitive state.

### 3.3 Classifying nodes based on NNA reconfigurations

We next quantified the cross-state and cross-subject entropies for every node in the brain (Figure 3a). We observed that the two entropies were significantly correlated with each other ($r=0.78$, $p<1\times10^{-16}$), suggesting that the nodes with high cross-subject variation also tend to have high cross-state variance.

Next, we used the two entropy measures to categorize nodes into three entropy classes: 1) *steady* nodes, those with consistent cross-state and cross-subject NNAs (displayed as red); 2) *flexible* nodes, those with flexible cross-state NNAs but consistent cross-subject NNAs (displayed as yellow); and 3) *transient* nodes, those with flexible cross-state and cross-subject NNAs (displayed as green). Figure 3b illustrates the localization of these entropy classes on the brain. We observed that steady nodes were mainly located in the superior frontal cortex, dorsolateral prefrontal cortex (dlPFC), precuneus, post-central gyrus, superior and middle temporal gyrus, and the occipital lobe. Flexible regions were located mainly in the inferior frontal gyrus, pre-central gyrus, anterior cingulate cortex, posterior parietal cortex, inferior parietal lobule, and cerebellum (Figure 3b). Finally, transient nodes were located mainly in subcortical areas such as the hippocampus/parahippocampus, thalamus, and caudate. Figure 3c illustrates NNAs for the steady nodes.

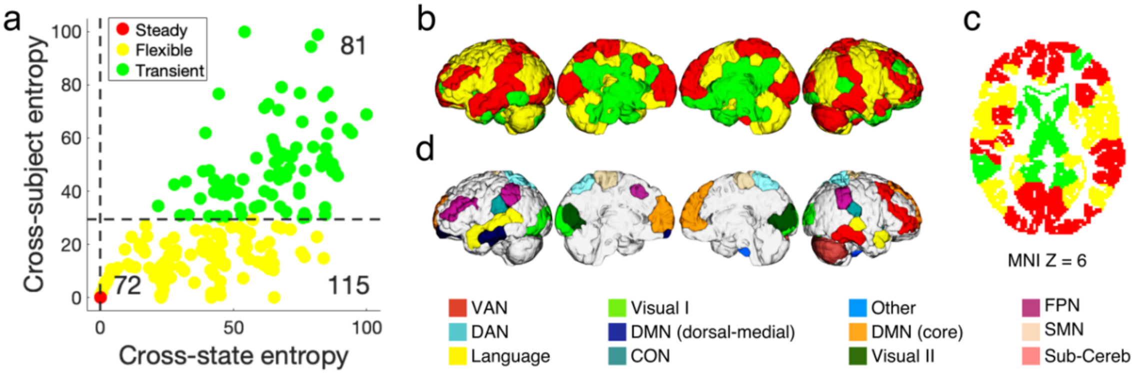

**Figure 3. Classifying nodes based on NNA reconfigurations.** a) For every node, we calculated two measures of variability: the NNA entropy across states (Entropy$_{\text{cross-state}}$) and the NNA entropy across subjects (Entropy$_{\text{cross-subject}}$). The two measures are significantly correlated (r=0.78, $p<1\times10^{-10}$). We grouped nodes into three entropy classes based on their position in the 2-dimensional entropy space: stable nodes (red, n=72/268); flexible nodes (yellow, n=115/268); and transient nodes (green, n=81/268). b) The three entropy classes are visualized on the brain. The steady class includes areas in superior frontal cortex, a large portion of the temporal and occipital lobe. The flexible class revolves around higher-order association areas including a large portion of the frontal and parietal lobes. The transient class is mainly localized in the subcortex and cingulate cortex. Also see Figure S5 for a more detailed illustration of these entropy measures, and Figure S6 for visualization of these flexibility classes as a continuum. c) The three entropy classes projected on the axial plane. d) Functional network assignments for the steady nodes, which are fixed across all functional conditions.

The current classification assigns nodes into discrete flexibility classes, which is suitable for visualization purposes. Figure S6 provides an illustration of these flexibility classes as a continuum, and Figure S5 visualizes the cross-state and cross-subject entropy values for all nodes using both population-based and subject-based analyses (see Methods for details).

### 3.4 Localization of the node entropy classes in networks

We next investigated the association between the three node entropy classes and our functional networks (Figure 4). We observed that all networks—excluding the visual I network—contained nodes of each class. However, the distribution of the three entropy classes differed considerably across networks (Figure 4a). For example, the dorsal attention network (76.6%), the cingulo-opercular network (56.5%), and the sensorimotor network (70.0%) had the majority of flexible nodes; the ventral attention network (57.7%), and the dorsal-medial portion of the default mode network (64.1%) had the most transient nodes; and the visual I network (62.5%), visual II network (64.3%), and the core portion of the default mode network (72.7%) contained most of the steady nodes. Finally, we observed that all functional networks contributed evenly to the steady class, but did not contribute evenly to the flexible or transient classes (Figure 4b). For example, the dorsal attention network contributed the most to the flexible class, while the dorsal-medial portion of the default mode network and the subcortical/cerebellum contributed the most to the transient class. This suggests that individual brain networks have varying levels of stability, as defined by NNA entropy across task states and subjects.

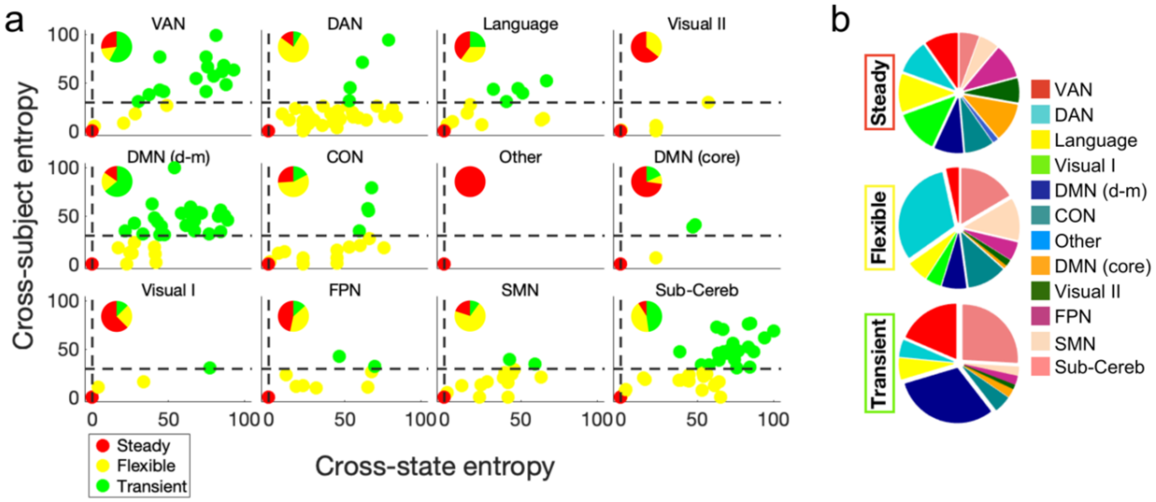

**Figure 4. Localization of the node entropy classes in networks.** a) The $\text{Entropy}_{\text{cross-state}}$ versus $\text{Entropy}_{\text{cross-subject}}$ diagram is displayed for the nodes within each network. The within-network distribution of the entropy classes is further demonstrated. Primary and secondary visual networks and the core portion of default network are mainly comprised of steady nodes. The DAN, CON, and SMN mainly include flexible regions, and the VAN, dorsal-medial DMN and subcortical/cerebellum networks are mainly comprised of transient regions. The FPN has a large portion of both steady and flexible regions, the subcortical/cerebellum network has a large portion of both flexible and transient nodes, and the Language network has close to an even distribution of all three entropy classes. b) The distribution of functional networks within each entropy class is visualized. While the contribution of functional networks to the steady class is evenly distributed across all networks, the DAN, CON, SMN, and subcortical/cerebellum have the largest contribution to the flexible class, and the VAN, dorsal-medial DMN, and subcortical/cerebellum have the largest contribution to the transient class. DMN (d-m): DMN (dorsal-medial).

### 3.5 Functional characterization of entropy classes

Next, we demonstrate that these entropy classes are functionally meaningful divisions. To this end, we provide evidence for the distinct functionality of these entropy classes.

#### 3.5.1 Flexible nodes inform state-decoding

We hypothesized that nodes in the flexible entropy class should have the highest predictive power for decoding states. We support this hypothesis with two lines of logic: first, because flexible nodes change their network organization according to task state, how the NNA reconfigures should therefore reflect the underlying state; second, because flexible nodes have consistent NNAs across subjects, they should generalize across novel subjects. To test our hypothesis, we computed how important each NNA feature was in the employed predictive pipeline. Node importance (i.e. "feature importance") was calculated by the GBM classifier as the number of times that each node was used to make a key decision that improved the classifier's performance measure (Friedman, 2001). We examined the distribution of the importance scores within each node entropy class. Consistent with our hypothesis, the flexible regions had the largest importance scores on average (Figure 5a; two-tailed t-test Bonferroni corrected for multiple comparisons, t(flexible, steady)=26.8, $p<3\times10^{-16}$, t(flexible, transient)=30.4, $p<3\times10^{-17}$, t(steady, transient)=0.34, $p<0.79$). That the task state can be decoded based on the NNA of the brain's flexible nodes suggests that these nodes are strongly modulated by the task demands such that the behavior of flexible nodes forms a robust and reliable signature of task state.

### 3.5.2 Steady nodes display the strongest behavioral associations

Similarly, we investigated the behavioral profiles of different entropy classes. We hypothesized that "steady" nodes should be more strongly associated with specific behavioral domains or paradigm classes. This is because steady nodes are less likely to "switch" association with task state and subjects, and therefore are more likely to be consistently reported among a significant number of task-based fMRI and PET experiments in BrainMap database, indicating a more specific brain-behavior association. We observed that indeed, there was a trend for steady nodes to have stronger behavioral domain and paradigm class loadings than either flexible or transient nodes (One-tailed t-test Bonferroni corrected for multiple comparisons, behavioral domain: t(steady, flexible)=1.5, $p<0.15$, t(steady, transient)=3.3, $p<0.001$; paradigm class: t(steady, flexible)=1.7, $p<0.09$, t(steady, transient)=4.6, $p<8\times10^{-6}$; Figure 5b, c). The same trend was observed at the network-level (Figure S7). This further suggests that entropy classes provide information that is relevant to how the brain executes its functional repertoire, as defined by behavioral domains and paradigm classes.

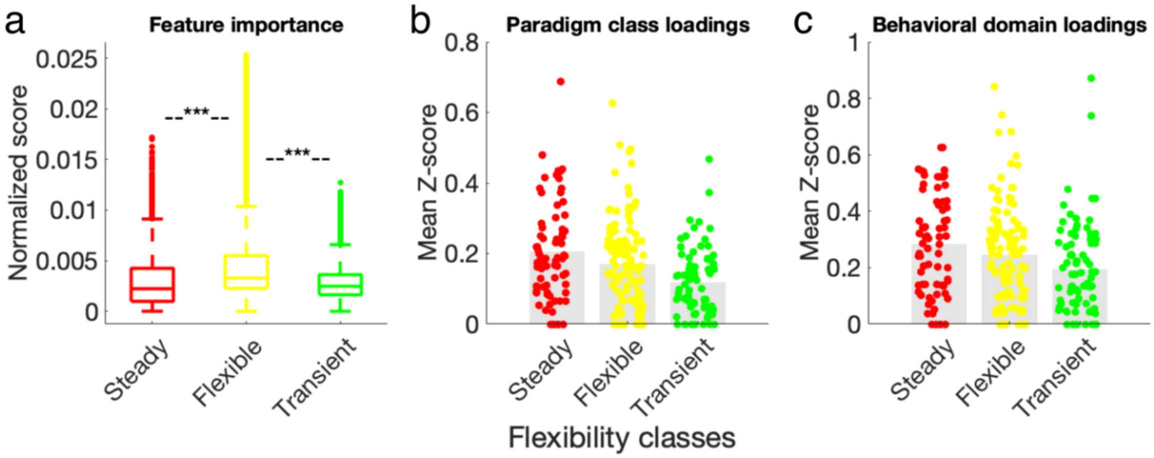

**Figure 5. Functional characterization of entropy classes.** a) The distribution of the feature (node) importance in predicting states is displayed across the three entropy classes. The flexible class contains the most important nodes for state decoding. Two-tailed t-test between every pair of steady, flexible, and transient class was performed and Bonferroni corrected for multiple comparisons, *** $p<1\times10^{-16}$. b) The distribution of paradigm classes across entropy classes is displayed. Steady nodes have higher values of significant paradigm class loadings than either flexible ($t=1.7$, $p<0.09$) or transient nodes ($t=4.6$, $p<8\times10^{-6}$). c) The distribution of behavioral domains across entropy classes is displayed. Steady nodes have higher values of significant behavioral domain than either flexible ($t=1.5$, $p<0.15$) or transient nodes ($t=3.3$, $p<0.001$). Analyses in (b) and (c) represent node-level analyses of BrainMap Functional Database's experimental meta-data representing 55 behavioral domains and 108 paradigm classes. Profiles were computed within each node as the forward inference likelihood z-scores, thresholded at $z>1.96$, and summed over all behavioral domains or paradigm classes. Colored circles represent the 268 nodes, colored according to their entropy classes. See Figure S7 for network-level functional relevance.

### 3.6 Graph visualization of the functional network reconfiguration across states

Finally, we evaluated the functional connectivity changes associated with changes in state, for comparison with studies that investigate connectivity changes as a proxy for network reorganization. Figure 6 (the right side of each panel) visualizes the connectivity matrices as force-directed graphs with nodes colored according to the functional network to which they belong and edges colored according to the entropy class of their two end points. Node size is proportional to the graph theory measure degree. We observed that nodes from different but related networks were integrated for all states. For instance, the core and dorsal-medial portions of the default mode network as well as the visual I and visual II networks showed strong integration with each other (Figure 6). However, the magnitude of this integration differed across different states. For example, while the visual I and visual II networks displayed strong integration across all states, the functional connectivity between them was weaker during the Social task. Finally, networks changed their segregation and integration patterns as a function of state. For instance, the dorsal attention network differed considerably across states, with higher segregation during the task states and higher integration during the rest states. Conversely, the subcortical/cerebellum and SMN presented higher segregation during the rest state and higher integration during the task states.

We next quantified the reorganization of functional networks across states (Figure 6; the bar plots on the bottom-left of each panel). The distribution of functional networks changed considerably across states, with more uniform distribution during rest than task. In this sense, the dorsal attention network was the largest network (with the maximum number of nodes) during GAMBLING, WM, EMOTION, LANGUAGE, and

RELATIONAL states. The subcortical/cerebellum network was considerably larger during rest than task, suggesting that nodes in the subcortical/cerebellum network tend to integrate with other networks during the execution of a task, while displaying segregated behavior during rest (consistent with our observation from the graph). The sensorimotor network lost the maximum number of nodes during the MOTOR task, while the language network lost the maximum number of nodes during the LANGUAGE task. While this observation may seem contradictory to the demands of the tasks, it could be an indication that these networks tend to integrate most with other networks when the task strongly engages them, potentially facilitating the information flow across networks (Cole et al., 2016; Ito et al., 2017). This is also consistent with our earlier finding (Figure 1c) that these networks displayed the maximum reconfiguration in changing to these tasks.

To investigate the role of entropy classes in the context of functional connectivity, we quantified the connectivity (i.e. edge strength) between nodes from the same or different entropy classes (Figure 6; the error bars on the top-left of each panel): steady-steady edges (colored as red), steady-flexible edges (colored as dark gray), steady-transient edges (colored as gray), flexible-flexible edges (colored as yellow), flexible-transient edges (colored as light gray), and transient-transient edges (colored as green). We observed that the edge strength between steady nodes was significantly stronger than the edge strength between other node classes ($p<2\times10^{-16}$, Bonferroni corrected for multiple comparisons). Among edges between nodes of the same class, edges between steady nodes were significantly stronger than the edges between flexible nodes ($p<2\times10^{-16}$, Bonferroni corrected for multiple comparisons) and edges between flexible nodes were significantly stronger than edges between transient nodes ($p<1\times10^{-2}$, Bonferroni corrected for multiple comparisons). Among edges between nodes of differing class, the edges between steady and flexible nodes were the strongest ($p<2\times10^{-2}$, Bonferroni corrected for multiple comparisons). Edges associated with transient nodes had the lowest strength.

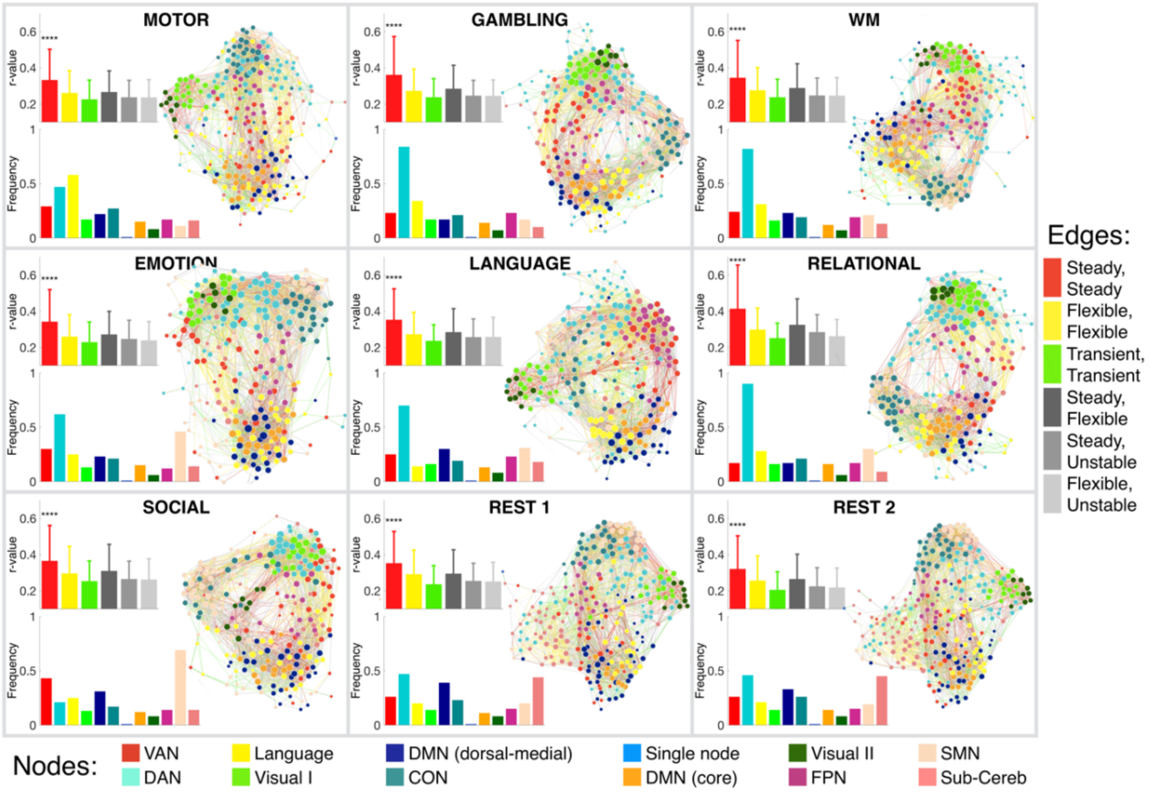

**Figure 6. Graph visualization of the functional connectivity reconfiguration across states.** Each panel corresponds to one functional condition. For each functional condition, three diagrams are displayed. The right diagram is the force-directed graph visualization of the functional connectivity matrix, thresholded to retain the top 10% of the edges, and averaged over all subjects. Nodes are colored according to their population-level network assignments (see the bottom legend bar), and are sized proportional to their degree values (that is, the sum of edge weights connected to the node). Edges are colored according to the entropy class of their two end points (see the right legend bar), for example the connections between two steady nodes are colored red, connections between a steady node and a flexible node are colored light gray, and so on. The graphs are structured such that nodes with stronger connections are spatially closer to each other. As expected, nodes within the same network are clustered spatially close to each other. Further, nodes from different but related networks are also spatially integrated (for example, the core and dorsal-medial portions of DMN, or primary and secondary visual networks are integrated). However, there is a difference in the integration and segregation level of different functional networks across different task states. The bottom left diagram within each panel represents the frequency of the nodes assigned to each functional network. The top left diagram within each panel represents the strength of different types of connections. Error bars represent the standard deviation across connections. The connections within steady nodes are significantly stronger than the rest of the connections. These steady nodes form a strong core organization in the brain with large interconnections, similar to the "rich club" pattern (see Figure S8).

## 4. Discussion

In this work, individualized and state-specific functional networks were used to study network reconfigurations evoked by different task and rest states. Our findings highlight that the brain's network organization is not fixed, but rather, that networks reconfigure as a function of cognitive state. We showed that this holds true across different task states and across individuals. We measured network reconfiguration as node-to-network assignments (NNA) and demonstrated that a novel subject's current state could be predicted with 66-97% accuracy based on NNA alone (Figure 2). This finding demonstrates that networks reconfigure in a meaningful and reproducible manner across cognitive states. This also highlights the robustness of state-evoked network reconfigurations across subjects. Further, we showed that nodes group into three entropy classes based on their NNAs across states and subjects (Figure 3). Steady nodes exhibit consistent NNAs across both states and subjects and are primarily located in visual and medial cortical regions. Flexible nodes change their network assignments according to state in a consistent manner across subjects and are primarily located in higher order cognitive regions. Transient nodes exhibit variable network assignments across both states and subjects and are primarily located in the sub-cortex and cerebellum; regions known to have lower reliability in functional connectivity studies (Noble et al., 2017). A direct assessment of reliability however, showed that there are no significant differences in reliability values across the three flexibility classes (Supplementary Materials, Figure S11).

Despite these trends of anatomical locations of nodal classes, all networks contain some nodes from each class. We further demonstrated the functional relevance of these nodal classes: flexible nodes contribute the most to models predicting cognitive state and steady nodes display the largest behavioral loadings based on a large-scale meta-analysis of task activation studies (Figure 5 and Figure S7). Together, these results demonstrate that instead of simply changing their functional connectivity due to task, nodes reorganize into specific networks associated with specific tasks, with the networks including or dropping nodes according to the demands of a particular task state. In a recent work, Bijsterbosch et al. have shown that cross-subject changes in spatial configuration of functional regions can be interpreted as changes in functional connectivity (Bijsterbosch et al., 2018). Similarly, the observed cross-state changes in network definitions are likely to be interpreted as changes in connectivity values across states. As such, the state-evoked reconfigurations need to be considered when interpreting changes in connectivity or any graph theory measures across states.

### 4.1 Functional network organization is state-dependent

Our results suggest that a similar core network structure is observed across many different cognitive states, but that nodes can be included or dropped from these networks as a function of the specific state. These results suggest that there is not a fixed functional network structure across states. Furthermore, our BrainMap meta-analysis shows that nearly all nodes were significantly associated with multiple paradigm classes, supporting the flexibility of nodes in executing functions. Altogether, our results suggest that the node-to-network assignments are inherently probabilistic as the functional organization of the brain is constantly reconfiguring with task (Figure 1 and Figure S1).

Numerous previous studies have focused on fixed networks in various ways, typically by using coordinates from previous resting-state studies (Cassady et al., 2019; Cole et al., 2014; Cole et al., 2016; Gratton et al., 2018; McMenamin et al., 2014; Mohr et al., 2016; Satterthwaite et al., 2015; Schultz and Cole, 2016; Telesford et al., 2016). Several studies, for example (Cole et al., 2013; Zanto and Gazzaley, 2013), have examined a single network, the frontoparietal network, and measured changes within this network as a function of task-state. While such an approach can be justified, it potentially misses additional information that could be gleaned in considering how the network itself reconfigures. The sensorimotor network has also been studied as an implicitly defined fixed functional network (Cassady et al., 2019; Power et al., 2011). While it is true that there is a core network, the findings presented here illustrate that the sensorimotor network is highly variable across states (Figure 4a), revealing substantial state-evoked reconfigurations that have not been previously characterized. Perhaps as one would expect, the largest reconfiguration of the sensorimotor network occurs when subjects are executing the motor task (Figure 1c). The sensorimotor network had the lowest portion of transient regions, suggesting the low cross-subject variability, consistent with previous observations (Finn et al., 2015; Mueller et al., 2013).

In addition, our results also replicate previously observed network reconfigurations. Consistent with Andrews-Hanna et al. (Andrews-Hanna et al., 2014), we highlight that there is a core sub-system of the default mode network (including medial superior frontal gyrus) and dorsal-medial and medial-temporal sub-systems of the default mode network (including ventral medial prefrontal cortex (vmPFC), inferior temporal gyrus and temporal pole). Our results extend this characterization by suggesting that the two subsystems of the default mode network display distinct reconfiguration behaviors, with the core subsystem remaining unchanged across states, and the dorsal-medial subsystem reconfiguring more flexibly across states and subjects (Zhang et al., 2016). Highlighting the complex role of the default mode network in both tasks and rest, our observations indicate that the default mode network nodes change their network assignments from the default mode network to other higher-order association networks (such as the frontoparietal network and the cingulo-opercular network). Overall, these observations support the notion that all networks reorganize due to task, likely in a task specific manner.

### 4.2 Functional networks contain steady, flexible, and transient nodes

Our findings reveal that all networks contain a combination of steady, flexible, and transient nodes (Figure 4a). This suggests that networks consist of a core set of steady nodes that then recruit and dismiss flexible and transient nodes according to the task demands. This result suggests that within each network there are specialized regions associated with distinct functionalities, altering both intra-network communication and inter-network integration. These functionalities are corroborated by the theory of “local” versus “distributed” neural communication (Cole et al., 2016; Ito et al., 2017), where transitions across states require the segregated processing units (distinct functional networks) to integrate information through flexibly changing network assignments.

The contribution to the steady nodes was evenly distributed across all networks (Figure 4b), indicating all networks retained some core functional network configuration. The majority of steady nodes were located in regions previously been implicated as part of the “rich club”, a common organization in complex systems where important (or “rich”) nodes connect preferentially to other important nodes (Grayson et al., 2014; Park and Friston, 2013; Van Den Heuvel and Sporns, 2011) (see Supplementary Materials, Figure S8). Additionally, we observed strong functional connectivity between nodes in the steady class (Figure 6, the top left error bars) in a similar vein to the connections between rich club nodes. These observations were consistent across all nine states (Figure S8), implying that the “rich club” organization of the brain is independent of task-state. This also mirrors a core-periphery structure in the brain with core nodes strongly and mutually interconnected and periphery nodes sparsely connected to each other and to the core nodes (Bassett et al., 2013; Fedorenko and Thompson-Schill, 2014; Park and Friston, 2013).

In contrast, the distribution of flexible nodes varies across networks, suggesting that some networks reconfigure more with cognitive state changes than others (Figure 4b). For example, the DAN, CON, and FPN had a larger proportion of flexible nodes (Figure 4a, b). Previous studies report that these networks rapidly update their connectivity patterns according to the task context (Anderson et al., 2013; Cole et al., 2013; Crossley et al., 2013; Krienen et al., 2014; Mennes et al., 2012), facilitating the information flow across networks (Cole et al., 2016; Ito et al., 2017). The large number of nodes that change their network assignments to the DAN, CON, and FPN may underlie these previously observed changes in connectivity patterns. Our results are also consistent with the role of these networks as functional hubs (Cole et al., 2013; Power et al., 2013) and their strong out-of-network connections (Ito et al., 2017).

Finally, the majority of the transient nodes were assigned to the subcortical/cerebellum, dorsal-medial subsystem of the default mode network, and the ventral attention network (Figure 4a, b). These nodes, mainly located in heteromodal association and limbic cortex, actively change their network assignment across both subjects and states. High cross-subject variability in these regions has been previously reported (Mueller et al., 2013; Zhang et al., 2016), and linked to differences in intelligence (Li et al., 2009), attention (Cohen et al., 2015; Rosenberg et al., 2016), and personality traits (Adelstein et al., 2011). This suggests that these regions promote a more personalized reconfiguration in the brain to adapt to the task at hand. Future work could explore the extent to which differences in the network reconfiguration in these regions could be indicative of, or be inferred from, the individual differences in task performance.

### 4.3 Functional network configuration predicts state

Individual cognitive states can be predicted from network node assignments, suggesting that the state-specific information of network organization is sufficiently robust and reliable to form a signature for a given state. Prior work on task decoding has typically employed binarized classification of tasks. The pair-wise accuracies (binarized classification) achieved here (ranging from 84% to 97% with average of 94%; Figure 2a) were significantly higher than those previously reported (Cole et al., 2011; Heinzle et al., 2012; Woolgar et al., 2011). Our accuracies for the 8-class classification (ranging from 66% to 87% with average of 78%; Figure 2b) were also higher than the accuracies reported in the literature, even for easier decoding problems such as classification with lower number of classes (Cole et al., 2013) or within-subject classification paradigms (Haxby et al., 2014). Nevertheless, the main focus of the predictive analysis was to demonstrate that state-evoked demands lead to different network reconfigurations, and functional connectivity studies should consider these reconfigurations when studying within- and between-network connectivity patterns. In fact, even though the network node assignment vector (of length 268) contains only integer values (i.e., 1, …, K for K networks), this minimal information was sufficient to predict the cognitive state (task) of novel subjects. This highlights the reliability of these state-specific reconfigurations of the brain's large-scale network organization.

### 4.4 Individualized functional networks are needed

In addition, our results highlight the need for individualized approaches to systems neuroscience and medicine (Perez Velazquez, 2017; Satterthwaite and Davatzikos, 2015) by demonstrating that node-to-network assignments contain large individual variations across states (i.e. our transient node class). Previous work has shown promise in defining network boundaries at the individual level (Braga and Buckner, 2017; Braga et al., 2019; Gordon et al., 2017b; Kong et al., 2018; Laumann et al., 2015; Wang et al., 2015). Individualized networks are shown to be highly reliable (Laumann et al., 2015), closely correspond with task activations (Gordon et al., 2017b), and behaviorally meaningful (Kong et al., 2018). Our work builds on this literature and extends it by highlighting that even within an individual the functional networks reliably reconfigure as a function of cognitive state. Such state-specific individualized networks could provide input for real-time fMRI neurofeedback paradigms (Emmert et al., 2016; Hartwell et al., 2016; Koizumi et al., 2017), and brain stimulation therapies (Fitzgerald, 2011; Plow et al., 2016), focusing on the disruption of networks at the single subject level. The conventional approach of defining fixed networks via population-level parcellations may limit the efficacy of such treatments. State-specific individualized networks also may provide higher specificity when considered as predictive features in machine learning algorithms (Bzdok and Meyer-Lindenberg, 2017; Vu et al., 2018).

In a recent work, Gratton et al. have studied the magnitude of functional network variability across subjects, tasks, and sessions (Gratton et al., 2018). Their study of brain reconfiguration was restricted to changes in functional connectivity matrices as a whole (i.e., correlation of the vectorized connectivity matrices) with the core finding that connectivity matrices are remarkably stable. Our work differs in that we study the large-scale network reconfigurations by probing the changes in node-to-network assignments across subjects and states. These two divergent but complementary views shed light on different aspects of the brain dynamics. While their analysis has shown that functional connections are largely stable with subtle cross-state changes (in terms of correlating vectorized matrices), here we demonstrate that about 73% (i.e., 196 out of 268) of the nodes actually change their network assignment as a function of state. In addition, their observations of functional connectivity matrices indicate that the state-evoked modulations are largely individual-specific (i.e., the consistent state-specific changes are subtle), whereas we demonstrate that state-specific changes in NNAs are sufficiently robust and reliable across subjects to predict the cognitive state. Combined, these observations suggest that while gross matrix level connectivity patterns change only minimally with state (Gratton et al., 2018) consistent with earlier functional fingerprinting findings across state (Finn et al, 2015), individual node to network assignments are flexible, allowing the brain to reconfigure its functional organization as the task demands change. Such reconfigurations are easily measured and change in a significant and meaningful manner.

### 4.5 Further considerations

Although cognitive-state dependent network reconfigurations were investigated here, we did not investigate the temporal dynamics within a given task. This was partly because the length of the sessions in HCP data vary (Uğurbil et al., 2013; Van Essen et al., 2013) and because of methodological concerns surrounding dynamic brain-state definitions, that is, we currently lack a strict way of identifying and validating within-task cognitive states which are consistent across all subjects. However, it is likely that the same reconfiguration is observed dynamically, on a moment-to-moment basis within each condition, specifically for the HCP tasks, which were designed to tap different cognitive processes even within a single task run (Barch et al., 2013). Future work could extend this framework to study the within session dynamics of network reconfiguration.

We employed a functional atlas consisting of 268 nodes, generated from an independent group of healthy subjects (Finn et al., 2015; Shen et al., 2013). This atlas has proven to be reproducible and reliable (Shen et al., 2013), beneficial in understanding cross-subject variability, and useful for developing predictive models of behavior (Finn et al., 2015; Rosenberg et al., 2013; Rosenberg et al., 2017; Shen et al., 2017). Given this past performance, this 268-node atlas represents a reasonable way to operationalize the brain's sub-units in network analyses. However, this atlas was defined at the group-level and therefore might not account for the individual variability in node definitions. To verify that the observed network reconfigurations are not confounded by the group-defined nodes, we repeated our analyses, using individualized node atlases as our starting point (see Supplementary Materials). Similar results were obtained (Figure S12), suggesting that the state-evoked network reconfigurations are not derived by the incongruency of signals within a node, but are robust observations of brain function. We would anticipate similar results with other atlases of a similar scale.

Just as the core functional connectivity patterns across the whole-brain are stable, underlying these functional patterns is a relatively fixed structural connectivity infrastructure. That is, the changing functional patterns on a moment to moment or task to task basis, are not due to rapidly changing structure. While there is a large literature where investigators have performed network analysis on structurally defined networks these two approaches are not mutually exclusive and this flexible functional organization is acknowledged to sit on top of a relatively fixed structural network.

Moving forward it will be important to relate state-evoked network reorganization with individual differences in behavior and/or clinical measures. Future work could quantify the extent to which subjects with similar node to network reconfiguration patterns are also similar in state, behavioral, or clinical measures; potentially grouping homogeneous subjects based on their NNA phenotypes. This methodology could be particularly compelling when applied to subjects with different clinical, neurodevelopmental, and aging categories.

## 5. Conclusion

We showed that the functional network organization of the brain is not fixed, but rather that functional networks reconfigure dynamically, in predictable ways, as a function of cognitive state. Future work should consider state-evoked reconfigurations when studying changes in connectivity across different task-states. Such an approach will hopefully allow us to begin to relate individual differences in network reconfiguration to individual differences in behavior or clinical symptoms.

## 6. Code availability

The 268-node functional parcellation is available online on the BioImage Suite NITRC page (https://www.nitrc.org/frs/?group_id=51). MATLAB, R (for graph visualization), and Python (for predictive modeling) scripts were written to perform the analyses described; these code are available on GitHub at https://github.com/YaleMRRC/Network-Parcellation-Tasks.git. The graph visualization is released separately under the terms of GNU General Public License and can be found here: https://github.com/YaleMRRC/Network-Visualization.git.

# 7. Supplemental Materials

## 7.1 Local versus global exemplar set

In the main analysis, exemplars were identified such that they were common across all functional states (denoted as “global exemplars”). This was to assure that the exemplars are agnostic to the states, and thus the proceeding predictive analyses were not biased by the prior knowledge of the state. In order to assure that the results were not confounded by introducing common exemplars across all states, we repeated the exemplar-based parcellation, this time restricting the analysis to each state separately. This approach yielded nine different sets of exemplars (denoted as “local exemplars”), one for each functional state. We compared the parcellation results from global and local exemplars using two separate measures. First, we calculated the Dice coefficient between the two group-level parcellation vectors (Figure S2a) for $K = 2 - 50$. Second, we computed the Hamming distance between the community co-membership matrices derived from each parcellation. For every subject in each state, a network co-membership matrix ($A$) was constructed from the two parcellation schemes: (i) using global exemplars ($A_{\text{global}}$), and (ii) using local exemplars ($A_{\text{local}}$). The elements of this matrix were calculated as follows:

$$a_{ij} = \begin{cases} 1, & \text{if node } i \text{ and node } j \text{ were assigned to the same network} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

The similarity between the two co-membership matrices ($A_{\text{global}}$ and $A_{\text{local}}$) were calculated as 1 minus the Hamming distance between the vectorized form of the two matrices. The result is displayed in Figure S2b for $K = 2 - 50$.

### 7.2 Number of networks ($K$) selected for the analysis

To choose the optimum number of networks, we examined the changes in network assignments by adding a new exemplar (or equivalently, adding a new network). For every state, we calculated the number of nodes that change their network membership as the number of exemplars, $K$, increases from 2 to 50. This was employed for every subject and state separately (Figure S3). Among the $K$s for which the number of changes were locally maximum, i.e. $K = \{12, 17, 31, 45\}$, we chose $K = 12$ as the minimum number of networks after which the memberships stabilized with significantly less changes in the NNAs.

### 7.3 Distribution of behavioral domains and paradigm classes across entropy classes

Behavioral domain and paradigm class profiles were defined by referencing the BrainMap Functional Database's experimental meta-data [www.brainmap.org], which represents 62,038 x-y-z foci from 10,467 functional human brain imaging task-activation experiments representing 27,820 subjects. These task-activation experiments represent 55 behavioral domains and 108 paradigm classes (Barron and Fox, 2015; Fox et al., 2005). Behavioral domains include categories and subcategories of mental processes isolated by the experimental contrasts. They comprise five main categories: cognition, action, perception, emotion, and interception. Paradigm classes include the experimental tasks isolated by an experimental contrast (see http://brainmap.org/scribe for more information on the BrainMap taxonomy). Profiles were computed within each node (or network) as a z-score of the forward inference likelihood that a particular behavioral domain or paradigm class is reported within that node (or network) [*P(Activation* | *Behavioral*

*Domain) or P(Activation | Paradigm Class)*] normalized by the likelihood expected if they were uniformly distributed throughout the brain (Lancaster et al., 2012). A high z-score indicates a high specificity of a particular behavioral domain or paradigm class for that node (or network). To create a profile for each node (or network), z-scores were thresholded at z>1.96 (associated with 95% confidence interval or p<0.05) and summed over all behavioral domains or paradigm classes. Profiles were computed using MATLAB.

### 7.4 Rich club organization of the brain across states

Given the observed strong connectivity between the steady nodes (Figure 6), we hypothesized that steady nodes contribute most to the rich club organization of the brain (Van Den Heuvel and Sporns, 2011). To test our hypothesis quantitatively, we computed and visualized the rich club regime of the brain network for every functional state. For every state, we computed to what extent high-degree nodes link preferentially to each other, as measured by rich club coefficient (Colizza et al., 2006; Van Den Heuvel and Sporns, 2011), $\phi(k)$:

$$\phi(k) = \frac{2E_{>k}}{N_{>k}(N_{>k} - 1)}, \tag{2}$$

where $E_{>k}$ is the number of edges in the subgraph containing all the nodes with degree greater than $k$, and $N_{>k}$ represents the number of nodes with degree greater than $k$. Consequently, we compared the $\phi(k)$s with the rich club coefficients obtained from an ensemble of 1000 random networks ($\phi_{\mathrm{rand}}(k)$), computed by shuffling the edges of the original network while preserving the degree distributions. We computed and visualized the normalized rich club coefficients by dividing the empirical $\phi(k)$ with the average of $\phi_{\text{rand}}(k)$:

$$\phi_{\text{norm}}(k) = \frac{\phi(k)}{\text{avg}\,[\phi_{\text{rand}}(k)]}. \tag{3}$$

Rich club coefficients, $\phi(k)$, were calculated using "rich_club_wu" function, and random networks were calculated using "randmio_und" function (average number of rewiring = 50) from Brain Connectivity Toolbox (Rubinov and Sporns, 2010).

For every state, we computed the rich club regime (Figure S8a, the shaded gray area) as the range of $k$s for which the rich club coefficient of the empirical network was significantly larger than the null distribution (permutation test; $p<0.05$). To illustrate the rich club regions in every state, we used the average $k$ within the rich club regime of every state ($k_{\text{threshold}}$: Figure S8a, the triangle marker).

To demonstrate that steady nodes contribute most to rich club regime, we computed the ratio of the rich club nodes ($k > k_{\text{threshold}}$) within each entropy class (Figure S8b). As displayed in Figure S8b, the steady class has the maximum ratio of the rich club nodes, highlighting the significant contribution of the steady nodes to the rich club organization of the brain in all states.

### 7.5 Individualized node atlases

To rule out the possibility that the network reconfigurations are simply driven by group-defined nodes with mixed signals, we repeated our network delineation analysis with individualized node definitions. We created the individualized atlas for every individual in HCP data set by employing our recently developed parcellation algorithm using data from REST1 and REST2 runs with left-right phase-encoding (LR). The algorithm starts from the initial group atlas and individualizes it using a variant of exemplar-based clustering method (see (Salehi et al., 2018) for more details). One advantage of this algorithm is that correspondence of nodes across individuals and with the initial group atlas is maintained, making for straightforward comparisons.

We repeated our network delineation algorithm with individualized nodes. Similar network reconfigurations were observed (Figure S12a). To verify that these network reconfigurations are specific to each state, we repeated our cross-validated predictive model to predict (or decode) the cognitive state of each individual based solely on the NNAs, which are generated from individualized nodes. Again, significant prediction accuracies were obtained (Figure S12b), suggesting that the network reconfigurations are reliable and state specific, even when starting from individualized nodes.

### 7.6 Reliability comparison across flexibility classes

We investigated the relationship between the flexibility classes and the test-retest reliability of functional connectivity. Reliability was measured by the intraclass correlation coefficient (ICC; (Shrout and Fleiss, 1979). Following the work by Noble et al. (Noble et al., 2017), we estimated the reliability values for every node in the atlas by taking the mean test-retest reliability of all edges associated with that node. We then compared the reliability values between every pair of Steady, Flexible, and Transient classes using two-tailed Wilcoxon rank sum test and corrected for multiple comparisons using Bonferroni correction (Figure S11). We observed that there is no significant difference in the ICC values across the three flexibility classes (Figure S11).

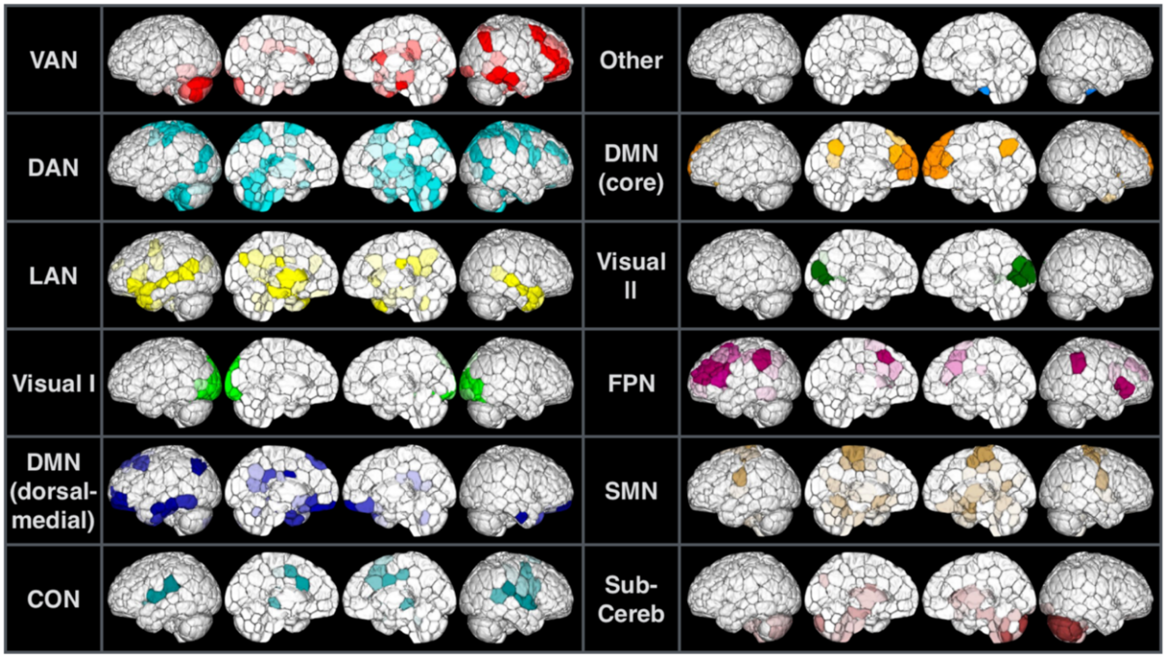

**Figure S1. The fuzzy (probabilistic) network definitions for $K = 12$ networks.** Every network is displayed by opacity of its nodes reflecting the likelihood of their NNA. The likelihood of each NNA is defined as the number of times this assignment is observed across all population-level state-specific parcellations, divided by the number of states (i.e., 9).

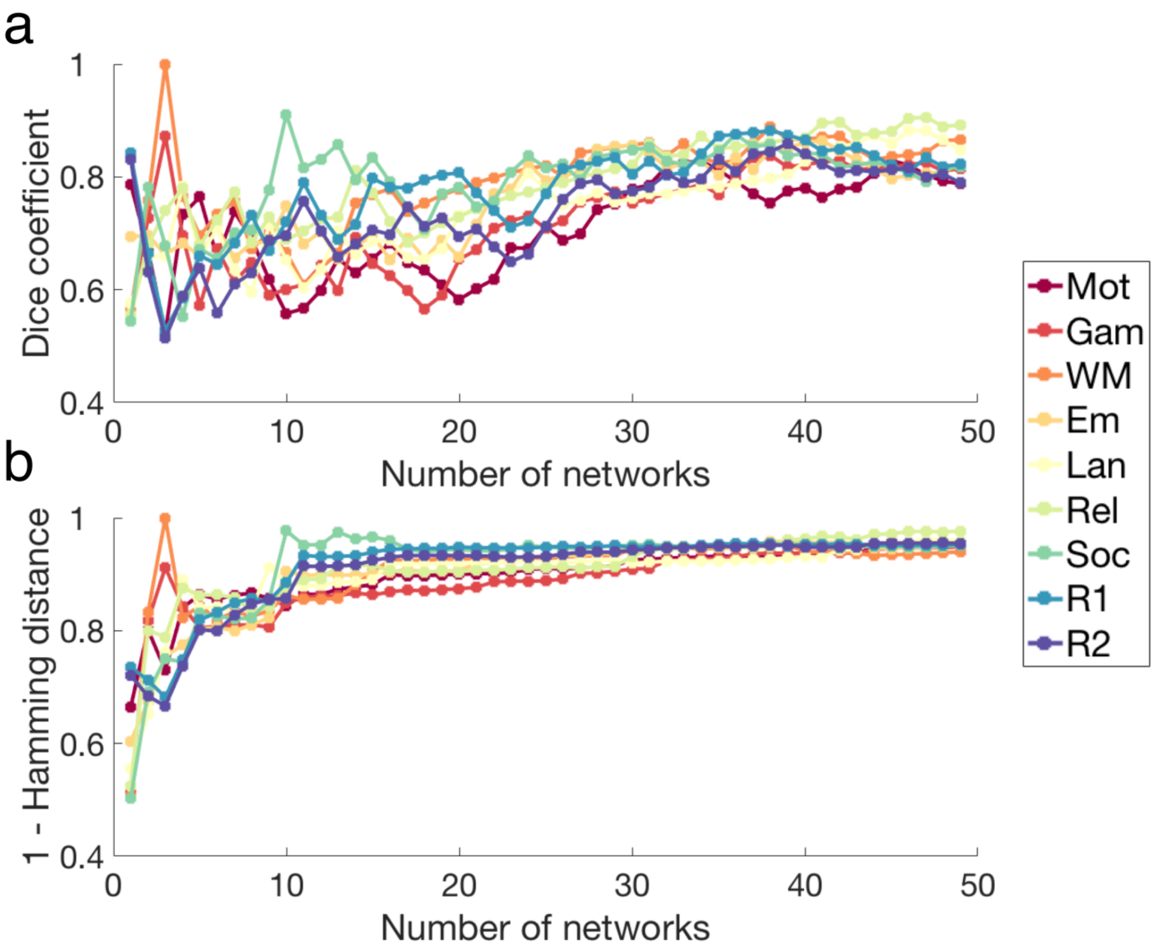

**Figure S2. Similarities between the parcellation schemes calculated from global and local exemplars.** a) Dice coefficient between the parcellation schemes, computed by local and global exemplars, are displayed for every state. b) 1 – Hamming distance between the node-to-network co-membership matrices, computed by the global and local parcellation schemes, are displayed for every state.

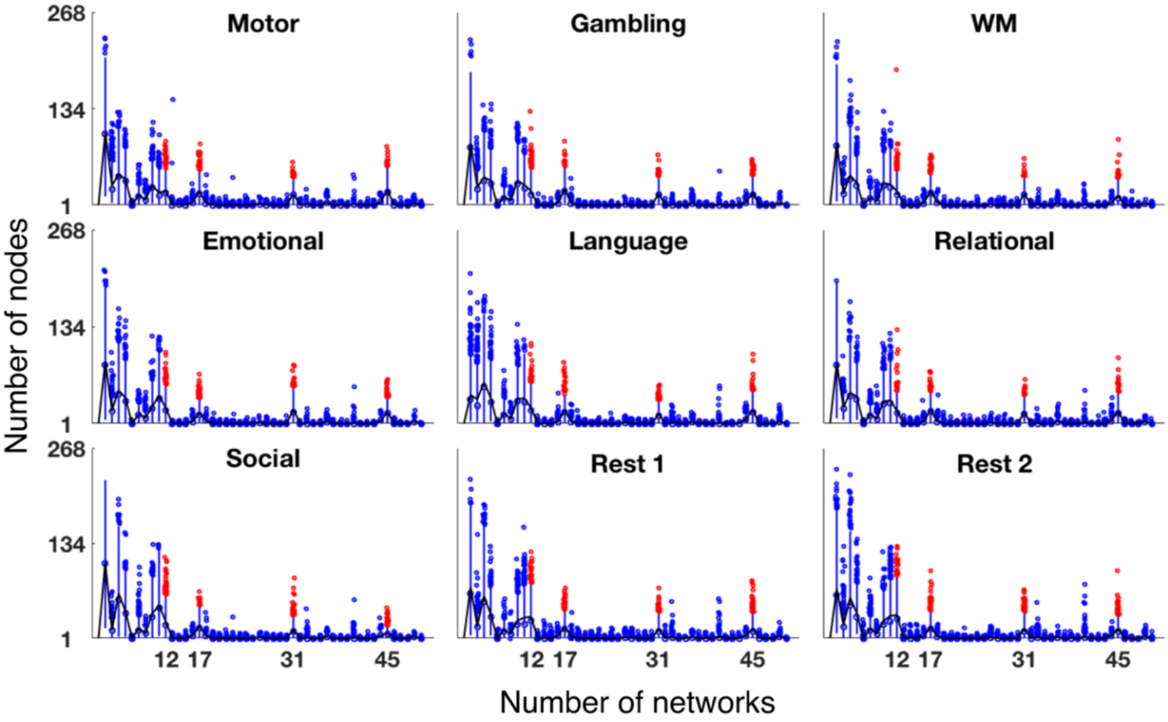

**Figure S3. The number of changes in the network memberships for different $K$ values.** For each functional state, the number of nodes that change their NNA when a new exemplar (representing a new network) is calculated. The number of changes is equal to the number of nodes that are assigned to the newly defined network.

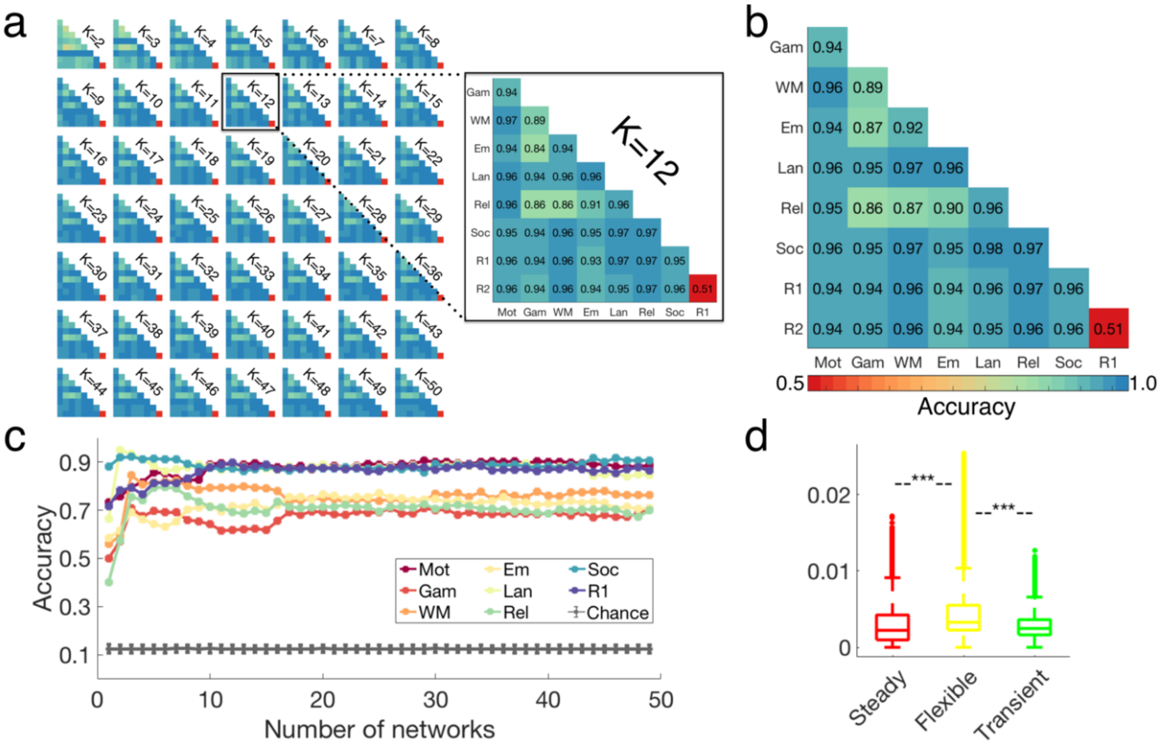

| | Mot | Gam | WM | Em | Lan | Rel | Soc | R1 |
|---|---|---|---|---|---|---|---|---|
| Gam | 0.94 | | | | | | | |
| WM | 0.97 | 0.89 | | | | | | |
| Em | 0.94 | 0.84 | 0.94 | | | | | |
| Lan | 0.96 | 0.94 | 0.96 | 0.96 | | | | |
| Rel | 0.96 | 0.86 | 0.86 | 0.91 | 0.96 | | | |
| Soc | 0.95 | 0.94 | 0.96 | 0.95 | 0.97 | 0.97 | | |
| R1 | 0.96 | 0.94 | 0.96 | 0.93 | 0.97 | 0.97 | 0.95 | |
| R2 | 0.96 | 0.94 | 0.96 | 0.94 | 0.95 | 0.97 | 0.96 | 0.51 |

| | Mot | Gam | WM | Em | Lan | Rel | Soc | R1 |
|---|---|---|---|---|---|---|---|---|
| Gam | 0.94 | | | | | | | |
| WM | 0.96 | 0.89 | | | | | | |
| Em | 0.94 | 0.87 | 0.92 | | | | | |
| Lan | 0.96 | 0.95 | 0.97 | 0.96 | | | | |
| Rel | 0.95 | 0.86 | 0.87 | 0.90 | 0.96 | | | |
| Soc | 0.96 | 0.95 | 0.97 | 0.95 | 0.98 | 0.97 | | |
| R1 | 0.94 | 0.94 | 0.96 | 0.94 | 0.96 | 0.97 | 0.96 | |
| R2 | 0.94 | 0.95 | 0.96 | 0.94 | 0.95 | 0.96 | 0.96 | 0.51 |

**Figure S4. State-decoding accuracies and important features.** The predictions accuracies for a 10-fold cross-validated predictive model using the parcellation schemes derived from the exemplars (nodes with fixed NNAs both across states and subjects, representing each network) calculated using the entire population (not just the training set as in Figure 2). After parcellating the entire population, a GBM is trained and tested in a 10-fold cross-validated setting where the GBM is trained on 9 folds of the data and used to predict the state for the left-out fold. a) The binary classification accuracies are displayed for all number of networks ($K = 2 - 50$). For every pair of tasks (task1, task2) we restricted our analysis to the data derived from the two tasks, yielding a population of 718 × 2 subjects with two outputs, task1 or task2. b) The binary classification accuracies averaged over number of networks ranging from 2 to 50. c) The 8-class classification accuracies are displayed for all number of networks ($K = 2 - 50$). The gray error bars represent the chance accuracy, computed as the mean and standard deviation of 1000 permutations. d) The distribution of the feature (node) importance across the three entropy classes. Interestingly, the flexible class comprises of the most important nodes for state decoding. Two-tailed t-test between every pair of steady, flexible, and transient class was performed and Bonferroni corrected for multiple comparisons, *** $p<1\times10^{-16}$. As expected, the result is similar to Figure 2, despite the difference in the calculated exemplars and the predictive pipeline.

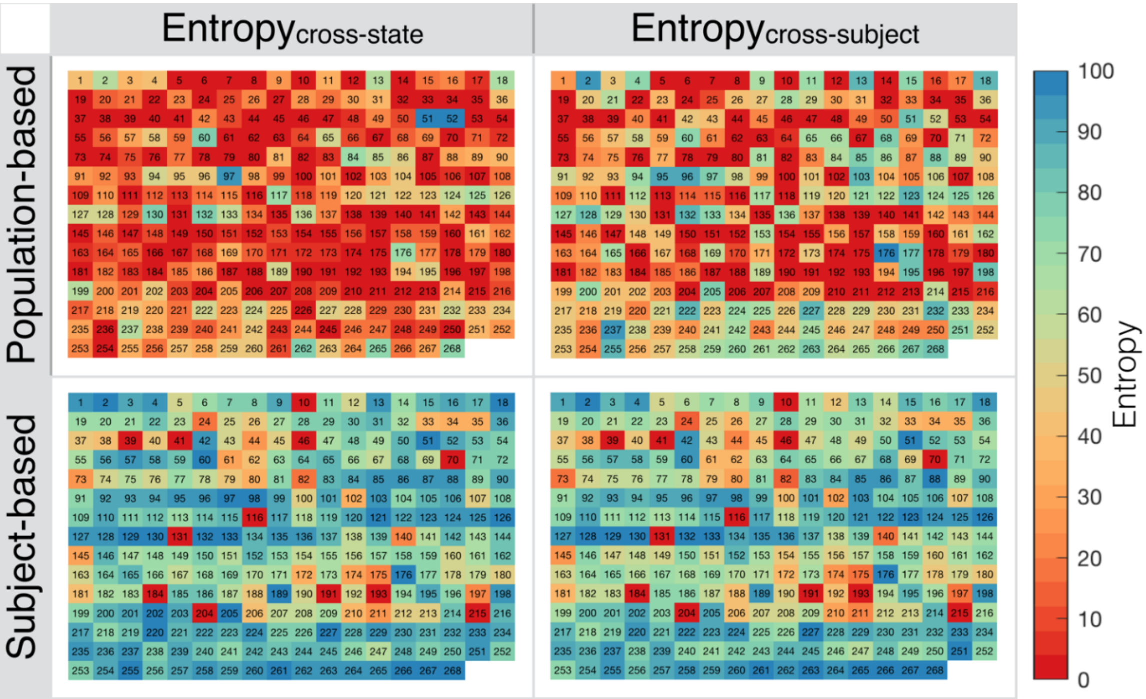

**Figure S5. Entropycross-state and Entropycross-subject for all 268 nodes using both population-based and subject-based methods for $K = 12$.** For every node, the two entropy measures (Entropy$_{\text{cross-state}}$ and Entropy$_{\text{cross-subject}}$) are calculated using histograms derived from the population-based analysis (with 100 iterations) and the subject-based analysis (across 718 subject). In the population-based analysis, at every step the entire population is divided into two equal size sets (each with 359 subjects) and the population-level parcellation is calculated for each set using the winner-takes-all strategy. This is employed 50 times generating a distribution of 100 population-level parcellations. As expected, the entropies calculated from the two population- and subject-based approaches are significantly correlated (cross-state entropy, $r_s = 0.79$, $p<1\times10^{-58}$; corss-subject entropy $r_s=0.77$, $p<1\times10^{-53}$; $r_s$, Spearman’s correlation coefficient). However, the entropies calculated from the subject-based approach are significantly higher on average due to their larger sensitivity to network variations (two-tailed Wilcoxon rank sum test, $p<2\times10^{-16}$).

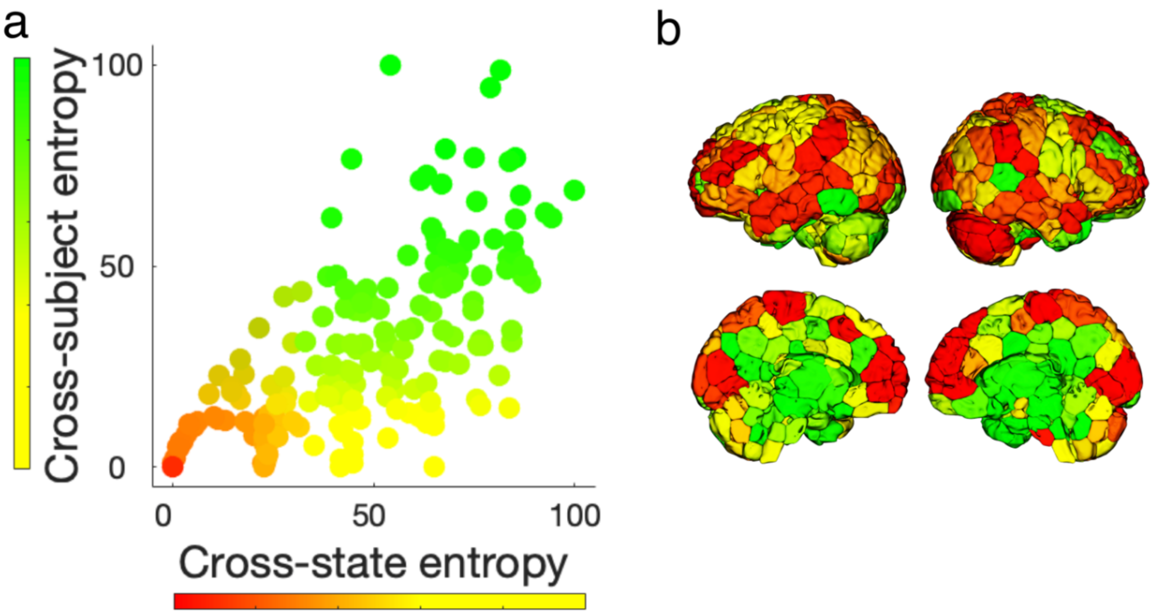

**Figure S6. Visualization of flexibility classes, defined based on NNA reconfigurations, as a continuum.** This is the continuous illustration of the flexibility classes defined in Figure 3 in the main manuscript. a) For every node, we calculated two measures of variability: the NNA entropy across states ($\mathrm{Entropy}_{\text{cross-state}}$) and the NNA entropy across subjects ($\mathrm{Entropy}_{\text{cross-subject}}$). We used two different colormaps for the two dimensions of entropy variation. Nodes are colored based on the linear combination of these two colormaps, such that warmer colors (red and orange) indicate lower cross-subject and cross-state entropy values; as we move towards yellow colors, we increase the cross-state entropy values, and as we progress towards green colors, we increase the cross-subject entropy values. b) Visualization of the entropy values [from part (a)] on the brain, using the same colormaps as part (a).

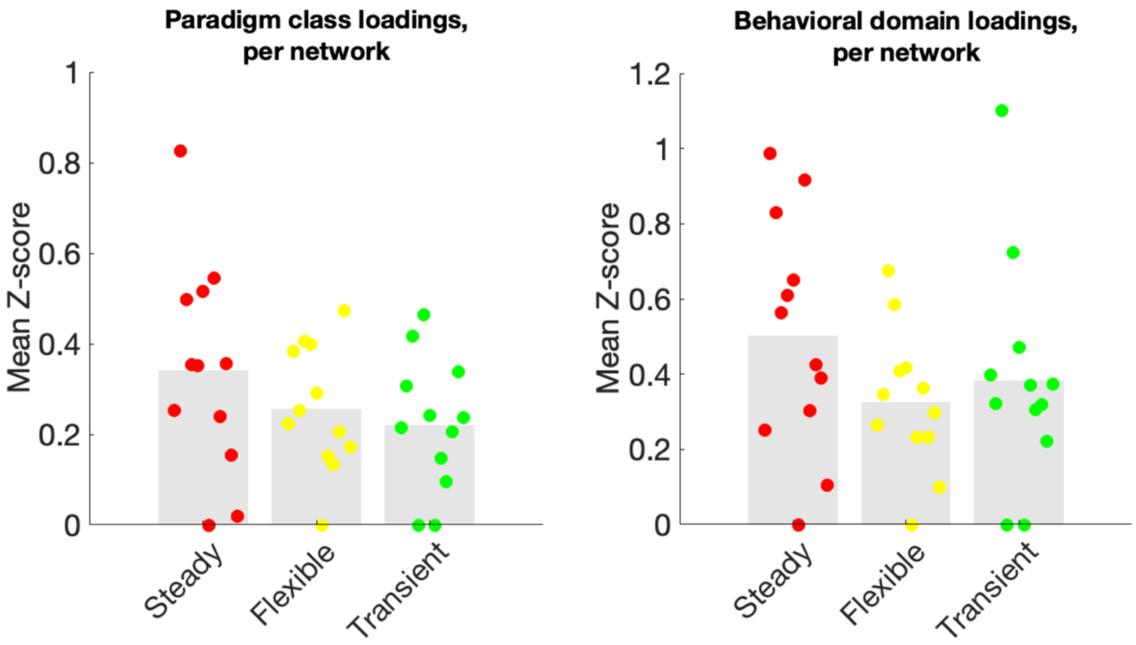

**Figure S7. Functional characterization of entropy classes at the network level.** The distribution of behavioral domains and paradigm classes are calculated across entropy classes. A network-level meta-analysis was performed using BrainMap Functional Database's experimental meta-data representing 55 behavioral domains and 108 paradigm classes. Profiles were computed within each of the 36 class-network pairings (3 entropy classes × 12 functional networks) as the forward inference likelihood z-scores, thresholded at z>1.96, and summed over all behavioral domains or paradigm classes. Colored circles represent the parts of each network that are associated with each entropy class (colored accordingly). The shaded bars display the mean within each entropy class.

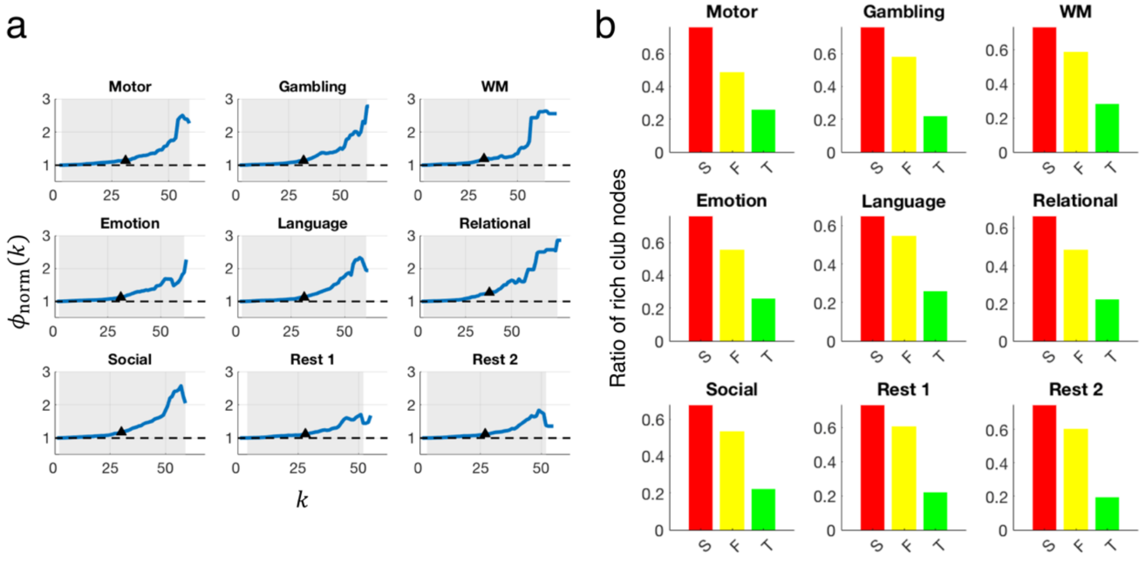

**Figure S8. Rich club organization of the brain across states.** a) Normalized rich club coefficient curves ($\phi_{\mathrm{norm}}$) are visualized for all functional states. The dashed line displays $\phi_{\mathrm{norm}} = 1$. The shaded gray area displays the rich club regime: $k$s for which the empirical rich club coefficients are significantly higher than an ensemble of 1000 null networks (permutation test; p<0.05). The triangle marker represents the average $k$ in the rich club regime, which was used as the threshold for defining rich club nodes ($k_{\mathrm{threshold}}$). b) The ratio of the rich club nodes ($k > k_{\mathrm{threshold}}$) within each entropy class is visualized for all functional states. S: Steady, F: Flexible, T: Transient. Consistent with our hypothesis, the steady nodes have the largest portion of rich club nodes, contributing most to the rich club organization of the brain in all states.

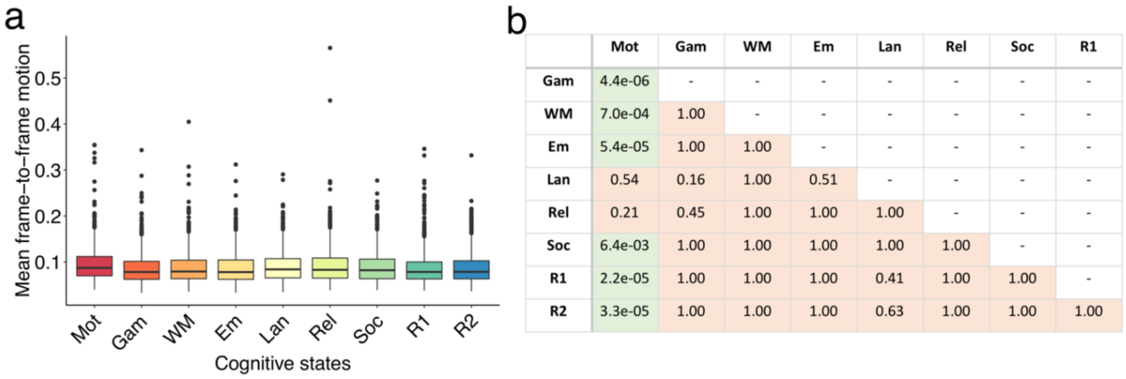

|  | Mot | Gam | WM | Em | Lan | Rel | Soc | R1 |
|---|---|---|---|---|---|---|---|---|
| Gam | 4.4e-06 | - | - | - | - | - | - | - |
| WM | 7.0e-04 | 1.00 | - | - | - | - | - | - |
| Em | 5.4e-05 | 1.00 | 1.00 | - | - | - | - | - |
| Lan | 0.54 | 0.16 | 1.00 | 0.51 | - | - | - | - |
| Rel | 0.21 | 0.45 | 1.00 | 1.00 | 1.00 | - | - | - |
| Soc | 6.4e-03 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | - | - |
| R1 | 2.2e-05 | 1.00 | 1.00 | 1.00 | 0.41 | 1.00 | 1.00 | - |
| R2 | 3.3e-05 | 1.00 | 1.00 | 1.00 | 0.63 | 1.00 | 1.00 | 1.00 |

**Figure S9. Head motion comparison across functional conditions.** Mean frame-to-frame displacement is computed for every functional condition, displayed as a box plot over sessions, where the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually. A two-tailed Wilcoxon rank sum test was performed on the mean frame-to-frame displacement between every pair of functional condition, and corrected for multiple comparisons using Bonferroni correction. b) P-values are for the pairwise Wilcoxon signed-rank test, performed on the mean frame-to-frame displacement, and corrected for multiple comparisons using Bonferroni correction.

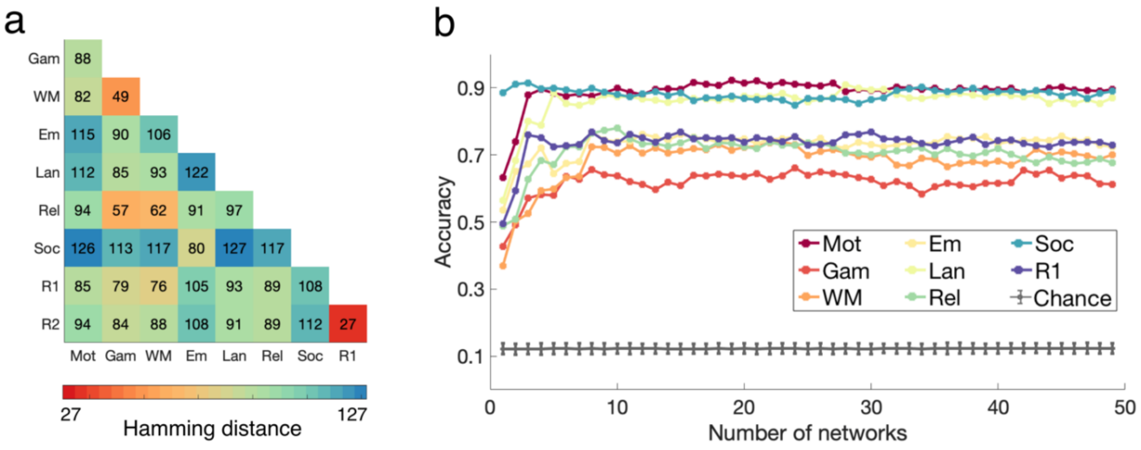

| | Mot | Gam | WM | Em | Lan | Rel | Soc | R1 |
|---|---|---|---|---|---|---|---|---|
| Gam | 88 | | | | | | | |
| WM | 82 | 49 | | | | | | |
| Em | 115 | 90 | 106 | | | | | |
| Lan | 112 | 85 | 93 | 122 | | | | |
| Rel | 94 | 57 | 62 | 91 | 97 | | | |
| Soc | 126 | 113 | 117 | 80 | 127 | 117 | | |
| R1 | 85 | 79 | 76 | 105 | 93 | 89 | 108 | |
| R2 | 94 | 84 | 88 | 108 | 91 | 89 | 112 | 27 |

**Figure S10. Cross-state network reconfigurations and state predictions after trimming all task data into 176 time points.** To rule out the possibility that the results are derived by differences in the number of time frames across conditions, we repeated our analysis this time trimming all the runs to have equal length (equal to the length of the shortest task, SOCIAL, with 176 time points). We constructed the individualized state-specific parcellation for each individual in each functional condition using the matched data. a) The network reconfiguration estimated by the number of nodes that change their network assignment across every pair of states. The observed pattern is very similar to that of the original analysis (Figure 1 in the manuscript). b) Accuracy of the state predictions based on the NNAs derived from the parcellations with matched data.

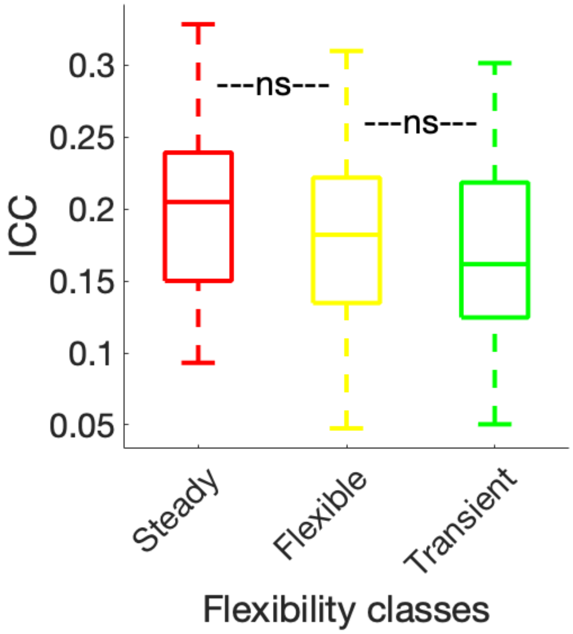

**Figure S11. Distribution of the functional connectivity reliability (ICC) across flexibility classes.** Reliability is measured by intraclass correlation coefficient (ICC). For every node in the atlas, the reliability is estimated as the mean test-retest reliability of all edges associated with that node. For each flexibility class, the reliability values are displayed as a box plot, where the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually. A two-tailed Wilcoxon rank sum test was performed between every pair of steady, flexible, and transient class and corrected for multiple comparisons using Bonferroni correction; ns=not significant.

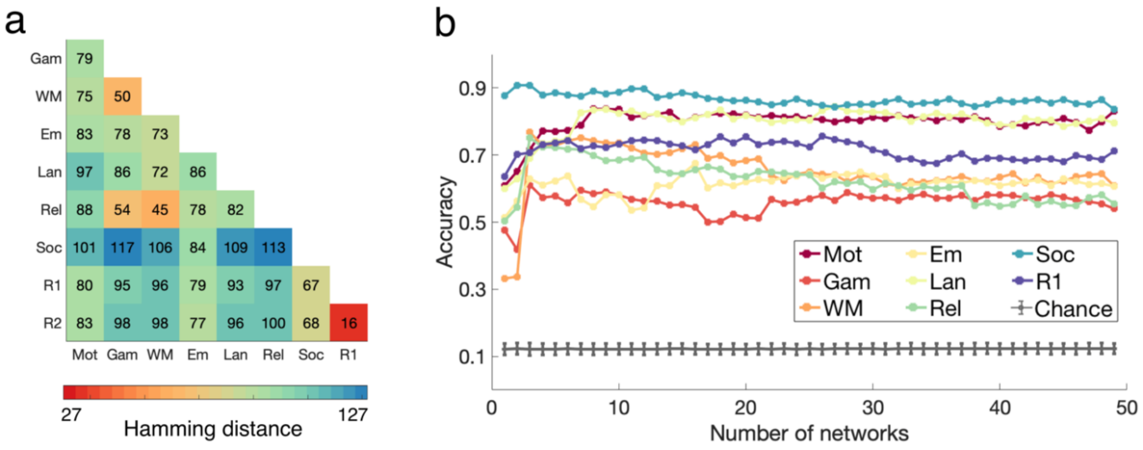

**Figure S12. Cross-state network reconfigurations and state predictions starting from individualized node atlases.** We constructed the individualized state-specific parcellation for each individual in each functional condition, this time starting from individualized nodes. For every individual, we constructed an individualized atlas using data from REST1 and REST2 runs. a) The network reconfiguration estimated by the number of nodes that change their network assignment across every pair of states. The observed pattern is very similar to that of the original analysis (Figure 1 in the manuscript). b) Accuracy of the state predictions based on the NNAs derived from the parcellations computed from the individualized atlas.

# Chapter 5: A Submodular Approach to Create Individualized Parcellations of the Human Brain

## Abstract

Recent studies on functional neuroimaging (e.g. fMRI) attempt to model the brain as a network. A conventional functional connectivity approach for defining nodes in the network is grouping similar voxels together, a method known as functional parcellation. The majority of previous work on human brain parcellation employs a group-level analysis by collapsing data from the entire population. However, these methods ignore the large amount of inter-individual variability and uniqueness in connectivity. This is particularly relevant for patient studies or even developmental studies where a single functional atlas may not be appropriate for all individuals or conditions. To account for the individual differences, we developed an approach to individualized parcellation. The algorithm starts with an initial group-level parcellation and forms the individualized ones using a local exemplar-based submodular clustering method. The utility of individualized parcellations is further demonstrated through improvement in the accuracy of a predictive model that predicts IQ using functional connectome.

## 1. Introduction

Functional connectivity analysis, as derived from BOLD fMRI, has shown promising role in establishing a better understanding of the functional organization of the human brain. Recent studies have shown that functional connectivity patterns of individuals are unique with large inter-subject variability (Finn et al., 2015). However, to date, the majority of previous work has ignored this inter-subject variability and created functional parcellations that are on average optimal for a group of subjects, rather than for any individual subject (Craddock et al., 2012; Shen et al., 2013). Such approaches have three key limitations. First, they require the usage of the entire dataset to construct an atlas, resulting in significant computational needs in the case of large datasets. Second, in the case when a new parcellation is created for the subjects at hand, there may not be known correspondence between the nodes in the new parcellations and previous ones, making it difficult to replicate and validate results across populations. Finally, they do not capture individual differences in the functional organization for each subject. This is particularly relevant for interventions that are focused on improvements in single patient, rather than group-level effects. Creating functional parcellation that accounts for an individual's functional organization while retaining explicit node correspondence to previously validated models remains a need for personalized applications.

Here, we propose a novel method to individualize the whole-brain functional parcellations consisting of hundreds of nodes (or parcels). This approach starts with a generic group parcellation and morphs it to account for the individual's functional organization. To individualize an existing parcellation, we use exemplar-based clustering and identify local exemplars for each node in the parcellation by optimizing a monotone submodular function using greedy algorithm. By starting with a group-level parcellation and the usage of exemplars, we retain known correspondence between nodes across individuals and across other studies. The individualization is applied to each subject independently, providing an online algorithm that uses only one individual data in a streaming fashion.

The remainder of the manuscript is organized as follows. First, we review previous group-level and individual-level parcellation approaches. Second, we formally introduce our proposed submodular approach. Third, we evaluate our approach against a non-individualized group-level parcellation using internal clustering evaluation measures and behavior prediction of IQ. Last, we offer concluding remarks.

## 2. Previous Work

Delineating functional boundaries either at the network-level (~5-20 clusters) or the node-level (~100-500 clusters) is a rich area of research with many proposed solutions. A non-exhaustive list of previous methods include *k*-means (Wang et al., 2015), hierarchical clustering (Meunier et al., 2010), spectral clustering (Craddock et al., 2012; Shen et al., 2013), and boundary mapping (Laumann et al., 2015). However, most of these approaches include one of the weaknesses stated above. Conceptually, the most similar previously published approach to our proposed work is the individualized network parcellation algorithm proposed in Wang et al. (Wang et al., 2015). This approach uses a *k*-means algorithm to individualize an existing network-level parcellation of 18 clusters.

Although *k*-means algorithms yield satisfactory results for problems with a small number of clusters, such as 18 networks, they generally suffer from sensitivity to the initialization and are more likely to converge to local optima. Thus, they may not be generalizable to problems with a large number of clusters such as individualization of a 200-500 node functional parcellation. In contrast, our submodular exemplar-based approach restricts the selection of the exemplars to the actual observed data points. By doing so, instead of minimizing a continuous loss function, we maximize a discrete submodular function for which the classical greedy algorithm provides the best approximation to the optimal solution. This approach is empirically more robust to noise and outliers than *k*-means methods.

## 3. Proposed Submodular Approach

Our proposed algorithm is comprised of three steps, illustrated in Figure 1. In the first step, an already defined group-level parcellation is warped to the individual's data, while accounting for the individual-specific gray-matter masks. Thus, every voxel in the individual's brain is assigned to one of the $N$ parcels defined by the group. In the second step, an *exemplar* is identified for each parcel by employing a *submodular* function optimization applied locally to each individual. In the third step, each voxel is assigned to the closest exemplar, where the closeness can be defined using any nonnegative similarity measure.

### 3.1 Exemplar-based clustering

Exemplar-based clustering is a method of data-summarization through identifying the most representative elements in the massive data, known as exemplars (Mirzasoleiman et al., 2016). One classic approach for identifying such exemplars is solving the $k$-medoids problem, by minimizing the following loss function:

$$L(S) = \frac{1}{|V|} \sum_{v \in V} \min_{e \in S} d(v, e), \tag{1}$$

where, $d: V \times V \rightarrow R$ defines the pair-wise dissimilarity between the elements of the ground set $V$. $L(S)$ measures how much information is lost by representing the entire dataset with its exemplars, i.e., the subset $S$. Except for special cases, the $k$-medoids problems are NP-hard. By introducing an appropriate auxiliary element $v_0$, the minimization of (1) can be turned into the maximization of a monotone submodular function $f$, as follows:

$$f(S) = L(v_0) - L(S \cup v_0), \tag{2}$$

for which general greedy algorithms provide an efficient $1 - 1/e \approx 0.63$ approximation of the optimal solution. For the choice of auxiliary element, any vector $v_0$ whose distance to every data point is greater than the pairwise distances between the data points, can serve the purpose. In the sequel, a formal definition of submodularity, and greedy algorithms are covered.

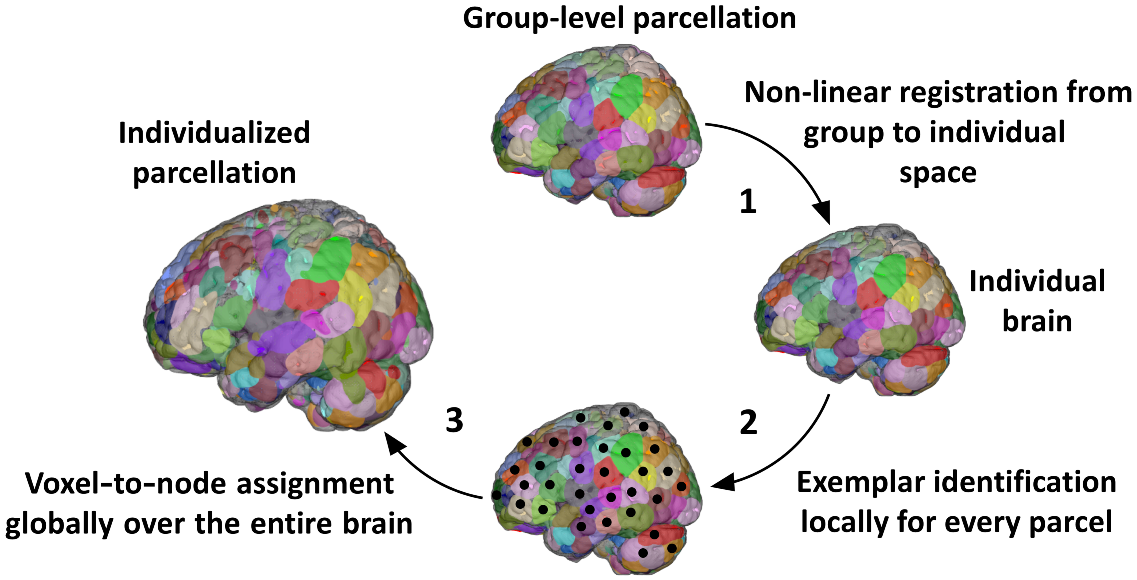

**Figure 1. Individualized parcellation pipeline.** Step 1: Registration of group-level parcellation to the individual space. Step 2: Identification of local exemplars for every node in individual brain. Step 3: Assignment of every voxel in individual brain to one of the exemplars spanning the entire brain.

### 3.2 Submodular functions

A function $f\colon 2^V \rightarrow R$ is submodular if for every $A \subseteq B \subseteq V$ and $e \in V\backslash B$ it holds that $f(A \cup e) - f(A) \geq f(B \cup e) - f(B)$. In other words, adding an element $e$ to a set $A$ increases the utility more than (or equal to) adding it to $A$'s superset, $B$, suggesting a natural diminishing returns. An important subclass of submodular functions (used in the proposed algorithm) are those which are non-negative and monotone. A function $f\colon 2^V \rightarrow R$ is monotone if for every $A \subseteq B \subseteq V$, $f(A) \leq f(B)$.

### 3.3 The greedy algorithm for optimization of the submodular function

Maximizing a non-negative monotone submodular function subject to a cardinality constraint, i.e.,

$$\begin{gathered}\max_{S \subseteq V} f(S) \\ \text{s.t.}\, |S| \leq K,\end{gathered} \tag{3}$$

is NP-hard (Feige, 1998). However, Nemhauser et al. proved that a simple greedy algorithm provides the best approximation (Nemhauser et al., 1978). The greedy algorithm starts with an empty set $S_0 = \emptyset$, and at each iteration $i$, it adds the element $e_i \in V$ such that the marginal gain is maximized, i.e.,

$$e_i^* = \arg\max_{e \in V} \Delta f(e|S_{i-1}) \coloneqq \arg\max_{e \in V} f(S_{i-1} \cup e) - f(S_{i-1}), \tag{4}$$

$$S_i = S_{i-1} \cup e_i^* \tag{5}$$

The algorithm continues until the cardinality constraint is reached, i.e., $|S| = K$. Herein, we attempt to represent each predefined parcel by one exemplar, thus $K = 1$. Note that by increasing the number of exemplars per parcel, one can derive a finer-grained parcellation scheme. Similarly, by skipping some of the parcels in the exemplar identification, one would derive a coarser parcellation, yielding to identification of functional networks (rather than nodes).

Local exemplar-search algorithm (i.e., searching for exemplars over each parcel independently, rather than the entire brain) can have global interpretations by defining a local dissimilarity measure such that for every node $v_i \in \text{cluster}_i$, the dissimilarity is defined as:

$$\begin{cases} d\left(v_i, v_j\right) = \infty, d\left(v_j, v_i\right) = 0, & \text{for } v_i \in \text{cluster}_i \text{ and } v_j \notin \text{cluster}_i \\ d\left(v_i, v_j'\right) = \|v_i - v_i'\|^2, & \text{for } v_i \in \text{cluster}_i \text{ and } v_i' \in \text{cluster}_i \end{cases} \tag{6}$$

This rich choice of dissimilarity measure is possible for exemplar-based methods (unlike many other classical algorithms). Dissimilarity measures only need to be nonnegative and do not require symmetry or triangle inequality properties. Finally, every voxel in the brain is assigned to the closest exemplar, yielding to an individualized parcellation with the same number of nodes as the initial group-level atlas (in this case, $N = 268$). Of note, the multi-step exemplar-based parcellation algorithm preserves a straightforward mapping between the parcellation of individuals to each other and to the group, as each node is represented by an exemplar derived from the same node in the group. Thus, we do not have to run another algorithm to retrieve the correspondences. This facilitates direct comparison of individuals with each other and with the group. Moreover, the greedy algorithm starts from an empty set and gradually adds elements, thus, there is no stochastic initialization process to affect the results. Herein, we used an accelerated version of the greedy algorithm, called lazy greedy (Minoux, 1978).

# 4. Methods

## 4.1 Dataset

For evaluation, we used a set of 200 randomly selected subjects (106 females, age 22-35) obtained from the Human Connectome Project (HCP) S900 release (Van Essen et al., 2013). Resting-state fMRI scans obtained on day 1 using the left-right phase-encoding direction were used. All data were preprocessed with the HCP minimal preprocessing pipeline. For calculating parcellation schemes, we further employed a Gaussian smoothing kernel with FWHM = 12 mm to each individual data. However, the rest of the analyses (i.e., forming the functional connectivity and the evaluation) were performed over the non-smoothed data. To relate these parcellations to individual behavior, we used fluid intelligence (IQ) as assessed using a form of Raven's progressive matrices with 24 items (Bilker et al., 2012) ($\mathrm{PMAT24\_A\_CR}$, $\mathrm{mean} = 16.28,\ \mathrm{s.d.} = 5.01,\ \mathrm{range} = 4 - 24$). Two subjects were excluded from evaluation involving behavior due to the missing IQ data.

## 4.2 Evaluation

We used the Shen functional parcellation (Finn et al., 2015; Shen et al., 2013)as the starting parcellation for our individualization approach and evaluated our approach against this parcellation. Note that here we are testing whether an individualized parcellation is better than the original group-level parcellation, supporting the utility of our approach. We do not aim to quantify which initial atlas results in the best individualized parcellation. We expect an individualized parcellation starting from any generic group-level parcellation to do better than the group-level parcellation, and that our results using the Shen parcellation will generalize to others.

We used two independent techniques for evaluation. First, the clustering quality was validated using two internal validation measures: Homogeneity and Davies-Bouldin index (DB). Second, we investigated whether the individualized parcellation increased the performance accuracy of predictive models derived from functional connectivity matrices.

### 4.3 Internal clustering validation

An ideal parcellation algorithm defines functionally homogenous nodes by grouping voxels with similar timecourses. Here, we assessed the homogeneity by calculating the average cross-correlations within each node, and averaging over all the nodes in the parcellated brain. We also calculated the DB index to assess the clustering ability in maximizing the intra-node compactness and the inter-node separation.

### 4.4 Predicting IQ

Functional connectivity matrices were computed using the proposed individualized parcellation scheme and the group-level parcellation, using standard connectivity methods.

After preprocessing, the average of the timecourses within each node was calculated, and the correlation between the averaged timecourses was computed, yielding to a $N \times N$ connectivity matrix for each subject, where $N = 268$ for both parcellations.

We employed the prediction framework used in Finn et al. (Finn et al., 2015) to build a predictive model of IQ based on $n - 1$ subjects (training set). The model was then tested over the left-out-subject (testing set) by predicting the IQ score. To assess the performance of the predictive model, we reported the Pearson correlation and the root mean squared error (RMSE) between the predicted scores and the observed IQ values.

### 4.5 Inter-subject variability

We further investigated to what extent the proposed individualized parcellation accounts for the inter-subject variability. We quantified this variation using normalized Hamming distance between each individual and the group parcellation.

## 5. Results

### 5.1 Internal clustering validation

Figure 2A displays the homogeneity scores (Left) and DB indices (Right) derived from both parcellation methods. The individualized parcellations have significantly higher level of homogeneity (two-tailed $t = 4.29, p < 2.3e - 5$), and significantly lower level of DB index (two-tailed $t = 12.5, p < 2.2e - 16$), both of which indicate a higher clustering quality. These results suggest that nodes derived from individualized parcellations are functionally more coherent, and, hence, can better capture individual differences in the functional organization of the brain.

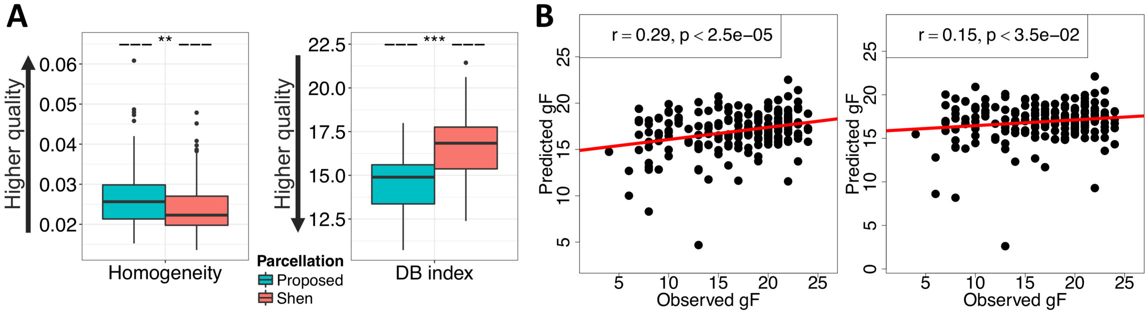

**Figure 2. Quantitative results for assessing the quality of parcellations.** A) Homogeneity (Left) and DB index (Right) comparison between the two parcellation methods. $** \, p < 2.3e-5$, $*** \, p < 2.2e-16$, two-tailed t-test. B) Comparing the predicted and observed IQ scores (gF) ($n = 198$ subjects) for the proposed individualized parcellation (Left), and Shen parcellation (Right). $p < 1e-15$, Hotelling's t-squared test.

### 5.2 Predicting IQ

Predicted IQ scores are plotted versus the observed IQ values in Figure 2B. Both group-level and individual-level parcellations generated significant predictions (group-level parcellation, $r = 0.15, p < 3.5e - 2,$ RMSE $= 5.07$; individualized parcellation, $r = 0.29, p < 2.5e - 5,$ RMSE $= 4.76$). However, the predictions using the individualized parcellations were significantly more accurate (Hotelling's $t = 6.0, p < 1e - 15$, comparison of two overlapping correlations based on dependent groups). The increased predictive power of the individualized parcellations indicates that our method preserves inter-subject variability and better accounts for individual's behavior.

### 5.3 Inter-subject variability

The average Hamming distance (normalized) between the individualized parcellations and the initial group parcellation is $\bar{d} = 0.45\,(\mathrm{s.d.} = 0.03)$. Since the Hamming distance captures the number of mismatches between the voxel-to-node assignment vectors, this result suggests that the proposed method on average changes ~45% of the initial group-level parcellation to account for the individual-specific functional node organization.

## 6. Conclusion

We developed a novel approach to the individualized whole-brain functional parcellation, starting from a generic population-level atlas. Our approach translates from the group to the individual by identifying local exemplars, or voxels, within each node, that best represent the timecourses of the remaining voxels in that node. The local exemplars are efficiently determined by optimizing a submodular utility function with greedy algorithm. After defining the local exemplars, the remaining voxels are assigned to the closest exemplar, yielding to a personalized parcellation that retains the explicit node correspondence to the population-level parcellation.

We showed that our proposed individualized parcellation yields a more homogenous node definition, and improves the predictive power of the resulting functional connectivities. Together, these findings highlight the need for more personalized parcellations, even at the node-level.

Future work will include further evaluation of our proposed approach including testing other existing group-level functional parcellations. We expect that our results using the Shen parcellation will generalize to parcellations created by other methods. Though it is possible that different parcellations with different number of nodes may see greater or lesser improvements with individualizations. Finally, our evaluations only included healthy controls. We will also begin to apply this approach to clinical populations, where we expect a different group-level functional organization than controls, and hence, such personalized approaches may show even greater improvements.

# Chapter 6: Individual parcellations are unique and reliable, forming an identifying fingerprint

## Abstract

Human brain is a network consisting of spatially distributed but functionally connected nodes. Advances in Neuroimaging techniques such as fMRI have enabled the reconstruction of this network at different scales. A conventional functional network analysis attempts to define the nodes in the brain by grouping voxels with similar time series, an approach known as parcellation. Traditional approaches to human brain parcellation have taken a group-level direction by collapsing data from all individuals in a group and generating a parcellation that is 'on average' optimal for everyone. While these techniques hold a great potential in developing a general blueprint for the brain functional organization, they do not capture the substantial inter-individual variability in brain function. Using data from multiple subjects and multiple sessions, here we establish that individual parcellations are unique and reliable, such that they can act as an identifying 'fingerprint'. We further show that that while all functional networks possess identification power of some sort, dorsal attention network, frontoparietal network, and language network emerged as the most distinct across individuals.

## 1. Introduction

Recent advances in Neuroimaging techniques have extended our understanding of the human brain functional organization. One widely used functional Neuroimaging technique, known as functional connectivity (FC) analysis, attempts to model the brain as a network, where different brain regions (or nodes) are connected by an edge if there is a synchrony in their temporal signals. Identifying nodes in this pipeline is a crucial step. Building the network at the voxel-level, with each voxel representing a node, would result in a noisy and high-dimensional model, making the subsequent connectivity analysis intractable (Thirion et al., 2014). An alternative approach is to define nodes by grouping voxels with similar time series, a technique known as human brain functional parcellation (or clustering). Since node definition is the first step in the functional connectivity analysis pipelines, it is crucial to define nodes that are functionally homogeneous. Failing to do so could propagate error to the subsequent steps and result in erroneous outcomes. Despite the recent advances in this field, the definition of the network nodes in the brain remains an open question (Eickhoff et al., 2015; Sporns, 2011).

Typically, human brain parcellation studies take a group-level inclination by collapsing data from all the subjects in a group to generate the best definition of the nodes that are optimal 'on average'. These group-level analyses have been accomplished either by averaging the subjects connectivity matrices (Power et al., 2011; Thomas Yeo et al., 2011) or by concatenating time courses from all subjects (Beckmann et al., 2005; Smith et al., 2009). There are four key limitation to these group-level parcellations: First, they require the usage of the entire dataset to construct an atlas. In addition to a significant computational and space burden, this limits their usage in clinical settings where often subjects are scanned gradually in a streaming fashion. Second, when averaging data, the importance of the correct registration is highlighted, as any mis-registration would result in mis-alignment across subjects, potentially yielding erroneous group averages. Third, when a new parcellation is created for the available subjects, there may not be precise correspondence between the new definition of the nodes and the old ones, making it challenging to validate and replicate results across population. Finally, they are limited in capturing individual-specific information and preserving individual differences in brain functional organization. This is in particular relevant for clinical interventions where the focus is on improvements in single patients, rather than group-level effects.

To overcome the above-mentioned limitations of group-level parcellations, individualized parcellation approaches have recently received increased attention. A number of parcellation algorithms have been introduced using advanced extension of the standard clustering techniques such as k-means (Kahnt et al., 2012), hierarchical clustering (Blumensath et al., 2013; Moreno-Dominguez et al., 2014), and spectral clustering (Craddock et al., 2012; Shen et al., 2013). Next section provides an overview of the individualized parcellation techniques in the field.

We have recently developed an individualized parcellation algorithm, which starts from a group-based template and warps it to adapt to each individual data (Salehi et al., 2017a). Here, we build up on our previous work to demonstrate that the individualized parcellations are reliable and reproducible across sessions of resting-state. Our results show that the parcellations are unique enough to each individual and reliable enough across sessions to act as an identifying ‘fingerprint’. Further analysis shows that while the entire brain contributes to the individual identifications, dorsal attention network, frontoparietal network, and language network emerge as the most distinctive ones, suggesting the relevance of the individual-specific topographic features to the higher-level cognitive functions.

## 2. Previous Work

Previous work has introduced a number of different approaches to human brain functional parcellation at the individual-level. Most of these approaches use resting-state functional connectivity data to delineate the brain into functionally coherent regions (Eickhoff et al., 2015; Thomas Yeo et al., 2011). Resting-state functional connectivity has been attractive for parcellation studies as (i) it is known to reflect the intrinsic functional organization of the brain which is not modulated by any behavioral task, and (ii) it does not require active engagement of subjects in any task, thus unburdening the experimental design and subject training demands. Nonetheless, there are also studies that have considered data from task-based fMRI, and other neurological properties such as architectural measures of cortical thickness (Glasser et al., 2016) or topographic organization (Felleman and Van, 1991).

For a functional parcellation to be clinically useful, there are certain constraints that needs to be fulfilled: (i) Functional parcellations should have high reliability and reproducibility within each individual, (ii) they should be sensitive to inter-individual differences in brain function, and (iii) they need to preserve correspondence across individuals to facilitate comparison across subjects and populations. While previous work has attempted to address these constraints to some extent, the field still lacks a fully automated, reproducible, and behaviorally cross-validated method of functional brain parcellation.

Cohen et al. proposed a boundary mapping parcellation technique to define functional areas based on the sharp transition boundaries (Cohen et al., 2008). Their work was complemented by Wig. et al, using snowball sampling method to highlight the individual-specific architectonics and connectivity patterns (Wig et al., 2014). Using the same parcellation method, Laumann et al. studied the reliability of a single subject parcellation based on 14 hours of resting-state fMRI data accumulated over more than a year (Laumann et al., 2015). One conceptual limitation of boundary mapping approaches is that exclusively rely on sharp transitions of a single property, and thus (i) are inherently limited by the sensitivity of the measuring technique to that property, and (ii) may miss the subdivisions developed based on other properties.

Another well-established parcellation strategy is based on agglomerative hierarchical clustering (Rokach and Maimon, 2005). This approach builds a hierarchy of clusters in a bottom-up fashion in which voxels with similar time series are merged to build larger functional units. While this approach has proved successful in identifying reliable parcels that are aligned with task-induced fMRI clusters (Blumensath et al., 2013), their ability to capture cytoarchitectonic boundaries is not clear. Van Oort et al. employed a contrasting strategy by dividing the large-scale functional units into smaller sub-regions using a top-down approach. Such top-down approach has shown perfect alignment with the underlying neurobiological features and proved useful for brain areas with convoluted folding patterns (van Oort et al., 2018); however, its performance is determined by the initial functional network definition.

Conceptually, the most similar previously published approach to our proposed parcellation is the individualized network parcellation algorithm proposed in Wang et al. (Wang et al., 2015). This approach uses a *k*-means algorithm to individualize an existing network-level parcellation of 18 clusters. Although *k*-means algorithm yields satisfactory results for the small-scale problems (such as 18 clusters), due to its sensitivity to initialization and the inherent risk of converging to local optima, it is not clear how they generalize to large-scale parcellation problems (with 200-500 nodes). In contrast, our submodular exemplar-based parcellation approach (Salehi et al., 2017a) restricts the selection of exemplars to the data points themselves, which in turn reduces the sensitivity of the results to noise and outliers. Our parcellation algorithm is a three-step pipeline which starts from a group-based parcellation and warps it to adapt to each individual data.

## 3. Theory

### 3.1 Localized exemplar-based parcellation

Our individualized parcellation algorithm is comprised of three steps. In the first step, an off-the-shelf group-level parcellation is applied to each individual's data, assigning each voxel to a node defined by the group parcellation. In the second step, one or more representative voxels are identified within each group-defined node. These voxels, named exemplars, are identified by maximizing a monotone non-negative submodular function (see Eq. 2). In the final step, each voxel in the brain is assigned to the functionally closest exemplar, such that the spatial contiguity of nodes is preserved. To account for the spatial contiguity, we had previously employed a Gaussian spatial smoothing with FWHM = 12 mm. However, spatial smoothing with large kernels reduces the spatial resolution, potentially blurring out meaningful information. We recently extended this method to have a built-in guarantee for the spatial contiguity of nodes. Details can be found elsewhere (Salehi et al., 2018a), but briefly, we start from each exemplar and gradually expand it such that at every assignment the spatial contiguity of the nodes are preserved.

### 3.2 exemplar-identification

As previously explained (Salehi et al., 2018a; Salehi et al., 2017a), the exemplar identification problem can be view as a data summarization problem where we the goal is to identify $k$ data points to represent the entire data such that minimum information is lost. More formally, the problem can be written as minimizing the following loss function subject to the cardinality constraint $|S| \leq k$:

$$L(S) = \frac{1}{|V|} \sum_{v \in V} \min_{e \in S} d(v, e), \tag{1}$$

where, $V$ is the ground set consisting of all data points, $d: V \times V \to R$ is a dissimilarity function (here, squared Euclidean distance), and $S$ is the objective exemplar set. Intuitively, $L(S)$ measures how much information we lose if we summarize $V$ to the exemplar set $S$ by representing each data point with its closest exemplar.

Minimizing (1) is generally NP-hard. The work by (Krause and Golovin, 2012) has shown that minimizing $L(S)$ can be transformed into maximizing a non-negative submodular function $f(S)$ by introducing an appropriate auxiliary exemplar $v_0$. More formally:

$$\max_{S \subseteq V} f(S), \quad s.t. \; |S| \leq k, \tag{2}$$

where:

$$f(S) = L(v_0) - L(S \cup v_0) \tag{3}$$

Greedy algorithm provides an efficient $1 - 1/e \approx 0.63$ approximation to the optimal solution (Nemhauser et al., 1978). Any vector $v_0$ whose distance to every data point is greater than the pairwise distances between data points can be used as an auxiliary exemplar.

### 3.3 Definition of Submodularity

A function $f: 2^V \rightarrow R$ is submodular if for every $A \subseteq B \subseteq V$ and $e \in V \backslash B$ it holds that $f(A \cup e) - f(A) \geq f(B \cup e) - f(B)$. That is, adding an element $e$ to a set $A$ increases the utility at more than (or at least equal to) adding it to $A$'s superset, $B$, suggesting a natural diminishing returns property.

## 4. Materials and Methods

### 4.1 Participants

Analyses were employed using resting-state fMRI data obtained from the Human Connectome Project (HCP) (Van Essen et al., 2013). HCP data includes two resting-state sessions acquired in two different days, referred to as Rest 1, and Rest 2. We limited the analyses to subjects for whom data were available for all functional conditions (with left-right (LR) and right-left (RL) phase encoding). Data from LR encoding was used in this work. To mitigate the substantial effects of head motion on functional parcellations, we further excluded subjects with excessive head motion (defined as mean frame-to-frame displacement > 0.1 mm and maximum frame-to-frame displacement > 0.15 mm), leaving 514 subjects (284 females; age = 22 - 36+) for analysis.

### 4.2 Imaging Parameters and Preprocessing

Starting with the minimally preprocessed HCP data (Glasser et al., 2013), further preprocessing steps were performed using BioImage Suite (Joshi et al., 2011) and included regressing 24 motion parameters, regressing the mean time courses of the white matter and cerebrospinal fluid as well as the global signal, removing the linear trend, and low-pass filtering (as previously described in (Finn et al., 2015a)).

### 4.3 Individualized functional parcellation

All data points were normalized into a unit norm sphere centered at the origin, and a point with norm greater than two was used as the auxiliary exemplar ($v_0$ in Eq. 3). This was to ensure the distance between auxiliary exemplar $v_0$ to every data point is greater than the pairwise distances between the data points. For every individual in each fMRI session (Rest 1 and Rest 2), the individualized parcellations were computed using the exemplar-based parcellation algorithm. In particular, each individual data was first parcellated using a group-level atlas (Shen et al., 2013) consisting of 268 nodes, defined based on a separate population of healthy individuals (Finn et al., 2015b). Next, an exemplar was identified within each group-defined node, maximizing Eq. 2 using greedy algorithm. Last, each voxel was assigned to the closest exemplar while ensuring spatial contiguity (see (Salehi et al., 2018a) for more details), resulting in two parcellation for each individual, one for Rest 1 and the other for Rest 2.

### 4.4 Individualized parcellation reproducibility

To establish the individualized parcellations are unique and reliable across sessions, we examined whether parcellations are more similar within an individual across sessions than across individuals. Parcellation similarities were quantified using $r_{\text{Hamming}}$ = 1-normalized Hamming distance, which measures the percentage of voxels that are consistent in their node assignments.

### 4.5 Identifying individuals based on their parcellations

Next, we investigated weather individualized parcellations are distinct enough to identify the individual from a large group of subjects (n=514). To this end, we considered individualized parcellations computed during Rest 1 as the ‘target’ set, and the ones computed during Rest 2 as the ‘database’ from which the individual is to be identified. In an iterative process, we selected one individual’s parcellation from the target set and compared it against every individual parcellation in the database to find the one with maximum similarity, where similarity was measured by $r_{\text{Hamming}}$. Once an identity was predicted, it was assigned a score of 1, if it matched the true identity, and 0, otherwise. The iterative process continued until all individuals were considered as the target. The identification success rate was measured as the percentage of individuals whose identity was correctly predicted. Next, we reversed the setting by considering Rest 1 as the ‘database’ and Rest 2 as the ‘target’ set, and repeated the iterative pipeline.

### 4.6 Identification power of each functional network

To test the hypothesis that certain functional networks contribute more to individuals’ discriminability than others, we repeated our identification pipeline, this time restricting our analysis to the nodes within each functional network. Accordingly, similarity was defined as the percentage of voxels within the nodes of that particular network which remained the same across parcellations. These networks (1-12) were identified in a previous work (Salehi et al., 2018b) based on the same data set used here. To ensure that parcellations were the only factor contributing to the identification power, we fixed the networks to be the same across individuals and sessions. To this end, we took the majority vote over all individuals in Rest 1 and Rest 2 (see (Salehi et al., 2018b) for more details).

## 5. Experiments and Results

### 5.1 Individualized parcellation reproducibility

We computed parcellation similarity between Rest 1 and Rest 2 both within and across individuals, resulting in a squared matrix of size $514 \times 514$, where every element $(i, j)$ in this matrix measures the parcellation similarity between individual $i$ in Rest 1 and individual $j$ in Rest 2 (Figure 1a). Diagonal elements in Figure 1a indicate parcellation similarities within an individual across sessions (Rest 1 and Rest 2), whereas off-diagonal elements show the similarity across individuals. The normalized distribution of the within-individual and cross-individual similarity values are depicted as histograms (Figure 1b), and compared using the non-parametric Kolmogorov–Smirnov test. We observe that parcellations are significantly more similar within an individual across resting-state sessions than across individuals (Figure 1b; K-S test; $p<10^{-16}$), suggesting that individualized parcellations are reproducible across sessions, and significantly rearrange from one individual to another.

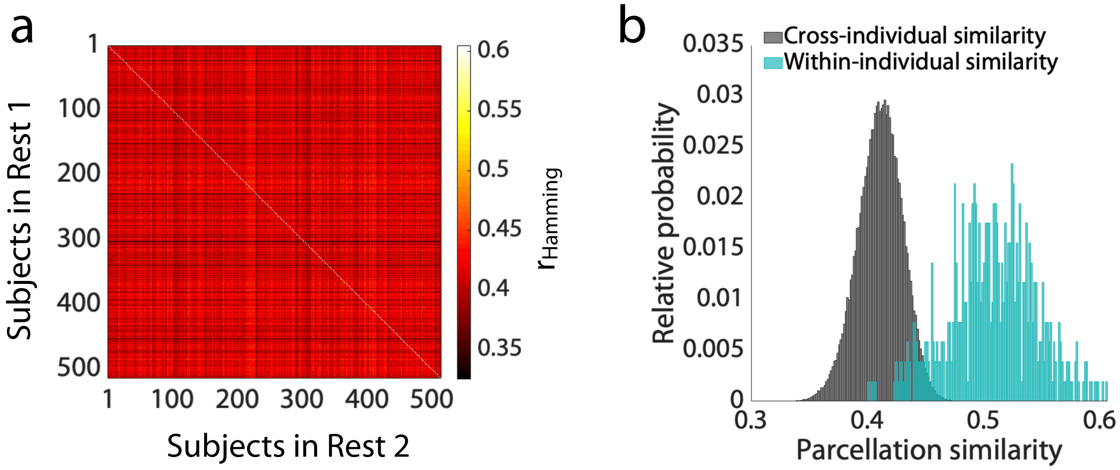

**Figure 1. Individualized parcellation reproducibility.** For every individual (n=514) data were collected during two resting-state sessions in two different days (Rest 1 and Rest 2). Parcellation was performed for each individual data and each session separately, and compared across the two sessions. a) Parcellation similarity was calculated between Rest 1 and Rest 2 sessions, both within an individual (diagonal elements) and across individuals (off-diagonal elements). Every element $(i,j)$ in this matrix represents the similarity between the parcellation of individual $i$ in Rest 1 and individual $j$ in Rest 2. Similarity was assessed by $r_{\text{Hamming}} = 1 -$ normalized Hamming distance. b) A histogram of all pair-wise parcellation similarities between Rest 1 and Rest is depicted both for within-individual (cyan; diagonal elements in [a]) and cross-individual (gray; off-diagonal elements in [a]) cases. The within- and cross- distributions are significantly different (K-S test; $p<10^{-16}$).

## 5.2 Identifying individuals based on their parcellations

We employed our individual identification paradigm (Figure 2a) based on the whole-brain parcellations (with 268 nodes). We could significantly identify individuals with 99.03% (509/514) and 98.64% (507/514) accuracy having Rest 1 and Rest 2 as our target session, respectively (Figure 2b). To assess the statistical significance of the identification accuracy, we performed nonparametric permutation testing. In each iteration, we randomly permuted the individuals' identity and repeated the identification pipeline and computed the success rate using the randomized identities. We repeated this procedure 1000 times. The highest success rate achieved was 1.17 % (6/514), yielding a p-value of 0.

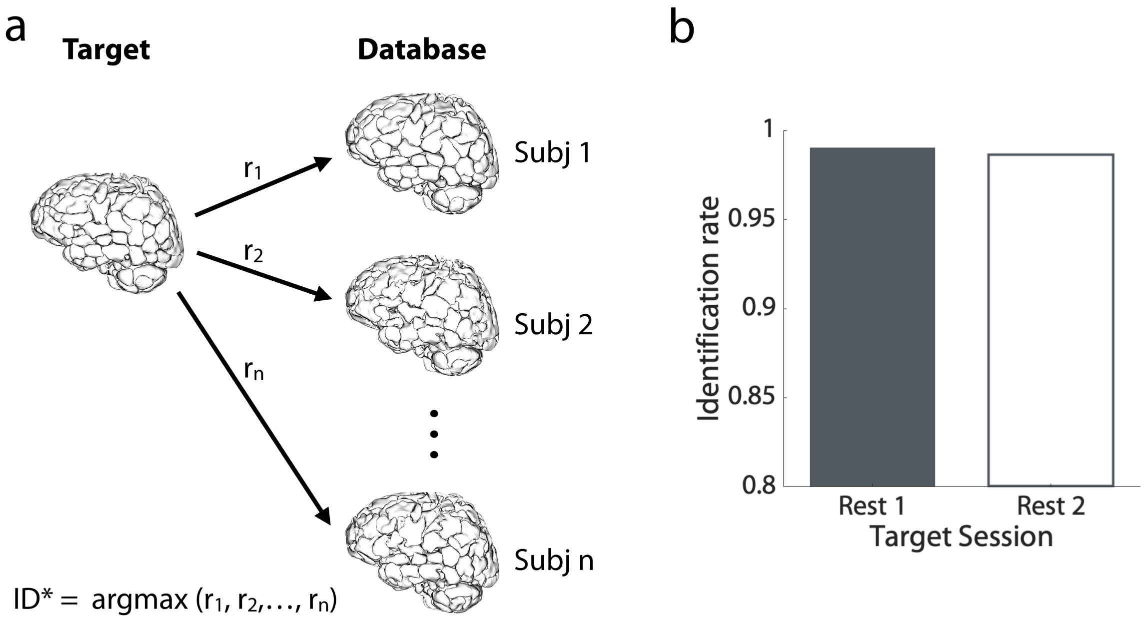

**Figure 2. Individual identification paradigm and accuracy.** a) For identification, we used the parcellation schemes from one session (Rest 1 or Rest 2) as the database and the parcellation schemes from the other session (Rest 2 or Rest 1) as the target set. Given a query parcellation scheme from the target set, we computed the similarity ($r_{\text{Hamming}}$) between this parcellation and all the parcellation schemes in the database, and selected the individual with highest parcellation similarity (argmax) as the predicted identity (ID*). b) For each of the two sessions considered as the target set, we computed the identification success rate as the percentage of individuals whose identity matched the true individual identity.

### 5.3 Identification power of each functional network

We next investigated the identification power for each of the twelve functional networks (Figure 3b). Dorsal attention network (DAN) emerged as the most discriminable network (identification rate = 94.55%), followed by frontoparietal network (89.11%), and language network (83.46%), all of which comprise higher-order association areas spanning the frontal, parietal and temporal lobes. Networks in primary sensory regions, on the other hand, displayed the least predictive power, with Visual I having the minimum identification power (41.05%) followed by sensorimotor network (41.25%), and subcortical-cerebellar network (51.36%). While these networks had the minimum success rate, they were still significantly higher than chance (maximum chance rate = 1.17%, p=0, permutation testing).

To assure that the identification power of networks is not confounded by the network size (i.e., the amount of data), we computed the Spearman correlation coefficient between the network sizes and success rates. We observe no significant correlation in either cases of Rest 1 ($r_s$ =0.48, p=0.12) or Rest 2 ($r_s$=0.45, p=0.15) considered as target session.

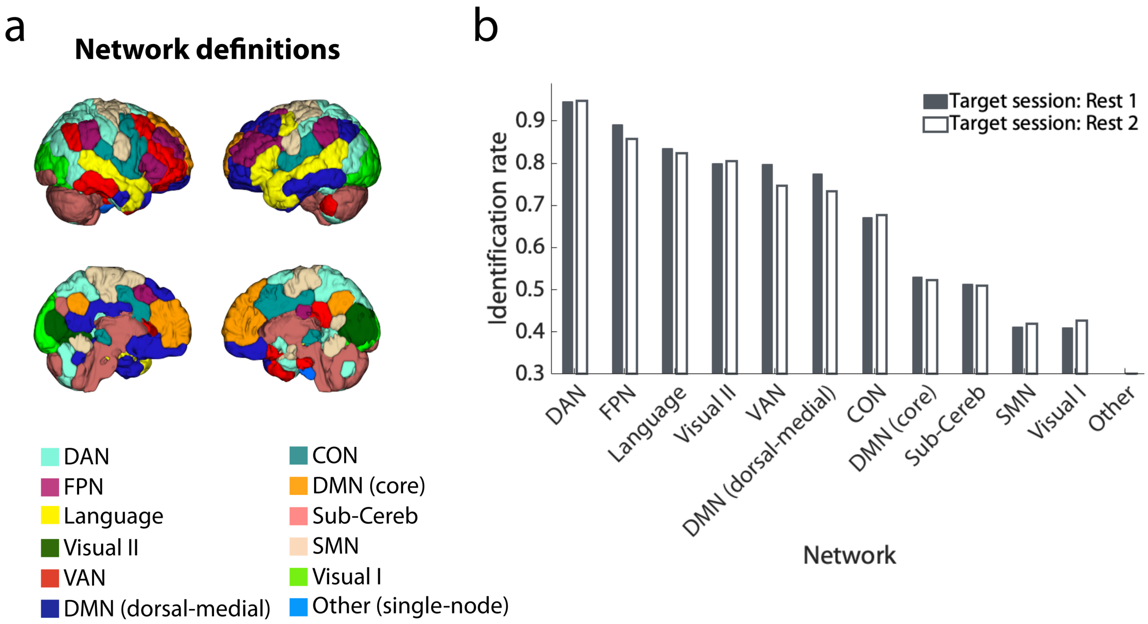

**Figure 3. Individual identification accuracy across networks.** a) Functional network definitions. Using a 268-node functional atlas defined on an independent data set of healthy subjects, we were further grouped nodes into networks (1–12) using our previously developed network delineation algorithm (Salehi et al., 2018b). We fixed network definitions across individuals to make sure node configuration is the only contributing factor to the identification rate. For this reason, we took the majority vote over all individuals and sessions. Networks are named according to their correspondence to other existing resting-state network definitions. DAN: dorsal attention network, FPN: frontoparietal network, VAN: ventral attention network, DMN: default mode network, CON: cingulo-opercular network, Sub-Cereb: subcortical/cerebellar network, and SMN: sensorimotor network. b) Identification success rate divided by the networks, sorted from high to low. For every network, identification paradigm (Figure 2) was repeated using only the nodes in that network to compute the parcellation similarity. Bar shading (black or gray) indicates which session was used as target.

## 6. Discussion

Here we show that the individualized functional parcellations, derived from our proposed submodular-based parcellation algorithm, are both unique and reliable across sessions. We demonstrate that the node spatial topography is unique enough to act as a fingerprint, identifying the individual from a pool of subjects with 99% accuracy. Our findings also suggest that nodes with high identifying power overlap with areas in higher-order association cortex, potentially suggesting the cognitive relevance of such node reconfigurations.

### 6.1 Implications for functional connectivity

While traditional connectivity studies have focused on group-level templates to define nodes (Sporns, 2011), our findings here suggest that there is considerable amount of inter-individual variability and uniqueness in parcellations. Considering such variability can further inform efforts to understand the link between brain function and behavior.

Previous work has shown that changes in spatial configuration of nodes can be interpreted as changes in connectivity (Bijsterbosch et al., 2018). As such, it is critical for connectivity studies to take these reconfigurations into account when interpreting the results. However, it is often challenging to isolate variabilities induced by meaningful individual-specific characteristics form noise. Given that functional atlases are the very beginning of the pipeline, it is critical for connectivity approaches to define functionally coherent nodes. Our finding that individuals can be identified based solely on the spatial topography of their nodes demonstrates that the observed node variabilities are not driven by noise, but are robust and reliable reflection of the underlying idiosyncrasies.

We have previously established the clinical utility of individualized parcellations (Salehi et al., 2017a), by showing that connectivity matrices built on individualized parcellations yield higher predictive power (for predicting IQ) than the ones built on group-level atlases. This work extends previous findings by showing that the individualized parcellations are reliable across sessions, and as such, can form a more powerful alternative to the commonly used group-based atlases. With these findings, we hope to inspire future connectivity approaches to consider personalized atlases in their analysis, which may in turn increase the efficacy of the subsequent analyses.

### 6.2 Localizing contributors: network-based analysis

Networks located in higher-order association areas displayed the most identifying power, whereas network in the primary sensory areas emerged as the least distinct (Figure 3). These findings are in line with previous reports from our group (Finn et al., 2015b; Salehi et al., 2017b) and others on the inter-individual variability in networks (Laumann et al., 2015; Wang et al., 2015) and connectivity profiles (Mejia et al., 2016; Miranda-Dominguez et al., 2014; Mueller et al., 2013). The emergence frontoparietal network as one of the most distinct systems, is in line with its role in human brain cognition. Nodes in frontoparietal network, are known as flexible hubs, shifting their connections based on the demands of the task (Cole et al., 2013). These regions have been associated with higher order cognitive processes (Dosenbach et al., 2008; Zanto and Gazzaley, 2013), and are also the most evolutionarily recent (Zilles et al., 1988).

Of note, while identification based on each independent network was successful, highest accuracy was achieved when the whole brain was considered (Figure 2). This finding mirrors recent characterization of cognitive processes arising from coordinated activity across the entire brain (Bassett et al., 2015; Finn et al., 2015b; Rosenberg et al., 2016). The nature of such variability across individuals and their relevance to differences in behavior remains an important question for future investigations.

## 7. Conclusion

Together, our findings suggest that individualized functional parcellations act as an identifying fingerprint, establishing that individual variability in the spatial topography of nodes are both substantial and reliable. Such variability is most prominent in areas associated with higher-order cognition, suggesting the potential cognitive relevance of such reconfigurations. Integrating these individualized atlases into the current connectivity pipelines can further inform future efforts to understand the link between brain function and behavior.

# Chapter 7: There is no single functional atlas even for a single individual: Parcellation of the human brain is state dependent

## Abstract

The goal of human brain mapping has long been to delineate the functional subunits in the brain and elucidate the functional role of each of these brain regions. Recent work has focused on whole-brain parcellation of functional Magnetic Resonance Imaging (fMRI) data to identify these subunits and create a functional atlas. Functional connectivity approaches to understand the brain at the network level require such an atlas to assess connections between parcels and extract network properties. While no single functional atlas has emerged as the dominant atlas to date, there remains an underlying assumption that such an atlas exists. Using fMRI data from a highly sampled subject as well as two independent replication data sets, we demonstrate that functional parcellations based on fMRI connectivity data reconfigure substantially and in a meaningful manner, according to brain state.

## 1. Introduction

Human neuroscience research has long endeavored to assign specific functions to clearly demarcated brain regions. Neuroimaging has offered deep insights into functional specialization in the human brain, and has permitted the localization of functional regions. Brain mapping approaches have been focused on assigning specific roles to such regions and through this approach much has been learned about the functional organization of the brain. Recently, there has been significant interest in whole-brain parcellation approaches for deriving brain atlases (Eickhoff et al., 2017). The parcellation problem itself is of great interest, and whole-brain parcellation is particularly relevant for defining the human connectome, which characterizes the interactions between brain regions. These developments have fueled a growing interest in functional connectivity-based analyses. There has been no consensus to date, however, on how to define the underlying atlas that best reflects the brain's functional organization.

The widely used Brodmann areas (Brodmann, 1909), defined by cytoarchitectural boundaries, were among the earliest attempts to subdivide the brain into functionally meaningful units, and there have been numerous, varied approaches to generate such atlases since (Craddock et al., 2012; Downing et al., 2001; McIntosh et al., 1994; Nieuwenhuys et al., 2015; Schubotz et al., 2010; Tzourio-Mazoyer et al., 2002). Neuroimaging-based parcellation approaches are attractive because they allow whole-brain parcellations in individuals (Blumensath et al., 2013; Chong et al., 2017; Salehi et al., 2017; Smith et al., 2013b) or groups of subjects (Fan et al., 2016; Glasser et al., 2016; Gordon et al., 2014; Power et al., 2011; Shen et al., 2013; Thomas Yeo et al., 2011b; Van Essen, 2013) (for a review see Eickhoff et al. (Eickhoff et al., 2018)). Most recent neuroimaging-based parcellation algorithms have been based entirely on functional connectivity data (Gordon et al., 2014; Power et al., 2011; Shen et al., 2013; Thomas Yeo et al., 2011b) or combinations of anatomical and functional data (Fan et al., 2016; Glasser et al., 2016). In addition, meta-analytic databases such as BrainMap (Fox and Lancaster, 2002) and NeuroSynth (Yarkoni et al., 2011) have attempted to collate information from thousands of studies to provide behavioral context on any region of the brain, and clustering methods have been developed to translate these findings into homogeneous regions (Eickhoff et al., 2011). Yet all of these commonly adopted parcellations, whether at the individual or group level, define a single functional atlas with the underlying assumption that parcels are homogenous in function and invariant in size, shape or position regardless of brain state.

In this work, we provide evidence that there is not a single functional parcellation atlas but rather that the flexible brain reconfigures these functional parcels depending upon what it is doing. While there are many timescales on which brain states can be measured and many ways in which they can be defined (see Discussion), here we use tasks to elicit discrete, distinct brain states, and demonstrate that parcel boundaries change across task-induced states, yet are reliably reproducible within a state. Further, we show that the particular configuration of the parcels provides meaningful information on brain state: that is, a measure as coarse as parcel size for a given atlas can significantly predict the task condition under which the data were acquired, as well as the within-condition task performance. Using a single, highly sampled subject, where we know there are no anatomic differences across conditions or sessions, and wherein one would expect the parcels to be consistent from session to session, we demonstrate that the parcels are indeed consistent for a given condition, but reproducibly reconfigure across conditions, even when starting with the same initial atlas each time. These results hold across two additional independent data sets, both within and across subjects, and at various parcellation spatial resolutions, suggesting that they are the result, not simply of individual idiosyncrasies in functional brain organization or of systematically varying noise in boundary estimates, but rather of robust, generalizable, state-dependent reorganization of functional areas. This suggests that a single functional parcellation of the human brain is neither attainable nor desirable, it emphasizes the importance of considering brain state when drawing functional boundaries on the brain, and offers functional parcellations as a tool to study state-dependent changes in functional organization of the human brain. These findings provide another dimension with which to understand the organization of the flexible brain under changing brain state or cognitive conditions.

## 2. Materials and Methods

Three independent data sets are used in this work. In the first, 30 sessions of fMRI data were obtained from a single subject; each session was approximately 60 minutes long and included 6 task conditions (*n*-back (Rosenberg et al., 2015), gradual-onset continuous performance task [gradCPT] (Esterman et al., 2012; Rosenberg et al., 2013; Rosenberg et al., 2016), stop-signal [SST] (Verbruggen et al., 2008), card guessing (Delgado et al., 2000), Reading the Mind in the Eyes (Baron-Cohen et al., 1997), and movie watching) and 2 rest conditions. Data were acquired from a single subject to eliminate the potential confound of inter-individual variations in anatomy, which would contribute to the variance in parcel boundaries (Bijsterbosch et al., 2018). The high number of fMRI sessions acquired in a single subject allows us to demonstrate consistent within-condition parcellations across sessions, while also demonstrating significantly different cross-condition parcellations, both within and across sessions. Since a single subject was used, the null hypothesis would be consistent parcellations across all sessions and conditions.

A second, independent data set was used from Midnight Scan Club (MSC) (Gordon et al., 2017b) to replicate these findings and demonstrate their generalizability. The MSC data are particularly well suited for this analysis, as they include task-based and resting-state fMRI data from 10 individuals, each of whom was scanned 10 times.

Finally, rather than measuring consistency within a subject across sessions, we used the Human Connectome Project (HCP) (Van Essen et al., 2013) data (n=514) to demonstrate that even when collapsing across subjects (rather than sessions), we observe that different conditions lead to reproducibly different functional parcellations. These three data sets are described in detail below.

## 2.1 Yale Data

### 2.1.1 Participant and processing

The primary subject R.T.C. is a healthy left-handed male, aged 56 years old at the onset of the study. The subject provided written informed consent in accordance with a protocol approved by the Human Research Protection Program of Yale University.

The subject was scanned at Yale University 33 times (that is, 33 sessions) over ten months. Scans were typically performed on Wednesdays at 8:30 am and Fridays at 2:00 pm. Functional MRI data were acquired on 2 identically configured Siemens 3T Prisma scanners equipped with a 64-channel head coil at the Yale Magnetic Resonance Research Center. Three sessions were excluded from analysis: the first two sessions were excluded because of the considerable adjustment in task design after session 2, and the ninth session was excluded due to interruptions in task presentation due to interruptions in task presentation.

The first session was used to acquire structural MRI data. High-resolution T1-weighted 3D anatomical scans were performed using a magnetization prepared rapid gradient echo (MPRAGE) sequence with the following parameters: 208 contiguous slices acquired in the sagittal plane, repetition time (TR) = 2400 ms, echo time (TE) = 1.22 ms, flip angle = 8°, slice thickness = 1 mm, in-plane resolution = 1 mm × 1 mm, matrix size = 256 × 256. A T1-weighted 2D anatomical scan was acquired using a fast low angle shot (FLASH) sequence with the following parameters: 75 contiguous slices acquired in the axial-oblique plane parallel to AC-PC line, TR = 440 ms, TE = 2.61 ms, flip angle = 70°, slice thickness = 2 mm, in-plane resolution = 0.9 mm × 0.9 mm, matrix size = 256 × 256.

Functional scans were performed using a multiband gradient echo-planar imaging (EPI) pulse sequence with the following parameters: 75 contiguous slices acquired in the axial-oblique plane parallel to AC-PC line, TR = 1000 ms, TE = 30 ms, flip angle = 55°, slice thickness = 2 mm, multiband acceleration factor = 5, in-plane resolution = 2 mm × 2 mm, matrix size = 110 × 110.

Data were analyzed using BioImage Suite (Joshi et al., 2011) and custom scripts in MATLAB (MathWorks). Motion correction was performed using SPM (https://www.fil.ion.ucl.ac.uk/spm/). White matter and CSF masks were defined in MNI space and warped to the single-subject space using a series of linear and non-linear transformations (see Scheinost et al. (Scheinost et al., 2015)). The following noise covariates were regressed from the data: linear, quadratic, and cubic drift, a 24-parameter model of motion (Satterthwaite et al., 2013), mean cerebrospinal fluid signal, mean white matter signal, and mean global signal. Finally, data were temporally smoothed with a Gaussian filter ($\sigma = 1.55$).

### 2.1.2 Dimensional task battery design

Functional scans were 6 minutes 49 seconds each, including initial shim and 8s discarded acquisitions before the start of each task. Tasks varied slightly in length (see below), but were all approximately 6 minutes in duration. A fixation cross was displayed after the end of each task and lasted until the beginning of the next task. Each task, with the exception of movie watching, was preceded by instructions and practice, after which the subject had the opportunity to ask questions before the scan began. All responses were recorded using a 2×2 button box.

Each session consisted of two resting-state runs and six task runs. The first and last functional runs (runs 1 and 8) were resting-state runs, during which the participant was instructed to stay still with his eyes open. Runs 2 – 7 were task runs, with the order counterbalanced across sessions.

#### 2.1.2.1 N-Back Task

The *n*-back task was adapted from that used in Rosenberg et al. (Rosenberg et al., 2015). In this task, the participant was presented with a sequence of images and was instructed to respond via button press if the image was different than the image presented two before, and to withhold response if it was the same. Images were presented for 1 second, followed by a 1-second inter-trial interval (ITI; fixation cross). The target (i.e., matching image) probability was 10%. There were two blocks, each with 90 trials. One block used images of emotional faces and the other block used images of scenes (Cohen et al., 2016; Conley et al., 2018). Block and stimulus order were randomized for each session. Task performance was assessed by *sensitivity* (*d′*), defined as hit rate relative to false alarm rate (Rosenberg et al., 2016).

#### 2.1.2.2 Gradual-onset Continuous Performance Task (gradCPT)

The gradCPT task was adapted from that described in Esterman et al. (Esterman et al., 2012) and Rosenberg et al. (Rosenberg et al., 2013; Rosenberg et al., 2016). In this task, the participant viewed a sequence of 450 scenes (city or mountain) that gradually transitioned via linear pixel-by-pixel interpolation from one to the next over 800ms. The participant was instructed to respond via button press to cities and to withhold response to mountains. Stimulus order was randomized, and 10% of images were mountains. Task performance was assessed by *sensitivity* (*d′*).

#### 2.1.2.3 Stop-Signal Task (SST)

The stop-signal task was adapted from that implemented in Verbruggen et al. (Verbruggen et al., 2008). In this task, the participant was required to determine via button press whether a presented arrow was pointing left (right index finger) or right (right middle finger). On 25% of trials, the arrow turned blue after some delay, indicating that the participant should withhold response. This stop-signal delay (SSD) was initially set to 250ms, and was continuously adjusted via the staircase tracking procedure (50ms increase after correct inhibition trials; 50ms decrease after each failure to inhibit). The arrow was presented for 1.5 seconds, followed by 0.5 seconds of fixation; there were 176 trials in total, with stimulus order randomized within block. Task performance was assessed by *missing probability*, defined as the percentage of missed responses on no-signal trials (Verbruggen et al., 2008).

#### 2.1.2.4 Card Guessing Task

The card guessing task was adapted from that originally developed by Delgado et al. (Delgado et al., 2000) and subsequently extended (Barch et al., 2013; Speer et al., 2014). In this task, the participant was presented with a card and asked to guess if the number on the back was lower than 5, or greater than 5 but less than 10. The question mark card was displayed for 1.5 seconds, or until the participant responded (right index finger for "lower," right middle finger for "higher"). The card then "flipped over" to reveal the number. The number was displayed for 0.5 seconds, followed by an arrow for 0.5 seconds to indicate accuracy (green and up for correct, red and down for incorrect), which was in turn followed by a 1-second inter-trial interval (fixation cross). There were 10 blocks, each with 10 trials, and guess accuracy was deterministic, such that in half of the blocks ("high win"), the participant was correct 70% of the time, while in the other half of the blocks ("high loss") he was correct 30% of the time; block (high win/loss) and trial (correct/incorrect) orders were randomized. Task performance was assessed by *RT variability*, defined as standard deviation of reaction time (May et al., 2004).

#### 2.1.2.5 Reading the Mind in the Eyes Task (“Eyes Task”)

The Eyes Task was adapted from that originally described in Baron-Cohen et al. (Baron-Cohen et al., 1997). In this task, the participant viewed a series of photographs of an individual’s eyes with four “mental state terms” (Baron-Cohen et al., 1997), one in each corner of the image, and was instructed to select via button press (with each button corresponding to one corner, and thus one term) the term that best described what the individual was thinking or feeling. There were 36 images in total. Each was presented once, in random order, for 9.25 seconds or until the participant responded; the remainder of each 10-second trial consisted of a fixation cross. Task performance was assessed by *RT variability*, defined as standard deviation of reaction time for correct trials.

#### 2.1.2.6 Movies Task

In this task, three movie clips were presented in continuous series; each was approximately 2 minutes long. The first clip was a trailer for “Inside Out,” the second clip was the wedding scene from “Princess Bride,” and the third clip was a trailer for “Up;” order was fixed across sessions. The participant was instructed to relax and enjoy the movies; no responses were required. No task performance was recorded.

## 2.2 Midnight Scan Club (MSC) Data

### 2.2.1 Participants and processing

The MSC data set (Gordon et al., 2017a) includes data from 10 healthy individuals (5 females; age = 24 – 34); each underwent 1.5 hours of functional MRI scanning on 10 consecutive days, beginning at midnight. For details of the data acquisition parameters and sample demographics see Gordon et al (Gordon et al., 2017a). Two individuals were excluded from the analysis: MSC08 was excluded because of excessive head motion and self-reported sleep (Gordon et al., 2017a); MSC10 was excluded for insufficient data (missing one session of incidental memory task).

All data were preprocessed using BioImage Suite (Joshi et al., 2011). Data were transformed to MNI space to facilitate analysis across multiple subjects. Preprocessing steps included regressing 24 motion parameters, regressing the mean time courses of the white matter and cerebrospinal fluid as well as the global signal, removing the linear trend, and low pass filtering.

### 2.2.2 Task battery design

Each scanning session started with a 30-min resting-state fMRI scan, followed by three separate task-based fMRI scans: motor task (2 runs per session, 7.8 min combined), incidental memory task (3 runs per session, 13.1 min combined), and semantic-coherence task (2 runs per session, 14.2 min combined).

#### 2.2.2.1 Motor Task

The motor task was adapted from that used in the Human Connectome Project (HCP) (Barch et al., 2013). In this task, participants were cued to perform one of the following movements: closing/relaxing their hands, flexing/relaxing their toes, or wiggling their tongue. Each block started with a 2.2 s cue indicating which movement to perform, followed by a central caret (flickering every 1.1s) to signal the movement. Each run consisted of 2 blocks of each movement type and 3 blocks of resting-fixation (15.4 s total).

#### 2.2.2.2 Incidental Memory Task

The incidental memory task consisted of three different types of stimuli (scenes, faces, and words), each presented in a separate run. For scene runs, participants were asked to decide if the presented scene was indoors or outdoors. For face runs, participant made male/female judgments. For word runs, participants made abstract/concrete judgments. Each run consisted of 24 stimuli, each repeating 3 times. Stimuli were presented for 1.7 s with a jittered 0.5-4.9 s inter-stimulus interval. All stimuli were taken from publicly available sources (see Gordon et al. (Gordon et al., 2017a) for details).

#### 2.2.2.3 Semantic-Coherence Task

The semantic-coherence task had a mixed block/event-related design, consisting of two different conditions (“semantic” and “coherence”). In the “coherence” task, participants viewed a concentric dot pattern (Glass, 1969) with 0% or 50% coherence, and made binary decisions whether the pattern was concentric or random. In the semantic task, participants viewed a word and indicated whether the word is a noun or verb. Each run consisted of two blocks of each task, separated by 44 s of rest. Each block started with a 2.2 s cue indicating which task was to be performed. Blocks consisted of 30 trials. Stimuli were presented for 0.5 s with a variable 1.7-8.3 s ISI. Each block finished with a 2.2 s cue indicating the end of the task block.

## 2.3 Human Connectome Project (HCP) Data

### 2.3.1 Participants and processing

The HCP data set includes data from 897 healthy individuals (S900) scanned during nine functional conditions (seven tasks and two rest). For details of the data acquisition parameters see Uğurbil et al. (Uğurbil et al., 2013) and Smith et al. (Smith et al., 2013a). Analyses were restricted to subjects for whom data were available for all nine functional conditions (with left-right (LR) and right-left (RL) phase encoding). To mitigate the substantial effects of head motion on functional parcellations, we further excluded subjects with excessive head motion (defined as mean frame-to-frame displacement > 0.1 mm and maximum frame-to-frame displacement > 0.15 mm), leaving 514 subjects (284 females; age = 22 – 36+) for analysis.

The HCP minimal preprocessing pipeline was employed (Glasser et al., 2013), which includes artifact removal, motion correction and registration to MNI space. Further preprocessing steps were performed using BioImage Suite (Joshi et al., 2011) and included standard preprocessing procedures (Finn et al., 2015) including regressing 24 motion parameters, regressing the mean time courses of the white matter and cerebrospinal fluid as well as the global signal, removing the linear trend, and low pass filtering.

### 2.3.2 Task battery design

Functional MRI scans were acquired during two different days: Day 1 included two runs (LR and RL) of the working memory (WM) task (5:01 min per run), incentive processing (gambling) task (3:12 min), motor task (3:34 min), and rest (14:33 min); day 2 included two runs of the language processing task (3:57 min), social cognition (theory of mind) task (3:27 min), relational processing task (2:56 min), emotion processing task (2:16 min), and rest (14:33 min).

The details of task design have been previously described (Barch et al., 2013; Van Essen et al., 2013). We provide a brief description of each task and an overview of the relevant aspects below.

#### 2.3.2.1 Working Memory Task

In this task, participants performed a visual *n*-back task, with blocked 0-back and 2-back conditions using four stimulus categories (faces, places, tools, body parts). Each run consisted of 8 task blocks (10 trials each), with each stimulus category used twice, and 4 fixation blocks. Each block started with a 2.5 s cue indicating the task type (0-back versus 2-back) and the target (for 0-back).

#### 2.3.2.2 Gambling Task

In this task, participants were presented with a mystery card and asked to guess if the number on the back was lower than 5, or greater than 5 but less than 10. On reward trials, participants were shown the number on the card, a green up arrow, and “$1”; on loss trials, participants were shown the number on the card, a red down arrow, and “-$0.50”; on neutral trials, participants were shown the number 5 and a gray, double-headed arrow. Each run consisted of 4 task blocks (8 trials each) and 4 fixation blocks. In half of the blocks (“mostly reward”), subjects were correct in 6 out of 8 trials (the remaining 2 trials were either neutral or loss), while in the other half of the blocks ("mostly loss") they were incorrect in 6 out of 8 trials (the remaining 2 trials were either neutral or reward).

#### 2.3.2.3 Motor Task

Participants were presented with visual cues that asked them to tap their left or right fingers, squeeze their left or right toes, or move their tongue. Each block started with a 3 s cue indicating which movement to perform. Each run consisted of 2 blocks of tongue movements, 4 blocks of hand movements (2 left and 2 right), 4 blocks of foot movements (2 left and 2 right), and 3 blocks of resting-fixation.

#### 2.3.2.3 Language Task

In this task, participants were aurally presented with 4 blocks of a story task and 4 blocks of a math task. In the story task, they heard brief fables (5-9 sentences) and completed two-alternative forced-choice questions about the topic of the story. In the math task, they completed addition and subtraction problems in a two-alternative forced-choice setting.

#### 2.3.2.4 Social Task

In this task, participants were presented with 20-s video clips of objects (squares, circles, triangles) either interacting (theory-of-mind) or moving randomly. Participants were asked to choose between three potential responses ("mental interaction", "no mental interaction", and "not sure"). Each run consisted of 5 video blocks (2 mental and 3 random in one run, 3 mental and 2 random in the other run) and 5 fixation blocks.

### 2.3.2.5 Relational Task

The relational task consisted of two different conditions (“relational” and “matching”). In the relational condition, participants were presented with two pairs of objects with one pair at the top and the other pair at the bottom of the screen. They were asked to decide whether the bottom pair of objects differed along the same dimension (i.e., shape or texture) as the top pair. In the control matching condition, they were presented with two objects at the top, one object at the bottom, and a word (“shape” or “texture”) in the middle of the screen. They were asked to determine whether the bottom object matched either of the top two objects on the dimension specified by the word. Each run consisted of 3 relational blocks (4 trials each), 3 matching blocks (5 trials each) and 3 fixation blocks.

### 2.3.2.6 Emotion Task

In this task participants were presented with blocks of “face” and “shape” tasks, and were asked to determine which of the two faces (or shapes) presented at the bottom of the screen matched the face (or shape) at the top of the screen. Faces had either angry or fearful expressions. Each block started with a 3 s cue indicating which task to perform. Each run included 3 face blocks (6 trials each) and 3 shape blocks (6 trials each), with 8 s fixation at the end of each run.

## 2.4 Individualized and state-specific parcellation algorithm

Here, we extended our previously developed individualized parcellation algorithm (Salehi et al., 2017) to account for the spatial contiguity of the parcels at the individual level. The presented algorithm is a priority-based submodular method that defines functional parcels in a streaming fashion, for every individual in every functional state. A key factor in this algorithm is that each parcellation begins with an atlas obtained from a separate group of subjects and finds an exemplar time-course for each parcel and then grows the parcels (see Figure 1a for a visual illustration). This exemplar-based approach has many advantages but the most important advantage for this work is that correspondence of parcels from different parcellations is maintained, making for straightforward comparisons of the resulting atlases.

Our algorithm runs in three steps:

**i. Registration of the initial group-level parcellation.** In the first step, an off-the-shelf group-level parcellation is applied to each individual's data, assigning each voxel to a parcel defined by the group parcellation. At this step, all individuals in every state have the same parcel definitions.

**ii. Exemplar identification.** In the second step, for every group-defined parcel in an individual brain, an exemplar is identified by maximizing a monotone nonnegative submodular function (see Eq. 2).

**iii. Spatially-constrained voxel-to-parcel assignment.** Finally, the third step assigns every voxel in each individual brain to the functionally closest exemplar while ensuring the spatial contiguity of the resulting parcel.

A visual illustration of the parcellation algorithm is provided in Figure 1a.

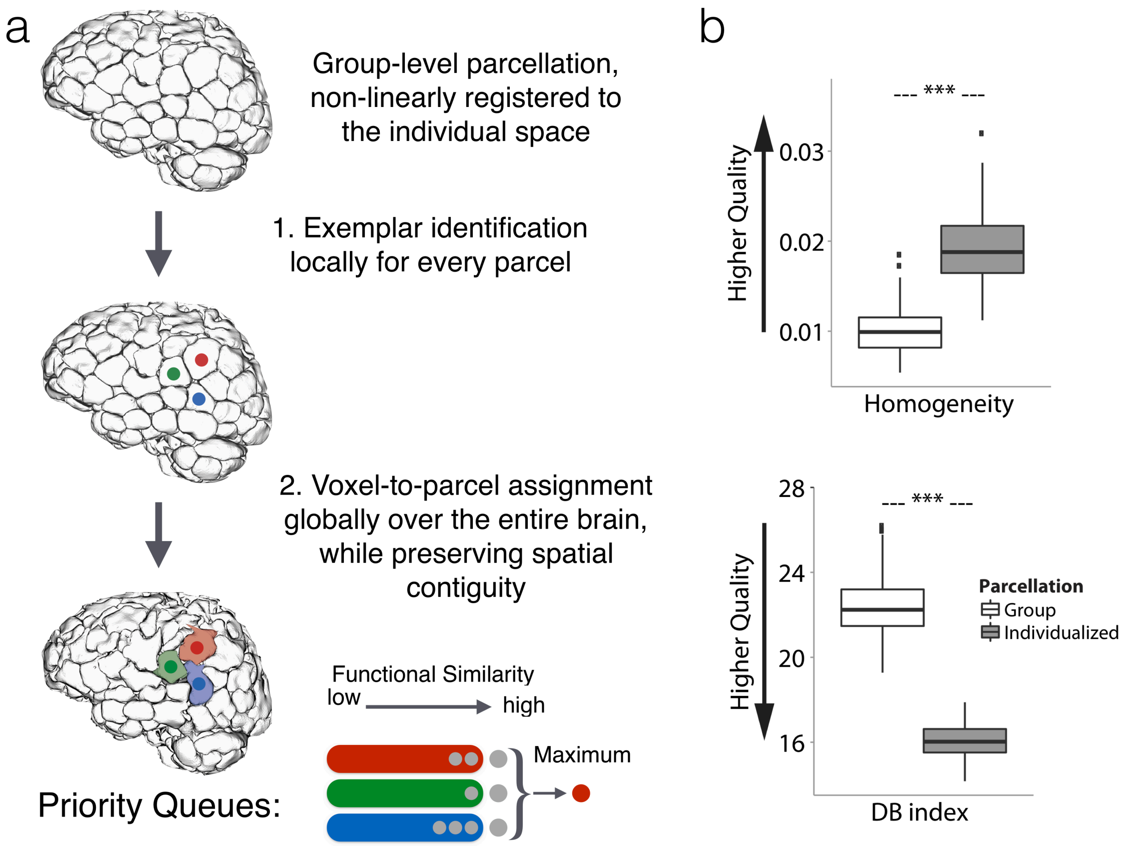

**Figure 1. Individualized state-specific parcellation pipeline and evaluation.** a) Three-step parcellation pipeline: Step 1: Starting from data in common space, a group-level functional atlas is first applied to the individual subject. Step 2: For every group-defined parcel in the individual brain, an exemplar is identified by maximizing a monotone nonnegative submodular function; here, three exemplars are shown in red, green, and blue. Step 3: Every voxel in the individual brain is assigned to the functionally closest exemplar while taking the spatial contiguity of the parcel into account. The spatial contiguity is assured by utilizing priority queues. Every exemplar $i$ is assigned a priority queue (denoted as $q_i$), here depicted as red, green, and blue queues which correspond to red, green, and blue exemplars, respectively. Initially, all the queues are empty. In the first round, spatial neighbors of exemplar $i$ are pushed into $q_i$. The voxels in each queue are sorted according to their functional distance to the corresponding exemplar such that the voxel with minimum functional distance (maximum similarity) is in front. Next, the front voxel in each $q_i$ is considered as a potential candidate for being assigned the label $i$. Among all these candidates, the one with minimum distance to its corresponding exemplar is selected and assigned the exemplar's label; here the candidate voxel from the 'red' queue is selected and labeled 'red'. Next, this voxel is popped out of the queue, and all of its spatial neighbors are pushed into the same queue. The algorithm continues until all the voxels are assigned a label. Note that at every step of the algorithm the labeled voxel is ensured to be spatially connected to its exemplar (either directly or through other previously labeled voxels). b) Quantitative results for assessing the quality of parcellation. Homogeneity (top) and DB index (bottom) comparison between the individualized, state-specific parcellations and the initial group-level parcellation, represented as

box plots over sessions in Yale data, with the central mark indicating the median, and the bottom and top edges of the box indicating the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually. *** $p < 2.2e-16$, two-tailed Mann-Whitney test.

### 2.4.1 Exemplar identification

The exemplar identification algorithm can be viewed as a data summarization step, where the goal is to summarize a massive amount of data by fewer representative points or exemplars. A classic way of defining such exemplars is by finding the set $S$ that minimizes the following loss function, subject to the constraint $|S| = k$ (known as $k$-medoid problem).

$$L(S) = \frac{1}{|V|} \sum_{v \in V} \min_{e \in S} d(v, e). \tag{1}$$

In this equation, $V$ is the ground set consisting of all data points, $d: V \times V \rightarrow R$ is a dissimilarity function defined on every pair of data points, and $S$ is the objective exemplar set. Intuitively, $L(S)$ measures how much information we lose if we summarize the entire ground set to the exemplar set by representing each data point with its closest exemplar.

Minimizing this loss function (1) is NP-hard, as it requires exponentially many inquiries. Using an appropriate auxiliary exemplar $v_0$, we transform the minimization of (1) into the maximization of a non-negative monotone submodular function (Gomes and Krause, 2010), for which general greedy algorithms provide an efficient $1 - 1/e \approx 0.63$ approximation to the optimal solution (Nemhauser et al., 1978):

$$\begin{gathered} \max_{S \subseteq V} f(S), \\ s.t.\ |S| \leq k, \end{gathered} \tag{2}$$

where:

$$f(S) = L(v_0) - L(S \cup v_0). \tag{3}$$

In practice, the greedy algorithm provides a considerably closer approximation to the optimal solution (see Salehi et al. (Salehi et al., 2018)). For the choice of auxiliary exemplar, any vector $v_0$ whose distance to every data point is greater than the pairwise distances between data points can be used.

**Definition 1 (Submodularity).** A function $f: 2^V \to R$ is submodular if for every $A \subseteq B \subseteq V$ and $e \in V \backslash B$ it holds that $f(A \cup e) - f(A) \geq f(B \cup e) - f(B)$. That is, adding an element $e$ to a set $A$ increases the utility at more than (or at least equal to) adding it to $A$'s superset, $B$, suggesting natural diminishing returns.

### 2.4.2 Spatially-constrained voxel-to-parcel assignment

After identification of all exemplars (one per parcel), every voxel in the individual brain is assigned to the functionally closest exemplar while taking the spatial contiguity of the parcel into account. The spatial contiguity is assured by utilizing priority queues. Every exemplar $i$ is assigned a priority queue (denoted as $q_i$). Initially, all the queues are empty. In the first round, spatial neighbors of exemplar $i$ are pushed into $q_i$. The voxels in each queue are sorted according to their functional distance to the corresponding exemplar such that the voxel with minimum functional distance (maximum similarity) is in front. Next, the front voxel in each $q_i$ is considered as a potential candidate for being assigned the label $i$. Among all these candidates, the one with minimum distance to its corresponding exemplar is selected and assigned the exemplar's label. Next, this voxel is popped out of the queue, and all of its spatial neighbors are pushed into the same queue. The algorithm continues until all the voxels are assigned a label. Note that at every step of the algorithm the labeled voxel is ensured to be spatially connected to its exemplar (either directly or through other previously labeled voxels).

### 2.5 Parcellation evaluation

The quality of the proposed parcellation approach was assessed and compared with the initial group-level parcellation, using two internal clustering validation methods: Homogeneity (Figure 1b, top panel) and Davies-Bouldin index (DB) (Davies and Bouldin, 1979) (Figure 1b, bottom panel). Homogeneity was assessed by calculating the average cross-correlations within each node, and averaging over all the nodes in the parcellated brain. DB index was used to assess the clustering ability in maximizing the intranode compactness and the inter-node separation. Higher homogeneity and lower DB index values indicate higher clustering quality.

### 2.6 Implementation details

Here we set $k = 1$, as we attempt to identify one exemplar per parcel (see Salehi et al. (Salehi et al., 2017) for details of interpretation and alternative approaches). For the choice of dissimilarity measure, we used squared Euclidean distance, after normalizing all the voxel-level time courses to a unit norm sphere centered at the origin. A point with the norm greater than 2 was used as the auxiliary exemplar. The parcellation algorithm was applied to each fMRI run (each individual in each state) independently, and thus was efficiently employed through parallelization. For HCP data, we restricted our analysis to left-right (LR) phase encoding.

### 2.7 Initial group-level parcellations

As the initial group-level parcellation, we primarily used a 268-parcel atlas, previously defined in our lab using a spectral clustering algorithm on resting-state data of a healthy population (Finn et al., 2015; Shen et al., 2013). We replicated the results with more fine-grained atlases including 368, 1041, and 5102 parcels. The 368-parcel atlas was defined by integrating the parcellation of cortex from Shen et al. (Shen et al., 2013), subcortex from the anatomical Yale Brodmann Atlas (Lacadie et al., 2008), and cerebellum from Yeo et al. (Thomas Yeo et al., 2011a). Similarly, the 1041-parcel parcellation was defined by integrating the subcortical and cerebellum portion of the 368-parcel parcellation with the 1000-parcel cortex parcellation from Yeo et al. (Thomas Yeo et al., 2011a). To define the 5102-parcel atlas, we started from the 1041-parcel parcellation and randomly divided all the parcels until the number of voxels per parcel reached approximately 25.

### 2.8 Statistical voting-based ensemble analysis

To estimate the similarity of parcellations within and across states, we employed a statistical voting-based ensemble analysis. The reason for this analysis is two-fold: First, to help rule out noise and session effects as potential confounds; and second, to aid with the statistical comparison across states by generating a large distribution of state-specific parcellations. For each condition, we divided all sessions (for Yale and MSC data) or individuals (for HCP data) into two equal-size groups: group 1 and group 2. We took the relative majority vote over the parcellations of each group by assigning each voxel to the parcel for which the maximum number of sessions voted. This resulted in two parcellations for each functional state, one for each group. Next, we assessed the similarity between every pair of the non-overlapping parcellations both within and across states. For instance, if there are $m$ functional states, this analysis generates an $m \times m$ matrix, where each element $(i, j)$ represents the similarity between the parcellation of group 1 for state $i$ and the parcellation of group 2 for state $j$. Thus, the diagonal elements represent the within-state similarities while the off-diagonal elements represent the cross-state similarities. We repeated the entire analysis 1000 times, generating an ensemble of $m \times m$ similarity matrices. The normalized distribution of the within-state and cross-state similarity values were depicted as histograms, and compared using the non-parametric Kolmogorov–Smirnov test (Figures 2d,e, 4d,e, and 5d,e). The averages of these similarity matrices were also displayed (Figures 2a,b, 4a,b, and 5a,b).

### 2.9 Similarity measures between parcellations

We compared parcellations at two different scales: at the fine scale, we studied the ratio of voxels that change their parcel assignment across different parcellations. The fine-scale similarity was calculated using 1 – normalized Hamming distance (Figures 2a, 4a, 5a). At the coarse scale, we studied the changes in the parcel sizes across parcellations. The coarse-scale similarity was computed using Spearman correlation between parcel-size vectors (Figures 2b, 4b, 5b).

### 2.10 Functional state decoding using parcel size as feature

We established a cross-validated predictive model that predicts the functional state of each unseen sample solely based on the size of parcels in that parcellation. Using a $k$-fold cross-validated approach, we trained and tested a gradient boosting classifier (GBM; with 300 estimators and learning rate = 0.1) using parcel sizes as features and the functional state as output. We randomly divided the entire data into *k* folds (*k*=*n* for the Yale and MSC data sets, where *n* is the total number of runs, and *k*=10 for the HCP data set). At each step, the model was trained on *k*-1 folds and tested to predict the state of the left-out fold. The predictive power of the model was estimated using precision (also known as positive predictive value) and recall (also known as sensitivity), calculated separately for each state. Both precision and recall measures range between 0 and 1, with higher values indicating higher predictive power. Precision calculates what fraction of the retrieved instances were actually relevant, while recall expresses what fraction of relevant instances were retrieved. In the case of 10-fold cross-validation (i.e., for the HCP data set), we repeated the predictive analysis 100 times to account for the randomness of the folds, and reported the precision and recall measures as the mean and standard deviation across all iterations.

To evaluate the significance of the results, we employed non-parametric permutation testing: we randomly permuted the output vector (here the functional states) 1000 times, and each time ran the permuted values through the same predictive pipeline and calculated the precision and recall measures of the permuted states.

### 2.11 Task performance prediction using parcel size as feature

To establish the relevance of parcellation boundary changes to within-task variation in brain state, we developed a cross-validated predictive pipeline that predicts task performance during each Yale data session from the parcel sizes in the parcellation for that session. Using a leave-one-out cross-validated approach, we trained and tested a ridge regression model (with regularization parameter = 1), using parcel sizes as features and the performance scores as output. For each task, a model was trained on *n*-1 sessions (of that task) and used to predict performance for the left-out session. Analyses were performed for each task independently. To measure performance, we used *d'* for *n*-back and gradCPT tasks, *RT variability* for eyes and card-guessing tasks, and *missing probability* for SST. The predictive power was estimated using r-values, defined as the square root of the percentage of the explained variance. To estimate the significance of the results, we employed non-parametric permutation testing, where we randomly permuted the behavioral scores 1000 times, and each time ran the permuted scores through the predictive pipeline and calculated the r-value.

### 2.12 Voxel uncertainty analysis

Next, we demonstrate that the state-evoked parcel reconfigurations are not driven by voxels with low certainty in their parcel assignment. To this end, we built upon our voting-based ensemble analysis and quantified voxel uncertainty as the proportion of times across 1000 iterations that a voxel was assigned to different parcels between the two groups.

### 2.13 Effects of parcellation algorithm on the state-evoked parcel reconfigurations

One of the main advantages of the proposed parcellation algorithm is that correspondence of parcels from different parcellations is maintained, making for straightforward comparisons of the resulting atlases. Most other parcellation approaches could yield completely different atlases, and in that case, it can be difficult to match parcels for quantitative comparison. However, to ensure that the observed state-evoked reconfigurations did not depend on the choice of parcellation, we employed two additional experiments. First, we repeated our analysis using a slight modification of our exemplar-based algorithm, where instead of employing a data-driven exemplar selection we manually fixed the exemplar to be a random voxel in the parcel, with a particular exemplar for a given parcel fixed across all sessions and conditions. This means the initial reference voxel was identical for a given parcel across all sessions and conditions. In this case, only the time-course data could shift the parcellation since all parcellations started from the same exemplar voxels across conditions and sessions. Results remained largely unchanged (Figure S4). Second, we replicated our results using Wang et al.'s individualized parcellation algorithm, which is based on an iterative *k*-means clustering approach (Wang et al., 2015). Wang's method has an averaging step for which we used a standard uniformly weighted averaging. For results, see Figure S5.

### 2.14 Effects of task activation on the state-evoked parcel reconfigurations

To address the question of whether the parcellation reconfigurations were driven by task activation, we performed two experiments. First, we tested whether the mean task activation (calculated per parcel) correlates with differences in parcel size relative to rest (Figure S6). Next, we eliminated parcels with significant task activation for any of the tasks and repeated the entire analysis (Figure S7). Significance was defined as $|z|>1.96$, associated with 95% confidence interval or $p<0.05$. These two experiments were performed on both Yale and HCP data sets. For Yale data, because all the tasks were continuous performance tasks, we approximated the task activations per parcel by computing the difference between the average temporal signal during task and rest. For HCP data, we generated task effect size maps using all available individuals' volume-based, FEAT-analyzed, first-level GLM output (COPE files) from the 1200 Subjects Release (S1200) for a given task to generate, using FSL FEAT's FLAME (FMRIB's Local Analysis of Mixed Effects (Smith et al., 2004)), cross-subject, voxel-wise Cohen's d effect-size contrast maps from t-statistic maps. We used the following contrasts for each task: REWARD-PUNISH (Gambling, cope6), 2BK-0BK (WM, cope11), FACES-SHAPES (Emotion, cope3), STORY-MATH (Language, cope4), AVG (Motor, cope7), REL-MATCH (Relational, cope4), and TOM-RANDOM (Social, cope6). We then applied the initial 268-parcel group-level parcellation (Shen et al., 2013) to these voxel-level maps to calculate a mean task effect size per parcel for each of the tasks.

### 2.15 Effects of head motion on the state-evoked parcel reconfigurations

To rule out the possibility that state-evoked reconfigurations were simply driven by characteristic head motion patterns specific to each functional condition, we tested whether there is significant difference in head motion between different functional conditions in Yale data. We performed pairwise Wilcoxon signed-rank test (which is a non-parametric statistical hypothesis test to compare two dependent samples) on the mean frame-to-frame displacement, and corrected for multiple comparison using Bonferroni correction (Figure S8).

## 3. Results

In the first data set, which we refer to as Yale data, we applied our parcellation algorithm to each functional run, generating one parcellation atlas for each condition and each scanning session (8 conditions × 30 sessions = 240 atlases total). By taking the relative majority vote over all sessions of each condition, we generated condition-specific parcellations. Figure 2a visualizes these parcellations using a force-directed graph, with edge weights indicating the similarity between parcellations, measured by $r_{\mathrm{Hamming}}$, which estimates the percentage of voxels with similar parcel assignment. Parcels are colored by the magnitude of reconfiguration in the given condition relative to Rest 1, where reconfiguration is defined as the percentage of voxels that change their parcel assignment. Figure 2a demonstrates that parcels with high reconfiguration are broadly distributed and condition specific. Increased similarity of reconfiguration maps is observed (Figure 2a) among condition-specific parcellations with increased pairwise similarity (Figure 2b). Consistent with our expectations, Rest 1 and Rest 2 parcellations are highly similar to each other, while the parcellation for the movie-watching condition is the most distinct (Figure 2). For visualization and interpretation purposes, we created a video of these state-specific parcellations (Video S1) as well as an interactive brain visualization platform (see http://htmlpreview.github.io/?https://github.com/YaleMRRC/Node-Parcellation/blob/master/Parcellation_visualization.html).

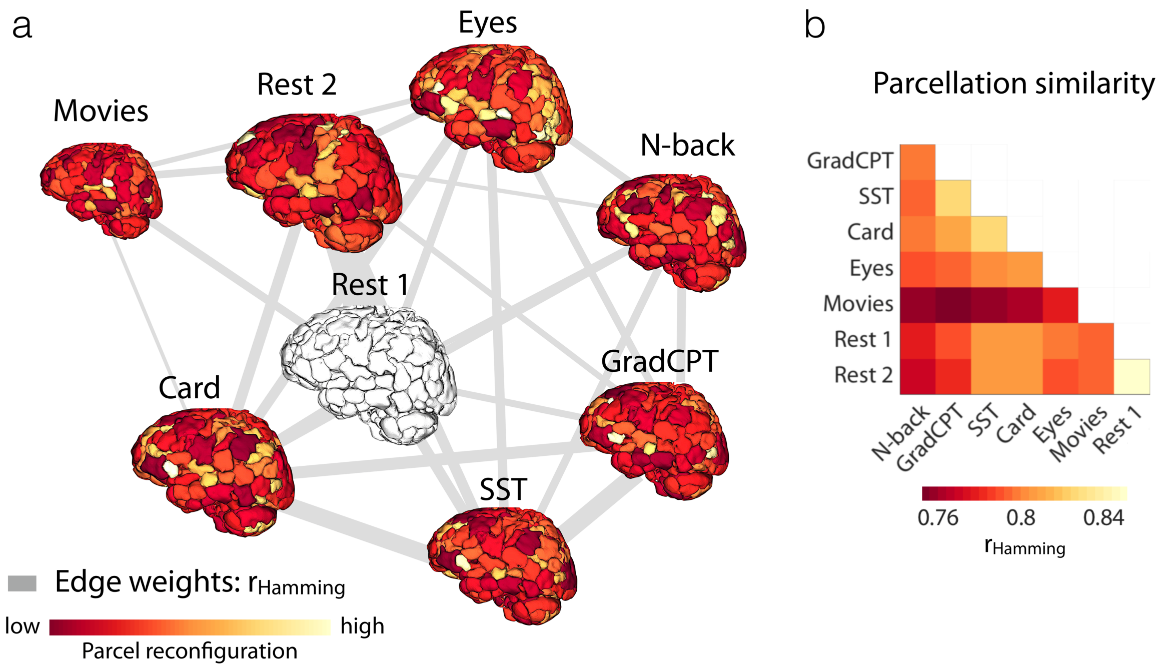

**Figure 2. Visualization of the condition-specific functional atlases.** (a) Condition-specific functional atlases are visualized in a force-directed graph, with edge weights indicating the similarity between parcellations, measured by $r_{\mathrm{Hamming}}$ = 1 - normalized Hamming distance. Force-directed graphs attempt to visually organize networks such that the energy of the graph as a whole is minimized. This is accomplished by assigning both repulsive and attractive forces to each pair of nodes such that the nodes with stronger interconnections are displayed closer to each other and the ones with weaker connections are more distant. Brain size is proportional to the graph theory measure degree. Edge thickness is proportional to the edge weights. Parcels are colored by the magnitude of reconfiguration in the given condition relative to Rest 1. (b) Cross-condition parcellation similarity measured by $r_{\mathrm{Hamming}}$.

### 3.1 Statistical voting-based ensemble analysis

We quantitatively examined these parcellation changes across conditions using the statistical voting-based ensemble method (see Methods). The results demonstrate that parcellations are significantly more similar within each functional condition than across different conditions, both for fine-scale (Figure 3a, d; K-S test; $p<0.001$) and coarse scale similarity measures (Figure 3b, e; K-S test; $p<0.001$). This finding suggests that parcellations are highly reproducible within a functional condition, and rearrange significantly between conditions. These changes are reflected even in a coarse-scale feature of spatial topography, i.e., parcel-size.

### 3.2 Functional state decoding using parcel size as feature

Next, we demonstrate that the observed task-induced parcellation reconfiguration is consistent across sessions and specific to each condition. To this end, we built a cross-validated predictive model that predicts the functional condition under which novel runs were acquired based solely on the parcel size vector. Prediction accuracies—measured as precision and recall for each condition—were significantly higher than random for all conditions (mean accuracy = 71%; Figure 3c). That a measure of spatial topography as coarse as parcel size can significantly predict which task a subject is performing suggests that functional brain atlases reconfigure with task-induced brain state in a consistent manner, forming a generalizable and robust signature of brain organization during that condition. Successful prediction also demonstrates that cross-condition reconfigurations are not driven by noise.

### 3.3 Task performance prediction using parcel size as feature

While the overall stability of the functional atlases within each condition was high, subtle variations were observed across different sessions of the same condition. To examine if parcel reconfigurations contain meaningful session-specific information, we attempted to predict task performance using parcel size as the model feature. Task performance was used as a proxy for level of engagement in the task, reflecting more fine-grained variations in brain state than cannot be captured by simply considering task as a state (see Methods for details of the behavioral measures that were used for each task). Task performance was successfully predicted from parcel sizes (Figure 3e), indicating that reconfigurations from session-to-session within a task condition reflect meaningful changes in brain function. The success of these predictions again rules out the possibility that even session-to-session reconfigurations are driven by noise.

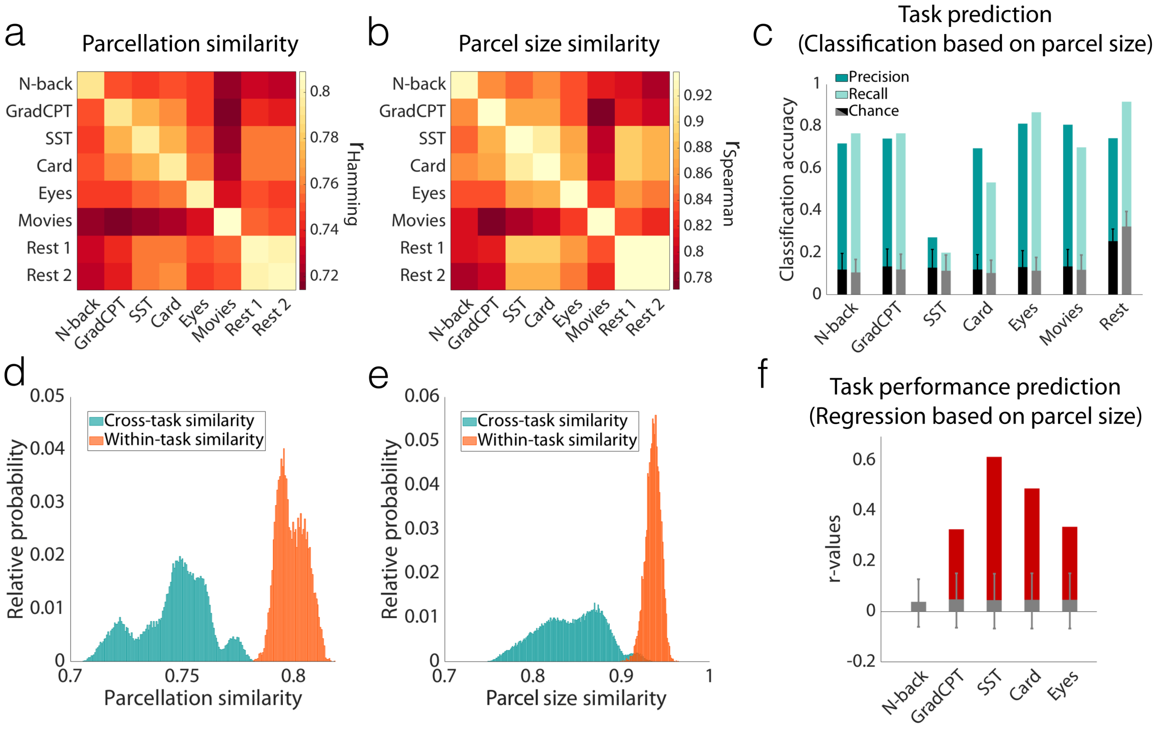

**Figure 3. Parcel definitions change with task condition; Yale single-subject data.** a) Pairwise parcellation similarity was calculated within and across functional conditions (8 conditions, N=30 sessions), using voting-based ensemble analysis with 1000 iterations. The matrix represents the average over all iterations. Similarity was assessed by $r_{\text{Hamming}} = 1 -$ normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{\text{Spearman}}$) as a proxy for parcellation similarities. c) The bars reflect the accuracy of predicting the functional condition using a leave-one-out cross-validated GBM classifier with parcel sizes as features. The predictive power is measured by precision (dark cyan) and recall (light cyan) values for each condition. The precision (black) and recall (gray) values of 1000 null models are also reported (error bars represent ±s.d.). d) A histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. Rest 1 and Rest 2 are grouped into one condition. The within- and cross- distributions are significantly different (K-S test; $p<0.001$). e) A histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. Again, the two distributions are significantly different (K-S test; $p<0.001$). f) The bars report the accuracy of predicting task performance using a leave-one-out cross-validated linear regression model with parcel sizes as features. Predictive power is measured by square root of coefficient of determination (r-value). The r-values of 1000 null models (for permutation testing) are also reported (error bars represent ±s.d.).

### 3.4 Voxel uncertainty analysis

Figure 4a shows that the majority of the voxels have less than 10% uncertainty in their parcel assignment, suggesting that the parcellations are reliable within a functional condition. Figure 4b visualizes the spatial distribution of voxel uncertainty in the brain. From this visualization, it is clear that the shape of the white areas in each parcel, representing voxels with low uncertainty (uncertainty < 10%), changes substantially between conditions, suggesting that the state-evoked reconfigurations are not simply driven by noisy boundaries, but that stable voxels form parcels of different shapes.

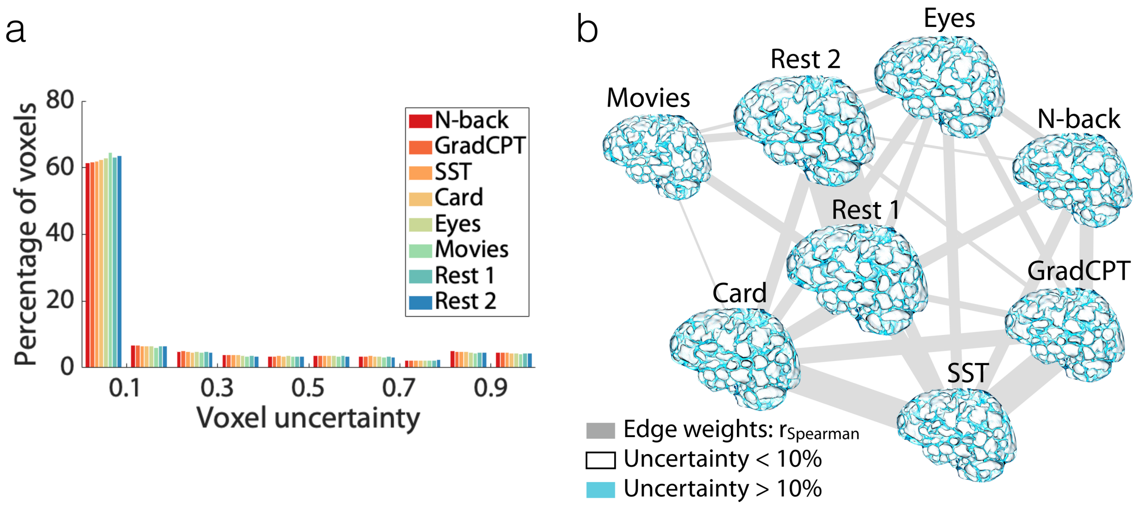

**Figure 4. Voxel uncertainty analysis.** Voxel uncertainty was computed for each functional condition as the proportion of times across 1000 iterations of the voting-based ensemble analysis that the voxel was assigned to different parcels between group 1 and group 2 (see Methods). a) The histogram of voxel uncertainty for each functional condition. b) The spatial distribution of voxel uncertainty on the brain, visualized in a force-directed graph with edge weights indicating the similarity between voxel uncertainty distributions across conditions, measured by $r_{Spearman}$. Brain size is proportional to the graph theory measure degree. Edge thickness is proportional to the edge weights. Voxels with a low uncertainty (uncertainty < 10%) are shown in white, and the parcels defined by these clearly change shape between conditions. The rest of the voxels (uncertainty > 10%) are shown in blue.

### 3.5 Replication of the state-evoked atlas reconfiguration across data sets

#### 3.5.1 Midnight Scan Club (MSC) data set

We replicated our main findings by leveraging a publicly available data set: Midnight Scan Club (MSC) (Gordon et al., 2017b). For each individual, we repeated the voting-based ensemble analysis described above and calculated the similarity between each task- (or rest-) based pair of parcellations. Figure 5 demonstrates the pairwise parcellation similarities, averaged over individuals (Figure 5a, d: fine-scale similarity; Figure 5b, e: coarse-scale similarity). Figure S1 shows the similarity matrices of each individual separately. Confirming the finding from the single-subject data described above, we observed that functional atlases are significantly more similar within condition than across conditions (K-S test, $p<0.001$). Similarly, parcel sizes are significantly more similar within a condition (K-S test, $p<0.001$). Predictive modeling based on parcel size could significantly predict condition for a novel run (mean accuracy = 66%; Figure 5c) and generalized across individuals. That all the predictions were successful, and consistent across individuals, is a strong indication of the significance of functional parcel reconfigurations associated with different tasks. The findings from this independent data set replicate the single-subject findings.

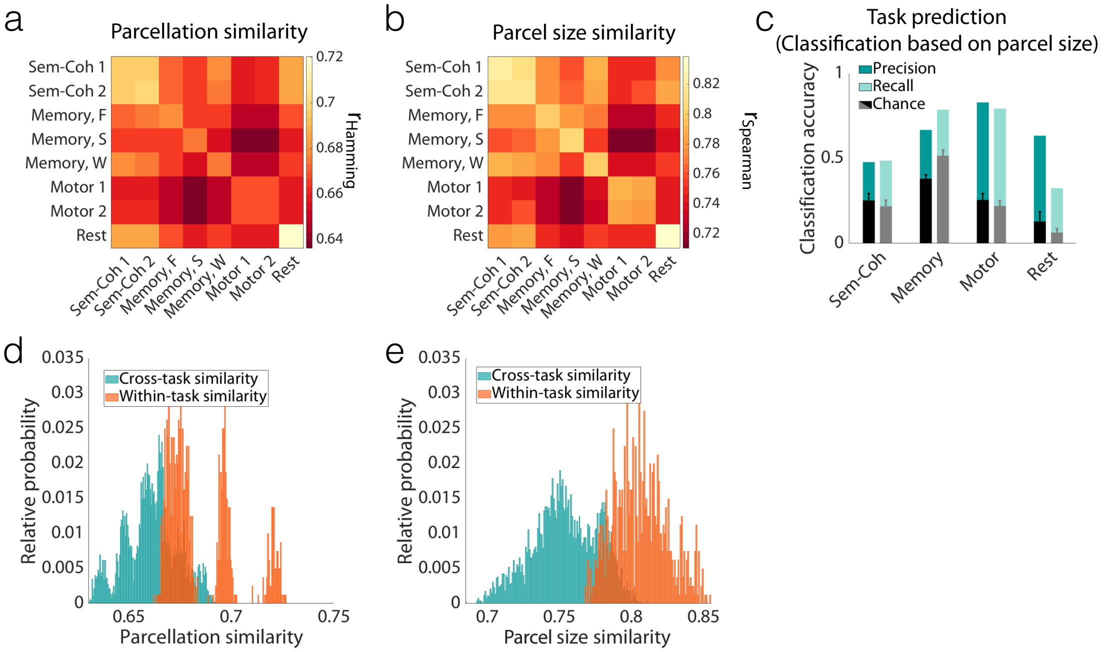

**Figure 5. Replication of the finding that parcel definitions change with task condition; Midnight Scan Club (MSC) data.** a) Pairwise parcellation similarity was calculated within and across functional conditions, averaged over all individuals (8 conditions, N=10 sessions, n=8 subjects). For every individual, voting-based ensemble analysis was used with 100 iterations. The matrix represents the average over all iterations. Similarity was assessed by $r_{Hamming}$ = 1 – normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities. c) The bars report the accuracy of predicting the functional condition using a leave-one-out cross-validated GBM classifier with parcel sizes as features. The predictive power is measured by the precision (dark cyan) and recall (light cyan) values for each condition. The precision (black) and recall (gray) values of 1000 null models are also reported (error bars represent ±s.d.). d) A histogram of parcellation similarities for all 100 iterations (averaged over all individuals) is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. Sem-Coh 1 and Sem-Coh 2 are grouped into one condition, as are Motor 1 and Motor 2 conditions. The within- and cross- distributions are significantly different (K-S test; p<0.001). e) A histogram of parcel size similarities for all 100 iterations (averaged over individuals) is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; p<0.001). Sem-Coh, semantic-coherence task; Memory F, S, and W, incidental memory task with faces, scenes, and words stimuli, respectively.

### 3.6 Human Connectome Project (HCP) data set

The results from the previous two data sets relied upon the construction of individual atlases across both sessions and conditions and were evaluated in terms of how the atlas *within an individual* changed between conditions. Next, we used the Human Connectome Project (HCP) 900 Subjects release (S900) data (Van Essen et al., 2013) to determine if such condition-dependent reconfigurations could be observed when measured across multiple subjects introducing additional variance through inter-subject anatomic differences. We repeated the voting-based ensemble analysis described above, this time replacing sessions with subjects, such that we considered multiple subjects in multiple conditions and calculated the similarity between each pair of parcellations. Figure 6 demonstrates the pairwise parcellation similarities across subjects. Despite the introduction of inter-individual anatomic and functional organization variance, the primary finding that the parcellation map changes with condition remains highly significant (Figure 6a, d: fine-scale similarity [K-S test, $p<0.001$]; Figure 6b, e: coarse-scale similarity [K-S test, $p<0.001$]). As an even stronger assessment of generalizability, we repeated our previous analysis to predict, based on parcel size, the task during which the data were collected for previously unseen subjects. Given that we had a larger sample size ($n=514$), we further challenged our model by employing a more rigorous 10-fold cross-validated pipeline (rather than leave-one-out). We observed that our model could significantly predict, in novel subjects, the task administered while the data were acquired (Figure 6c, mean accuracy = 73%) based on parcel size, alone. This observation is an even stronger replication of our main findings; that is, the observed state-evoked parcellation reconfigurations are robust and reliable not only across different sessions, but also across distinct individuals.

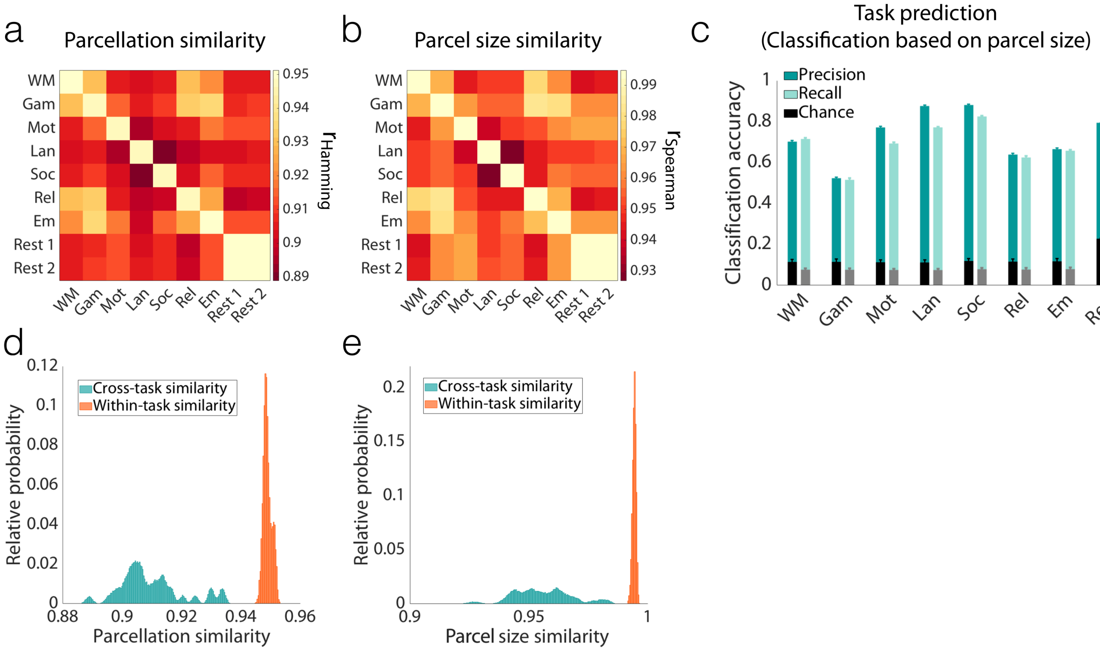

**Figure 6. Replication of the finding that parcel definitions change with task condition, even when considered across individuals; Human Connectome Project (HCP) data.** a) Pairwise parcellation similarity was calculated within and across functional conditions (9 conditions, n=514 subjects), using voting-based ensemble analysis with 1000 iterations. The matrix represents the average over all iterations. Similarity was assessed by $r_{Hamming}$ = 1 – normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities. c) The bars report the accuracy of predicting the functional condition using 10-fold cross-validated GBM classifier with parcel sizes as features, iterated 100 times. The predictive power is measured by precision (dark cyan) and recall (light cyan) values for each condition (reported as mean and s.d. across iterations). The precision (black) and recall (gray) values of the 1000 null models are also reported (error bars represent ±s.d.). d) A histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. Rest 1 and Rest 2 are grouped into one condition. The two distributions are significantly different (K-S test; $p<0.001$). e) A histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). WM, working-memory task; Gam, gambling task, Mot, motor task; Lan, language task; Soc, social task; Rel, relational task; Em, emotion task.

### 3.7 Robustness of the state-evoked atlas reconfiguration across scales

In our final analysis, we sought to address the question of parcellation scale, asking if these reconfigurations are simply due to the rather large parcel sizes typically targeted in current parcellation approaches. If the parcel resolution were too low, then it could be the case that our parcels are composed of smaller subunits that do not change their shape but potentially exhibit altered connectivity with the other subunits dependent upon the functional state (hence modifying how they are grouped into a single parcel). If this is the case, then making atlases with more parcels should get to a level where the parcels remain unchanged even with changes in brain state. In the limit of parcels reduced to the size of a voxel, there can be no condition-induced change in parcel definition since the parcel is defined by the voxel size. Here we attempted to determine the critical resolution at which parcellations stabilize. To address this question, we repeated our analyses with Yale data, using atlases containing greater numbers of parcels, from 368, to 1041, to 5102 parcels (see Methods). In each case, parcel reconfigurations were observed, even when the number of parcels was increased to 5102, more than ten times the typical number of parcels used in functional connectivity analysis (Bellec et al., 2015; Nieuwenhuys, 2013; Van Essen et al., 2012) (Figure 7). Results with 368-parcel and 1041-parcel atlases are reported in Figure S2 and Figure S3, respectively. To date, the human neuroscience community has not considered atlases with 5000 or more parcels.

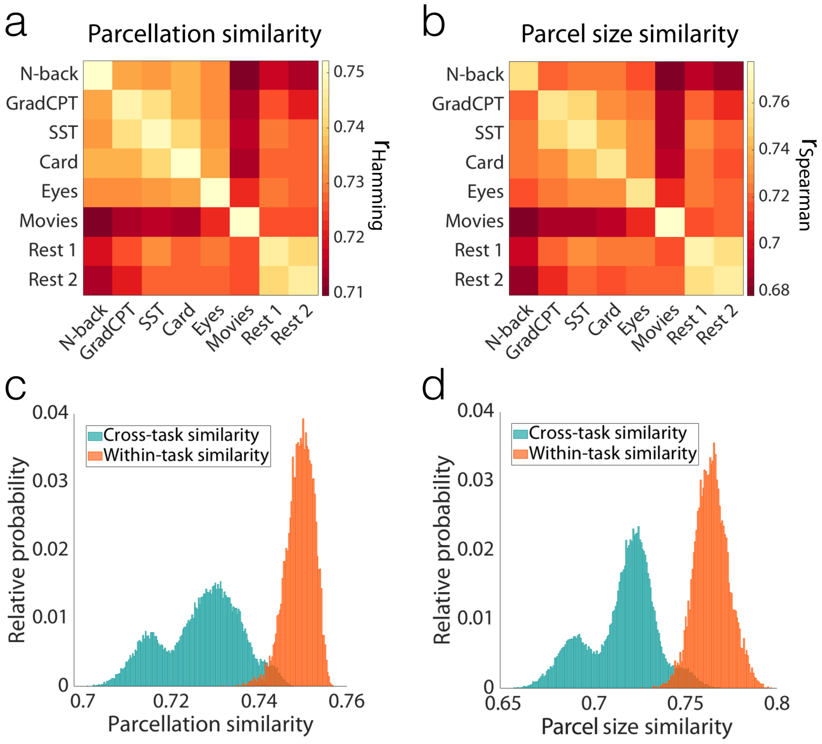

**Figure 7. Parcel size effects:** Even for an atlas with 5102 parcels (on average 25 voxels per parcel), parcel definitions change with task condition, with high reliability within conditions. a) Pairwise parcellation similarity was calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by $r_{\text{Hamming}} = 1 -$ normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{\text{Spearman}}$) as a proxy for parcellation similarities. c) The histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). d) The histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$).

## 4. Discussion

Together, these findings suggest that there is no single functional atlas for the human brain. The boundaries of functionally defined parcels change with changes in brain state in a consistent and reproducible manner. These reconfigurations appear to be cognitively meaningful, as evidenced by their utility in predictive models of task condition and within-condition task performance.

### 4.1 Reconfiguration of the connectome

These results are consistent with the extensive evidence for reconfigurations of the functional connectome. There is growing evidence that patterns of functional connectivity change with changing brain states (e.g., as induced by distinct tasks), and that these changes are functionally significant (Greene et al., 2018). Recent work has highlighted individual differences in parcel spatial configuration and their impact on brain-behavior relationships (Bijsterbosch et al., 2018), but the common assumption that parcel boundaries are fixed for a given individual remains unchallenged. In fact, we show that while boundary reconfigurations may be relatively modest overall, they are highly task-specific, consistent, and functionally significant both within and across individuals, suggesting that tasks induce reliable perturbations of a core functional architecture of the human brain. Taking such reconfigurations into account may further inform efforts to relate changing patterns of functional connectivity to behavior, clinical symptoms, and cognition.

The stability of the state-evoked boundary reconfigurations across a large cohort of individuals, as demonstrated by the HCP data (Figure 6), is important. Previous work has shown that patterns of functional connectivity are unique to each individual (Finn et al., 2015), and any state-evoked reconfigurations are also individual-specific (Gratton et al., 2018). That the average of these state-evoked reconfigurations across individuals does not mask out the effect of brain state is a significant finding, suggesting that task-induced parcel reconfigurations are not idiosyncratic and subject-specific, but rather robust and generalizable; that is, despite significant individual differences in brain functional organization, task-induced changes in this organization, as reflected in parcel boundary shifts, are highly similar across individuals.

### 4.2 Composition of functional subunits

The cause of such parcel reconfigurations remains an important open question. At the atlas sizes considered here, any individual functional parcel will contain hundreds of thousands of neurons and may span multiple cortical subareas (Van Essen and Glasser, 2018), and a growing literature suggests that the BOLD signal can be used to identify inter-digitated neural representations, even within a circumscribed area with well-defined functional specialization (Haxby, 2012; Haxby et al., 2014; Norman et al., 2006). And while it is possible that focal task activations affect parcel boundaries, they do not fully explain the present results, as parcel reconfiguration is broadly distributed across the brain, and change in parcel size relative to rest is not related to parcel task activation (Figure S6). Furthermore, the same results were obtained even after exclusion of voxels with significant task activations (Figure S7). Together, this suggests a more complex picture of parcels in flux across the brain, wherein different groups of neurons synchronize to execute different tasks resulting in different parcel definitions. It is possible that invariant boundaries may not be found until the acquisition resolution is of the order of a fundamental cortical unit, such as microcolumns (300 – 500 μm) (Molnár, 2013; Mountcastle, 1957) or less. This presents an exciting opportunity for future work with high-field fMRI to interrogate the underpinnings of boundary shifts at high resolution.

The existence and nature of such a fundamental unit—and its potential impact on the BOLD signal—are open questions, but regardless of their answer, the findings here of substantial parcel reconfiguration even at atlas sizes of 5000 parcels (consisting of, on average, 25, 2 $mm^3$ voxels, or ~5.8 $mm^3$ of tissue), suggest that current human functional parcellation atlases are far from this limit. Nevertheless, our goal does not have to be identifying parcels that correspond to a minimal functional or computational unit. In practice, a parcel that has a homogeneous, distinct identity in a given state is sufficient for many analyses. In other words, as long as the atlas is consistent for the state of interest, atlases with 200 – 400 parcels are entirely appropriate.

### 4.3 Additional considerations

The exemplar-based parcellation approach used in this work imposes a constraint that holds the total number of parcels fixed and eliminates the problem of establishing correspondence between parcels under different conditions. Maintaining correspondence across parcellations facilitates quantitative assessment of the change in each parcel. It is possible that parcels not only reconfigure, but also blend and/or split, leading to a change in the total number of parcels with a change in brain state. This additional degree of freedom would make quantification of parcel changes across atlases obtained under different conditions more difficult due to the correspondence problem, while also potentially amplifying our main finding that there is no fixed functional atlas at the scale typically used in fMRI studies. In other words, without an exemplar-based approach, the reconfiguration of these parcels would likely be more extensive than that shown here.

The algorithm presented also constrains parcels to be contiguous. Without this constraint, it is possible for connectivity-based parcellations to group voxels, that may be widely separated spatially, into single parcels with very high homogeneity in their time-courses. If parcels are allowed to be non-contiguous, it becomes an open question as to how much distance is allowed between the different parts of a parcel. With no consideration given to spatial contiguity, a successful parcellation algorithm could simply rank voxel time-courses by similarity and then split this ranking into $n$ parcels, which would provide maximal parcel homogeneity (a common measure of parcellation success), but would identify broadly distributed brain networks rather than functionally specialized, circumscribed brain regions.

The present work focuses on defining parcels with hard boundaries. However, one can use the voxel uncertainties to define soft borders, as displayed in the blue boundaries in Figure 4. One can also define a probabilistic soft parcellation by considering the consistency of voxel-to-parcel assignments across different functional conditions. While defining such soft borders is feasible, most parcellation approaches to date have focused on hard boundaries. One application of such functional atlases is in building connectivity matrices, and with soft boundaries the question of how to weight the time-courses across these boundaries is currently unanswered. Given the use of parcellations with hard boundaries for validated, commonly used functional connectivity analysis pipelines, we chose to focus here on such hard boundaries.

States are defined in this work as task conditions with acquisitions spanning a series of continuous performance, event-related, and blocked tasks. It is likely that parcel reconfiguration occurs over considerably shorter periods of time than these minutes-long task intervals, particularly given the growing literature on the dynamic nature of functional brain organization (Cohen, 2018) and the present finding that there are meaningful, state-induced changes in parcellation boundaries within a given task (i.e., that task performance can be predicted from parcel size). Future work may seek to characterize how parcellation boundaries shift to reflect the dynamic reconfiguration of macroscale neural circuitry underlying moment-to-moment changes in brain state.

This work also does not invalidate multimodal parcellation approaches that combine both anatomical and functional data (including functional data across a wide range of task and resting states), but suggests that such approaches, by defining a mean atlas across states, may mask meaningful and informative functional reconfigurations associated with specific brain states.

### 4.4 Implications

Essentially all publications to date that employ parcellations with hard boundaries have used fixed atlases at the group or individual level, with parcel boundaries defined anatomically or functionally (or through a combination of methods). Such work assumes that parcel boundaries do not change meaningfully for a given individual or group, but the findings here suggest that this is not the case for atlases defined using functional MRI data. While functional parcel reconfigurations would not affect results generated using data collected during a single state, they may affect investigations of connectivity changes across states. More generally, it suggests that a single functional parcellation applicable to all brain states is not appropriate. Further, these findings suggest that state-dependent changes in connectivity could be attributable in part to reconfiguration of the underlying parcellation, and this should be considered as part of the interpretation of changes in connectivity that occur across different cognitive conditions.

It has been posited that imposing an atlas on the brain is simply a data-reduction strategy to reduce the connectivity matrix to a manageable size. This may indeed be a useful feature (for a discussion, see Eickhoff et al. (Eickhoff et al., 2018)), and this work does not impact its utility. The findings here, however, support the notion that function-based parcellations are not only a means of data dimensionality reduction but also a way to reveal meaningful patterns of, and changes in, brain functional organization. The results demonstrating prediction of state and performance based on the particular parcel size support this conclusion.

New approaches to individualized parcellation, such as the approach used in this work, could potentially lead to custom state-dependent atlases for individuals or groups, which may in turn provide further insight into understanding state-dependent functional reorganization of the human brain. Much work needs to be done to understand the relationship between functional edge strength measures and these variable parcel configurations, and the impact of these variables on functional connectivity measures in both health and disease. It is an open question whether fixed functional subunits, invariant to brain state, can be defined in the human cortex with current state-of-the-art neuroimaging methods. We do know from these results however, that such an atlas would need to have more than 5000 parcels. Nevertheless, state-dependent changes in parcellation boundaries offer an important tool to study neural representation and dynamic interactions in the human brain.

## 5. Conclusion

This work demonstrates that there is no single functional atlas for the human brain at the 200 – 5000 parcel resolution level, but rather that parcels reconfigure depending on task-induced state in a robust and reliable manner. Such reconfigurations are distinct and reliable enough to use quantitative parcellation characteristics (i.e., parcel size) to predict task condition across multiple tasks, as well as within-task performance. That functional parcel definitions are fluid must be considered when interpreting changes in functional connectivity patterns across states. Such parcel reconfigurations may be leveraged to better understand dynamic changes in functional organization of the human brain. These results therefore provide another mechanism by which to understand the flexible brain and how it reconfigures to perform the task at hand. The derivation of state-specific, individualized functional atlases could provide an important tool for human neuroscience, and this work calls for the continued development and validation of such approaches.

## 6. Code and Data Availability

All the functional parcellations are available online on the BioImage Suite NITRC page (https://www.nitrc.org/frs/?group_id=51). An interactive visualization of the state-specific parcellations can be found here: http://htmlpreview.github.io/?https://github.com/YaleMRRC/Node-Parcellation/blob/master/Parcellation_visualization.html. MATLAB, C++ (for parcellation algorithm), and Python (for predictive modeling) scripts were written to perform the analyses described; these codes are available on GitHub at https://github.com/YaleMRRC/Node-Parcellation.git. The force-directed graph visualization (R script) is released separately under the terms of GNU General Public License and can be found here: https://github.com/YaleMRRC/Network-Visualization.git. Three data sets were used to support the findings of this study. Yale data are publicly available at International Neuroimaging Data-sharing Initiative (INDI) [http://fcon_1000.projects.nitrc.org/]. Midnight Scan Club (MSC) data are publicly available through the Open fMRI data repository at https://openneuro.org/datasets/ds000224/versions/00002. Human Connectome Project (HCP) data (S900) are publicly available at https://www.humanconnectome.org/study/hcp-young-adult/document/900-subjects-data-release.

## 7. Supplemental Materials

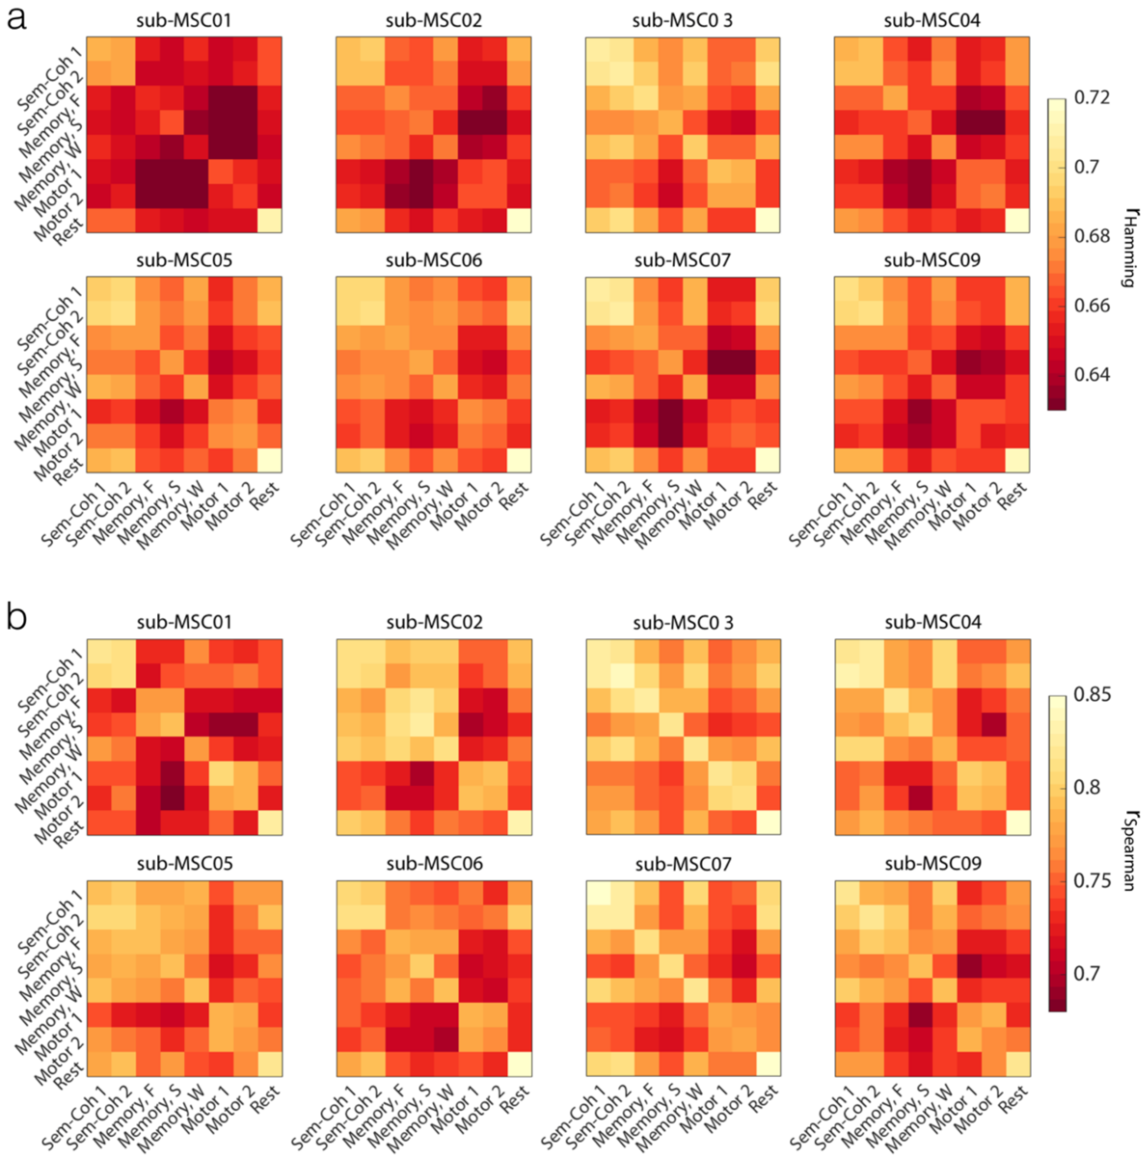

**Figure S1. Individual level replication of the finding that parcel definitions change with task condition; Midnight Scan Club (MSC) data.** Pairwise parcellation similarity was calculated for each individual, within and across functional conditions. For every individual, voting-based ensemble analysis was used with 100 iterations. The matrix represents the average over all iterations. Similarity was assessed by $r_{Hamming} = 1 -$ normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities.

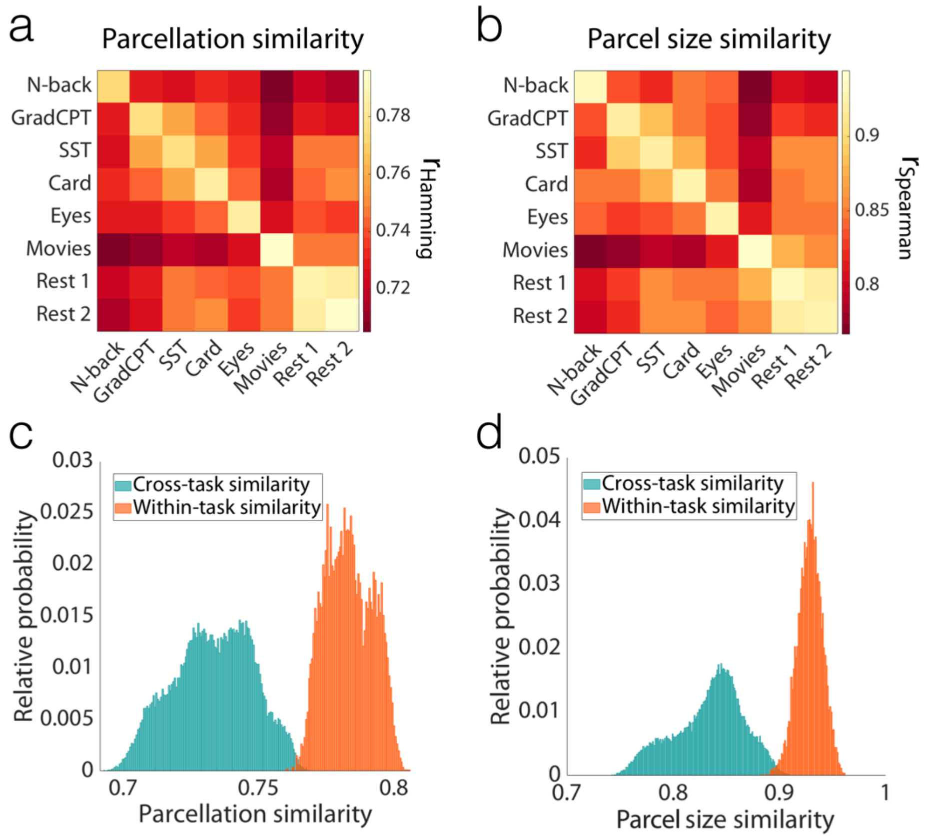

**Figure S2. Replication of the finding that parcel definitions change with task condition, using parcellation of size 368 parcels; Yale data.** a) Pairwise parcellation similarity was calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by $r_{\text{Hamming}} = 1 -$ normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{\text{Spearman}}$) as a proxy for parcellation similarities. c) The histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). d) The histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$).

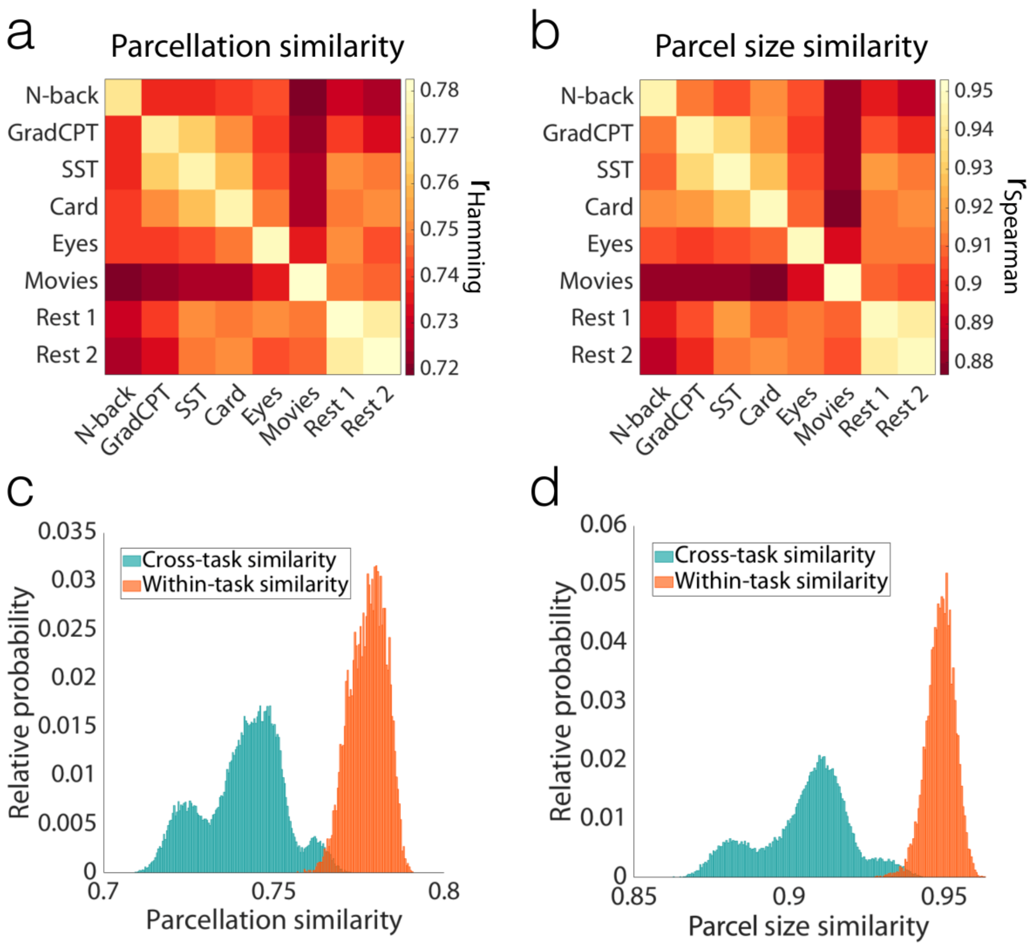

**Figure S3. Replication of the finding that parcel definitions change with task condition, using parcellation of size 1041 parcels; Yale data.** a) Pairwise parcellation similarity was calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by $r_{Hamming}$ = 1 – normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities. c) The histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). d) The histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$).

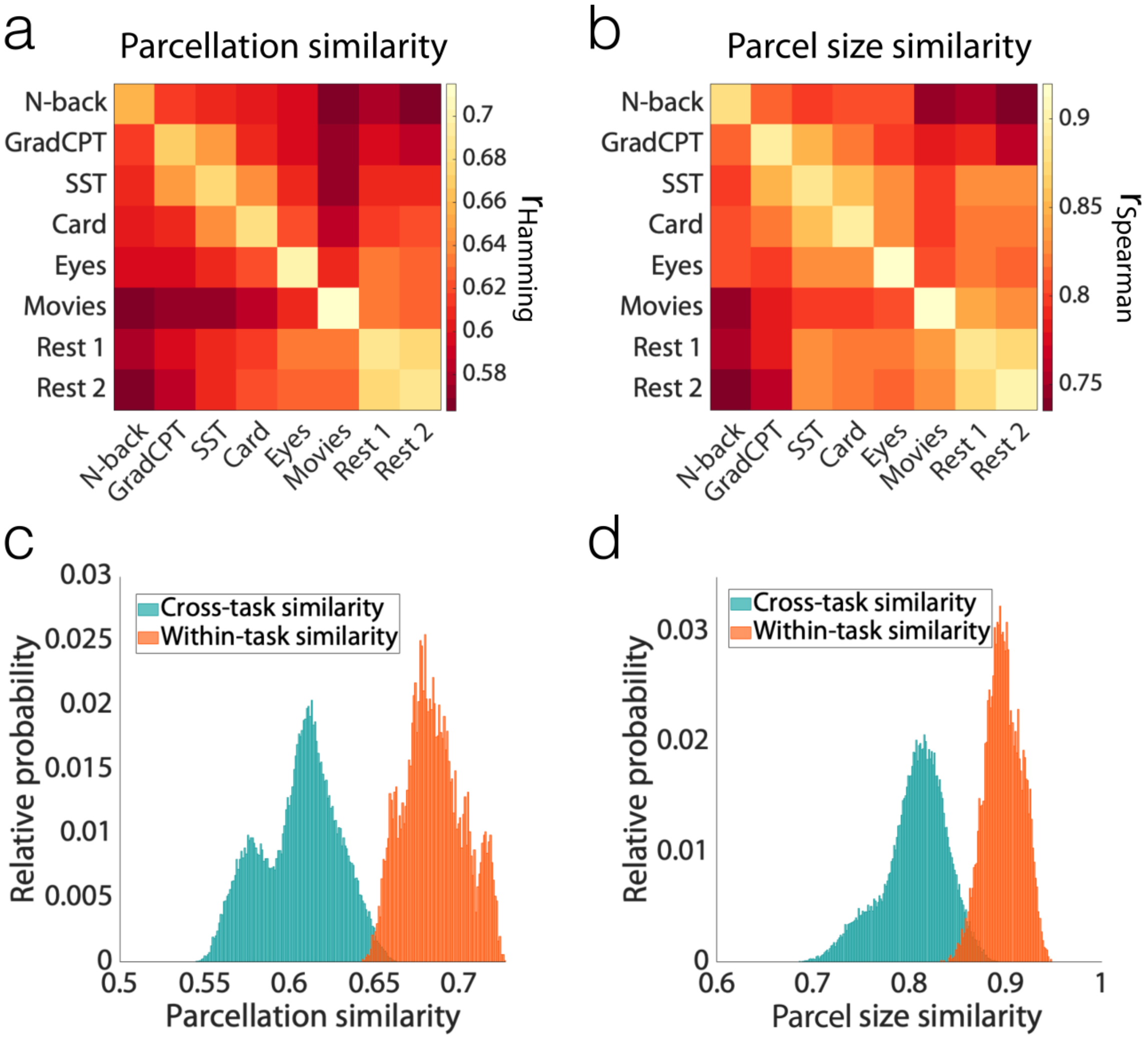

**Figure S4. Replication of the finding that parcel definitions change with task condition, even when exemplar for each parcel is fixed across sessions and conditions; Yale data.** Indication of robustness of the results to the choice of exemplar, such that even with identical exemplars across sessions and conditions, parcels reconfigure reliably based on the brain state. a) Pairwise parcellation similarity was calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by $r_{Hamming}$ = 1 – normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities. c) The histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). d) The histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$).

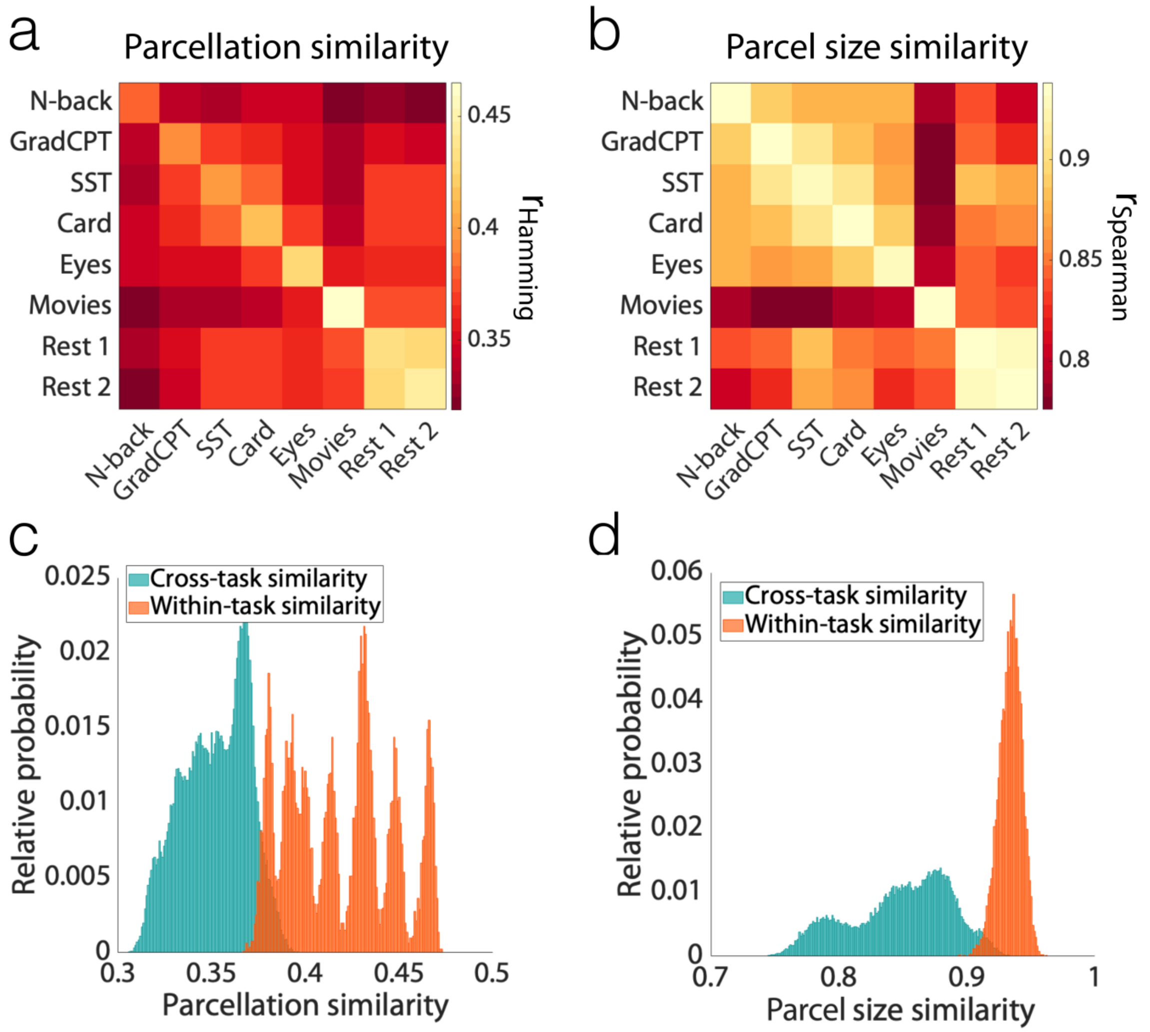

**Figure S5. Replication of the finding that parcel definitions change with task condition, even with a different parcellation algorithm: Wang's iterative parcellation; Yale data.** a) Pairwise parcellation similarity was calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by $r_{Hamming}$ = 1 – normalized Hamming distance. b) The same analysis as (a) was performed, this time using rank correlation of parcel-size vectors ($r_{Spearman}$) as a proxy for parcellation similarities. c) The histogram of the parcellation similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). d) The histogram of the parcel size similarities for all 1000 iterations is depicted for within-condition (diagonal elements in [b]) and cross-condition (off-diagonal elements in [b]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$).

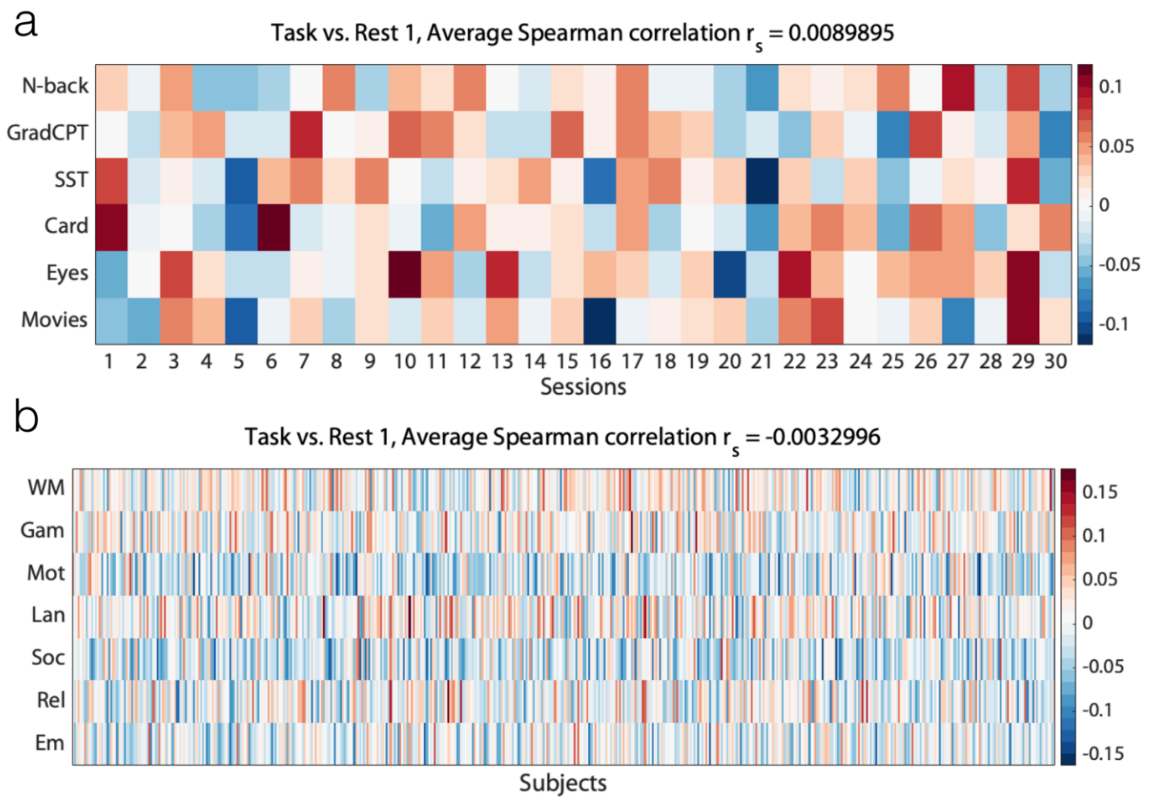

**Figure S6. Correlation between the task activation maps and parcel size differences relative to rest**. The Spearman correlation between the task activation per parcel and the difference in the parcel size (from every task run to the first rest run [Rest 1]) is displayed. (a) Yale data; task activation was approximated for each parcel as the difference between the average temporal signal during each task run and the first rest run (Rest 1). (b) HCP data; task effect size maps were generated using all available individuals' volume-based, FEAT-analyzed, first-level GLM output (COPE files) from the 1200 Subjects Release (S1200) for a given task to generate, using FSL FEAT's FLAME (FMRIB's Local Analysis of Mixed Effects), cross-subject, voxel-wise Cohen's d effect-size contrast maps from t-statistic maps. These values were then averaged for each parcel to yield a mean task effect size per parcel for each task.

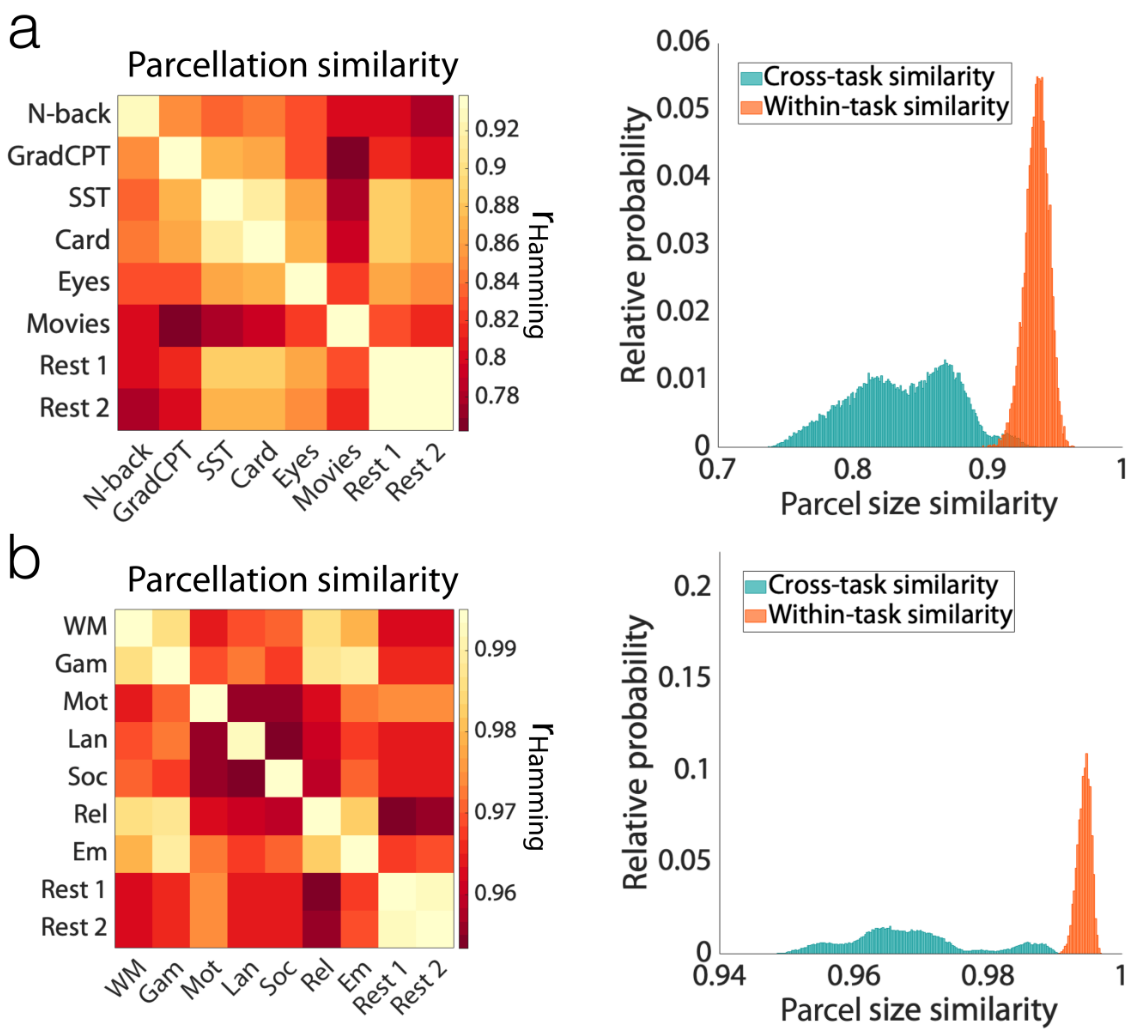

**Figure S7. Replication of the finding that parcel definitions change with task condition, even after eliminating parcels with significant task activation, for (a) Yale data, and (b) HCP data.** Left panels display the pairwise parcellation similarity, calculated within and across functional conditions, using voting-based ensemble analysis with 1000 iterations. Similarity was assessed by rank correlation of parcel-size. Right panels display the histogram of the parcel size similarities for all 1000 iterations for within-condition (diagonal elements in [a]) and cross-condition (off-diagonal elements in [a]) comparisons. The two distributions are significantly different (K-S test; $p<0.001$). a) For Yale data, task activations were approximated per parcel by computing the difference between the average temporal signal during task and rest. b) For HCP data, task effect size maps were generated using all available individuals' volume-based, FEAT-analyzed, first-level GLM output (COPE files) from the 1200 Subjects Release (S1200) for a given task to generate, using FSL FEAT's FLAME (FMRIB's Local Analysis of Mixed Effects), cross-subject, voxel-wise Cohen's d effect-size contrast maps from t-statistic maps. These values were then averaged for each parcel to yield a mean task effect size per parcel for each task.

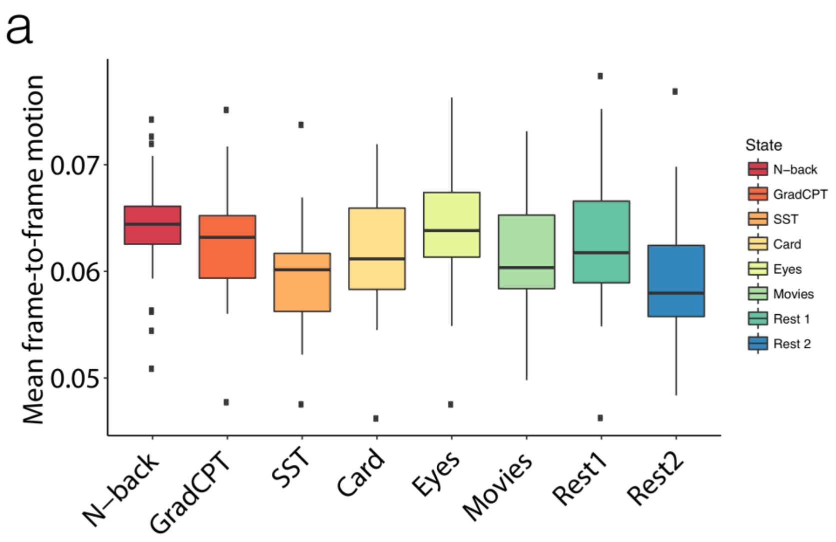

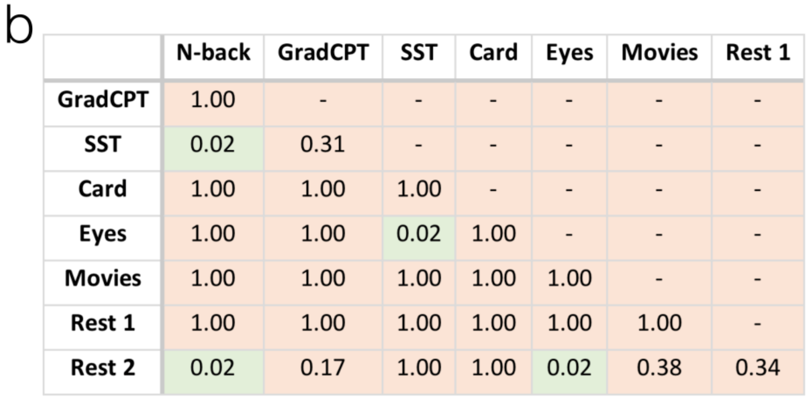

| | N-back | GradCPT | SST | Card | Eyes | Movies | Rest 1 |
|---|---|---|---|---|---|---|---|
| **GradCPT** | 1.00 | - | - | - | - | - | - |
| **SST** | 0.02 | 0.31 | - | - | - | - | - |
| **Card** | 1.00 | 1.00 | 1.00 | - | - | - | - |
| **Eyes** | 1.00 | 1.00 | 0.02 | 1.00 | - | - | - |
| **Movies** | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | - | - |
| **Rest 1** | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | - |
| **Rest 2** | 0.02 | 0.17 | 1.00 | 1.00 | 0.02 | 0.38 | 0.34 |

**Figure S8. Statistical comparison of head motion across functional conditions; Yale data.** a) Mean frame-to-frame displacement is computed for every functional condition, displayed as a box plot over sessions, where the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually. b) A pairwise Wilcoxon signed-rank test was performed on the mean frame-to-frame displacement. P-values are reported after Bonferroni correction for multiple comparisons.

# Chapter 8: Conclusion and Future work

Here, I presented a data-driven, submodularity-based methodology to human brain data summarization, which leads to increased insight about brain function and its link to behavior. All the projects described here are united in their emphasis on individual and state variability in in human brain functional organization. I hope the presented findings inspire future research on developing individualized and state-dependent approaches to understanding the brain. Below, I offer an overall summary of the main findings, and how they may help shape the future direction of human brain research.

## 1. Summary

In chapter 2, I discussed a novel data-driven, predictive approach to characterizing human brain attention, which is based on higher-order connectivity measure. Results show that a model based on the flexibility of the connections during sessions of rest and task can predict individual differences in attentional ability, and generalize across novel individuals, scanning sites, attentional tasks, and behavioral measures. Further investigation of the parameters of the model revealed that individuals with higher stability during task sessions and higher flexibility during resting-sates had higher attentional ability. The extracted information from the state differences and individual variability, motivated the study presented in the subsequent chapters.

In particular, while in chapter 2 we started form global atlas definitions, the substantial variability of connectivity measures raised two questions: 1) Whether the underling atlas varies across individuals in a functionally relevant manner, and 2) whether the functional atlas should be defined based on the underlying state or whether there is a single functional atlas that operates across all brain states. Chapter 3 and 4 attempted to answer the first and the second questions (respectively) at the network level, while Chapter 5, 6, and 7 extended it to the node level atlases.

In chapter 3, I developed a parcellation algorithm for identifying individualized functional networks, and showed that the individualized differences in network definitions are substantial and systematic such that they reveal information about subgroups in the populations, here differences in individuals' sex. The information extracted by the proposed parcellation approach is reliable and substantial enough to predict novel individual's sex, suggesting that (i) network organization may differ across individuals and global configurations should not be assumed, and (b) the proposed approach can extract meaningful signals from the noisy individual variabilities. If a basic characteristic such as sex in a cohort of healthy subjects can be identified by network topography, it is reasonable to expect other characteristics such as cognition and psychiatric symptoms could also lead to individualized networks. Future work can build upon this method to investigate the link between differences in networks to differences in behavioral phenotypes or clinical symptoms.

Chapter 4 extends this analysis to extract state-evoked differences in brain networks. I showed that even within an individual there is state-evoked differences in functional network organization, considering which may lead to increased insight about the underlying cognitive processes. The observed difference, while subtle, are consistent across subjects, such that a simple vector of node-to-network assignment could significantly predict the underlying cognitive state of a novel individual brain.

While the two previous chapters revealed promising results about the network organization of the brain, we ultimately wanted to move to node-level analysis to summarize the data from the ground up by grouping the voxels into functionally homogeneous nodes based on the synchrony of the time series. Chapter 5 proposes a novel submodular-based algorithm to individualized functional parcellation, starting from a generic group-level atlas. The proposed individualized method showed increased improvements in functional coherence and homogeneity of nodes. Further, functional connectivity matrices constructed from the individualized atlas showed increased predictive power as compared with the ones built on the group-level atlas. This raises the question that to what extent the individual differences in functional nodes are reliable and driven by individual's idiosyncrasy rather than noise.

In chapter 6, I answer this question demonstrating that an individual's functional node atlas is unique and distinct enough from that of other individuals to identify him or her from the group. While nodes from the entire brain contributed to the identification power of the model, nodes from dorsal attention network and frontoparietal network emerged as the most distinctive, suggesting the potential relevance of the individual differences in node configuration to cognition. These findings suggest that individualized functional atlases are reliable enough to be integrated into the current connectivity pipelines which attempt to link individual differences in connectivity to differences in behavior. This work could potentially inspire future research on personalized approaches to human brain parcellation. The uniqueness of individualized atlases (shown in Chapter 6), together with its utility in increasing the predictive power of the connectome-based predictive models (shown in chapter 5), suggest that the individualized atlases are functionally meaningful, and considering them instead of the generic group-averaged maps could enhance our understanding of the brain-behavior link and inspire future efforts in neuromarker discovery.

Unlike fingerprints which remain essentially unchanged throughout lifetime, brain fingerprints dynamically change with the brain state. In Chapter 7, I show that as individuals engage in different cognitive challenges, their brain's functional nodes reconfigure accordingly. I demonstrate that such reconfigurations are robust and reliable both within and across individuals, such that a measure of spatial topography as coarse as node size can predict the underlying cognitive state, approximated by the functional task, or the level of engagement in the task. This finding, that the brain functional atlas reorganizes with task, is different from what is traditionally assumed in the field. How to best characterize the human brain's underlying functional atlas has been a fundamental question in neuroscience. There have been numerous efforts to advance our parcellation algorithms and extend our data limits to better approximate the underlying atlas, assuming that there is one, at least for a single individual. Yet, here, grounded on rigorous methodology and empirical evidence, we demonstrate that the functional boundaries of the brain change reliably by task.

The work presented here builds the foundation for future research into individualized and state-specific approaches to human brain parcellation, both at the network- and node-level. Establishing that brain maps are unique to each individual and display reliable state-evoked reconfigurations is an indication of their behavioral relevance, and their potential utility for personalized and data-driven approaches to psychiatric disorders.

## 2. Looking ahead: Future work and potential implications

Human brain appears to be governed by two fundamental principles: functional specialization, and functional integration (within and across these specialized regions). Much research has been devoted to understanding each principle, and how their balance can give rise to a wide spectrum of cognitive processes and behavior.

The functional specialization of the brain has received supports from extensive evidence (Kanwisher 2010; Pauli, et al. 2016; Yeo, et al. 2014). Traditionally, these function-location mappings were conducted through studies of patients with brain damage (e.g., Broca's area (Dronkers, et al. 2007) and Wernicke's area (Wernicke 1995)). More recently, studies based on task-based neuroimaging techniques have identified process-specific brain regions associated with different subsystems of cognitive processes. For example, in the attention community, the association of alerting with fronto-parietal cortex and thalamus, orienting with superior parietal lobe, and executive control with anterior cingulate and frontal areas was uncovered (Fan, et al. 2005). While these centralized function-location mappings provide a systematic approach to study human brain function, there is significant evidence that features of higher-order cognitive abilities such as attention and working memory are likely encoded in distributed neural systems involving networks of many regions (Finn, et al. 2015c). Accordingly, restricting our measurements to a specific region of the brain, activated during a task paradigm, may not provide a holistic view of the complex cognitive mechanisms ranging from attention and memory to language and reasoning (Sternberg and Sternberg 2016).

Neuroimaging-based functional parcellation approaches advocate a more holistic view as they allow parcellation of whole brain (Blumensath, et al. 2013; Chong, et al. 2017; Fan, et al. 2016; Glasser, et al. 2016; Gordon, et al. 2014a; Power, et al. 2011; Salehi, et al. 2017a; Shen, et al. 2013; Smith, et al. 2013a; Thomas Yeo, et al. 2011a; Van Essen 2013). The present work falls in this category, and contributes to the rich body of literature on whole-brain functional parcellation, by extending these approaches to individualized and state-specific settings.

The robustness and reliability of the functional parcellation reconfigurations is promising. That these reconfigurations contain information both about individuals and the underlying cognitive state is an important finding, calling for more personalized and state-dependent approaches to human brain mapping. However, the cause of such reconfigurations remains an important open question. One hypothesis is that the functional boundaries are changing to reflect the synchronization of the underlying subset of neurons in adaptation with the cognitive demands of the ongoing task. While this hypothesis could not be tested with current state-of-the-art fMRI acquisition resolutions (Uğurbil, et al. 2013), it presents an exciting opportunity for future work with high-field fMRI to interrogate the underpinnings of these boundary shifts.

Animal studies can form an ideal candidate for testing such hypotheses as they provide substantially higher spatial resolution (at the scale of micrometers) and enable the use of invasive techniques (such as electrophysiological recordings and Calcium ($Ca^{2+}$) imaging), integration of which could result in a more direct measurement of the neural activity with higher precision and specificity. We have recently established a novel methodology for simultaneous mesoscopic $Ca^{2+}$ imaging and fMRI in mice (Lake, et al. 2019). $Ca^{2+}$ imaging provides a powerful tool for measuring the neural activity at high spatiotemporal resolution with selective specificity to each cell-type. A multi-modal parcellation based on this dual-imaging data can offer new insight into mesoscopic reconfiguration of the functional nodes with task. Also, favored by the cell-type specificity of $Ca^{2+}$ imaging, future work can address the contribution of each specific cell-type (e.g., pyramidal cells (PYR), parvalbumin (PV), somatostatin (SOM), vasoactive intestinal peptide (VIP)) to the underlying node reconfigurations.

Another important future direction includes extending the current methodology to clinical populations. The current work can contribute to this central area of research in several ways: First, it is critical to investigate the brain map differences between healthy controls and patients. Previous work has addressed this question at the population level, highlighting the coarse differences in functional brain maps between the healthy controls and patients with Alzheimer's Disease (Seeley, et al. 2007), schizophrenia (Jafri, et al. 2008), psychosis (Schreiner, et al. 2017), autism (Nunes, et al. 2018), and major depressive disorder (Liu, et al. 2012). Yet, the substantial heterogeneity within each group, and its link to differences in symptomology and level of progression into the disease remains an important question for future investigations. In a recent work, we discovered different sub groups of depression can be predicted based on the individuals' connectivity profiles (Barron, et al. 2018). Future work can investigate the existence of similar subgroups based on functional parcellation patterns. Furthermore, employing a group-level parcellation based on a cohort of healthy individuals may not be applicable for a population of patients with brain lesions (e.g., brain tumors). Individualized approaches, like the work presented here, would facilitate the application of parcellation across a wider range of clinical populations. By enhancing the subject and state specificity of the brain maps, the proposed parcellation could empower the study of idiosyncratic disorders and complex neuropsychiatric conditions with high variability in individuals' symptomology and disease trajectory.

That the current summarization technique predicts a novel individual's cognitive state from a pool of 8 tasks is promising, indicating the power of the model to extract relevant information. Yet, the current results are far from real-time decoding of the entire cognitive space. Future work can extend this methodology towards settings closer to real-time thought inferences by taking advantage of streaming summarization algorithms (e.g., (Badanidiyuru, et al. 2014; Kazemi, et al. 2019)), and a more comprehensive task ontology (Poldrack, et al. 2011). Although a monolithic leap, if successful, it can have far-reaching scientific, clinical, and social implications, including the advancement of novel brain-controlled prosthetics, human-robot interactions, and the study of ephemeral phenomena, such as dreaming and meditation.

Finally, the proposed state-specific individualized method can increase the information efficiency for the subsequent models built on these atlases. Recently, with advances in machine-learning techniques and increased accessibility of neuroimaging data sets, there has been increasing interest in developing predictive models that relate the human brain to behavioral phenotypes (Scheinost, et al. 2019; Yarkoni and Westfall 2017). Such approaches generally rely on the functional connectomes built on fixed population-level atlases. Findings presented here suggests that one single functional atlas may not be applicable across all individuals and states. Such population-based models could reduce specificity of the results, potentially masking real effects. Incorporating the individualized approaches from the beginning of the predictive pipeline could increase the amount of information related to each individual's behavior.

## 3. Data and Code Availability

Following upon the publication of these papers, I released our method for individualized and state-specific functional parcellation as part of our open-source software, BioImage Suite (https://bioimagesuiteweb.github.io/webapp/). All codes are provided online on our group's GitHub page at https://github.com/YaleMRRC/. It is intended that the proposed algorithm help other investigators to implement customized atlases for their functional connectivity analysis, based on each individual's underlying functional state. The proposed technology is attempted to facilitate usage of submodular based algorithms for a community of researchers across disciplines. It is made modular and simple to execute and requires little or zero machine learning and programming backgrounds, intended to broaden the range of audience who can take advantage of the proposed methodology.

Together, the present work highlights the need and potential for individualized and state-specific parcellations in current fMRI pipelines. I hope this work motivates future studies to incorporate customized atlases into their analysis, which in turn, increases the efficacy of existing models, and increases the predictive power of the connectivity-based models.

# Bibliography

Arslan, S., Rueckert, D., 2015. Multi-Level Parcellation of the Cerebral Cortex Using Resting-State fMRI. International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 47-54.

Badanidiyuru, A., Mirzasoleiman, B., Karbasi, A., Krause, A., 2014. Streaming submodular maximization: Massive data summarization on the fly. Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 671-680.

Barron, D.S., Salehi, M., Browning, M., Harmer, C.J., Constable, R.T., Duff, E., 2018. Exploring the prediction of emotional valence and pharmacologic effect across fMRI studies of antidepressants. NeuroImage: Clinical 20, 407-414.

Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philosophical Transactions of the Royal Society of London B: Biological Sciences 360, 1001-1013.

Bijsterbosch, J.D., Woolrich, M.W., Glasser, M.F., Robinson, E.C., Beckmann, C.F., Van Essen, D.C., Harrison, S.J., Smith, S.M., 2018. The relationship between spatial configuration and functional connectivity of brain regions. Elife 7, e32992.

Blumensath, T., Jbabdi, S., Glasser, M.F., Van Essen, D.C., Ugurbil, K., Behrens, T.E., Smith, S.M., 2013. Spatially constrained hierarchical parcellation of the brain with resting-state fMRI. Neuroimage 76, 313-324.

Braun, U., Schäfer, A., Walter, H., Erk, S., Romanczuk-Seiferth, N., Haddad, L., Schweiger, J.I., Grimm, O., Heinz, A., Tost, H., 2015. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. Proceedings of the National Academy of Sciences 112, 11678-11683.

Brodmann, K., 1909. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Barth.

Bush, G., 2011. Cingulate, frontal, and parietal cortical dysfunction in attention-deficit/hyperactivity disorder. Biological psychiatry 69, 1160-1167.

Chen, H., Li, K., Zhu, D., Jiang, X., Yuan, Y., Lv, P., Zhang, T., Guo, L., Shen, D., Liu, T., 2013. Inferring group-wise consistent multimodal brain networks via multi-view spectral clustering. IEEE transactions on medical imaging 32, 1576-1586.

Chong, M., Bhushan, C., Joshi, A., Choi, S., Haldar, J., Shattuck, D., Spreng, R., Leahy, R., 2017. Individual parcellation of resting fMRI with a group functional connectivity prior. Neuroimage 156, 87-100.

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, 238-251.

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, 1914-1928.

Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S.M., Beckmann, C., 2006. Consistent resting-state networks across healthy subjects. Proceedings of the National Academy of Sciences 103, 13848-13853.

Das, A., Dasgupta, A., Kumar, R., 2012. Selecting diverse features via spectral regularization. Advances in neural information processing systems, pp. 1583-1591.

Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. Neuroimage 31, 968-980.

Dosenbach, N.U., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2007. Distinct brain networks for adaptive and stable task control in humans. Proceedings of the National Academy of Sciences 104, 11073-11078.

Dronkers, N.F., Plaisant, O., Iba-Zizen, M.T., Cabanis, E.A., 2007. Paul Broca's historic cases: high resolution MR imaging of the brains of Leborgne and Lelong. Brain 130, 1432-1441.

Du, Y., Pearlson, G.D., Yu, Q., He, H., Lin, D., Sui, J., Wu, L., Calhoun, V.D., 2016. Interaction among subsystems within default mode network diminished in schizophrenia patients: a dynamic connectivity approach. Schizophrenia research 170, 55-65.

Fan, J., McCandliss, B.D., Fossella, J., Flombaum, J.I., Posner, M.I., 2005. The activation of attentional networks. Neuroimage 26, 471-479.

Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., 2016. The human brainnetome atlas: a new brain atlas based on connectional architecture. Cerebral cortex 26, 3508-3526.

Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. Nature neuroscience 18, 1664-1671.

Flandin, G., Kherif, F., Pennec, X., Malandain, G., Ayache, N., Poline, J.-B., 2002. Improved detection sensitivity in functional MRI data using a brain parcelling technique. International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 467-474.

Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2006. Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. Proceedings of the National Academy of Sciences 103, 10046-10051.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., 2016. A multi-modal parcellation of human cerebral cortex. Nature 536, 171.

Gonzalez-Castillo, J., Hoy, C.W., Handwerker, D.A., Robinson, M.E., Buchanan, L.C., Saad, Z.S., Bandettini, P.A., 2015. Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. Proceedings of the National Academy of Sciences 112, 8762-8767.

Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2014. Generation and evaluation of a cortical area parcellation from resting-state correlations. Cerebral cortex 26, 288-303.

Henssen, A., Zilles, K., Palomero-Gallagher, N., Schleicher, A., Mohlberg, H., Gerboga, F., Eickhoff, S.B., Bludau, S., Amunts, K., 2016. Cytoarchitecture and probability maps of the human medial orbitofrontal cortex. cortex 75, 87-112.

Jafri, M.J., Pearlson, G.D., Stevens, M., Calhoun, V.D., 2008. A method for functional network connectivity among spatially independent resting-state components in schizophrenia. Neuroimage 39, 1666-1681.

Kahnt, T., Chang, L.J., Park, S.Q., Heinzle, J., Haynes, J.-D., 2012. Connectivity-based parcellation of the human orbitofrontal cortex. The Journal of neuroscience 32, 6240-6250.

Kanwisher, N., 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. Proceedings of the National Academy of Sciences 107, 11163-11170.

Kazemi, E., Mitrovic, M., Zadimoghaddam, M., Lattanzi, S., Karbasi, A., 2019. Submodular Streaming in All Its Glory: Tight Approximation, Minimum Memory and Low Adaptive Complexity. In: Kamalika, C., Ruslan, S. (Eds.), Proceedings of the 36th International Conference on Machine Learning. PMLR, Proceedings of Machine Learning Research, pp. 3311--3320.

Kempe, D., Kleinberg, J., Tardos, É., 2003. Maximizing the spread of influence through a social network. Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 137-146.

Krause, A., Guestrin, C.E., 2012. Near-optimal nonmyopic value of information in graphical models. arXiv preprint arXiv:1207.1394.

Laird, A.R., Fox, P.M., Eickhoff, S.B., Turner, J.A., Ray, K.L., McKay, D.R., Glahn, D.C., Beckmann, C.F., Smith, S.M., Fox, P.T., 2011. Behavioral interpretations of intrinsic connectivity networks. Journal of cognitive neuroscience 23, 4022-4037.

Lake, E.M., Ge, X., Shen, X., Herman, P., Hyder, F., Cardin, J.A., Higley, M.J., Scheinost, D., Papademetris, X., Crair, M.C., 2019. Simultaneous mesoscopic Ca2+ imaging and fMRI: Neuroimaging spanning spatiotemporal scales. bioRxiv, 464305.

Lee, M.H., Hacker, C.D., Snyder, A.Z., Corbetta, M., Zhang, D., Leuthardt, E.C., Shimony, J.S., 2012. Clustering of resting state networks. PloS one 7, e40370.

Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., Faloutsos, C., VanBriesen, J., Glance, N., 2007. Cost-effective outbreak detection in networks. Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 420-429.

Lin, H., Bilmes, J., 2011. A class of submodular functions for document summarization. Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1. Association for Computational Linguistics, pp. 510-520.

Lin, Q., Rosenberg, M.D., Yoo, K., Hsu, T.W., O'Connell, T.P., Chun, M.M., 2018. Resting-state functional connectivity predicts cognitive impairment related to alzheimer's disease. Frontiers in aging neuroscience 10, 94.

Liu, L., Zeng, L.-L., Li, Y., Ma, Q., Li, B., Shen, H., Hu, D., 2012. Altered cerebellar functional connectivity with intrinsic connectivity networks in adults with major depressive disorder. PloS one 7, e39516.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature 412, 150.

Mennes, M., Kelly, C., Colcombe, S., Castellanos, F.X., Milham, M.P., 2013. The extrinsic and intrinsic functional architectures of the human brain are not equivalent. Cerebral cortex 23, 223-229.

Meunier, D., Lambiotte, R., Fornito, A., Ersche, K., Bullmore, E.T., 2009. Hierarchical modularity in human brain functional networks. Frontiers in neuroinformatics 3, 37.

Meunier, D., Lambiotte, R., Fornito, A., Ersche, K.D., Bullmore, E.T., 2010. Hierarchical modularity in human brain functional networks. Hierarchy and dynamics in neural networks 1, 2.

Moreno-Dominguez, D., Anwander, A., Knösche, T.R., 2014. A hierarchical method for whole-brain connectivity-based parcellation. Human brain mapping 35, 5000-5025.

Niendam, T.A., Laird, A.R., Ray, K.L., Dean, Y.M., Glahn, D.C., Carter, C.S., 2012. Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. Cognitive, Affective, & Behavioral Neuroscience 12, 241-268.

Nieuwenhuys, R., Broere, C.A., Cerliani, L., 2015. A new myeloarchitectonic map of the human neocortex based on data from the Vogt–Vogt school. Brain Structure and Function 220, 2551-2573.

Nunes, A.S., Peatfield, N., Vakorin, V., Doesburg, S.M., 2018. Idiosyncratic organization of cortical networks in autism spectrum disorder. Neuroimage.

OpenStax, Brodmann's cytotechtonic map. By OpenStax - https://cnx.org/contents/FPtK1zmh@8.25:fEI3C8Ot@10/Preface, CC BY 4.0, https://commons.wikimedia.org/w/index.php?curid=30147951.

Pauli, W.M., O'Reilly, R.C., Yarkoni, T., Wager, T.D., 2016. Regional specialization within the human striatum for diverse psychological functions. Proceedings of the National Academy of Sciences 113, 1907-1912.

Poldrack, R.A., Kittur, A., Kalar, D., Miller, E., Seppa, C., Gil, Y., Parker, D.S., Sabb, F.W., Bilder, R.M., 2011. The cognitive atlas: toward a knowledge foundation for cognitive neuroscience. Frontiers in neuroinformatics 5, 17.

Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., Petersen, S.E., 2011. Functional network organization of the human brain. Neuron 72, 665-678.

Price, T., Wee, C.-Y., Gao, W., Shen, D., 2014. Multiple-network classification of childhood autism using functional connectivity dynamics. International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 177-184.

Rosenberg, M.D., Finn, E.S., Scheinost, D., Papademetris, X., Shen, X., Constable, R.T., Chun, M.M., 2016. A neuromarker of sustained attention from whole-brain functional connectivity. Nat Neurosci 19, 165-171.

Salehi, M., Greene, A.S., Karbasi, A., Shen, X., Scheinost, D., Constable, R.T., 2018a. There is no single functional atlas even for a single individual: Parcellation of the human brain is state dependent. bioRxiv, 431833.

Salehi, M., Karbasi, A., Barron, D.S., Scheinost, D., Constable, R.T., 2018b. State-specific individualized functional networks form a predictive signature of brain state. bioRxiv, 372110.

Salehi, M., Karbasi, A., Scheinost, D., Constable, R.T., 2017a. A submodular approach to create individualized parcellations of the human brain. International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 478-485.

Salehi, M., Karbasi, A., Shen, X., Scheinost, D., Constable, R.T., 2017b. An exemplar-based approach to individualized parcellation reveals the need for sex specific functional networks. Neuroimage, Submitted 142.

Scheinost, D., Noble, S., Horien, C., Greene, A.S., Lake, E.M., Salehi, M., Gao, S., Shen, X., O'Connor, D., Barron, D.S., 2019. Ten simple rules for predictive modeling of individual differences in neuroimaging. Neuroimage.

Schreiner, M., Forsyth, J.K., Karlsgodt, K.H., Anderson, A.E., Hirsh, N., Kushan, L., Uddin, L.Q., Mattiacio, L., Coman, I.L., Kates, W.R., 2017. Intrinsic connectivity network-based classification and detection of psychotic symptoms in youth with 22q11. 2 deletions. Cerebral cortex 27, 3294-3306.

Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., Greicius, M.D., 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. Journal of Neuroscience 27, 2349-2356.

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, 403-415.

Shine, J.M., Bissett, P.G., Bell, P.T., Koyejo, O., Balsters, J.H., Gorgolewski, K.J., Moodie, C.A., Poldrack, R.A., 2016. The dynamics of functional brain networks: integrated network states during cognitive task performance. Neuron 92, 544-554.

Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. Cerebral cortex 22, 158-165.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., 2009. Correspondence of the brain's functional architecture during activation and rest. Proceedings of the National Academy of Sciences 106, 13040-13045.

Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E.C., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., Ugurbil, K., Van Essen, D.C., 2013. Functional connectomics from resting-state fMRI. Trends Cogn Sci 17, 666-682.

Smitha, K., Akhil Raja, K., Arun, K., Rajesh, P., Thomas, B., Kapilamoorthy, T., Kesavadas, C., 2017. Resting state fMRI: A review on methods in resting state connectivity analysis and resting state networks. The neuroradiology journal 30, 305-317.

Sporns, O., 2011. The human connectome: a complex network. Annals of the New York Academy of Sciences 1224, 109-125.

Sporns, O., Betzel, R.F., 2016. Modular brain networks. Annual review of psychology 67, 613-640.

Stern, E.R., Fitzgerald, K.D., Welsh, R.C., Abelson, J.L., Taylor, S.F., 2012. Resting-state functional connectivity between fronto-parietal and default mode networks in obsessive-compulsive disorder. PloS one 7, e36356.

Sternberg, R.J., Sternberg, K., 2016. Cognitive psychology. Nelson Education.

Thirion, B., Flandin, G., Pinel, P., Roche, A., Ciuciu, P., Poline, J.B., 2006. Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fMRI datasets. Human brain mapping 27, 678-693.

Thomas Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of Neurophysiology 106, 1125-1165.

Toga, A.W., Clark, K.A., Thompson, P.M., Shattuck, D.W., Van Horn, J.D., 2012. Mapping the human connectome. Neurosurgery 71, 1-5.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273-289.

Uğurbil, K., Xu, J., Auerbach, E.J., Moeller, S., Vu, A.T., Duarte-Carvajalino, J.M., Lenglet, C., Wu, X., Schmitter, S., Van de Moortele, P.F., 2013. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. Neuroimage 80, 80-104.

Van Den Heuvel, M., Mandl, R., Pol, H.H., 2008. Normalized cut group clustering of resting-state FMRI data. PloS one 3, e2001.

Van Essen, D.C., 2013. Cartography and connectomes. Neuron 80, 775-790.

Weber, E.M.G., Hahn, T., Hilger, K., Fiebach, C.J., 2017. Distributed patterns of occipito-parietal functional connectivity predict the precision of visual working memory. Neuroimage 146, 404-418.

Wernicke, K., 1995. The aphasia symptom-complex: A psychological study on an anatomical basis (1875). Reader in the History of Aphasia: From (Franz Gall to) 4, 69-89.

Yarkoni, T., Westfall, J., 2017. Choosing prediction over explanation in psychology: Lessons from machine learning. Perspectives on Psychological Science 12, 1100-1122.

Yeo, B.T., Krienen, F.M., Eickhoff, S.B., Yaakub, S.N., Fox, P.T., Buckner, R.L., Asplund, C.L., Chee, M.W., 2014. Functional specialization and flexibility in human association cortex. Cerebral cortex 25, 3654-3672.

Zhu, X., Wang, X., Xiao, J., Liao, J., Zhong, M., Wang, W., Yao, S., 2012. Evidence of a dissociation pattern in resting-state default mode network connectivity in first-episode, treatment-naive major depression patients. Biological psychiatry 71, 611-617.

Zilles, K., Palomero-Gallagher, N., 2001. Cyto-, myelo-, and receptor architectonics of the human parietal cortex. Neuroimage 14, S8-S20.