Network variants are similar between task and rest states

Brian T. Kraus[a], **Diana Perez**[a], **Zach Ladwig**[b], **Benjamin A. Seitzman**[c], **Ally Dworetsky**[d], **Steven E. Petersen**[c,d,e,f,g], **Caterina Gratton**[a,b,h,*]

[a]Department of Psychology, Northwestern University, Evanston, IL 60208, United States

[b]Interdepartmental Neuroscience Program, Northwestern University, Chicago, IL 60611, United States

[c]Department of Neurology, Washington University in St. Louis, St. Louis, MO 63110, United States

[d]Department of Radiology, Washington University in St. Louis, St. Louis, MO 63110, United States

[e]Department of Neuroscience, Washington University in St. Louis, St. Louis, MO 63110, United States

[f]Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO 63110, United States

[g]Department of Psychological and Brain Sciences, Washington University in St. Louis, St. Louis, MO 63110, United States

[h]Department of Neurology, Northwestern University, Chicago, IL 60611, United States

Abstract

Recent work has demonstrated that individual-specific variations in functional networks (termed "network variants") can be identified in individuals using resting state functional magnetic resonance imaging (fMRI). These network variants exhibit reliability over time, suggesting that they may be trait-like markers of individual differences in brain organization. However, while networks variants are reliable at rest, is is still untested whether they are stable between task and rest states. Here, we use precision data from the Midnight Scan Club (MSC) to demonstrate that (1) task data can be used to identify network variants reliably, (2) these network variants show substantial spatial overlap with those observed in rest, although state-specific effects are present, (3) network variants assign to similar canonical functional networks in task and rest states, and (4) single tasks or a combination of multiple tasks produce similar network variants to rest. Together, these findings further reinforce the trait-like nature of network variants and demonstrate the utility of using task data to define network variants.

*Corresponding author at: Department of Psychology, Northwestern University, Evanston, IL 60208, United States. caterina.gratton@northwestern.edu (C. Gratton).

1. Introduction

An important issue in contemporary cognitive neuroscience is how to best identify individual differences in human brain organization that may be relevant to individual differences in behavior. It has long been established that measuring functional brain networks via fMRI can identify regions with correlated activity (i.e., "functional connectivity") that support common functions such as motor processing and cognitive control (Biswal et al., 1995; Dosenbach et al., 2007). This organization of human brain networks has been mapped in multiple large groups (Power et al., 2011; Yeo et al., 2011) and more recently has been extended to individual-level brain networks (Braga and Buckner, 2017; Finn et al., 2015; Gordon et al., 2017; Gratton et al., 2019; Greene et al., 2019; Kong et al., 2019; Laumann et al., 2015, 2016; Marek et al., 2018; Miranda-Dominguez et al., 2014; Mueller et al., 2013; Poldrack et al., 2015; Seitzman et al., 2019; Sylvester et al., 2020). These networks can be measured during tasks or during a resting-state, where participants are asked to lie in the scanner without performing any particular task.

While there are many commonalities in network organization across individuals, some locations show large variations (Braga and Buckner, 2017; Gordon et al., 2017 a, 2017b, 2017c; Gratton et al., 2018; Kong et al., 2019; Mueller et al., 2013; Seitzman et al., 2019). This variability across individuals is consistently larger than that found within individuals, suggesting that these individual variations represent meaningful "trait-like" differences in brain organization (Gratton et al., 2018; Seitzman et al., 2019). Individual differences in network organization are even sufficiently large to identify individuals based on their functional connectivity profile alone (Finn et al., 2015; Miranda-Dominguez et al., 2014).

A recent paper investigated these "trait-like" characteristics of individual differences in brain networks in additional detail (Seitzman et al., 2019). The authors found that particular punctate locations showed very large individual differences in functional network organization relative to the group (with correlations $r < 0.3$), that they termed "network variants". All individuals showed evidence of network variants, although these variants differed in location, size, number, and network association across individuals. Importantly, not only were network variants common, but they were also quite reliable within individuals across resting-state scans, given sufficient data.

However, one outstanding question concerning the "trait-like" nature of network variants is whether they are stable across different states, such as during task performance. In order to have utility as biomarkers, ideally network variant characteristics would not be dependent on ongoing cognition, such as what a person is thinking about during a scan session (Gratton et al., 2018). To date, network variants have only been identified using resting state fMRI data (Seitzman et al., 2019), and it is unknown whether similar variants can also be identified using task fMRI data. Recent evidence suggests that functional connectivity as a whole is stable between states with only small observable task-dependent changes (Cole et al., 2014; Gratton et al., 2018; Krienen et al., 2014), although individually-specific task-dependent changes were somewhat larger (Gratton et al., 2018). Thus, we sought to test whether these findings would also generalize to network variants.

A related practical consideration is that most existing datasets do not have sufficient resting-state fMRI data to achieve high reliability at the individual level (most have 5–10 min. of data, where reliability is poor; Elliott et al., 2019; Gordon et al., 2017; Laumann et al., 2015; Noble et al., 2017). However, many of these datasets with insufficient rest data also collect task data for other purposes. Moreover, in populations where excessive motion is an issue, it can be difficult to collect resting state fMRI whereas task fMRI may be more feasible (Greene et al., 2018; Vanderwal et al., 2019). One possibility to address these issues is to combine rest and task fMRI data together to achieve higher reliability for network variants (as has been suggested by Elliot et al. (2019) for the connectome as a whole), which seems like a promising approach. However, for this approach to be effective, it is important to understand the degree to which the specific network properties under investigation are state dependent. While Gratton et al. (2018) demonstrated that task effects on functional connectivity as a whole were modest, they were present and significant. Thus, an important question for practical identification of network variants in other datasets, especially those of clinical interest, is to what degree their locations and network properties are state-dependent and can be computed from task-based measures instead. If task data can be used to identify similar network variants as observed in resting state data, this would make datasets with large amounts of task fMRI data available for this individual-specific analysis.

To address these theoretical and practical questions, we utilized data from the Midnight Scan Club (MSC). The MSC dataset is well-suited to examine these issues as it includes data from 10 individuals across 10 different sessions with approximately 10.5 total hours of fMRI data across 4 task states and rest. With such a large amount of data for task and rest states, we can compare (1) reliability of network variants identified using task data, (2) similarity in the *locations* of network variants between task and rest states, (3) similarity of network variant *connectivity profiles* between task and rest states, and (4) whether network variants identified using individual tasks differ from each other. Jointly, these analyses suggest that network variants are largely trait-like and that task data can reasonably be used for their identification and analysis.

2. Methods

2.1. Overview

The feasibility of using task data to measure network variants was investigated using several analyses. (1) In the first analysis, we quantified the reliability (within-state consistency) of network variants identified with task data (Fig. 2). (2) Next, we examined the stability (between-state similarity) of network variant locations during task and rest (Fig. 4). (3) Third, we measured the stability (between-state) of patterns of functional connectivity for network variants (Fig. 5). (4) To examine whether network variants show greater stability than expected relative to other areas of cortex, we compared the stability of network variant locations to that of other vertices on the cortex (Fig. 6). (5) Finally, we examined whether network variants showed task-specific patterns of deviation from group functional connectivity (Fig. 7). These analyses quantify both the degree to which network variants show trait-like stability and the feasibility of using task data to identify network variants.

2.2. Datasets

The primary data used in this manuscript are from the Midnight Scan Club (MSC) dataset (Gordon et al., 2017). This dataset contains 5 h of resting state data and approximately 5.5 h of task data across 3 tasks (semantic/dot coherence, motor, and implicit memory) for 10 participants, collected across 10 separate fMRI sessions. One participant was excluded from analysis due to excessive head motion and drowsiness (Gordon et al., 2017; Laumann et al., 2016). A separate secondary dataset comprised of 120 neurotypical adults was used as the "group average" reference and is referred to as the WashU 120 (Power et al., 2013). These datasets are described in more detail elsewhere (Gordon et al., 2017; Power et al., 2013). The MSC and WashU MRI parameters are described briefly below. All data collection was approved by the Washington University Internal Review Board.

2.3. MRI parameters

2.3.1.   MSC—The acquisition parameters for the MSC dataset have been fully described elsewhere (Gordon et al., 2017). In brief high-resolution T1-weighted (224 slices, isotropic 0.8 mm^3 voxels, TE = 3.74 ms, TR = 2400 ms, TI = 1000 ms, flip angle = 8°), T2-weighted, and BOLD data (both task and rest; gradient-echo EPI sequence, isotropic 4 mm^3 voxels, TE = 27 ms, and TR = 2200 ms, whole brain acquisition) were collected on a Siemens 3T Magnetom Tim Trio with a 12-channel head coil.

2.3.2.   WashU 120—The acquisition parameters for this dataset have been fully described elsewhere. In brief, high-resolution T1-weighted (176 slices, isotropic 1 mm^3 voxels, TE = 3.06 ms, TR = 2400 ms, TI = 1000 ms, flip angle = 8°), and BOLD data (gradient-echo EPI sequence, isotropic 4 mm^3 voxels, TE = 27 ms, and TR = 2500 ms) were collected on a Siemens 3T Magnetom Tim Trio with a 12-channel head coil. This dataset contains eyes-open resting state data for each participant.

2.4. Task designs and analysis

The MSC dataset includes five different conditions. The WashU 120 contains only resting state data. The conditions in the MSC dataset are described in more detail elsewhere (Gordon et al., 2017; Gratton et al., 2018), but are briefly outlined below.

2.4.1.   Resting state—In the MSC, each session began with a 30-minute scan where participants were instructed to focus on a white fixation superimposed on a black background.

2.4.2.   Task data—The MSC dataset contained data from four different tasks presented in counterbalanced order. Briefly, the motor task included two runs in each session (15.6 min total) of a blocked motor task adapted from the Human Connectome Project (Barch et al., 2013) where participants were cued to move either hands, feet, or tongue. The semantic and coherence tasks were presented as separate counter-balanced blocks in the same run two times in each session (14.2 min total). In the semantic task, participants responded to whether each word was a verb or a noun (both presented at 50% frequency). In the coherence task, Glass patterns were presented as white dots on a black screen (Glass, 1969). These dots varied by how concentrically they were arranged, with either a 50% or 0%

coherence to a concentric arrangement (both presented at 50% frequency). The semantic and coherence tasks were analyzed together as a mixed block-event related analysis (Gordon et al., 2017) and are referred to collectively as the "semantic/coherence" task in this manuscript. The memory task included three runs in each session (~15 min total) with a different stimulus type (faces, scenes, words) presented in each run. In each run, participants were presented with 24 images, each of which was presented 3 times. Participants were asked to indicate whether the faces were male or female, the scenes were indoor or outdoor, and whether the words were abstract or concrete. Task activations were not directly analyzed in this manuscript, but task-related activity was calculated and regressed out from the timeseries via a general linear model as reported in Gratton et al. (2018).

2.5. Data and code availability

All of the data have been made publicly available (MSC: https://openneuro.org/datasets/ds000224/versions/00002; WashU 120: https://legacy.openfmri.org/dataset/ds000243/). Code for analysis related to network variants in MATLAB is available at: https://github.com/GrattonLab/SeitzmanGratton2019PNAS; other code related to MSC processing can be found at: https://github.com/MidnightScanClub. Code related specifically to the analyses in this paper will be located at https://github.com/GrattonLab/Kraus_2021_TaskRestVariants.

2.6. MRI processing

Data processing for the MSC dataset (Gordon et al., 2017) and WashU 120 (Power et al., 2013) are explained in detail at the linked references. The relevant details of the data processing are outlined below.

2.6.1.   Structural MRI processing—For both datasets, the T1-weighted images were processed via automatic segmentation of the white matter, gray matter, and ventricles in Freesurfer 5.3 (Fischl et al., 2002). The default recon-all command in Freesurfer was then applied to create the anatomical surface for each participant (Dale et al., 1999). In the MSC dataset, these surfaces were subsequently hand edited to improve the quality of the registration. The surfaces were registered to the fs_LR_32k surface space using the procedure described in Glasser et al. (2013).

2.6.2.   Functional MRI processing—The BOLD fMRI data from different runs in each session were concatenated within participants and pre-processed in the volume. First, a slice timing correction was applied, and the data were aligned to the first frame of the first run via rigid body transforms and then normalized to mode 1000 (Miezin et al., 2000). The functional data for the MSC dataset was then registered to the T2 image and subsequently to the T1 image which was previously registered to template space. In the WashU 120 dataset, the images were registered directly to the T1 image. Both datasets were then resampled to 3 mm isotropic resolution and registered to the 711–2B atlas (Smith et al., 2004). For the MSC dataset, distortion correction was also applied in this step (Gordon et al., 2017). This correction was not applied to the WashU 120 dataset as no field maps were collected.

The task fMRI data were processed in the volume via a finite impulse response general linear model (GLM) approach as described elsewhere (Gratton et al., 2018). The residuals

from the GLM were used to compute task functional connectivity using the "background connectivity" approach (Al-Aidroos et al., 2012; Fair et al., 2007). Subsequently, task residual processing was identical to the resting state data (i.e., completing all steps in the following section). As a supplemental analysis to assess whether regressing out the task activations altered functional connectivity between states, we also performed several analyses without removing task activations from the data (referred to as non-residualized data here). However, unless otherwise noted all references to task data in this paper were performed on the residuals from the task data.

2.6.3. Functional connectivity processing—In the volume, additional preprocessing steps were applied to the data to remove artifacts from the data as outlined elsewhere (Gordon et al., 2017; Power et al., 2014). These steps included (1) the removal of nuisance signals via regression from the white matter, ventricles, global signal, motion parameters, as well as their derivatives, (2) removal of frames with high motion (FD > 0.2 mm, along with sequences containing less than 5 contiguous low motion frames, the first 30 s of each run, and runs with < 50 low motion frames (Power et al., 2014), and (3) bandpass filtering of the data (0.009 Hz to 0.08 Hz). For MSC03 and MSC10, the motion parameters were low-pass filtered (below 0.1 Hz) before FD calculations to address respiratory activity in the motion traces (Fair et al., 2020; Gordon et al., 2017; Gratton et al., 2018, 2020; Laumann et al., 2016). After this preprocessing, the cortical functional data were registered to the surface. The cortical surfaces were transformed into the CIFTI data format in Connectome Workbench (Marcus et al., 2011). Lastly, a geodesic Gaussian smoothing kernel was applied (FWHM = 6 mm, sigma = 2.55) using 2-D smoothing on the cortical surface.

2.7. Network variants

2.7.1. Functional connectivity: creation of temporal correlation matrices— Next, correlation matrices were created for each participant (for task and rest data separately) based on timeseries correlations after censoring high motion frames. The amount of data (number of timepoints) used for correlations was matched in length across participants and states, approximately equally sampling across sessions (see Table S1 for more details). The only exception to equally sampling across sessions were the reliability analyses (Section 3.1), where data was sampled consecutively from the first session until the total amount of required data (100 min) was reached rather than sampling an equal amount of data from each session. This sampling method was selected for the reliability estimates to provide a more ecologically valid assessment of reliability.

Based on these equally sampled timeseries, a pairwise temporal correlation matrix was calculated for each participant between every pair of cortical surface vertices (59,412 × 59,412). The correlation matrix was then Fisher transformed. Thus, each row of this matrix represents a seed map for a given seed (vertex) to all other vertices on the surface (Fig. 1A; (Seitzman et al., 2019). We also created a similar correlation matrix for the WashU 120, but in this case, the individual subject correlation matrices were averaged together to create a single group average matrix (Fig. 1B).

2.7.2. Spatial correlations to the group average—In order to identify regions with strong dissimilarity from the typical functional connectivity profile (i.e., network variants), each MSC participant's correlation matrix was contrasted with the group average correlation matrix from the WashU 120 dataset using a spatial correlation. To compute the correlation, the seed map for each vertex (i.e., a row of the correlation matrix) in a given individual was correlated with the corresponding vertex's seed map in the WashU 120 group data, yielding one spatial correlation value per vertex (Fig. 1C). This was repeated for each individual in the MSC (Seitzman et al., 2019). Vertices with low signal (mean BOLD signal < 750 after mode 1000 normalization) were masked (Gordon et al., 2016; Ojemann et al., 1997).

2.7.3. Calculation of network variants—Next, in order to select regions that were most different from the group (network variants), the spatial correlation maps were binarized to keep only the lowest 5% of correlations to the group for each participant (values in the lowest 5% were set to 1, all other vertices were set to 0; Fig. 1D; note this threshold differs slightly from Seitzman et al. (2019) due to small changes in the processing stream of the MSC data). Variants composed of less than 50 contiguous vertices and vertices with an SNR < 750 were excluded from analysis (Seitzman et al., 2019). The surviving vertices in the lowest 5% of correlations to the group were then considered for further analysis. This procedure was used to create network variants for each individual in both task and rest data separately. For the task data, data from all 3 tasks was used unless otherwise noted (see Table S1). To confirm the findings were not dependent on threshold, the analyses reported here are also shown at a 2.5% threshold in the *Supplementary Results*. For the 2.5% threshold, variants were required to be larger than 25 contiguous vertices. In addition, we also performed several analyses using an absolute threshold ($r < 0.3$) rather than a relative threshold. For these analyses, variants of less than 50 contiguous vertices were excluded from analysis.

2.8. Within-state reliability of network variants in task data

If task data is to be considered a suitable substitute for rest in defining network variants, it must show similar within-state reliability as rest (as well as similar topography and network properties – tackled in the next sub-sections). To test how much task data is necessary to achieve reliable estimates of functional connectivity within an individual, each participant's data was divided into split-halves composed of either odd or even numbered sessions. As above, data was equally sampled across all sessions for task and rest.

Following Laumann et al. (2015) and Gordon et al. (2017), functional connectivity for each even numbered session was estimated using all of the available task data and this split-half was treated as the best estimate of "true" functional connectivity. In each odd numbered session, data was consecutively sampled starting with the first session in 5 min increments (Seitzman et al., 2019). When all of the task data in a session was exhausted, data was then sampled from the next session until either 100 min of task data was sampled, or no more task data was available (see Table S1 for details on the data sampling of each analysis). The data from these sessions was used as a test split-half.

For each functional connectivity matrix in both segments of data, a spatial correlation map versus the group average was computed following the steps outlined above. The similarity of the individual-to-group spatial correlation maps was then compared between the "true" half and progressive amounts of data from the test split-half via a Pearson correlation. To compare the spatial topograpghy of network variants, their locations were defined using the procedure outlined in Section 2.7.3 only no exclusion for variant size was applied for these analyses. These binarized spatial location maps were then compared between the "true" and test halves via a Dice-Sorenson correlation (Dice, 1945; Sorensen, 1948).

2.9. Comparing variant spatial locations

2.9.1. Comparison of variant locations—To inspect the similarity of network variant locations across task and rest states, spatial overlap maps were also created for each participant. As above, these maps were created by binarizing the vertices on the surface that were considered variants in either the task or rest states within each participant. Then, the similarity in variant spatial locations within and between states for each participant was determined via a Dice correlation. In this analysis, variants were separately defined for rest and task states in each split half resulting in 4 variant maps (35.2 min total for each split-half matrix per participant; see Table S1 - data was approximately equally sampled across sessions in task and rest).

2.9.2. Between state comparison—To compare the spatial overlap of variants between task and rest states, each split-half of rest data was separately compared to each split-half of task data. This process was repeated for both pairs of split-halves (odd rest split-half to even task split-half and even rest split-half to odd task split-half) to produce two comparisons per participant. The spatial location overlap of the split-halves between states was computed via a Dice correlation. Paired t-tests were then performed using the averaged value of the between-state comparisons within each participant.

2.9.3. Within state comparison—To estimate the spatial overlap of variants within task and rest states, the split-halves of rest and task were both separately compared (odd rest split-half to even rest split-half and odd task split-half to even task split-half) to produce two comparisons per participant. The spatial location overlap of the split-halves within states was computed via a Dice correlation. Paired t-tests were reported using the averaged values of the within each state comparisons within each participant.

2.9.4. Across subject comparison—To better evaluate the relative similarity of variant spatial locations within participants (within and between states), we created a comparison benchmark of the similarity of variant spatial locations across subjects. This benchmark described the likelihood of observing variants in the same location across subjects, which we used as a comparison for the probability of observing variants in the same location within subjects (either within or between states). For this comparison, the spatial locations for a given participant's variants (computed from a single split-half) were compared to the spatial locations for variants from all other participants (also computed from a single split half, from the opposite state). For instance, one split-half of the task data for MSC01 was compared with the corresponding split-half of rest data for MSC02. The

similarity in variant spatial location maps was computed with a Dice correlation. This was repeated for all combinations of subjects and split-halves, yielding 128 comparisons. To compute t-statistics for this comparison, the values from all of the comparisons for each subject's data versus every other subject were averaged into one value for that subject.

2.9.5. Comparison of variant magnitude differences across states—To compare the magnitude of individual differences between states, variants were defined using the procedure outlined in Section 2.7.3. Then, the spatial correlation magnitude at each of these vertices was calculated in both states. Finally, the mean absolute difference between rest and task spatial correlation magnitudes was calculated across all vertices. This provided a measure of the absolute deviation of spatial correlation values across states for network vertices.

To test whether these values were bigger than what would be expected in other areas of the cortex, the variants in each participant were randomly rotated, thereby matching for variant size and shape in the null distribution (Gordon et al., 2017). Within each participant, 1000 rotations were randomly generated and performed within each hemisphere. In the case that a rotation resulted in parcels that overlapped with the medial wall, this rotation was recalculated until none of the rotated parcels overlapped with each other or the medial wall. For each of the 1000 rotations, the mean absolute difference between task and rest spatial correlation magnitudes was calculated. These values were then compared to the mean absolute difference in spatial correlation magnitudes for variant locations in each participant.

2.10. Network variants versus other vertices

Another outstanding issue in the measurement of network variants is whether they are more similar between states than would be expected from regions with more similarity to the group. Such a finding would provide evidence for their trait-like stability, as posited elsewhere (Seitzman et al., 2019). In order to examine this issue, we thresholded maps at 2.5% increments starting with the lowest 2.5% of spatial correlation values versus the group average for each split-half within each participant (e.g., bin 1: 0–2.5% similarity to the group, bin 2: 2.5–5%, bin 3: 5–7.5%, …, etc.). For each 2.5% increment, we compared their spatial locations via Dice correlation for within-state, between-state, and across-subjects comparisons. The average value for each of these comparisons was then plotted with its associated standard error (across subjects). No size exclusion was applied to this analysis.

2.11. Similarity of network variants in individual tasks versus multiple tasks

We also evaluated the degree to which the spatial location of variants defined at rest is similar to variants defined from a single task versus variants defined from multiple tasks combined. As above, the amount of data was matched across individual tasks and participants (see Table S1). Data were sampled into split-halves using the same procedure previously outlined. This same procedure was also applied to sample an equal amount of data from all three tasks combined. The amount of data for each individual task and for all three tasks combined was 11.3 min total per split-half for each task and participant. MSC09 was not used in this analysis due to the small amount available for this participant (1.25 min in the smaller split-half). Dice correlations were used to quantify the spatial location overlap

of network variants within-state, between-state, and across-subjects. Within subject ANOVAs were used to test for differences in spatial overlap. Due to the relatively small amount of individualized data available for this analysis, the results should be interpreted with some caution.

2.12. Similarity of network variant assignments between states

We next assessed whether variants also showed similar network assignments between states. To test this, each contiguous variant unit was assigned to a network, and the similarity of these networks assignments was evaluated between states and across subjects.

2.12.1. Assigning network variants to functional networks—In this analysis, each variant was assigned to a functional network via a winner-take-all algorithm. This algorithm assigned variants to one of 14 group average template networks (Gordon et al., 2017). The templates for matching the network variants were the spatial maps of the group average canonical networks from the group average of the WashU 120 dataset. First, a seed map was generated for every vertex composing a given variant. These seed maps were then averaged together to form one seed map for each network variant (Seitzman et al., 2019). For each map and the template, the highest 5% of correlations were binarized. Next, we computed the similarity (Dice correlation) between the binarized spatial location map for each network variant and the spatial location map for each canonical network template. Vertices within 30 mm geodesic distance from any vertex in a given variant were excluded from that variant's seed map and the template seed map (Gordon et al., 2017). Each network variant was then assigned to the "winning" template with which it showed the greatest spatial overlap (highest Dice correlation). If no good "winning" template was found (as determined by the lowest 5% of dice correlation values across subjects, Dice < 0.1428 at 5% threshold) or multiple networks tied for the highest value, then the variant was assigned to an "unknown" network. In addition, network variants were also removed from further analysis if their assigned network spatially overlapped at least 50% with the group average network template (9/123 = 7% in rest, 11/123 = 9% in task; (Seitzman et al., 2019).

2.12.2. Evaluating variant network assignment consistency between states—
We then investigated how likely variants are to show consistent network assignments between states and across subjects. To do this, first variants in task and rest states were assigned to networks according to the procedure above. Next, the same vertices were selected in the other state (e.g., the variant locations from task data were selected in the resting state data), and the network assignment procedure was repeated for these vertices in the other state. A Dice correlation was then used to quantify the likelihood of the same vertices between states being assigned to the same network (with matching network assignments set to 1 and non-matches set to 0). To test whether variant assignment was more stable between states than across subjects, we also used this method to quantify how often the same vertices assigned to the same network in all other subjects.

Additionally, given that contiguous variant units may not be homogenous, we also examined whether assigning the averaged seedmaps for contiguous variant units (> 50 vertices) to networks systematically differed from assigning the seedmaps for each individual vertex

composing a variant to a network. The network assignment procedure used was identical, but instead of using the average seedmap for each variant, the seedmaps for each individual vertex composing each variant were assigned to a network. We then compared the distributions for the network assignments of individual vertices versus contiguous variants, as well as the similarity of the vertex-level network assignments between states and across subjects.

We also examined the consistency of network variants' functional connectivity patterns between states via the correlation of the full seed map for each variant across states. For each variant, a seedmap was calculated for each vertex and then all of the seedmaps contained in a variant were averaged together. This same process was then performed for those same vertices in the opposite state. As a comparison, this was also repeated for the same location in the opposite state across all the other subjects. The similarity of these maps was computed via a spatial correlation.

2.13. Relationship between task activations and functional connectivity

In an additional analysis, we explored whether the amount of task activation relative to baseline was associated with the likelihood of a variant showing inconsistent patterns of functional connectivity across task states. If task activations systematically alter functional connectivity, then it could be that areas which show larger magnitude task activations may be identified as variants in task states but not during rest. To test this question, the average task activation values were calculated separately for vertices which were identified as variants in both task and rest states (overlapping vertices), vertices which were identified as variants only in rest states, and variants which were identified as variants only in task states (see Fig. 3B). The average task activation for each subject was then calculated for variants identified in both states (overlapping), rest states only, and task states only. Each parcel of contiguous vertices in each state had to be composed of at least 30 vertices to be included in this analysis. These values were then averaged within subjects and tasks and compared via *t*-tests.

3. Results

3.1. High amounts of task data can produce reliable estimates of network variants

The overarching question of this manuscript is whether network variants are similar between task and rest states, indicating that they are trait-like and that task data may be used to identify variants. As a first step, we asked if task data can produce reliable (i.e., high within-state similarity) estimates of network variants, as has been shown for rest data. Following Gordon et al. (2017) and Laumann et al. (2015), reliability was estimated using a split-half procedure. Specifically, task data from each participant was divided into two halves. Variants were created from one half of the data (mean = 109 min, range = 75.9–123.9 min); these variants were treated as our best estimate of "true" variants. The other half was incrementally sampled in 5 min increments of data; each incremental sample was used to generate a new estimate of network variants. This incremental sample estimate was then compared to the other full half (the "true" estimate). The same procedure was conducted both for the comparisons of the continuous spatial correlation maps (Fig. 2A, compared

using Pearson correlations) and for the binarized spatial locations of network variants (Fig. 2B, compared using Dice correlations).

The general trend is that as data quantity increases (x-axis), the similarity between the two halves of data also increases (y-axis) before reaching an asymptote, typically around 40 min. of data. In particular, with > 40 min the continuous spatial correlation maps reach reliabilities with $r > 0.8$ (Fig. 2A). In addition, the binary overlap between spatial locations for network variants (the lowest 5% of correlations with the group average) also becomes stable with high amounts of data (Fig. 2 B; note that however these values typically asymptote at a lower level, around Dice correlation = 0.7, likely due to instability caused by the thresholding operation). However, reaching stability for the binarized data requires additional time beyond what is necessary to achieve reliable spatial correlation maps (> 70 min). Overall, this task reliability profile is similar to what is seen with rest (Fig. S1). Similar results are also obtained using non-residualized task data (i.e., data without task activation removal, treated the same as rest; Fig. S2) and when using an absolute ($r < 0.3$) rather than percentile threshold (Fig. S3; though the results are more variable across subjects than when a percent threshold is used). Thus, similar to rest, high amounts of task data can produce reliable measures of network deviation in individuals with respect to the group average.

3.2. Network variants occur in relatively similar locations across states

Next, we examined the similarity of network variants between states. The number of variants (Table S2) and vertices (when thresholded at $r < 0.3$; Table S3) in both task and rest states are summarized in the supplemental tables. We did not see significant differences in the quantity of variants for either measure between states (all $p > .05$).

We next focused on whether network variants occur in the same locations between task and rest. If network variants are trait-like, then their vertices would be expected to overlap at a significantly greater rate between states (within a person) than across people. Consistent with this pattern, Fig. 3A,B shows that network variants appear in similar spatial locations between task and rest states in single individuals (see Figure S4 for all participants).

To quantify the degree of similarity in spatial locations for network variants between states, we calculated the Dice coefficient between task and rest state network variants. We benchmarked this between-state overlap against two comparisons: (a) within-state overlap (how similar the spatial locations of network variants are across two different samples from the same state, e.g., between two split-halves of resting state data from MSC01) and (b) across-subject overlap (how similar the spatial locations of network variants are between subjects, e.g., between a task split-half of MSC01 and a rest split-half from MSC02). These comparisons allowed us to determine if the spatial locations of network variants were more trait-like (showing relatively similar variants between states, but distinct variants across people) or state-like (showing relatively distinct variants between states, closer to the level seen across people). Importantly, data quantities were matched for all comparison samples within and across individuals (see Table S1).

The results showed that variants were more likely to spatially overlap between states ($M = 0.533$, $SD = 0.103$) than across subjects ($M = 0.068$, $SD = 0.025$, $t(8) = 14.96$, $p < .0001$, $d = 4.99$), suggesting a large trait-like component to network variants. Variants were also significantly more likely to spatially overlap within states ($M = 0.641$, $SD = 0.082$) than between states ($M = 0.533$, $SD = 0.103$, $t(8) = 6.38$, $p = .0002$, $d = 2.13$; see Fig. 4), suggesting some state-dependence. However, the magnitude of the difference between states was substantially smaller than the magnitude of the difference across subjects. Similar results are also observed at a 2.5% threshold (see Fig. S5) as well as when using an absolute ($r < 0.3$) rather than percentile threshold (see Fig. S6). Non-residualized task data (where task activations were not removed) showed a similar overall pattern, but significantly lower between-state stability ($t(8) = 8.09$, $p < .0001$, $d = 2.7$; Fig. S7), suggesting that it may be important to address task activations to improve the correspondence of network variant locations between task and rest. Jointly, these findings indicate that variant locations are affected by state, but are relatively more trait-like than state-like.

In light of these results, we were also interested in whether network variants exhibited larger state-dependent changes in their underlying spatial correlations to the group map than would be expected from other similarly sized clusters on the cortex. For this, we quantified the spatial correlations to the group map between states for network variants versus other areas of the cortex (Fig. 3C). On average, the differences in correlation values between states are relatively small for network variants ($M = 0.102$, $SD = 0.016$, $Skew = 0.824$). However, the tail of the distribution for these magnitude differences is somewhat larger for network variants than would be expected by chance in other areas of the cortex ($M = 0.064$, $SD = 0.012$, $Skew = 0.522$), indicating that a subset of locations show larger magnitude differences than expected for other areas of cortex (see Figure S8). This is consistent with some network variant locations showing an interaction between individual differences in functional connectivity and task state.

Finally, we asked whether task activations were related to inconsistencies in network variant identification between states. To examine this, each variant was divided up into sub-units depending on whether its vertices were identified as variants in both task and rest states (overlapping), in rest only, or in task only. The average task activation was then calculated separately for each variant sub-unit (Fig. S9). The mean task activations did not differ by variant sub-units for any of the tasks ($p > .05$). This suggests that state-dependent differences in network variants are not simply driven by average task activations. Future studies will be needed to examine if more nuanced relationships link specific activations to specific network variants.

3.3. Network variants associate with similar networks in task and rest

While the spatial locations of network variants are relatively similar between states, it is possible that the variants may show different underlying functional correlation maps. To test this possibility, each variant was assigned to a template network using a winner-take-all procedure in each state (see Section 2.12.1). A Dice correlation was then used to quantify how often variants were assigned to the same network between states. As shown in Fig. 5B, the distribution of variant network assignments between states was similar. In addition, the

likelihood of obtaining a matching network assignment for a variant using the same vertices between states was high ($M = 0.814$, $SD = 0.109$; All Variants in Fig. 5C). This was significantly greater than the likelihood of variants matching across subjects ($M = 0.159$, $SD = 0.056$), $t(8) = 20.03$, $p < .0001$, $d = 6.68$. This shows that network variants tend to be assigned to the same network between states (versus across subjects), providing evidence for their trait-like stability.

In addition to network assignments, we also looked at the similarity of whole-brain seedmaps for each variant (rather than using a winner-take-all procedure to assign to one of 14 networks). The purpose of this analysis was to examine the similarity of seedmaps between states and across subjects while avoiding the non-linearities inherent in the network assignment procedure. Variant seedmaps showed a very high similarity between states ($M = 0.824$, $SD = 0.033$), much higher than across subjects ($M = 0.034$, $SD = 0.041$, $t(8) = 33.07$, $p < .0001$, $d = 11.02$; see Fig. S10), suggesting that the network profiles of variant regions also exhibit trait-like properties.

The analyses reported above assume that contiguous variants represent single homogenous regions with the same network assignment. However, this may not always be the case and thus would distort any observed differences in network assignment across states. To test if treating contiguous variants as a single unit contributes to state-level differences in variants' network assignments, we assigned individual variant vertices to a network and repeated the assignment procedure (Fig. S11A). The distribution of the network assignments was extremely similar across states and to that seen with contiguous unit assignment. Additionally, variant vertices were much more likely to assign to the same network between states ($M = 0.819$, $SD = 0.031$) than across subjects ($M = 0.193$, $SD = 0.028$, $t(8) = 41.69$, $p < .0001$, $d = 13.9$; Fig. S11B). These results confirm that variants show similar network assignment profiles between task and rest states, whether they are considered homogenous units or as individual vertices.

3.4. Network variants are highly stable compared to vertices more similar to the group

Next, we looked at whether network variants are unique trait-like clusters relative to other subsets of vertices selected based on their similarity to the group. To test this, we binned vertices into sub-units starting with those showing a low similarity to the group average map, (e.g., 2.5–5%, 5–7.5%, 7.5–10%, etc.), up through the most similar vertices to the group (e.g., 95–97.5%, 97.5–100%).

The results of this analysis are shown in Fig. 6. Interestingly, only bins on the extreme low and high ends of the curve contained vertices which formed contiguous clusters (Fig. 6A). Thus, regions of functional connectivity that are most similar and most dissimilar to the group average tend to represent focal regions. Not surprisingly given this pattern, network variant locations (the left side of the curve; Fig. 6B) show more stability between states, with higher within- and between-state similarity than almost any other percentile bin. The only other bins that showed as strong of stability were at the highest end of the spectrum: those with the most consistency to the group. This is likely driven by the fact that network variants (as well as the most group consistent regions) agglomerate into contiguous parcels, whereas other bins show a more scattered representation across cortex. This provides evidence for the

distinctive nature of network variants relative to other locations selected for their group similarity.

3.5. Network variant stability was similar when examined for single tasks

The last comparison tested was whether the stability in variant spatial locations differed between any of the 3 tasks collected in the MSC. A significant effect would indicate that there were task-specific effects of how stable variants spatial locations were across states. To test this effect, similar overlap analyses as those presented in Fig. 4 were carried out, but with data separated by task rather than concatenating across all tasks. Note that, given the separation, this analysis was conducted on relatively less data (11.3 min per split-half) and thus the results should be interpreted with caution due to their lower reliability.

Overall, the results suggested that variant similarity between task and rest does not substantially differ between tasks (Fig. 7). To test this, a one-way within subjects ANOVA was run on each subject's values from the rest-to-task comparisons for each task. Greenhouse-Geisser corrections were applied where appropriate. The effect of task was not significant, ($F(1.631,11.418) = 1.565$, $p = .248$, $d = 0.95$), indicating that there were no significant differences in stability for any individual task versus rest. This suggests that there is no difference in the stability of network variants when using a single task versus multiple tasks compared to rest. Similar results are also seen when using a 2.5% threshold to identify network variants (see Fig. S13).

4. Discussion

This study examined whether network variants (locations where an individual's brain network organization differs markedly from the group) were susceptible to state-level changes induced by task demands. This question has bearing on both the trait-like aspects of network variants and on practical questions of whether task data can reasonably be used to identify network variants. To examine this question, we used resting state and task fMRI data from the Midnight Scan Club to define network variants. These network variants were then compared to each other in various ways to evaluate whether network variants showed trait-like stability between states. We found that (1) network variants defined from task data show similar reliability to resting state data, (2) network variants appear in similar spatial locations between task and rest data, although with some small state-dependent effects, (3) the network assignments of variants are similar between task and rest, and (4) that there is similar stability for network variants identified from multiple or single tasks. Thus, our analyses suggest that while network variants show some differences that are due to state-level factors, they are predominantly stable between states, consistent with their trait-like properties. Importantly, these findings suggest that a reasonable approximation of network variants can be identified from task as well as rest.

4.1. Network variants are largely trait-like

The current findings add an additional piece of evidence to suggest that network variants are relatively trait-like in nature by demonstrating their consistency across task states. We found that individuals' network variants tend to occur in similar locations across different states as

well as have similar network labels. These findings add to previous evidence from Seitzman et al. (2019) showing (a) that within states, network variants are reliable across sessions, even over a year, (b) that variants defined at rest and assigned to the default mode network can be used to predict deactivations during task performance, and (c) that variants can be used to separate individuals into sub-groups related to behavioral measures acquired outside of the scanner. This consistency across states, high reliability, and relationship to brain activations and behavioral measures collected outside the scanner jointly suggest that network variants may be trait-like properties of brain networks, rather than functions of ongoing processing.

Taken together, these results suggest that the neurobiological underpinnings of network variants may stem from a systematic change in the function of underlying brain tissue from that which is typically observed in the population, is als altering the Hebbian associations between brain regions (Gratton et al., 2019; Seitzman et al., 2019). By altering functional association between brain regions, functional connectivity will be altered as well. Given their stability, these individual-level differences in brain network organization may be driven by longer-term factors such as genetics, prolonged life experiences, or other environmental factors. Under this framework, network variants may be viewed as prime candidates for neural correlates of trait-like individual differences in behavior, such as cognitive ability and risk for psychopathology.

Although network variants primarily exhibit trait-like properties, it is also clear that they show some minor but significant state dependence. This is consistent with past work showing that the functional organization of the brain is largely stable, but can show subtle modifications in adjusting to different task states (Cole et al., 2014; Gratton et al., 2016; Krienen et al., 2014). In the context of individual differences, recent work has emphasized that stable group and individual factors dominate brain networks, but more modest task-based effects are also present in functional organization between task and rest states across the full connectome (Gratton et al., 2018). These state-dependent changes in functional connectivity likely also contribute to the observed differences in network variants between different task states.

One possible explanation for this observed state-dependence is that if network variants are areas which have shifted in function from their 'typical' role, this may also compound differences during particular task states. The exact nature of this shift would thus be dependent on how the task affects this atypical neuronal population and its coherence with other regions. While there is some evidence that functional connectivity changes more in task active regions (Gratton et al., 2016), we did not find that state-dependent shifts in network variant connectivity were related to overall task activations (see Fig. S9). It is possible that variant state changes are related to more nuanced relationships between specific variant connectivity profiles and particular task processes. A full examination of this question will benefit from future studies adopting specialized tasks with contrasts targeting discrete networks (e.g., Braga et al., 2019; DiNicola et al., 2020). Future studies with a larger sample of tasks and participants will also elucidate whether these state-dependent changes in variants are systematically related to unique strategies or task performance

(performance on the current tasks, which were all relatively simple, was at ceiling and all MSC individuals are high IQ).

4.2. Practical considerations for identifying network variants with task data

Due to the similarities of network variants between task and rest states, it seems reasonable that network variants can be estimated using task data. Network variants show substantial stability in spatial location and network assignments between rest and task states ($M = 0.631$; see Figure S13). In practice, there is approximately a 17% reduction in the stability of network variant spatial location when examined between states ($M = 0.533$) versus within states ($M = 0.641$). The underlying correlation patterns and network assignments show an even higher stability between states ($M = 0.814$). While variants in diverse states are not as consistent as those within the same state, the results suggest that it is reasonable to conduct network variant analyses across tasks if needed for data quantity reasons. Notably, spatial location differences in network variants between states are much smaller than the differences observed from using small amounts of data with low reliability (approximately 56% lower stability between 5 min and 70 min of data).

Critically, the finding that variants can be identified during task states that have good correspondence to rest opens up new datasets for variant analysis that would not be available otherwise. While most currently available datasets only have small amounts of rest fMRI data (5–10 mins) which are insufficient to achieve high reliability (Fig. 2B, S1B), many have substantial additional task fMRI data. Moreover, while datasets from clinical populations are of great interest, it may be particularly difficult to obtain large amounts of high quality resting state data in many of these populations (due to exacerbated head motion, compliance issues, etc.; Greene et al., 2018; Hodgson et al., 2017; Vanderwal et al., 2019). Thus, combining data across task and rest may be an effective strategy to increase functional connectivity reliability sufficient for network variant analyses without sacrificing substantial variability due to state differences. An important caveat is that properly accounting for task activations may be critical to ensuring higher similarity between variants in rest and task states. Because of this, recent work suggesting movie watching (where it is difficult to remove evoked signals) as an alternative to rest may not be ideal for combining data across states.

Elliot et al. (2019) have shown that combining resting state data with background connectivity during task performance increases the reliability of functional connectivity data as a whole. Critically, this increase in reliability was also accompanied by an increase in the predictive utility of functional connectivity for cognitive measures and the heritability of functional connectivity. Additional work has also incorporated this strategy of analyzing functional connectivity data in developmental samples (Cui et al., 2020), demonstrating its utility in a dataset with a limited amount of resting-state fMRI data. Our results suggest that this method can also be used to identify network variants, which may yield new insights into neurologic and psychiatric disease.

4.3. Important methodological considerations for identifying network variants

Despite this promise for using task data to identify network variants, some caveats for this approach are discussed below. First, while the reliability of cortical functional connectivity overall becomes stable with > 30 min of data (Gordon et al., 2017; Laumann et al., 2015), the definition of network variants appear to require more data to achieve an adequate reliability. From our estimates, it appears that > 70 min of total (rest or task) fMRI data are required to achieve asymptotic reliability for network variants. This added need for data is likely due to the thresholding procedure used in defining network variants, as spatial correlations maps (Fig. 2A, S1A) achieve high reliability more quickly. Combining task and rest data may be an effective strategy to achieve this amount of data in available datasets with insufficient rest data alone.

In addition to combining rest and task data, our initial evidence suggests that it is possible to either use data from a single task or combine data from multiple tasks to generate network variants. We found no significant differences in the ability to identify network variants across 3 different tasks versus rest, indicating that all of the tasks performed relatively similarly. However, for this analysis we were restricted to using a small amount of data which complicates the interpretation of our analysis. It may be possible that with larger amounts of data or with different tasks, task-specific effects for network variants could be observed.

A final caveat of our analyses is that it is important to match data in terms of both quantity and quality when comparing the reliability of fMRI data. Because reliability of a scan steadily increases with the amount of low-motion data (Elliott et al., 2019; Gordon et al., 2017; Laumann et al., 2015; Noble et al., 2017), it is necessary to match these attributes in order to make accurate comparisons. Factors such as head motion, sampling across sessions, or processing may systematically influence estimates of functional connectivity. In the current study, we found similar levels of motion between task (Mean FD = 0.129) and rest (Mean FD = 0.133), and data was approximately equally sampled across sessions (see Table S1). Recent work also highlights the importance of accounting for the presence of task activations when using task data for functional connectivity analyses (Cole et al., 2019). We also found evidence in line with this conclusion, demonstrating that the presence of task activations makes network variants less similar between task and rest states. Therefore, it is important to account for these task activations to improve similarity between task and rest variants (although overall within-state reliability was similar regardless of the precesen of evoked signals). Thus, it is critical to control for these factors when comparing data across states, participants, and studies. Our results also suggest that using a percent threshold to define variants may work better than an absolute cutoff. We found more stability in variants across states defined with a percentile rather than an absolute cutoff. This may be an important consideration to maximize the reliability of variants across individuals.

5. Limitations

We also note several limitations to the current study. First, although the MSC dataset has a large amount of resting state fMRI data, it was more limited in the amount of available task fMRI data after removal of fixation periods, especially for single tasks. This is particularly

true of the motor task, which had a high number of frames lost due to movement. Therefore, we only draw limited inferences about task-specific effects on network variants from the current study.

Second, although we found some state-dependent changes in network variants, it was unclear what drove these differences. We did not observe systematic patterns of between-state variations associated with particular networks or particular tasks. It is possible that future work with a broader set or more targeted tasks may help to identify these differences and shed light on state-dependent mechanisms.

Third, several of the analysis decisions in this paper are based on thresholding data, which has the potential to systematically change patterns in the data. One example of this is the thresholding of network variants in individuals as the lowest 5% of correlations with the group average networks. This decision was based on examining histograms of correlations to the group for each individual and finding the approximate tail of the distribution. In addition, we also demonstrated that the reported results replicate at other thresholds (2.5%, $r < 0.3$) showing their robustness to this thresholding choice. Another thresholding choice was in the assignment of variants to networks, which used a winner-take-all network assignment procedure based on the spatial overlap of the top 5% of correlations between a seed region and a group average network. It should be noted that the network assignments mirrored past work (Seitzman et al., 2019), with higher-level association networks making up the majority of network variants and seed map-based comparisons (free of thresholding or winner-take-all selection) were highly stable across states (Fig. S10). To alleviate these concerns, future studies may consider employing newer techniques such as bagging (Nikolaidis et al., 2020) to improve reliability of thresholded data or "soft" multi-network assignments (Cui et al., 2020) for networks with an ambiguous pattern of connectivity.

Fourth, we defined network variant "units" based on their contiguity. However, it is possible that contiguous variant units contain adjacent, but not homogenous deviant tissue – that is, they may not index a single underlying brain area. Moreover, some variant units may index shifts in the borders of nearby systems, while others may index more drastic functional alterations in an entire isolated region. We conducted analyses at the vertex-level (see Figures S11, S13) to address potential inhomogeneities in network variant units. In future work, we hope to expand our understanding of homogenous variant units and their properties.

Fifth, it is possible that the stable trait-like characteristics of network variants are linked to mis-localizations in anatomical alignment. In previous work, we have found that areas with gross anatomical deformations during registration are not associated with network variants (Seitzman et al., 2019). However, it is possible that network variants occur, in part, due to fine-scale anatomical variations (e.g. cell populations) in given regions of the cortex. These anatomical variations would then contribute to the observed variability in functional connectivity across individuals. Additional investigations of network variants using different methodologies will be needed to further examine their underlying neurobiology.

6. Conclusion

Overall, these results suggest that brain network variants show trait-like stability between multiple states. Network variants measured with task data were reliable (given sufficient data), appeared in similar spatial locations, and showed similar patterns of functional connectivity to those identified using rest data. There were also smaller state-dependent differences associated with network variants. Jointly this work suggests that combining rest and task data may be a reasonable strategy to identify network variants from datasets with insufficient rest.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

References

Al-Aidroos N, Said CP, Turk-Browne NB, 2012. Top-down attention switches coupling between low-level and high-level areas of human visual cortex. Proc. Natl. Acad. Sci 109 (36), 14675–14680. [PubMed: 22908274]

Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C, 2013. Function in the human connectome: task-fMRI and individual differences in behavior. Neuroimage 80, 169–189. [PubMed: 23684877]

Biswal B, Zerrin Yetkin F, Haughton VM, Hyde JS, 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34 (4), 537–541. [PubMed: 8524021]

Braga RM, Buckner RL, 2017. Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. Neuron 95 (2), 457–471. [PubMed: 28728026]

Braga RM, DiNicola LM, Becker HC, Buckner RL, 2019. Situating the left-lateralized language network in the broader organization of multiple specialized large-scale distributed networks. J. Neurophysiol 124 (5), 1415–1448.

Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE, 2014. Intrinsic and task-evoked network architectures of the human brain. Neuron 83 (1), 238–251. [PubMed: 24991964]

Cole MW, Ito T, Schultz D, Mill R, Chen R, Cocuzza C, 2019. Task activations produce spurious but systematic inflation of task functional connectivity estimates. Neuroimage 189, 1–18. [PubMed: 30597260]

Cui Z, Li H, Xia CH, Larsen B, Adebimpe A, Baum GL, Cieslak M, Gur RE, Gur RC, Moore TM, 2020. Individual variation in functional topography of association networks in youth. Neuron.

Dale AM, Fischl B, Sereno MI, 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. Neuroimage 9 (2), 179–194. [PubMed: 9931268]

Dice LR, 1945. Measures of the amount of ecologic association between species. Ecology 26 (3), 297–302.

DiNicola LM, Braga RM, Buckner RL, 2020. Parallel distributed networks dissociate episodic and social functions within the individual. J. Neurophysiol 123 (3), 1144–1179. [PubMed: 32049593]

Dosenbach NU, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RA, Fox MD, Snyder AZ, Vincent JL, Raichle ME, 2007. Distinct brain networks for adaptive and stable task control in humans. Proc. Natl. Acad. Sci 104 (26), 11073–11078. [PubMed: 17576922]

Elliott ML, Knodt AR, Cooke M, Kim MJ, Melzer TR, Keenan R, Ireland D, Ramrakha S, Poulton R, Caspi A, 2019. General Functional Connectivity: shared features of resting-state and task fMRI drive reliable and heritable individual differences in functional brain networks. Neuroimage 189, 516–532. [PubMed: 30708106]

Fair DA, Miranda-Dominguez O, Snyder AZ, Perrone A, Earl EA, Van AN, Koller JM, Feczko E, Tisdall MD, van der Kouwe A, 2020. Correction of respiratory artifacts in MRI head motion estimates. Neuroimage 208, 116400. [PubMed: 31778819]

Fair DA, Schlaggar BL, Cohen AL, Miezin FM, Dosenbach NU, Wenger KK, Fox MD, Snyder AZ, Raichle ME, Petersen SE, 2007. A method for using blocked and event-related fMRI data to study "resting state" functional connectivity. Neuroimage 35 (1), 396–405. [PubMed: 17239622]

Finn ES, Shen X, Scheinost D, Rosenberg MD, Huang J, Chun MM, Papademetris X, Constable RT, 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. Nat. Neurosci 18 (11), 1664. [PubMed: 26457551]

Fischl B, Salat DH, Busa E, Albert M, Dieterich M, Haselgrove C, Van Der Kouwe A, Killiany R, Kennedy D, Klaveness S, 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33 (3), 341–355. [PubMed: 11832223]

Glass L, 1969. Moire effect from random dots. Nature 223 (5206), 578. [PubMed: 5799528]

Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, 2013. The minimal preprocessing pipelines for the Human Connectome Project. Neuroimage 80, 105–124. [PubMed: 23668970]

Gordon EM, Laumann TO, Adeyemo B, Gilmore AW, Nelson SM, Dosenbach NU, Petersen SE, 2017a. Individual-specific features of brain systems identified with resting state functional correlations. Neuroimage 146, 918–939. [PubMed: 27640749]

Gordon EM, Laumann TO, Adeyemo B, Huckins JF, Kelley WM, Petersen SE, 2016. Generation and evaluation of a cortical area parcellation from resting-state correlations. Cerebral Cortex 26 (1), 288–303. [PubMed: 25316338]

Gordon EM, Laumann TO, Adeyemo B, Petersen SE, 2017b. Individual variability of the system-level organization of the human brain. Cerebral Cortex 27 (1), 386–399. [PubMed: 26464473]

Gordon EM, Laumann TO, Gilmore AW, Newbold DJ, Greene DJ, Berg JJ, Ortega M, Hoyt-Drazen C, Gratton C, Sun H, 2017c. Precision functional mapping of individual human brains. Neuron 95 (4), 791–807. [PubMed: 28757305]

Gratton C, Dworetsky A, Coalson RS, Adeyemo B, Laumann TO, Wig GS, Kong TS, Gratton G, Fabiani M, Barch DM, 2020. Removal of high frequency contamination from motion estimates in single-band fMRI saves data without biasing functional connectivity. Neuroimage, 116866. [PubMed: 32325210]

Gratton C, Kraus BT, Greene DJ, Gordon EM, Laumann TO, Nelson SM, Dosenbach NU, Petersen SE, 2019. Defining individual-specific functional neuroanatomy for precision psychiatry. Biological Psychiatry.

Gratton C, Laumann TO, Gordon EM, Adeyemo B, Petersen SE, 2016. Evidence for two independent factors that modify brain networks to meet task goals. Cell Reports 17 (5), 1276–1288. [PubMed: 27783943]

Gratton C, Laumann TO, Nielsen AN, Greene DJ, Gordon EM, Gilmore AW, Nelson SM, Coalson RS, Snyder AZ, Schlaggar BL, 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. Neuron 98 (2), 439–452. [PubMed: 29673485]

Greene DJ, Koller JM, Hampton JM, Wesevich V, Van AN, Nguyen AL, Hoyt CR, McIntyre L, Earl EA, Klein RL, 2018. Behavioral interventions for reducing head motion during MRI scans in children. Neuroimage 171, 234–245. [PubMed: 29337280]

Greene DJ, Marek S, Gordon EM, Siegel JS, Gratton C, Laumann TO, Gilmore AW, Berg JJ, Nguyen AL, Dierker D, 2019. Integrative and network-specific connectivity of the basal ganglia and thalamus defined in individuals. Neuron 105 (4), 742–758. [PubMed: 31836321]

Hodgson K, Poldrack RA, Curran JE, Knowles EE, Mathias S, Göring HH, Yao N, Olvera RL, Fox PT, Almasy L, 2017. Shared genetic factors influence head motion during MRI and body mass index. Cerebral Cortex 27 (12), 5539–5546. [PubMed: 27744290]

Kong R, Li J, Orban C, Sabuncu MR, Liu H, Schaefer A, Sun N, Zuo X–N, Holmes AJ, Eickhoff SB, 2019. Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion. Cerebral Cortex 29 (6), 2533–2551. [PubMed: 29878084]

Krienen FM, Yeo BT, Buckner RL, 2014. Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. Philos. Trans. R. Soc.B: Biol. Sci 369 (1653), 20130526.

Laumann TO, Gordon EM, Adeyemo B, Snyder AZ, Joo SJ, Chen M–Y, Gilmore AW, McDermott KB, Nelson SM, Dosenbach NU, 2015. Functional system and areal organization of a highly sampled individual human brain. Neuron 87 (3), 657–670. [PubMed: 26212711]

Laumann TO, Snyder AZ, Mitra A, Gordon EM, Gratton C, Adeyemo B, Gilmore AW, Nelson SM, Berg JJ, Greene DJ, 2016. On the stability of BOLD fMRI correlations. Cerebral Cortex 27 (10), 4719–4732.

Marcus D, Harwell J, Olsen T, Hodge M, Glasser M, Prior F, Jenkinson M, Laumann T, Curtiss S, Van Essen D, 2011. Informatics and data mining tools and strategies for the human connectome project. Front Neuroinform. 5, 4. [PubMed: 21743807]

Marek S, Siegel JS, Gordon EM, Raut RV, Gratton C, Newbold DJ, Ortega M, Laumann TO, Adeyemo B, Miller DB, 2018. Spatial and temporal organization of the individual human cerebellum. Neuron 100 (4), 977–993. [PubMed: 30473014]

Miezin FM, Maccotta L, Ollinger JM, Petersen SE, Buckner RL, 2000. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. Neuroimage 11 (6), 735–759. [PubMed: 10860799]

Miranda-Dominguez O, Mills BD, Carpenter SD, Grant KA, Kroenke CD, Nigg JT, Fair DA, 2014. Connectotyping: model based fingerprinting of the functional connectome. PLoS ONE 9 (11), e111048. [PubMed: 25386919]

Mueller S, Wang D, Fox MD, Yeo BT, Sepulcre J, Sabuncu MR, Shafee R, Lu J, Liu H, 2013. Individual variability in functional connectivity architecture of the human brain. Neuron 77 (3), 586–595. [PubMed: 23395382]

Nikolaidis A, Heinsfeld AS, Xu T, Bellec P, Vogelstein J, Milham M, 2020. Bagging improves reproducibility of functional parcellation of the human brain. Neuroimage, 116678. [PubMed: 32119986]

Noble S, Spann MN, Tokoglu F, Shen X, Constable RT, Scheinost D, 2017. Influences on the test–retest reliability of functional connectivity MRI and its relationship with behavioral utility. Cerebral Cortex 27 (11), 5415–5429. [PubMed: 28968754]

Ojemann JG, Akbudak E, Snyder AZ, McKinstry RC, Raichle ME, Conturo TE, 1997. Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. Neuroimage 6 (3), 156–167. [PubMed: 9344820]

Poldrack RA, Laumann TO, Koyejo O, Gregory B, Hover A, Chen M–Y, Gorgolewski KJ, Luci JM, Joo SJ, Boyd RL, 2015. Long-term neural and physiological phenotyping of a single human. Nat. Commun 6 (1), 1–15.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, 2011. Functional network organization of the human brain. Neuron 72 (4), 665–678. [PubMed: 22099467]

Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE, 2014. Methods to detect, characterize, and remove motion artifact in resting state fMRI. Neuroimage 84, 320–341. [PubMed: 23994314]

Power JD, Schlaggar BL, Lessov-Schlaggar CN, Petersen SE, 2013. Evidence for hubs in human functional brain networks. Neuron 79 (4), 798–813. [PubMed: 23972601]

Seitzman BA, Gratton C, Laumann TO, Gordon EM, Adeyemo B, Dworetsky A, Kraus BT, Gilmore AW, Berg JJ, Ortega M, 2019. Trait-like variants in human functional brain networks. Proc. Natl. Acad. Sci 116 (45), 22851–22861. [PubMed: 31611415]

Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, 2004. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23, S208–S219. [PubMed: 15501092]

Sorensen TA, 1948. A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on Danish commons. Biol. Skar 5, 1–34.

Sylvester CM, Yu Q, Srivastava AB, Marek S, Zheng A, Alexopoulos D, Smyser CD, Shimony JS, Ortega M, Dierker DL, 2020. Individual-specific functional connectivity of the amygdala: a substrate for precision psychiatry. Proc. Natl. Acad. Sci 117 (7), 3808–3818. [PubMed: 32015137]

Vanderwal T, Eilbott J, Castellanos FX, 2019. Movies in the magnet: naturalistic paradigms in developmental functional neuroimaging. Dev. Cogn. Neurosci 36, 100600. [PubMed: 30551970]

Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol 106 (3), 1125–1165. [PubMed: 21653723]
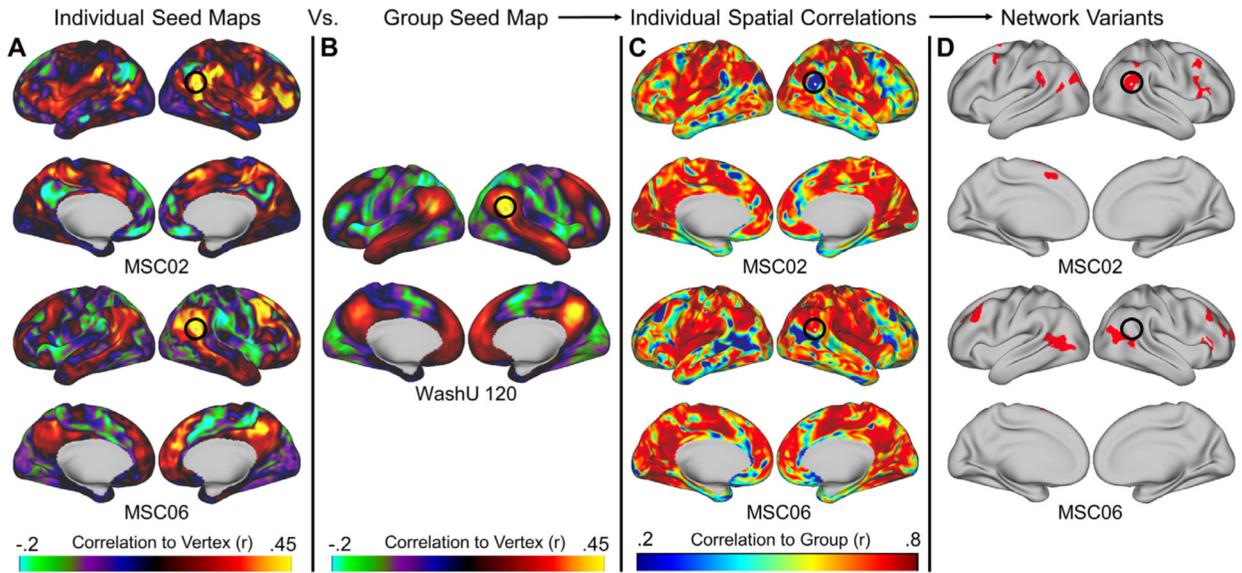
Fig. 1.

The process for generating network variants. First, a seed map is generated containing the correlations of every vertex to every other vertex within each individual during either task or rest (A) and the group average (B). Next, a spatial correlation between these maps is computed and the degree of similarity between each individual and the group is generated (C). Finally, regions with the lowest correlations (5%) with the group are selected as variants (D; variants are shown in red and small variants and locations in low signal regions are excluded). The seed maps shown here are for the vertex inside the black circles. This vertex corresponds to a network variant in MSC02 but not in MSC06.
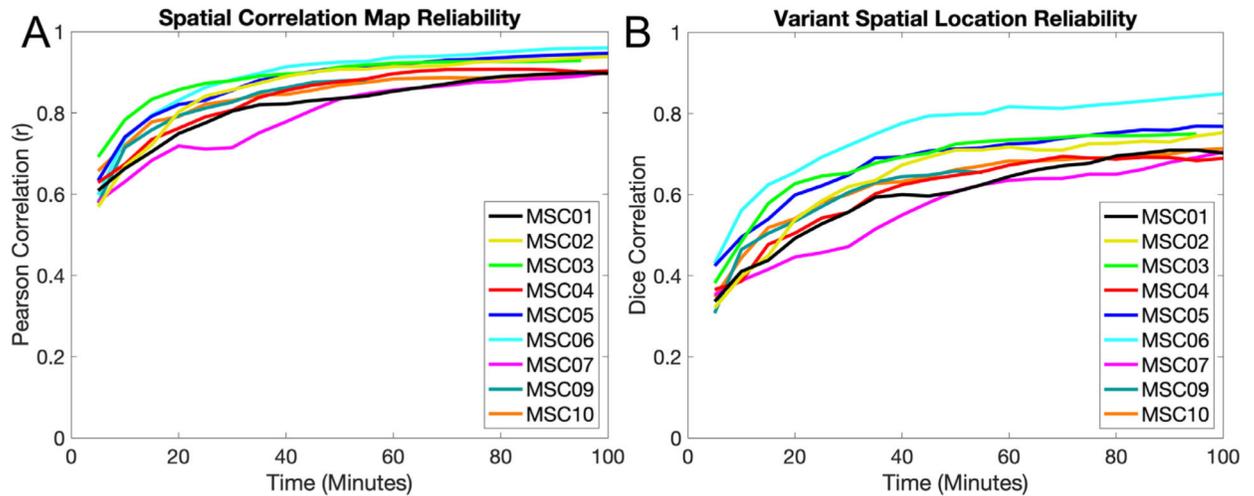
Fig. 2.

Reliability of network variants measured from task data. (A) Comparisons of individual-to-group spatial correlation maps from task data are shown for 5 min increments for each participant. The similarity between maps is calculated via Pearson correlation. Spatial correlation maps begin to plateau around 30–40 min at $r > 0.8$. (B) Comparisons of network variant maps (after binarization of the spatial correlations to the bottom 5% of spatial locations) from task data are shown in 5 min increments for each participant. The similarity between binarized spatial locations is calculated via Dice correlations. Binarized maps take longer (50–70 min.) and reach a slightly lower asymptote (Dice > 0.65), likely due to instability in the thresholding operation. Note that MSC09 (55 min) and MSC03 (95 min) did not have 100 min of data in their respective split-halves. A similar pattern was observed for the reliability of network variants in rest data (Fig. S1). Similar analyses were also performed using the non-residualized task data (Fig. S2) and data thresholded at $r < 0.3$ (Fig. S3).
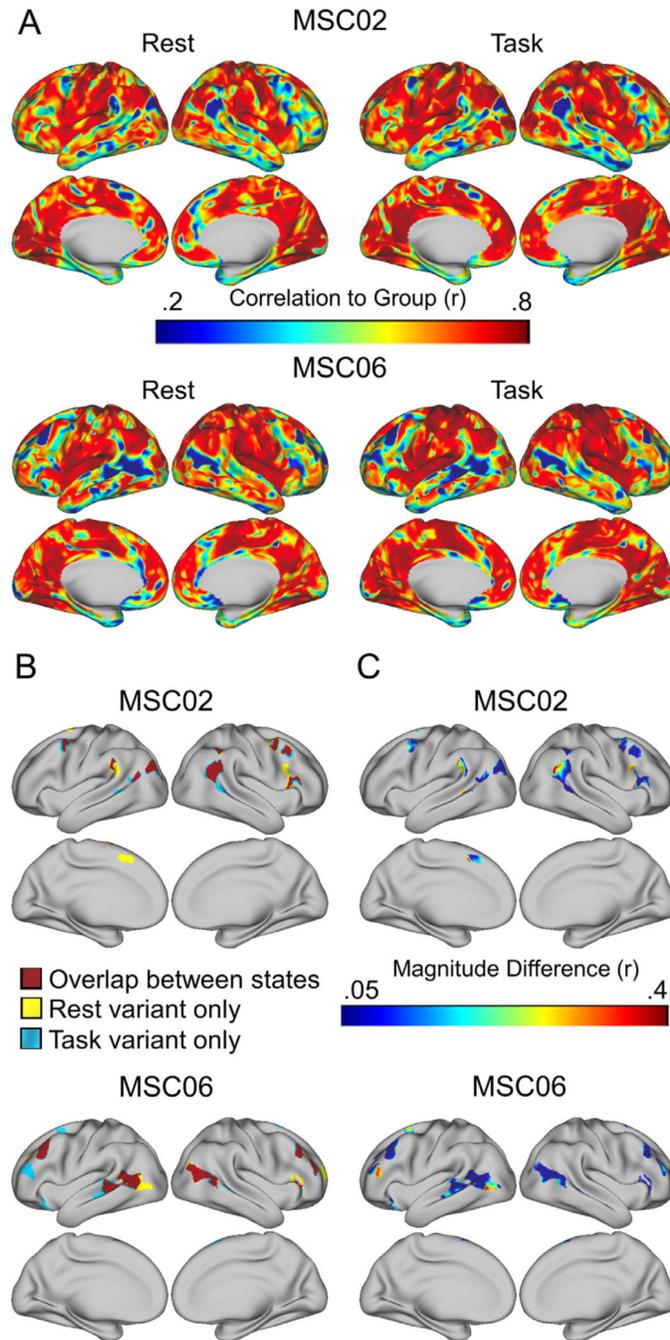
Fig. 3.

Comparison of network variant locations between task and rest states. (A) Individual-to-group spatial correlation maps are shown separately for task and rest states in MSC02 and MSC06; note the strong similarity in map profiles. (B) These maps are binarized (bottom 5% of correlations to group) to generate network variants for each state. The overlap between states is shown in red, and variant locations seen in only one state are shown in yellow (for rest) and blue (for task). (C) For each location labeled as a variant in either the task or rest states, we have plotted the absolute difference in magnitude of the spatial

correlations to the group between states. Most locations show small magnitude differences, but a small number show more sizeable differences.
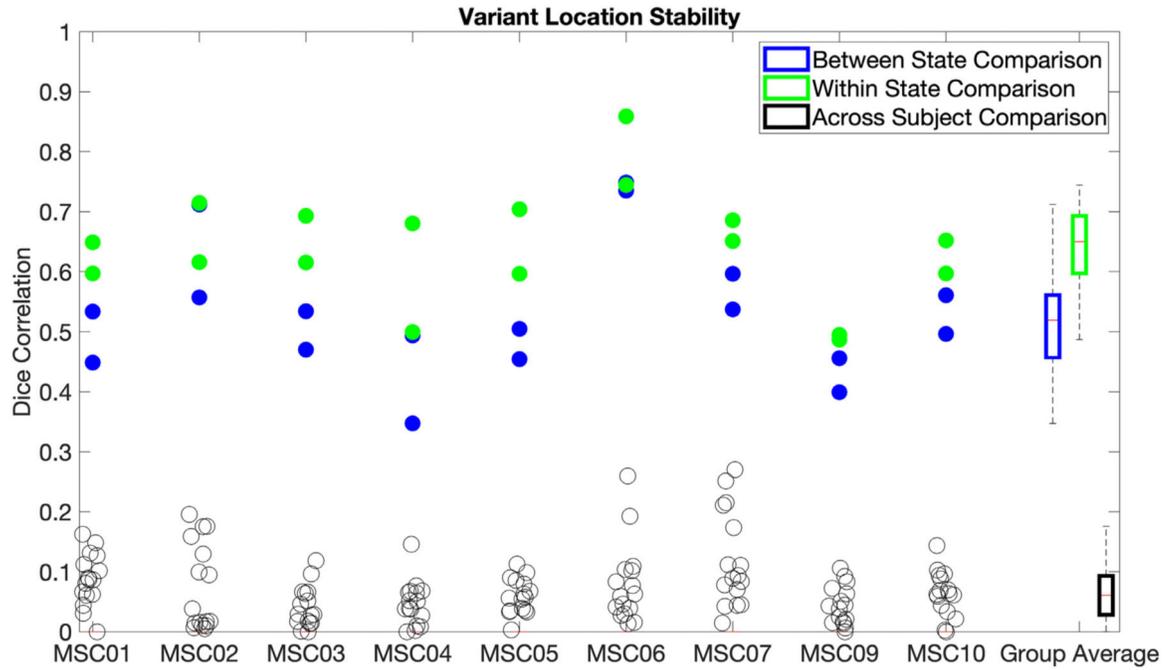
Fig. 4.
Network variant spatial locations were compared using a Dice correlation, for network variants between states (blue), within states (green), and across subjects (black). Values are plotted for each subject (columns) with a summary boxplot at the far right. Within and between state comparisons are shown as two dots, one for each pair of split-halves tested (see Sections 2.9.2 and 2.9.3). The same results are also reported using a 2.5% threshold (Fig. S6), for non-residualized data (Fig. S7), and for data using an absolute threshold ($r <$ 0.3; Fig. S8).
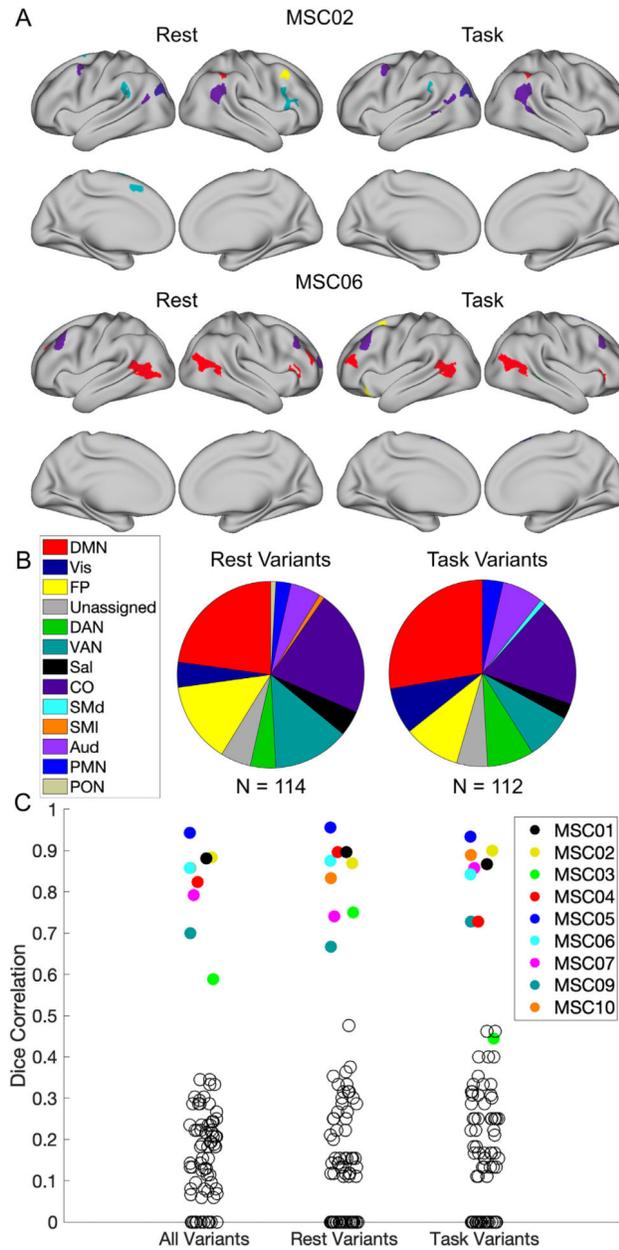
Fig. 5.
Comparison of variant network assignments between task and rest states. (A) The network assignments for variants are shown for task and rest states in MSC02 and MSC06. (B) The distributions for the networks that variants are assigned to in rest and task states are also shown across all subjects. (C) The likelihood of a variant in one state being assigned to the same network in the opposite state within a subject is plotted for each MSC subject. The empty black circles represent the likelihood of the variants in one subject being assigned to the same network in another subject. Comparisons were also made for whole-brain seedmaps (Fig. S10) and for network assignments performed at the vertex-level rather than on contiguous variants (Fig. S11), where again strong consistency was seen across states relative to across people. *Note that MSC03 has fewer variants than other participants in both*

states, and thus their data is much more susceptible to noise (see Fig. S11 *for evidence of high stability of MSC03 variants at the vertex level).*
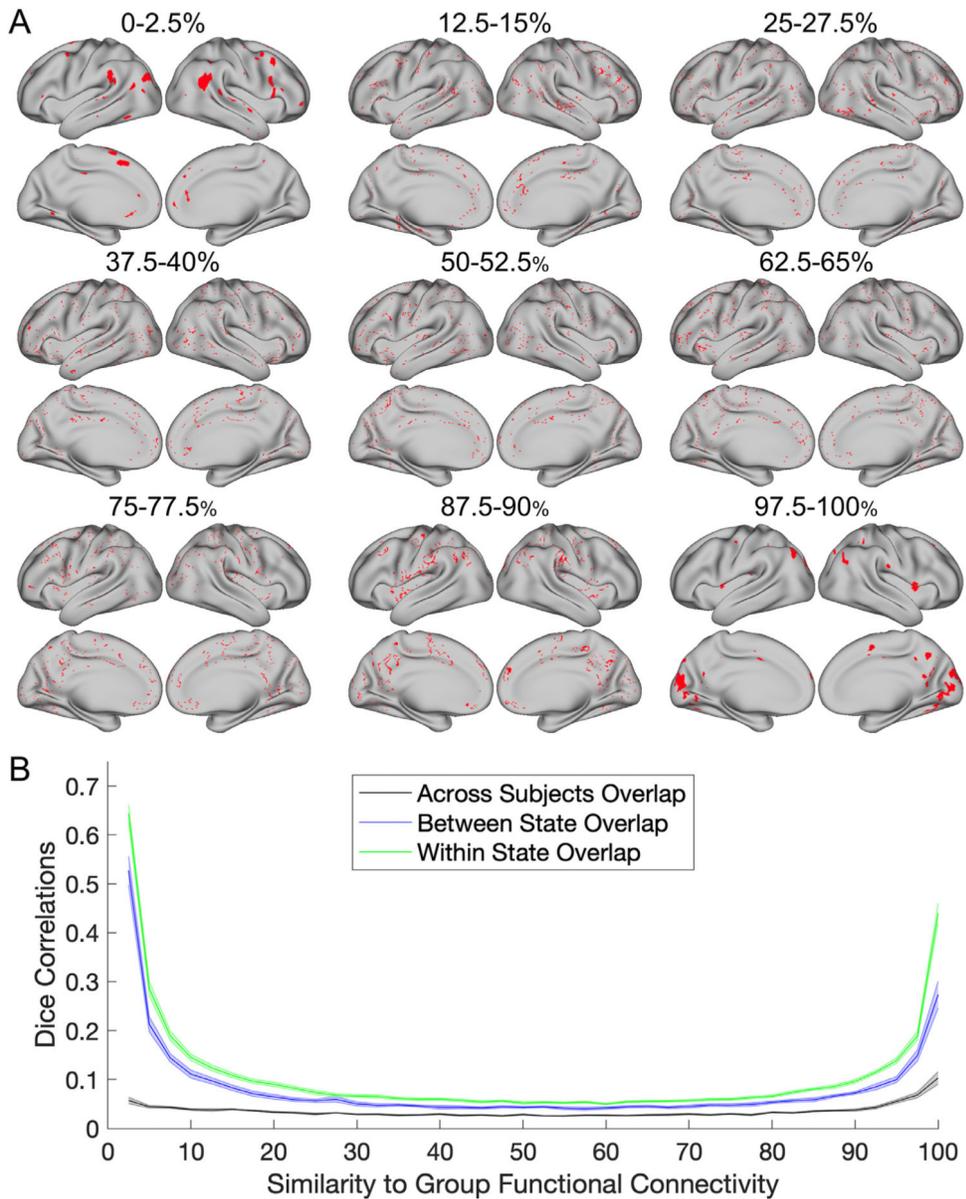
Fig. 6.
State/trait comparisons for locations binned by their similarity to the group (in 2.5% size bins). (A) The spatial distribution of the vertices in each bin are shown in red for MSC02. (B) At each bin, split-halves of the data were compared within state (green), between state (blue), and across subjects (black). Standard error bars are shown for each line. The areas which are the most similar and dissimilar to the group average functional connectivity show the most consistency (and thus exhibit a trait-like nature).
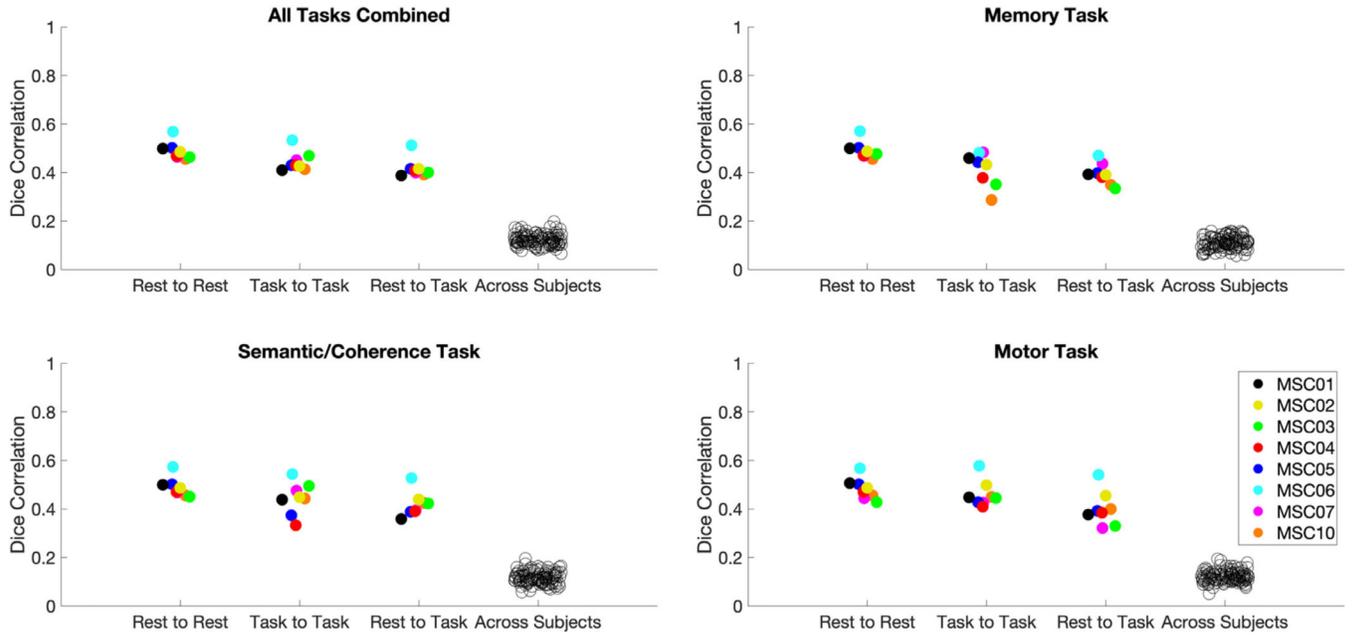
Fig. 7.
Comparisons for the similarity of network variant spatial locations in individual tasks relative to all tasks and rest. Dice correlations were used to compare network variant spatial locations using data from all 3 tasks combined (upper left), the memory task (upper right), the semantic/coherence tasks (bottom left), and the motor task (bottom right). Each comparison shows the rest-to-rest, task-to-task (for the relevant task), rest-to-task (for the relevant task), and across subject comparisons of network variants. All comparisons were matched on the amount of data within and across participants (see Table S1). Note that MSC09 was excluded from this analysis due to small amounts of high quality data in the motor task.