Quantifying the idiosyncrasy in brain activations

Johan Nakuci[1*], Jiwon Yeon[2,3], Kai Xue[1], Ji-Hyun Kim[4], Sung-Phil Kim[4] and Dobromir Rahnev[1]

[1]School of Psychology, Georgia Institute of Technology, Atlanta, Georgia, 30332, USA.

[2]Department of Psychology, Stanford University, Stanford, California, 94305, USA.

[3]Department of Ophthalmology, Stanford University, School of Medicine, Stanford, California, 94305, USA.

[4]Department of Biomedical Engineering, Ulsan National Institute of Science and Technology, Ulsan, South Korea.

* **Email**: jnakuci3@gatech.edu

Acknowledgments

This work was supported by the National Institutes of Health (grant R01MH119189 to DR) and the Office of Naval Research (grant N00014-20-1-2622 to DR).

Competing interests: Authors declare that they have no competing interests.

Author contributions:

Conceptualization: JN, DR

Methodology: JN, JY, KX, DR

Data Curation: JY, JHK, SPK

Visualization: JN, DR

Funding acquisition: DR

Writing – original draft: JN, DR

Writing – review & editing: JN, JY, KX, JHK, SPK, DR

Abstract

Research in neuroscience often assumes universal neural mechanisms, but increasing evidence points towards sizeable individual differences in brain activations. What remains unclear is the extent of the idiosyncrasy and whether different types of analyses are associated with different levels of idiosyncrasy. Here we compute the within-subject reliability and subject-to-group similarity of activations for task, reaction time (RT), and confidence at both the single trial and block levels. We find that maps of RT and confidence effects at the block level feature substantially lower subject-to-group similarity values compared to RT and confidence effects at the trials level, and that this effect occurs despite high within-subject reliabilities for both analyses. Critically, we develop a computational model that can quantify the relative strength of group- and subject-level factors in observed brain activations. The model suggests that subject-level factors carry similar weight as group-level factors for analyses at the trial level, but are over six times more important for RT and confidence analyses at the block level. These results show that different analyses are associated with different levels of idiosyncrasy and specifically suggest that behavioral changes that occur at the block level are markedly more idiosyncratic compared to other mental processes.

Highlights

- Brain activation among people differ.
- The extent of these differences at trial- and block-level analyses is unknown.
- Results indicate that block-level analyses are more idiosyncratic.
- These findings suggest the need to examine individual-subject results.

Keywords: individual differences, fMRI, perceptual decision making, block vs. event-related designs

1. Introduction

Human behavior is idiosyncratic: what elicits a certain behavior in one person is often very different from what elicits that same behavior in another (Eilam, 2015; Forkosh et al., 2019). Yet, research in neuroscience often tacitly assumes that behavior is homogeneous in the population and that the same neural correlates of a given behavior should emerge across individuals (Friston et al., 1999). This assumption is implicit in the common practice – enshrined in popular tools for functional MRI (fMRI) analyses such as SPM, AFNI and FSL – of performing and reporting second-level results as the true neural correlate of a given behavioral outcome.

However, increasing amount of evidence points towards the existence of substantial idiosyncrasy in brain activations, such that the same task can elicit different patterns of activity in different subjects (Miller et al., 2012, 2009; Seghier et al., 2008). Yet, it remains unclear how to precisely quantify the strength of the observed idiosyncrasy, as well as whether different types of analyses are associated with different levels of idiosyncrasy. To address these questions, here we compute the within-subject reliability and subject-to-group similarity of activations for task, reaction time (RT), and confidence at both the single trial and block levels. The idea is that the subject-to-group similarity can inform us about how different each person's activation map is from the group. However, this information has to be interpreted in the context of the noisiness of each individual map, which can be quantified by assessing its within-subject reliability. Finally, a computational model that takes both of these quantities into account can be used to assess the relative contribution of group- and subject-level factors to each activation map.

We collected data from a perceptual decision-making task inside an MRI scanner where subjects (N = 50) judged whether a briefly presented display featured more red or blue dots and provided a confidence rating (**Fig. 1A**). The experiment was organized in 96 blocks of 8 trials each (see Materials and Methods for full details). We performed trial-level analyses to assess the activation maps associated with task trials, as well as with RT

and confidence (by comparing trials with higher- vs. lower-than-median RT and confidence). In addition, we performed block-level analyses to assess the activation maps associated with task blocks, as well as with block-level RT and confidence (by comparing blocks with higher- vs. lower-than-median RT and confidence). To anticipate, we successfully quantified the level of idiosyncrasy in each of these maps and found that block-level activations are associated with substantially higher levels of idiosyncrasy than trial-level activations.

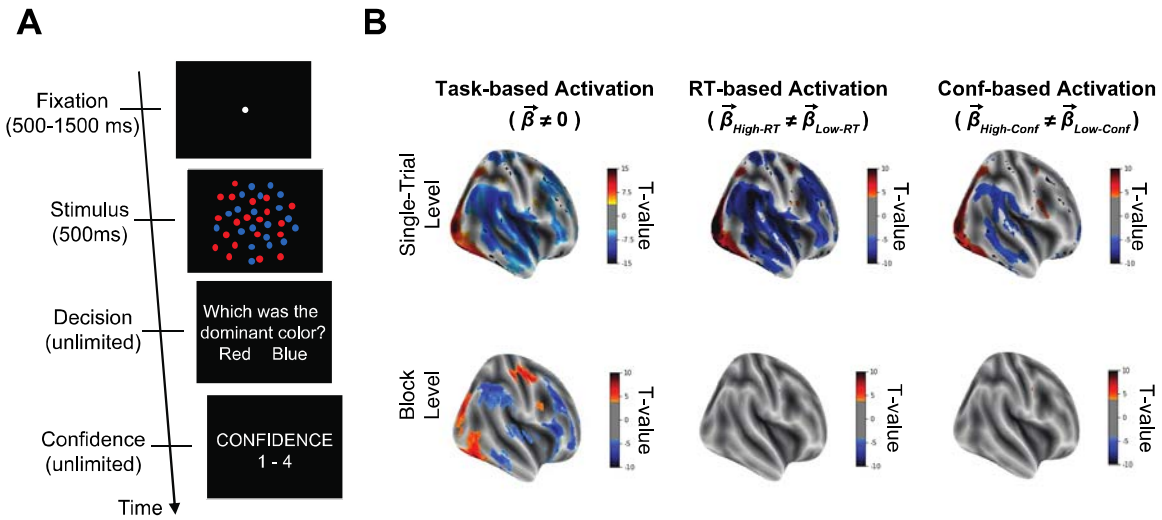

Figure 1. Task and results of standard group analyses. (A) Task. Subjects performed a simple perceptual decision-making task that required them to judge the dominant color in a display of colored dots and rate their confidence. (B) Results of standard second analyses for task-, RT-, and confidence-based contrasts. Trial-level analyses showed strong increases and decreases in activation across a range of brain regions for task- (top left), RT- (top middle) and confidence-based (top right) analyses. Conversely, while block-level analyses also showed strong increases and decreases in the task-based contrast (bottom left), they produced no significant activations for RT- (bottom middle) and confidence-based (bottom right) analyses. A common interpretation of these results is that block-level activations for high vs. low RT or confidence do not differ from each other. All maps thresholded at $p < 0.001$ uncorrected for display purposes.

2. Results

We first performed standard group fMRI analyses by conducting t-tests across all subjects for each voxel. We implemented these analyses both at the single trial and the block levels. At the trial level, we found that contrasts related to the task (Task >

Background), RT (Fast RT > Slow RT), and confidence (High confidence > Low confidence) all produced regions of strong activation and de-activation (**Fig. 1B**). However, a very different pattern emerged at the block level. At that level, we still observed strong activations and deactivations related to the task (Task > Background; **Fig 1B**) that bore a resemblance to the corresponding trial-level analyses. Yet, the block-level contrasts related to RT (Fast RT > Slow RT) and confidence (High confidence > Low confidence) produced very different results from the corresponding trial-level analyses. Specifically, both contrasts revealed no activations anywhere in the brain after whole-brain correction (**Fig. 1B, Fig. S1**).

These results raise the question about what is different about the RT and confidence block-level analyses. The customary interpretation of these null group results would be that brain activity does not differ for blocks with high vs. low average RT (or high vs. low confidence) in our study. However, examination of individual subject maps demonstrates that such conclusion would be misguided. To demonstrate that, we inspected the activations for task, RT, and confidence at both the trial and block levels in Subjects 1-3 (**Fig. 2**). We observed that at the trial level, all three subjects demonstrated very similar patterns of activation and deactivation across the brain for each of the task, RT, and confidence contrasts (**Fig. 2A**). At the block level, we observed a similar pattern for the task contrast with results that generally matched across the three subjects.

However, very different results emerged for the RT and confidence contrasts performed at the block level. These contrasts still revealed areas that consistently tracked RT and confidence for each individual subject, but the individual maps were highly dissimilar across subjects. For example, high-RT blocks were associated with stronger visual cortex activations in Subject 1 but weaker visual cortex activations in Subject 2, whereas the opposite pattern of results was observed in the somato-motor cortex of these two subjects. One may want to dismiss such subject-by-subject variability as being purely due to noise. However, examining the equivalent analyses for odd and even blocks only (see smaller brain maps in **Fig. 2**) reveals that each individual map is highly reliable and therefore the

stark differences across subjects in the block-level RT and confidence contrasts maps do not appear to be simply due to noise.

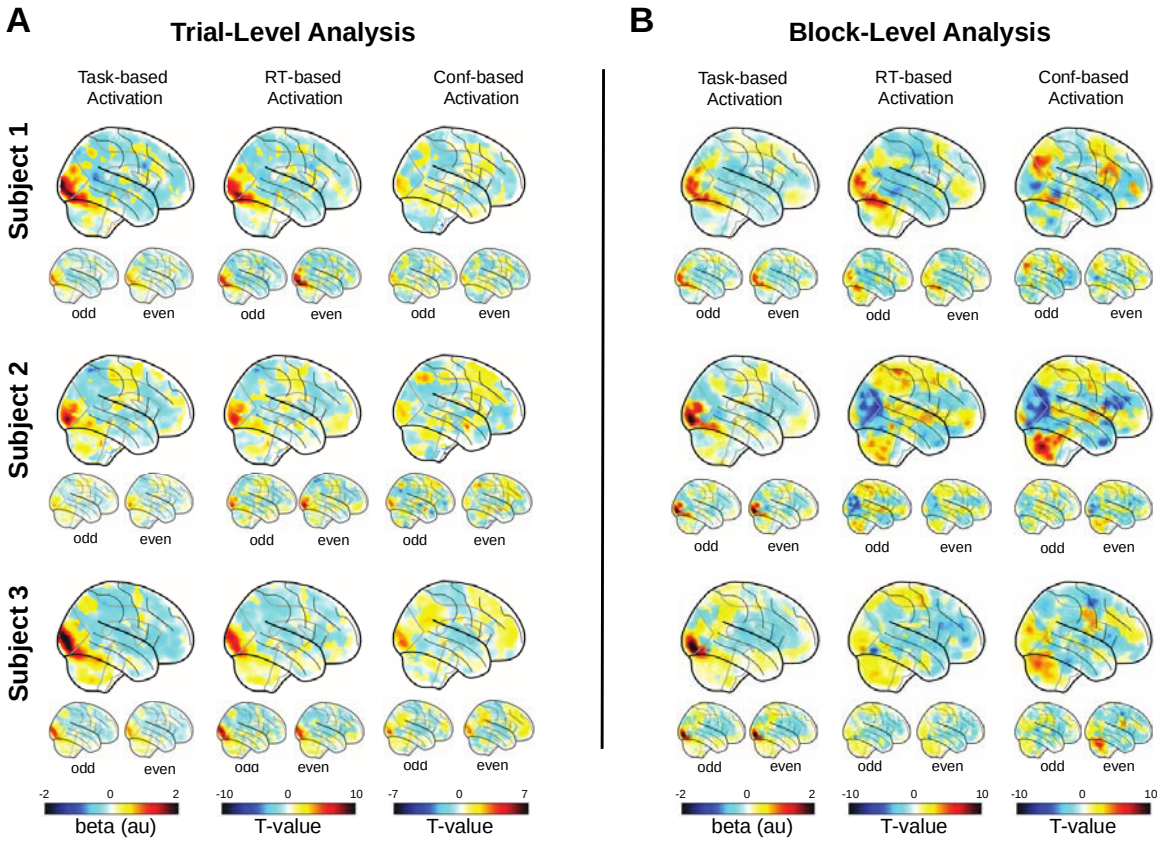

Figure 2. Trial- and block-level activations for task, RT, and confidence for three example subjects. The activation maps for task, RT, and confidence contrasts from the first three subjects are displayed for (A) trial- and (B) block-level analysis. Small brains underneath represent the same contrasts conducted only on odd or even trials or blocks. Similar activations for all three subjects appear for all trial-level contrasts, as well as for the task block-level contrast. However, the RT and confidence contrasts in the block-level analyses show a very different pattern with strong and consistent within-subject results that markedly vary across subjects.

To formally test these impressions, we first examined both the within-subject reliability and subject-to-group similarity of the whole-brain maps for the task, RT, and confidence contrasts at both the trial and block levels. We computed within-subject reliability by performing Pearson correlations between the activations obtained when examining only the odd or only the even trials or blocks. We computed subject-to-group similarity by

correlating each subject's brain map with the group map obtained by averaging the maps of the remaining 49 subjects.

As may be expected from Figure 2, we found high within-subject reliability and subject-to-group similarity in task activations at both the trial and block levels (**Fig. 3A**). Interestingly, the within-subject reliability was higher at the block level (average correlation between odd and even, $r_{trial} = 0.81$; $r_{block} = 0.97$; $p = 1.2 \times 10^{-17}$), whereas the subject-to-group similarity was higher at the trial level ($r_{trial} = 0.72$; $r_{block} = 0.46$; $p = 8.9 \times 10^{-13}$). One potential concern with these types of analyses could be that they may be biased by voxels that are either particularly noisy or not involved in the task in any way. Therefore, to test whether these results are robust, we repeated them by first selecting the top 75, 50, 25, or 10% of the most strongly activated voxels for each subject (see Methods). These analyses showed that selecting smaller percentages of the most highly activated voxels generally increased both the within-subject reliability and subject-to-group similarity (all p's < .05), but the pattern of results remained essentially unchanged.

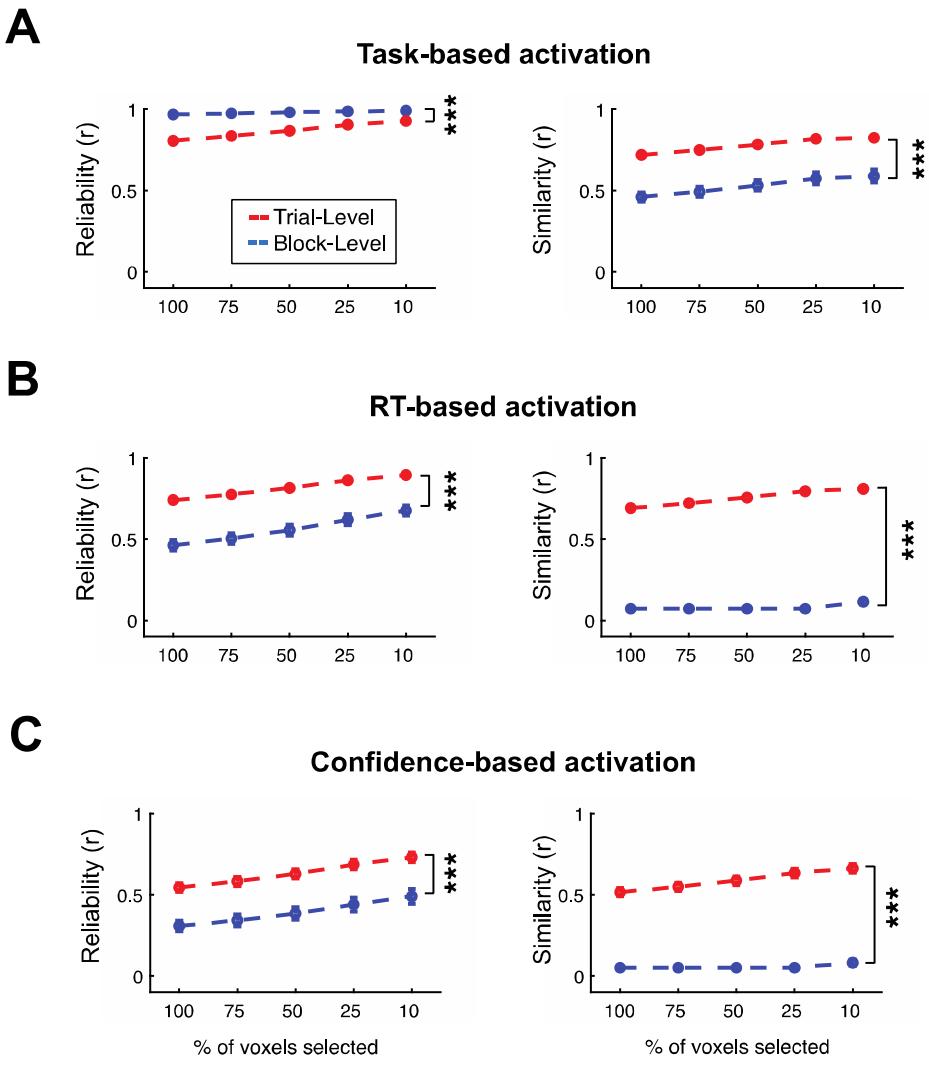

Figure 3. Within-subject reliability and subject-to-group similarity for each map. Within-subject reliability and subject-to-group similarity values as a function of the percent of most activated voxels selected for trial- and block-level (A) task- (B) RT- and (C) confidence-based activation. The within-subject reliability was highest for block-level task-based activation. For all voxel selection levels, the subject-to-group was highest for trial-level task-, RT, and Confidence-based activation. Subject-to-group similarity is computed as the average similarity between the maps of each person and the group map of the remaining subject. Error bars show SEM; ***, $p < 0.001$. All p-values are Bonferroni corrected.

Critically, we examined the within-subject reliability and subject-to-group similarity for the RT and confidence maps. Echoing the qualitative impressions from **Fig. 2**, we found high within-subject reliability for all RT and confidence maps, though the reliability values were slightly higher at the trial level for both RT ($r_{trial} = 0.74$; $r_{block} = 0.46$; $p = 2.6$

x 10^{-10}; **Fig. 3B**) and confidence (r_{trial} = 0.55; r_{block} = 0.31; p = 3.2 x 10^{-7}; **Fig. 3C**). However, we found vast differences in subject-to-group similarity at the trial and block levels. Specifically, subject-to-group similarity values were relatively high at the trial level for both RT (r_{trial} = 0.69) and confidence (r_{trial} = 0.52), but became extremely small at the block level (RT: r_{block} = 0.07; confidence: r_{block} = 0.05), with the differences between trial and block level similarity values being highly significant (RT: p = 4.2 x 10^{-35}; confidence: p = 3.1 x 10^{-21}). Again, equivalent results were obtained when the top 75, 50, 25, or 10% of voxels were considered. These results demonstrate that RT and confidence maps at the block level are strongly idiosyncratic because the differences across subjects cannot be attributed to overly low within-subject reliability of these maps.

To gain further intuition for the underlying effects, we conducted two additional analyses. First, we tested the consistency of the sign of voxel activations (whether they were positive or negative) across subjects. We found that for the task maps, there were large portions of the brain that showed consistently positive or consistently negative activations at both the trial and block levels (**Fig. 4A left**). Indeed, the maximal overlap across subjects at the trial level was 100% for both positive and negative activations, whereas at the block level the maximum overlap was 96% and 94% for positive and negative activations. On the other hand, the trial and block levels produced very different consistency maps for RT and confidence (**Fig. 4A middle and right**). At the trial level, we found many areas of strong consistency with maximal overlap of 100% and 98% for positive and negative activations in the RT maps, and maximal overlap of 100% and 88% for positive and negative activations in the confidence maps. However, at the block level, the consistency across subjects was substantially lower with maximal overlap of 82% and 80% for positive and negative activations in the RT maps, and maximal overlap of 76% and 82% for positive and negative activations in the confidence maps. (Note that the expected values in random data are 80% for positive and negative activations, so the consistency scores for RT and confidence at the block level are close to what would be expected by chance.) Overall, consistency values for the RT and confidence maps at the

trial level were much higher than the corresponding maps at the block level (all p's < 10^{-125}).

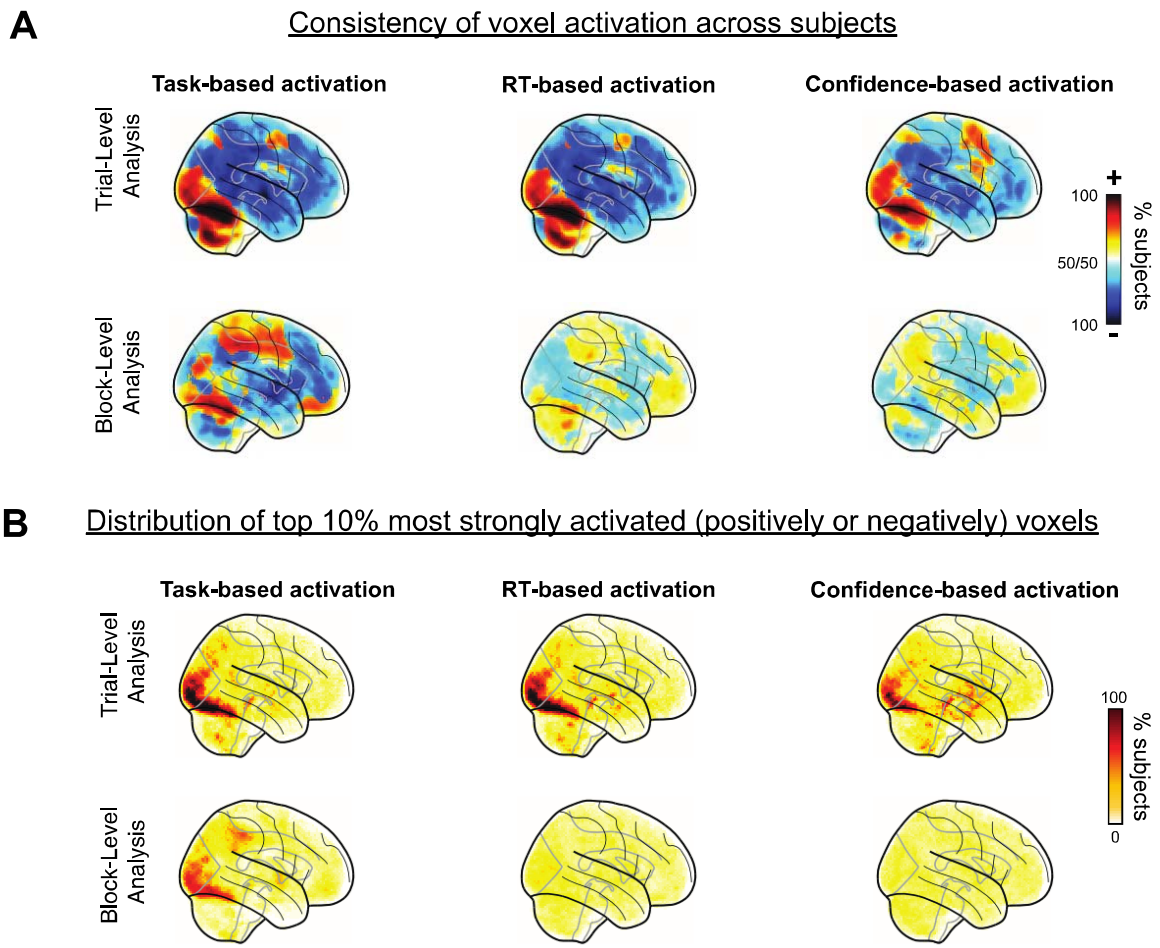

Figure 4. Maps of the activation consistency and distribution of the top-10% most activated voxels across subjects. (A) Maps of voxel consistency computed as the proportion of subjects showing a positive or negative relationship between voxel activity and behavior. Task activations maps, as well as RT and confidence maps at the trial level show a high level of consistency. Conversely, RT and confidence maps at the block level show low levels of consistency. (B) Maps of the distribution of the top-10% most activated voxels. Again, the RT and confidence analyses at the block level show the smallest level of consistency.

Second, we examined the distribution of the locations of the top-10% most strongly activated voxels for each subject (both positive and negative activations were considered). We found that at both the trial and block levels, the most strongly activated

voxels clustered in the occipital and parietal lobes (**Fig. 4B left**). The maximum overlap among the 10% most activated voxels across subjects was 98% at the trial and 76% at the block levels. However, there was again a stark difference between RT and confidence results at the trial vs. the block levels (**Fig. 4B middle and right**). At the trial level, there were again areas of strong clustering of the most activated voxels (mostly in the occipital lobe) for both RT and confidence maps (maximal overlap: 98% and 78%, respectively). Yet, at the block level, there was very little clustering for either the RT or confidence maps (maximal overlap: 32% and 30%, respectively). Note that these values were only slightly higher than the maximum expected value in random data (28%). Altogether, both additional analyses further underscore the very high level of idiosyncrasy in RT and confidence maps at the block level, whereas less idiosyncrasy occurs for RT and confidence maps at the trial level as well as task maps at either the trial or block levels.

The block-level maps of Subjects 1-3 (**Fig. 2**) suggest that the low subject-to-group similarity in RT and confidence maps at the block level are likely due to large-scale, rather than fine-grained, differences in the activation maps. To confirm this impression, we repeated the same analyses above with a wide range of smoothing levels (from 5 to 20 mm) and obtained very similar results (**Fig. S2-7**; note that the results above come from 10 mm smoothing). Further, rather than performing the correlations on a voxel-per-voxel level, we did so on the level of 200 large regions of interest obtained from the Schaefer atlas (Schaefer et al., 2018) and still obtained the same results (**Fig. S8**). These findings clearly indicate that the low subject-to-group similarity values in RT and confidence maps at the block level are due to genuine, large-scale differences in the maps rather than issues of misalignment of individual brain maps (Haxby et al., 2020, 2011; Nieto-Castañón and Fedorenko, 2012).

Having established the existence of a stark difference between the level of idiosyncrasy for different types of analyses, we sought to precisely quantify the level of idiosyncrasy by building a simple computational model that jointly generates behavior and brain activity maps. The critical idea behind the model is to separately model group-level

factors (i.e., factors that are identical for all subjects) and subject-level factors (i.e., factors that are different for each subject). Specifically, the model produces activation maps for each individual trial or block based on three group-level factors (group task map, group RT map, and group confidence map), three subject-level factors (subject-specific task map, subject-specific RT map, and subject-specific confidence map), and one noise factor (**Fig. 5**). To keep the model simple, both behavior and individual voxel activation for each group- and task-level factor were generated randomly by ignoring known temporal and inter-regional dependencies. The weight of the noise factor was fixed to 1, leaving the model with a total of six free parameters (one for the weight of each group- and subject-level factor). We then fit the model to the observed within-subject reliability and subject-to-group consistency values computed using 100% of the voxels. Despite its simplicity, the model was able to provide excellent fit to the data from **Fig. 3** by capturing closely the observed patterns of within-subject reliability (**Fig. 6A-C, left** and subject-to-group similarity (**Fig. 6A-C, right**) at both the trial and block levels.

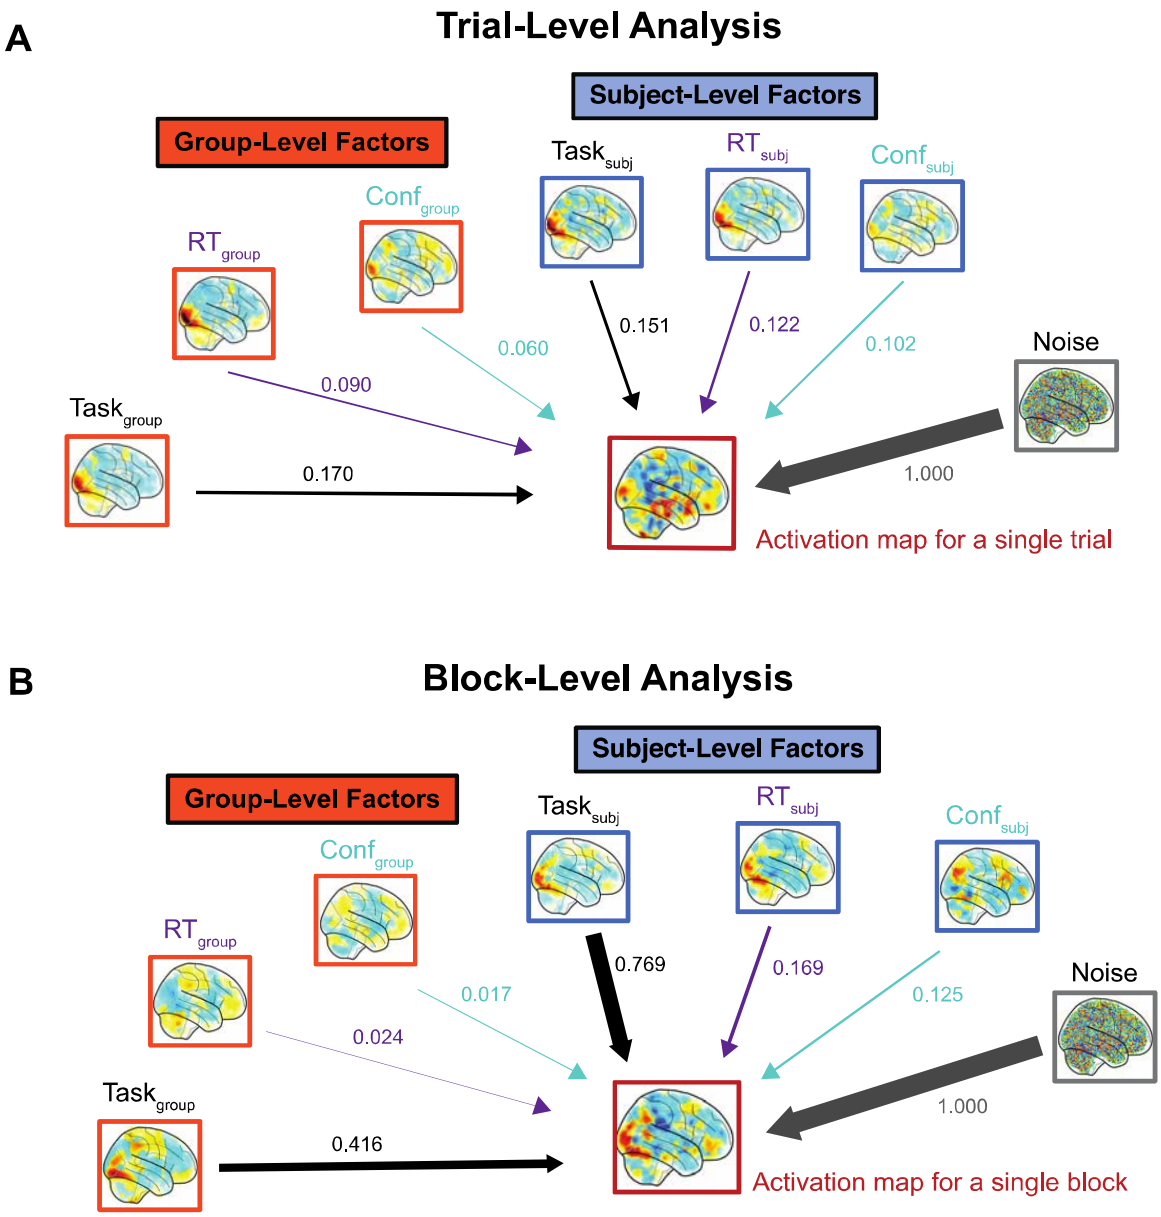

Figure 5. Model structure. A graphical depiction of the model at the (A) trial and (B) block level. The activation map for an individual block or trial is modeled as the confluence of three group-level, three subject-level, and one noise factor. The thickness of the arrows and associated numbers correspond to the weights obtained from fitting the model to the data. The brain maps displayed in the figure were produced as follows: the group brain maps show the actual maps obtained in our data, the subject-level brain maps show the maps for subject 1, and the activation map at the bottom shows the activation map for trial 1 or block 1 of subject 1.

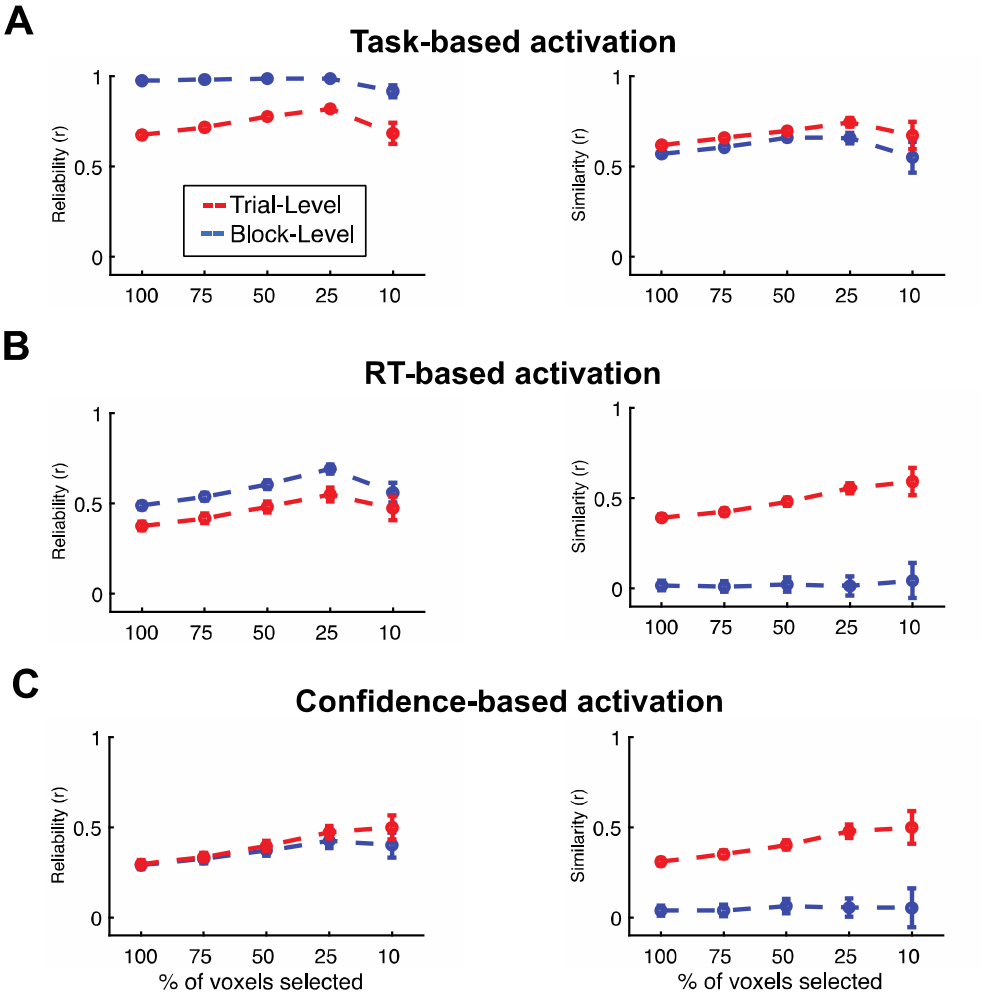

Figure 6. Model fits. Model fits to the within-subject reliability and subject-to-group similarity values at the trial and block levels for (A) task, (B) RT, and (C) confidence analyses. The model was fit only to the data where 100% of voxels selected. Despite its simplicity, the model is able to reproduce the empirical data for the remaining analyses with smaller percentages of selected voxels (see Fig. 3) very well.

Critically, the model allowed us to examine the weights of the group- and subject-level factors, thus providing insight into the relative contribution of each. We found comparable weights for the subject- and group-level task factors at both the trial level (subject-level factor weight = 0.151, group-level factor weight = 0.170, ratio = 0.88) and the block level (subject-level factor weight = 0.769, group-level factor weight = 0.416, ratio = 1.84; **Fig. 7A**). Thus, the subject-level factors were only slightly higher than the group-level factors, pointing to a balance between influences that are common across all

subjects and influences that are specific to each individual. We observed slightly higher relative weights for the subject-level factors for the RT and confidence maps at the trial level (RT: subject-level factor weight = 0.122, group-level factor weight = 0.09, ratio = 1.35; Confidence: subject-level factor weight = 0.102, group-level factor weight = 0.06, ratio = 1.7). Critically, substantially higher subject-level factors (indicative of high levels of idiosyncrasy) emerged for both the RT and confidence analyses at the block level (RT: subject-level factor weight = 0.169, group-level factor weight = 0.024, ratio = 7.0; Confidence: subject-level weight = 0.125, group-level weight = 0.017, ratio = 7.4). In other words, our model suggests that group- and subject-level factors have relatively similar influence on task activation maps, which corresponds well to recent findings about group- and subject-level influences on brain connectivity (Gratton et al., 2018). However, the model reveals that the influence of subject-level factors is about seven times larger than the influence of group-level factors for RT and confidence maps at the block level, showing a substantial amount of idiosyncrasy for those analyses (**Fig. 7B**).

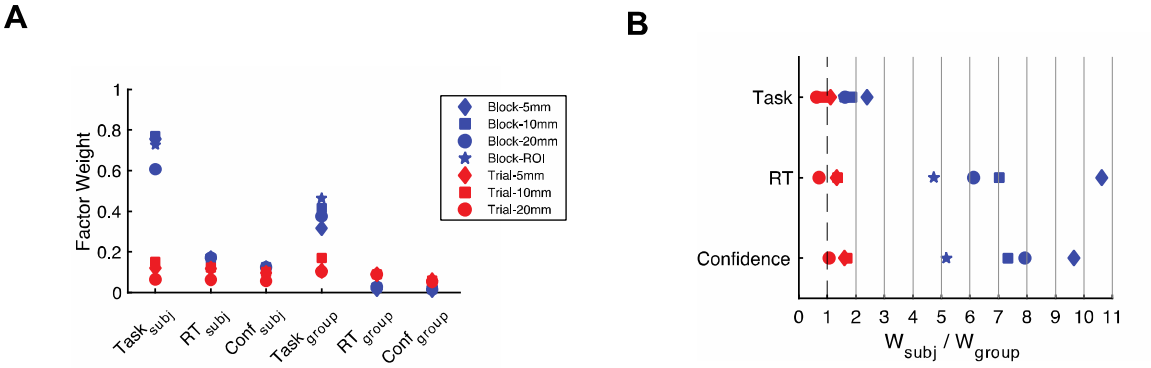

Figure 7. Model weights and ratios. A) Model weights. Subject- and group-level weights obtained from fitting the model separately to each analysis. The subject-level weights were consistently higher than the group-level weights across all analyses in the block-level analysis. B) Weight ratios. Relative weights of the subject-level and corresponding group-level factors from each analysis. The analyses at the block level produced uniformly higher ratios (demonstrating larger levels of idiosyncrasy) and this effect was especially pronounced for the RT and confidence analyses.

To examine the robustness of the modeling results, we repeated the model fitting on data with 5 to 20 mm smoothing, as well as on the data from the 200 Schaefer atlas ROIs. These three additional analyses produced similar results: the weights ratio between the

subject- and group-level factors was between 0.6 and 2.4 for the task factors in all cases, between 0.7 and 1.4 for RT at the trial level, and between 1.1 and 1.7 for confidence at the trial level. However, much higher ratios – indicating high levels of idiosyncrasy – emerged at the block level for both RT (factor weight ratios 4.7 and 10.6) and confidence (factor weight ratios 5.1 and 9.6). As can be seen in Fig. 7B, higher levels of smoothing generally decreased the ratio of subject- to group-level factor weights, but generally preserved the large discrepancy between trial- and block-level analyses for RT and confidence.

3. Discussion

A major goal of neuroscience research has been to understand the neural correlates of behavior. Behavior is a complex phenomenon that is often idiosyncratic to a person (Eilam, 2015; Forkosh et al., 2019). Idiosyncratic behavior is ubiquitous in social situations (Durlauf, 2001), economic decisions (Kable and Glimcher, 2007), judgments of beauty (Martinez et al., 2020), confidence ratings (Navajas et al., 2017), response bias (Rahnev, 2021), and even low-level perception (Afraz et al., 2010). However, an implicit assumption in much of neuroscience research is that the neural correlates of behavior are the same across individuals (Friston et al., 1999). Here, we test and reject this assumption. We find that RT and confidence maps computed at the block level are very consistent within an individual, and yet remarkably different across subjects. On the other hand, task, RT, and confidence maps at the trial level are both consistent within an individual and similar across subjects. A computational model explains these results and suggests that for analyses at the trial level, the influence of subject-level factors is only slightly stronger than the influence of group-level factors, whereas for RT and confidence analyses at the block level, the influence of subject-level factors is about seven times larger than the influence of group-level factors.

There are at least two important conclusions that one can draw from the current results. First, across all analyses, subject-level factors were at least as important as group-level factors. While this effect could be at least partly driven by issues such as misalignment

across different brains, the results were remarkably stable whether they were computed using 5-, 10-, or 20-mm smoothing, or even at the level of 200 ROIs. If brain misalignment were the main source of the observed idiosyncrasy here, one would expect that larger smoothing or analyses at the levels of large ROIs would produce different results. Therefore, it appears that many fMRI maps are likely substantially influenced by idiosyncratic, subject-level factors. Second, some internal processes are associated with a greater degree of idiosyncrasy than others. Specifically, RT and confidence contrasts computed at the block level are substantially more idiosyncratic than similar contrasts computed at the trial level. Our analyses here do not permit us to paint a more comprehensive picture of what types of analyses are likely to be more or less idiosyncratic, but they do demonstrate that the brain correlates of different mental processes have widely varying degrees of idiosyncrasy.

Our results have implications about the common way of reporting the result of fMRI studies. The majority of papers in the field only report second-level, group maps (Bandettini and Ungerleider, 2001; Belliveau et al., 1991; le Bihan et al., 1993; Poldrack, 2011). The implicit assumption is that the group map represents the true "neural correlates" of a given behavior. However, our results demonstrate that in some cases such group-level results only have a small relationship to the neural correlates for any actual person (i.e., they explain just 1% of the variance in single-subject maps). This is not to say that second-level maps are somehow wrong – instead, they represent our best approach for uncovering what is common across all subjects. The issue is that focusing exclusively on the relatively weak commonality across people can distract from the idiosyncratic effects within each individual, which can sometimes be much stronger than the group-level effects. Our results thus highlight the need for a renewed focus on investigating the brain-behavior relationship at the level of single subjects (Gilmore et al., 2021; Gordon and Nelson, 2021; Naselaris et al., 2021; Song and Rosenberg, 2021).

Our findings also suggest novel ways for finding robust biomarkers for various mental disorders (Elliott et al., 2018; Kaufmann et al., 2017; Li et al., 2020; Parkes et al., 2020).

Most research in the field has focused on biomarkers unrelated to behavior such as functional connectivity patterns at rest(Drysdale et al., 2017; Woodward and Cascio, 2015). An exciting possibility is that subject-level activations maps for disease-relevant behaviors could serve as much more powerful biomarkers because of their high reliability and clear differences among people. Focusing directly on the relationship between one's behavior and one's brain activations may help to delineate the intricate relationship between the brain and psychopathology (Gratton et al., 2020). For example, one potential implication of our results is that internal processes related to mental illness that change at a slower rate (and thus may be more similar to our RT and confidence analyses at the block level) might have neural correlates that are largely unique to an individual. For example, in one individual, fluctuation in positive mood or thoughts might depend on activity in the frontal cortex, but in another individual, they could depend on the parietal cortex. Similar effects have already been suggested in the context of pain perception (Kohoutová et al., 2022). Therefore, subject-level effects could be crucial to diagnosing and treating different mental illnesses. Additionally, an analysis that is focused on subject-level variability might be more informative since between-subject analyses ignore the large degree of within-subject variability (Fisher et al., 2018; Lebreton et al., 2019).

Our model represents one of relatively rare attempts to build process models that jointly generate behavior and brain activations. There are rich literatures of building process model in cognitive science that focuses exclusively on behavior (Ratcliff, 1978; Rescorla and Wagner, 1972; Zhang and Luck, 2008) and in cognitive neuroscience that focuses exclusively on brain activity (Breakspear, 2017; Coombes et al., 2007; Wilson and Cowan, 1972). Yet, understanding the brain-behavior relationship clearly requires the development of process models that jointly generate both types of data, thus explicitly clarifying the links between the two. Future work in the field should increasingly emphasize process models that specify the mechanisms that generate behavior and brain activity.

In conclusion, we find a stark level of idiosyncrasy in RT and confidence analyses at the block level such that single-subject maps are remarkably reliable, yet very different across subjects. These results have strong implications about the common practice of only reporting second-level analyses and suggest the need to more frequently examine individual-subject results.

4. Materials and Methods

4.1 Subjects

Fifty-two healthy subjects were recruited for this study. Two subjects were excluded because one had metal braces in their teeth and one decided to stop the experiment after the second run. All analyses were thus based on the remaining 50 subjects (25 females; Mean age = 26; Age range = 19-40; Compensated 20,000 KRW or approximately 18 USD). All subjects were screened for any history of neurological disorders or MRI contraindications. The study was approved by Ulsan National Institute of Science and Technology Review Board (UNISTIRB-20-30-C) and all subjects gave written consent.

4.2 Task

Subjects had to determine which set of colored dots (red vs. blue) is more frequent in a cloud of dots (Fig. 1A). Each trial began with a white fixation dot presented for a variable amount of time between 500-1500 ms at the center of the screen on a black background. Then, the stimulus was shown for 500 ms, followed by untimed decision and confidence screens. The stimulus consisted of between 140 and 190 red- and blue-colored dots (dot size = 5 pixels) dispersed randomly inside an imaginary circle with a radius of 3° from the center of the screen. Four different dot ratios were used – 80/60, 80/70, 100/80, and 100/90, where the two numbers indicate the number of dots from each color. The experiment was organized in blocks of 8 trials each, with each dot ratio presented twice in a random order within a block. The more frequent color was pseudo randomized so that there were equal number of trials where red and blue were the correct answer within a run (consisting of 16 blocks). Subjects used an MRI-compatible button box with their right hand to indicate their decision and confidence responses. For the decision response, the index finger was used to indicate a “red” response and the middle finger for a “blue” response. Confidence was given on a 4-point scale, where 1 is the lowest and 4 is the highest, with the rating of 1 mapped to the index finger and the rating of 4 mapped to the little finger.

Subjects performed 6 runs each consisting of 16 blocks of 8 trials (for a total of 768 trials per subject). Three subjects completed only half of the 6th run and another three subjects completed only the first 5 runs due to time constraints. The remaining 44 subjects completed the full 6 runs. Subjects were given 5 seconds of rest between blocks, and self-paced breaks between runs.

4.3 MRI recording

The MRI data was collected on a 64-channel head coil 3T MRI system (Magnetom Prisma; Siemens). Whole-brain functional data were acquired using a T2*-weighted multi-band accelerated imaging (FoV = 200 mm; TR = 2000 ms; TE = 35 ms; multiband acceleration factor = 3; in-plane acceleration factor = 2; 72 interleaved slices; flip angle = 90°; voxel size = 2.0 x 2.0 x 2.0 mm^3). High-resolution anatomical MP-RAGE data were acquired using T1-weighted imaging (FoV = 256 mm; TR = 2300 ms; TE = 2.28 ms; 192 slices; flip angle = 8°; voxel size = 1.0 x 1.0 x 1.0 mm^3).

4.4 MRI preprocessing and general linear model fitting

MRI data were preprocessed with SPM12 (Wellcome Department of Imaging Neuroscience, London, UK). We first converted the images from DICOM to NIFTI and removed the first three volumes to allow for scanner equilibration. We then preprocessed with the following steps: de-spiking, slice-timing correction, realignment, segmentation, coregistration, normalization, and spatial smoothing with 10 mm full width half maximum (FWHM) Gaussian kernel. In control analyses, we used 5 and 20 mm FWHM smoothing to investigate whether the results are due to fine-grained differences in the activations maps between subjects, given that local differences would be substantially reduced by larger smoothing kernels. Despiking was done using the 3dDespike function in AFNI. The preprocessing of the T1-weighted structural images involved skull-removal, normalization into MNI anatomical standard space, and segmentation into gray matter, white matter, and cerebral spinal fluid, soft tissues, and air and background.

We fit a general linear model (GLM) that allowed us to estimate the beta values for each voxel in the brain for trial-level and block-level analysis of the experiment. For trial-level analysis, the model consisted of separate regressors for trials that had greater or smaller than the median RT or confidence, inter-block rest periods, as well as linear and squared regressors for six head movement (three translation and three rotation), five tissue-related regressors (gray matter, white matter, cerebrospinal fluid, soft tissues, and air and background), and a constant term per run. For the block-analysis, the model consisted of regressors for each individual block, inter-block rest periods, as well as linear and squared regressors for six head movement (three translation and three rotation), five tissue-related regressors (gray matter, white matter, cerebrospinal fluid, soft tissues, and air and background), and a constant term per run.

4.5 Standard group-level analyses

We first performed a standard group analysis by conducting t-tests across all subjects for each voxel. A task-based analysis compared the obtained beta values with zero to identify regions of activation and de-activation. Two behavior-based analyses compared the beta values for trials or blocks with higher- vs. lower-than-median average reaction times (RT) and higher- vs. lower-than-median average confidence. Significance was assessed using $p < 0.05$ after Bonferroni correction for multiple comparisons. For display purposes, Fig. 1 and Fig. S1 used the more liberal threshold of $p < 0.001$ uncorrected.

4.5 Within-subject reliability analyses

We examined the within-subject reliability of the whole-brain maps produced by the task, RT, and confidence analyses. To do so, we first re-did each analysis by only using the odd trials or blocks, as well as by only using the even trials or blocks. We then compared the similarity between the maps obtained for odd and even trials and blocks using Pearson correlation. We performed the analysis five times based on the top 10, 25, 50, 75, or 100% of most strongly activated voxels in the following way. We first identified the X% most strongly activated voxels (i.e., the voxels with highest absolute activation values) when only examining the data from the odd trials or blocks. The activation values

used were the average beta value for task-based analyses, and the t-value (obtained by using a t-test to compare the beta values for trials or blocks with above- vs. below-median RT or confidence) for the RT and confidence analyses. This selection procedure ensured that both positively and negatively activated voxels were selected and that an equal number of voxels were selected each time. The activations in the selected top X% of voxels from the odd trials or blocks were then correlated with the activations in the same voxels in the even trials or blocks, thus obtaining an "odd-to-even" correlation value. Then, using an equivalent procedure, we identified the top X% of most activated voxels in the even trials or blocks, and correlated their activations with the activations in the corresponding voxels in the odd trials or blocks, thus obtaining an "even-to-odd" correlation value. Finally, we computed the overall within-subject reliability as the average of the odd-to-even and even-to-odd correlation values.

4.6 Subject-to-group similarity analyses

Critically, we examined the subject-to-group similarity in the maps produced by each analysis. For each subject, we correlated their individual task-, RT-, and confidence-based activation maps with the corresponding group map obtained by averaging the maps of the remaining 49 subjects. Similar to the within-subject reliability analyses, we conducted these analyses separately for the top 10, 25, 50, 75, or 100% of most activated voxels. These voxels were selected in the same way as for the within-subject reliability analyses using all of the data in a given subject; the activations in the voxels identified for a given subject were then correlated with the average activations in the same voxels for the remaining 49 subjects.

4.7 Consistency in activation analysis

As another test of the across-subject similarity, we computed the consistency in the sign of activation. Our main analyses relied on taking correlations, but it is possible that just considering the sign of activation (rather than the strength of activation) would produce different results. To investigate this possibility, we examined the consistency of the sign of voxel activations (positive or negative) across subjects. To do so, we first set all voxels

values that were equal to zero to not-a-number value (NaN). This applied to regions that are outside the brain. We then binarized the voxel activation values $activation_i$ such that:

$$binary_i = \begin{cases} 1, & activation_i \geq 0 \\ 0, & activation_i < 0 \end{cases}$$

The consistency of the sign of a voxel's activation across subjects (C_i) was then calculated as percentage of subjects for which a voxel i was positively or negatively activated using the formula:

$$C_i = 100 \times \frac{1}{50} \sum_{i=1}^{50} binary_i$$

As defined, C_i goes from 0 (all subjects having negative activation for that voxel) to 100 (all subjects having positive activations for that voxel), with a value of 50 indicating that half of the subjects had positive and half had negative activation. However, when reporting the values of C_i, we flipped values under 50 using the formula $C_{i,flipped} = 100 - C_i$, so that these values represent the percent of subjects with negative activations. The analysis was performed separately for the trial- and block-level task-based activation maps, RT-based activation maps, and confidence-based activation maps. The activation values were the average beta value (for task-based analyses) or t-value (for RT and confidence analyses).

Low across-subject similarity in these analyses would result in most voxels having consistency, C_i, values close to 50 (corresponding to the voxel activation having positive sign in half the subjects and negative sign in the other half). However, due to chance, the consistency values are bound to sometimes be higher. Therefore, to enable the appropriate interpretation of the obtained results, we computed the expected consistency values in the maps of 50 subjects whose maps have no relationship to each other. Specifically, we generated a random set of voxel activation values for each of 50 sample

subjects. Maximal consistency from the random data was calculated in the same manner as the empirical values and the procedure was repeated 1000 times. This analysis revealed that completely random data would produce a maximal consistency of 80% (for both positive and negative activations) given the number of voxels and number of subjects that we had, which was close to the empirically observed values for RT and confidence analyses.

4.8 Distribution of top-10% most strongly activated voxels

As a final test of the across-subject similarity for the different maps, we sought to identify the consistency of the location of the most strongly activated brain regions across subjects. For each subject, we selected the top-10% most strongly activated voxels by considering the absolute value of either the average beta value (for task-based analyses) or t-value (for RT and confidence analyses). Note that this procedure selected positive and negative activations. We then estimated, for each voxel, the percent of subjects for which the voxel was selected as one of the top-10% most strongly activated voxels. As before, the analysis was performed separately at the trial and block levels for task-based activation maps, RT-based activation maps, and confidence-based activation maps.

Low across-subject similarity in these analyses would result in most voxels being selected about 10% of the time. However, due to chance, some voxels are bound to be selected more than 10% of the time. Therefore, to enable the appropriate interpretation of the obtained results, we computed the expected level of maximal overlap in the maps of 50 subjects whose maps have no relationship to each other. Specifically, for each of the 50 subjects, we generated a random set voxel activation values. We then selected the top-10% of the highest absolute values from each subject and calculated the overlap across subjects. The expected value from random data was computed as the average maximal overlap after 1000 iterations. This analysis revealed that completely random data would produce a maximal overlap of 28% given the number of voxels and number of subjects that we had, which was only a little less than the empirically observed values for RT and

confidence analyses (32% for RT-based analyses and 30% for confidence-based analyses).

4.9 ROI-based within-subject reliability and subject-to-group similarity analyses

All of the above analyses were performed at the level of individual voxels. However, to ensure that the results obtained were not due to the fine-grained misalignment of individual maps, we performed the within-subject reliability and subject-to-group similarity analyses on the level of 200 large regions of interest (ROIs) obtained from the Schaefer atlas (Schaefer et al., 2018). For these analyses, we averaged all the beta values within an ROI and repeated both the within-subject reliability and subject-to-group similarity analyses in the same way as for the voxel-based analyses above. These ROI-based analyses produced very similar results to the main voxel-based analyses (see Fig. S7), indicating that the low subject-to-group similarity in RT and confidence analyses at the block level is due to genuine, large-scale differences in the maps rather than issues of misalignment of individual brain maps.

4.10 Model specification

Our analyses revealed that task-based activation maps at both the trial and block levels are largely consistent across subjects, whereas RT and confidence maps specifically at the block level are largely idiosyncratic. We sought to precisely quantify the contributions of group- and subject-level factors in all activations maps by building and fitting a simple computational model.

The model jointly generates behavior and brain activity maps using minimal assumptions. Specifically, RT and confidence for a trial or block of trials were generated randomly from a standard normal distribution (with a mean of 0 and standard deviation of 1). Critically, the model assumes that the activation map for each trial or block is a function of seven different factors. The first three are group-level factors (i.e., factors common to all subjects) for the task itself, the influence of the trial or block-specific RT, and the influence of the trial or block-specific confidence. The next three factors are subject-level

factors (i.e., factors specific to each subject) for the task itself, the influence of the trial or block-specific RT, and the influence of the trial or block-specific confidence. Finally, the 7th factor is simply Gaussian noise. Critically, each factor is weighed by a corresponding factor weight that determines the strength of influence of that factor to the final voxel activation values, such that the activation strength (β) for a given voxel on a given trial or block is:

$$\beta = w_{task_{group}} * f_{task_{group}} + w_{rt_{group}} * f_{rt_{group}} * RT + w_{conf_{group}} * f_{conf_{group}} * conf + w_{task_{subj}} * f_{task_{subj}} + w_{rt_{subj}} * f_{rt_{subj}} * RT + w_{conf_{subj}} * f_{conf_{subj}} * conf + w_{noise} * f_{noise}$$

where RT and $conf$ are the reaction time and confidence on that trial or block, the w's are the weights associated with each factor, and the f's are the factors that influence the voxel activity for a given trial or block. Without loss of generality, the weight of the noise factor (w_{noise}) was fixed to 1. The value of each factor f was randomly sampled from a standard normal distribution such that group-level factors were randomly sampled for each voxel, subject-level factors were randomly sampled for each voxel and subject, and the noise factor was randomly sampled for each voxel, subject, and trial or block.

We note that, in the model, both the behavioral measures (RT and $conf$) and the factors (f) that control the individual voxel activations were sampled from a standard normal distribution, which makes the simulated values of RT, confidence, and voxel activations (β's) not match their observed values. Additional parameters could be easily used to match the actual distributions of the empirical RT, confidence, and β values, but this added level of complexity would not affect the model's ability to explain the quantities of interest here, which are the within-subject reliability values, subject-to-group similarity values, and fingerprinting results. This is because all three of these analyses are based on correlations, and correlations are insensitive to additive and multiplicative changes of the underlying variables. Therefore, we chose not to fit the actual observed values of RT, confidence, and voxel β values so as to keep the model as simple as possible.

4.11 Model fitting

We first fit the model to the empirically observed within-subject reliability and subject-to-group similarity values. The model had six free parameters corresponding to the weights, w, of the group- and subject-level factors that determined the simulated β value for each voxel in each block. For a given set of weights, we simulated a complete experimental dataset by generating simulated data for 50 subjects with 768 trials or 96 blocks per subjects (each trial or block had corresponding RT, confidence, and per voxel beta value). Based on these data, we then computed the within-subject reliability and subject-to-group similarity values in the same way as for the empirical data. When simulating the model, we observed that the exact number of voxels used made no systematic difference to the observed values of the obtained within-subject reliability and subject-to-group similarity values. Therefore, we used 10,000 voxels, which allowed for stable values to be obtained on different iterations. The fitting minimized the mean squared error (MSE) between the simulated and empirically observed within-subject reliability and subject-to-group similarity values calculated using the top-100% most activated voxels (that is, using all voxels). Once the fitting was completed, we also generated the predictions of the best-fitting model for the within-subject reliability and subject-to-group similarity values calculated using the top 10, 25, 50, and 75% most activated voxels. The fitting itself was carried out using the Bayesian Adaptive Direct Search (BADS) toolbox(Acerbi and Ma, 2017). We fit the model 10 times are reported the best fitting model among the 10 iterations.

4.12 Data and code availability

All preprocessed data and code are available at https://osf.io/xe8r5/.

References

Acerbi, L., Ma, W.J., 2017. Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search. In Advances in Neural Information Processing Systems 30, 1834–1844.

Afraz, A., Pashkam, M.V., Cavanagh, P., 2010. Spatial heterogeneity in the perception of face and form attributes. Current Biology 20. https://doi.org/10.1016/j.cub.2010.11.017

Bandettini, P.A., Ungerleider, L.G., 2001. From neutrons to BOLD: new connections. Nat Neurosci 4.

Belliveau, J.W., Kennedy, D.N., McKinstry, R.C., Buchbinder, B.R., Weisskoff, R.M., Cohen, M.S., Vevea, J.M., Brady, T.J., Rosen, B.R., 1991. Functional mapping of the human visual cortex by magnetic resonance imaging. Science (1979) 254. https://doi.org/10.1126/science.1948051

Breakspear, M., 2017. Dynamic models of large-scale brain activity. Nat Neurosci. https://doi.org/10.1038/nn.4497

Coombes, S., Venkov, N.A., Shiau, L., Bojak, I., Liley, D.T.J., Laing, C.R., 2007. Modeling electrocortical activity through improved local approximations of integral neural field equations. Phys Rev E Stat Nonlin Soft Matter Phys 76. https://doi.org/10.1103/PhysRevE.76.051901

Drysdale, A.T., Grosenick, L., Downar, J., Dunlop, K., Mansouri, F., Meng, Y., Fetcho, R.N., Zebley, B., Oathes, D.J., Etkin, A., Schatzberg, A.F., Sudheimer, K., Keller, J., Mayberg, H.S., Gunning, F.M., Alexopoulos, G.S., Fox, M.D., Pascual-Leone, A., Voss, H.U., Casey, B.J., Dubin, M.J., Liston, C., 2017. Resting-state connectivity biomarkers define neurophysiological subtypes of depression. Nat Med 23, 28–38. https://doi.org/10.1038/nm.4246

Durlauf, S.N., 2001. A framework for the study of individual behavior and social interactions. Sociol Methodol 31. https://doi.org/10.1111/0081-1750.00089

Eilam, D., 2015. The cognitive roles of behavioral variability: Idiosyncratic acts as the foundation of identity and as transitional, preparatory, and confirmatory phases. Neurosci Biobehav Rev 49, 55–70. https://doi.org/https://doi.org/10.1016/j.neubiorev.2014.11.023

Elliott, M.L., Romer, A., Knodt, A.R., Hariri, A.R., 2018. A Connectome-wide Functional Signature of Transdiagnostic Risk for Mental Illness. Biol Psychiatry 84, 452–459.

Fisher, A.J., Medaglia, J.D., Jeronimus, B.F., 2018. Lack of group-to-individual generalizability is a threat to human subjects research. Proc Natl Acad Sci U S A 115, 6106–6115.

Forkosh, O., Karamihalev, S., Roeh, S., Alon, U., Anpilov, S., Touma, C., Nussbaumer, M., Flachskamm, C., Kaplick, P.M., Shemesh, Y., Chen, A., 2019. Identity domains capture individual differences from across the behavioral repertoire. Nat Neurosci 22, 2023–2028. https://doi.org/10.1038/s41593-019-0516-y

Friston, K.J., Holmes, A.P., Price, C.J., Büchel, C., Worsley, K.J., 1999. Multisubject fMRI Studies and Conjunction Analyses. Neuroimage 10, 385–396.

Gilmore, A.W., Nelson, S.M., McDermott, K.B., 2021. Precision functional mapping of human memory systems. Curr Opin Behav Sci. https://doi.org/10.1016/j.cobeha.2020.12.013

Gordon, E.M., Nelson, S.M., 2021. Three types of individual variation in brain networks revealed by single-subject functional connectivity analyses. Curr Opin Behav Sci. https://doi.org/10.1016/j.cobeha.2021.02.014

Gratton, C., Kraus, B.T., Greene, D.J., Gordon, E.M., Laumann, T.O., Nelson, S.M., Dosenbach, N.U.F., Petersen, S.E., 2020. Defining Individual-Specific Functional Neuroanatomy for Precision Psychiatry. Biol Psychiatry.

Gratton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2018. Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. Neuron. https://doi.org/10.1016/j.neuron.2018.03.035

Haxby, J. v., Guntupalli, J.S., Connolly, A.C., Halchenko, Y.O., Conroy, B.R., Gobbini, M.I., Hanke, M., Ramadge, P.J., 2011. A Common, High-Dimensional Model of the Representational Space in Human Ventral Temporal Cortex. Neuron 72, 404–416.

Haxby, J. v, Guntupalli, J.S., Nastase, S.A., Feilong, M., 2020. Hyperalignment: Modeling shared information encoded in idiosyncratic cortical topographies. Elife 9, e56601.

Kable, J.W., Glimcher, P.W., 2007. The neural correlates of subjective value during intertemporal choice. Nat Neurosci 10. https://doi.org/10.1038/nn2007

Kaufmann, T., Alnæs, D., Doan, N.T., Brandt, C.L., Andreassen, O.A., Westlye, L.T., 2017. Delayed stabilization and individualization in connectome development are related to psychiatric disorders. Nat Neurosci 20, 513–515.

Kohoutová, L., Atlas, L.Y., Büchel, C., Buhle, J.T., Geuter, S., Jepma, M., Koban, L., Krishnan, A., Lee, D.H., Lee, S., Roy, M., Schafer, S.M., Schmidt, L., Wager, T.D., Woo, C.-W., 2022. Individual variability in brain representations of pain. Nat Neurosci 25, 749–759. https://doi.org/10.1038/s41593-022-01081-x

le Bihan, D., Turner, R., Zeffiro, T.A., Cuenod, C.A., Jezzard, P., Bonnerot, V., 1993. Activation of human primary visual cortex during visual recall: A magnetic

resonance imaging study. Proc Natl Acad Sci U S A 90. https://doi.org/10.1073/pnas.90.24.11802

Lebreton, M., Bavard, S., Daunizeau, J., Palminteri, S., 2019. Assessing inter-individual differences with task-related functional neuroimaging. Nat Hum Behav.

Li, A., Zalesky, A., Yue, W., Howes, O., Yan, H., Liu, Y., Fan, L., Whitaker, K.J., Xu, K., Rao, G., Li, J., Liu, S., Wang, M., Sun, Y., Song, M., Li, P., Chen, J., Chen, Y., Wang, Huaning, Liu, W., Li, Z., Yang, Y., Guo, H., Wan, P., Lv, L., Lu, L., Yan, J., Song, Y., Wang, Huiling, Zhang, H., Wu, H., Ning, Y., Du, Y., Cheng, Y., Xu, J., Xu, X., Zhang, D., Wang, X., Jiang, T., Liu, B., 2020. A neuroimaging biomarker for striatal dysfunction in schizophrenia. Nat Med 26, 558–565.

Martinez, J.E., Funk, F., Todorov, A., 2020. Quantifying idiosyncratic and shared contributions to judgment. Behav Res Methods 52. https://doi.org/10.3758/s13428-019-01323-0

Miller, M.B., Donovan, C.L., Bennett, C.M., Aminoff, E.M., Mayer, R.E., 2012. Individual differences in cognitive style and strategy predict similarities in the patterns of brain activity between individuals. Neuroimage. https://doi.org/10.1016/j.neuroimage.2011.05.060

Miller, M.B., Donovan, C.L., van Horn, J.D., German, E., Sokol-Hessner, P., Wolford, G.L., 2009. Unique and persistent individual patterns of brain activity across different memory retrieval tasks. Neuroimage. https://doi.org/10.1016/j.neuroimage.2009.06.033

Naselaris, T., Allen, E., Kay, K., 2021. Extensive sampling for complete models of individual brains. Curr Opin Behav Sci. https://doi.org/10.1016/j.cobeha.2020.12.008

Navajas, J., Hindocha, C., Foda, H., Keramati, M., Latham, P.E., Bahrami, B., 2017. The idiosyncratic nature of confidence. Nat Hum Behav 1. https://doi.org/10.1038/s41562-017-0215-1

Nieto-Castañón, A., Fedorenko, E., 2012. Subject-specific functional localizers increase sensitivity and functional resolution of multi-subject analyses. Neuroimage 63, 1646–1669.

Parkes, L., Satterthwaite, T.D., Bassett, D.S., 2020. Towards precise resting-state fMRI biomarkers in psychiatry: synthesizing developments in transdiagnostic research, dimensional models of psychopathology, and normative neurodevelopment. Curr Opin Neurobiol 65, 120–128.

Poldrack, R.A., 2011. Inferring mental states from neuroimaging data: From reverse inference to large-scale decoding. Neuron. https://doi.org/10.1016/j.neuron.2011.11.001

Rahnev, D., 2021. Response Bias Reflects Individual Differences in Sensory Encoding. Psychol Sci 32. https://doi.org/10.1177/0956797621994214

Ratcliff, R., 1978. A theory of memory retrieval. Psychol Rev 85. https://doi.org/10.1037/0033-295X.85.2.59

Rescorla, R.A., Wagner, A.R., 1972. A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement. Clasical conditioning II: current research and theory.

Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T., 2018. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. Cerebral Cortex 28, 3095–3114.

Seghier, M.L., Lee, H.L., Schofield, T., Ellis, C.L., Price, C.J., 2008. Inter-subject variability in the use of two different neuronal networks for reading aloud familiar words. Neuroimage. https://doi.org/10.1016/j.neuroimage.2008.05.029

Song, H., Rosenberg, M.D., 2021. Predicting attention across time and contexts with functional brain connectivity. Curr Opin Behav Sci. https://doi.org/10.1016/j.cobeha.2020.12.007

Wilson, H.R., Cowan, J.D., 1972. Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. Biophys J 12. https://doi.org/10.1016/S0006-3495(72)86068-5

Woodward, N.D., Cascio, C.J., 2015. Resting-State Functional Connectivity in Psychiatric Disorders. JAMA Psychiatry 72, 743–744. https://doi.org/10.1001/jamapsychiatry.2015.0484

Zhang, W., Luck, S.J., 2008. Discrete fixed-resolution representations in visual working memory. Nature 453. https://doi.org/10.1038/nature06860