



Published in final edited form as:

*Neuroimage*. 2021 August 15; 237: 118137. doi:10.1016/j.neuroimage.2021.118137.

## Relating whole-brain functional connectivity to self-reported negative emotion in a large sample of young adults using group regularized canonical correlation analysis

Leonardo Tozzi<sup>a</sup>, Elena Tuzhilina<sup>b</sup>, Matthew F. Glasser<sup>c</sup>, Trevor J. Hastie<sup>b</sup>, Leanne M. Williams<sup>a,d,\*</sup>

<sup>a</sup>Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, USA

<sup>b</sup>Department of Statistics, Stanford University, Stanford, USA

<sup>c</sup>Departments of Radiology and Neuroscience, Washington University, St. Louis, USA

<sup>d</sup>Sierra-Pacific Mental Illness Research, Education and Clinical Center (MIRECC) Veterans Affairs Palo Alto Health Care System, Palo Alto, CA, USA

### Abstract

The goal of our study was to use functional connectivity to map brain function to self-reports of negative emotion. In a large dataset of healthy individuals derived from the Human Connectome Project ( $N = 652$ ), first we quantified functional connectivity during a negative face-matching task to isolate patterns induced by emotional stimuli. Then, we did the same in a complementary task-free resting state condition. To identify the relationship between functional connectivity in these two conditions and self-reports of negative emotion, we introduce group regularized canonical correlation analysis (GRCCA), a novel algorithm extending canonical correlations analysis to model the shared common properties of functional connectivity within established

---

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

\*Corresponding author at: Department of Psychiatry and Behavioral Sciences, Stanford University School of Medicine, Stanford, CA, 94305, USA. leawilliams@stanford.edu (L.M. Williams).

#### Data/code availability statement

The code used for the generation of the connectivity matrices and visualizations is available at: [https://github.com/leotozzi88/rcca\\_HCP\\_emotion](https://github.com/leotozzi88/rcca_HCP_emotion). The code for the canonical correlations analysis is available at: <https://github.com/ElenaTuzhilina/Connectome>. All data used is available at <https://db.humanconnectome.org> and is accessible in compliance with the WU-Minn HCP Consortium Open Access Data Use Terms (<https://www.humanconnectome.org/study/hcpyoung-adult/document/wu-minn-hcp-consortium-open-access-datause-terms>).

#### Ethics statement

Data from the Human Connectome Project young healthy adults used in this study were acquired in accordance with the Washington University Institutional Review Board (Van Essen et al., 2013).

#### Declaration of Competing Interest

The authors declare no conflicts of interest.

#### Credit authorship contribution statement

**Leonardo Tozzi:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Project administration, Visualization, Writing original draft, Writing review & editing. **Elena Tuzhilina:** Conceptualization, Data curation, Formal analysis, Methodology, Software, Visualization, Writing original draft, Writing review & editing. **Matthew F. Glasser:** Resources, Data curation, Writing review & editing. **Trevor J. Hastie:** Methodology, Software, Supervision, Writing review & editing. **Leanne M. Williams:** Funding acquisition, Resources, Supervision, Writing review & editing.

#### Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2021.118137.

brain networks. To minimize overfitting, we optimized the regularization parameters of GRCCA using cross-validation and tested the significance of our results in a held-out portion of the data set using permutations.

GRCCA consistently outperformed plain regularized canonical correlation analysis. The only canonical correlation that generalized to the held-out test set was based on resting state data ( $r = 0.175$ , permutation test  $p = 0.021$ ). This canonical correlation loaded primarily on Anger-aggression. It showed high loadings in the cingulate, orbitofrontal, superior parietal, auditory and visual cortices, as well as in the insula. Subcortically, we observed high loadings in the globus pallidus. Regarding brain networks, it loaded primarily on the primary visual, orbito-affective and ventral multimodal networks.

Here, we present the first neuroimaging application of GRCCA, a novel algorithm for regularized canonical correlation analyses that takes into account grouping of the variables during the regularization scheme. Using GRCCA, we demonstrate that functional connections involving the visual, orbito-affective and multimodal networks are promising targets for investigating functional correlates of subjective anger and aggression. Crucially, our approach and findings also highlight the need of cross-validation, regularization and testing on held out data for correlational neuroimaging studies to avoid inflated effects.

## Keywords

Canonical correlations; Emotion; Functional connectivity; Resting state; Face-matching; Negative valence

## 1. Introduction

The recent Research Domain Criteria (RDoC) initiative is an effort spearheaded by the United States National Institutes of Mental Health to advance mental health research (Cuthbert and Insel, 2013). One core assumption of RDoC is that units of analysis measuring the same underlying construct can be integrated. In the present study, we use the Negative Valence Systems as a case study to test this assumption. Negative Valence Systems are primarily responsible for responses to aversive situations or context, such as fear, anxiety, and loss (Cuthbert and Kozak, 2013). In particular, we search for hidden constructs that underlie negative valence as measured both by the “self-report” and “circuit” units of analysis. Our choice is motivated by the fact that, in the context of mental health, mood disorders are the leading cause of disability and share as a core feature impairments in the regulation of Negative Valence Systems, which manifest clinically as persistent self-reported negative emotion (Whiteford et al., 2013; Woody and Gibb, 2015). To advance diagnosis and treatment of these disorders, there is a case to be made that quantitative measures are needed that do not rely exclusively on patient self-reports (Williams, 2016). Therefore, our findings shed light into the brain representation of self-reported negative emotion and inform future translational studies on mood disorders using the RDoC framework.

Many discoveries on how the functioning of human brain circuits correlates with self-reports in humans have been made possible because of advances in functional Magnetic Resonance

Imaging (fMRI). In particular, several recent crucial advances have been enabled by the collection and release of very large datasets leveraging cutting-edge methods for the acquisition and processing of images. One such dataset is the one collected by the Human Connectome Project (HCP) (Van Essen et al., 2013). Among the measures comprised in this dataset, functional connectivity, the correlation of blood oxygen level dependent signal timeseries between two or more brain regions, has proven an invaluable tool to explore the relationship between brain activity and constructs such as intelligence, sustained attention, impulsivity and behavior (Hearne et al., 2016; Li et al., 2013; Rosenberg et al., 2016; Smith et al., 2015). Accumulating evidence suggests that, similarly to these other complex phenomena, representations of functional connectivity are needed to map human brain function onto the emotional states that accompany it (see (Hamann, 2012; Lindquist et al., 2012) for comprehensive reviews).

The goal of our study was to evaluate functional connectivity as a target for mapping brain function to self-reports of negative emotion. To meet this objective, we incorporated several novel design and analytic strategies. Regarding the design of our experiment, first we quantified functional connectivity under a task-evoked condition in order to isolate any specific connectivity patterns induced by emotional stimuli. We build upon a robust foundation of studies attempting to relate brain activity to self-reports of negative emotion as participants undergo a task that involves processing of negative stimuli, for example faces expressing negative emotions (Mauss and Robinson, 2009; Murphy et al., 2003). Functional connectivity between brain areas is likely dynamic and changes depending on the current brain state, especially during a task. We hypothesized that the measures from an fMRI paradigm and a self-report questionnaire designed to probe the same psychological construct should be related and RDoC, along with a substantial body of work, proceeds on this assumption. In particular, we quantified functional connectivity during a task involving matching of faces expressing fear and anger which has been related to negative emotional states and disorders of these states (Prater et al., 2013; Westlund Schreiner et al., 2017). Second, we also considered that task-evoked functional connectivity might be a more fluctuating state-like measure compared to self-reports of emotional states experienced over of a period of days. As such, it might capture more transient features of functional connectivity, which do not necessarily relate to negative valence in the days preceding the scan or to personality traits. To address this point, we quantified functional connectivity in a complementary task-free resting state condition. The rationale for this is that resting state captures patterns of functional connectivity that are shared across states (Cole et al., 2014). We wanted to test if these patterns relate to feelings of negative emotion experienced throughout the days before the scan or to trait-like measures of personality related to negative emotion.

Regarding our analytic strategies, first of all we used the entire HCP Healthy Young Adult data release, thus guaranteeing a well-powered sample acquired and preprocessed with cutting-edge methods. We then used canonical correlation analysis (CCA), a well-established method for exploring correlations between functional connectivity and self-report measures (Drysdales et al., 2017; Smith et al., 2015; Xia et al., 2018). Compared to previous implementations, we made novel modifications to the algorithm which allowed us to address key challenges in the investigation of brain functional connectivity derived

from fMRI. First of all, we addressed a common issue with analyzing connectivity at the whole-brain level, namely the very high number of features compared to the number of measurements. To minimize the risk of overfitting, we used a regularized form of CCA combined with cross-validation to optimize our regularization parameters (González et al., 2008). Furthermore, we developed a novel implementation of CCA which allowed us to model shared common properties of functional connectivity within established brain networks. Finally, to ensure that the correlations detected were not spurious, we tested significance of our results in a held-out portion of the data set by permutation testing.

## 2. Methods

### 2.1. Dataset

Our sample is derived from the HCP Healthy Young Adult release, a large public dataset of 1200 subjects aged between 22 and 35 years without any psychiatric or neurological disorder (Van Essen et al., 2013). The acquisition parameters and minimal preprocessing of these data are described in (Glasser et al., 2013) and comprised high spatial (2 mm isotropic) and temporal (TR = 0.72 s) resolution multi-band fMRI. Briefly, participants underwent a large number of MRI scans, that included T1 and T2 weighted structural imaging, diffusion tensor imaging, and nearly 2 h of resting-state and task fMRI. For the present study, we used fMRI data from the second day of acquisition: 4:32 min. of Emotion task (2 runs acquired with RL and LR phase encoding respectively, 2:16 min. and 176 time-points each) and 30 min. of resting-state (2 runs acquired with RL and LR phase encoding respectively, 15 min and 1200 time-points each).

To select our sample, we accessed the data at <https://db.humanconnectome.org>. Using the online filtering options, we selected only participants who had completed the full task and resting state scanning protocol, had no known quality issues and had completed the NIH toolbox battery (see below). This returned a total of 652 subjects. For these, we downloaded fMRI data denoised for spatially specific artifacts from head motion, subject physiology, and MR physics sources using ICA-FIX (Salimi-Khorshidi et al., 2014), in particular single run sICA+FIX in the case of the resting state data and multi-run sICA+FIX in the case of the Emotion data. These latter data were provided by MFG (Glasser et al., 2018). All analyses were conducted in greyordinate space, i.e. they were constrained to the gray matter by using files in the CIFTI format, thus taking full advantage of HCP preprocessing and minimizing non-neuronal signal and blurring due to cross-subject misalignment or over smoothing across tissue boundaries (Glasser et al., 2013).

### 2.2. fMRI emotion task

This task is described in detail in (Barch et al., 2013) and has been widely used to engage the neural circuits underlying negative emotion. Briefly, participants are presented with blocks of trials that either ask them to decide which of two angry or fearful faces presented on the bottom of the screen match the one at the top, or which of two shapes presented at the bottom of the screen match the one at the top. Trials are presented in blocks of 6 trials of the same task (face or shape), with the stimulus presented for 2 s and a 1 s interstimulus

interval. Each block is preceded by a 3 s task cue (“shape” or “face”), so that each block is 21 s including the cue. Each of the two runs includes 3 face blocks and 3 shape blocks.

### 2.3. Negative emotion self-reports

The NIH Toolbox Emotion battery is a self-report developed to measure the full spectrum of emotional health (Babakhanyan et al., 2018). Given the focus of our study on negative emotion we selected the specific set of six measures defined as assessments of negative emotion out of the total 17. These are: fear-affect (e.g. feelings of anxiety, worry, fright), fear-somatic (e.g. nausea, dizziness, heart racing), anger-affect (e.g. feeling of wanting to break things or yelling at someone), anger-hostility (e.g. jealousy or envy towards others), anger-aggression (e.g. getting into fights, threatening others), and sadness (e.g. hopelessness, depression, guilt). In the NIH Toolbox Emotion battery, the questions assessing fear-affect, fear-somatic, anger-affect and sadness are all about experience of those emotions in the 7 days preceding the scan (“In the past 7 days...”). The questions assessing anger-hostility and anger-aggression ask the participant to rate how much a statement related to the emotion refers to them in general (“How true of you is this statement?”). The “raw” scores of the HCP Healthy Young Adult release for each participant for each of these negative emotion measures were used in our study. These scores are normed to a healthy population with mean=50 and standard deviation=10.

### 2.4. Bias field correction and parcellation

These analyses were conducted in Matlab R2018a (9.4.0.949201) for Mac (The MathWorks, Inc.) or using connectome workbench 1.3.2 for Mac (<https://www.humanconnectome.org/software/connectomeworkbench>).

First, an improved gradient echo and spin echo B1-receive field bias field correction was applied to each dense denoised timeseries as described in (Glasser et al., 2016). Briefly, after removing transmit field effects and excluding dropout regions, the low spatial frequency intensity variations ( $\sigma = 5$  mm) within gray matter were used to compute a more accurate, smoother receive bias field, which was then scaled to a volume mean of 10,000. Then, the improved scaled field was used to replace the bias field correction applied during minimal HCP preprocessing by multiplying the data by the original bias field map and then dividing it by the new field map. This produces data that are bias corrected in the same way as the current version of the HCP Pipelines (e.g. HCP Healthy Young Adult 7T data or HCP Lifespan data releases). After this bias field correction, each dense denoised bias-corrected timeseries was parcellated using connectome workbench (*wb\_command -cifti-parcellate*) to obtain the mean timeseries in each region of the HCP’s multi-modal cortical parcellation version 1.0 (Glasser et al., 2016). Since the parcellation did not include subcortical structures, these were derived from the Freesurfer segmentation (Fischl et al., 2002) and added to the CIFTI dense label file using connectome workbench (*wb\_command -cifti-createdense-from-template*).

### 2.5. Connectome construction

For each subject, parcellated timeseries were demeaned and the two timeseries for each condition (Emotion task and Rest) were concatenated. Connectivity matrices were

constructed by using L2 regularized partial correlations as implemented in FSLNets (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets>) to account for cross-subject differences in global respiratory noise. To determine the optimal regularization parameter  $\lambda$ , we proceeded as follows separately for the Emotion task and Rest conditions (Bijsterbosch et al., 2018). First, we concatenated the timeseries of all subjects and computed a group connectivity matrix  $C^{group}$  using partial correlations without regularization. Then, for the fixed regularization parameter  $\lambda$  and for each subject  $s = 1, \dots, 652$  we generated a connectivity matrix  $C^s$  and computed Pearson correlations between the upper triangle of  $C^s$  with that of the group connectivity matrix  $C^{group}$ , i.e.  $r^s = cor(C_{upper}^s, C_{upper}^{group})$ . We repeated this for all  $\lambda$  in the grid  $\lambda = 10^{-3}, 10^{-2}, \dots, 10^3$  and chose as the optimal value of  $\lambda$  the one that maximized this correlation on average across subjects, in other words, that maximized  $r = \frac{1}{652} \sum_{s=1}^{652} r^s$ . To further fine-tune the value of  $\lambda$ , we then repeated this process varying it in the interval surrounding the first optimal  $\lambda$ . At the end of this process, two matrices  $C^s$  were generated for each subject using the optimal  $\lambda$ : Emotion task matrix  $C_{Emotion}^s$  and Rest task matrix  $C_{Rest}^s$ . The values of  $\lambda$  maximizing the average correlation of individual connectivity matrices to the group matrix were respectively  $\lambda = 8$  for the Emotion task condition and  $\lambda = 0.7$  for the Rest condition (Figs. S1–,S2), which is expected due to the much fewer timepoints in the emotion task (352) than the resting state data (2400). In the present study, we did not model brain responses to the Emotion task explicitly. This was motivated by the following two considerations. First of all, previous studies have shown that differences in connectivity between rest and tasks is detectable even when the modeled response to task stimuli has been removed. This is because, even though they share some functional networks, resting state and task fMRI data are fundamentally different and because the effects of task fMRI are not fully modeled by task fMRI designs (Glasser et al., 2018). Secondly, we wanted to compare resting state and task data directly. Modeling the conditions of the task data would have made this comparison difficult, because our measure of interest (“functional connectivity”) would have been computed differently in the two conditions.

## 2.6. Data matrices

From here on, the analysis was conducted in *R* version 3.6.1. For each participant  $s = 1, \dots, 652$ , our imaging features were the upper triangles of the connectivity matrices  $C_{Emotion}^s$  and  $C_{Rest}^s$  described above (71,631 features in total), whereas our questionnaire features were the negative emotion raw scores from the NIH Toolbox Emotion (6 features in total). We stored fMRI and questionnaire features in  $X$  and  $Y$  matrices of sizes  $652 \times 71,631$  and  $652 \times 6$ , respectively. This resulted in two data pairs:  $(X_{Rest}, Y)$  and  $(X_{Emotion}, Y)$ . Further analyses were done separately for each of them.

## 2.7. Feature preprocessing

For a pair of data  $(X, Y)$  we first regressed the effects of sex from both the questionnaire and connectivity sides. In other words, for each feature we computed the mean value across the female and male sex categories and for each participant we removed the corresponding mean

from the feature value. Since the data was quite homogeneous in terms of the age, we did not correct for this variable. We denote the adjusted data  $(\tilde{X}, \tilde{Y})$ .

## 2.8. Canonical correlations overview

The steps adopted to run canonical correlation analysis for  $(\tilde{X}, \tilde{Y})$  are summarized in Fig. 1 and outlined in detail in the following sections.

Since the number of imaging features was significantly larger than the number of participants, to avoid overfitting, we applied regularization to the fMRI features. In this study we considered two regularization schemes: the standard  $\ell_2$ -norm regularization (RCCA method (Leurgans et al., 1993 ; Vinod, 1976)) and a novel group regularization (GRCCA method (Tuzhilina et al., 2020)) that takes into account the organization of brain connectivity into networks. First of all, we randomly divided the data into a train (80%) and test (20%) set, i.e.  $(\tilde{X}_{train}, \tilde{Y}_{train})$  and  $(\tilde{X}_{test}, \tilde{Y}_{test})$ . The decision of an uneven split was made because most of the model fitting and tuning of the hyper parameters is conducted on the training data, whereas the test data is used only to validate the canonical correlations. Thus, we wanted to make sure that the model would be able to use a large number of subjects ( $N = 515$ ) to capture all the relevant features of the data even if only small correlations were present. At the same time, our split also allowed us to leave a large enough set to test the correlation ( $N = 137$ ). The regularization parameters  $\lambda$  and  $\mu$  (see below) were chosen by 10-fold cross-validation in the training set  $(\tilde{X}_{train}, \tilde{Y}_{train})$ . Since there were some close relatives (e.g. siblings and twins) among participants, to reduce the potential correlation between train and test data splits, we ran block cross-validation including all family members in the same fold. Let us denote the corresponding 10 folds by  $(\tilde{X}_{train}^{(1)}, \tilde{Y}_{train}^{(1)}), \dots, (\tilde{X}_{train}^{(10)}, \tilde{Y}_{train}^{(10)})$ . Next, for each combination of  $\lambda$  and  $\mu$  and for fold number  $k = 1, \dots, 10$  we used  $(\tilde{X}_{train}^{(-k)}, \tilde{Y}_{train}^{(-k)})$ , i.e. all but  $k$ -th fold, to fit the canonical correlation models. Thus, we obtained the coefficient vectors  $\alpha_{train}^{(-k)}$  and  $\beta_{train}^{(-k)}$  for the first canonical pair. Using these coefficients and the left out fold  $(\tilde{X}_{train}^{(k)}, \tilde{Y}_{train}^{(k)})$  we computed the questionnaire and connectivity first canonical variates  $\tilde{X}_{train}^{(k)} \alpha_{train}^{(-k)}$  and  $\tilde{Y}_{train}^{(k)} \beta_{train}^{(-k)}$ , then we measured the correlation between them. Thus, for each fold we calculated the correlation  $\rho_{train}^{(k)} = \text{cor}(\tilde{X}_{train}^{(k)} \alpha_{train}^{(-k)}, \tilde{Y}_{train}^{(k)} \beta_{train}^{(-k)})$ . As the optimal  $\lambda$  and  $\mu$ , we chose the  $\lambda$  and  $\mu$  maximizing the average first canonical correlation  $\rho_{train} = \frac{1}{10} \sum_{k=1}^{10} \rho_{train}^{(k)}$ . Once the regularization parameters were chosen, we ran the canonical correlation algorithm with those parameters set to the chosen values on the whole training set  $(\tilde{X}_{train}, \tilde{Y}_{train})$ . We obtained the canonical coefficient vectors  $\alpha_{train}$  and  $\beta_{train}$ , then we used these coefficients and the independent test set  $(\tilde{X}_{test}, \tilde{Y}_{test})$  to calculate the test imaging canonical variate  $\tilde{X}_{test} \alpha_{train}$  as well as the test questionnaire canonical variate  $\tilde{Y}_{test} \beta_{train}$ . Finally, we measured test canonical correlation as  $\rho_{test} = \text{cor}(\tilde{X}_{test} \alpha_{train}, \tilde{Y}_{test} \beta_{train})$ .

## 2.9. Correlation validation via permutation test

To check the significance of the obtained correlation  $\rho_{test}$  we applied a permutation test as follows. First we permuted the rows of  $\tilde{Y}_{train}$  thereby breaking the relation between  $\tilde{X}_{train}$  and  $\tilde{Y}_{train}$ ; we denote the resulting matrix by  $\tilde{Y}_{train}^{perm}$ . We used  $(\tilde{X}_{train}, \tilde{Y}_{train}^{perm})$  and, as in the previous section, we applied GRCCA with the chosen  $\lambda$  and  $\mu$  to compute  $\alpha_{train}^{perm}$  and  $\beta_{train}^{perm}$ . Then we permuted the rows of  $\tilde{Y}_{test}$ . Denoting the result by  $\tilde{Y}_{test}^{perm}$  we evaluated the permuted test correlation as  $\rho_{test}^{perm} = cor(\tilde{X}_{test}\alpha_{train}^{perm}, \tilde{Y}_{test}^{perm}\beta_{train}^{perm})$ . We repeated this procedure 1000 times, thus obtaining the null distribution for the test correlation. We positioned the actual test correlation  $\rho_{test}$  in this distribution and computed the  $p$ -value as the relative frequency of the event  $\{\rho_{test} < \rho_{test}^{perm}\}$ . We considered as our final results only correlations for which this  $p$ -value was  $< 0.05$ .

## 2.10. Evaluating canonical pairs beyond the first

For correlations having a  $p$ -value of  $\rho_{test} < 0.05$ , we also investigated all canonical pairs beyond the first (six in total). To check if they were informative, we computed an aggregated measure of correlation between all canonical variates. Specifically, if  $(\alpha_{train}^1, \beta_{train}^1), \dots, (\alpha_{train}^6, \beta_{train}^6)$  represent canonical coefficients corresponding to six canonical pairs, then  $V_{test} = (\tilde{X}_{test}\alpha_{train}^1, \dots, \tilde{X}_{test}\alpha_{train}^6)$  and  $U_{test} = (\tilde{Y}_{test}\beta_{train}^1, \dots, \tilde{Y}_{test}\beta_{train}^6)$  represent matrices containing six imaging and six questionnaire canonical variates, respectively. We measured the aggregated correlation as the first correlation obtained by CCA conducted for the pair  $(V_{test}, U_{test})$ . Hereafter we denote this aggregated correlation measure by  $\rho_{test}^{CCA}$ . Note that no regularization is required for CCA this time, since the data contains only 6 features on both sides. This approach is grounded upon two considerations. First, it is a generalization of the method we used to measure the performance on the test set based on the correlation between first variates. Second, computing the first CCA correlation is equivalent to measuring the cosine of the principal angle between canonical subspaces (Zhu and Knyazev, 2013). In other words, it measures how close the two canonical subspaces spanned by the six pairs of canonical variates are. To test the significance of the aggregated correlation we once again obtained the null distribution for  $\rho_{test}^{CCA}$  via permutation testing as outlined above and computed a  $p$ -value. We considered the canonical pairs beyond the first informative if this  $p$ -value was  $< 0.05$ .

## 2.11. Canonical correlations implementation

To conduct RCCA we used the `rcc()` function from the CCA R package. Since this function is not able to handle matrices with a large number of features ( $\tilde{X}$  has 71,631 fMRI features in our case), we could not apply it directly to the data. To get around the dimensionality issue we applied the ‘‘Kernel Trick’’ (Tuzhilina et al., 2020) but see (Haroon et al., 2004 ; Mihalik et al., 2020 ) for similar implementations): rather than computing canonical coefficients in 71,631 column space we did all the computations in 652 row space. To achieve this, we proceeded as follows. First, we decomposed the matrix  $\tilde{X}$  of size 652



$\times 71,631$  into the product of a square matrix  $R$  (of size  $652 \times 652$ ) and a rectangular matrix  $V^T$  (of size  $652 \times 71,631$ ). Here  $V$  is a matrix with orthogonal columns. There is mathematical proof that the RCCA problem for the original data  $(\tilde{X}, \tilde{Y})$  can be reduced to the RCCA problem for the low-dimensional data pair  $(R, \tilde{Y})$ . Specifically, one can show that the canonical variates calculated for  $(\tilde{X}, \tilde{Y})$  are equal to the ones computed for  $(R, \tilde{Y})$ . Moreover, the original RCCA coefficient vector  $\alpha$  can be easily recovered from  $\alpha_R$ , i.e. the coefficient vector calculated for the pair  $(R, \tilde{Y})$ , via the linear relation  $\alpha = V\alpha_R$ . Therefore, running regularized CCA for the original wide fMRI matrix  $\tilde{X}$  is equivalent to applying the *rcc()* function to a small square matrix  $R$ . For this approach during the cross-validation (see above), we tested 9 different values of the regularization parameter:  $\lambda = 10^{-2}, 10^{-3}, \dots, 10^6$  for the Rest task data and  $\lambda = 10^{-6}, 10^{-5}, \dots, 10^2$  for the Emotion task data.

### 2.12. RCCA accounting for shared properties between brain regions (group RCCA)

Note that standard RCCA utilizing  $\ell_2$ -norm as a regularization term has a property of shrinking CCA coefficients to zero with the growth of the penalty factor  $\lambda$ . It treats all the features equally thereby ignoring any underlying data structure. We developed a novel algorithm by exploiting the presence of a group structure of the fMRI features (induced by the brain networks). The procedure, that we call group RCCA (GRCCA), is a modification of Regularized Canonical Correlation Analysis that integrates the group structure into the regularization scheme (Tuzhilina et al., 2020). Specifically, we divided our original 379 regions into 14 networks. To do so, we grouped the original cortical regions into 12 established networks: Visual primary, Visual secondary, Somatomotor, Cingulo-opercular, Dorsal-attention, Language, Frontoparietal, Auditory, Default, Posterior-multimodal, Ventral-multimodal, Orbito-affective (Ji et al., 2019). To these, we added a Subcortical and a Cerebellar network. This resulted in dividing 71,631 fMRI connectivity features in 105 groups (each corresponding to a pair of networks). Further, in the standard RCCA problem we replaced the  $\ell_2$ -norm constraint, which limits the deviation of all the CCA coefficients from zero, by constraints on within and between group variations. The first constraint restricted the deviation of the CCA coefficients from the corresponding group means thereby stimulating the homogeneity of coefficients inside each group. The second constraint limited the deviation of the group means from zero thereby encouraging the sparsity on a group level. These two constraints resulted in two hyperparameters: a penalty factor  $\lambda$  (controls within group variation) and a penalty factor  $\mu$  (controls between group variation). During the cross-validation (see above), we tested the following grid of the regularization parameters:  $\lambda = 10^{-4}, 10^{-3}, \dots, 10^4$  and  $\mu = 10^{-4}, 10^{-3}, \dots, 10^4$ .

### 2.13. Interpretation and visualization of canonical pairs

The RCCA and GRCCA analyses returned a weight (canonical coefficient) for each edge of the fMRI connectivity matrix for each canonical pair. To interpret these weights, we provide six visualizations for the canonical pair surviving the permutation testing. The first represents the correlation  $\rho_{test}$  between the questionnaire and connectivity scores (i.e. canonical variates  $\tilde{X}_{test}\alpha_{train}$  and  $\tilde{Y}_{test}\beta_{train}$ ), with a scatterplot representing the score values per participant. The second is the weights  $\beta_{train}$  that were given to the questionnaires by this

procedure. The third visualization is a  $379 \times 379$  matrix showing the fMRI weights  $\alpha_{train}$  given to each edge of the connectivity matrix, ordered by network. The fourth is a  $14 \times 14$  matrix showing the same weights averaged by network. For the fifth and sixth visualizations, we used the Brain Connectivity Toolbox to calculate the sum of the positive and negative weights of connections involving each region (a measure known as “strength” of a network node (Rubinov and Sporns, 2010)). These are then shown on an inflated brain.

### 3 Code and data availability

The code used for the generation of the connectivity matrices and visualizations is available at: [https://github.com/leotozzi88/rcca\\_HCP\\_emotion](https://github.com/leotozzi88/rcca_HCP_emotion). The code for the canonical correlations analysis is available at: <https://github.com/ElenaTuzhilina/Connectome>. All data used is available at <https://db.humanconnectome.org> and is accessible in compliance with the WU-Minn HCP Consortium Open Access Data Use Terms (<https://www.humanconnectome.org/study/hcpyoung-adult/document/wu-minn-hcp-consortium-open-access-data-use-terms>).

## 4. Results

### 4.1. Sample characteristics

The final sample characteristics are presented in Table 1. Results in the sample for the NIH Emotion Toolbox battery were in line with those of a healthy population as expected, with means ranging from 46.324 to 51.783 (reference population=50) and standard deviations from 7.948 to 8.913 (reference population=10).

### 4.2. Canonical correlations

GRCCA consistently outperformed RCCA (i.e. returned higher correlations in the cross-validation), therefore we only tested GRCCA results in the held-out test set.

**4.2.1. Emotion task**—When selecting the regularization parameters as the ones returning the highest correlation between self-report and imaging features for the first canonical pair, for RCCA, the optimal regularization parameter was  $\lambda = 0.001$  (mean  $r = 0.095$  across cross-validation folds, Fig. S3). Using GRCCA, the best regularization parameters were  $\lambda = 0.01$  and  $\mu = 100$  (mean  $r = 0.136$  across cross-validation folds) (Fig. S4). The correlation obtained by GRCCA did not generalize to the held-out test data ( $r = 0.075$ , permutation test  $p = 0.17$ , Fig. S5).

**4.2.2. Rest**—When selecting the regularization parameters as the ones returning the highest correlation between self-report and imaging features for the first canonical pair, for RCCA, the optimal regularization parameter was  $\lambda = 100$  (mean  $r = 0.148$  across cross-validation folds, Fig. 2). Using GRCCA, the best regularization parameters were  $\lambda = 100$  and  $\mu = 104$  (mean  $r = 0.158$  across cross-validation folds, Fig. 2). The correlation between the first canonical variates returned by GRCCA generalized to the held-out test data ( $r = 0.175$ , permutation test  $p = 0.021$ , Fig. S6). However, our permutation measuring

the aggregate correlation between canonical subspaces was not significant, so we did not consider as results canonical correlations beyond the first ( $r = 0.371$ ,  $p = 0.325$ , Fig. S7).

This canonical variate loaded primarily on anger-aggression ( $-0.101$ ) (Table 2). The brain regions with the highest node positive strength were the anterior cingulate, orbitofrontal cortex, superior parietal cortex, auditory cortex and left globus pallidus. The brain regions with the highest negative strength were the visual cortex, anterior cingulate, posterior cingulate, orbitofrontal cortex and insula (Fig. 3). The brain networks with the highest strength was the primary visual network and the brain networks with the highest average negative strength were the orbito-affective and ventral multimodal networks (Table 3, Fig. 3).

## 5. Discussion

In the present work, we identified a linear relationship between self-reports and function of brain circuits relevant to the RDoC Negative Valence System in a large sample of healthy individuals. Our goal was to identify biomarkers mapping emotion state to quantifiable brain function measures. We make the case that this is necessary to identify precise targets to manipulate in subsequent studies of emotion circuits and for future clinical studies in populations suffering from mood disorders (Finn et al., 2017).

In this study, we also present the first neuroimaging application of GRCCA, a novel algorithm for regularized canonical correlation analyses that takes into account grouping of the variables during the regularization scheme. In our specific case, this allowed us to group functional connections between brain regions into known functional networks. Also, the use of the “kernel trick” makes GRCCA able to handle a much larger number of features compared to previous CCA implementations. In our study, this allowed us to extend our analyses to the entire functional connectome. Importantly, GRCCA consistently outperformed RCCA, i.e. always identified higher test correlations in the cross-validation procedure. This indicates that GRCCA is a promising new tool for linking functional connectivity measures with data from other units of analysis.

From a methodological point of view, another crucial consideration is that adopting a cross-validated and regularized implementation of the CCA algorithm mitigated the inflation of results due to overfitting. Indeed, at low values of our regularization parameter the correlation between self-report and brain circuit data was extremely high ( $r > 0.90$ ) in the training set but consistently  $\sim 0$  in the validation set. Even with an adequate regularization, the correlations we detected were weak ( $r < 0.20$ ), suggesting that we were able to detect them only because of our well-powered large sample. Furthermore, only one of these correlations generalized to a held-out test set. As previously shown by others, our results highlight the fundamental need of cross-validation, regularization and testing on held-out data for correlational neuroimaging studies to avoid the erroneous reporting of unrealistically high effects (Dinga et al., 2019). Indeed, the size of our correlations might be in the range of what should be expected from such analyses in the field of Psychology, at least in healthy populations (Schäfer and Schwarz, 2019). More in detail, the small effect size of our results could be due to several factors. First of all, it is possible that functional

connectivity mostly contains information about the current functional brain state and not about past states. This is in contrast with the measures of the NIH Emotion toolbox, in which some questionnaires refer to experiences extending up to several days before the scan and other to stable personality traits. It could also be that the relationship between functional connectivity and emotional states is complex or indirect and that its quantification requires more sophisticated modeling techniques compared to what is achieved with an exploratory linear method like GRCCA. Finally, our sample mean and standard deviation on the NIH Emotion toolbox was perfectly in line with the negative emotion values of a healthy reference population. A sample containing clinical participants (e.g., suffering from mood disorders) might have a larger variability of self-reports of negative valence as well as more extreme values. Therefore, future studies could use GRCCA to link functional connectivity and self-reported negative emotion in such populations. We speculate that in that case, stronger effects could be detected.

Using GRCCA, we identified only one canonical pair relating emotional states to brain function at Rest that generalized to the held-out test set. This canonical pair provides novel insights into the brain circuitry underlying negative emotion. Interestingly, it had strong negative loading especially on Anger-aggression, which was one of the two scales of the NIH Emotion toolbox assessing negative emotion over a much longer period of time (“in general”), compared to over the past 7 days. This might suggest that what is being captured in resting state functional connectivity is more similar to a “trait” characteristic which relates to emotional phenomena lasting for a longer time before the scan, which is in line with our rationale for including Rest data in our analyses. The final canonical pair had high loadings in the connections involving the cingulate cortex, the orbitofrontal cortex and insula. Among subcortical regions, the globus pallidus was also highly loaded. These regions are all known to be involved in emotional regulation and our result might reflect this aspect of their function (Frank et al., 2014). When considering brain networks more broadly, the canonical pair had positive loadings on the primary visual network as well as negative loadings in the orbito-affective and the ventral multimodal networks. The high loadings in the visual network were surprising and suggest that connectivity of this network is a potentially understudied target for future research into the functional correlates of negative emotion. Concerning the orbito-affective and ventral multimodal networks, they were only recently identified (Ji et al., 2019). The orbito-affective network encompasses cortical and subcortical regions associated with reward processing and is modulated differentially by rewarding stimuli (Ji et al., 2019). This suggests that it might have a role in the processing of positive emotion, which might explain its negative loading in our analysis focused on finding correlates of negative emotion. The ventral multimodal network, on the other hand, consists of the ventral surface of the temporal lobe and extends to the right ventral striatum and hippocampus. The function of this network is unclear, but it has been hypothesized to represent higher-order semantic categories (Ji et al., 2019). We suspect that, in our results, the high negative loading of this network might be rather related to the known function of these regions in the processing of emotion (Frank et al., 2014). In particular, it might suggest that dysregulation of this network is related to excessive negative emotion.

Concerning the Emotion task, we found no canonical pairs in this condition that survived our permutation testing on the held-out test set. One possible explanation of these null

findings might lie in the fact that functional connectivity measures are more reliable the longer the duration of the scan. Good reliability of a measure is needed to test whether it is related to other ones that vary substantially across individuals (Anderson et al., 2011; Birn et al., 2013; Elliott et al., 2019; Noble et al., 2019; Termenon et al., 2016; Tozzi et al., 2020). In the HCP data we used for the current analysis, the Emotion task was comparatively much shorter than the Rest scan (4:32 versus 30 min.). Thus, it is possible that longer scans are needed to obtain reliable functional connectivity estimates that can be linked to individual characteristics.

Our study was not without limitations. First of all, we were unable to validate our GRCCA results in a completely independent sample, since we could not find a public dataset containing the same self-report. Nevertheless, we used regularization and cross-validation to mitigate the effects of overfitting and tested our final results in a held-out portion of the dataset. Also, our study dataset is much larger than typically reported in the literature. Future neuroimaging studies using the NIH Emotion battery should consider making their data public to enable replication of our results. Another limitation was that we computed functional connectivity during our task in the same way we did it for resting state data. We adopted this strategy because it is commonly used and because it allowed us to easily compare the same measure between the two conditions. However, a very recent study has claimed that task functional connectivity computed in this way may produce false positives compared to computational models that allow for a flexible task-evoked BOLD response shape (Cole et al., 2019). Thus, it is possible that computing functional connectivity during specific conditions or comparing functional connectivity between different conditions might have returned results for the Emotion task. Doing so constitutes a promising avenue for future research which could also be facilitated by our novel GRCCA algorithm. We want to acknowledge the small effect size of the correlation we detected with our exploratory approach in this sample. Caution is warranted in the interpretation of such small effects and, in any case, they suggest that future studies wishing to investigate correlations between functional connectivity and self-reports of negative emotion might need very large samples. Finally, we also find that the complex multivariate patterns returned by canonical correlations make a direct neurobiological interpretation of our findings challenging. Nevertheless, we hope our results will inform future studies even just by showing that functional connectivity measures related to self-reported negative valence are distributed in the brain, involving sometimes even understudied or newly discovered networks.

In sum, we show that task-free resting state functional connectivity might be a promising target to identify correlates of self-reported negative emotion. We do so by using a novel algorithm for canonical correlation analyses that accounts for the grouping of functional connectivity measures into known brain networks. Functional connections involving the visual, auditory, cingulate, the orbitofrontal cortices and insula as well as the globus pallidus relate to self-reports of negative emotion in the days preceding the scan. Furthermore, the visual, orbito-affective and multimodal networks are promising targets for future research investigating negative valence across units of analysis in the resting state. Crucially, we highlight the need of cross-validation, regularization and testing on held-out data for correlational neuroimaging studies to avoid inflated effects.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments

Funding sources

This work was supported by the National Institutes of Health [grant number U01MH109985 under PAR-14-281].

## References

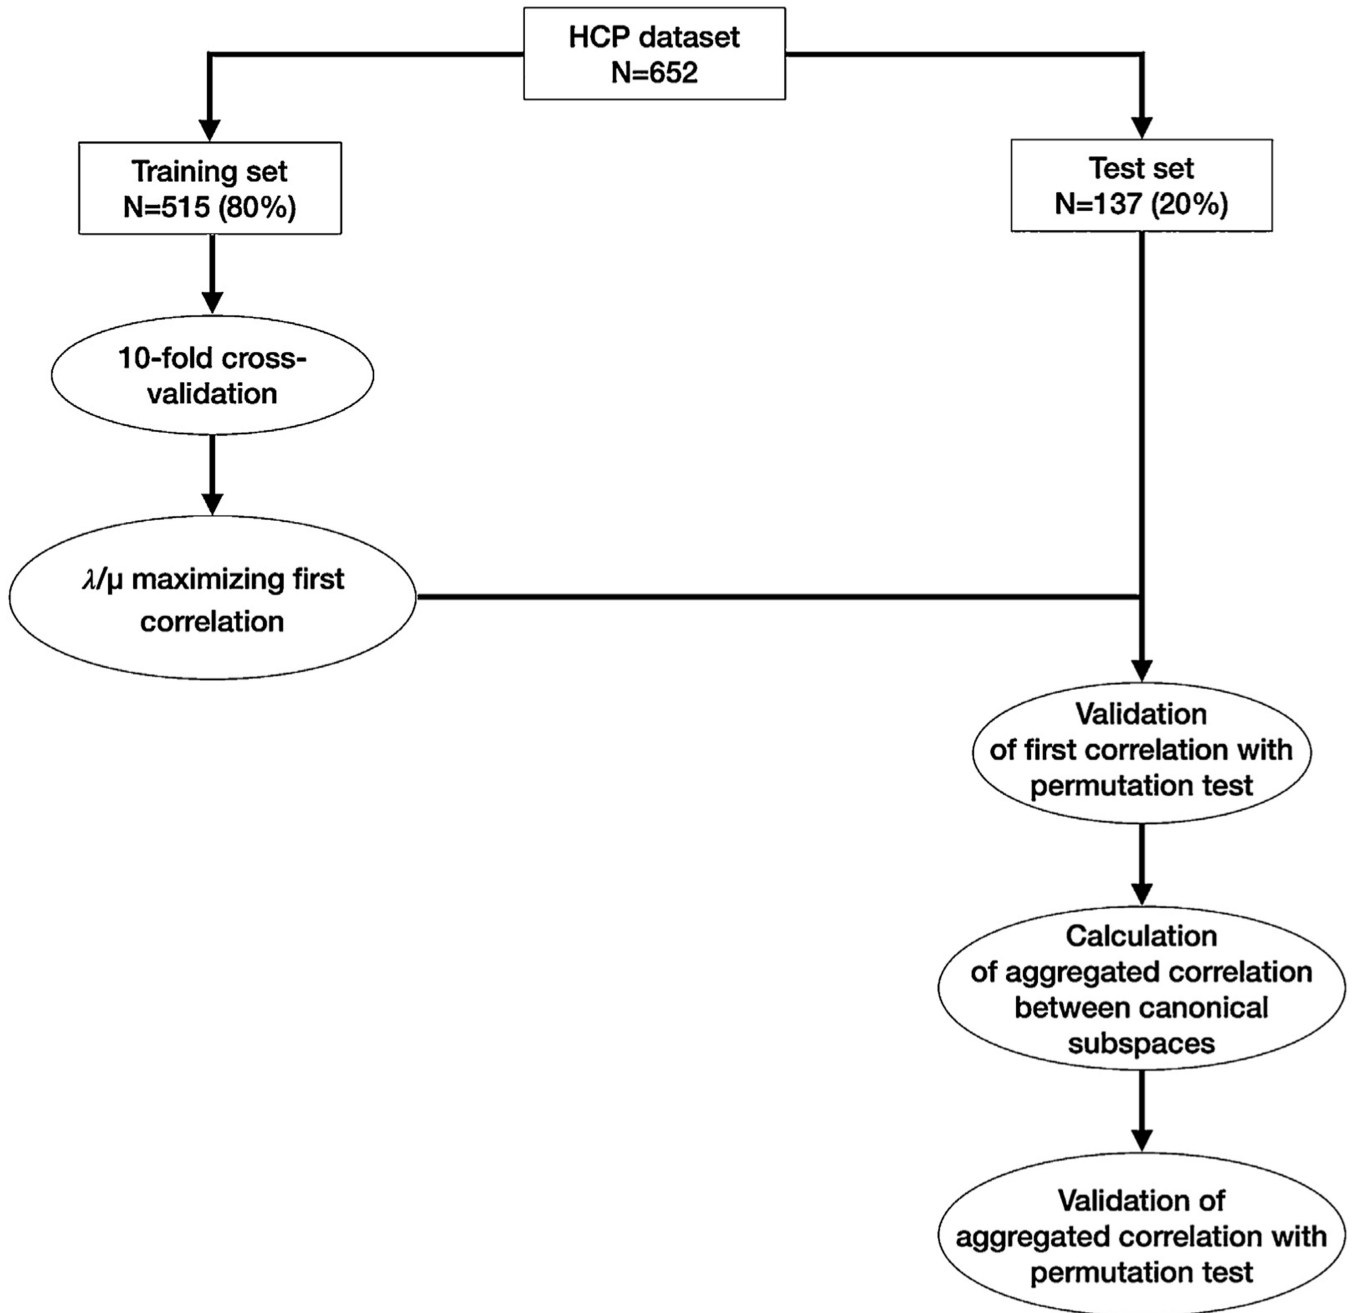
- Anderson JS, Ferguson MA, Lopez-Larson M, Yurgelun-Todd D, 2011. Reproducibility of single-subject functional connectivity measurements. *AJNR Am. J. Neuroradiol* 32, 548–555. doi:10.3174/ajnr.A2330. [PubMed: 21273356]
- Babakhanyan I, McKenna BS, Casaletto KB, Nowinski CJ, Heaton RK, 2018. National institutes of health toolbox emotion battery for English and Spanish-speaking adults: normative data and factor-based summary scores. *Patient Relat. Outcome Meas* 9, 115–127. doi:10.2147/PROM.S151658. [PubMed: 29588623]
- Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C, Nolan D, Bryant E, Hartley T, Footer O, Bjork JM, Poldrack R, Smith S, Johansen-Berg H, Snyder AZ, Van Essen DC, 2013. Function in the human connectome: task-fMRI and individual differences in behavior. *NeuroImage Mapp. Connect* 80, 169–189. doi:10.1016/j.neuroimage.2013.05.033.
- Bijsterbosch JD, Woolrich MW, Glasser MF, Robinson EC, Beckmann CF, Van Essen DC, Harrison SJ, Smith SM, 2018. The relationship between spatial configuration and functional connectivity of brain regions. *elife* 7, e32992. doi:10.7554/eLife.32992.
- Birn RM, Molloy EK, Patriat R, Parker T, Meier TB, Kirk GR, Nair VA, Meyerand ME, Prabhakaran V, 2013. The effect of scan length on the reliability of resting-state fMRI connectivity estimates. *Neuroimage* 83, 550–558. doi:10.1016/j.neuroimage.2013.05.099. [PubMed: 23747458]
- Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE, 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83, 238–251. doi:10.1016/j.neuron.2014.05.014. [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. doi:10.1016/j.neuroimage.2018.12.054. [PubMed: 30597260]
- Cuthbert BN, Insel TR, 2013. Toward the future of psychiatric diagnosis: the seven pillars of RDoC. *BMC Med.* 11, 126. doi:10.1186/1741-7015-11-126. [PubMed: 23672542]
- Cuthbert BN, Kozak MJ, 2013. Constructing constructs for psychopathology: the NIMH research domain criteria. *J. Abnorm. Psychol* 122, 928–937. doi:10.1037/a0034028. [PubMed: 24016027]
- Dinga R, Schmaal L, Penninx BWJH, van Tol MJ, Veltman DJ, van Velzen L, Mennes M, van der Wee NJA, Marquand AF, 2019. Evaluating the evidence for biotypes of depression: methodological replication and extension of Drysdale et al. (2017). *NeuroImage Clin.* 22, 101796. doi:10.1016/j.nicl.2019.101796.
- Drysdale AT, Gosenick L, Downar J, Dunlop K, Mansouri F, Meng Y, Fetcho RN, Zebly B, Oathes DJ, Etkin A, Schatzberg AF, Sudheimer K, Keller J, Mayberg HS, Gunning FM, Alexopoulos GS, Fox MD, Pascual-Leone A, Voss HU, Casey B, Dubin MJ, Liston C, 2017. Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nat. Med* 23, 28–38. doi:10.1038/nm.4246. [PubMed: 27918562]
- Elliott ML, Knodt AR, Cooke M, Kim MJ, Melzer TR, Keenan R, Ireland D, Ramrakha S, Poulton R, Caspi A, Moffitt TE, Hariri AR, 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. doi:10.1016/j.neuroimage.2019.01.068. [PubMed: 30708106]

- Finn ES, Scheinost D, Finn DM, Shen X, Papademetris X, Constable RT, 2017. Can brain state be manipulated to emphasize individual differences in functional connectivity? *NeuroImage Funct. Archit. Brain* 160, 140–151. doi:10.1016/j.neuroimage.2017.03.064.
- Fischl B, Salat DH, Busa E, Albert M, Dieterich M, Haselgrove C, van der Kouwe A, Killiany R, Kennedy D, Klaveness S, Montillo A, Makris N, Rosen B, Dale AM, 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33, 341–355. [PubMed: 11832223]
- Frank DW, Dewitt M, Hudgens-Haney M, Schaeffer DJ, Ball BH, Schwarz NF, Hussein AA, Smart LM, Sabatinelli D, 2014. Emotion regulation: quantitative meta-analysis of functional activation and deactivation. *Neurosci. Biobehav. Rev* 45, 202–211. doi:10.1016/j.neubiorev.2014.06.010. [PubMed: 24984244]
- Glasser MF, Coalson TS, Bijsterbosch JD, Harrison SJ, Harms MP, Anticevic A, Van Essen DC, Smith SM, 2018. Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data. *Neuroimage* 181, 692–717. doi:10.1016/j.neuroimage.2018.04.076. [PubMed: 29753843]
- Glasser MF, Coalson TS, Robinson EC, Hacker CD, Harwell J, Yacoub E, Ugurbil K, Andersson J, Beckmann CF, Jenkinson M, Smith SM, Van Essen DC, 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178. doi:10.1038/nature18933. [PubMed: 27437579]
- Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, Van Essen DC, Jenkinson MWU-Minn HCP Consortium, 2013. The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124. doi:10.1016/j.neuroimage.2013.04.127. [PubMed: 23668970]
- González I, Déjean S, Martin PGP, Baccini A, 2008. CCA: an R package to extend canonical correlation analysis. *J. Stat. Softw* 23, 1–14. doi:10.18637/jss.v023.i12.
- Hamann S, 2012. Mapping discrete and dimensional emotions onto the brain: controversies and consensus. *Trends Cogn. Sci* 16, 458–466. doi:10.1016/j.tics.2012.07.006, (Regul. Ed.). [PubMed: 22890089]
- Hardoon DR, Szedmak S, Shawe-Taylor J, 2004. Canonical correlation analysis: an overview with application to learning methods. *Neural Comput.* 16, 2639–2664. doi:10.1162/0899766042321814. [PubMed: 15516276]
- Hearne LJ, Mattingley JB, Cocchi L, 2016. Functional brain networks related to individual differences in human intelligence at rest. *Sci. Rep* 6, 32328. doi:10.1038/srep32328. [PubMed: 27561736]
- Ji JL, Spronk M, Kulkarni K, Repovš G, Anticevic A, Cole MW, 2019. Mapping the human brain's cortical-subcortical functional network organization. *Neuroimage* 185, 35–57. doi:10.1016/j.neuroimage.2018.10.006. [PubMed: 30291974]
- Leurgans SE, Moyeed RA, Silverman BW, 1993. Canonical correlation analysis when the data are curves. *J. R. Stat. Soc. Ser. B* 55, 725–740. doi:10.1111/j.2517-6161.1993.tb01936.x, (Methodological).
- Li N, Ma N, Liu Y, He XS, Sun DL, Fu XM, Zhang X, Han S, Zhang DR, 2013. Resting-state functional connectivity predicts impulsivity in economic decisionmaking. *J. Neurosci* 33, 4886–4895. doi:10.1523/JNEUROSCI.1342-12.2013. [PubMed: 23486959]
- Lindquist KA, Wager TD, Kober H, Bliss-Moreau E, Barrett LF, 2012. The brain basis of emotion: a meta-analytic review. *Behav. Brain Sci* 35, 121–143. doi:10.1017/S0140525X11000446. [PubMed: 22617651]
- Mauss IB, Robinson MD, 2009. Measures of emotion: a review. *Cognit. Emot* 23, 209–237. doi:10.1080/02699930802204677. [PubMed: 19809584]
- Mihalik A, Ferreira FS, Moutoussis M, Ziegler G, Adams RA, Rosa MJ, Prabhu G, de Oliveira L, Pereira M, Bullmore ET, Fonagy P, Goodyer IM, Jones PB, Hauser T, Neufeld S, Romero-Garcia R, St Clair M, Vértes PE, Whitaker K, Inkster B, Ooi C, Toseeb U, Widmer B, Bhatti J, Villis L, Alrumaithi A, Birt S, Bowler A, Cleridou K, Dadabhoy H, Davies E, Firkins A, Granville S, Harding E, Hopkins A, Isaacs D, King J, Kokorikou D, Maurice C, McIntosh C, Memarzia J, Mills H, O'Donnell C, Pantaleone S, Scott J, Fearon P, Suckling J, van Harmelen AL, Kievit R, Shawe-Taylor J, Dolan R, Mourão-Miranda J, 2020. Multiple Holdouts with stability: improving the generalizability of machine learning analyses of brain-behavior relationships. *Biol.*

- Psychiatry Innov. Clin. Neurosci. Tools Tech. Transform. Framew 87, 368–376. doi:10.1016/j.biopsych.2019.12.001.
- Murphy FC, Nimmo-Smith I, Lawrence AD, 2003. Functional neuroanatomy of emotions: a meta-analysis. *Cogn. Affect. Behav. Neurosci* 3, 207–233. doi:10.3758/CABN.3.3.207. [PubMed: 14672157]
- Noble S, Scheinost D, Constable RT, 2019. A decade of test-retest reliability of functional connectivity: a systematic review and meta-analysis. *Neuroimage* 203, 116157. doi:10.1016/j.neuroimage.2019.116157.
- Prater KE, Hosanagar A, Klumpp H, Angstadt M, Phan KL, 2013. Aberrant amygdala-frontal cortex connectivity during perception of fearful faces and at rest in generalized social anxiety disorder. *Depress. Anxiety* 30, 234–241. doi:10.1002/da.22014. [PubMed: 23184639]
- Rosenberg MD, Finn ES, Scheinost D, Papademetris X, Shen X, Constable RT, Chun MM, 2016. A neuromarker of sustained attention from whole-brain functional connectivity. *Nat. Neurosci* 19, 165–171. doi:10.1038/nn.4179. [PubMed: 26595653]
- Rubinov M, Sporns O, 2010. Complex network measures of brain connectivity: uses and interpretations. *NeuroImage Comput. Models Brain* 52, 1059–1069. doi:10.1016/j.neuroimage.2009.10.003.
- Salimi-Khorshidi G, Douaud G, Beckmann CF, Glasser MF, Griffanti L, Smith SM, 2014. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468. doi:10.1016/j.neuroimage.2013.11.046. [PubMed: 24389422]
- Schäfer T, Schwarz MA, 2019. The meaningfulness of effect sizes in psychological research: differences between sub-disciplines and the impact of potential biases. *Front. Psychol* 10. doi:10.3389/fpsyg.2019.00813.
- Smith SM, Nichols TE, Vidaurre D, Winkler AM, Behrens TEJ, Glasser MF, Uğurbil K, Barch DM, Van Essen DC, Miller KL, 2015. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nat. Neurosci* 18, 1565–1567. doi:10.1038/nn.4125. [PubMed: 26414616]
- Termenon M, Jaillard A, Delon-Martin C, Achard S, 2016. Reliability of graph analysis of resting state fMRI using test-retest dataset from the human connectome project. *Neuroimage* 142, 172–187. doi:10.1016/j.neuroimage.2016.05.062. [PubMed: 27282475]
- Tozzi L, Fleming SL, Taylor Z, Raterink C, Williams LM, 2020. Test-retest reliability of the human functional connectome over consecutive days: identifying highly reliable portions and assessing the impact of methodological choices. *Netw. Neurosci* doi:10.1162/netn\_a\_00148.
- Tuzhilina E, Tozzi L, Hastie T, 2020. Canonical correlation analysis in high dimensions with structured regularization. arXiv:2011.01650 [stat].
- Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Uğurbil K, Consortium WMH, 2013. The WU-minn human connectome project: an overview. *Neuroimage* 80, 62–79. doi:10.1016/j.neuroimage.2013.05.041. [PubMed: 23684880]
- Vinod HD, 1976. Canonical ridge and econometrics of joint production. *J. Econom* 4, 147–166. doi:10.1016/0304-4076(76)90010-5.
- Westlund Schreiner M, Klimes-Dougan B, Mueller BA, Eberly LE, Reigstad KM, Carstedt PA, Thomas KM, Hunt RH, Lim KO, Cullen KR, 2017. Multi-modal neuroimaging of adolescents with non-suicidal self-injury: amygdala functional connectivity. *J. Affect. Disord* 221, 47–55. doi:10.1016/j.jad.2017.06.004. [PubMed: 28628767]
- Whiteford HA, Degenhardt L, Rehm J, Baxter AJ, Ferrari AJ, Erskine HE, Charlson FJ, Norman RE, Flaxman AD, Johns N, Burstein R, Murray CJL, Vos T, 2013. Global burden of disease attributable to mental and substance use disorders: findings from the global burden of disease study 2010. *Lancet* 382, 1575–1586. doi:10.1016/S0140-6736(13)61611-6. [PubMed: 23993280]
- Williams LM, 2016. Precision psychiatry: a neural circuit taxonomy for depression and anxiety. *Lancet Psychiatry* 3, 472–480. doi:10.1016/S2215-0366(15)00579-9. [PubMed: 27150382]
- Woody ML, Gibb BE, 2015. Integrating NIMH research domain criteria (RDoC) into depression research. *Curr. Opin. Psychol* 4, 6–12. doi:10.1016/j.copsyc.2015.01.004. [PubMed: 25642446]



- Xia CH, Ma Z, Ciric R, Gu S, Betzel RF, Kaczkurkin AN, Calkins ME, Cook PA, de la Garza AG, Vandekar SN, Cui Z, Moore TM, Roalf DR, Ruparel K, Wolf DH, Davatzikos C, Gur RC, Gur RE, Shinohara RT, Bassett DS, Satterthwaite TD, 2018. Linked dimensions of psychopathology and connectivity in functional brain networks. *Nat. Commun* 9, 1–14. doi:10.1038/s41467-018-05317-y. [PubMed: 29317637]
- Zhu P, Knyazev AV, 2013. Angles between subspaces and their tangents. *J. Numer. Math* 21, 325–340. doi:10.1515/jnum-2013-0013.




**Fig. 1.**

Summary of the canonical correlation analysis. We divided the data into train and test sets. Then, we chose regularization parameters by 10-fold cross-validation on the training set as the ones maximizing the first canonical correlation across folds. Then, we validated the canonical correlation weights obtained from the training set in the test set using a permutation test. Finally, we investigated canonical pairs beyond the first by calculating an aggregated correlation between canonical subspaces and also validating it using permutation testing. See methods for details. HCP=Human Connectome Project.


### Regularized canonical correlation analysis

max correlation = 0.148 for  $\lambda=100$



### Group regularized canonical correlation analysis

max correlation = 0.158 for  $\mu=10000$  and  $\lambda=100$



**Fig. 2.** RCCA and GRCCA cross-validation for Rest data to select hyperparameters maximizing the first canonical correlation. Top: mean correlation and standard error confidence intervals (y-axis) for the first canonical component for 9 different values of  $\lambda = 10^{-6}, 10^{-5}, \dots, 10^2$  (x-axis) across the train (left) and test (right) splits generated by the cross-validation procedure using RCCA. Bottom: mean correlation and standard error confidence intervals (y-axis) for the first canonical component for 9 different values of  $\lambda = 10^{-4}, 10^{-3}, \dots, 10^4$  (x-axis) and 9 different values of  $\mu = 10^{-4}, 10^{-3}, \dots, 10^4$  (line colors) across the train (left) and test (right) splits generated by the cross-validation procedure using GRCCA.

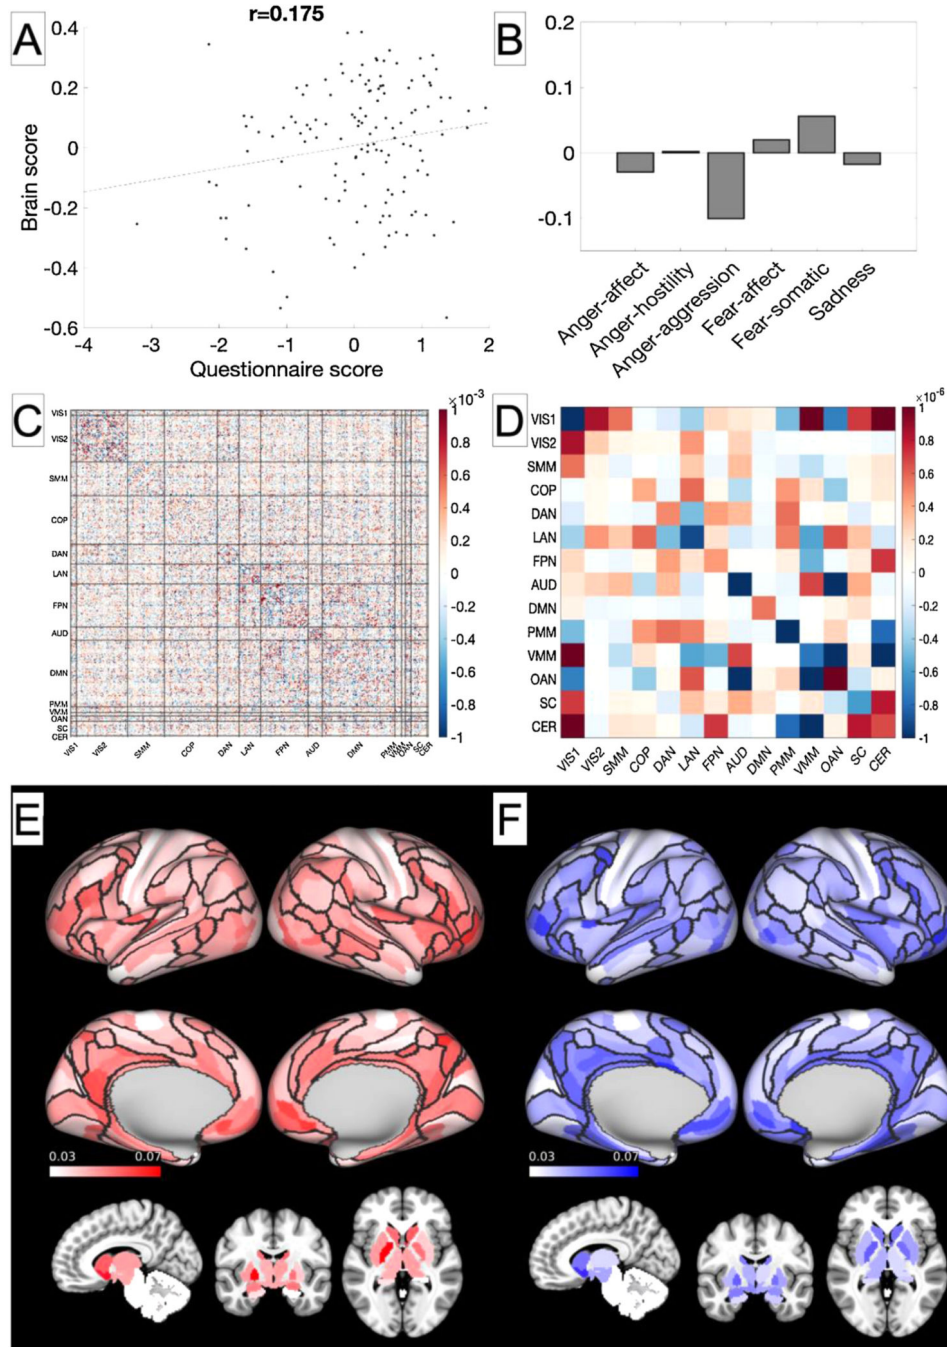
GRCCA=group regularized canonical correlation analysis, RCCA=regularized canonical correlation analysis.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



**Fig. 3.** Visualization of Rest canonical pair 1. A: correlation between the scores derived for the brain and questionnaire features. B: weights given to the questionnaires. C: matrix of the weights given to each edge of the connectivity matrix, ordered by network. D: the same weights averaged by network. E: sum of the positive weights of connections involving each region (node strength). F: sum of the negative weights of connections involving each region (node strength). VIS1=visual primary, VIS2=visual secondary, SMM=somatomotor, COP=cingulo-opercular, DAN=dorsal attention, LAN=lateral attention, FPN=fronto-

parietal, AUD=auditory, DMN=default mode, PMM=posterior multimodal, VMM=ventral multimodal, OAN=orbito-affective, SC=subcortical, CER=cerebellum. The 14 functional network boundaries are overlaid in black.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 1**

Sample characteristics.

<b>N</b>	<b>652</b>
Sex	<i>F</i> = 359, <i>M</i> = 293
Age	28.802 ± 3.663
Anger-affect	47.743 ± 8.197
Anger-hostility	49.885 ± 8.501
Anger-aggression	51.711 ± 8.913
Fear-affect	50.082 ± 8.015
Fear-somatic	51.783 ± 8.250
Sadness	46.324 ± 7.948

Mean and standard deviation are given for each continuous measure.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Table 2**

Questionnaire weights for the Rest canonical variate generalizing to the held-out test data.

Questionnaire	Weight
Anger-affect	-0.029
Anger-hostility	0.002
Anger-aggression	-0.101
Fear-affect	0.020
Fear-somatic	0.056
Sadness	-0.018

Loadings of the questionnaires of the NIH Emotion toolbox returned by GRCCA for the only component that survived permutation testing in the held-out test set. GRCCA=group regularized canonical correlation analysis.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



**Table 3**

Network weights for the Rest canonical variate generalizing to the held-out test data.

	Positive network strength	Negative network strength
Visual primary	5.55E-06	1.86E-06
Visual secondary	1.83E-06	3.72E-07
Somatomotor	1.60E-06	8.46E-07
Cingulo-opercular	1.64E-06	9.81E-07
Dorsal attention	1.52E-06	1.16E-06
Lateral attention	3.03E-06	2.03E-06
Fronto-parietal	2.01E-06	7.56E-07
Auditory	2.23E-06	2.30E-06
Default mode	4.45E-07	6.62E-07
Posterior multimodal	1.94E-06	1.77E-06
Ventral multimodal	2.32E-06	4.06E-06
Orbito-affective	1.20E-06	4.31E-06
Subcortical	2.95E-06	5.25E-07
Cerebellum	3.69E-06	2.60E-06

Loadings of the functional brain connections returned by GRCCA for the only component that survived permutation testing in the held-out test set. For brain data, we computed the average weights of connections within and between each network and summed the weights of the resulting connections involving each network (strength). GRCCA=group regularized canonical correlation analysis.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript