

1 **Full Title:**

2 Decoding task-specific cognitive states with slow, directed functional networks in the
3 human brain

4 **Short title:** Directed fMRI functional networks for decoding cognitive states

5 **Authors:**

6 Devarajan Sridharan^{1,2}, Shagun Ajmera¹, Hritik Jain¹ and Mali Sundaresan¹

7 **Affiliations:**

8 ¹Centre for Neuroscience, Indian Institute of Science, Bangalore 560012, India

9 ²Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560012, India

10 **Corresponding author:**

11 Devarajan Sridharan

12 E-mail: sridhar@iisc.ac.in (DS)

13 **Keywords:**

14 functional connectivity; Granger causality; machine learning; support vector machines; fMRI decoding;

15 emergent dynamics; cognitive score prediction

16 Abstract

17 Flexible functional interactions among brain regions mediate critical cognitive functions. Such interactions
18 can be measured from functional magnetic resonance imaging (fMRI) data with either instantaneous (zero-
19 lag) or lag-based (time-lagged) functional connectivity; only the latter approach permits inferring directed
20 functional interactions. Yet, the fMRI hemodynamic response is slow, and sampled at a
21 timescale (seconds) several orders of magnitude slower than the underlying neural dynamics
22 (milliseconds). It is, therefore, widely held that lag-based fMRI functional connectivity, measured with
23 approaches like as Granger-Geweke causality (GC), provides spurious and unreliable estimates
24 of underlying neural interactions. Experimental verification of this claim has proven challenging because
25 neural ground truth connectivity is often unavailable concurrently with fMRI recordings. We address this
26 challenge by combining machine learning with GC functional connectivity estimation. We estimated
27 instantaneous and lag-based GC functional connectivity networks using fMRI data from 1000 participants,
28 drawn from the Human Connectome Project database. A linear classifier, trained on either instantaneous or
29 lag-based GC, reliably discriminated among seven different task and resting brain states, with over 80%
30 cross-validation accuracy. With network simulations, we demonstrate that instantaneous and lag-based
31 GC exploited interactions at fast and slow timescales, respectively, to achieve robust classification. With
32 human fMRI data, instantaneous and lag-based GC identified distinct, cognitive core networks. Finally,
33 variations in GC connectivity explained inter-individual variations in a variety of cognitive scores. Our
34 findings show that instantaneous and lag-based methods reveal complementary aspects of functional
35 connectivity in the brain, and suggest that slow, directed functional interactions, estimated with fMRI,
36 provide robust markers of behaviorally relevant cognitive states.

38 Author Summary

39 Functional MRI (fMRI) is a leading, non-invasive technique for mapping networks in the human brain. Yet,
40 fMRI signals are noisy and sluggish, and fMRI scans are acquired at a timescale of seconds, considerably
41 slower than the timescale of neural spiking (milliseconds). Can fMRI, then, be used to infer dynamic
42 processes in the brain such as the direction of information flow among brain networks? We sought to
43 answer this question by applying machine learning to fMRI scans acquired from 1000 participants in the
44 Human Connectome Project (HCP) database. We show that directed brain networks, estimated with a
45 technique known as Granger-Geweke Causality (GC), accurately predicts individual subjects' task-specific
46 cognitive states inside the scanner, and also explains variations in a variety of behavioral scores across
47 individuals. We propose that directed functional connectivity, as estimated with fMRI-GC, is relevant
48 for understanding human cognitive function.

49 Introduction

50 Coordinated activity among brain regions underlies a variety of cognitive functions [1]. Mapping functional
51 coupling among brain regions is, therefore, key to mapping brain function and for understanding how the
52 brain produces behavior. Human fMRI studies have commonly investigated such functional coupling with
53 correlation-based measures, including the Pearson correlation coefficient [2,3] and partial correlations [4]
54 between pairs of brain regions. These measures frequently incorporate regularization penalties to estimate
55 sparse functional networks [5]. Correlation-based measures are ideal for characterizing “instantaneous”
56 functional coupling, representing functional interactions among brain regions that occur at timescales faster
57 than the sampling rate of the measurement [6]. In contrast, comparatively few studies, have examined lag-
58 based measures of functional connectivity to examine time-lagged interactions among brain regions [7,8].

59

60 Measures of linear dependence and feedback, based on Granger-Geweke causality (GC) [9,10] represent
61 a powerful approach for estimating both instantaneous and lag-based functional connectivity. These
62 measures are firmly grounded in information theory and statistical inferential frameworks [9–11]. GC
63 measures have been widely applied to estimate functional connectivity in recordings of brain activity made
64 with electroencephalography (EEG; [12]), magnetoencephalography (MEG; [13]) and electrocorticography
65 (ECoG; [14]). However, the application of GC measures to brain recordings made with functional magnetic
66 resonance imaging (fMRI) has provoked significant controversy [15–18]. Because the hemodynamic
67 response is produced and sampled at a timescale (seconds) several orders of magnitude slower than the
68 underlying neural processes (milliseconds), previous studies have argued that lag-based measures,
69 particularly lag-based GC, produce spurious and unreliable estimates of functional connectivity, when
70 applied to fMRI data (fMRI-GC) [18–21].

71

72 Three primary confounds have been identified with inferring connectivity with fMRI-GC. First, systematic
73 differences in hemodynamic lags across regions could yield spurious directionality of GC connections
74 [16,18]. Second, in simulations, measurement noise added to the signal during fMRI acquisition
75 significantly degrades GC functional connectivity estimates [19]. Finally, downsampling recordings to the

76 typical fMRI sampling rate (seconds), three orders of magnitude slower than the timescale of neural spiking
77 (milliseconds), effectively eliminates all traces of functional connectivity inferred by GC [19].

78
79 The controversy regarding the application of GC to fMRI data continues to date, primarily because of the
80 lack of access to ground-truth in neural data. On the one hand, claims regarding the efficacy of GC
81 estimates are primarily based on simulations [11,22], and are only as valid as the underlying model of
82 neural activity and hemodynamic responses. Because the precise mechanism by which neural responses
83 generate hemodynamic responses is an active area of research, strong conclusions cannot be drawn
84 based on fMRI simulations alone. On the other hand, establishing ground-truth validity for fMRI functional
85 connectivity requires invasive neurophysiological recordings across many brain regions, concurrently
86 during fMRI scans, a prohibitive enterprise.

87
88 We seek to address this controversy by applying machine learning to fMRI-GC networks, which works
89 around these challenges. We estimated instantaneous and lag-based GC connectivity with fMRI data
90 drawn from 1000 human subjects, recorded under seven different task conditions and in the resting state
91 (~8000 functional scans drawn from the Human Connectome Project database; [23]). We trained a linear
92 classifier, based on GC connectivity features alone, to discriminate among the different task and resting
93 conditions, and assessed classifier accuracy with cross validation. The results show that instantaneous and
94 lag-based GC connectivity, estimated from fMRI data, can decode task-specific cognitive states with
95 superlative accuracies. Next, with simulations, we show that slow, multi-second timescales emerge in
96 sparse, random networks despite individual neurons operating at fast, millisecond timescales – a result that
97 explains why directed functional connectivity can be reliably estimated with GC in slowly sampled fMRI
98 data. Finally, we demonstrate that GC connectivity features can be used as predictors [24] to explain inter-
99 individual variations in behavioral scores across a variety of cognitive tests. The results suggest that
100 instantaneous and lag-based GC measures applied to fMRI data permit mapping slow, emergent and
101 behaviorally relevant functional interactions in the human brain.

Results

GC estimated from slowly sampled fMRI data suffices to distinguish task and resting states

We asked if instantaneous GC (iGC) and directed GC (dGC) (SI Mathematical Note Section S1) connectivity would flexibly reconfigure with task demand, by testing if GC connectivity sufficed to accurately classify among seven different task states or the resting state (SI Table S1A; Methods; [9,10]). Data were obtained from 1000 participants from the Human Connectome Project (HCP) database [25] (RRID: SCR_008749). We used connection weights among brain regions in each network (iGC or dGC) as feature vectors in a linear classifier based on Support Vector Machines (SVM) for high dimensional predictor data. Accuracies for classifying resting state from a working memory task (WM task) are described first; accuracies for other tasks are presented subsequently.

Both iGC and dGC connectivity were able to distinguish the working memory task from resting state significantly above chance (Fig. 1B; $p < 0.001$, permutation test). Maximum accuracy (median, [95% CI]) was 97.3% [96.3 - 98.0%] with iGC and 92.0% [90.5 - 93.2%] with dGC (SI Fig. S1B Yeo Parcellation, iGC: precision= 97.2, recall= 97.4; dGC: precision= 90.9, recall= 93.2). k-fold ($k=10$) cross-validation accuracy was comparable (iGC: 97.1% [96.2 - 97.9%], dGC: 91.7% [90.3 - 93.0%]). These numbers correspond to maximum cross validation accuracy across all five parcellations tested (SI Table S3; SI Fig. S1A); accuracies with each parcellation are shown in the Supporting Information (SI Fig. S1B). Non-linear classifiers, such as SVMs based on radial basis function kernels produced similar results, with comparably above chance classification accuracy for both iGC and dGC connectivity (SI Fig. S1C).

Figure 1. Discriminating between task and resting state with instantaneous and directed GC networks.

A. Schematic of task state classification based on instantaneous (iGC) and directed (dGC) Granger-Geweke Causality with fMRI data from 1000 subjects (see text for details).

B. Two-way classification accuracies (leave-one-out) for each of seven tasks versus resting state based on GC. Red unfilled bars and blue filled bars: accuracies based on dGC and iGC features, respectively (task

129 key in SI Table S1). Error-bars: Clopper-Pearson binomial confidence intervals. Chance accuracy: 0.5 (not
130 shown).

131 **C.** Two-way task versus resting state classification accuracies based on dGC (red dots) and iGC (blue -
132 dots), as a function of number of task scan time points (volumes). Dashed lines: linear fits.

133 **D.** Two-way task versus resting state classification accuracies based on dGC after averaging dGC matrices
134 over different numbers of subjects (x-axis). Each task is represented with a different color. Colored dashed
135 lines: biexponential fits. Black dashed horizontal and vertical lines: 95% accuracy and n=5 subjects'
136 average, respectively.

137 **E.** Two-way classification accuracies across each pair of tasks. Cells: classification accuracies for each pair
138 of tasks based on dGC (lower triangular matrix) or iGC (upper triangular matrix). Diagonal cells: number of
139 task scan timepoints. Highlighted cells: lowest (dashed-line border) and highest (solid-line border)
140 accuracies achieved with dGC (red) and iGC (blue).

141 **F.** N-way classification accuracies among all seven tasks. Dashed line: chance accuracy (14.3%). Other
142 conventions are the same as in panel B.

143 **G.** Two-way sub-task classification accuracies (SI Table S1B) based on GC. ns.: accuracy not significantly
144 above chance. Other conventions are the same as in panel B.

145 **H.** (Left) Two-way task versus resting state classification accuracies obtained with regional time series sub-
146 sampled at 2x (filled symbols) and 3x (open symbols) of the TR (720 ms) (y-axis) plotted against
147 accuracies obtained with the original data (1x, x-axis) for each of 7 tasks. Red: dGC, Blue: iGC. Dashed
148 diagonal line: Line of equality ($x=y$). (Right) N-way classification accuracies among all seven tasks with data
149 sampled at 1x, 2x, 3x of the original TR. Other conventions are the same as in panel F. For panels B,E,F:
150 accuracies correspond to highest values, across all parcellations tested, and hyperparameter optimization
151 was done for panel B. For panels C,G,H: accuracies correspond to Shirer et al [26] 14-network parcellation.
152 For panel D: accuracies correspond to Shirer et al [26] 90-node parcellation.

153

154 We repeated these analyses by classifying the six other tasks (SI Table S1) versus resting state. iGC and
155 dGC connectivity could accurately classify each task from resting state significantly above chance. For iGC,
156 maximum classification accuracies ranged from 90.1%, for emotion task vs. resting state classification, to
157 97.1%, for language task vs. resting state classification. Similarly, for dGC, accuracies ranged from 78.1%,

158 for emotion task vs. resting state classification, to 92.8%, for language task vs. resting state classification
159 (Fig. 1B). In general, classification accuracy increased with more scan timepoints for each task versus
160 resting state classification (Fig. 1C), consistent with GC being an information theoretic measure; we
161 confirmed this result with simulations also (SI Fig. S1D).

162
163 In these analyses, classification accuracies based on dGC were systematically lower than those based on
164 iGC. We asked if dGC accuracies were poorer due to noise corrupting the fit of the autoregressive model,
165 and if a more consistent estimate could be obtained by averaging dGC connectivity features, to remove
166 uncorrelated noise, across subjects. We addressed this question by partitioning the data into two groups --
167 a training (T) group and a test (S) groups – with 500 subjects each. We trained the classifier on group T
168 and tested the classifier prediction by averaging GC matrices across several folds of S, each fold
169 containing a few ($m=2, 4, 5, 10, 25$ or 50) subjects; the procedure was repeated by exchanging training and
170 test datasets (see Methods). For the vast majority of tasks (6/7), dGC's classification accuracy was more
171 than 95% with as few as $m=5$ subjects within each fold of the test set (Fig. 1D). These results suggest that
172 averaging dGC matrices across a few subjects, yielded reliable estimates of dGC connectivity.

173
174 We considered other factors that, in addition to intrinsic connectivity differences, could have produced
175 these superior classification accuracies. First, GC-based accuracies for classifying task versus resting state
176 scans might arise from differences in brain regions activated during each of these scans. In addition to
177 task-relevant sensory input, overt motor responses always occurred during task scans but were absent
178 during resting state scans [23,27]. Could GC features discriminate among more subtle connectivity
179 variations across/within tasks? Second, scan data from the HCP database was sampled at a TR (repetition
180 time) of 720 ms, considerably faster than the TR for conventional fMRI scans. Would GC accuracies
181 degrade if the data were sampled at much slower sampling rate (~2000 ms), in line with conventional fMRI
182 TR?

183
184 We addressed the first question in two stages. First, we asked if GC connectivity features would be able to
185 classify which of the seven tasks each subject was performing in the scanner. First, we performed a
186 pairwise classification of each task from the other. Maximum classification accuracies for iGC (dGC) ranged

187 from 87% (67%) for the emotion vs. gambling task classification to 98% (91%) for the language vs. social
188 task classification. Again, the number of timepoints for each task proved to be a strong indicator of
189 classification accuracies (Fig. 1E): average inter-task classification accuracies were highest for the
190 language task (iGC: 97%, dGC:88%, n=316 timepoints) and lowest for the emotion task (iGC: 91%, dGC:
191 77%, n=176 timepoints). Next, we performed an n-way classification analysis across all 7 tasks, again
192 using linear SVM (Methods). Accuracies were significantly above chance (14.3% for 1-in-7 classification)
193 for classifying among the seven tasks (Fig. 1F; maximum accuracy, iGC: 74.4% [73.3%-75.4%]; dGC:
194 47.6% [46.4%-48.7%]; $p < 0.001$, permutation test). These results indicate that functional connectivity was
195 consistently estimated with GC, and reliably different across tasks.

196
197 Second, each of the different tasks in the HCP database comprised of blocks of contiguous trials, each
198 corresponding to one of (at least) two different sub-tasks ([27]; SI Table S1B). For example, the motor task
199 comprised of blocks of movements of the right or left hand interleaved with blocks of trials involving
200 movement of the right or left foot. Similarly, the working memory task comprised of interleaved blocks of 0-
201 back and 2-back tasks. We asked, therefore, if GC connectivity could distinguish among subtler variations
202 in brain states across sub-tasks within each task. We sought to classify across two sub-tasks for each of six
203 tasks (SI Table S1B). In all cases, except one, both iGC and dGC connectivity discriminated between each
204 pair of sub-tasks with higher than chance accuracies (Fig. 1G; maximum accuracy, iGC: 89.2% [87.6% -
205 90.7%]; dGC: 80.1% [78.9% - 82.9%]; $p < 0.05$ permutation test). These results indicate that GC functional
206 connectivity could accurately distinguish among sub-tasks within each task as well.

207
208 Next, we tested whether GC connectivity estimated from slowly sampled fMRI data could accurately
209 classify task and resting states. We downsampled the data to either one half (2xTR=1440 ms) or one third
210 (3x TR=2160 ms) of its original sampling rate, by decimation, while also concatenating the decimated data
211 to the end of the sub-sampled timeseries to preserve the overall number of timepoints (Methods). We
212 repeated both of the previous classification analyses – pairwise task versus resting state classification (Fig.
213 1H left), as well as n-way inter-task classification (Fig. 1H right). Following downsampling, we observed that
214 classification accuracies were marginally higher than accuracies in the original data in the case of dGC (2x:
215 $p = 0.02$; 3x: $p = 0.06$; Wilcoxon one-tailed signed rank test), and were even higher than those in the original

216 data in the case of iGC (2x: $p=0.01$; 3x: $p=0.01$), across tasks. These results indicate that the superlative
217 sampling rate of the HCP fMRI data was not the primary reason for these high classification accuracies for
218 GC-based classification.

219
220 We performed three control analyses to further confirm these results: i) by performing stationarity tests on
221 the data prior to GC estimation and classification; ii) using single full regression to estimate GC [28,29],
222 instead of estimating with separate full and reduced regressions; and iii) incorporating motion scrubbing
223 [30] to ensure that the classification accuracies were not driven by head motion artifacts. These controls
224 are described in SI Results, section on “Control analyses”. In every case, we obtained equivalent or
225 superior classification accuracies (SI Fig. S2), confirming that the results were not due to data non-
226 stationarity, biases in GC estimation or head motion artifacts.

227
228 These results demonstrate that both iGC and dGC yielded task-specific signatures of functional
229 connectivity even with slowly sampled fMRI data ($TR\sim 2000$ ms): these estimates were consistent across
230 subjects and reliably different across tasks to permit successful classification. Furthermore, these
231 superlative classification accuracies were obtained despite widely held caveats concerning the application
232 of GC to fMRI data [28].

233 234 235 **Correlation-purged GC connectivity suffices for accurate task-state classification**

236 Correlation-based (zero-lag) connectivity measures (e.g. partial correlations or PC) have been widely
237 applied to estimate functional connectivity from fMRI data [5,31]. In fact, several previous studies [18,19]
238 have argued that correlation-based measures are more reliable and should be preferred to lag-based
239 measures like GC [11], for estimating functional connectivity with fMRI data. We tested this claim here with
240 a three-fold analysis approach.

241
242 First, we asked how classification accuracies based on PC connectivity would compare with those reported
243 above, based on GC connectivity. Maximum classification accuracies with PC connectivity ranged from 96-
244 99% for task versus resting state classification, and were consistently higher than accuracies with GC

connectivity (Fig. 2A). These results are along expected lines: estimators based on same-time covariance, such as PC, are less susceptible to noise than those based on lagged covariance, such as GC (derived analytically in the Methods, section on *Functional connectivity estimation and classification with partial correlations*). In addition, as mentioned previously, GC is an information theoretic measure: classification accuracy with iGC and dGC increased systematically with more scan time points, asymptotically matching PC accuracies (SI Fig. S1D).

Figure 2. Classification accuracies with GC purged of instantaneous correlations.

A. Two-way task versus resting state classification accuracies, based on partial correlations (PC; grey unfilled bars). Numbers reported correspond to highest leave-one-out classification accuracies across parcellations, obtained with hyperparameter optimization. Corresponding accuracies for dGC (red dots) and iGC (blue dots) are shown for comparison. Other conventions are as in Fig. 1B.

B. Schematic illustrating procedure for purging data of instantaneous correlations. fMRI regional timeseries were purged of instantaneous correlations by either whitening the data with zero-phase component analysis (ZCA), separately for each task and resting state scan, or by projecting data into a space spanned by the generalized eigenvectors (GEV), common to both task and resting state scans. GC and PC were then estimated with the ZCA or GEV projections of the timeseries data, followed by classification analysis based on GC or PC connection strength as features.

C. (Top) Two-way task versus resting state classification accuracies following ZCA-based decorrelation. Gray circles: Classification accuracies based on PC. Other conventions are as in Figure 1B. Dashed line: chance accuracy (50%).(Bottom) Same as in top panel, but for classification following GEV-based decorrelation.

D. (Top) Schematic showing unweighted directed graph obtained from dGC; this digraph representation encodes only the dominant direction of connectivity, and not its magnitude. (Bottom) Two-way task versus resting state classification accuracies based on dGC digraph representations. Secondary ordinate (y-axis on the right): number of scan timepoints for each task.(Panels C-D). GC features were estimated with the Shirer et al [26] 14-network parcellation.

273 Second, we asked if lag-based connectivity could accurately classify task from resting state, once the data
274 were purged of all instantaneous correlations. To accomplish this, we adopted two approaches: i) zero-
275 phase component analysis (ZCA) and ii) generalized eigenvalue decomposition (GEV) (Methods). Briefly,
276 ZCA (or the Mahalanobis transformation) produces whitened time series data that is closest, in a least
277 squares sense, to the original regional time series data. As an alternative approach, we decorrelated both
278 task and resting state time series jointly by projecting them onto a single set of generalized eigenvectors
279 (GEV). These approaches provided empirical upper and lower bounds on GC's performance on correlation-
280 purged data (Methods).

281
282 GC connectivity features sufficed to successfully classify all tasks from resting state, even in correlation-
283 purged data. With ZCA, iGC accuracies ranged from 84% to 96% whereas dGC accuracies ranged from
284 82% to 96% across tasks. With GEV, iGC accuracies ranged from 60% to 71% whereas dGC accuracies
285 ranged from 56% to 76% across tasks; in each case, classification accuracies were significantly above
286 chance ($p < 0.001$, permutation test). We confirmed that performance in each case was not an artifact of the
287 decorrelation procedure (ZCA/GEV) by randomly interchanging task and resting state labels for each pair of
288 datasets across subjects (Methods); shuffling labels reduced classification accuracy to chance. Note that in
289 every case, classification performance based on PC connectivity was at chance (Fig. 2C), a direct
290 consequence of removing instantaneous correlations from the data. Despite this, classification accuracies
291 based on iGC connectivity were not at chance; in the next section, we discuss potential reasons for these
292 differences between iGC and PC classification accuracies.

293
294 Third, we asked if an unweighted directed graph (digraph) network representation – whose edges indicated
295 the dominant direction, but not magnitude, of connectivity (Fig. 2D) – would suffice to distinguish task from
296 resting brain states (Methods). Again, dGC directed graphs successfully distinguished each task from
297 resting state well above chance. Classification accuracies ranged from 56% for the motor task versus
298 resting state classification to 68% for the relational task versus resting state; for each task, classification
299 accuracies were significantly above chance ($p < 0.001$; permutation test). Interestingly, we did not see a
300 strong influence of the number of data points on classification accuracy in this case (Fig. 2D, purple dots).
301 For instance the emotion task ($n=176$ timepoints) was classified with an accuracy of 62% from resting state,

302 which was comparable to the classification accuracy of working memory (n=405 timepoints) from resting
303 state (64%). Both iGC and PC, which are symmetric connectivity measures, could provide no directed
304 connectivity information.

305
306 These results demonstrate that lag-based connectivity contained sufficient information to classify task from
307 resting state even when instantaneous correlations were entirely purged from the data. Moreover,
308 unweighted directed connectivity graphs alone, indicating the direction, but not scalar magnitude, of GC
309 connectivity, sufficed to accurately classify task from resting brain states. These findings indicate that
310 directed functional connectivity measures, like dGC, provide connectivity information that is distinct from,
311 and complementary to, what can be obtained with undirected functional connectivity measures, like PC.

312 313 **Instantaneous and directed GC identify complementary aspects of functional connectivity**

314 What characteristics of functional connectivity are respectively identified by instantaneous and lag-based
315 connectivity? And how can lag-based connectivity be reliably estimated with fMRI data, which is sampled at
316 time scales orders of magnitude slower than neural timescales? We addressed both of these questions,
317 first, with simulations (this section) and, then, with real data (next section).

318
319 First, we tested the ability of GC to reliably recover functional interactions in simple, two-node feedforward
320 networks operating at different timescales (Fig. 3A). We simulated fMRI data using a two-stage model
321 (Methods): i) a latent variable model that describes the dynamics of the nodes (vector Ornstein-Uhlenbeck
322 process; [32]); ii) a convolution of these neural dynamics with a hemodynamic response function to obtain
323 the simulated fMRI time series [18,19]. Based on this model, we simulated activity in two 2-node networks.
324 In the first network, individual node decay timescales were set to 50 ms, whereas in the second network,
325 these were set to 1000 ms (parameters in SI Table S6A). For convenience, we refer to these two network
326 timescales as “fast” (50 ms) and “slow” (1000 ms). We then varied the sampling interval (T_s) of the
327 simulated data from 50 ms to 1450 ms in steps of 100 ms. Connections at both “fast” and “slow” timescales
328 were generally discovered by iGC regardless of sampling interval, although connections at slow timescales
329 were less robustly detected than those at fast timescales (Fig. 3A). On the other hand, the connection in

330 the “fast” timescale network was not discovered by dGC when the sampling interval was higher than 50 ms,
331 in line with the results of Smith et al [18]. However, the connection in the “slow” timescale network was
332 reliably discovered by dGC across a wide range of sampling intervals, upto, and exceeding 1000 ms. In
333 each case, dGC failed to discover the underlying interaction when the sampling interval was much higher
334 than the slowest timescale in each network, consistent with recent theoretical results [6]. These findings
335 suggest that dGC can detect slow neural processes, which operate at a timescale slower than TR, in fMRI
336 data.

337
338 **Figure 3. Robustness of GC estimates depend on network timescales in simulated hemodynamic**
339 **data.**

340 **A.** (Top) Two-node networks with fast (50 ms; left) or slow (1000 ms; right) decay timescales of individual
341 nodes. Each subpanel shows ground truth connectivity either as a schematic (left) or connectivity matrix
342 (right). In the matrix, a non-zero entry at cell (i, j) corresponds to a connection from node j (source) to node
343 i (destination). (Bottom) dGC (red), iGC (blue), and PC (black) connection strengths as a function of
344 sampling intervals. Filled circles and solid lines: Strengths of true connections and curve (biexponential)
345 fits, respectively. Open circles and dashed lines: Strengths of spurious connections and curve fits,
346 respectively. Dashed vertical line: Sampling interval of 750 ms, mimicking the TR of the fMRI data. Matrices
347 to the right of each plot show GC connection strengths estimated at sampling interval of 750 ms. Black
348 squares surrounding matrix cells denote significant connections (Methods). For iGC and PC (symmetric
349 connectivity), only the lower triangular matrix is shown, for clarity.

350 **B.** (Top left) Schematic showing a cluster of neurons, each with timescale 50ms, connected with sparse,
351 random, net excitatory connectivity. Matrix: Connectivity among the 100 neurons in a representative cluster.
352 Red: excitatory connections; blue: inhibitory connections. Each such cluster forms one of the nine nodes in
353 the simulated network. (Top right) Connectivity among the nine nodes in the network. (Bottom left)
354 Eigenspectrum (upper panel) of a representative 100 neuron cluster, showing one slow emergent timescale
355 corresponding to the real-part of one eigenvalue close to zero. Histogram (lower panel) showing timescales
356 of all eigenmodes, with the slowest eigenmode at >2000ms. (Bottom right) Eigenspectrum (upper panel) of
357 sub-network DEF exhibits multiple slow emergent timescales. Histogram (lower panel) showing timescales
358 of all eigenmodes, with three slow eigenmodes at ~1000-6000 ms.

359 **C.** Same as in A, but for simulated 9-node networks. (Left) Sub-network ABC, (middle) sub-network DEF,
360 (right) sub-network GHI. Other conventions are as in panel A.

361
362 How might such slow timescales, orders of magnitude slower than spike times and membrane time
363 constants, arise in fMRI data? To answer this question, we availed of established results in random matrix
364 theory. Connectivity in randomly connected E-I networks of neurons can produce slow timescales, without
365 fine-tuning of network parameters [32,33]. We modeled sparse, random, net excitatory connectivity in a
366 small network of (N=100) neurons with connection parameters drawn from previous studies (SI Table S6B;
367 [32,34,35]). The eigenspectrum of the network revealed that each network exhibited one eigenvalue close
368 to zero corresponding to a slow timescale (~1000 ms or greater, Fig. 3B bottom left); the latter constitutes
369 an emergent timescale associated with the dominant eigenmode that is a property of network connectivity
370 (Methods).

371
372 We modeled nine such networks, organized into three, non-interacting, clusters (Fig. 3B top right): a) a
373 cluster with a purely feedforward connection across two networks, b) a cluster with recurrent excitatory (E-
374 E) feedback connections among two networks and c) a cluster with recurrent excitatory-inhibitory (E-I)
375 feedback connections among two networks. In each case, connectivity across networks was mediated by a
376 small proportion (5%) of neurons in each network (parameters in SI Table S6B). This configuration mimics
377 “small-world” connectivity in brain networks [36], with locally-connected brain regions interacting through
378 sparse, long-range connections [37]. The eigenspectra revealed that dynamics in all clusters operated at
379 timescales of around 6000 ms, comparable to or slower than the individual network timescales (Fig. 3B
380 bottom right). To simulate fMRI data we averaged the activity across all 100 neurons in each network and
381 convolved it with a canonical HRF. As before, these nine timeseries were then sampled at various sampling
382 intervals, including a 750 ms interval mimicking the scan TR, and analyzed with GC to detect significant
383 connections.

384
385 iGC and dGC identified complementary aspects of connectivity with these simulated data (Fig. 3C). iGC
386 robustly identified feedforward and excitatory (E-E) feedback connections. dGC also estimated these

387 connections, albeit with the following differences: First, in the feedforward network dGC occasionally
388 identified a spurious connection, albeit much weaker in magnitude, in the direction opposite to the true
389 connection (Fig. 3C, left column, red dashed line). Second, when the E-E feedback connections were
390 precisely balanced in strength (symmetric), dGC also failed to identify the connection reliably (SI Fig. S3A).
391 Yet, when these connections were of different strengths dGC reliably identified both connections, and their
392 relative strengths (Fig. 3C, middle column, red). In contrast, when the connections were of different signs
393 (E-I feedback) dGC robustly identified both connections, whereas iGC failed to reliably detect this
394 connection (Fig. 3C, right column, blue). Yet, taken together, iGC and dGC identified all three connection
395 types reliably.

396
397 Next, we compared the efficacy of connectivity estimation with partial correlations (PC). While PC robustly
398 identified both feedforward and feedback E-E connections (Fig. 3C left and middle columns, black), it,
399 surprisingly, failed to estimate feedback E-I connections, particularly when these were balanced in strength
400 (Fig. 3C right column, black). When the E and I connection strengths were not balanced, but were strongly
401 biased in favor of the E or the I connection, PC estimates varied with the sign of the more dominant
402 connection (SI Fig. S3B, right top). These results generalize beyond these particular simulations; in SI
403 Mathematical Note, Section S2 and S3, we identify, analytically, network configurations for which PC
404 estimates systematically deviate from ground-truth connectivity (see also SI Results, section on “Analytical
405 relationship between PC and iGC”).

406
407 Taken together, these results indicate that instantaneous and lag-based connectivity measures can reveal
408 complementary aspects of brain connectivity. In addition, the results challenge the notion that correlation-
409 based measures, like PC, should be favored over lag-based measures, like dGC for measuring functional
410 connectivity in the brain [18]. Rather, the strengths and weaknesses of each measure (GC and PC) must
411 be recognized when seeking to apply these to brain imaging data.

412 413 **Identifying a cognitive core system and predicting behavioral scores with GC connectivity**

414 Our classification analyses and simulations suggested that iGC and dGC reliably recover task-specific brain
415 networks, the latter when slow-timescale processes occur within the network. We asked whether iGC and

dGC connectivity merely reflected reliable statistical patterns of brain activity, or whether it would be relevant for understanding the nature of information flow in the brain, and its relationship to behavior. To answer this question, we first investigated whether each measure would identify brain networks with consistent outflow and inflow hubs across tasks. Next, we asked whether GC connectivity would be relevant for predicting brain-behavior relationship.

First, we sought to identify a common core of “task-generic” connections across cognitive tasks. For this, we applied a feature selection approach – recursive feature elimination (Methods) – a technique that identifies a minimal set of features that provide maximal cross validation accuracy (generalization performance) [38]. Prior to analysis of real data, we validated RFE by applying it to estimate connectivity differences in two simulated networks (Fig. 4A,B). RFE accurately identified connections that differed in simulation ground-truth: specifically, differences in fast timescale connections were reliably identified by iGC, and in slow timescale connections by dGC (Fig. 4B bottom). RFE based on dGC and iGC also accurately identified the relevant connections, but not always their directionality, even when systematic variations in hemodynamic lag occurred across regions; the results are described in SI Results, section on “Effects of regional variations in hemodynamic lag”.

Figure 4. Recursive feature elimination (RFE) identifies task-generic and task-discriminative networks based on GC connectivity.

A. Schematic showing two simulated networks each with fast (50 ms; ABC) and slow (1000 ms; DEF) sub-networks, with distinct connectivity patterns. Network activity was simulated for 375 seconds with a sampling interval of 5 ms, convolved with the hemodynamic response function and sub-sampled at 750 ms to yield 500 simulated time points.

B. (Top) RFE curves, with classification accuracy as a function of remaining features, for classification based on dGC (left) and iGC (right). (Bottom) Maximally discriminative features following RFE based on dGC (left) and iGC (right). Entries denote average beta weights across RFE iterations.

C. RFE curve for two-way classification of each of six tasks (all tasks except Motor) versus rest, based on dGC (top) and iGC (bottom). Color conventions are as in Figure 1D. Data points: RFE accuracies; solid lines: piecewise linear fits. Vertical dashed line: location of the elbow for each RFE curve.

445 **D.** Task-generic connections following task-versus-resting RFE, based on dGC (left) and iGC (right)
446 features, using Shirer et al (2012) 14-network parcellation [26] (SI Table S4); each network is indicated with
447 a different color and a label. Directed dGC connections are shown as tapered links, broad at the source
448 node and narrow at the destination node. Undirected iGC connections are shown as bidirectional links
449 between the respective pair of nodes. Colors of the connections represent the color of the destination node.
450 **E.** Same as in panel D, but for n-way classification across the six tasks. Color conventions are as in panel
451 B.
452 **F.** Same as in panel E, but for task-discriminative connections, which maximally discriminated each task
453 from the five others, following n-way RFE, based on dGC features (left) and iGC features (right). Other
454 conventions are the same as in panel C.

455
456 We applied RFE to classify tasks versus resting state; we chose these six tasks (all tasks except the motor
457 task) as being the most likely to engage common cognitive control mechanisms (Fig. 4C). For these RFE
458 analyses we employed a 14 network functional parcellation [26], as it consistently gave good classification
459 accuracies with both iGC and dGC connectivity (SI Fig. S1B). Following RFE, we applied a binomial test
460 across tasks (Methods) to identify a common core of task-generic connections, separately for iGC and
461 dGC.

462
463 RFE identified distinct task-generic networks with iGC and dGC, which comprised of connections that
464 distinguished a majority of tasks from resting state. The iGC task-generic network revealed a visuospatial
465 network hub, which connected with the anterior salience, dorsal DMN, higher visual and posterior salience
466 networks (Fig. 4D, right). The dGC task-generic network confirmed the hub-like connectivity of the
467 visuospatial network but, in addition, revealed consistent directed information outflow from the visuospatial
468 network to the other networks (Fig. 4D, left). In addition, dGC revealed consistent inflow into the higher-
469 visual network across tasks, including from the visuospatial, right executive control, and auditory networks,
470 consistent with the ability of top-down inputs from these networks to strongly modulate sensory encoding in
471 higher visual cortex [39]. Finally, the higher-visual network projected consistently to the sensorimotor
472 network, suggesting a final common pathway, across these tasks, for influencing behavior. Interestingly,
473 the only network providing inflow into the visuospatial network hub was the anterior salience network, in line

474 with a previous study that indicated a role for the salience network in controlling other task positive
475 networks [7].

476
477 Similarly, we also identified connections that were maximally discriminative across tasks; again iGC and
478 dGC showed distinct sets of these “task-discriminative” connections (Fig. 4E-F; SI Results, section on
479 “Identifying task discriminative networks with GC”).

480
481 To address GC’s relevance for understanding brain-behavior relationships we tested whether the strength
482 of functional connections estimated with iGC and dGC could predict inter-individual variations in behavioral
483 scores as measured by a standard cognitive battery (Methods; SI Table S7). We employed a leave-one-out
484 prediction analysis based on multilinear regression followed by robust correlations of predicted and
485 observed scores (Fig. 5A; $p < 0.05$ with Benjamini-Hochberg correction; Methods).

486
487 **Figure 5. GC connectivity explains inter-individual variations in behavioral scores.**

488 **A.** (Left) Schematic of behavioral score prediction analysis. GC connectivity strengths for each task were
489 used as independent factors to predict behavioral scores using linear regression with a leave-one-out
490 approach. 51 different behavioral scores were predicted (SI Table S7), compared against observed scores
491 (upper right), and their correlation values plotted as a matrix (lower right).

492 **B.** Exemplar score predictions based on dGC (left panels) and iGC (right panels). In order (from left to
493 right): List Sorting score predicted from Working memory task dGC connectivity, Anger Emotion
494 Recognition score from Emotion task dGC connectivity, Endurance score from Motor task iGC connectivity
495 and Reading score from Language task iGC connectivity.

496 **C.** (Top) Prediction statistics for selected scores based on dGC connectivity. Correlation coefficients (r
497 values) between the predicted and observed scores are plotted in the top half of each stem plot, and
498 significance (p values) are plotted in the bottom half. Each score is denoted by a different color, and each
499 sub-panel shows predictions based on GC connectivity for a different task; Stems with open symbols
500 represent non-significant correlation coefficients, whose corresponding p -values are not shown. p values
501 are floored at 10^{-4} for ease of visualization. (Bottom) Same as in top panel, but predictions based on iGC
502 connectivity.

503 **D.** (Top) Inter-subject correlation matrix of composite behavioral scores. Row and column indices: subjects.
504 (Bottom) Cumulative distributions (solid lines) and density function estimates (filled area) of correlation
505 coefficients between observed and predicted composite scores, for the same subject (yellow) or across
506 different subjects (grey). Predictions were based on GC estimates from the relational and working memory
507 tasks. p-value: Kolmogorov-Smirnov test.

508
509 Both iGC and dGC predicted key behavioral scores. Several scores were predicted uniformly well by iGC
510 across tasks (Fig. 5B, right; Fig. 5C, bottom; SI Fig S5B). Scores of fluid intelligence (Penn progressive
511 matrices), grip strength, endurance, and language (reading and picture-vocabulary) (Fig. 5B right; $r: 0.077 -$
512 0.363 ; $p < 0.02$), were all well predicted by iGC, in addition to scores of spatial orientation (Penn line
513 orientation test) and dexterity (SI Fig. S5B, $r: 0.081 - 0.243$; $p < 0.0125$). On the other hand, dGC-based
514 predictions were more selective, in that several behavioral scores were best predicted by dGC based on
515 specific tasks alone (Fig. 5B, left; Fig. 5C, top; SI Fig. S5A). For instance, dGC in the emotion task alone
516 predicted positive affect (Fig. 5C, top; $r = 0.094$, $p = 0.028$) and anger emotion recognition (Fig. 5B, left, cyan;
517 $r = 0.106$, $p = 0.001$) scores, dGC in the gambling task alone predicted self-report scores of perceived social
518 support ($r = 0.101$, $p = 0.002$), fear ($r = 0.139$, $p < 0.001$) and sadness ($r = 0.107$, $p = 0.001$) and dGC in the motor
519 task alone predicted median reaction time in the fluid intelligence test ($r = 0.123$, $p < 0.001$). In addition, dGC
520 in the working memory task predicted a range of scores in the “cognition” category including list sorting (Fig.
521 5B, left, pink; $r = 0.119$, $p = 0.000$), fluid intelligence, picture discrimination speed, spatial orientation and self
522 regulation (discounting of delayed reward; Fig. 5C top, SI Fig. S5A).

523
524 A variety of behavioral scores were also successfully predicted based on PC connectivity (SI Fig. S5C);
525 several of these (~60%) overlapped with those predicted with GC connectivity (SI Results, section on
526 “Predicting behavioral scores with PC connectivity”). Yet, PC connection features that led to successful
527 predictions overlapped strongly with iGC connection features than with dGC connection features, as
528 quantified by the mean correlation between their regression weights in the respective prediction
529 models (PC vs. iGC: $r = 0.38 \pm 0.01$, mean \pm std; PC vs. dGC: $r = 0.03 \pm 0.01$, $p < 0.001$, ranksum test).

531 Finally, we tested whether GC connectivity could predict a combined set of behavioral scores unique to
532 each subject. For this, we created a vector of all independent behavioral scores (composite score;
533 Methods), and confirmed that this composite behavioral score uniquely identified each subject in the
534 database, as evidenced by the highest values along the main diagonal of the inter-subject correlation
535 matrix (Fig. 5D top). Following this, we performed the leave-one-out prediction, as before, except that we
536 used dGC and iGC connectivity features from two of the tasks alone (working memory and relational; also
537 see SI Fig. S5D). We then tested whether each subject's predicted composite score would correlate best
538 with her/his own observed composite scores. Although we did not observe the highest correlation values
539 consistently along the main diagonal, the distribution of correlation coefficients along the diagonal were
540 significantly different (and higher) than the distribution of off-diagonal correlation coefficients (Fig. 5D
541 bottom; $p < 10^{-15}$, Kolmogorov-Smirnov test). Inter-individual variation GC connectivity, therefore, contained
542 sufficient information to accurately identify subject-specific behavioral scores in this cohort of subjects.

543
544 In summary, the ability to successfully predict subject-specific behavioral scores suggests that GC
545 functional connectivity is relevant for understanding brain-behavior relationships. Moreover, connection
546 features that were relevant for behavioral predictions with PC overlapped highly with iGC, but not with dGC,
547 thereby validating our simulation results regarding the complementarity of iGC and dGC connectivity
548 estimates.

550 Discussion

551 Neural processes in the brain range from the timescales of microseconds to milliseconds, for extremely
552 rapid processes (e.g. sound localization), to timescales of several seconds to minutes, for processes that
553 require coordination across diverse brain networks (e.g. when having a conversation), and hours to days,
554 for processes that involve large-scale neuroplastic changes (e.g. when learning a new language).
555 Coordinated activity among brain regions that mediate each of these cognitive processes should manifest
556 in the form of functional connectivity among these regions at the corresponding timescales. Our results
557 indicate that applying Granger-Geweke Causality (GC) with fMRI data permits estimating behaviorally
558 relevant functional connectivity at a timescale corresponding to the sampling rate of fMRI data (seconds).

559
560 The application of GC to neuroscience is a contentious topic, for a variety of reasons [15–17,19,28]. One
561 particular challenge stems from the use of the word “causality”: the notion of causality in GC is different
562 from the notion of interventional causality [40]. Our use of the term Granger causality, here, purely reflects
563 its application as a marker of information flow among brain networks [19,41], and is not meant to indicate
564 causality in a physical, interventional sense.

565
566 With this understanding, our results contain three key insights. First, we show that, either iGC or dGC
567 connectivity suffices to reliably classify task-specific cognitive states with superlative accuracies (Fig. 1B).
568 Instantaneous and directed GC – both measures of conditional linear dependence and feedback [9] – were
569 able to robustly estimate task-specific functional interactions even with slowly sampled fMRI data. Our
570 application of machine learning and classification analysis circumvents the lack of access to ground truth
571 connectivity, and our simulations suggest that GC connectivity is relevant for estimating slow, emergent
572 interactions among brain networks [15–19].

573
574 Second, we show that functional connections identified by iGC and dGC carry complementary information,
575 both in simulated and in real fMRI recordings, and we demonstrate key caveats with employing correlation-
576 based measures of functional connectivity like partial correlations, despite superior classification accuracies
577 with these latter measures. First, PC fails to correctly infer reciprocal excitatory-inhibitory interactions, which

578 can be accurately inferred with lag-based methods like dGC. Second, PC may yield incorrect estimates of
579 functional connectivity that do not match ground truth (SI Fig. S3C). In particular, when the data are well
580 described by an autoregressive model framework our results suggest that instantaneous connectivity
581 measures, like iGC, provide more accurate descriptions of functional connectivity than PC. Third, even with
582 data completely purged of partial correlations, dGC connectivity was sufficient to classify task-specific
583 cognitive states (Fig. 2C). In fact, unweighted directed connectivity alone sufficed to produce accurate
584 classification at accuracies significantly above chance (Fig. 2D). These results indicate that information flow
585 mapped by GC connectivity can be complementary to that of PC, and highlights the need for examining
586 diverse measures, both instantaneous and lag-based, to obtain a complete picture of functional connectivity
587 in the brain.

588
589 Third, differences in inter-individual iGC and dGC connectivity were able to successfully explain inter-
590 individual variation in behavioral scores on various cognitive tasks, and to identify an individual-specific
591 composite marker of behavioral scores, with high accuracies. Because these behavioral scores were
592 acquired in a separate testing session outside the scanning session [27], the results suggest that GC
593 connectivity was both individual-specific, and stable over timescales exceeding the scan session, to permit
594 accurate prediction. Moreover, in our analysis, each subject's behavioral score was predicted based on
595 her/his GC connectivity, whereas the regression beta weights – representing the relationship between GC
596 connectivity and behavior – were computed from the population of all subjects excluding that subject (Fig.
597 5A). Successful predictions, therefore, indicate a consistent mapping between GC connectivity and
598 behavioral scores across the population of subjects. These findings complement recent results showing
599 that dynamic, resting-state functional connectivity, based on correlations, can explain significant variance in
600 human behavioral data [42], and indicate the relevance of lag-based connectivity measures for
601 understanding brain-behavior relationships.

602
603 Does GC's discriminatory power rely on directed functional connectivity in the underlying neural response
604 or systematic distortions of this connectivity induced by subsampling [19] and hemodynamic filtering
605 [20,21]? While our findings cannot completely rule out the latter hypothesis, we address, next, three key

606 caveats raised by previous studies for estimating functional connectivity with fMRI-GC, and argue why our
607 results support the former hypothesis.

608
609 First, several studies have shown that sub-sampling of neural time series, at the scale of fMRI TR, renders
610 functional connections undetectable with GC [11,18–20]. In these studies, GC was estimated with
611 simulated fMRI time series, sampled at an interval (TR) of seconds, and failed to recover underlying neural
612 interactions, which occur at millisecond timescales (e.g. [18]). However, these claims depended strongly on
613 the nature and timescale of the connectivity in the networks employed in these simulations. For instance, a
614 widely cited study [18] employed purely feedforward connectivity matrices with a 50 ms neural timescale in
615 their simulations, and argued that functional connections are not reliably inferred with GC applied to
616 simulated fMRI data. In addition to being neurally implausible, such purely feedforward network
617 configurations yield eigenmodes whose slowest timescales are identical with the timescales of individual
618 nodes [43]. Therefore, such a configuration rendered lag-based measures like GC, irrelevant for estimating
619 neural interactions from slowly sampled fMRI data [18,19]. Furthermore, such connectivity precludes the
620 occurrence of slower, behaviorally relevant timescales of seconds, which readily emerge in the presence of
621 feedback connections, both in simulations [32,33] and in the real brain [44,45]. Our simulations show that
622 slow timescale interactions emerge in networks with sparse, random, net excitatory connectivity, mimicking
623 connectivity in the neocortex [32,33,35]. While earlier studies have employed large-scale, biologically
624 plausible models [46,47] to demonstrate the emergence of slow (<0.1 Hz) emergent functional interactions
625 among brain networks, our results build upon these previous findings and show that such emergent
626 functional interactions at slow timescales can be readily inferred from simulated fMRI data with GC. In fact,
627 GC connectivity continued to robustly classify distinct task states even when data were sampled at 2x or 3x
628 the original sampling interval of the fMRI data. Thus, while it is likely that GC applied to fMRI data is unable
629 to detect connections at timescales faster than TR, our results show that sufficient distinguishing
630 information occurs in slow-timescale connections to enable accurate inter-task classification.

631
632 Sub-sampling alone, may also produce spurious GC causality; the precise conditions under which spurious
633 GC arises for continuous time vector autoregressive processes is an area of active research, and must be
634 addressed in future studies [6,20].

635

636 Second, previous studies have shown that systematic differences in hemodynamic (HRF) lags (e.g. time to
637 onset, or time to peak) among brain regions may produce spurious dGC estimates [19,21]. With simulations
638 we demonstrated that fMRI-GC could identify differences in slow-timescale network connectivity, despite
639 systematic differences and heterogeneity in HRF onset latencies across nodes (SI Fig. S3D-E). In all
640 cases, applying recursive feature elimination with either dGC or iGC features identified the precise subset
641 of connections that distinguished distinct network configurations. In a majority of cases, dGC also correctly
642 identified the directionality of these connections. In our simulations, the only scenario in which dGC
643 features failed to identify the directionality of connections correctly, was when the onset latency in the
644 “destination” nodes were biased to be systematically earlier than those in the “source” nodes. Nevertheless,
645 in the real data it is unlikely that systematic inter-regional HRF differences were responsible for the
646 observed superior classification accuracies. Variations in HRF delays would indeed confound dGC
647 connectivity estimates – if they occurred consistently between brain regions across subjects and tasks (e.g.
648 SI Fig. S3D, red curves). Yet, such a scenario cannot account for the high classification accuracies among
649 tasks and sub-tasks based on dGC connectivity alone. In other words, even if HRF latency differences
650 systematically biased dGC connectivity estimates, these estimates were sufficiently and reliably different
651 across task cognitive states to permit accurate classification among them. Finally, network properties of key
652 regions identified with fMRI-GC were consistent with their known functional properties of these regions. For
653 instance, dGC identified the visuospatial network as an information outflow hub, across all six cognitive
654 tasks (Fig. 4D left). The visuospatial network comprises frontal cortex regions, including the frontal eye
655 field, as well as posterior parietal cortex, which are both widely implicated in visuospatial attention control
656 [48–51]. In addition, the only network that provided task-generic incoming connections to the visuospatial
657 network was the anterior salience network comprising the anterior fronto-insular cortex and the anterior
658 cingulate cortex [52,53], regions implicated in feature-based attention and the suppression of distractors
659 [54]. Information outflow from these key networks identified by dGC is consistent with their role in attention
660 and executive control.

661

662 Third, simulations and theoretical results indicate that scanner noise can degrade, or even obliterate GC
663 connectivity estimates [19]. On the other hand, our classification accuracies suggest that GC estimates

664 were sufficiently robust to scanner noise to permit accurate task and sub-task classification in these data. In
665 fact, we show that averaging dGC connectivity across as few as 5 subjects' data improves classification
666 accuracy to over 95%, for nearly all tasks (Fig. 1D). Such superlative classification accuracies are unlikely
667 to have occurred if scanner noise were to significantly degrade GC estimates.

668
669 In sum, these results strongly indicate that slow functional interactions in the brain can be meaningfully
670 inferred with GC from fMRI data. While the directionality of interactions measured by GC may need to be
671 interpreted with care [11,21], our results suggest that fMRI-GC may be useful for formulating hypothesis
672 about the role of particular brain regions in providing “top-down” control signals, for modulating activity in
673 other brain regions [7,8], as well as for investigating the nature of information flow in cortical microcircuits
674 with slow sampling rate techniques, such as calcium imaging [55]. The causal role of these brain regions in
675 behavior can then be directly tested with interventional approaches such as transcranial magnetic
676 stimulation, optogenetic inactivation or by examining patient populations with lesions in specific brain
677 regions [56]. Such a systematic analysis will pave the way for a mechanistic understanding of how flexible
678 functional interactions among brain regions mediate complex cognitive behaviors.

681 Materials and Methods

682 Ethics statement.

683 The scanning protocol for the Human Connectome Project was approved by the Human Research
684 Protection Office at Washington University at St. Louis' (IRB # 201204036). Only de-identified, publicly
685 released data were used in this study.

687 Data and code availability statement.

688 Data used in the study is available in the public domain at the Human Connectome Project database
689 (<https://db.humanconnectome.org/>). Code used for analyses are available at the following link:

690 <https://figshare.com/s/9d9131a6780fc8197cf1>

691 Data sharing permissions can be found at the HCP website. Code may be shared or re-used upon
692 requesting the corresponding author. These data and code sharing policies comply with the requirements
693 of all funding agencies supporting this research and comply with institutional ethics protocols.

695 fMRI data, parcellation and time-series extraction.

696 We analyzed minimally preprocessed brain scans of 1000 subjects, drawn from the Human Connectome
697 Project (HCP) database (S1200 release) (age range: 22-35 years; 527 females); fMRI acquisition and
698 preprocessing details are described elsewhere [23,25]. SI Table S2 shows the identifiers of the subjects
699 from whom data were analyzed. Data were analyzed from resting state and seven other task conditions (SI
700 Table S1): Emotion processing, Gambling, Language, Motor, Relational processing, Social cognition and
701 Working memory; in most figures, these tasks are referred to with their initial letters. fMRI scans for the
702 relational task were not available for 9/1000 subjects; therefore, we analyzed a total of 7991 scans across
703 all tasks and subjects.

704
705 We employed five different brain parcellations based one anatomical atlas and four functional atlases (SI
706 Table S3). For the tasks versus resting-state classification based on GC connectivity (first section of
707 Results), all 5 parcellations were used. Based on the classification performance in this analysis, we picked
708 the three parcellations with the highest accuracies (90 node and 14 network parcellations of [26] and 96

709 network parcellation of [57]) and these were used for the pairwise classification analysis of each task versus
710 the other as well as the n-way task classification analyses. Analysis with averaging GC features across
711 subjects (Fig. 1D) was performed with a 90 node parcellation [26]. Classification analyses with data purged
712 of instantaneous correlations and unweighted digraph representations (second section of Results) were
713 performed with the Shirer et al [26] 14 network parcellations. Analyses involving identifying task-generic
714 and task-discriminative networks, as well as behavioral score predictions, based on GC features (last
715 section of the Results) were performed with the Shirer et al [26] 14 network parcellation. Voxel time series
716 were extracted using Matlab and SPM 8 [58], and regional and network time series were computed by
717 averaging the time series across all voxels in the respective region or network.

718 We employed parcellations with fewer, more coarse-grained regions, rather than fine-grained parcellations
719 because Granger Causality estimates were more reliable when the number of regions was fewer than the
720 number of timepoints. Both task and resting scans were of sufficient duration (~200-300 volumes) to permit
721 robust GC estimation. Finally, we noticed that in some parcellations, there were overlapping voxels
722 between some of the regions. To avoid mixing of signals, we assigned each overlapping voxel to the region
723 to whose centroid it was closest, based on Euclidean distance.

724 **Estimating functional connectivity with GC.**

725 We modeled instantaneous and lag-based functional connectivity between brain regions using conditional
726 Granger-Geweke Causality [9]. The linear relationship between two multivariate signals \mathbf{x} and \mathbf{y} conditioned
727 on a third multivariate signal \mathbf{z} can be measured as the sum of linear feedback from \mathbf{x} to \mathbf{y} ($F_{\mathbf{x} \rightarrow \mathbf{y} | \mathbf{z}}$), linear
728 feedback from \mathbf{y} to \mathbf{x} ($F_{\mathbf{y} \rightarrow \mathbf{x} | \mathbf{z}}$), and instantaneous linear feedback ($F_{\mathbf{x} \circ \mathbf{y} | \mathbf{z}}$) [9,41]. To quantify these linear
729 relationships, we model the future of each time series in terms of their past values, using multivariate
730 autoregressive (MVAR) modeling (SI Mathematical Note, Section S1, equation 1). MVAR model order was
731 determined with the Akaike Information Criterion (AIC) for each subject, and was typically 1. The MVAR
732 model fit was used to estimate both an instantaneous connectivity matrix using iGC ($F_{\mathbf{x} \circ \mathbf{y} | \mathbf{z}}$) and a lag-
733 based connectivity matrix using dGC ($F_{\mathbf{x} \rightarrow \mathbf{y} | \mathbf{z}}$). Details are provided in SI Mathematical Note, Section S1.

735 Briefly, $F_{x \rightarrow y|z}$ is a measure of the improvement in the ability to predict the future values of y given the
736 past values of x , over and above what can be predicted from the past values of z and y , itself (and vice
737 versa for $F_{y \rightarrow x|z}$). $F_{x \circ y|z}$, on the other hand, measures the instantaneous influence between x and y
738 conditioned on z (see SI Mathematical Note, Section S1). We refer to $F_{x \circ y|z}$, as *instantaneous* GC (iGC),
739 and $F_{x \rightarrow y|z}$ and $F_{y \rightarrow x|z}$ as lag-based GC or *directed* GC (dGC), with the direction of the influence (x to y
740 or vice versa) being indicated by the arrow. The “full” measure of linear dependence and feedback $F_{x,y|z}$ is
741 given by: $F_{x,y|z} = F_{x \rightarrow y|z} + F_{y \rightarrow x|z} + F_{x \circ y|z}$. $F_{x,y|z}$ measures the complete conditional linear dependence
742 between two time series. If, at a given instant, no aspect of one time series can be explained by a linear
743 model containing all the values (past and present) of the other, $F_{x,y|z}$ will evaluate to zero [41].
744

745 **Classification with linear SVM based on GC connectivity.**

746 The connection strengths of the estimated GC functional connectivity matrices were used as feature
747 vectors with a linear classifier based on SVM for high dimensional predictor data. For a parcellation with n
748 ROIs, the number of features for iGC-based classification was $n(n-1)/2$ (upper triangular portion of the
749 symmetric $n \times n$ iGC matrix) and for dGC-based classification it was $n^2 - n$ (all entries of the $n \times n$ dGC matrix,
750 excluding self-connections on the main diagonal). Based on these functional connectivity features, we
751 asked if we could reliably distinguish each task condition from resting state (e.g. language versus resting)
752 or each task condition from the other
753

754 For pairwise classification of resting state scans versus each task we used Matlab's `fitclinear` function,
755 optimizing hyperparameters using a 5-fold approach: by estimating hyperparameters with five sets of 200
756 subjects in turn, and measuring classification accuracies with the remaining 800 subjects. Classification
757 performance was assessed with leave-one-out and 10-fold cross-validation. We also assessed the
758 significance of the classification accuracy with permutation testing (see Methods). In simulations, we
759 observed that the magnitude of GC estimates varied based on the number of timepoints used in the
760 estimation. To prevent this difference in number of timepoints from biasing classification performance, each
761 scan was truncated to a common minimum number of time samples across the respective scans being
762 classified (task, resting) before estimating GC. For each subject, GC connectivity was estimated
763 independently for the two scan runs (left-to-right and right-to-left phase encoding runs), and averaged

764 across the runs. Hyperparameters optimized included the regularization parameter, regularization method
765 (ridge/lasso) and the learner (linear regression model, svm/logistic) using the
766 `OptimizeHyperparameters` option to the `fitclinear` function. Hyperparameter optimization was
767 performed only for task vs. rest classifications, but not for subject feature averaging, task vs. task, or N-way
768 classification analyses.

769
770 For pairwise classification of each task versus the other, default hyperparameters were used in the
771 `fitclinear` function and classification performance was assessed with leave-one-out cross-validation.
772 For n-way classification, we used MATLAB's `fitcecoc` function, which is based on error-correcting output
773 codes, and fits multiclass models for SVMs. Briefly, the function implemented a one-vs-all coding design,
774 for which seven (number of classes in multiclass classification) binary learners were trained. For each
775 binary learner, one class was assigned a positive label and the rest were assigned negative labels. This
776 design exhausts all combinations of positive class assignments. Classification performance in n-way
777 classification was assessed with leave-one-out cross-validation. For each classification analysis mentioned
778 above, task scans were truncated to the common minimum number of time samples across each set of
779 scans, before estimating GC.

780 781 **Classification based on GC connectivity across sub-tasks and with sub-sampled data.**

782 Tasks in the HCP data were run as a block design, alternating between various conditions (sub-tasks). We
783 tested whether GC connectivity would be able to classify among sub-tasks within each task (SI Table S1B).
784 fMRI time series corresponding to each sub task was obtained by concatenating blocks of fMRI task time
785 series pertaining to the respective sub task; the temporal order across blocks was preserved while
786 concatenating the data. We also ensured that data at the conjunction of two successive blocks, which
787 represented non-contiguous time points, were not used for GC estimation. The two sub tasks to be
788 classified were then truncated to have same number of time points. GC estimation and pair-wise
789 classification across sub-tasks was performed with the procedure described in the previous section. The
790 Shirer et al [26] 14-network parcellation was used for these analyses. For the motor task, time series for the
791 left and right finger movement blocks were combined into a "hand" movement sub-task, and left and right
792 toe movement blocks were combined into a "foot" movement sub-task.

793

794 We also tested whether GC on fMRI data sampled at slower rates would suffice to classify among task and
795 resting states. We obtained time series downsampled at 2x the original sampling interval by removing data
796 at even numbered sample points, and retaining data at odd numbered sample points ($k=1, 3, 5\dots$). The
797 even-sample point data were appended the end of odd-sample data series, thereby retaining the overall
798 number of data points in the original time series. Again, we ensured that data at the conjunction of the odd-
799 and even-sampled data series (last odd-sampled point and first even sampled point), which represented
800 non-contiguous data points, were not used for GC estimation. Similarly, we obtained time series
801 downsampled at 3x the original sampling interval by removing every third data point, starting with the
802 second or third data point, and concatenating these timeseries to retain the overall number of data points in
803 the original timeseries. As before, GC estimation and pair-wise classification was performed with the
804 procedure described in the previous section

805

806 **Permutation testing of classifier accuracies.**

807 We performed permutation tests for evaluating the statistical significance of classifier performance, using
808 the method outlined in [59]. The test involved permuting task labels independently for each subject and
809 computing a null distribution of 10-fold cross-validation accuracy. We employed 1000 surrogates and
810 assessed significance of each empirically estimated 10-fold cross-validation accuracy values for dGC and
811 iGC, based on the proportion of samples in the null distribution which were greater than the cross-validation
812 accuracy estimated from the data. We conducted these analyses for the tasks versus resting state
813 classifications, n-way task classification, classification analyses after purging instantaneous correlations
814 and those based on digraph features, separately for the two metrics (dGC and iGC).

815

816 **Testing for data stationarity and goodness of MVAR model fit.**

817 Computing GC based on VAR modeling assumes that the timeseries represent a stationary process. Four
818 different tests were performed to test whether the MVAR model provided a valid and adequate fit to the
819 data (SI Table S5). We performed these tests for parcellated time-series using scripts provided in the
820 Multivariate Granger Causality (MVGCC) toolbox [60]. First, we checked for the stability of the MVAR model
821 fit by computing logarithm of the spectral radius using the *var_speccrad()* function. A negative value was

822 taken to indicate a stable fit. Second, we assessed consistency of the model fit, which quantifies what
823 proportion of the correlation structure in data is accounted for by the VAR model, using the *consistency()*
824 function. We adopted a threshold of 80% (or above) for both task and resting timeseries to consider the
825 data to have passed the test for consistency [60]. Third, we evaluated the whiteness of residuals based on
826 the Durbin-Watson test for absence of serial correlation of VAR residuals, using the *whiteness()* function.
827 Values of the Durbin-Watson statistic less than 1 or greater than 3 signify a strong positive or negative
828 correlation, respectively among the residuals [60]. Subjects for whom the Durbin-Watson statistic lay
829 between 1 and 3 for more than 90% of the regional timeseries, for both task and resting state data, were
830 considered to have passed the test. Fourth, we checked for stationarity based on the augmented Dicky-
831 Fuller unit-root test (ADF), using the *mvgc_adf()* function. As in the previous case, subjects for whom the
832 ADF test statistic was less than its critical value for more than 90% of the regional timeseries, for both task
833 and resting state data, were considered to have passed the test.

834

835 **Control for motion artifacts.**

836 We checked whether systematic differences in motion artifacts could contribute to the superlative
837 classification accuracies observed with GC. For this, we calculated Frame-wise Displacement (FD) [30] as
838 the sum of temporal derivatives of translational and rotational displacement along the three (x,y,z) axes in
839 mm, with the estimated motion parameters provided by HCP. Frames with $FD > 0.5\text{mm}$ were considered
840 “misaligned” and were discarded (“scrubbed”) while estimating GC values. Because dGC is estimated
841 based on lagged correlations, we also discarded one frame before and after every misaligned frame (AR
842 model order was typically 1 for these data). We then repeated the SVM-based two-way classification of
843 resting state from the seven different task states, with GC features estimated on the “motion scrubbed”
844 data; we also repeated n-way classification among the 7 tasks. Comparison of classification (cross-
845 validated) accuracies with and without motion scrubbing, across all 1000 subjects, is shown in SI Fig. S2C.

846

847

848 **Functional connectivity estimation and classification with partial correlations.**

849 We compared the performance of classification based on GC measures with that based on partial
850 correlations (PC). Partial correlations were computed based on the inverse of the covariance matrix as

851 outlined previously [4,5]. Like iGC, the PC connectivity matrix is undirected and symmetric. Therefore, only
852 the upper triangular portion of the matrix, including $(n*(n-1)/2)$ PC weights, was used as features in the
853 classification analyses. Classification and cross-validation analyses followed the procedures described in
854 the Methods section on “*Classification with linear support vector machines based on GC connectivity*”.

855
856 PC connectivity performed consistently better than GC connectivity for classifying task from resting state
857 (Fig. 2A). We propose the following analytical explanation for this observation: PC, an estimator based on
858 instantaneous covariance, is less susceptible to noise than GC, which is based on lagged covariance. This
859 is due to the fact that the estimation of lagged-covariance is susceptible to errors from noise at multiple
860 time-points. For illustration, consider a timeseries generated by a VAR(1) model: $\mathbf{x}(t) = A \mathbf{x}(t - 1) + \mathbf{e}(t)$.
861 The lagged (lag-1) covariance matrix (Σ_1) is estimated from the data as:

$$862 \quad E[\mathbf{x}(t) \mathbf{x}(t - 1)^T] = E[(A\mathbf{x}(t - 1) + \mathbf{e}(t)) \mathbf{x}(t - 1)^T] = A E[\mathbf{x}(t - 1) \mathbf{x}(t - 1)^T] + E[\mathbf{e}(t) \mathbf{x}(t - 1)^T]$$

863 Thus, when estimating the lagged covariance, the variance of the interaction term $E[\mathbf{e}(t) \mathbf{x}(t - 1)^T]$ (second
864 term in the right hand side) contributes to the variance of Σ_1 in addition to the variance in computing the
865 instantaneous covariance $E[\mathbf{x}(t - 1) \mathbf{x}(t - 1)^T]$ (first term on the right hand side).

866 867 **Classification based on GC connectivity in zero-lag correlation purged data.**

868 To test the complementarity of information conveyed by GC functional connectivity versus functional
869 connectivity based on instantaneous correlations we decorrelated the regional time series data to purge
870 them of instantaneous correlations. We adopted two approaches for this purpose: i) zero-phase component
871 analysis (ZCA) and ii) generalized eigenvalue decomposition (GEV).

872
873 *i) Zero-phase component analysis (ZCA).* Consider demeaned $t \times r$ data matrix \mathbf{X} of regional timeseries with t
874 timepoints and r regions, with covariance matrix \mathbf{C} . Decorrelating the data, to remove correlations among
875 the columns of \mathbf{X} , is achieved with a whitening transformation. A common whitening transformation is
876 based on principal components analysis (PCA): $\mathbf{Y} = \mathbf{W}_{\text{PCA}} \mathbf{X}$, with $\mathbf{W}_{\text{PCA}} = \mathbf{D}^{-1/2} \mathbf{E}^T$ where \mathbf{D} is a diagonal
877 matrix, with the eigenvalues of \mathbf{C} on its diagonals, and the columns of \mathbf{E} contain the eigenvectors of \mathbf{C} .

878 While the PCA transformation effectively decorrelates regional timeseries, there is no way to ensure one-to-
879 one correspondence of the whitened dimensions across subjects, rendering subsequent classification
880 analysis challenging. Consequently, here we chose a different whitening transformation based on zero-
881 phase component analysis (ZCA), also known as the Mahalanobis transformation. Based on this
882 transformation, whitening is achieved as: $\mathbf{Y} = \mathbf{W}_{\text{ZCA}}\mathbf{X}$, with $\mathbf{W}_{\text{ZCA}} = \mathbf{E}\mathbf{D}^{-1/2}\mathbf{E}^T = \mathbf{C}^{-1/2}$. A particular advantage
883 of the ZCA transformation is that it yields whitened data that is as close as possible to the original data, in a
884 least-squares sense [61]. Therefore, each subject's data is projected on to a set of dimensions are most
885 closely aligned with the underlying regional timeseries dimensions. Because the regions exhibit spatial
886 correspondence across subjects (due to fMRI spatial normalization), the ZCA dimensions possess a
887 natural, one-to-one correspondence across subjects, permitting subsequent classification. Before
888 classification analysis ZCA dimensions were identified for each subject, separately for task and resting
889 datasets. Regional time series for task and resting data were independently decorrelated by projecting onto
890 their respective ZCA dimensions. GC (and PC) functional connectivity was estimated based on the these
891 decorrelated timeseries, followed by classification analysis, as described previously (Methods section on
892 "*Classification with linear support vector machines based on GC connectivity*"). As proof that the ZCA
893 transformation was working effectively, classification accuracy based on PC (an instantaneous correlation
894 measure) computed from ZCA components was at chance across all tasks (Fig. 2C top).

895
896 *ii) Generalized Eigenvalue Decomposition (GEV).* Although ZCA effectively purged correlations from the
897 data, for the subsequent classification analyses task and resting state data were projected onto different,
898 respective ZCA dimensions. Thus, the above-chance task versus resting state classification accuracy with
899 GC features derived from ZCA components (Fig. 2C top) could perhaps be explained by, for example,
900 systematic differences with how reliably ZCA dimensions were estimated across task and resting-state
901 scans. We therefore sought an approach that could project both task and resting data into the same
902 dimension while simultaneously decorrelating both. Such joint decorrelation may be achieved by projecting
903 the data on to the generalized eigenvectors of the covariance matrices of the two datasets [62]. Let \mathbf{C}_T and
904 \mathbf{C}_R denote the covariance matrices of the task and resting datasets respectively. The generalized
905 eigenvectors of these two symmetric matrices are given by the columns of $\mathbf{G} = \mathbf{E}_T \mathbf{D}_T^{-1/2} \mathbf{E}_R$, where, as

906 before \mathbf{D}_T is a diagonal matrix, with the eigenvalues of \mathbf{C}_T on its diagonals, and the columns of \mathbf{E}_R and \mathbf{E}_T
907 contain the eigenvectors of \mathbf{C}_R and \mathbf{C}_T respectively. It can be readily verified that $\mathbf{G}^T \mathbf{C}_T \mathbf{G}$ and $\mathbf{G}^T \mathbf{C}_R \mathbf{G}$ are
908 both diagonal matrices. Therefore, \mathbf{G} is a matrix that jointly diagonalizes both \mathbf{C}_T and \mathbf{C}_R and projecting
909 either task or resting state data into the columns of \mathbf{G} decorrelates the respective timeseries. So, for these
910 analyses, the regional time series for the task and resting state conditions for each subject were jointly
911 decorrelated by projecting them onto a single space spanned the generalized eigenvectors. This was
912 followed by classification analysis with GC features obtained from the decorrelated time series. As before,
913 we confirmed the effectiveness of the decorrelation by computing classification accuracy based on PC from
914 GEV components, which was at chance across all tasks (Fig. 2C bottom).

916 **Classification based on unweighted digraph representations of GC connectivity.**

917 An unweighted directed graph (digraph) network representation shows the dominant direction (but not
918 magnitude) of functional connectivity among brain regions. Obtaining significant directed connections with
919 dGC is challenging due to number of multiple comparisons required for testing n^2-n connections. To identify
920 significant directed connections, overcoming the multiple comparisons problem, we first subtracted the dGC
921 connectivity matrix from its transpose and then applied the following two-stage procedure. In the first stage,
922 the 1000 subjects were divided into five folds. For each two-way task versus resting state classification,
923 recursive feature elimination (RFE, described in a later section titled “GC feature selection based on
924 *Recursive Feature Elimination*”) was performed based on dGC features of subjects from one fold (i.e. with
925 200 subjects). A minimal set of connection features identified by RFE, and their corresponding symmetric
926 counterparts were then employed in the subsequent analyses; we term these connections K ; the cardinality
927 of K (the number of significant connections) was typically in the range of 2 - 86 (2.5th - 97.5th percentile). In
928 the second stage, we identified statistically significant connections among these K features alone. For each
929 of the subjects in the four remaining folds (i.e. 800 subjects), a null distribution for the dGC values of the
930 features in K was obtained by estimating dGC following phase-scrambling the time series [8]. Next, we
931 identified significant connections based on dGC values that occurred at the tail of the null distribution; the
932 threshold for significant connections was determined based on a p-value of 0.05 with a Bonferroni
933 correction for multiple comparisons. Classification performance based on digraph features was assessed
934 with leave-one-out cross-validation.

935

936 **GC connectivity in simulated fMRI time series.**

937 To test the ability of GC measures to reliably recover functional interactions at different timescales, we
938 simulated fMRI time series for model networks. Simulated fMRI time series were generated using a two-
939 stage model. The first stage involved modeling latent neural dynamics with a stochastic, linear vector
940 differential equation given by:

$$941 \quad \tau dr/dt = -r + Wr + \epsilon$$

942 where \mathbf{r} is the multivariate neural state variable representing the state of each neuron (or node) in the
943 network (an $N \times 1$ vector, with N being the number of neurons), $d\mathbf{r}/dt$ is its temporal derivative, W is the
944 neural (“ground truth”) connectivity matrix (dimension $N \times N$), τ is the time constant of each neuron (or node)
945 and ϵ is i.i.d Gaussian noise ($N(0, \Sigma)$), with $\Sigma = I_N$ ($N \times N$ identity matrix). Although this model does not
946 explicitly incorporate signal propagation delays, such vector Ornstein-Uhlenbeck models rank, arguably,
947 among the most common models employed for simulating neural and fMRI time series, in many previous
948 studies [6,18,19]. The multivariate time series $\mathbf{r}(t)$, sampled at discrete time points $\mathbf{r}(k\Delta)$ with a sampling rate
949 of Δ , were generated based on the discrete time (1-lag) connectivity matrix $A(\Delta)$ and a residual noise
950 intensity $\Sigma(\Delta)$. Here:

$$951 \quad A(\Delta) = e^{\Delta A}; \quad \Sigma(\Delta) = (1/\Delta) (\Gamma(0) - e^{\Delta A} \Gamma(0) e^{\Delta A'})$$

952 where $A = (1/\tau) (W - I_N)$, e^A denotes the matrix exponential, A' is the transpose of A , and $\Gamma(0)$ is the zero
953 lag autocovariance which satisfies the continuous time Lyapunov equation $A\Gamma(0) + \Gamma(0)A' + \Sigma = 0$ [19]. In the
954 second stage, the latent neural dynamics were convolved with the hemodynamic response function (HRF)
955 to obtain the simulated fMRI time series: $\mathbf{y} = H \otimes \mathbf{x}$, where H is the canonical hemodynamic response
956 function (hrf; simulated with *spm_hrf* in SPM8), \otimes is the convolution operation and \mathbf{y} is the simulated fMRI
957 time series. Finally, following convolution with the hrf, the data were downsampled to 750 ms, to mimic the
958 repeat time (TR) of the HCP fMRI scans used in this study. The same model was used for the different
959 simulations used in the manuscript (third section of the Results). The parameters for the 2-node
960 simulations, and for the 9-node (100 neurons per node) simulations are described in SI Table S6.

961

962 For the 2-node simulations, iGC and dGC values were estimated by simulating the network for 200
963 timepoints, averaged across 25 repetitions. The 9-node simulations were performed with a 900 neuron
964 network, with 100 neurons per node. Each node had sparse, random excitatory/inhibitory connectivity
965 among its neurons (parameters in Table S6), whereas only 5% of neurons in each node were involved in
966 inter-node connections, to mimic sparse, long-range connectivity in the neocortex [63]. The network was
967 simulated for 200 timepoints, and timeseries from all (100) neurons in each node were averaged to
968 generate 9 node timeseries. iGC and dGC values were estimated from the node timeseries and averaged
969 across 10 independent repetitions. Significance was assessed with a bootstrap approach that involved
970 generating 1000 surrogates by phase scrambling the node timeseries to yield a null distribution of GC
971 values [8], followed by a Benjamini-Hochberg correction for multiple comparisons.

972
973 Simulations comparing PC and iGC connectivity were performed as follows: We simulated a 7-node
974 network with a 1-lag VAR model of the form: $\mathbf{X}_k = A \mathbf{X}_{k-1} + \boldsymbol{\varepsilon}_k$, where \mathbf{X}_k is the state of the discrete time
975 process at discrete timestep 'k', A is the connectivity matrix, and $\boldsymbol{\varepsilon}$ is Gaussian noise with covariance matrix
976 Σ_d . A was chosen to be a random matrix with spectral radius less than 1 to ensure stability. Σ was chosen
977 such that the covariance between every pair of residuals was zero (independent residuals) except for the
978 first two residuals. The correlation between these residuals, ε^1 and ε^2 , was parametrically varied between -
979 1.0 and 1.0 to systematically vary the strength of iGC connectivity. Note that, under this model, iGC
980 between X^1 and X^2 vanishes only if and only if ε^1 and ε^2 are uncorrelated [9].

981 **GC feature selection based on Recursive Feature Elimination (RFE).**

982 We performed features selection for analyses reported in Fig. 2D, 4B,C and S4B based on Recursive
983 Feature Elimination (RFE). RFE identifies a minimal set of features, which provide maximal cross-validation
984 accuracy [38]. Here, we implemented a two-level algorithm, described previously [16,64]. First, the data
985 were divided into N_1 (here, 10) folds. Of these, N_1-1 folds were used as "training" data, and one fold was
986 reserved as "test" data for quantifying the generalization performance of the classifier. Training data were
987 pooled and further divided into N_2 (here, 5) folds. The SVM classifier was then trained on N_2-1 folds
988 (leaving out one fold) and discriminative weights were obtained. The above procedure was repeated N_2
989 times by leaving out each fold, in turn. Average weights were then computed by averaging the absolute
990

991 values of the discriminative weights across the N_2 runs. Next, 10% of the features (connections)
992 contributing the lowest average weights were discarded, and the classifier was trained again with only the
993 retained set of features. This procedure of feature selection and training was repeated until no more
994 features remained. At this stage, the generalization performance for every set of retained features (each
995 “RFE level”) was assessed using the left out “test” data. The entire procedure was repeated N_1 times by
996 leaving out each fold of the original data, in turn, as test data. Final generalization performances and
997 discriminative weights of each RFE level were obtained as the average over N_1 folds. We selected the set
998 of connections at the RFE level at which the generalization performance reached an “elbow”: a minimal set
999 of connections at which generalization performance dipped dramatically below its maximal level. To identify
000 this elbow (e), we used a custom elbow fitting procedure, requiring a piecewise linear fit to the RFE curve,
001 based on two lines, one for “ $x > e$ ” and another for “ $x \leq e$ ”, with the first line required to have a higher slope
002 than the second. The first point in each RFE curve was excluded from the higher slope line fit (Fig. 4C, 4E,
003 SI Fig. S4B). RFE was typically repeated 5 times before determining peak accuracy and corresponding
004 features.

006 **Simulating hemodynamic lag variations across nodes.**

007 We simulated systematic differences in hemodynamic lags across nodes by varying the onset parameter of
008 the *spm_hrf* function (SPM8; [58]). For network configurations A and B described in Figure 4A, we
009 simulated 4 scenarios: a) same mean HRF onset ($\mu_L = 3s$) across nodes; b) source node HRF onset lagging
010 the destination node by 1s ($\mu_{L-src} > \mu_{L-dst}$); c) source node HRF onset leading destination node by 1s ($\mu_{L-src} >$
011 μ_{L-dst}); and d) mixed latencies of lead and lag across source and destination nodes (see next). GC was
012 estimated for 100 simulated participants, by sampling onset latencies for each of the 6 nodes (A-F) from
013 normal distributions (truncated to have only positive latency values), over a range of different standard
014 deviations ($\sigma_L = 0-1s$, in steps of 0.2s). Onset latencies were sampled independently across participants, but
015 were sampled such that the relative latency between each pair of source and destination nodes, across
016 corresponding network configurations, remained the same for each participant. For example, if the onset
017 latency difference between nodes A and B was 0.7s ($\mu_{L-B} - \mu_{L-A} = 0.7s$) for a particular subject, the same
018 difference in onset latency was also maintained between nodes B and C ($\mu_{L-C} - \mu_{L-B} = 0.7s$). For simulations
019 with mixed latencies (case d), 50% of simulated participants had onset latencies drawn from distributions

020 with the source node lagging the destination node (case b) and the remaining 50% with the source node
021 leading the destination node (case c). GC values were averaged over 5 runs for each simulated
022 participant. Finally, we performed RFE to identify key connections that distinguished the two network
023 configurations (same procedure as in Fig. 4B). Connections weights of the most discriminative connections
024 following RFE are shown in SI Fig. S3E (for $\sigma_L=0.4s$). Difference of dGC connections strengths as well as
025 iGC connection strengths, for various values of σ_L , are shown in SI Fig. S3D.

028 **Identifying “task-generic” and “task-discriminative” GC connections.**

029 To identify a minimal set of connections that occurred consistently across tasks (“task-generic”
030 connections), we adopted the following approach. We performed RFE analysis for task versus resting state
031 classification for each of the six tasks (all tasks except motor); we expected each of these tasks to recruit
032 common cognitive control mechanisms. We then performed a binomial test to identify connections that
033 were consistently activated across tasks. Briefly, the presence or absence of a connection in the set of RFE
034 features for a given task versus resting state classification was considered as a Bernoulli trial, with
035 probability of success (its presence) p being the mean number of RFE features identified across all six
036 classifications. The number of trials n was the number of tasks versus resting state classifications (here
037 $n=6$). The probability of a randomly picked connection being present in more than k such RFE sets is given
038 by the cumulative distribution function for the binomial distribution $F(k; n, p)$. Significant connections were
039 identified as those that occurred in k or more tasks, with threshold at the $p=0.05$ level.

040
041 To identify a minimal set of connections that maximally differed across tasks (“task-discriminative”
042 connections), we used RFE with an n -way classifier, to classify among all six tasks (again, except the motor
043 task). The n -way classifier is based on training n (here, 6) one-vs-all binary learners. At the second level of
044 the RFE procedure described above, average weights were computed for each of these n binary learners
045 by averaging the absolute values of the discriminative weights across the N_2 runs. Next, a set of features
046 obtained by taking union of 1% of the features (connections) contributing the lowest average weights in
047 each learner was discarded, and the classifier was trained again with only the retained set of features.

049 While quantifying the overlap between task-generic and task-discriminating connections identified
050 separately for dGC, iGC and PC, we converted the dGC matrix to a lower triangular matrix by reflecting all
051 connections about the main diagonal. The degree of overlap between PC and GC connections was
052 quantified as the number of overlapping connections as proportion of the total number of connections
053 identified by PC. We then computed a null distribution of the degree of overlap by randomly permuting the
054 connection identities within each matrix, while preserving the overall number of connections in each matrix,
055 and generating 1000 surrogate samples. The significance of the overlap of task-generic or task-
056 discriminating connections between each pair of metrics (PC-dGC or PC-iGC) was quantified as the
057 fraction of overlapping connections in the data that exceeded this null distribution.

059 **Predicting behavioral scores based on GC connectivity**

060 We asked whether inter-individual differences in GC connectivity would be relevant for predicting inter-
061 individual differences in behavioral scores. HCP provides a well-validated battery of behavioral scores
062 assessed with a wide range of cognitive tasks. The task battery is based on the NIH Toolbox for
063 Assessment of Neurological and Behavioral function [65], developed to create a uniform set of measures
064 for rapid data collection in large cohorts. The toolbox includes assessments of cognitive, emotional, motor
065 and sensory processing scores in healthy individuals. We pre-selected, based on domain knowledge, a
066 specific subset of 51 scores for these analyses, using age-adjusted scores, wherever available (listed in SI
067 Table S7). Next, we sought to predict subjects' behavioral scores based on GC connectivity with an
068 established leave-one-out approach [66]. Briefly, we used linear regression to predict behavioral scores
069 using, as features, GC estimates of functional connectivity, separately for iGC (91 features or connections)
070 and dGC (182 features). The leave-one-out analysis was performed such that the support vector regressor
071 was fit on all but one subject and the learned beta weights were used to obtain predictions of the left-out
072 subject's behavioral score, using that subject's own GC connectivity weights. Predicted scores were
073 correlated with the actual scores using robust correlations ("percentage-bend" correlations; [67]).

074
075 Next, we asked if GC connectivity could identify an individual based on a composite marker of her/his
076 behavioral scores. Because 40 subjects did not have a full complement of behavioral scores, data from the
077 remaining 960 subjects was included in this analysis. The 51 behavioral scores were, each, z-scored

078 across subjects and formatted into a “composite behavioral score” vector. This vector served as an
079 individual specific composite marker of behavioral scores, as revealed the weak off-diagonal values in the
080 covariance matrix of this vector across subjects (Fig. 5D top). dGC and iGC features of individual tasks, as
081 well as combination of tasks (Relational and Working memory), were used to then predict the composite
082 score marker for individual subjects, using the same leave-one-out procedure as described above. The
083 observed and predicted set of composite scores was correlated across subjects. The distribution of
084 observed versus predicted correlation values for each subject (values on main diagonal; Fig. 5D bottom
085 yellow) were compared against between-subject correlation values (off-diagonal values; Fig. 5D bottom
086 grey) using a Kolmogorov-Smirnov test.

087
088
089 **Acknowledgments.** The authors would like to thank Lionel Barrett and Catie Chang for their comments on
090 a preliminary version of this manuscript, and Govindan Rangarajan and Arshed Nabeel for helpful
091 discussions.

References

1. Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME. From The Cover: The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci*. 2005;
2. Vincent JL, Kahn I, Snyder AZ, Raichle ME, Buckner RL. Evidence for a Frontoparietal Control System Revealed by Intrinsic Functional Connectivity. *J Neurophysiol*. 2008;
3. Buckner RL, Sepulcre J, Talukdar T, Krienen FM, Liu H, Hedden T, et al. Cortical Hubs Revealed by Intrinsic Functional Connectivity: Mapping, Assessment of Stability, and Relation to Alzheimer's Disease. *J Neurosci*. 2009;
4. Marrelec G, Krainik A, Duffau H, Pélégrini-Issac M, Lehericy S, Doyon J, et al. Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage*. 2006;
5. Ryali S, Chen T, Supekar K, Menon V. Estimation of functional connectivity in fMRI data using stability selection-based sparse partial correlation with elastic net penalty. *Neuroimage*. 2012;59(4):3852–61.
6. Barnett L, Seth AK. Detectability of Granger causality for subsampled continuous-time neurophysiological processes. *J Neurosci Methods*. 2017;275:93–121.
7. Sridharan D, Levitin DJ, Menon V. A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci*. 2008;105(34):12569–74.
8. Ryali S, Supekar K, Chen T, Menon V. Multivariate dynamical systems models for estimating causal interactions in fMRI. *Neuroimage*. 2011;54(2):807–23.
9. Geweke JF. Measures of conditional linear dependence and feedback between time series. *J Am Stat Assoc*. 1984;79(388):907–15.
10. Geweke J. Measurement of linear dependence and feedback between multiple time series. *J Am Stat Assoc*. 1982;77(378):304–13.
11. Seth AK, Barrett AB, Barnett L. Granger Causality Analysis in Neuroscience and Neuroimaging. *J Neurosci*. 2015;
12. Dhamala M, Rangarajan G, Ding M. Analyzing information flow in brain networks with nonparametric Granger causality. *Neuroimage*. 2008;41(2):354–62.
13. Ding M, Wang C. Analyzing MEG Data with Granger Causality: Promises and Pitfalls. In: Supek S,

- 122 Aine CJ, editors. Magnetoencephalography: From Signals to Dynamic Cortical Networks. Berlin,
123 Heidelberg: Springer Berlin Heidelberg; 2014. p. 309–18.
- 124 14. Bastos AM, Vezoli J, Bosman CA, Schoffelen J-M, Oostenveld R, Dowdall JR, et al. Visual areas
125 exert feedforward and feedback influences through distinct frequency channels. *Neuron*.
126 2015;85(2):390–401.
- 127 15. Friston K, Moran R, Seth AK. Analysing connectivity with Granger causality and dynamic causal
128 modelling. Vol. 23, *Current Opinion in Neurobiology*. 2013. p. 172–8.
- 129 16. Chang C, Thomason ME, Glover GH. Mapping and correction of vascular hemodynamic latency in
130 the BOLD signal. *Neuroimage*. 2008;43(1):90–102.
- 131 17. Wen X, Rangarajan G, Ding M. Is Granger Causality a Viable Technique for Analyzing fMRI Data?
132 *PLoS One*. 2013;8(7):1–11.
- 133 18. Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, et al. Network
134 modelling methods for fMRI. *Neuroimage*. 2011;54(2):875–91.
- 135 19. Seth AK, Chorley P, Barnett LC. Granger causality analysis of fMRI BOLD signals is invariant to
136 hemodynamic convolution but not downsampling. *Neuroimage*. 2013;65:540–55.
- 137 20. Lin F-H, Hara K, Solo V, Vangel M, Belliveau JW, Stufflebeam SM, et al. Dynamic Granger-Geweke
138 causality modeling with application to interictal spike propagation. *Hum Brain Mapp*. 2009
139 Jun;30(6):1877–86.
- 140 21. Solo V, Poline J-B, Lindquist MA, Simpson SL, Bowman FD, Chung MK, et al. Connectivity in fMRI:
141 Blind Spots and Breakthroughs. *IEEE Trans Med Imaging*. 2018 Jul;37(7):1537—1550.
- 142 22. Solo V. State-Space Analysis of Granger-Geweke Causality Measures with Application to fMRI.
143 *Neural Comput*. 2016 May;28(5):914–49.
- 144 23. Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, et al. The minimal
145 preprocessing pipelines for the Human Connectome Project. *Neuroimage*. 2013;80:105–24.
- 146 24. Aiken LS, West SG, Pitts SC. Multiple Linear Regression. *Handb Psychol*. 2003;
- 147 25. Van Essen DC, Ugurbil K, Auerbach E, Barch D, Behrens TEJ, Bucholz R, et al. The Human
148 Connectome Project: A data acquisition perspective. Vol. 62, *NeuroImage*. 2012. p. 2222–31.
- 149 26. Shirer WR, Ryali S, Rykhlevskaia E, Menon V, Greicius MD. Decoding subject-driven cognitive
150 states with whole-brain connectivity patterns. *Cereb Cortex*. 2012;22(1):158–65.

- 151 27. Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, et al. Function in the
152 human connectome: task-fMRI and individual differences in behavior. *Neuroimage*. 2013;80:169–89.
- 153 28. Stokes PA, Purdon PL. A study of problems encountered in Granger causality analysis from a
154 neuroscience perspective. *Proc Natl Acad Sci*. 2017;
- 155 29. Barnett L, Barrett AB, Seth AK. Misunderstandings regarding the application of Granger causality in
156 neuroscience. *Proc Natl Acad Sci*. 2018;
- 157 30. Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE. Spurious but systematic correlations
158 in functional connectivity MRI networks arise from subject motion. *Neuroimage*. 2012;59(3):2142–54.
- 159 31. Liang X, Wang J, Yan C, Shu N, Xu K, Gong G, et al. Effects of Different Correlation Metrics and
160 Preprocessing Factors on Small-World Brain Functional Networks: A Resting-State Functional MRI
161 Study. *PLoS One*. 2012;7(3):1–16.
- 162 32. Ganguli S, Bisley JW, Roitman JD, Shadlen MN, Goldberg ME, Miller KD. One-dimensional
163 dynamics of attention and decision making in LIP. *Neuron*. 2008;58(1):15–25.
- 164 33. Rajan K, Abbott LF. Eigenvalue spectra of random matrices for neural networks. *Phys Rev Lett*.
165 2006;97(18):188104.
- 166 34. Holmgren C, Harkany T, Svennenfors B, Zilberter Y. Pyramidal cell communication within local
167 networks in layer 2/3 of rat neocortex. *J Physiol*. 2003;
- 168 35. Markram H. Organizing principles for a diversity of GABAergic interneurons and synapses in the
169 neocortex. *Science* (80-). 2000;
- 170 36. Bassett DS, Bullmore E. Small-world brain networks. *Neuroscientist*. 2006.
- 171 37. Sporns O, Chialvo DR, Kaiser M, Hilgetag CC. Organization, development and function of complex
172 brain networks. *Trends in Cognitive Sciences*. 2004.
- 173 38. Guyon I, Elisseeff A. An Introduction to Variable and Feature Selection. *J Mach Learn Res*. 2003;
- 174 39. Gilbert CD, Li W. Top-down influences on visual processing. *Nature Reviews Neuroscience*. 2013.
- 175 40. Pearl J. Causality: Models, reasoning, and inference, second edition. *Causality: Models, Reasoning,*
176 *and Inference, Second Edition*. 2011.
- 177 41. Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using Granger
178 causality and fMRI. *Neuroimage*. 2005;25(1):230–42.
- 179 42. Liégeois R, Li J, Kong R, Orban C, Van De Ville D, Ge T, et al. Resting brain dynamics at different

- 180 timescales capture distinct aspects of human behavior. *Nat Commun.* 2019;10(1):2317.
- 181 43. Sundaresan M, Nabeel A, Sridharan D. Mapping distinct timescales of functional interactions among
182 brain networks. 31st Conf Neural Inf Process Syst. 2017;
- 183 44. Runyan CA, Piasini E, Panzeri S, Harvey CD. Distinct timescales of population coding across cortex.
184 *Nature.* 2017;548(7665):92–6.
- 185 45. Vidaurre D, Smith SM, Woolrich MW. Brain network dynamics are hierarchically organized in time.
186 *Proc Natl Acad Sci.* 2017;114(48):12827–32.
- 187 46. Deco G, Jirsa V, McIntosh AR, Sporns O, Kötter R. Key role of coupling, delay, and noise in resting
188 brain fluctuations. *Proc Natl Acad Sci.* 2009;106(25):10302–7.
- 189 47. Krishnan GP, González OC, Bazhenov M. Origin of slow spontaneous resting-state neuronal
190 fluctuations in brain networks. *Proc Natl Acad Sci.* 2018;
- 191 48. Behrmann M, Geng JJ, Shomstein S. Parietal cortex and attention. *Current Opinion in Neurobiology.*
192 2004.
- 193 49. Corbetta M, Akbudak E, Conturo TE, Snyder AZ, Ollinger JM, Drury HA, et al. A common network of
194 functional areas for attention and eye movements. *Neuron.* 1998;
- 195 50. Thompson KG, Bichot NP. A visual salience map in the primate frontal eye field. *Progress in Brain*
196 *Research.* 2004.
- 197 51. Schall JD. On the role of frontal eye field in guiding attention and saccades. In: *Vision Research.*
198 2004.
- 199 52. Chen AC, Oathes DJ, Chang C, Bradley T, Zhou Z-W, Williams LM, et al. Causal interactions
200 between fronto-parietal central executive and default-mode networks in humans. *Proc Natl Acad Sci.*
201 2013;
- 202 53. Dosenbach NUF, Fair DA, Cohen AL, Schlaggar BL, Petersen SE. A dual-networks architecture of
203 top-down control. *Trends Cogn Sci.* 2008;
- 204 54. Li V, Michael E, Balaguer J, Herce Castañón S, Summerfield C. Gain control explains the effect of
205 distraction in human perceptual, cognitive, and economic decision making. *Proc Natl Acad Sci.* 2018;
- 206 55. Fallani FDV, Corazzol M, Sternberg JR, Wyart C, Chavez M. Hierarchy of Neural Organization in the
207 Embryonic Spinal Cord: Granger-Causality Graph Analysis of In Vivo Calcium Imaging Data. *IEEE*
208 *Trans Neural Syst Rehabil Eng.* 2015;23(3):333–41.

- 209 56. Gaillard R, Naccache L, Pinel P, Clémenceau S, Volle E, Hasboun D, et al. Direct Intracranial, fMRI,
210 and Lesion Evidence for the Causal Role of Left Inferotemporal Cortex in Reading. *Neuron*. 2006;
- 211 57. Thomas Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, et al. The
212 organization of the human cerebral cortex estimated by intrinsic functional connectivity. Vol. 106,
213 *Journal of Neurophysiology*. Bethesda, MD; 2011. p. 1125–65.
- 214 58. Penny W, Friston K, Ashburner J, J. Kiebel S, E. Nichols T. *Statistical Parametric Mapping: The*
215 *Analysis of Functional Brain Images*. 2007.
- 216 59. Ojala M, Garriga GC. Permutation tests for studying classifier performance. *J Mach Learn Res*.
217 2010;11:1833–63.
- 218 60. Barnett L, Seth AK. The MVGC multivariate Granger causality toolbox: A new approach to Granger-
219 causal inference. *J Neurosci Methods*. 2014;223:50–68.
- 220 61. Kessy A, Lewin A, Strimmer K. Optimal Whitening and Decorrelation. *American Statistician*. 2018;
- 221 62. Karampatziakis N, Mineiro P. Discriminative Features via Generalized Eigenvectors. *Proc 31st Int*
222 *Conf Mach Learn*. 2014;
- 223 63. Knösche T, Tittgemeyer M. The Role of Long-Range Connectivity for the Characterization of the
224 Functional–Anatomical Organization of the Cortex. *Front Syst Neurosci*. 2011;5:58.
- 225 64. Gel’fand IM, Yaglom AM. Calculation of the amount of information about a random function
226 contained in another such function. *Am Math Soc Transl*. 1959;12(1):199–246.
- 227 65. Gershon RC, Wagster M V., Hendrie HC, Fox NA, Cook KF, Nowinski CJ. NIH Toolbox for
228 Assessment of Neurological and Behavioral Function. *Neurology*. 2013;
- 229 66. Tavor I, Parker Jones O, Mars RB, Smith SM, Behrens TE, Jbabdi S. Task-free MRI predicts
230 individual differences in brain activity during task performance. *Science*. 2016;
- 231 67. Wilcox RR. The percentage bend correlation coefficient. *Psychometrika*. 1994;
- 232
- 233

234 Supporting Information

235 **SI Text. Results and SI Figure Legends**

236 **SI Figure S1. GC classification accuracies for different parcellations and alternative classifiers**

237 **SI Figure S2. Stationarity tests and 1-stage versus 2-stage GC**

238 **SI Figure S3. Relationship between network connectivity, GC and partial correlations**

239 **SI Figure S4. Task generic and discriminative connections based on partial correlations (PC)**

240 **SI Figure S5. Behavioral score predictions based on GC connectivity strengths**

241 **SI Table S1A. Task descriptions.**

242 **SI Table S1B. Description of sub-tasks**

243 **SI Table S2. Subject identifiers**

244 **SI Table S3. Parcellations used in the analysis**

245 **SI Table S4. Network labels in the Shirer et al (2012) 14-network parcellation**

246 **SI Table S5. Number of subjects passing stationarity tests**

247 **SI Table S6. Parameters of simulated networks**

248 **SI Table S7. Behavioral scores and descriptions**

249 **SI Text. Mathematical Note**

Figure 1

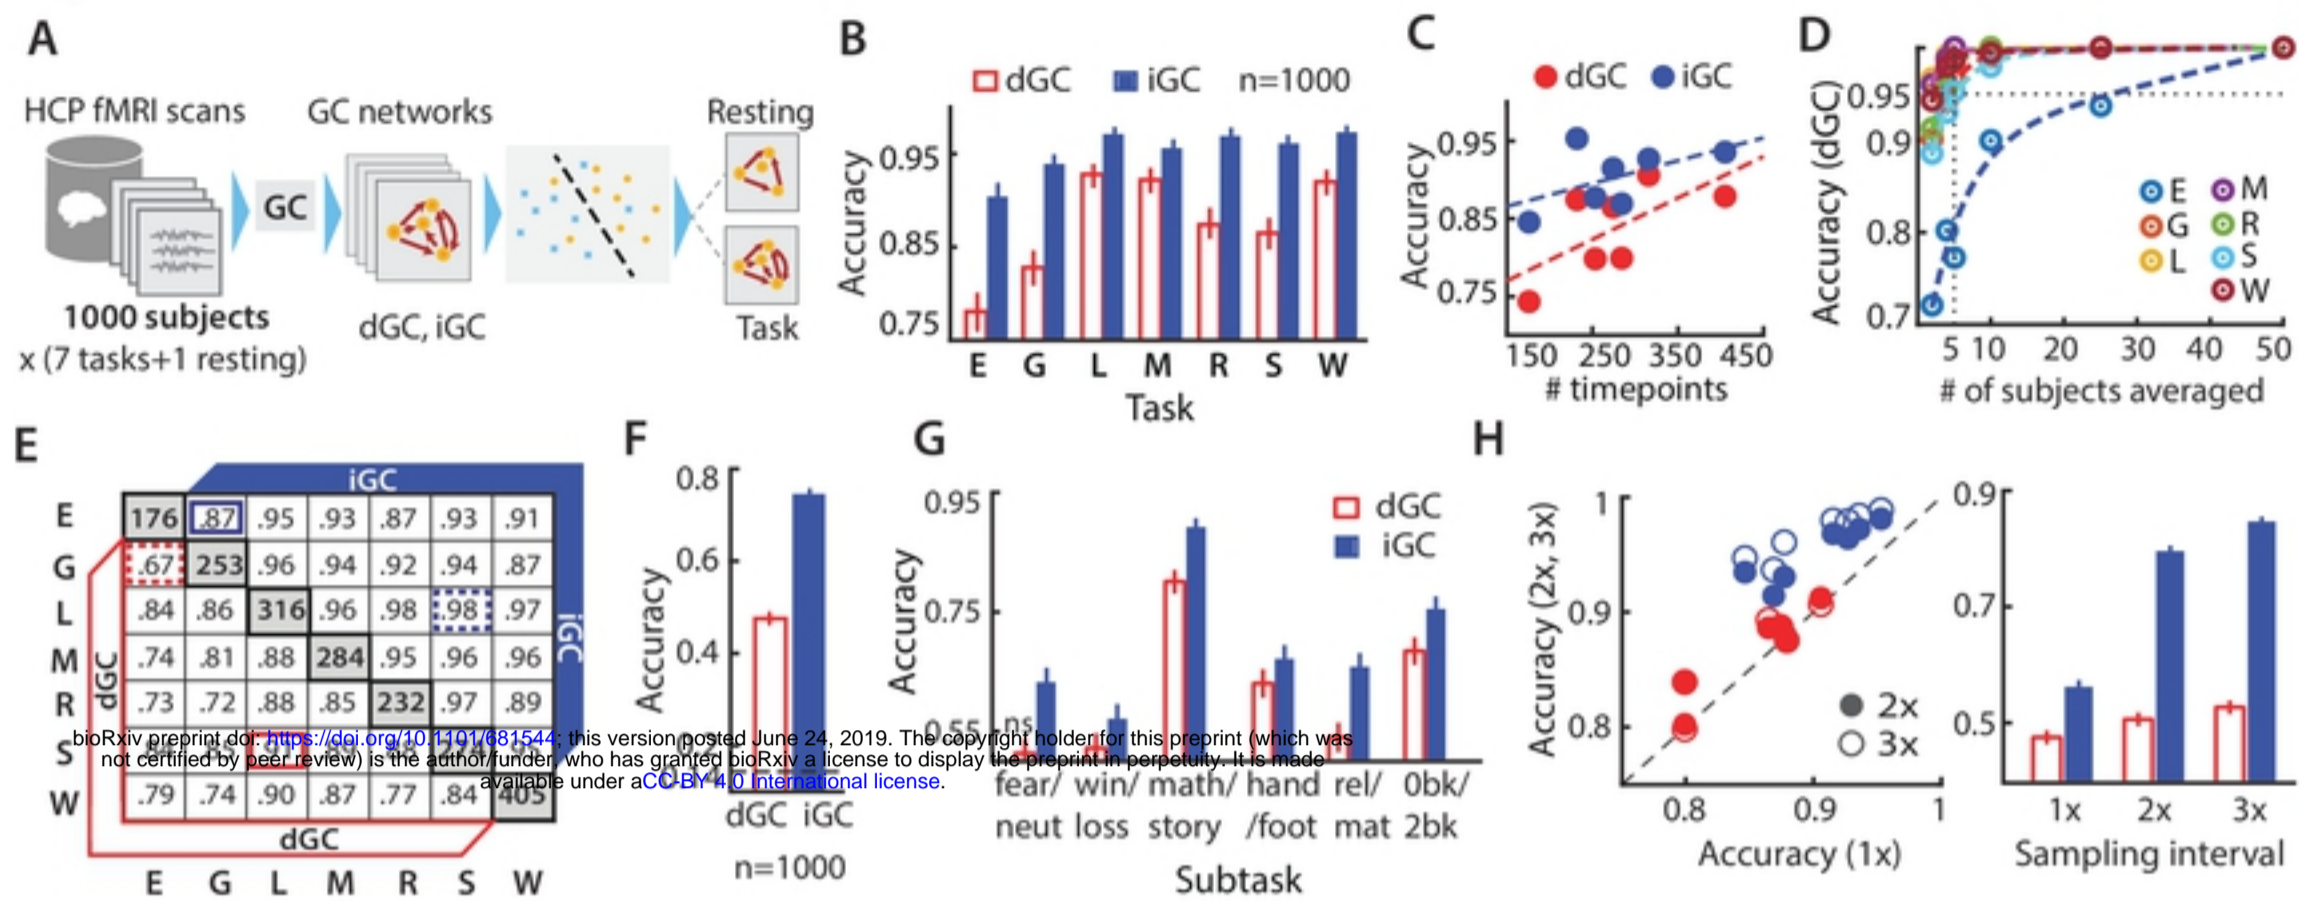
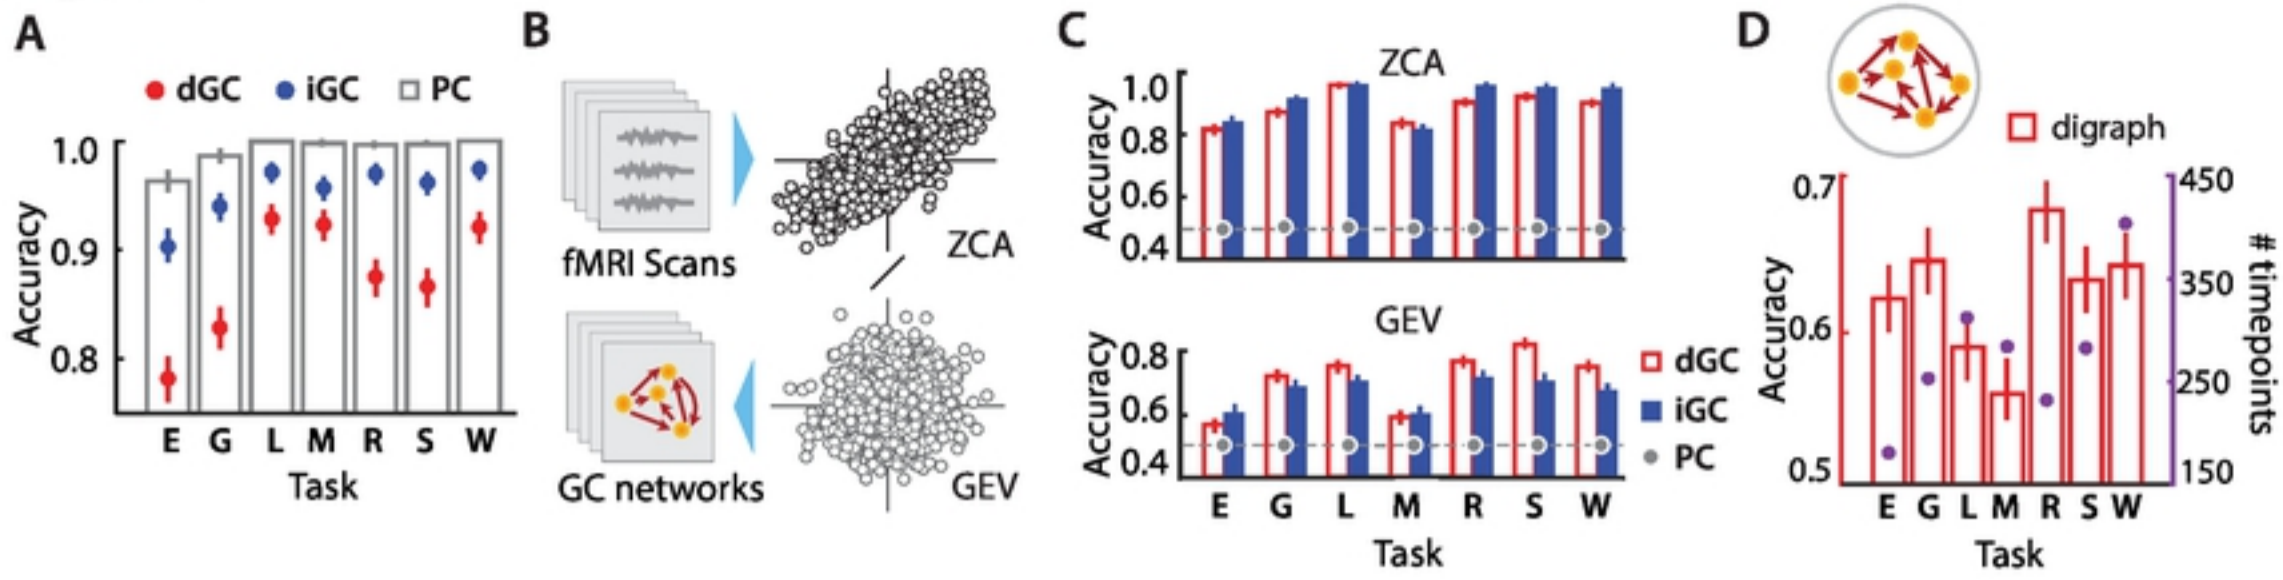


Figure 1

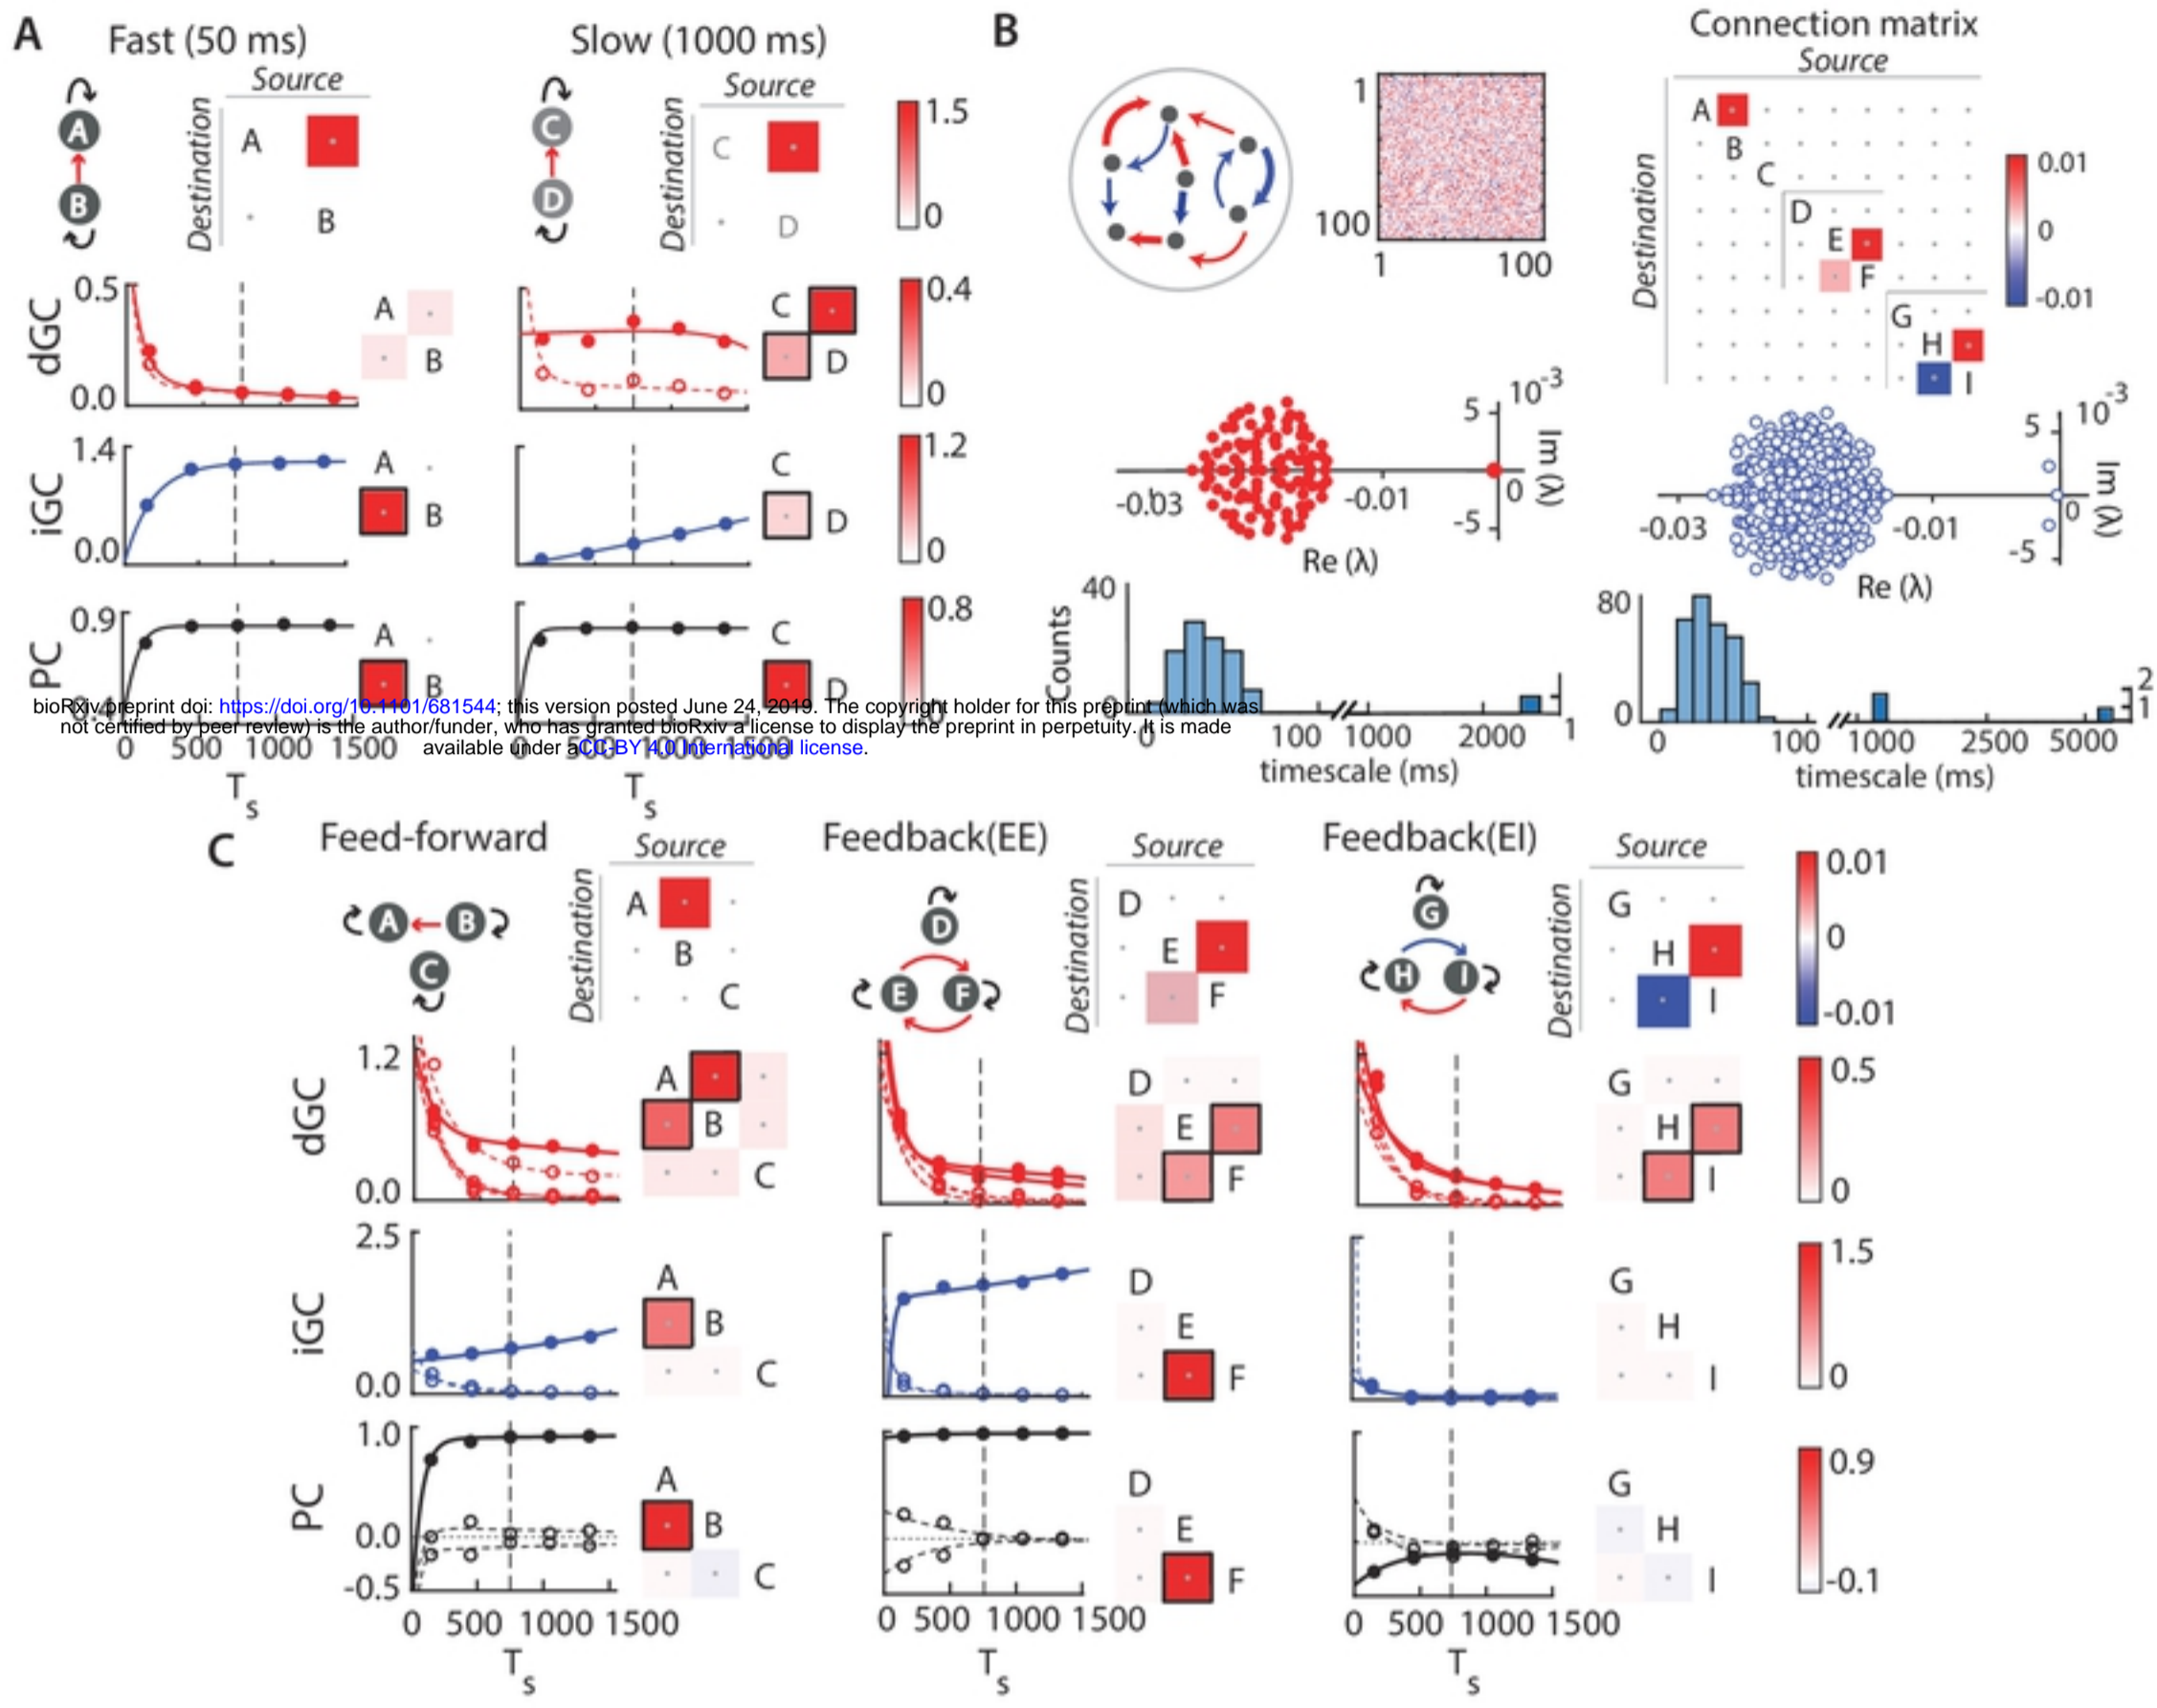
Figure 2



bioRxiv preprint doi: <https://doi.org/10.1101/681544>; this version posted June 24, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figure 2

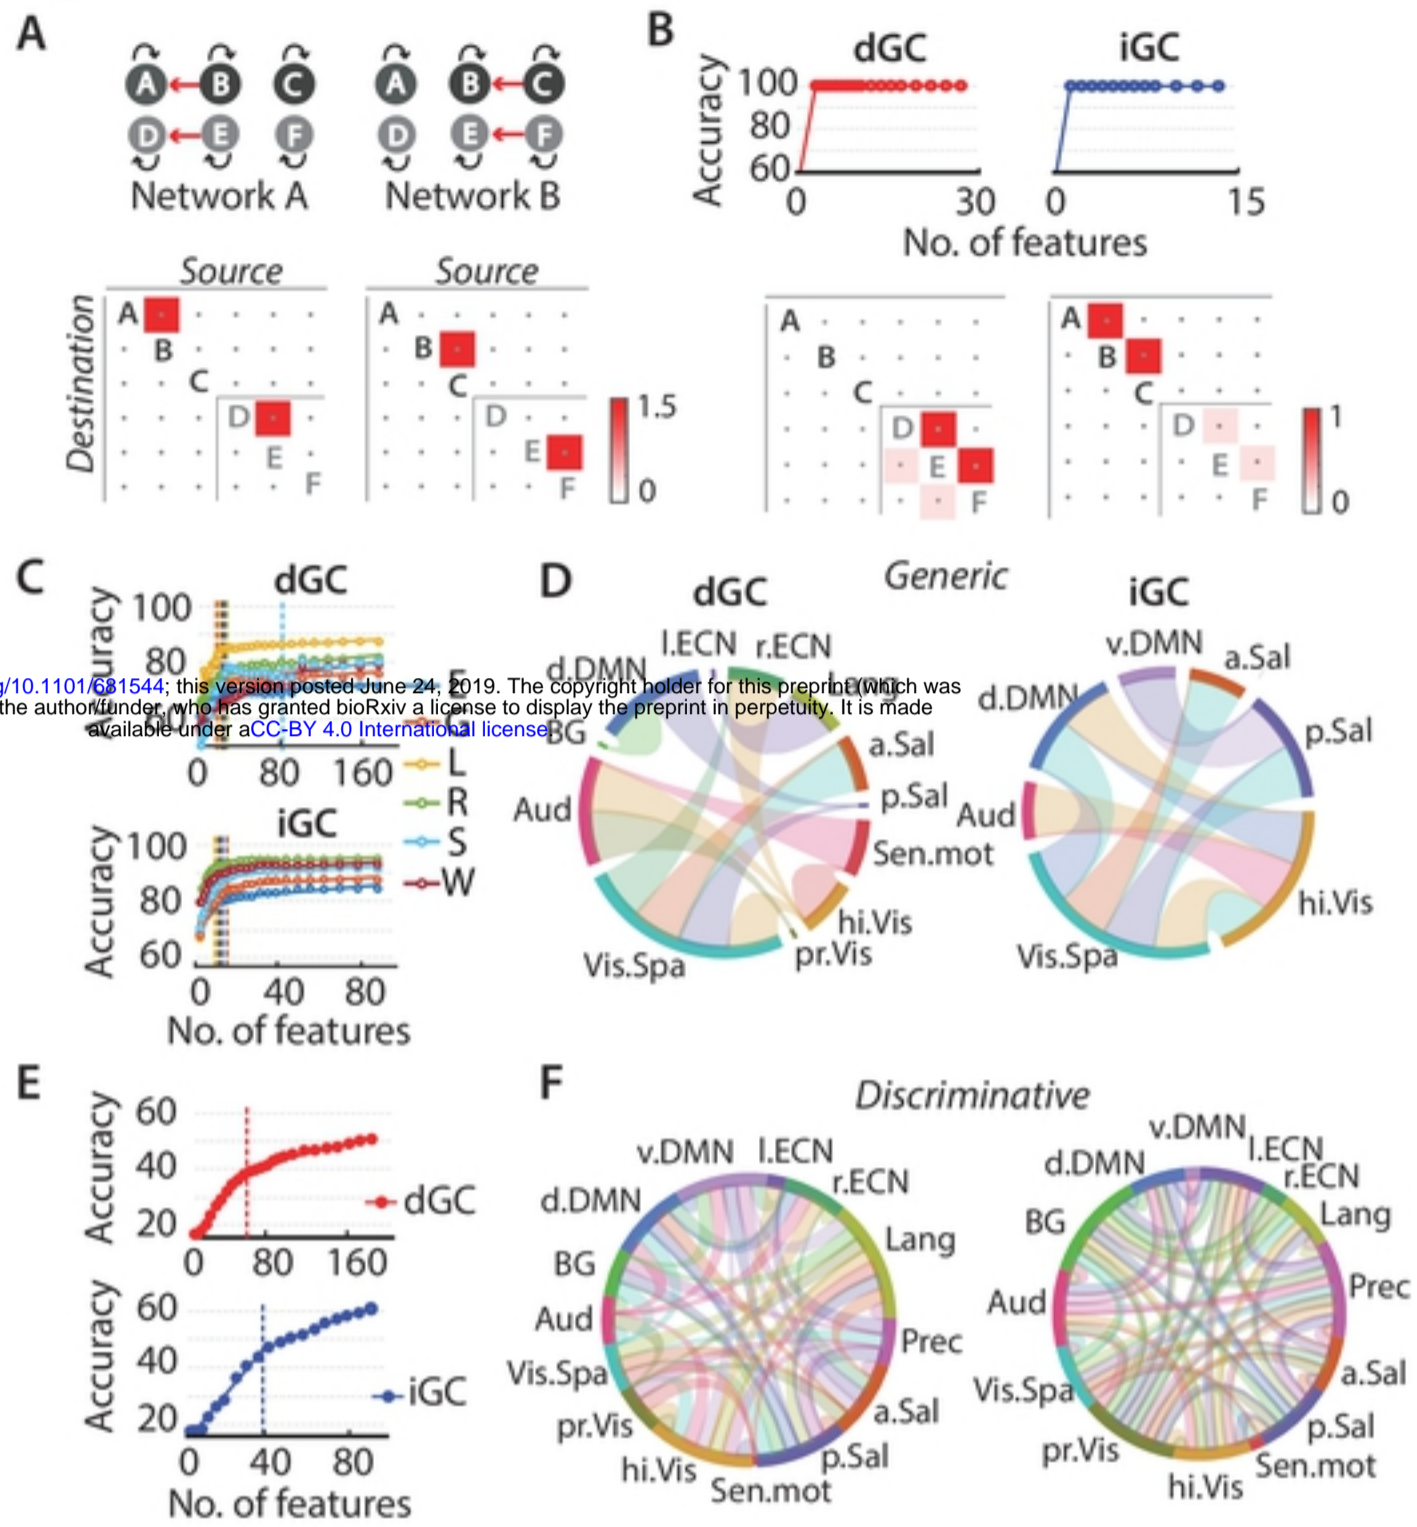
Figure 3



bioRxiv preprint doi: <https://doi.org/10.1101/681544>; this version posted June 24, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figure 3

Figure 4



bioRxiv preprint doi: <https://doi.org/10.1101/631544>; this version posted June 24, 2019. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figure 5

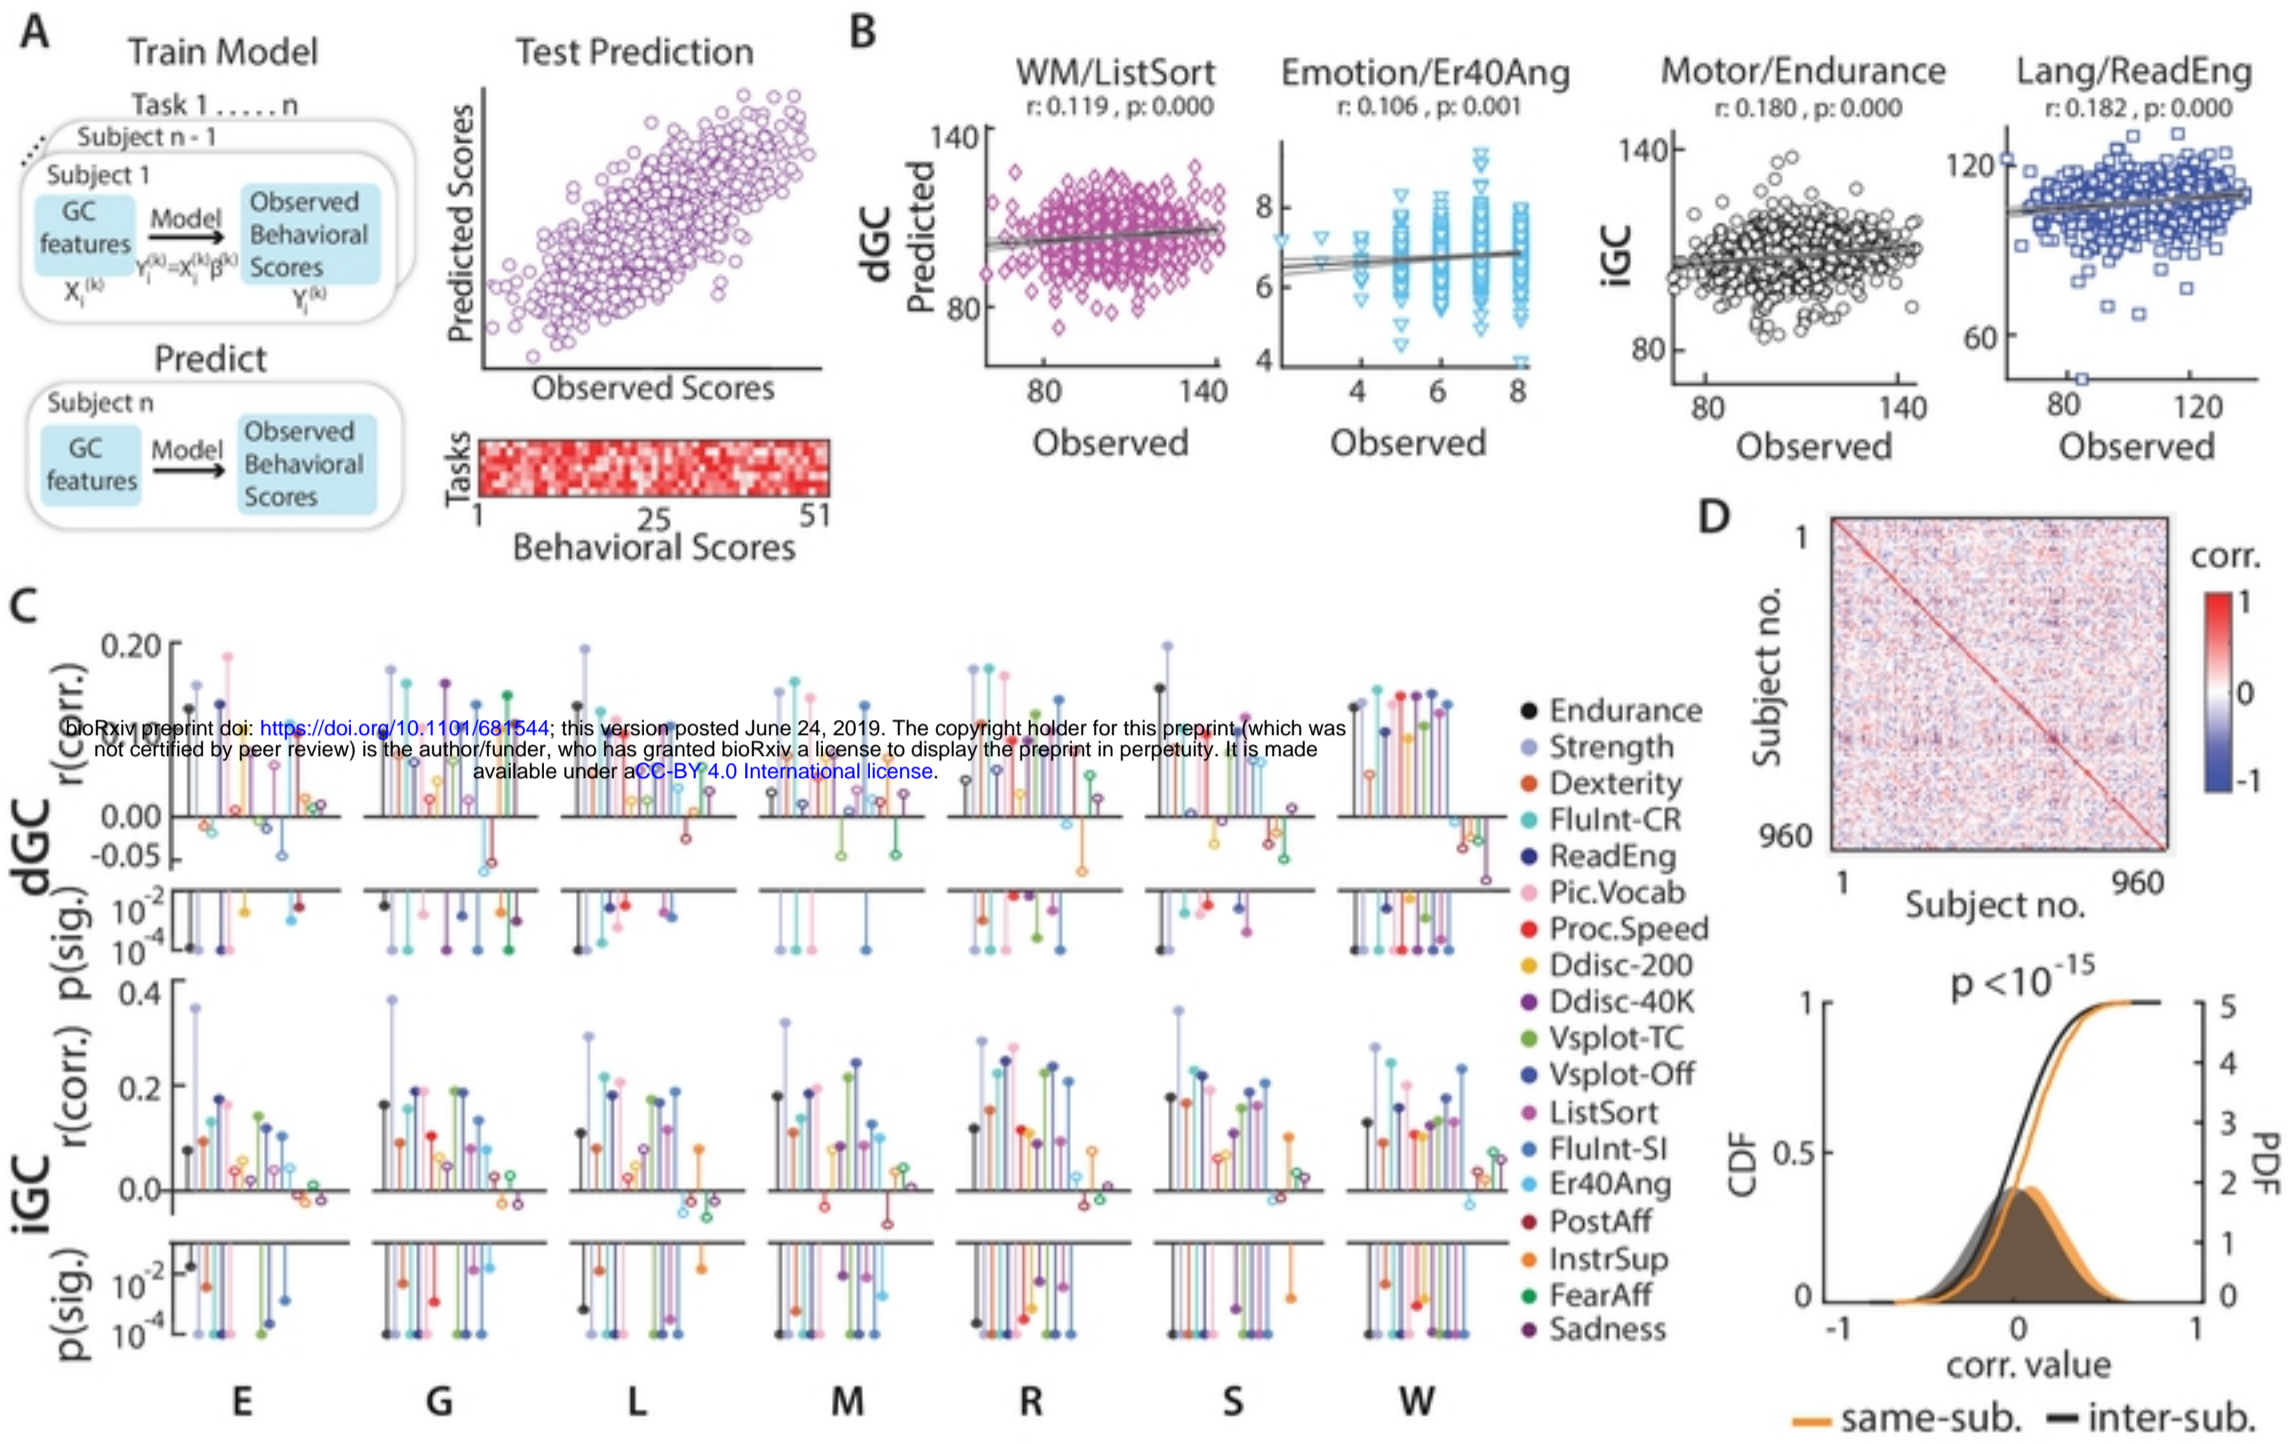


Figure 5