Analysis of the Human Connectome Data Supports the Notion of A "Common Model of Cognition" for Human and Human-Like Intelligence

Andrea Stocco[1], Zoe Steine-Hanson[2,§], Natalie Koh[3,†], John E. Laird[4], Christian J. Lebiere[5], and Paul Rosenbloom[6]

[1] Department of Psychology, University of Washington, Seattle, WA 98195

[2] Department of Computer Science, Oregon State University, Corvallis, OR 97331

[3] Department of Radiology, University of Washington, Seattle, WA 98195

[4] Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109

[5] Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213

[6] Institute for Creative Technologies, University of Southern California, Los Angeles, CA 90089

[§] Now at School of Computer Science and Engineering, University of Washington, Seattle, WA 98195

[†] Now at Department of Bioengineering, Northwestern University, Chicago, IL 60208

Corresponding Author: Andrea Stocco, Campus Box 351525, University of Washington, Seattle, WA, 98195. Phone: +1 206 543 4159; Email: stocco@uw.edu

Abstract

The Common Model of Cognition (CMC) is a recently proposed, consensus architecture intended to capture decades of progress in cognitive science on modeling human and human-like intelligence. Because of the broad agreement around it and preliminary mappings of its components to specific brain areas, we hypothesized that the CMC could be a candidate model of the large-scale functional architecture of the human brain. To test this hypothesis, we analyzed functional MRI data from 200 participants and seven different tasks that cover the broad range of cognitive domains. The CMC components were identified with functionally homologous brain regions through canonical fMRI analysis, and their communication pathways were translated into predicted patterns of effective connectivity between regions. The resulting dynamic linear model was implemented and fitted using Dynamic Causal Modeling, and compared against four alternative brain architectures that had been previously proposed in the field of neuroscience (two hierarchical architectures and two hub-and-spoke architectures) using a Bayesian approach. The results show that, in all cases, the CMC vastly outperforms all other architectures, both within each domain and across all tasks. The results suggest that a common, general architecture that could be used for artificial intelligence effectively underpins all aspects of human cognition, from the overall functional architecture of the human brain to higher level thought processes.

Introduction

The fundamental organizational principle of a complex system is often referred to as its "architecture," and represents an important conceptual tool to make sense of the relationship between a system's function and structure. For instance, the von Neumann architecture describes the organizing principle of modern digital computers; it can be used both to describe a computer at a functional level of abstraction (ignoring the specific wiring of its motherboard) and, conversely, to conduct diagnostics on an exceedingly complicated piece of hardware (properly identifying the components and pathways on a motherboard and the function of their wiring).

The stunning complexity of the human brain has inspired a search for a similar "brain architecture" that, akin to von Neumann's, could relate its components to its functional properties. Succeeding in this quest would lead to a more fundamental understanding of brain function and dysfunction and, possibly, to new principles that could further the development of artificial intelligence (Hassabis et al., 2017).

Most attempts in this direction have been "bottom-up," that is, driven by the application of dimensionality-reduction and machine-learning methods to large amounts of connectivity data, with the goal of identifying clusters of functionally connected areas (Cole et al., 2013; Gorgolewski et al., 2014; Huntenburg et al., 2018). Although these models can be used to predict task-related activity, they rely on large-scale connectivity and are fundamentally agnostic as to the function of each network node. The results of such approaches are also dependent on the type of data and the methods applied. For instance, one researcher might focus on purely functional measures, such as task-based fMRI and the co-occurrence of activity across brain regions and domains; a second researcher, instead, might focus on spontaneous, resting-state activity and slow-frequency time series correlations.

As recently pointed out (Jonas and Kording, 2017), none of these methods is guaranteed to

converge and provide a functional explanation *from* the data. However, the same methods can be successfully used to test a two or more models *against* the data via a "top-down" approach (Jonas and Kording, 2017). That is, given a candidate functional model of the brain, traditional connectivity methods can provide reliable answers as to its degree of fidelity to the empirical data and its performance compared to other models. A top-down approach, however, critically depends on having a likely and theoretically-motivated functional proposal for a brain architecture.

The Common Model of Cognition

A promising candidate architecture is the Common Model of Cognition (CMC: Laird et al., 2017). The CMC is an architecture for general intelligence that reflects the historical convergence of multiple computational frameworks (developed over the course of five decades in the fields of cognitive psychology, artificial intelligence, and robotics) that agree on a common set of organizing principles. Because of its generality and consensus, it has been used as a "reference architecture" for cognitive agents (Mohan, submitted) The CMC assumes that agents exhibiting human-like intelligence share five functional components: A feature-based, declarative *long-term memory*, a buffer-based *working memory*, a reinforcement-learning-based set of state-action patterns represented in *procedural memory*, and dedicated *perception* and *action* systems. Working memory acts as the hub through which all of the other components communicate, with one additional connection between perception and action (Figure 1A). The CMC also includes a set of constraints on the mechanisms and representations that characterize each component's functional properties.

The CMC's components and assumptions distill lessons learned over the last fifty years in the development of computational cognitive models and artificial agents with general human-like abilities. Surprisingly, these lessons seem to cut across the specific domains of application. For

instance, the cognitive architecture Soar (Laird, 2012) is predominantly used in designing autonomous artificial agents and robots, while the cognitive architecture ACT-R (Anderson, 2007) is predominantly used to simulate psychological experiments and predict human behavior (Kotseruba and Tsotsos, 2018); yet, they separately converged on many of the CMC assumptions (Laird et al., 2017). Similarly, the SPAUN large-scale brain model (Eliasmith et al., 2012) and the Leabra neural architecture (O'Reilly et al., 2016) are independently designed to simulate brain function through artificial neurons; despite making different assumptions in terms of neural coding, representation, and learning algorithms, they agree on the use of high-level modules (including ones for working memory, procedural memory, and long-term memory) that are similar to the CMC. Even recent AIs that are made possible by advances in artificial neural networks employ, at some level, the same components. DeepMind's AlphaGo, for example, includes a monte-carlo search tree component for look-ahead search and planning (working memory) and a policy network (i.e., procedural memory), in addition to dedicated systems for perception and action (Silver et al., 2016). Similarly, the Differentiable Neural Computer uses supervised methods to learn optimal policies (procedural memory) to access an external memory (symbolic long-term memory: Graves et al., 2016).

Because the CMC reflects the general organization of systems explicitly designed to achieve human-like flexibility and intelligence, the CMC should also apply to the human brain. Therefore, it provides an ideal candidate for a top-down examination of possible brain architecture.

Assuming that the CMC is a valid candidate, how can its viability as a model of the human brain architecture be assessed? Operationally, a candidate model should successfully satisfy two criteria. The first is the *generality* criterion: The same cognitive architecture should account for brain activity data across a wide spectrum of domains and tasks. The second is the *comparative superiority* criterion: An ideal architecture should provide a superior fit to experimental brain data

compared to competing architectures of similar complexity.

To test the CMC against these two criteria, we conducted a comprehensive analysis of task-related neuroimaging data from 200 young adult participants in the Human Connectome Project (HCP), the largest existing repository of high-quality human neuroimaging data. Although the HCP project contains both fMRI and MEG data, fMRI was chosen because it allows for unambiguous identification of subcortical sources of brain activity, which is crucial to the CMC and problematic for MEG analysis. The HCP includes functional neuroimages collected while participants performed seven psychological tasks. These tasks were taken or adapted from previously published influential neuroimaging studies and explicitly selected to cover the range of human cognition (Van Essen et al., 2013), therefore making it an ideal testbed for the *generality* criterion. Specifically, the tasks examine language processing and mathematical cognition (Binder et al., 2011), working memory, incentive processing and decision making (Delgado et al., 2000), emotion processing (Hariri et al., 2002), social cognition (Wheatley et al., 2007), and relational reasoning (Smith et al., 2007). The seven tasks were collected from six different paradigms (language processing and mathematical cognition were tested in the same paradigm).

To properly translate the CMC into a brain network architecture, its five components need to be identified with an equal number of spatially-localized but functionally homologous Regions of Interest (ROIs). This initial identification was based on well established findings in the literature and is also consistent with the function-to-structure mappings that had been proposed in other neurocognitive architectures, such as the mappings suggested for ACT-R's module-specific buffers (Anderson, 2007; Borst et al., 2015; Borst and Anderson, 2013) and the functional components employed in large-scale models of the brain (Eliasmith et al. 2012; O'Reilly et al. 2016). At this level, the working memory (WM) component was identified with the fronto-parietal network comprising the dorsolateral prefrontal cortex (PFC) and posterior parietal cortex; the long-term memory (LTM) component with regions in the middle, anterior, and superior

temporal lobe; the procedural knowledge component with the basal ganglia; the action component with the premotor and primary motor cortex; and perception with sensory regions, including the primary and secondary sensory and auditory cortices, and the entire ventral visual pathway, the latter of which comprises the occipital and inferior temporal lobes (Figure 1B).

To properly characterize each individual, a processing pipeline was designed to progressively identify the relevant corresponding ROIs for each task and, within each task, for each of the ~200 participants, thus accounting for individual differences in functional neuroanatomy (Figure 1C; See Materials & Methods).

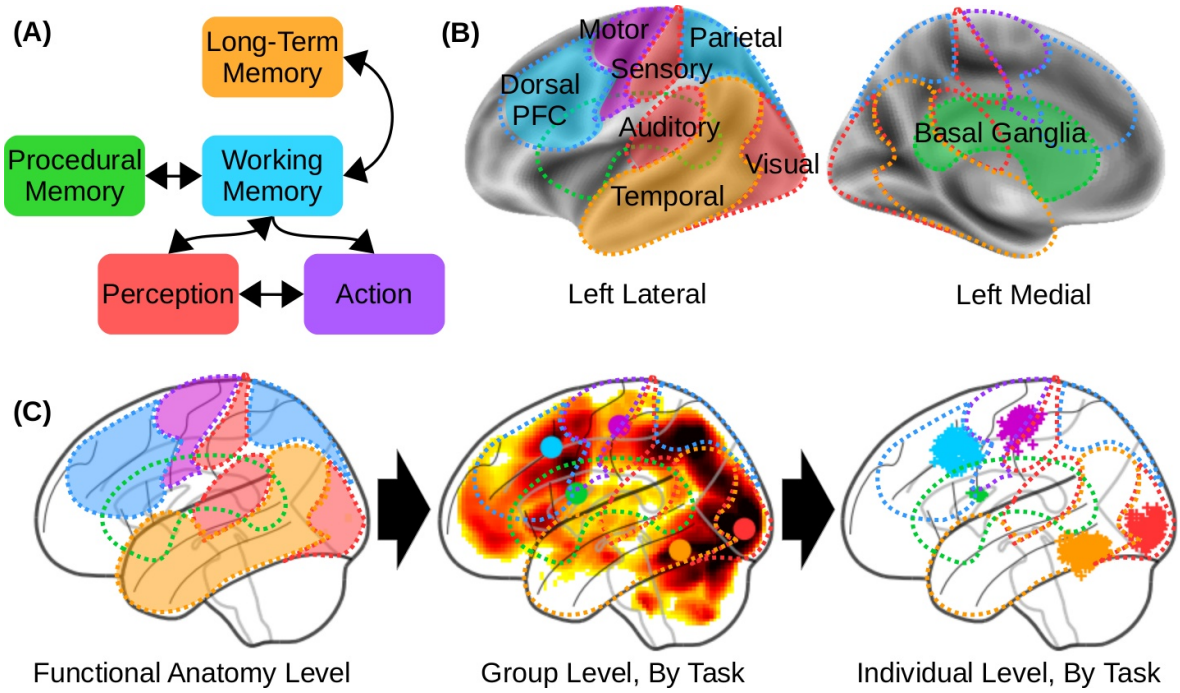

Figure 1. (A) Architecture of the Common Model of Cognition, as described by (Laird et al., 2017). (B) Theoretical mapping between CMC components and homologous cortical and subcortical regions, as used in this study's pipeline to identify the equivalent Regions of Interest (ROIs). (C) Progressive approximation of the ROIs, from high-level functional mappings (left) to task-level group results (middle, with group-level centroid coordinated marked by a color circle) to the individual functional centroids of the regions in our sample (right; each individual centroid represented by a "+" marker). Group-level and individual-level data come from the Relational Reasoning task.

Alternative Architectures

To address our second criterion of *comparative superiority*, the CMC dynamic model was compared against other DCM models that implement alternative brain architectures. Because the space of possible models is large, we concentrated on four models that are representative of theoretical neural architectures previously suggested in the neuroscientific literature (Figure 2). These four models can be divided into two families. In the "Hierarchical" family, brain connectivity implements hierarchical levels of processing that initiate with Perception and culminate with Action. In this family, the brain can be abstracted as a feedforward neural network model with large-scale gradients of abstraction (Huntenburg et al., 2018).

Within this hierarchical structure, each ROI represents a different level and projects both forward to the next level's ROI and backwards to the preceding level's ROI. In the *open-loop* variant, Perception gives rise to two distinct branches, one feeding the LTM component (representing the abstraction of perception into memory) and one ascending to procedural memory, WM and, eventually, Action components. The second, *closed-loop* hierarchical model also incorporates bidirectional connections between LTM and WM, thus reconnecting the two pathways into a loop (Figure 2B).

In the "Hub and Spoke" family (Figure 2C), a single ROI is singled out as the network's "Hub" and receives bidirectional connections from all the other ROI (the "Spokes"). With the exception of the Hub, no ROI is mutually connected to any other one. In the first variant, the role of the Hub is played by the WM component. Because, in our mapping, the WM component corresponds to the lateral PFC, this model captures the increasingly popular and supported view of PFC as a flexible hub for control (Cole et al., 2013, 2012). In the second variant, the role of the Hub is played by the Procedural Memory component, which reflects the centrality of procedural control in many production-system-based cognitive architectures (Anderson, 2007; Kieras and

Meyer, 1997; Laird, 2012). Because, in our mapping, Procedural Memory is identified with the basal ganglia, this architecture also reflects the centrality of these nuclei in action selection and in coordinating cortical activity (Eliasmith et al., 2012; Hazy et al., 2007; Stocco et al., 2010).

Like the CMC, these architectures are representative of how the five components could be organized in a large-scale conceptual blueprint for the brain architecture; they simply make different choices as to which connections between components are more fundamental and better reflect the underlying neural organization.In addition to representing plausible alternative architectures, these alternative models differ minimally from the CMC and can be easily generated by replacing at most five connections from the CMC architecture (dashed lines and red lines, Figure 2B-C). Thus, any resulting differences in fit are unlikely to arise because of differences in network complexity.

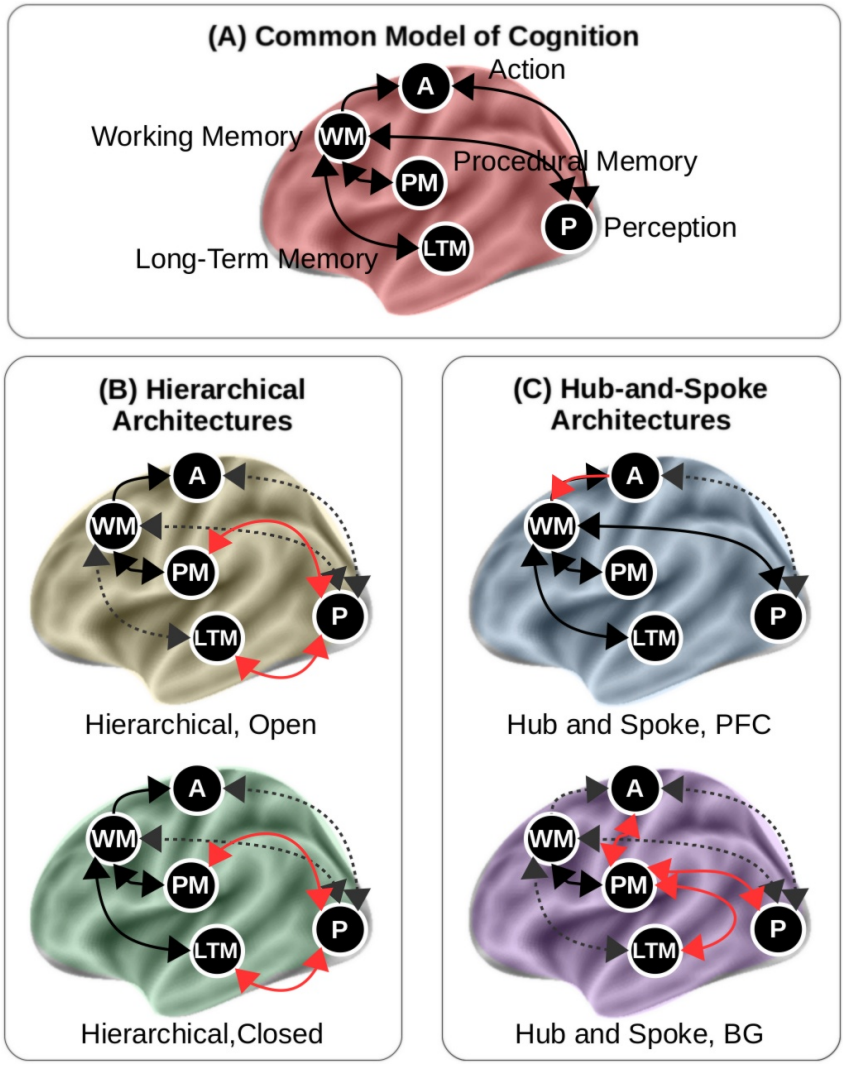

Figure 2: The five architectures tested in this study. (A) The CMC; (B) The "Open" and "Closed" versions of the hierarchical family; (C) The prefrontal (PFC) and basal ganglia (BG) versions of the hub-and-spoke architecture. In (B) and (C), pathways that are common to the CMC are shown in black; pathways that are present in the CMC but not included in the alternative models are shown as grey dashed arrows; and pathways that are present in the alternative models but not in the CMC are shown in red.

Modeling Network Dynamics

The link between the network of ROIs and their neural activity was provided through Dynamic Causal Modeling (DCM) (Friston et al., 2003), a neuronal-mass mathematical modeling technique that approximates the time-course of brain activity in a set of brain regions as a dynamic system that responds to a series of external drives. Specifically, the time course of the underlying neural activity y of a set of regions is controlled by the bilinear state change equation:

$$dy/dt = \mathbf{A}y + \sum_i x_i \mathbf{B}^i y + \mathbf{C}x \qquad (1)$$

where x represents the event vectors (i.e., the equivalent of a design matrix in traditional GLM analysis), **A** defines intrinsic connectivity between ROIs, **C** defines the ROI-specific effects of task events, and **B** defines the modulatory effects that task conditions have on the connectivity between regions. For simplicity, the modulatory effects in **B** were set to zero, reducing the equation to the form $\mathbf{A}y + \mathbf{C}x$. A predicted time course of BOLD signal was then generated by applying a biologically-plausible model (the balloon model: (Buxton et al., 1998; Friston et al., 2000)) of neurovascular coupling to the simulated neural activity y. The parameters of the full DCM model were estimated by applying the expectation-maximization procedure (Friston et al., 2003) to reduce the difference between the predicted and observed time course of the BOLD signal in each ROI.

Our preference for this technique was motivated by the existence of an integrated framework to design, fit, and evaluate models; by its ability to estimate the directional effects within a network (as opposed to traditional functional connectivity analysis); and by its underlying distinction between the modeling of network dynamics and the modeling of recorded imaging signals (as opposed to Granger causality), which makes it possible to apply the same neural models to different modalities (e.g., M/EEG data) in future work.

Materials and Methods

The study presented herein consists of an extensive analysis of a large sample (*N*=200) of neuroimaging data from the Human Connectome Project, the largest existing repository of young adult neuroimaging data. The analysis was restricted to the task fMRI subset, thus excluding both the resting state fMRI data, the diffusion imaging data, and all of the M/EEG data. The task fMRI data consisted of two sessions of each of seven paradigms, designed to span different domains.

Tasks fMRI Data

The HCP task-fMRI data encompasses seven different paradigms designed to capture a wide range of cognitive capabilities. Of these paradigms, six were included in our analysis. The Motor Mapping task was not included because it would have required the creation of multiple ROIs in the motor cortex, one for each effector (arm, leg, voice), thus making this model intrinsically different from the others. A full description of these tasks and the rationale for their selection can be found in the original HCP papers (Barch et al., 2013; Van Essen et al., 2013). This section provides a brief description of the paradigms, while Table 1 provides an overview.

Table 1: Overview of the seven task-fMRI paradigms used in the HCP dataset. Italics indicate tasks and conditions that were not included in our analysis; bold typeface marks experimental conditions that were selected as "Critical" (as opposed to "Baseline") in the design of the experimental matrices (see below, "DCM-specific GLM analysis" section)

Task (Representative Reference)	**Relevant Conditions (for GLM analysis)**	**Included in DCM analysis?**
Motor Mapping (Buckner et al., 2011)	*Hand, arm, foot, leg, voice responses*	*No*
Emotion Processing (Hariri et al., 2002)	Neutral shapes vs. **Fearful and angry faces**.	Yes
Incentive Processing (Delgado et al., 2000)	"Winning" vs. **"Losing"** blocks of choices	Yes
Language and Mathematical Processing (Binder et al., 2011)	Listening vs. **Answering questions** (in both Language and Math blocks)	Yes
Relational Reasoning (Smith et al., 2007)	Control Arrays vs. **Relational arrays**	Yes
Social Cognition (Wheatley et al., 2007)	Randomly moving shapes vs. **Socially interacting shapes**	Yes
Working Memory	0-Back vs. **2-Back blocks** of faces, places, tools, and body parts.	Yes

Emotion Processing Task. Participants are presented with 12 blocks of six consecutive trials. During each trial, they are asked to decide either which of two visual stimuli presented on the bottom of the screen match the stimulus at the top of the screen. In six of the blocks, all of the visual stimuli are emotional faces, with either angry or fearful expressions. In the remaining six blocks, all of the stimuli are neutral shapes. Each stimulus is presented for 2 s, with a 1 s inter-trial interval (ITI). Each block is preceded by a 3 s task cue ("shape" or "face"), so that each block is 21 s including the cue.

Incentive Processing Task. The task consists of four blocks of eight consecutive decision-making trials. During each trial, participants are asked to guess whether the number underneath a "mystery card" (visually represented by the question mark symbol "?") is larger or smaller than 5 by pressing one of two buttons on the response box within the allotted time. After each choice, the number is revealed; participants receive a monetary reward (+$1.00) for correctly guessed trials; a monetary loss (-$0.50) for incorrectly guessed trials; and receive no money if the number is exactly 5. Unbeknownst to participants, blocks are pre-designed to lead to either high rewards (6 reward trials, 2 neutral trials) or high losses (6 loss trials, 2 neutral trials), independent of their actual choices. Two blocks are designated as high-reward, and two as high-loss blocks. Each stimulus has a duration of up to 1.5 s, followed by a 1 s feedback, with a 1 s ITI, so that each block lasts 27 s.

Language and Mathematical Processing Task. The task consists of 4 "story" blocks interleaved with 4 "math" blocks. The two types of blocks are matched for duration, and adhere to the same internal structure in which a verbal stimulus is first presented auditorily, and a two-alternative question is subsequently presented. Participants need to respond to the question by pressing one of two buttons with the right hand. In the story blocks, the stimuli are brief, adapted Aesop stories (between 5 and 9 sentences), and the question concerns the story's topic (e.g., "Was the story about *revenge* or *reciprocity*?"). In the math blocks, stimuli are addition or subtraction

problems (e.g., “Fourteen plus twelve”) and the question provides two possible alternative answers (e.g., “*Twenty-nine* or *twenty-six*?”). The math task is adaptive to maintain a similar level of difficulty across the participants.

Relational Processing Task. The task consists of six “Relational” blocks alternated with six “Control” blocks. In relational blocks, stimuli consist of two pairs of figures, one displayed horizontally at the top of the screen and one pair displayed at the bottom. Figures consist of one of six possible shapes filled with one of six possible textures, for a total of 36 possible figures. Both pairs of figures differ along one dimension, either shape or texture; participants are asked to indicate through a button press if the top figures differ on the same dimension as the bottom figures (e.g., they both differ in shape). In the control blocks, the stimuli consist of one pair of figures displayed horizontally at the top of the screen, a third figure displayed centrally at the bottom of the screen, and a word displayed at the center of the screen. The central word specifies a stimulus dimension (either “shape” or “texture”) and participants are asked to indicate whether the bottom figure matches either of the two top figures along the dimension specified by the word. Both relational and control blocks have a total duration of 16 s, but they vary in the number of stimuli. Specifically, relational blocks contain four stimuli, presented for 3.5 s with a 500 ms ITI, while control blocks contain five stimuli presented for 2.8 s with a 400 ms ITI.

Social Cognition Task. The task consists of 10 videoclips of moving shapes (circles, squares, and triangles). The clips were either obtained or modified from previously published studies (Castelli et al., 2000; Wheatley et al., 2007). In five of the clips, the shapes are moving randomly, while in the other five the shapes’ movement reflects a form of social interaction. After viewing each clip, participants press one of three buttons to indicate whether they believed the shapes were interacting, not interacting, or whether they were unsure. All clips have a fixed duration of 20 s with an ITI of 15 s.

Working Memory. The task consists of eight 2-back blocks and eight 0-back blocks, with

each block containing 10 trials. Each trial presents the picture of a single object, centered on the screen, and participants have to press one of two button to indicate whether the object is a target or not. In the 2-back blocks, a target is defined as the same object that had been seen two trials before, so that participants have to maintain and update a "moving window" of the past two objects to perform the task correctly. In the 0-back blocks, a target is defined as a specific object, presented at the very beginning of the block, so that participants have to only maintain a single object in working memory throughout the block. The stimuli belong to one of four possible categories: faces, places, tools, and body parts. The category of the objects being used as stimuli changes from block to block, but is consistent within one block, so that there is an even number of face, place, tool, and body part blocks for each condition. Each block begins with a 2.5 s cue that informs the participant about the upcoming block type (2-back or 0-back). Each stimulus is presented for 2 s with a 500 ms ITI, for a total duration of 27.5 s per block.

Data Processing and Analysis

Imaging Acquisition Parameters As reported in (Barch et al., 2013), functional neuroimages were acquired with a 32-channel head coil on a 3T Siemens Skyra with TR = 720 ms, TE = 33.1 ms, FA = 52°, FOV = 208 × 180 mm. Each image consisted of 72 2.0mm oblique slices with 0-mm gap in-between. Each slice had an in-plane resolution of 2.0 x 2.0 mm. Images were acquired with a multi-band acceleration factor of 8x.

Image Preprocessing Images were acquired in the "minimally preprocessed" format (Van Essen et al., 2013), which includes unwarping to correct for magnetic field distortion, motion realignment, and normalization to the MNI template. The images were then smoothed with an isotropic 8.0 mm FWHM Gaussian kernel.

Canonical GLM Analysis Canonical GLM analysis was conducted on the smoothed minimally preprocessed data using a mass-univariate approach, as implemented in the SPM12 software package (Penny et al., 2011). First-level (i.e., individual-level) models were created for

each participant. The model regressors were obtained by convolving a design matrix with a hemodynamic response function; the design matrix replicated the analysis of (Barch et al., 2013), and included regressors for the specific conditions of interests described in Table 1. Second-level (i.e., group-level) models were created using the brain-wise parametric images generated for each participant as input.

DCM-specific GLM Analysis In parallel with the canonical GLM analysis, a second GLM analysis was carried out as part of the DCM analysis pipeline. The purpose of this analysis was two-fold. First, it defined the event matrix $\boldsymbol{x}$ that is used in the DCM equation (Eq. 1) to measure the parameter matrix **C**. Second, it provided a way to define the omnibus *F*-test that is used in the ROI definition (see below). Because these models are not used to perform data analysis, the experimental events and conditions are allowed to be collinear.

Like most cognitive neuroscience paradigms, each of our tasks includes at least two different conditions, under which stimuli must be processed in different ways. In all cases, the difference between conditions can be framed in terms of a more demanding, "critical" condition and an easier, "control" condition, with the more demanding events associated with greater mental elaboration of the stimuli. The critical condition of each task is marked in boldface in Table 1.

As is common in DCM analysis, these two task conditions were modeled in a layered, rather than orthogonal fashion. The difference is illustrated in Figure 3: While in traditional GLM analysis the two conditions are modeled as non-overlapping events in the design matrix, in the DCM-specific definition of the matrix all trials belong to the same "baseline" condition, which represents the basic processing of the stimulus across all trials. Stimuli from the critical condition form a subset of all stimuli presented in the baseline condition. The critical condition is therefore appended to the baseline condition in the design matrix to model the additional processes that are specifically related to it.

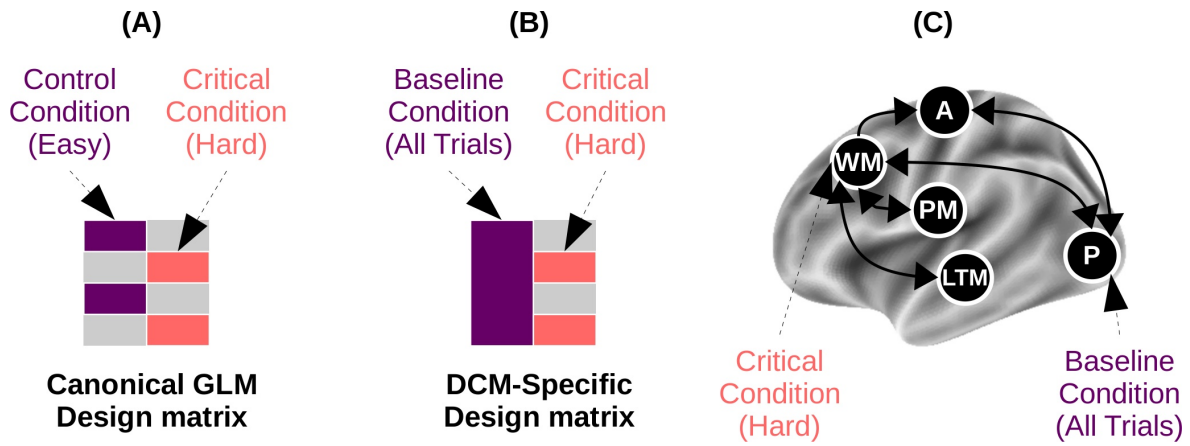

Figure 3: Difference between design matrix used for canonical GLM (A) and for DCM analysis (B). In all of the network models, the Baseline condition drives neural activity in perceptual areas, while the Critical condition drives neural activity in the Working Memory component (C).

In DCM, each condition can affect one or more ROIs independently. In our analysis, the association between conditions and ROIs was kept constant across all tasks. Specifically, the baseline conditions selectively affected the perceptual ROI, while the critical condition selectively affected the WM ROI. This choice reflects the greater mental effort that is common to all critical conditions, and is confirmed by the greater PFC activity found in all of the GLM analyses of the critical conditions (Barch et al., 2013).

Regions-of-Interest Definition

To objectively define the Regions-of-Interest (ROIs) for each task and participant, a processing pipeline was set up. The starting point of the pipeline was an *a priori*, theoretical identification of each CMC component with large-scale neuroanatomical distinctions. As noted in the main text, this initial identification was based on well established findings in the literature as well as the function-to-structure mappings proposed in other large-scale neurocognitive architectures (Anderson 2007; Borst et al. 2015; Borst and Anderson 2013; Eliasmith et al. 2012; O'Reilly et al. 2016). Specifically, working memory (WM) was identified with the fronto-parietal

network comprising the dorsolateral prefrontal cortex (PFC) and posterior parietal cortex; long-term memory (LTM) with regions in the middle, anterior, and superior temporal lobe; the procedural knowledge component with the basal ganglia; the action component with the premotor and primary motor cortex; and perception with sensory regions, including the primary and secondary sensory and auditory cortices, and the entire ventral visual pathway (Figure 1B).

Beginning with these macro-level associations, the pipeline progressively refined the exact ROI for each component through two consecutive approximations. Fig 1C provides a visual illustration of this procedure using the data from the relational reasoning task.

The first approximation was designed to account for group-level variability due to the different tasks and stimuli used in the four datasets. This was necessary because, for example, the different stimulus modalities determine which sensory area (e.g., auditory vs. visual areas) would be engaged and different task requirements would recruit different portions of the PFC. These differences were accounted for by conducting a separate group-level GLM analysis for each dataset, and identifying the coordinates of three points that have the highest statistical response within the anatomical boundaries of the visual areas (limited to the occipital lobe and the ventral portion of the temporal lobe), the dorso-lateral PFC, and the basal ganglia (limited to the striatum).

The second approximation was designed to account for individual-level variability in functional neuroanatomy. The group-level coordinates of each component, derived from the previous step, were then used as the starting point to search in 3D space for the closest active peak within the individual statistical parameter maps obtained from GLM models of each participant (see Figure 1C, right panel). For maximal sensitivity, the map was derived from an omnibus *F*-test that included all the experimental conditions. In practice, this *F*-test was designed to capture any voxel that responded to any experimental condition. The same *F*-contrast was also used to adjust (i.e., mean-correct) each ROI's timeseries (Ashburner et al., 2016; Penny et al., 2011).

The individual coordinates, thus defined, were then visually inspected; when the coordinates were outside the predefined anatomical boundaries, manually re-adjusted. Across over 1,200 coordinates examined, only 2 required manual adjustment (~ 0.2%). Figure 4 illustrates the distribution of the individual coordinates of each region for each task, overlaid over a corresponding group-level statistical map of task-related activity. Each individual coordinate is represented by a crossmark; the ~200 crossmark form a cloud that captures the spatial variability in the distribution of coordinates.

Finally, the individualized ROI coordinates were then used as the center of a spherical ROI. All voxels within the sphere whose response was significant at a minimal threshold of $p < 0.5$ were included as part of the ROI. For each ROI of every participant in every task, a representative time course of neural activity was extracted as the first principal component of the time series of all of the voxels within the sphere.

All the ROIs thus obtained were located in the left hemisphere: this simplifying approach was preferred to possible alternatives, such as including homologous regions in the right hemisphere (which would have required introducing additional assumptions about inter-hemispheric connectivity) or creating bi-lateral ROIs (which would have reduced the amount of variance captured in each ROI). Because all tasks show stronger activation in the left hemisphere than in the right, our results are still representative of brain activity in these domains.

Finally, network models were created by connecting all the individually-defined ROIs according to the specifications of each architecture (Figure 2). It should be noted that synaptic pathways exist that connect every pair of components; thus, this network model is designed to capture the fundamental layout of a brain architecture in terms of functionally necessary connections, rather than anatomical details.

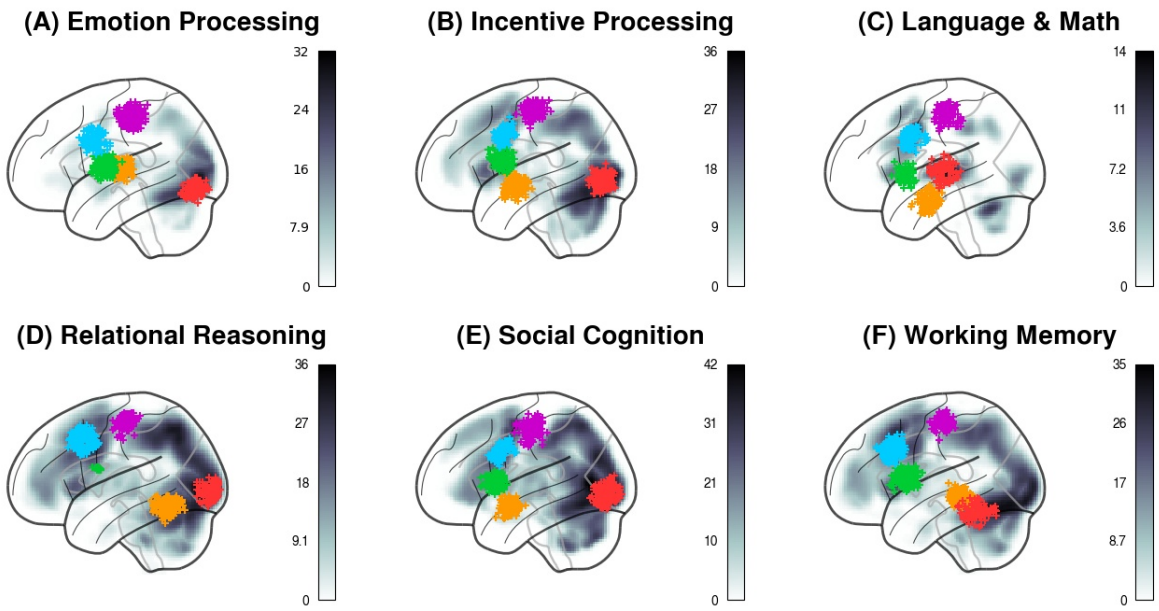

Figure 4: Lateral view of the distribution of the ROI centroids across individual participants and tasks. Each "+" marker represents the centroid of an ROI for one participant. Colors represent the components, following the conventions of Fig 1A-C. The background represents the statistical parametric map (in greyscale) of the corresponding group-level analysis used to identify the seed coordinates for each ROI (Step 2 in Figure 1C).

Results

Once the five DCM models were separately fitted to the functional neuroimaging data, they were compared against each other using a Bayesian random-effects procedure (Stephan et al., 2009). Like many other model comparison procedures, this approach provides a way to balance the complexity of a model (as the number of free parameters) versus its capacity to fit the data. Compared to popular log-likelihood-based measures (e.g., Akaike's information criterion (Akaike, 1974)), this procedure is more robust in the face of outlier subjects, and thus better suited for studies that, like the present one, include a large number of participants and deal with

considerable inter-individual variability (Stephan et al., 2010, 2009).

Specifically, the probability r_k that a model k would fit a random individual in a sample of participants is drawn from a Dirichlet distribution Dir(α_1, α_2,… α_K), and the distributions of probabilities of architectures 1, 2… k across n individuals are then drawn from multinomial distributions (see Figure 5A). The result of this modeling effort is a posterior distribution of the probabilities r_k for each model. These distributions can then be compared in terms of their relative *expected* and *exceedance* probability, that is, the mean probability of each model's r_k across the sample and the probability that r_k is larger than the competing models (Figure 5B). Expected probability is calculated as the mean of each distribution; the properties of the Dirichlet distribution guarantee that the sum of the means of all distributions is 1. The Exceedance probability can be calculated by sampling from a multinomial distribution generated from random samples of the original distributions, thus again guaranteeing that all probabilities sum up to 1. Figure 5, inspired by (Stephan et al., 2009), provides a graphical illustration of the procedure.

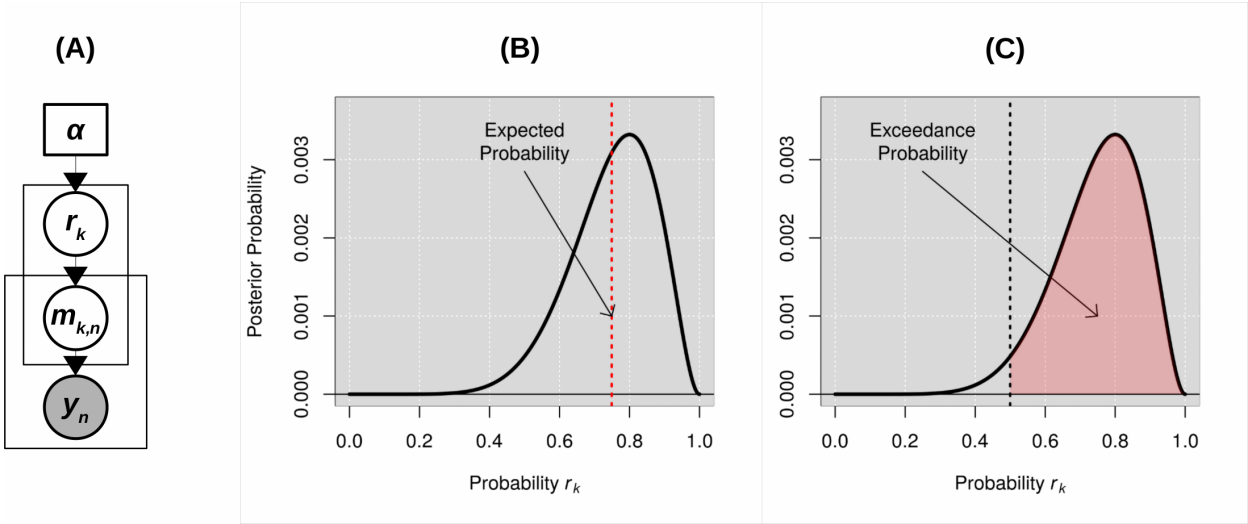

Figure 5. (A) Visual representation of the hierarchical Bayesian modeling procedure. (B) Visual representation of the relationship between a model's probability distribution r_k *(purple curve), its expected probability (left), and its exceedance probability (right) in the case of two possible models (k = 2). In this case, the exceedance probability reduces to the area to the right of* $r_k = 0.5$*). Modified from (Stephan et al., 2009).*

The model posterior distributions are visualized for each task in Figure 6A-F. The expected probabilities are represented as the colored vertical lines, while the exceedance probabilities are summarized as colored bars in Figure 6H. Table 2 provides a detailed list of model comparison metrics, including the ones derived from the hierarchical Bayesian procedure used in this study (Dirichlet's α, expected, and exceedance probabilities) as well as the group-level log-likelihood of each model.

Table 2: Results of Bayesian model comparison across tasks and models.

Task	Model	Dirichlet	Expected Probability	Exceedance Probability	Log-likelihood
All Tasks Combined	Common Model	134.23	0.7759	1.0000	-3,766,837.56
	Hierarchical Closed	21.8	0.1260	0.0000	-3,797,769.54
	Hierarchical Open	0.99	0.0057	0.0000	-4,292,199.54
	Hub-and-spoke BG	14.99	0.0866	0.0000	-4,250,392.03
	Hub-and-spoke PFC	0.99	0.0057	0.0000	-4,300,912.00
Emotion Processing	Common Model	101.37	0.5307	1.0000	-858,282.97
	Hierarchical Closed	3.68	0.0193	0.0000	-865,011.12
	Hierarchical Open	37.27	0.1951	0.0000	-860,955.74
	Hub-and-spoke BG	37.66	0.1972	0.0000	-861,523.46
	Hub-and-spoke PFC	11.02	0.0577	0.0000	-861,878.56
Incentive Processing	Common Model	144.77	0.7131	1.0000	-861,026.48
	Hierarchical Closed	1.36	0.0067	1.0000	-869,976.35
	Hierarchical	39.82	0.1962	0.0000	-866,148.11

	Open				
	Hub-and-spoke BG	14.8	0.0729	0.0000	-869,471.16
	Hub-and-spoke PFC	2.26	0.0111	0.0000	-866,795.76
Language & Math	Common Model	85.83	0.4494	0.7526	-894,992.94
	Hierarchical Closed	7.74	0.0405	0.0000	-903,067.62
	Hierarchical Open	76.84	0.4023	0.2474	-895,740.71
	Hub-and-spoke BG	13.65	0.0715	0.0000	-918,300.93
	Hub-and-spoke PFC	6.94	0.0363	0.0000	-902,568.43
Relational Reasoning	Common Model	77.8	0.4116	0.9505	-746,544.42
	Hierarchical Closed	8.35	0.0442	0.0000	-749,922.68
	Hierarchical Open	57.81	0.3059	0.0490	-747,385.57
	Hub-and-spoke BG	42.56	0.2252	0.0005	-672,161.79
	Hub-and-spoke PFC	2.48	0.0131	0.0000	-750,719.30
Social	Common Model	141.47	0.7368	1.0000	-712,975.64
	Hierarchical	1.37	0.0072	0.0000	-720,399.27

Cognition	Closed				
	Hierarchical Open	33.34	0.1737	0.0000	-716,591.85
	Hub-and-spoke BG	13.57	0.0707	0.0000	-719,374.15
	Hub-and-spoke PFC	2.24	0.0117	0.0000	-715,828.71
Working Memory	Common Model	142.87	0.7441	1.0000	-113,247.58
	Hierarchical Closed	46.13	0.2403	0.0000	-112,996.33
	Hierarchical Open	1.00	0.0052	0.0000	-680,320.81
	Hub-and-spoke BG	0.99	0.0051	0.0000	-680,396.45
	Hub-and-spoke PFC	1.01	0.0053	0.0000	-679,145.16

Both metrics provide evidence in favor of the CMC. As shown in Figure 6A-F, the CMC provides a better fit to the data than any alternative architecture, and its exceedance probabilities range from 0.75 to 1.0 (Figure 6H). Thus, the CMC uniquely satisfies both the generality and comparative superiority criteria. By contrast, all of the other architectures are consistently outperformed by the CMC in every domain (violating comparative superiority) and their relative rankings change from task to task (violating generality). It is interesting to note that, across tasks, the relative ranking of the architectures does not reflect their relative similarity, measured as the number of different connections. For instance, the PFC variant of the "Hub-and-Spoke" family (Figure 2C) is the one model most similar to the CMC, but is consistently outperformed by other models that, across all tasks, come closer to the CMC distributions (Figure 6A-F). This fact further suggests that the CMC superiority is due to the holistic nature of its connectivity (i.e., how the components "go together") rather than the sum of its specific connectivity elements.

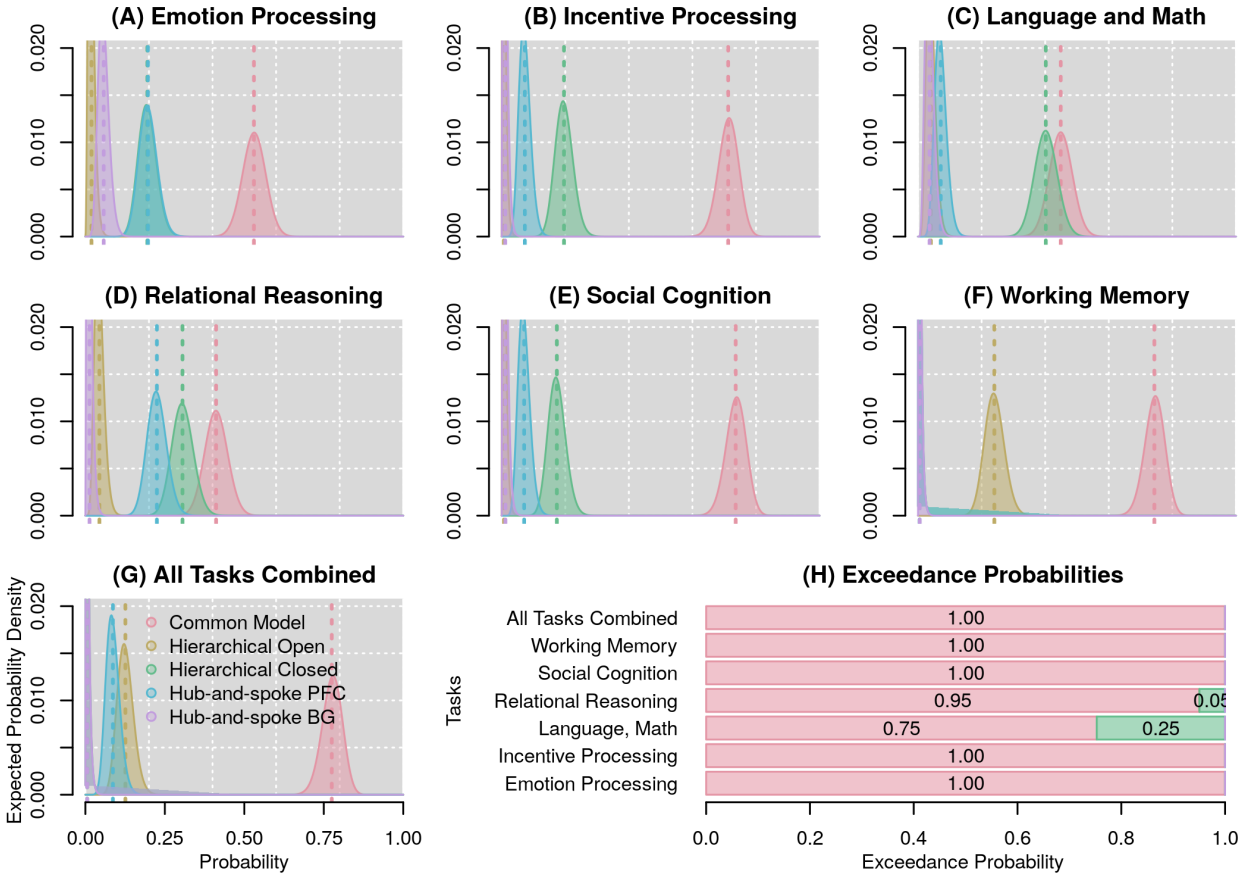

Figure 6: Results of the Bayesian model comparisons. In all plots, different colors represent different architectures.(A-G) Probability distributions that each of the five architectures is true, given the data within each task and across all tasks combined. Vertical dotted lines represent the mean of each distribution, i.e. the expected probability of each model. (H) Corresponding exceedance probabilities, represented as stacked bars for each task.

The only task in which another model comes close to the CMC in terms of fit was the Language and Mathematical cognition paradigm, in which the Hierarchical Closed model had a 0.25 exceedance probability against the CMC (Figure 6H). This paradigm was unique because it included two entirely different tasks of comparable difficulty, instead of a single task with two conditions of different difficulty, as was the case in all other tasks. This peculiarity raises a potential concern that the CMC's superiority could be an artifact of modeling each task in isolation, and that in conditions where multiple tasks were modeled simultaneously, a different model could potentially provide a superior fit. To examine this possibility, a second analysis was carried out, which included only the 168 participants for whom data for all seven tasks was available. In this analysis, the data from each of the six paradigms performed by the same individual is modeled as a different run from a "meta-task" performed by that individual. When such an analysis was performed, the CMC maintained its superiority, all other models having a combined exceedance probability $< 1.0 \times 10^{-10}$ (Figure 6G-H).

As noted earlier, although the competing architectures were chosen to represent current alternatives views, we cannot entirely rule out the existence of alternative architectures that explain the data better than the CMC. It is possible, however, to decide whether all of the connections in the CMC are necessary, or whether a simpler model could potentially fit the data equally well. To this end, a Bayesian parameter averaging procedure (Kasess et al., 2010) was conducted to generate the posterior distributions of the intrinsic connectivity parameter values (corresponding to matrix A in Eq. 1) across participants for each task. Figure 7 visually depicts the six task-specific connectivity matrices, indicating both the mean value (as the matrix cell color) and the associated posterior probability (as the overlaid number) for each CMC connection in each task. As the figure shows, the parameter values change significantly from task to task, implying that the CMC architecture is adaptively leveraged to meet the specific requirements of each paradigm. Nonetheless, virtually all parameters have a posterior probability $p \approx 1.0$ of being

different than zero (with just two out of 84 parameters having smaller posterior probabilities of $p = 0.75$ and $p = 0.98$), suggesting that all the components and their functional connections remain necessary across all domains.

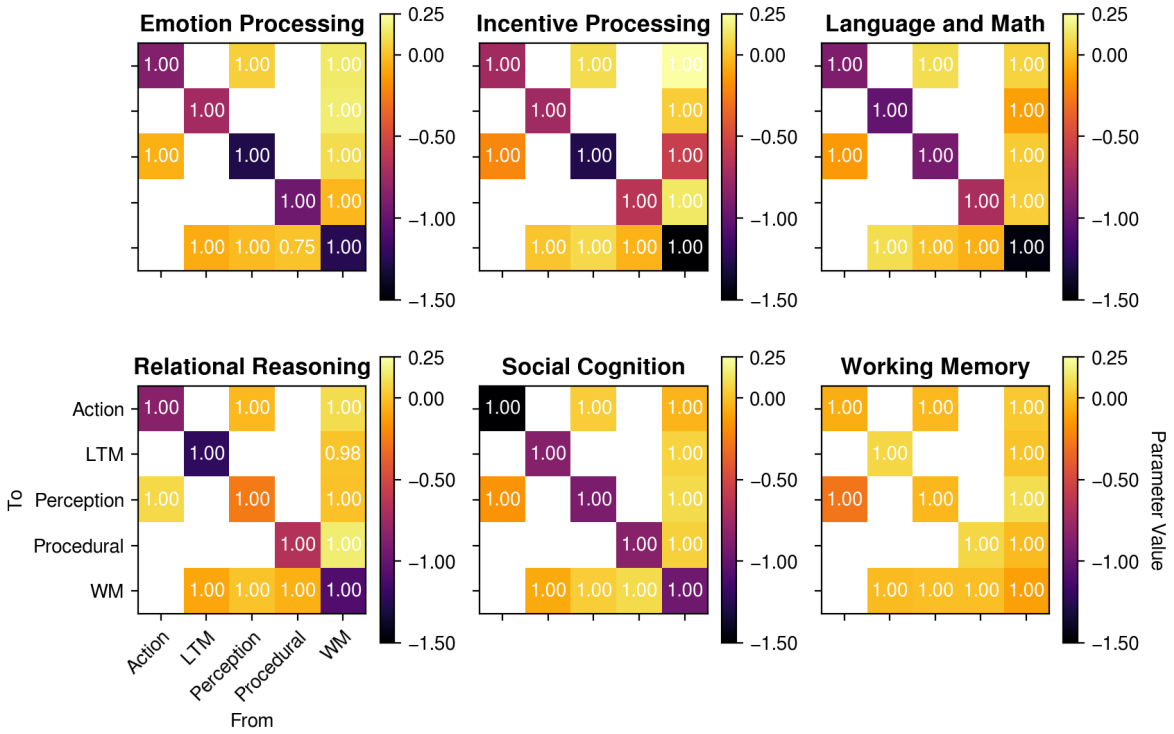

Figure 7. Estimated DCM intrinsic connectivity parameters for the CMC model. Each plot represents the intrinsic connectivity matrix (matrix A *in Eq. 1); the cell color indicates the parameter value, and the white text indicates the posterior probability that the parameter value is significantly different than zero. White matrix cells indicate connections that are not present in the CMC (see Figure 1A).*

Discussion

In this study, a comparative analysis was performed of the relative ability of five theoretical architectures to account for brain activity across seven different domains. These results provide overwhelming and converging evidence in favor of the Common Model of Cognition (CMC), a consensus architecture derived from the analysis of both human and artificial intelligent systems. Specifically, the CMC consistently outperforms all other architectures across all of the domains, thus jointly satisfying the two a-priori criteria of *generality* and *superiority*. Thus, the CMC emerges as a viable high-level blueprint of the human brain's architecture, potentially providing the missing unifying framework to relate brain structure and function for research and clinical purposes.

Although surprisingly robust and based on a large set of data, these results should be considered in light of three potential limitations. First, our conclusions are based on an analysis of task-related brain activity. Despite being established in the literature, the HCP tasks remain artificial, laboratory tasks, and their ecological validity is unknown. In contrast, many prominent studies have focused on task-free, resting-state paradigms. Thus, although the use of the task-related activity provides the most natural test for the generality criterion, the extent to which the CMC applies to resting-state fMRI remains to be explored.

Second, as noted above, our selection of alternative models was representative but not exhaustive. Although most distinct architectures that can be generated using just the five CMC components are likely to be unreasonable from a functional standpoint, some of them could, potentially, outperform the CMC.

Finally, it can be argued that our approach does not take full advantage of the possibilities of DCM, which makes it possible to accommodate non-linear, modulatory effects in the dynamic model. For example, the strategic role of procedural knowledge in the CMC (Assumption B3 of the original paper: Laird et al., 2017) is compatible with a "modulatory" view of the basal ganglia,

which has also been argued for from a theoretical standpoint (Stocco et al., 2010) and empirically observed in at least one study using effective connectivity analysis (Prat et al., 2016). In this study, modulatory connections were deliberately not included, so to level the field and make the five possible architectures more similar to each other in terms of overall complexity. Preliminary evidence, however, suggests that modulatory versions of the CMC might even outperform the standard version discussed herein (Steine-Hanson, Koh, & Stocco, A., 2019; Stocco et al., 2018).

These limitations notwithstanding, the fact that the CMC, which draws inspiration from high-level models of human cognition and *artificial* intelligent systems, also accounts for the neural activity of the human brain, which is a low-level *biological* intelligent system, is worthy of further consideration. A mundane explanation could simply be that humans reproduce the fundamental architecture of their intelligence when designing human-like models and intelligent systems. A more radical explanation is that the architectural space for general (or, at least, human-like) intelligence is inherently constrained and possibly independent of its physical realization, whether organic or artificial. In this sense, the CMC could be a model for *any* intelligent system, at any level of organization. Both hypotheses are worth exploring in future research.

Acknowldgements

This effort has been sponsored by award FA9550-19-1-0299 from the Air Force Office of Scientific Research (AFOSR) to AS, by award FA9550-19-0180 from AFOSR to JL, by award W911NF-14-D-0005 from the U. S. Army to PR, and by an award from the Defense Advanced Research Projects Agency (DARPA) to CL. Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred. Author contributions: Conceptualization: AS, CL, JL, and PR; Data curation; AS, NK, and ZSH; Formal analysis: AS, NK, and ZSH; Funding acquisition AS, CL, JL, and PR; Investigation: AS, CL, JL, and PR; Methodology: AS; Project administration: AS; Resources AS, CL, JL, and PR; Software: AS, NK and ZSH; Visualization: AS; Writing: AS, CL, JL, NK, PR

and ZSH.

All the raw imaging data is available through the Human Connectome Project (http://www.humanconnectome.org). All the pipeline scripts designed to preprocess the data, generate and fit the DCM models, extract and interpret the results, and generate the figures reported in this paper are available as an open-source code repository on GitHub (https://github.com/UWCCDL/CMC-DCM).

The authors would like to thank Jim Treyens and John R. Anderson for their comments on early drafts of this manuscript.

References

Akaike, H., 1974. A new look at the statistical model identification, in: Selected Papers of Hirotugu Akaike. Springer, pp. 215–222.

Anderson, J.R., 2007. How Can the Mind Occur in the Physical Universe? Oxford University Press.

Ashburner, J., Barnes, G., Chen, C., Daunizeau, J., Flandin, G., Friston, K., Kiebel, S., Kilner, J., Litvak, V., Moran, R., Others, 2016. SPM12 manual. URL: http://www.fil.ion.ucl.ac.uk/spm/doc/spm12 manual.pdf.

Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., WU-Minn HCP Consortium, 2013. Function in the human connectome: task-fMRI and individual differences in behavior. Neuroimage 80, 169–189.

Binder, J.R., Gross, W.L., Allendorfer, J.B., Bonilha, L., Chapin, J., Edwards, J.C., Grabowski, T.J., Langfitt, J.T., Loring, D.W., Lowe, M.J., Koenig, K., Morgan, P.S., Ojemann, J.G., Rorden, C., Szaflarski, J.P., Tivarus, M.E., Weaver, K.E., 2011. Mapping anterior temporal lobe language areas with fMRI: a multicenter normative study. Neuroimage 54, 1465–1475.

Borst, J.P., Anderson, J.R., 2013. Using model-based functional MRI to locate working memory updates and declarative memory retrievals in the fronto-parietal network. Proc. Natl. Acad. Sci. U. S. A. 110, 1628–1633.

Borst, J.P., Nijboer, M., Taatgen, N.A., van Rijn, H., Anderson, J.R., 2015. Using data-driven model-brain mappings to constrain formal models of cognition. PLoS One 10, e0119673.

Buckner, R.L., Krienen, F.M., Castellanos, A., Diaz, J.C., Yeo, B.T.T., 2011. The organization of the human cerebellum estimated by intrinsic functional connectivity. J. Neurophysiol. 106, 2322–2345.

Buxton, R.B., Wong, E.C., Frank, L.R., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn. Reson. Med. 39, 855–864.

Castelli, F., Happé, F., Frith, U., Frith, C., 2000. Movement and mind: a functional imaging study of perception and interpretation of complex intentional movement patterns. Neuroimage 12, 314–325.

Cole, M.W., Reynolds, J.R., Power, J.D., Repovs, G., Anticevic, A., Braver, T.S., 2013. Multi-task connectivity reveals flexible hubs for adaptive task control. Nat. Neurosci. 16, 1348–1355.

Cole, M.W., Yarkoni, T., Repovs, G., Anticevic, A., Braver, T.S., 2012. Global connectivity of prefrontal cortex predicts cognitive control and intelligence. J. Neurosci. 32, 8988–8999.

Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., Fiez, J.A., 2000. Tracking the hemodynamic responses to reward and punishment in the striatum. J. Neurophysiol. 84, 3072–3077.

Eliasmith, C., Stewart, T.C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., Tang, C., Rasmussen, D., 2012. A large-scale model of the functioning brain. Science 338, 1202–1205.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19, 1273–1302.

Friston, K.J., Mechelli, A., Turner, R., Price, C.J., 2000. Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. Neuroimage 12, 466–477.

Gorgolewski, K.J., Lurie, D., Urchs, S., Kipping, J.A., Craddock, R.C., Milham, M.P., Margulies, D.S., Smallwood, J., 2014. A correspondence between individual differences in the brain's intrinsic functional architecture and the content and form of self-generated thoughts. PLoS One 9, e97176.

Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., Colmenarejo, S.G., Grefenstette, E., Ramalho, T., Agapiou, J., Badia, A.P., Hermann, K.M., Zwols, Y., Ostrovski, G., Cain, A., King, H., Summerfield, C., Blunsom, P., Kavukcuoglu, K., Hassabis, D., 2016. Hybrid computing using a neural network with dynamic external memory. Nature 538, 471–476.

Hariri, A.R., Tessitore, A., Mattay, V.S., Fera, F., Weinberger, D.R., 2002. The amygdala response to emotional stimuli: a comparison of faces and scenes. Neuroimage 17, 317–323.

Hassabis, D., Kumaran, D., Summerfield, C., Botvinick, M., 2017. Neuroscience-Inspired Artificial Intelligence. Neuron 95, 245–258.

Hazy, T.E., Frank, M.J., O'reilly, R.C., 2007. Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. Philos. Trans. R. Soc. Lond. B Biol. Sci. 362, 1601–1613.

Huntenburg, J.M., Bazin, P.-L., Margulies, D.S., 2018. Large-Scale Gradients in Human Cortical Organization. Trends Cogn. Sci. 22, 21–31.

Jonas, E., Kording, K.P., 2017. Could a Neuroscientist Understand a Microprocessor? PLoS Comput. Biol. 13, e1005268.

Kasess, C.H., Stephan, K.E., Weissenbacher, A., Pezawas, L., Moser, E., Windischberger, C., 2010. Multi-subject analyses with dynamic causal modeling. Neuroimage 49, 3065–3074.

Kieras, D.E., Meyer, D.E., 1997. An overview of the EPIC architecture for cognition and

performance with application to human-computer interaction. Human-computer interaction 12, 391–438.

Kotseruba, I., Tsotsos, J.K., 2018. 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review.

Laird, J.E., 2012. The Soar Cognitive Architecture. MIT Press.

Laird, J.E., Lebiere, C., Rosenbloom, P.S., 2017. A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. AI Magazine 38.

Mohan, S., submitted, Exploring the Role of Common Model of Cognition in Designing Adaptive Coaching Interactions for Health Behavior Change.

O'Reilly, R.C., Hazy, T.E., Herd, S.A., 2016. The Leabra Cognitive Architecture: How to Play 20 Principles with Nature. The Oxford Handbook of Cognitive Science 91.

Penny, W.D., Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., 2011. Statistical Parametric Mapping: The Analysis of Functional Brain Images. Academic Press.

Prat, C.S., Stocco, A., Neuhaus, E., Kleinhans, N.M., 2016. Basal ganglia impairments in autism spectrum disorder are related to abnormal signal gating to prefrontal cortex. Neuropsychologia 91, 268–281.

Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D., 2016. Mastering the game of Go with deep neural networks and tree search. Nature 529, 484–489.

Smith, R., Keramatian, K., Christoff, K., 2007. Localizing the rostrolateral prefrontal cortex at the individual level. Neuroimage 36, 1387–1396.

Steine-Hanson, Z. K., Koh, N., & Stocco, A., 2019. Refining the Common Model of Cognition

through large neuroscience data, in: Alexei V. Samsonovich and Christian J. Lebiere (Eds.). (Ed.), Procedia Computer Science. Presented at the Postproceedings of 9th Annual International Conference on Biologically Inspired Cognitive Architectures, Elsevier.

Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009. Bayesian model selection for group studies. Neuroimage 46, 1004–1017.

Stephan, K.E., Penny, W.D., Moran, R.J., den Ouden, H.E.M., Daunizeau, J., Friston, K.J., 2010. Ten simple rules for dynamic causal modeling. Neuroimage 49, 3099–3109.

Stocco, A., Laird, J., Lebiere, C., Rosenbloom, P., 2018. Empirical Evidence from Neuroimaging Data for a Standard Model of the Mind, in: C. Kalish, M. Rau, J. Zhou, and T. T. Rogers (Ed.), Proceedings of the 40th Annual Meeting of the Cognitive Science Society. pp. 1094–1099.

Stocco, A., Lebiere, C., Anderson, J.R., 2010. Conditional routing of information to the cortex: A model of the basal ganglia's role in cognitive coordination. Psychol. Rev.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, 2013. The WU-Minn Human Connectome Project: an overview. Neuroimage 80, 62–79.

Wheatley, T., Milleville, S.C., Martin, A., 2007. Understanding animate agents: distinct roles for the social network and mirror system. Psychol. Sci. 18, 469–474.