AFRL-AFOSR-VA-TR-2023-0370

Testing a Common Model for Human and Human-Like Intelligence

Stocco, Andrea
UNIVERSITY OF WASHINGTON
4333 BROOKLYN AVE NE
SEATTLE, WA, 98195
USA

06/15/2023
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20230615	Final	START DATE 20190701	END DATE 20220930

4. TITLE AND SUBTITLE
Testing a Common Model for Human and Human-Like Intelligence

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA9550-19-1-0299	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Andrea Stocco

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
UNIVERSITY OF WASHINGTON
4333 BROOKLYN AVE NE
SEATTLE, WA 98195
USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 N. Randolph St. Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR RTA2

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-VA-TR-2023-0370

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The goal of this proposal was to test the viability of the Common Model of Cognition (CMC: Figure 1) as possible systems-level brain architecture. The CMC was initially proposed by John Laird, Christian Lebiere, and Paul Rosenbloom (who served as consultants for this project) as a synthesis of decades of progress in the field of cognitively inspired AI. As such, it provides a natural computational framework to explain the human mind but remains silent as to the nature of cognition's neural bases – the human brain. Thus, we set out to examine the degree to which the CMC could provide a reasonable account of the brain's intrinsic 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 and generality.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU	20

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
HAL GREENWALD	

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

Final Report:
"Testing a Common Model for Human and Human-Like Intelligence"

AFOSR Award FA9550-19-1-0299

PI: Dr. Andrea Stocco, University of Washington, Seattle

Background and Rationale

The goal of this proposal was to test the viability of the Common Model of Cognition (CMC: Figure 1) as possible systems-level brain architecture. The CMC was initially proposed by John Laird, Christian Lebiere, and Paul Rosenbloom (who served as consultants for this project) as a synthesis of decades of progress in the field of cognitively inspired AI. As such, it provides a natural computational framework to explain the human mind but remains silent as to the nature of cognition's neural bases – the human brain. Thus, we set out to examine the degree to which the CMC could provide a reasonable account of the brain's intrinsic 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 and generality.

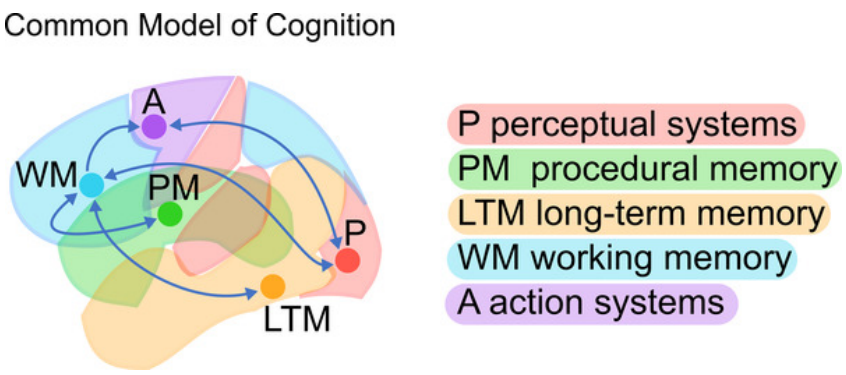

Figure 1: Components and network connectivity of the Common Model of Cognition (CMC). Colored areas represent putative associated brain regions

Goal 1: Viability of the Architecture

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 elements. Depending on the methods used, the number of anatomically identifiable areas in the human brain counts in the hundreds (Glasser et al., 2016; Power et al., 2011; Yeo et al., 2011), and thus does not provide a reliable starting point. The number of functionally distinct circuits, however, is recognized as being at least one order of magnitude smaller, as different brain areas form interconnected networks (Cole et al., 2016; Power et al., 2011; Yeo et al., 2011). This study takes, as a reference point, the influential estimate given in Yeo et al. (2011), which counts seven distinct functional networks—a number that is comparable to the number of components in the CMC.

An initial identification can be made between CMC components and some of these networks. 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 can be identified with the fronto-parietal network comprising the dorsolateral prefrontal cortex (PFC) and posterior parietal cortex. The long-term memory (LTM) component corresponds with regions involved in the encoding of episodic memories, such as the hippocampus and the surrounding medial temporal lobe regions (Moscovitch et al., 2005; Ranganath and Maureen, 2012; Squire, 2004), as well as with regions involved in memory retrieval, such as the medial frontal cortex and the precuneus; these regions are referred to as the default mode network (Raichle and Snyder, 2007; Ranganath and Maureen, 2012). The action components can be identified with the sensorimotor network (Power et al., 2011); the procedural knowledge component with the basal ganglia (Yin and Knowlton, 2006); and the perception modules with the dorsal and ventral visual networks, as well as, depending on the task, the auditory networks (Fig. 1B).

To test the CMC architecture, we used Dynamic Causal Modeling (DCM), 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:

$$d\mathbf{y}/dt = \mathbf{A}\mathbf{y} + \mathbf{C}\mathbf{x} \quad (1)$$

where **x** represents the event vectors (i.e., the equivalent of a design matrix in traditional GLM analysis), **A** defines intrinsic connectivity between ROIs and **C** defines the ROI-specific effects of task events. 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 model parameters (connectivity matrics **A** and **C** and the hemodynamic parameters) were then estimated using the expectation-maximization algorithm.

Accounting for Task-Based Brain Activity

The connectivity patterns were then tested against a large repository of human brain imaging data (the Human Connectome Project, containing data from 1,200 individuals performing seven different tasks) using Dynamic Causal Modeling.

The HCP task-fMRI data encompasses seven different paradigms designed to capture a wide range of cognitive capabilities. Of these paradigms, six (Emotion Recognition, Incentive Processing, Language and Math, Relational Reasoning Social Cognition, and Working Memory) 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).

To properly characterize each individual component, a processing pipeline was designed to progressively identify a relevant corresponding region of interest (ROI) for each task and, within each

task, for each of the ~200 participants, thus accounting for individual differences in functional neuroanatomy. Thiso provided task- and individual-specific ROIs, accounting for differences in functional neuroanatomy. The location of the ROIs corresponding to each component for each participant are shown in Figure 2.

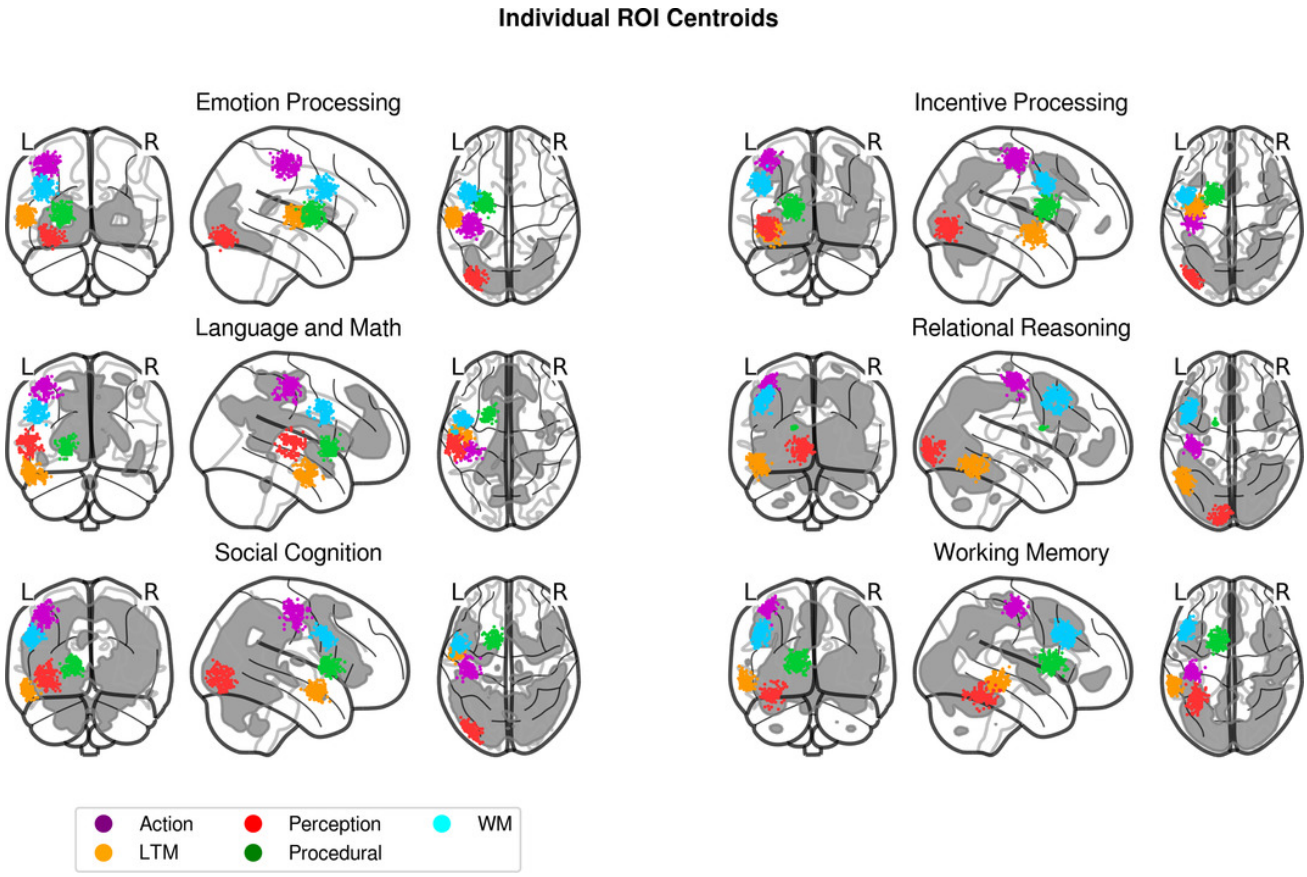

Figure 2: Location of the centroids of each individual-specific ROI for each HCP task. Grey areas represent the extent of significant brain activity elicited by task conditions (thresholded at $T > 5.00$, FWE $p < 0.01$).

To address our second criterion of comparative superiority, the CMC connectivity model was compared against other connectivity models that implement alternative brain architectures. Because the space of possible models is large, we concentrated on six examples that are representative of theoretical neural architectures previously suggested in the neuroscientific literature (Fig. 2). These six alternatives 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 conceptualized as a feedforward neural network model with large-scale gradients of abstraction. Three possible hierarchical architectures were created, in which the Procedural region falls between Perception and LTM (Hierarchical 1), or between LTM and WM (Hierarchical 2), or between WM and Action (Hierarchical 3).

In the “Hub-and-Spoke” family, a single ROI is singled out as the network's “Hub” and receives bidirectional connections from all the other ROIs (the “Spokes”). With the exception of the Hub, no ROI is mutually connected to any other one. Three different Hub-and-Spoke architectures were created by selecting as the Hub one of the regions, with the exception of Perception and Action.

Once the time-series for each ROI was extracted, different networks were created by connecting all of the individually-defined ROIs according to the specifications of each architecture . 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.

The predicted time course of the neural response for each network model was then generated by using Eq. (1) to simulate network activity as it unfolded over the course of the task. The 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 of each node in the network. The parameters of the full DCM, which include both the network connectivity parameters and the physiological parameters of the neurovascular coupling 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. As an example, Figure 3 illustrates the predicted and observed time courses of brain activity in one individual performing one of the tasks (the Relational Reasoning tasks) of the HCP.

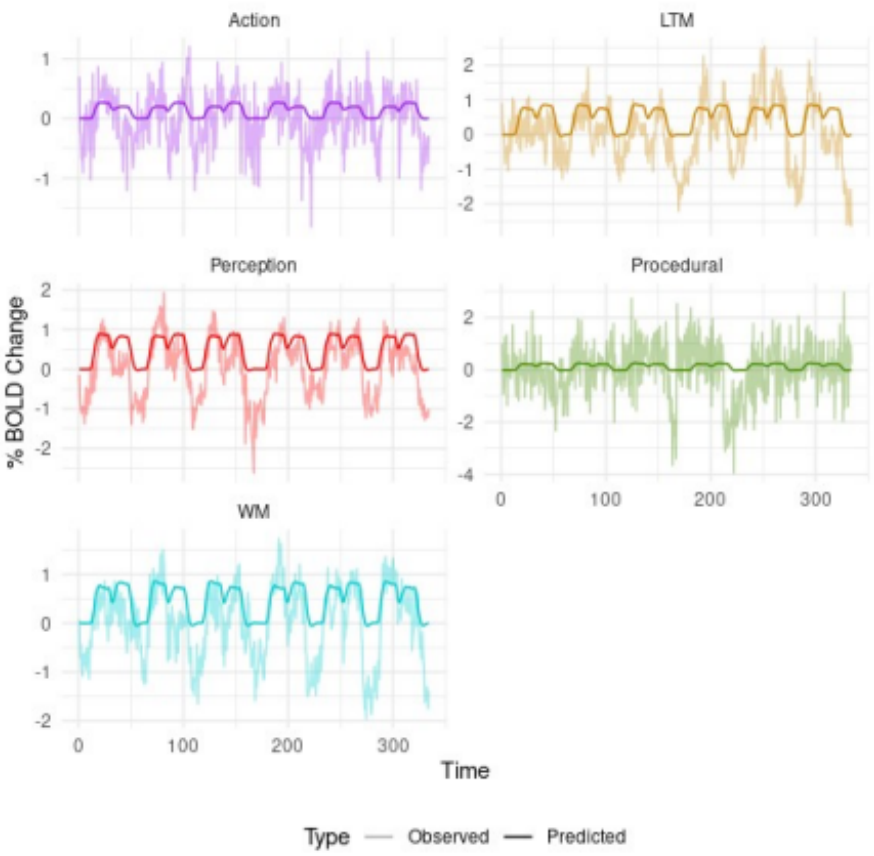

Figure 3: An example of predicted and observed BOLD responses in one individual performing the HCP Relational Reasoning task.

Once the seven DCM models were separately fitted to the functional neuroimaging data, they were compared against each other using a Bayesian random-effects procedure. 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), 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).

The posterior distributions of each architecture are visualized for each task in Fig. 4. The expected probability densities are represented as the colored bel-shaped distributions, while the exceedance probabilities are summarized as the corresponding labels.

Both metrics provide strong evidence in favor of the CMC. As shown in Fig. 4, the CMC provides a better fit to the data than any alternative architecture, and its exceedance probabilities range from 0.99 to 1.0. 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).

One other possibility is that the superiority of the CMC originates from some peculiarity of its network connectivity that was missing in the other architectures. The one notable difference, in this sense, is the presence of a direct link between the Perception and Action ROIs, which are bilaterally connected in the CMC but unconnected in the six other rival architectures. To examine the role of a direct perception-action link in fitting the data, the six alternative architectures were augmented with bilateral Perception-Action connectivity and a new Bayesian model comparison was run. Overall, the addition of the direct Perception-Action connectivity improved the fit of the alternative models, but the CMC maintained its superiority across tasks.

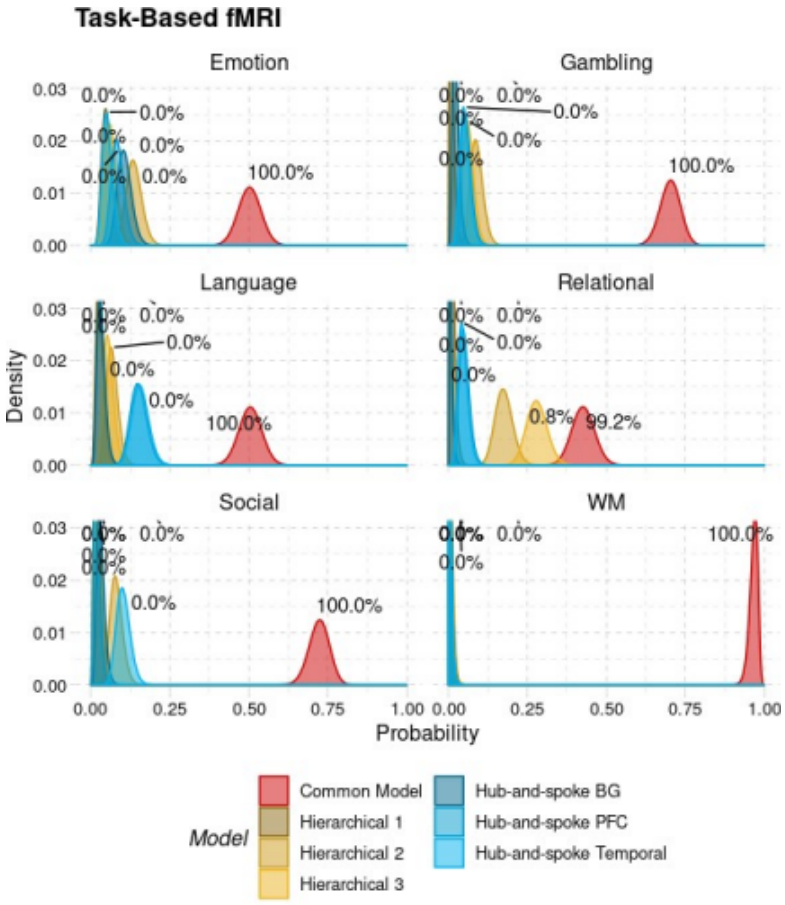

Figure 4: Posterior probabilities that each architecture fits an individual from the HCP dataset. Labels indicate the exceedance probabilities (i.e., the probability that a given architecture fits better than all of the others).

Accounting for Resting State Brain Activity

A challenge for all of the computational architectural network models is accounting for spontaneous, or *resting-state brain activity*, which occurs outside the constraints of task events that invoke activity in specific brain regions.

The spontaneous activity observed in the brain at rest exhibits many remarkable characteristics. It is oscillatory in nature (Fox et al., 2005) and concentrated in the low frequency band (<0.1 Hz). Because it is widespread and persistent, low frequency oscillations (LFOs) are believed to account for most of a neuron's energy expenditure, much more so than the localized bursts of activity that occur during task performance (Pezzulo et al., 2021). This burdensome metabolic cost suggests that LFOs must play a functionally important, if yet unknown, role in cognition. In addition, spontaneous activity at rest appears to be spatially and temporally organized. Specifically, LFOs are organized into different networks of regions, with regions within a network being more correlated than regions across networks (Fox et al., 2005; Power et al., 2011), and with rapid transitions that resemble task-related activity (Kang et al., 2019). Thus, in agreement with the goals of our proposal, the original work was extended from task-based fMRI data to task-free (i.e, resting-state) fMRI data.

Because ROIs could not be identified with task-based activation criteria, we created a new set of ROI based on Neurosynth-based meta-analysis, using each component's functional name (i.e., "Working Memory") as searches. The meta-analytic maps were equalized on the bases of the reported effect sizes and absolute volumes and then optimized to eliminate any overlap. The resulting ROIs

provide an a priori volume reference for the CMC component that can be used for any type of effective connectivity analysis, whether task-based or resting-state based. (Figure 5).

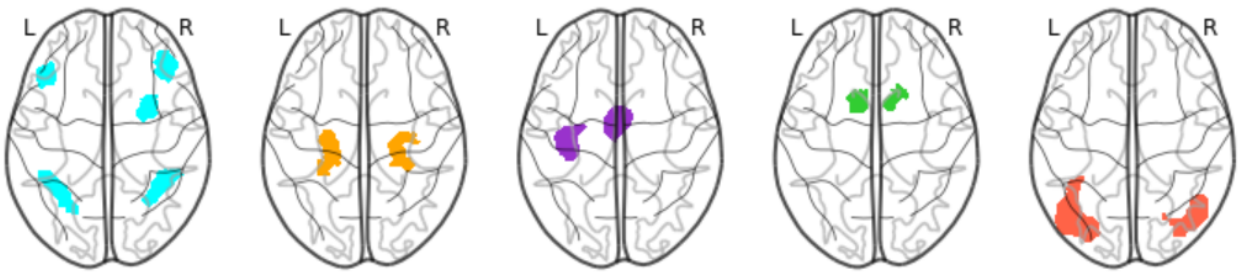

WM	LTM	Action	PM	Perception
1,608 vx Prefrontal Cortex Inferior Parietal Lobule Supramarginal Gyrus Middle Frontal Gyrus	758 vx Parahippocampal gyrus Hippocampus	1,319 vx Medial frontal gyrus Primary motor cortex Premotor Cortex Supp Motor Area Ant Cingulate Gyrus Sup Frontal Gyrus	440 vx Dorsal Striatum Lentiform Nucleus Pallidum	2,088 vx Middle Occipital Gyrus Middle Temporal Gyrus Fusiform Gyrus

Figure 5: Optimized meta-analytic ROIs.

Lacking any stimulus-based activity, local changes in each ROI timeseries were modeled as periodic square waves with different frequencies and phases, with each wave affecting all of the components (Figure 6).

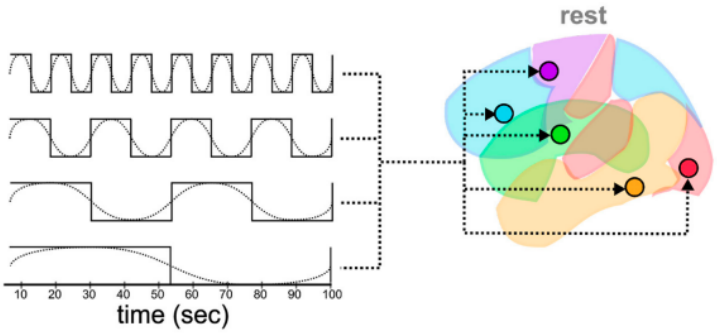

Figure 6: Square wave regressors model resting-state regional activity

After isolating spontaneous regional activity from the CMC network, we were able to clearly observe the patterns of regional connectivity within the CMC components. These patterns were, once more, best explained by the CMC architecture rather than other, comparable, architectures (Figure 7).

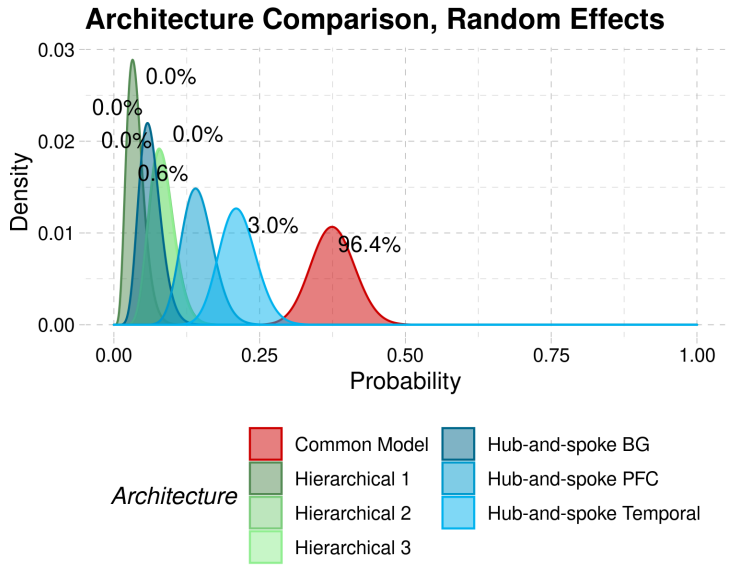

Figure 7: The CMC outperforms other architectures in explainaing resting-state connectivity data.

Importantly, the resulting pattern of connectivity data, once local activity was accounted for, was remarkably similar to that of task-based data (Figure 8):

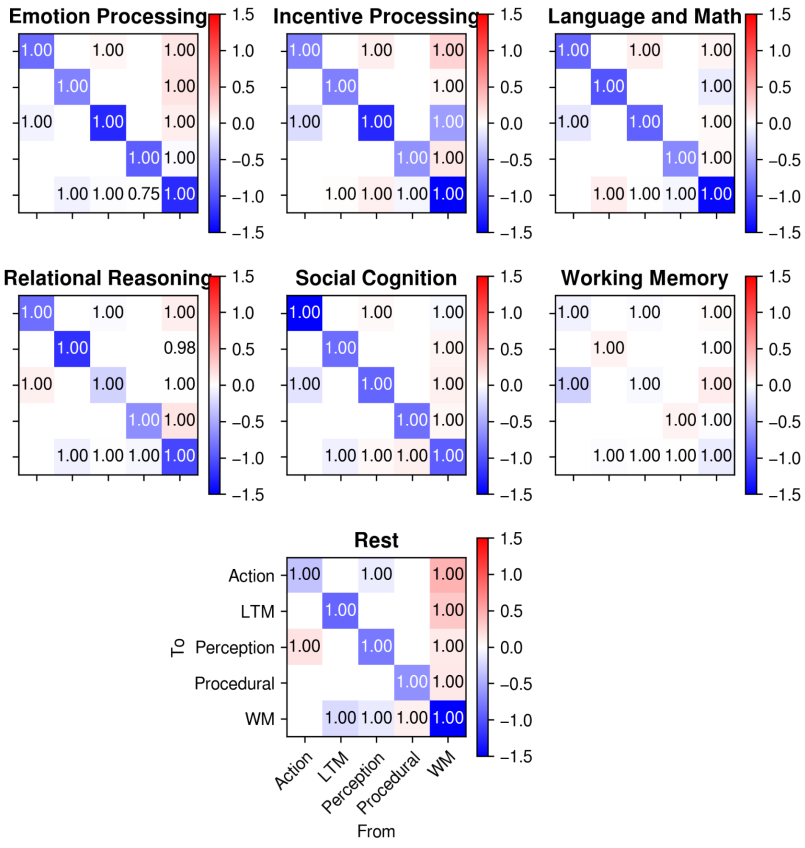

Figure 8: A comparison of intrinsic connectivity between regions (matrix A in Equation 1) during task-based activity (top) and during rest (bottom). Cell colors represent connection weights; numbers in the cells represent posterior probabilities that the weights are different from zero.

Goal 2: Convergent and External Validity

The second goal of this proposal was to verify the external validity of our conclusions by testing whether the same results could be obtained using a different methodology. To this aim, we applied Granger Causality Modeling (GCM) to the same task-based fMRI data and used it to infer a cognitive architecture from the data.

Convergent Validity: Granger Causality Modeling

In GCM, the existence of a causal effect between two time series x and y is established by comparing two models (Granger, 1969), one auto-regressive linear model in which the value of y at times t depends only on its past value at time t-1:

$$y(t) = \beta_0 + \beta_1 y(t\text{-}1)$$

and an alternative model that includes the effect of the past state of x:

$$y(t) = \beta_0 + \beta_1 y(t\text{-}1) + \beta_2 x(t\text{-}1)$$

If the second model is significantly better than the first, then x is said to *Granger-cause* y. The extent of the temporal window in the past (i.e., time t – 1) that is allowed to affect the present (at time t) is called the *temporal lag*. To analyze functional connectivity with Granger analysis, a multivariate Granger causality model was set up in which the BOLD response at time t across all regions, represented as the vector $\mathbf{x}(t)$, was modeled as the contribution of all of the regions (including itself) at lags 1, 2, … k:

$$\mathbf{x}(t) = \beta_0 + \beta_1 \mathbf{x}(t\text{-}1) + \ldots + \beta_k \mathbf{x}(t\text{-}k)$$

To determine the optimal lag value, 10 models were created by varying k from $k = 1$ to $k = 10$, and the value of k that gave rise to the model with the lowest Bayesian information criterion was selected. Across all participants and tasks, the maximum lag that was observed was $k = 6$, and the modal was $k = 2$. Note that, when $k > 1$, there are multiple different parameter estimates that quantify the directional effect of a region on another region, one for each lag. To reduce the dimensionality of these estimates, only the most significant lag (i.e., the one with the smallest p-value) was selected.

For each participant, a subject-level inferred architecture was then created by discretizing the matrix of connections and maintaining only directed links with $p < .05$. To infer a group-level architecture from the individual-level architectures, the most likely directed links between regions need to be inferred from the frequency of their distribution in the sample of participants. To determine the probability that each directed connection c is part of the group-level architecture, we modeled the probability of it appearing across all participants as a binomial distribution, with a prior probability of $p = .50$. This prior reflects the agnostic assumption that a functional connection between two regions might or might not exist with equal probability. To derive a *general* architecture from these six task-specific ones, we considered each task as an independent experiment to test these connections and used Fisher's (1932) method to combine the p-values from each task. According to this method,

the distribution of the log of p-values from independent tests follows a $\chi 2$ distribution with $2N$ degrees of freedom, and the p-value of each connection can be calculated from the $\chi 2$ cumulative distribution function as follows:

$$P_{global} = p(\chi^2_{2N} > \sum_{task} \log P_{task})$$

A connectivity matrix of a probable architecture inferred across participants and task domains can be seen in Fig. 9. Once discretized at a threshold of ($p > .95$), the inferred GCM architecture contains 17 of the total 25 possible connections, and of the 17 connections found, 14 of them overlap with the 14 connections of the CMC. Thus, the GCM-derived model can be seen not as an alternative, but as a sub-specification of the CMC — or, in statistical terms, as model that is *nested* within the CMC architecture.

The three additional connections suggested by the GCM architecture include: (1) A backward projection from Action to WM, (2) a projection from Action to LTM, and (3) a projection from Perception to LTM (Fig. 9). Together, these new connections could be considered reflective of an architecture that is similar to the CMC but with hubs of network convergence in both the WM and LTM regions. Of these connections, the backward link between Action and WM was already tested by Stocco et al. (2021) and was found to significantly improve the fit of the CMC to neuroimaging data across all tasks. The other two connections, however, had never been considered before.

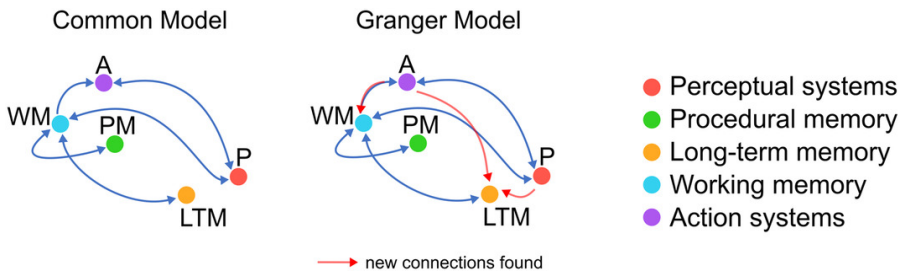

Figure 9: A visual representation of the architecture inferred from GCM. *Blue*: connections present in both; *Red*: connections unique to the Granger Model.

Re-Evaluation of DCM Results

To evaluate the inferred, data-driven Granger model, a model comparison similar to that of Goal 1 was performed. In a previous DCM-based analysis of architecture structures, we had not been able to incorporate data-driven inferences about connections. Instead, the plausibility of the CMC was evaluated by comparing its predictions against a set of representative alternative architectures across tasks. These architectures were divided into two categories, or "families," and represented the possible organizational structures of general purpose architecture: Hierarchical and Hub-and-Spoke.

Thus, the data-driven GCM architecture was evaluated by comparing its DCM predictions against both the CMC and the six alternate architectures tested in previous analyses. Fig. 10 visualizes the fit of each of the eight models within each task and across all tasks combined. The fit of each model was analyzed using the posterior distribution of the probability that architecture would fit a random individual (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). The model posterior distributions for each task, with vertical shaded lines representing the expected probabilities, or the means of each distribution. As shown in Fig. 10, the GCM provides a better fit than any alternative

architecture, including the CMC, across almost all cognitive domains with the exception being Language and Math. These data show that the new GCM can meet a high level of generality and comparative superiority. Note that, compared to the original results (Figure 4), the performance of the original CMC architectureis greatly reduced, which suggests that GCM's superiority is mostly due to the GCM outcompeting the CMC for the same share of participants.

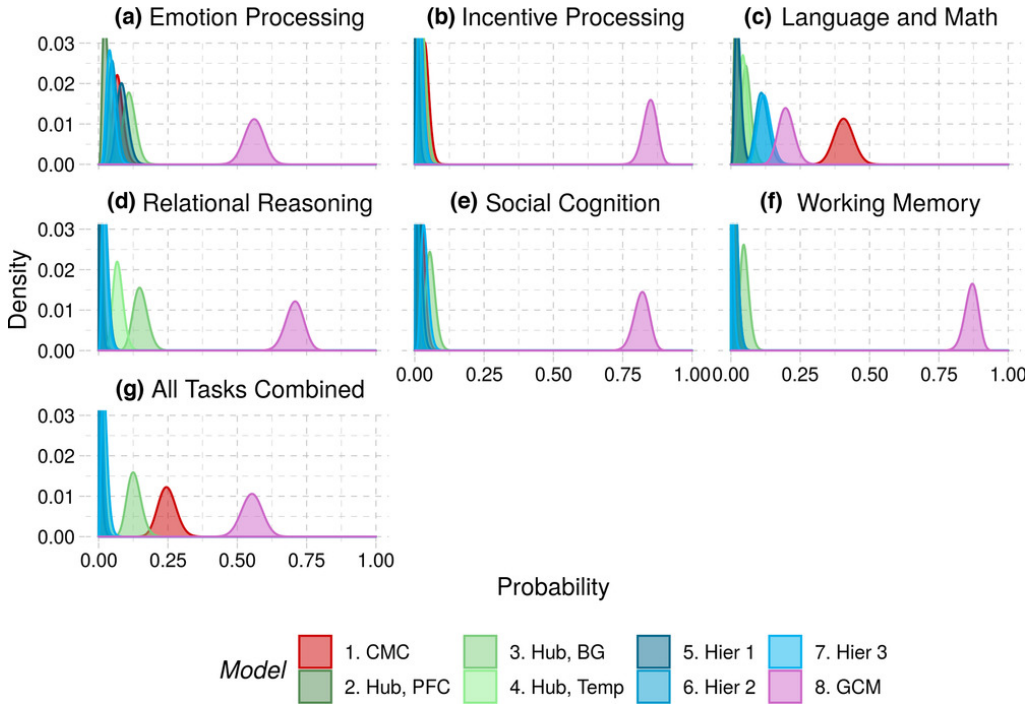

Figure 10: Results of the Bayesian model comparisons. (A–G) Probability distributions that each of the eight architectures is true, given the data within each task and across all tasks combined.

External Validity: Ground-Truth Testing With Neurodegenerative Patient Data

While the GCM analysis provided *convergent* validity, it did not provide *external* validity, that is, independent confirmation of a model when applied to different data and domains. We reasons that, if the CMC represents the fundamental architecture of the brain, then damage to any of its components should produce large-scale, network-level changes in brain activity. Conversely, it should be possible to link large-scale alterations in functional connectivity to specific damage to on component.

Parkinson's Disease (PD) privides an ideal case to test this hypothesis. Unlike other neurodegenerative diseases (such as Alzheimer's Disease), which interest multiple brain regions concurrently, PD selectively affects the basal ganglia, which correspond to the CMC's "Procedural Memory" component. Furthermore, despite resulting in significant cognitive and behavioral effects, changes in functional connectivity are unexpectedly difficult to find in PD, with many studies failing to find any significant differences between PD and age-match healthy controls.

Before applying the CMC to PD patients, it is necessary to further specify the role of the Procedural Memory component. In previous analysis, this component was simply seen as a node interacting exclusively with the Working Memory component (i.e., prefrontal cortex). This was,

however, a necessary simplification; the original framing of the CMC explicitly states that procedural memory is also responsible for coordinating the access of other components to Working Memory. This is also in line with the putative function of the basal ganglia, as suggested by existing neuroimaging findings and neurocomputational models.

Thus, we first examined three sub-variants of the CMC: one of which the Procedural Memory simply affects Working Memory; one in which Procedural Memory *modulates* the access of other components (Perception and Long-Term memory) to Working Memory; and a third one in which Procedural Memory has both direct and modulatory effects to Working Memory (Figure 11):

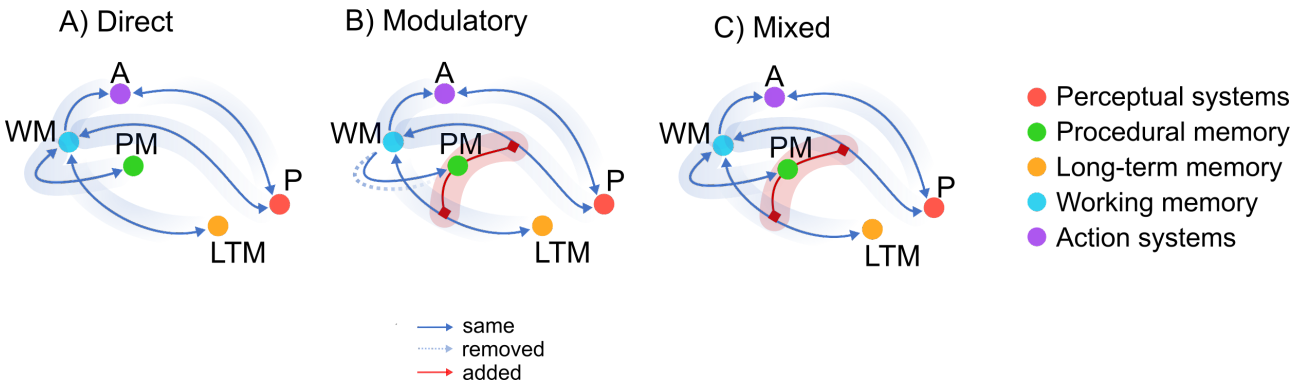

Figure 11: The structure of the three CMC versions implementing different functions of the basal ganglia (Procedural Memory, PM) on cognition.

The three variants of the CMC were compared using DCM. Because two of the models include second-order, modulatory effects, we used an expanded version Equation 1:

$$d\mathbf{y}/dt = \mathbf{A}\mathbf{y} + \mathbf{C}\mathbf{x} + \sum_i y_i \mathbf{D}^{(i)}\mathbf{y} \qquad (2)$$

where y_i is the activity in the i-th region, and $\mathbf{D}$ is $n \times n \times n$ matrix of the modulatory effects that every region has on the connections between every other region. The three architectures were then fit to the same Human Connectome data and tasks of Goal 1.

Because the three models are instantiations of the same architecture, we performed a fixed-effects, rather than mixed-effects, model analysis — that is, we computed the group-level log-likelihoods of each architecture explaining the participants'. As Figure 12 shows, the mixed model outperformed the other two models across all tasks. To ensure that the results in Figure 12 were not due to the fact that the mixed model has more parameters, we ran separate comparisons (mixed vs. direct and mixed vs. modulatory), using Wilks' formula to derive a p-value for the difference in log-likelihood when accounting for the difference in parameters. The results confirmed that, even when accounting for the additional parameters, the mixed model provided a significantly superior ($p < 0.0001$) fit to the data than the other two models across all of the six tasks.

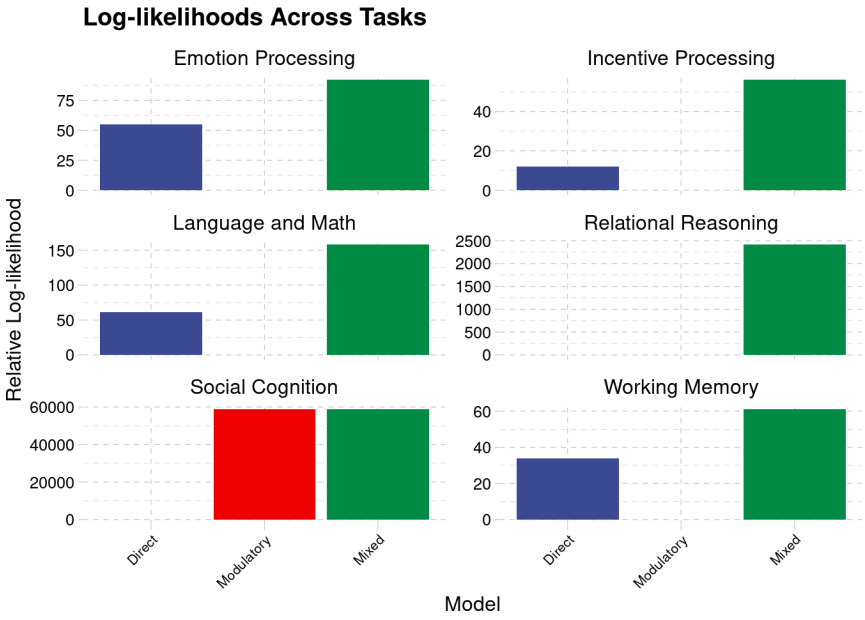

Figure 12: Relative log-likelihoods of the three CMC variants (Figure 11) across all of the HCP tasks.

Having established that the mixed CMC architecture provides a better way to capture the effect of the basal ganglia on the cognitive architecture, we set forward to test whether it could detect abnormal connectivity in PD.

To this aim, we took advantage of a collaboration with the Integrative Brain Imaging Center at the University of Washington, and fit the mixed model to a sample of N=70 PD patients and N=41 healthy controls, matched for age, gender, education, and cognitive abilities. The model was fit to resting-state data. The results confirmed the existence of abnormal connectivity patterns in PD (Figure 13). Specifically, the two modulatory connections (one affecting the connectivity between the temporal lobe and the profrontal cortex, and one affecting the connectivity between the visual cortex and the prefrontal cortex) were found to be significantly dampened in PD when compared to controls. The reduced modulatory connectivity is likely caused by the loss on dopaminergic neurons in teh basal ganglia circuit.

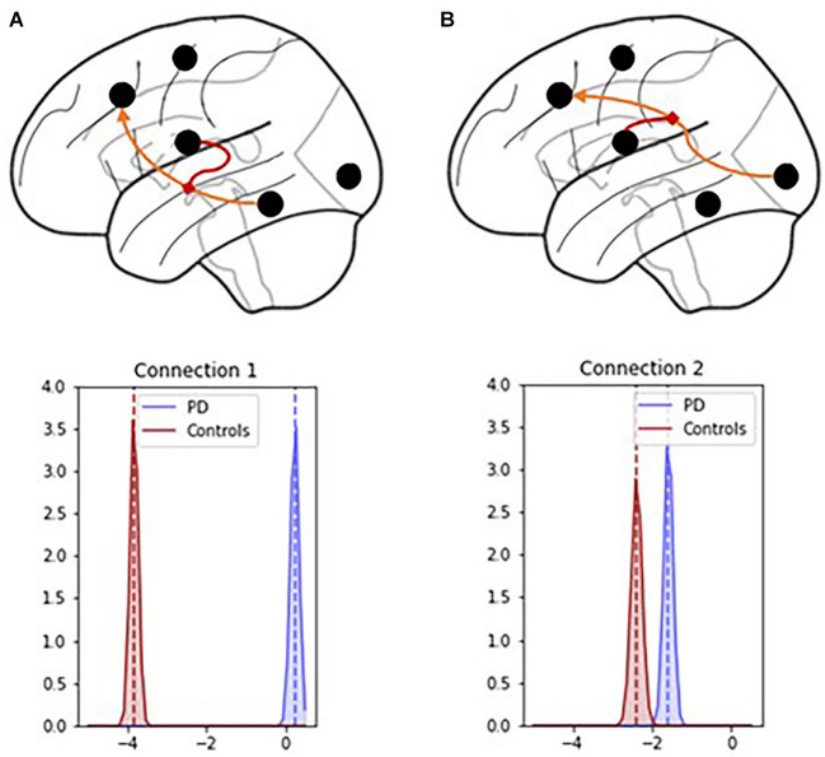

Figure 13: The two modulatory connections of the mixed CMC showed abnormally lower intensity in PD patients than in controls.

Goal 3: Individual Differences

The third goal of the proposal was to test whether the CMC could be used to capture individual differences in cognitive function as well as the invariant cognitive and brain architecture of different individuals. A successful neuro-cognitive architecture should be able to capture not only the invariant parts of brain structure and function across individuals but also relate variations in such structures to observable variations in behavior.

Correlations Between CMC Connectivity and Individual Abilities

To test this hypothesis, we took advantage of a collavoration with Chantel Prat of the University of Washington's Department of Psychology, who has been collecting psychometric and neuroimaging data predicting individual differences in learning how to program in Python. This is a real-life, complex case scenario that has been shown to be reliably associated with other outcomes (such as intelligence and scholastic achievement).

To investigate whether CMC connectivity would correlate with programming skill, we applied the DCM implementation of the CMC (Equation 1) to the resting-state fMRI data of N=37 undergraduate students involved in a Python programming course. The CMC model was fit to the timeseries data of the five meta-analytic regions (Figure 5) using the square-wave periodic regressors (Figure 6). Our results show that CMC connectivity does indeed correlate with several programming outcomes, with connectivity between perceptual and motor regiorns and between procedural and prefrontal regions being the strongest predictors across four relevant outcomes (learning rate, quiz accuracy, programming accuracy, and overall performance).

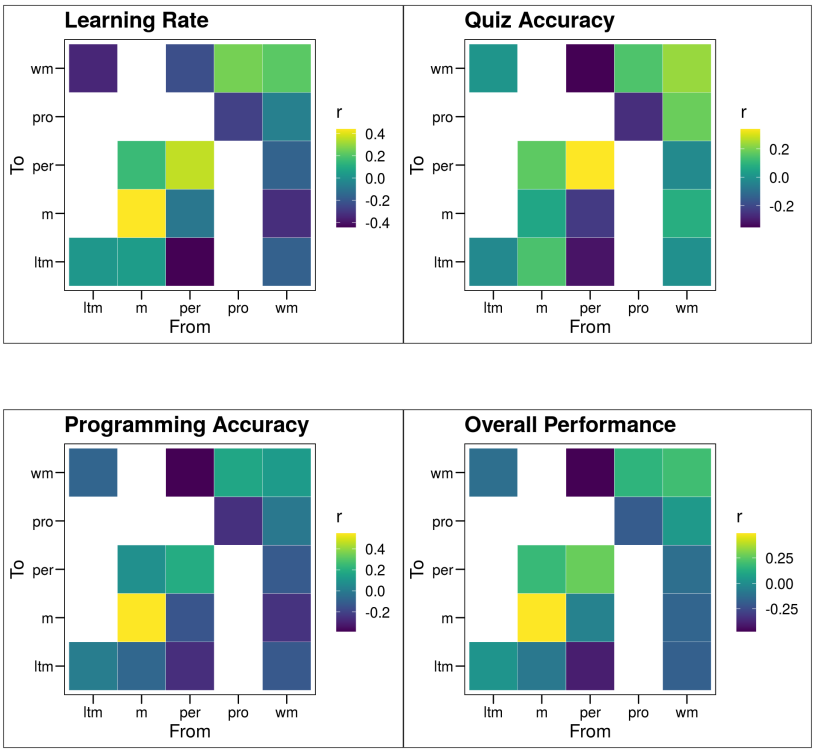

Figure 14: Relationship between different programming learning outcomes and CMC connectivity parameters.

Idiographic Relationship Between Architecture Parameters and Brain Activity

To a certain extent, it is unreasonable to expect a direct mapping between effective connectivity values in the CMC model and individual abilities. This is because the CMC was explicitly designed to capture a sparse, minimal architeecture underlying brain activity. Individual differences, however, are likely underpinned not only by these connections, but also by the functional computations carried out by the corresponding regions.

Importantly, the original CMC paper include not only a list of components and their functional connections, but also a list of minimal functional requirements for each components. These minimal functional requirements can be given a computational interpretation. For example, the procedural memory components, associated with the basal ganglia, is supposed to implement Reinforcement Learning, using the temporal difference method to learn the values of specific actions. Similarly, the long-term memory components, associated with the medial temporal lobe and the hippocampus, is supposed to carry out computations that keep track of multiple decaying traces associated with each memory. Thus, it should be possible, in principle, to find associations between brain activity and the parameters that underpin the workings of these components.

In our investigation, we decided to focus on the long-term memory component, whose computations are much less understood than the procedural memory's Reinforcement Learning. We also decided to focus on resting-state brain activity, which captures task-free characteristics of an individual and has been successfully used to study the nature of the CMC (Goal 1). Unlike our other work, we decided to use EEG, rathern than fMRI, data, which provides a complementary view of brain activity.

The computations of the long-term memory component were instantiated using a standard implementation of the multiple trace theory, pioneered by Anderson and Schooler (1991) and Pavlik

and Anderson (2005). In this implementationm, the activation $A(m)$ of a memory m at time t is described by the following two equations:

$$A_t(m) = \sum n \log (t - t_n)^{-d(n)} \quad (2)$$
$$d(n) = e^{A(m)} + \alpha$$

In this equation, t_n is the time at which the same memory has been experienced for the n-th time, $d(n)$ is a trace-specific forgetting rate, and α is an individual specific parameter that regulates the rate of forgetting. We collected 5-minutes of eyes-open and 5-mins of eyes-closed EEG recordings for N=50 undergraduates. The value of α, which we will refer to simply as the rate of forgetting, was measured for each participants using an adaptive fact learning paradigms, which adapted the rate of presentation of new stimuli to match a person's rate of forgetting.

Our results showed that individual differences in rate of forgetting were broadly correlated with brain activity across multiple frequency bands and channel locations. To better understand teh data, we use a machine-learning approach. Specifically, we fit a LASSO model to the matrix of neuroimaging features (mean power across 16 channels x 7 frequency bands for each participant) and used cross-validation to determine the degree to which individual rates of forgetting could be decided from resting state EEG data. Our results show that individual rate of forgetting could be decoded with high accuracy (r = .90: Figure 15) using the distributions of power spectra in the delta, alpha, and gamma band from frontal and temporal channels (Figure 16):

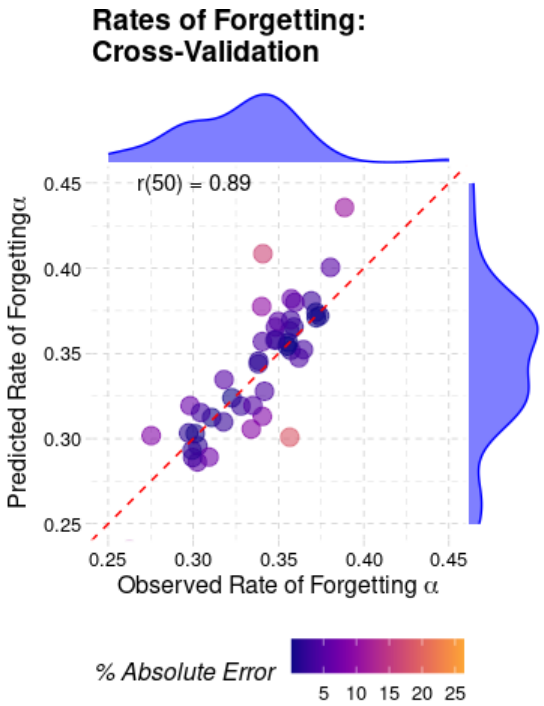

Figure 15: Cross-validation of the LASSO model predicting individual differences in rate of forgetting from resting-state EEG data.

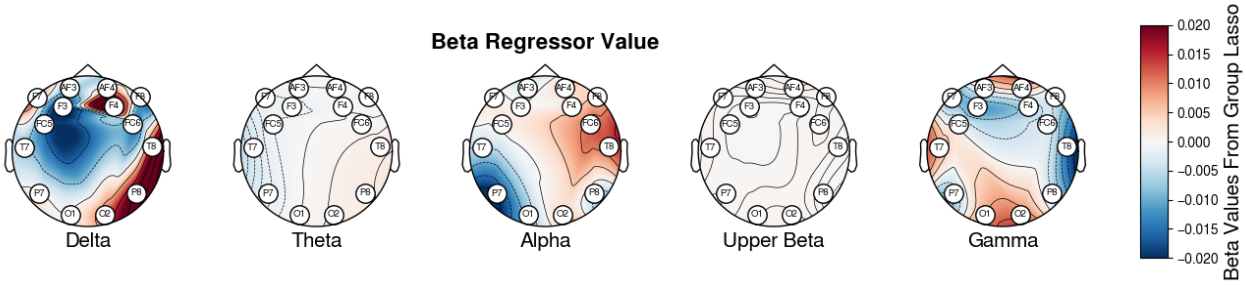

Figure 16: Distribution of the power signals that best predict individual differences in forgetting from EEG data.

Because EEG data does not easily allow for the spatial localization of the underlying signals, the same analysis was repeated using a subset of *N*=33 individuals for whom resorting-state fMRI data was available. To gether a higher-density connectivity matrix, we used the Power (2011) parcellation, and computed the mean correlation coefficients across all of the 264 regions. We applied the same machine learning technique (LASSO) and found that, once again, we were able to correctly decode individual differences in rate of forgetting ($r = 0.96$, Figure 17).

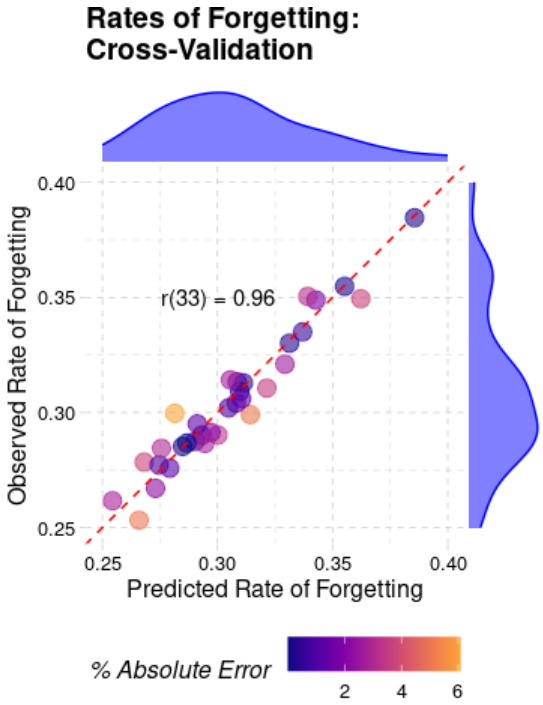

Figure 17: Cross-validation of the LASSO model predicting individual differences in rate of forgetting from resting-state fMRI data.

Using the high-density Power parcellation, we were also able to successfully idemtify which parts of the brain connectivity were the best predictors of individual differences in rate of forgetting. Specifically, a small set of functional connections between fronto-medial prefrontal regions, left lateral prefrontal regions, and temporal regions were sufficient to decode the value of the parameter (Figure 18).

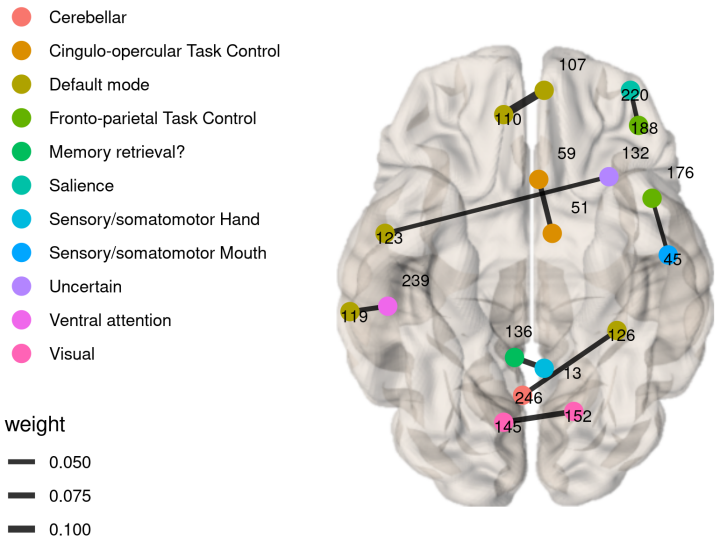

Figure 18: Distribution of the functional connectivity signals that are sufficient to predict an individual's rate of forgetting.

Data Availability

All of the data for this project has been made available on public repositories, as detailed in Table 1.

Table 1: Data and Code Availability for the Project

Repository	Contents
github.com/UWCCDL/CMC-DCM	DCM Analysis of Task-based data (Goal 1); DCM Analysis of resting-state data (Goal 1); Granger causality modeling (Goal 2)
github.com/IBIC/PDN-Pipelines	Application of CMC to Parkinson's Disease data (Goal 2)
github.com/UWCCDL/EEG_RateOfForgetting	Analysis of individual differences in LTM predicted by EEG data (Goal 3)
github.com/UWCCDL/Beluga	Analysis of individual differences in LTM predicted by fMRI data (Goal 3)
github.com/UWCCDL/ProcVsDecl	Analysis of individual difference in functional connectivity predicting LTM- or PM-based strategies (Goal 3)

List of Publications Attributed to the Grant

Peer-Reviewed Journals

Yang, Y. C., Sibert, C. L., & Stocco, A. (under review). *Reliance on Episodic vs. Procedural Systems in Decision-Making Depends on Individual Differences in Their Relative Neural Efficiency*.

Stocco, A., Rice, P. J., Thomson, R., Smith, B., Morrison, D., & Lebiere, C. (under review). *An Integrated Computational Framework for the Neurobiology of Memory Based on the ACT-R Declarative Memory System.*

Hake, H. S., Sibert, C., & Stocco, A. (2022). Inferring a cognitive architecture from multitask neuroimaging data: A data‐driven test of the Common Model of Cognition using Granger causality. *Topics in Cognitive Science, 14*(4), 845-859.

Sibert, C. L., Hake, H. S., & Stocco, A. (2022). The structured mind at rest: Low-frequency oscillations reflect interactive dynamics between spontaneous brain activity and a common architecture for task control. *Frontiers in Neuroscience 16*, 832503.

Wapstra, N. J., Ketola, M., Thompson, S., Lee, A., Madhyastha, T., Grabowski, T. J., & Stocco, A. (2022). Increased Basal Ganglia Modulatory Effective Connectivity Observed in Resting-State fMRI in Individuals With Parkinson's Disease. *Frontiers in Aging Neuroscience, 14*.

Stocco, A. (2021). Qualitative invariant effects arise from neural constraints: Common architecture and sources of individual differences. *Journal of Cognition, 4*(1), 53.

Smith, B. M., Chiu, M., Yang, Y., Sibert, C., & Stocco, A. (2021) When fear shrinks the brain: A computational model of the effects of post-traumatic stress on hippocampal volume. *Topics in Cognitive Science, 13*(3), 499–514.

Yang, Y. C., Karmol, A. M., & Stocco, A. (2021). Core Cognitive Mechanisms Underlying Syntactic Priming: A Comparison of Three Alternative Models. *Frontiers in Psychology, 12*, 2293.

Stocco, A., Sibert, C., Steine-Hanson, Z., Koh, N., Laird, J. E., Lebiere, C. J., & Rosenbloom, P. (2021). Analysis of the human connectome data supports the notion of a "Common Model of Cognition" for human and human-like intelligence across domains. *NeuroImage, 235*, 118035.

Zhou, P., Sense, F, van Rijn, H., & Stocco, A. (2021) Reflections of idiographic model parameters in resting-state neuroimaging data. *Cognition, 212*, 140660.

Peer-Reviewed Conference Proceedings

Lebiere, C., Cranford, E. A., Martin, M., Morrison, D., & Stocco, A. (2022). Cognitive Architectures and their Applications. In *2022 IEEE 8th International Conference on Collaboration and Internet Computing (CIC)* (pp. 55-60). IEEE.

Sibert, C., Hake, H., & Stocco, A. (2021). The structured mind at rest: Evidence for the "Common Model of Cognition" in resting state fMRI. *Proceedings of the 19th International Conference on Cognitive Modeling*

Rice, P. J., & Stocco, A. (2021). Estimating individual differences in working memory through ACT-R modeling and resting state connectivity. *Proceedings of the 18th International Conference on Cognitive Modeling.*

Hake, H., Sibert, C., & Stocco, A. (2021). Inferring a cognitive architecture from multi-task neuroimaging data: A data-driven test of the common model of cognition using Granger causality. *Proceedings of the 18th International Conference on Cognitive Modeling.*

Yang, Y., Sibert, C., & Stocco, A. (2021). Individual differences in decision making strategies can be predicted by resting-state functional connectivity. *Proceedings of the 18th International Conference on Cognitive Modeling*.

Xu, Y., Prat, C. S., Sense, F., van Rijn, H., & Stocco, A. (2021) Distributed brain connectivity predicts individual differences in forgetting: A neurocomputational analysis of resting-state fMRI. In *Proceedings of the 43nd Annual Meeting of the Cognitive Science Society, p. 3186*.

Sibert, C., Hake, H. S., Laird, J. L., Lebiere, C., Rosenbloom, P., & Stocco, A. (2021). The role of the basal ganglia in the human cognitive architecture: A Dynamic Causal Modeling comparison across tasks and individuals. In *Proceedings of the 43nd Annual Meeting of the Cognitive Science Society, pp. 418-424.*

Smith, B. M., Chiu, M., Yang, Y., Sibert, C., & Stocco, A. (2020) Modeling the effects of post-traumatic stress on hippocampal volume. *Proceedings of the 18th International Conference on Cognitive Modeling, pp. 263-269.* **[♛ Allen Newell Prize for Best Student Paper]**