

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Applying the Common Model of Cognition to Resting-State fMRI Leads to the Identification of Abnormal Functional Connectivity in Parkinson's Disease

Permalink

<https://escholarship.org/uc/item/3404d4c0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Ketola, Micah

Thompson, Shelby

Madhyastha, Tara

et al.

Publication Date

2020

Peer reviewed

Applying the Common Model of Cognition to Resting-State fMRI Leads to the Identification of Abnormal Functional Connectivity in Parkinson's Disease

Micah Ketola (ketolm@uw.edu)

Department of Radiology and Integrated Brain Imaging Center, University of Washington Medical Center,
Campus Box 357115, Seattle, WA 98195 USA

Shelby Thompson (shelbyelizabetht@gmail.com)

Department of Environmental and Health Sciences, Spelman College,
350 Spelman Lane SW, Atlanta, Georgia 30314 USA

Tara Madhyastha (taramad@amazon.com)

Amazon Web Services,
1007 Stewart St, Seattle, WA 98101 USA

Thomas J. Grabowski (tgrabow@uw.edu)

Department of Radiology and Integrated Brain Imaging Center, University of Washington,
Campus Box 357115, Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)

Department of Psychology and Institute for Learning and Brain Sciences (I-LABS), University of Washington,
Campus Box 351525, Seattle, WA 98195 USA

Abstract

A complete understanding of cognitive function in humans must incorporate a model of interactions between networked brain regions. Alterations to these network interactions underlie cognitive impairment in many neurodegenerative diseases, providing an important physiological link between brain structure and cognitive function. Cognitive architectures have often been used to explain how healthy brains function, typically using task-based activity. However, this description is incomplete. Most systems-level brain activity is spontaneous, or intrinsic, and occurs whether or not a subject is performing a task. Here, we provide evidence that the Common Model of Cognition, a consensus model derived from an analysis of existing cognitive architectures, can (a) be generalized to account for brain activity at rest, rather than during tasks, and (b) correctly identify differences in basal ganglia connectivity in Parkinson's Disease.

Keywords: Common Model of Cognition; Resting-state fMRI; Parkinson's Disease; Dynamic Causal Modeling

Introduction

Cognitive architectures aim to provide an explanation for cognition in terms of interacting brain modules. They support a broad framework for which specific models are created and in theory these models, provided the correct assumptions within a sufficient framework, replicate how the brain accomplishes a task. The modules that make up an architecture are functionally defined and in many cases have been successfully identified with brain regions during the analysis of task-based activity (Anderson, 2007; Anderson, Albert, & Fincham, 2005; Anderson, Fincham, Qin, & Stocco, 2008; Borst, Taatgen, Stocco, & Rijn, 2010). This has enabled architectures to form neuroimaging predictions, greatly expanding the space in which models can be compared and validated.

While cognitive architectures have been extremely successful in replicating healthy brains and even individual differences (Daily, Lovett, & Reder, 2001), they have been seldom applied to neurological diseases and, therefore, have had limited translational applications. We see this as a promising area of research in cognitive modeling. In accounting for deficiencies, we are able to provide a litmus test of sorts. Given that assumptions of the architecture do not hold in modeling brains with neurological diseases, we may better understand the impact of such diseases on the brain, and reevaluate the validity of assumptions made. This will increase the legitimacy of the architecture and provide a more holistic understanding of brain and disease.

Here, we show that the Common Model of Cognition (CMC), a consensus architecture derived from an analysis of existing cognitive architectures (Laird, Lebiere, & Rosenbloom, 2017), can be successfully applied to the the clinical problem of identifying abnormal functional connectivity in Parkinson's Disease (PD). This application is important as results of functional connectivity abnormality in PD have yielded inconsistent and contradictory results (Baggio, Segura, & Junque, 2015; Göttlich et al., 2013). To further compound the perplexity of the situation, the etiology of PD is well understood – its symptoms quite apparent and clear – and yet, its brain signatures difficult to catch.

Parkinson's Disease

Parkinson's Disease is a neurodegenerative disease that targets predominantly dopaminergic neurons in the basal ganglia. Contemporary accounts of basal ganglia function (Stocco, Lebiere, & Anderson, 2010; Frank, Seeberger, &

O'Reilly, 2004), suggest that they work by controlling or “gating” the influx of signals from other cortical areas to the prefrontal cortex. In PD, the loss of dopamine inputs to the basal ganglia causes an upregulation of their inhibitory pathways (and a downregulation of their excitatory pathways), resulting in more conservative gating (Albin, Young, & Penney, 1989). Symptoms include visible motor impairments such as shaking and slowness of movement, as well as non-motor complications such as depressive disturbances and cognitive impairment (Marsh, 2013; Watson & Leverenz, 2010).

We hypothesized that a potential problem with previous work is that it has been almost entirely data driven, rather than theory-driven. Thus, it is possible that the effects of PD are buried in subtle patterns of data, and that these patterns might not be observable through traditional analyses. For this reason, we decided to specifically model the loss of dopamine inputs, using an architecture from Cognitive Science, the CMC.

The Common Model of Cognition

The CMC is a consensus architecture for general intelligence that is the culmination of work from over five decades in artificial intelligence, cognitive science, neuroscience, and robotics (Laird et al., 2017). Rather than a formal implementation, it serves as a blueprint to understand the organization of a human-like mind. There are five functional components: long-term memory, working memory, procedural memory, perception systems, and action systems.

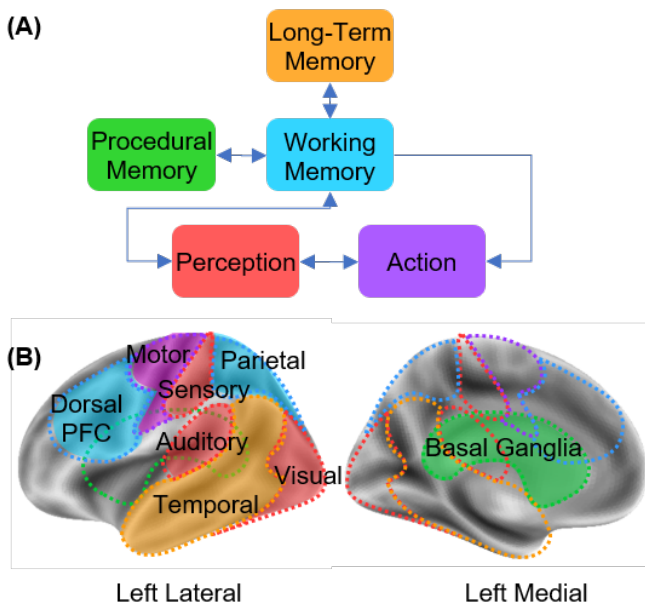


Figure 1: (A) Architecture of the Common Model of Cognition, as described by (Laird et al., 2017). (B) Theoretical mapping between CMC components and homologous cortical and subcortical regions.

Importantly, the CMC provides theory-driven hypotheses about the functional relationship between brain regions. Some of these hypotheses have a clear connection with PD,

since the loss of dopamine neurons has cognitive consequences that can be computationally characterized (Frank et al., 2004) and have been successfully modeled in cognitive architectures (Stocco, 2018).

Although the CMC is a purely *functional* architecture, whose components are characterized in terms of abstract computations, researchers have proposed, developed, and even tested methods to map the components of the CMC onto homologous brain region (1B) and to translate the relationships between CMC components into predicted patterns of functional connectivity, (Stocco, Laird, Lebiere, & Rosenbloom, 2018; Steine-Hanson, Koh, & Stocco, 2018). Ultimately, these efforts have shown that the CMC provides a remarkably good fit to fMRI data from over 200 participants, across a variety of representative tasks (Stocco et al., 2019).

The success of this approach suggests that the CMC might provide a new, theoretically-driven way to look at functional connectivity data in the human brain. Inspired by its success in accounting for functional connectivity in healthy young adults, we decided to extend this approach to the case of PD.

To do so, the original approach was extended to apply the CMC to *resting-state* fMRI (rs-fMRI), instead of *task-based* fMRI. While all previous applications of the CMC to neuroimaging data have focused on task-based fMRI, in contemporary clinical research the dominant approach is to use rs-fMRI.

Resting-state fMRI

rs-fMRI consists of continuous recordings of brain activity while participants are awake but not engaged in any particular task; the analysis of such recordings has become integral in the identification and investigation of networks within the brain. Although it is most intuitive to think about measuring network function during tasks (when cognitive networks are specifically recruited), it is now known that the majority of brain activity occurs spontaneously and in the absence of specific stimuli. Analysis of rs-fMRI data has shown that intrinsic activity has a rich spatiotemporal structure, reflecting how networks of cortical regions combine and recombine over time.

Investigations into intrinsic activity have led to the discovery of many innate networks, the best known being the Default Mode Network (DMN). This network, as well as many others, inform us about the intrinsic rhythms and oscillations in healthy and disordered brains. Irregularity in network functioning has been discovered in conjunction with neurodegenerative diseases, psychiatric disorders, aging, and more (Hohenfeld, Werner, & Reetz, 2018; Sambataro et al., 2010). While the extent to which network function impacts cognition is lesser understood, networks can nonetheless serve as biomarkers to help diagnose and track disease progress (Hohenfeld et al., 2018).

Further, rs-fMRI has been shown to be highly predictive of brain activity during cognitive tasks (Cole, Ito, Bassett, & Schultz, 2016). The same structural networks that are activated at rest provide a base for which cognitive activation

flows. In better understanding how resting networks are affected by PD, we are able to investigate potential abnormalities in the building blocks that give rise to cognition.

Dynamic Causal Modeling

As in previous investigations into the relationship between the CMC and brain data (Stocco et al., 2018; Steine-Hanson et al., 2018; Stocco et al., 2019), functional connectivity between brain regions is analyzed in terms of *effective* connectivity using Dynamic Causal Modeling (Henceforth, DCM: Friston, Harrison, & Penny, 2003). DCM estimates functional connectivity between pairs of regions by iteratively refining the parameter estimates of a dynamic system, in which brain regions are approximated as neural mass points. Thus, in contrast to most functional connectivity analyses, which are “bottom-up” and data-driven, DCM is a “top-down”, model-based technique, in which a theoretical model of network connectivity is fit to the data. As such, it provides a natural way to implement the CMC.

DCM was preferred over other approaches such as Structural Equation Modeling (SEM) and Granger Causality Modeling (GCM) due to convenience. This choice was made for three reasons. First, it can account for the temporal dynamics of an fMRI time-series, whereas SEM cannot (Friston, 2011). This is important as without temporal dynamics the data is effectively reduced by one dimension, leaving out a significant source of variability. Second, it allows us to better model directed causal influences, which are implied in the directed arrows between CMC components (Fig. 1A). Although GCM *can* disentangle the direction of influence, DCM has proven superior at dealing with the variable nature of the BOLD response timing (Friston, 2009). Finally, DCM, but neither SEM nor GCM, is capable of modeling second-order interactions between nodes in a network, i.e., cases in which a region modulates the connectivity between two other regions. This specific case, as it will be shown, is of particular interest in PD, as it plays a significant role in capturing the nature of basal ganglia function.

DCM is composed of both a *neural* model, that receives experimental stimuli and predicts the underlying dynamics of brain activity, and an *observational* model, that takes in the predicted underlying dynamics and outputs predictions of observed brain activity. In the case of fMRI, the neural model is given regions of interest (ROIs) and represents the time course of activity in each region i as a nonlinear state equation:

$$\dot{\mathbf{y}} = \mathbf{A}\mathbf{y} + \sum_i x_i \mathbf{B}(i)\mathbf{y} + \mathbf{C}\mathbf{x} + \sum_i \mathbf{y} \mathbf{D}(i)\mathbf{y}_i \quad (1)$$

In this equation \mathbf{A} defines intrinsic connectivity between different regions (fixed connectivity), \mathbf{B} defines the modulatory effects that task conditions have on the connectivity between regions (modulation of connectivity), \mathbf{C} defines effects by task inputs, \mathbf{D} defines the modulatory effects that regions have on the connections between other regions, \mathbf{x} defines task inputs, and \mathbf{y} defines brain activity.

The observational model is composed of a hemodynamic model that uses neural activity to cause changes in blood flow, which in turn causes changes in blood volume and the amount of deoxyhemoglobin. From there, the volume of blood and deoxyhemoglobin concentration are entered into an output nonlinearity, and give rise to an observed BOLD response (Friston et al., 2003).

DCM and Resting-State

Because during rs-fMRI there is no task to be performed nor significant external events driving brain activity, the application of DCM to rs-fMRI posed a significant challenge. Without any task conditions or external input to initiate network dynamics, and therefore a null \mathbf{C} matrix, the DCM would simply remain uninitialized with all parameters left at the default values. Friston, Kahan, Biswal, and Razi (2014) circumnavigated the problem by creating a deterministic DCM. Their version of resting-state DCM estimates effective connectivity based on second-order statistics rather than on the time-series of activation. This transforms analysis from the computationally expensive issue of estimating hidden neuronal states to the more efficient problem of estimating the spectral density of activity changes (Friston et al., 2014). While more computationally efficient, by using second-order statistics their method can no longer capture temporal dynamics in the estimation of effective connectivity. For this reason, we adopted an alternative procedure proposed by Di and Biswal (2014).

One of the characteristics of resting-state brain activity is the presence of spontaneous correlations at very low-frequencies, which organize brain networks along different rhythms (Fox et al., 2005). Di and Biswal (2014) explicitly modeled these low-frequency fluctuations (LFF) within the resting-state signal using deterministic inputs. They used periodic sine and cosine functions at 0.01, 0.02, 0.04, and 0.08 Hz (Di & Biswal, 2014) as task conditions for the \mathbf{C} matrix. As drivers for activity in DCM cannot be partial, the periodic functions were transformed into boxcar functions. There were two boxcar functions for each frequency with a 90 degree lag in between, at cycles of 100, 50, 25, and 12.5 s. The boxcars were used as input to every node in their analysis replacing traditional task-based inputs. To validate their procedure, they showed that an F -test based on their sinusoidal regressors correctly identified the main nodes of the DMN, and they successfully fitted DCMs based on the regions thus identified. (Di & Biswal, 2014).

Because this study similarly dealt with rs-fMRI, we borrowed the procedure developed by Di and Biswal (2014), making use of eight boxcar regressors derived from sine waves of different frequencies and phases (Fig. 2).

CMC Implementations for PD

As the correct interpretation of the role of the basal ganglia within the CMC is crucial to understanding PD, two different interpretations of the CMC were explored and tested against controls and PD patients.

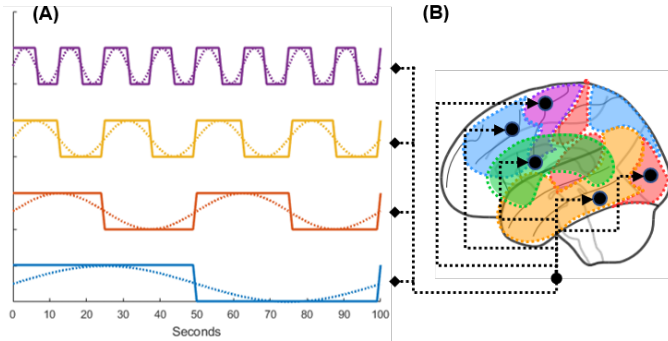


Figure 2: (A) Illustration of how low-frequency fluctuations were transformed into boxcar functions, which were then (B) used as drivers for the activity in all ROIs.

The first model, the *Direct* CMC, is perhaps the most straightforward translation of the CMC into patterns of functional connectivity. Specifically, in this model all connections between ROIs are implemented as patterns of direct connectivity (i.e., matrix A). Critically, Procedural Memory directly projects to Working Memory (Figure 3), reflecting the assumption that procedural knowledge directly manipulates the contents of working memory. This implementation has been previously tested in Stocco et al. (2019).

The second model, the *Modulatory* CMC, incorporates additional assumptions that capture our modern understanding of the basal ganglia, the brain region associated with the CMC's Procedural Memory component. As its name implies, this model replaces the direct connections from Procedural Memory to Working Memory with two second-level, modulatory connections that control projections from Perception to Working Memory, and from Long-Term Memory to Working Memory (Figure 3). Thus, in this model, the connectivity between ROIs involves both matrices A and D of Eq. 1. This implementation was previously tested in Stocco et al. (2018). The use of the modulatory, instead of direct, connections reflects a different functional view of the basal ganglia; according to this view, the basal ganglia do not directly manipulate the contents of working memory, but rather "gate" (Frank et al., 2004) or "route" (Stocco et al., 2010) information from other areas.

Experimental Predictions

Based on the existing literature, two predictions were made.

Our first prediction is that the modulatory model would provide a better fit than the direct model to the data across both groups of participants, PD patients and controls. This prediction is supported by the fact that the modulatory model better captures the functional role of the basal ganglia, as seen in contemporary models (Stocco et al., 2010; O'Reilly & Frank, 2006).

Our second prediction is that, within the modulatory model, a difference will be found between PD patients and controls in the parameters that regulate modulatory connec-

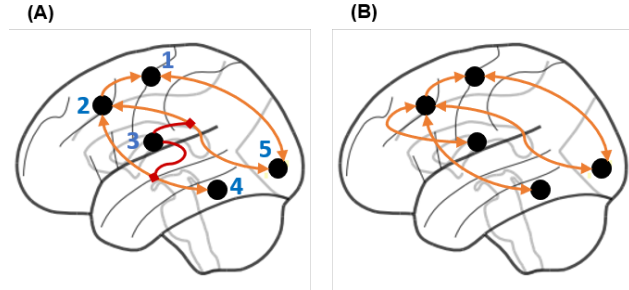


Figure 3: (A) In the Modulatory Model, the Procedural Memory component modulates the signals flowing through incoming connections to Working Memory. (B) In the Direct Model, Procedural Memory has a direct connection to Working Memory. (Regions are labelled as follows: 1. Action, 2. Working Memory, 3. Procedural Memory, 4. LTM, 5. Perception)

tivity (i.e., matrix D in Eq. 1). Specifically, we predict that modulatory signals will be lower in PD patients than in controls. This is because dopamine depletion in PD is typically understood as causing an increase in the filtering activity of the basal ganglia (the "brake" in the so-called "brake-accelerator" model: Albin, Young, and Penney (1989)), which, in DCM, would be reflected in a lower or negative value in the D matrix.

Materials and Methods

Participants

Participants were recruited from a larger, multimodal study of functional networks in PD. As part of a comprehensive protocol involving EEG and MRI, 111 participants received a resting fMRI scan. Participants were monitored using eye-tracking for wakefulness. Of these, 6 participants were excluded due to insufficient data quality. Among the 105 remaining participants (PD: $N = 67$, Age = 67.54 ± 8.03 , Female = 26; Controls: $N = 39$, Age = 69.41 ± 8.88 , Female = 15), 2 participants were diagnosed at consensus with dementia (PD = 2; Controls = 0), and 37 participants were diagnosed at consensus with Mild Cognitive Impairment (PD = 26; Controls = 11).

Image Acquisition and Processing

MRI data was acquired on a research-dedicated 3T Philips Achieva whole-body scanner (Philips Medical Systems, R5.1.7) with a 32-channel SENSE head coil at the Integrated Brain Imaging Center of the University of Washington, Seattle. Functional resting-state data was acquired while participants were instructed to lay quietly and focus on a fixation cross, using a gradient echo-planar multi-echo pulse sequence with TR = 2,500 ms, a 79° flip angle, and TE = 9.5/27.5/45.5 ms. Multiecho recordings allow for increased sensitivity and a reduced amount of artifacts. Each volume acquisition consisted of 37 oblique axial slices, each of which was 3.5 mm

thick with 0-mm gap and contained 64×64 voxels with an in-plane resolution of 3.5×3.5 mm.

In addition to functional images, a T1-weighted structural scan was acquired as an anatomical reference (1-mm isotropic multiecho MP-RAGE: Sagittal TR = 10.019 ms, TE = 4.61 ms, FoV = 260x260x189.6 mm, and an 8° flip angle).

Resting-state fMRI data were processed using a combination of FSL, AFNI, and SPM. Functional data underwent slice-timing correction, motion correction and realignment, removal of baseline drift, and then spike detection and removal. From there the data were co-registered to the skull-stripped T1-weighted structural scan, normalized to the MNI ICBM152 stereotactical space, and smoothed using a 3D Gaussian filter with 8-mm full-width half maximum.

Bilateral Regions of Interest

Departing from previous work that has focused on left-hemisphere regions, this study used bilateral, instead of unilateral ROIs. While unilateral ROIs are understandable for task-based fMRI due to the prevalence of increased activity in the left hemisphere (Barch et al., 2013; Stocco et al., 2019), it was not justifiable in our case. Bilateral ROIs were obtained by combining homologous regions in the left and right hemisphere.

The ROIs were selected following the procedure described in Stocco et al. (2019) to translate CMC components into anatomical regions. This procedure proceeds by first identifying the candidate regions in a group-level analysis of all data. To account for individual differences in functional neuroanatomy, the precise coordinates of each ROI are then localized on an individual brain by identifying the peak of functional activity that is closest to the centroid of the group-level ROIs. The distribution of the centroids of the individual-level ROIs are visually represented in Figure 4. All ROIs were bilateral with the exception of the Action ROI. The Action ROI was lateralized to the left hemisphere.




Figure 4: Location of the centroids of each individual ROI. Each point represents the centroid of one ROI; variations in the centroids account for individual differences in functional anatomy.

Model Comparison Procedure

Given two or more architectures designed to reproduce the timecourse of brain activity, it is possible to compare them by estimating a likelihood function $\mathcal{L}(m|y)$ that assigns a poste-

rior probability to a given model m , given the data y . Likelihood functions are typically chosen to balance between rewarding higher model fit (i.e., minimizing the residuals) and penalizing model complexity (i.e., reducing the number of parameters). Two common examples of these metrics are the Bayesian Information Criterion (BIC) and the Akaike Information Criterion. In DCM, it is customary to use a metric known as *Free Energy* (Kasess et al., 2010), which is similar to the BIC but, importantly, does not assume that parameters are independent of each other. Group-level likelihood values for a model m can then expressed as the product of the likelihood of that model fitting each participant p , i.e., $\prod_p \mathcal{L}(m|y_p)$. In terms of log-likelihood, this translates to the sum of all of the individual log-likelihoods: $\sum_p \log \mathcal{L}(m|y_p)$. Although more sophisticated model comparison procedures have been proposed (Stephan, Penny, Daunizeau, Moran, & Friston, 2009), the log-likelihood based metric used here is not only the most easily interpretable, but also the most relevant, as it specifically applies to cases in which it is assumed that the model is constant or architectural across individuals (Kasess et al., 2010).

Results

Prediction 1: Model Comparison

To test the first experimental prediction, the log likelihoods $\mathcal{L}(m|y_p)$ of the two models were compared in PD and controls. The comparisons provided unequivocal support for the modulatory version of the CMC architecture in both cases. The modulatory model's log-likelihood exceeded the direct model's by 90 in the case of controls, and by 20 in the case of PD patients (Fig. 5A-B). Since a difference in log-likelihoods represents a ratio of likelihoods, modulatory model is approximately e^{20} to e^{90} times more likely than the direct model, given the observed timecourses of neural activity in the five ROIs.

This difference in log-likelihoods can also be translated and interpreted as an equivalent p -value in statistical analysis using Wilk's theorem (Wilks, 1938). The theorem states that, given a likelihood ratio λ between two models, the value $-2\log(\lambda)$ approximates the corresponding statistic of a χ^2 distribution with the degrees of freedom corresponding to 1 + the difference in parameters between models. Thus, $\log(\lambda) = \log(\mathcal{L}(m_{\text{direct}})/\mathcal{L}(m_{\text{modulatory}})) = \log \mathcal{L}(m_{\text{direct}}) - \log \mathcal{L}(m_{\text{modulatory}}) = -90$ (for Controls) and -20 (for PD). The difference in the number of parameters between the two models can be set to $n = 2$, which is the number modulatory parameters in the matrix \mathbf{D} in Eq. 1; this matrix is present in the modulatory model but absent in the direct model. The p -value associated with $\chi^2_{(n=3)}$ is $p < 0.0001$ for both groups (Fig. 5C), thus suggesting almost zero chance that the alternative hypothesis (that is, the direct model) could better explain the data.




Figure 5: (A) Absolute log-likelihoods for models in the PD group and (B) in the Control group. (C) Corresponding values on a χ^2 distribution.

Prediction 2: Group Differences

Having established the superiority of the modulatory model, we analyzed whether there were significant differences between PD and Control modulatory connection strengths. The mean square error (MSE) between groups was computed using the D matrices of functional connectivity, $MSE = 0.0952$. To account for potentially non-normally distributed parameter values, we completed non-parametric permutation tests. For these tests, group membership was assigned randomly and the MSE was computed separately 10,000 times. Out of 10,000 tests, there were 192 instances where the computed MSE was equal to, or larger than, the original value, implying that the difference is significant ($p < 0.02$).

On average, the modulatory connectivity was negative in both groups, confirming our second prediction on the role of the basal ganglia in gating and filtering cortical signals to the prefrontal cortex (Stocco et al., 2010; Frank et al., 2004). As hypothesized, patients exhibited stronger (i.e., more negative) modulatory effects than controls, consistent with a downregulation of the direct, or excitatory, pathway of the basal ganglia.

Group Differences in Intrinsic Connectivity

In addition to the modulatory connectivity, group differences were also investigated across all other connectivity parameters in the modulatory model, that is, those encapsulated in matrices A and C of Eq. 1. The comparison was carried out using the same procedure outlined in the previous section. No significant differences were found in functional connectivity between PD and Controls in either model. This lack of effect is broadly consistent with the general lack of effect that is reported in the literature, and is consistent with the specificity of PD in targeting the basal ganglia

Discussion

In this paper, we have shown that cognitive architectures can be applied in clinical neuroscience to identify abnormal patterns of brain activity that are characteristic of neurological diseases, such as PD. Specifically, we have shown that by adapting the consensus architecture known as the Common Model of Cognition (Laird et al., 2017), we could successfully identify abnormal patterns of functional connectivity in PD patients, which have proven otherwise elusive when bottom-up approaches were attempted. The direction of the results, with PD patients exhibiting greater inhibition of cortico-cortical connectivity from the basal ganglia, are also consistent with the known etiology of the disease.

These results should be considered in light of a number of limitations. The first being that they were derived from a reasonable but still small number of individuals. Further, the method used for resting-state DCM does not allow for dynamic changes in effective connectivity. It uses data from the entire 10 minute resting-state scan to create a single account of brain connectivity. We understand that this account is in part naive, as the brain fluctuates throughout resting, and see this as an area for future research.

These limitations notwithstanding, we believe that our results have several implications. First, our results demonstrate the potential of using cognitive architectures in the domain of clinical neuroscience. By distilling decades of brain and cognitive research, cognitive architectures provide a fundamental tool to implement a-priori hypotheses and informed dimensionality reduction in the analysis of brain data for patients.

At the same, our findings provide further credibility to the Common Model of Cognition. Specifically, they show that it can be applied to resting-state as well as task-based neuroimaging data, significantly expanding its application and bringing it into the fold of contemporary approaches in neuroscience.

Finally, our results also suggest important modifications to the Common Model of Cognition. Specifically, they suggest that the CMC should be revised to explicitly include the modulatory role of the basal ganglia as part of the function of the Procedural Memory component.

Acknowledgments

This research was supported by grants from the National Institutes of Health, R01 NS099199-04 (Grabowski), and the Air Force Office of Scientific Research, FA9550-19-1-0299 (Stocco).

References

- Albin, R. L., Young, A. B., & Penney, J. B. (1989, jan). The functional anatomy of basal ganglia disorders. *Trends in Neurosciences*, 12(10), 366–375.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford University Press.

- Anderson, J. R., Albert, M. V., & Fincham, J. M. (2005). Tracing problem solving in real time: fmri analysis of the subject-paced tower of hanoi. *Journal of Cognitive Neuroscience*, *17*(8), 1261–1274.
- Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, *12*(4), 136–143.
- Baggio, H. C., Segura, B., & Junque, C. (2015). Resting-State Functional Brain Networks in Parkinson’s Disease. *CNS Neuroscience and Therapeutics*, *21*(10), 793–801.
- Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., . . . others (2013). Function in the human connectome: Task-fmri and individual differences in behavior. *Neuroimage*, *80*, 169–189.
- Borst, J. P., Taatgen, N. A., Stocco, A., & Rijn, H. V. (2010). The neural correlates of problem states: Testing fmri predictions of a computational model of multitasking. *PLoS ONE*, *5*(9).
- Cole, M. W., Ito, T., Bassett, D. S., & Schultz, D. H. (2016). Activity flow over resting-state networks shapes cognitive task activations. *Nature Neuroscience*, *19*(12), 1718–1726.
- Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: A source activation account. *Cognitive Science*, *25*(3), 315–355.
- Di, X., & Biswal, B. B. (2014). Identifying the default mode network structure using dynamic causal modeling on resting-state functional magnetic resonance imaging. *NeuroImage*, *86*, 53–59.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678.
- Frank, M. J., Seeberger, L. C., & O’Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science*, *306*(5703), 1940–1943.
- Friston, K. J. (2009). Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS Biology*, *7*(2), 0220–0225.
- Friston, K. J. (2011). Functional and Effective Connectivity: A Review. *Brain Connectivity*, *1*(1), 13–36.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic Causal Modelling. *Human Brain Function: Second Edition*, *19*, 1063–1090.
- Friston, K. J., Kahan, J., Biswal, B., & Razi, A. (2014). A DCM for resting state fMRI. *NeuroImage*, *94*, 396–407.
- Göttlich, M., Münte, T. F., Heldmann, M., Kasten, M., Hagenah, J., & Krämer, U. M. (2013). Altered Resting State Brain Networks in Parkinson’s Disease. *PLoS ONE*, *8*(10).
- Hohenfeld, C., Werner, C. J., & Reetz, K. (2018). Resting-state connectivity in neurodegenerative disorders: Is there potential for an imaging biomarker? *NeuroImage: Clinical*, *18*(February), 849–870.
- Kasess, C. H., Stephan, K. E., Weissenbacher, A., Pezawas, L., Moser, E., & Windischberger, C. (2010). Multi-subject analyses with dynamic causal modeling. *Neuroimage*, *49*(4), 3065–3074.
- 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*(4), 13–26.
- Marsh, L. (2013). Depression and parkinson’s disease: current knowledge. *Current neurology and neuroscience reports*, *13*(12), 409.
- O’Reilly, R. C., & Frank, M. J. (2006). Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural computation*, *18*(2), 283–328.
- Sambataro, F., Murty, V. P., Callicott, J. H., Tan, H. Y., Das, S., Weinberger, D. R., & Mattay, V. S. (2010). Age-related alterations in default mode network: Impact on working memory performance. *Neurobiology of Aging*, *31*(5), 839–852.
- Steine-Hanson, Z., Koh, N., & Stocco, A. (2018). Refining the Common Model of Cognition through large neuroscience data. *Procedia computer science*, *145*, 813–820.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017.
- Stocco, A. (2018). A biologically plausible action selection system for cognitive architectures: Implications of basal ganglia anatomy for learning and decision-making models. *Cognitive science*, *42*(2), 457–490.
- Stocco, A., Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2018). Empirical evidence from neuroimaging data for a standard model of the mind. In *Proceedings of the 40th annual meeting of the cognitive science society*.
- 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. *Psychological review*, *117*(2), 541.
- Stocco, A., Steine-Hanson, Z., Koh, N., Laird, J., Lebiere, C., & Rosenbloom, P. (2019). Analysis of the Human Connectome data supports the notion of a “Common Model of Cognition” for human and human-like intelligence. *BioRxiv*, 703777.
- Watson, G. S., & Leverenz, J. B. (2010). Profile of cognitive impairment in parkinson’s disease. *Brain Pathology*, *20*(3), 640–645.
- Wilks, S. S. (1938, 03). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Statist.*, *9*(1), 60–62.