The Structured Mind at Rest: Evidence for the "Common Model of Cognition" in Resting State fMRI

Catherine Sibert (sibert@uw.edu)

Department of Psychology, University of Washington, Seattle, WA 98195 USA

Holly Hake (hakehs@uw.edu)

Department of Psychology, University of Washington, Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)

Department of Psychology, University of Washington, Seattle, WA 98195 USA

Abstract

The Common Model of Cognition (CMC) has been proposed as a high level framework through which functional neuroimaging data can be predicted and interpreted. Previous work has found the CMC is capable of predicting brain activity across a variety of tasks, but it has not been tested on resting state data. This paper adapts a previously used method for comparing theoretical models of brain structure, Dynamic Causal Modeling (DCM), for the task-free environment of resting state, and compares the CMC against six alternate architectural frameworks. For a large sample of subjects from the Human Connectome Project (HCP), the CMC provides the best account of resting state brain activity, suggesting the presence of a general purpose structure of connections in the brain that drives activity when at rest and when performing directed task behavior.

Keywords: Brain architecture, Cognitive Architecture, Computational models, Dynamic Causal Modeling, fMRI, Resting state.

Introduction

Despite a shared goal of understanding the underlying mechanisms of the brain, research that focuses on high-level structural models of cognition remains largely isolated from efforts to interpret direct measurements of brain activity. Many neuroscientists are reluctant to rely on the results and conclusions from cognitive architectures because, while the behavior of the models often closely matches observed human data, the mechanisms driving that behavior are rooted in the principles of computer science and information theory. Efforts have been made to connect components of cognitive architectures to corresponding brain regions, but direct biological brain functions are rarely well captured by the more conceptual architecture modules, and architectures often make incompatible assumptions about the basic functional components that are needed to support cognition.

The Common Model of Cognition

One successful attempt to achieve consensus is represented by the so-called Common Model of Cognition (CMC; Laird et al., 2017). The CMC is a computational framework that can serve as a blueprint to understand the organization of a human-like mind. Abstract computations are categorized into five functional components (long-term memory, working memory, procedural memory, perception systems, and action systems) with specific directional relationships (Fig. 1A) between them.

Although it was not proposed specifically as a *brain* architecture, a number of studies have found that the CMC is surprisingly effective at modeling brain activity across tasks and individuals (Steine-Hanson et al., 2018; Stocco et al., 2018, 2021). In this interpretation, the CMC's functional components are mapped onto large-scale brain regions (Fig. 1B) and their relations are translated into predicted patterns of functional connectivity. In other words, the neural counterparts of the functional components and their connections serve as a simplified architecture for the human brain, not only the human mind.

Resting State Brain Activity

A secondary problem with cognitive architecture models is their focus on the brain *at work*. Virtually all mappings between cognitive architectures and brain activity have been carried out based on neural responses to specific tasks (Anderson et al., 2008; Eliasmith et al., 2012). This bias was inherited from the brain imaging analyses carried out to test the CMC, which, so far, have similarly focused on task-based activity.

In contrast, while many analyses of fMRI data compare differences in activity while subjects perform a variety of tasks, a lot of recent work has instead focused on the connectivity of the brain at rest. This line of research was spawned by the observation that even spontaneous brain activity shows a high degree of structure (Fox et al., 2005; Sherzhad et al., 2008), which is revealed in terms of correlations between the time courses of the activity of different brain regions. These patterns of correlations are fairly stable across individuals (Gratton et al., 2018), to the point that variations in the patterns of correlations can be used to reliably predict abnormal neurological conditions (Hohenfeld, Werner, & Reetz, 2018) and can even be used to successfully predict the patterns of brain activity during tasks (Cole et al., 2016; Yeo et al., 2011). These findings suggest the possibility of an underlying

structure to the brain that can be adapted to tasks as needed, but is still present even when resting.

This paper extends the work of Stocco et al. (2021) by testing the Common Model of Cognition on brain activity *at rest* using a pre-defined network of brain regions. Specifically, this paper adapts the framework of Dynamic Causal Modeling (DCM) and compares the Common Model of Cognition (CMC), against six other exemplar network structures that could capture the underlying structure of the mind.

Dynamic Causal Modeling

The DCM framework aims to identify the causal influences of neuronal systems by quantifying the dynamic fluctuations in brain activity (Friston et al., 2003).

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

In this equation, hemodynamic brain activity, represented by vector \mathbf{y} , is multiplied by matrix \mathbf{A} , which contains a set of parameters constituting a proposed structure of connectivity between regions. Thus, the structure of matrix \mathbf{A} can be adapted to test alternative connectivity architectures. \mathbf{C} is a matrix of the parameters that specify how external or driving inputs elicit changes in brain activity, and \mathbf{x} defines the matrix of task inputs. Since there are not any external inputs driving activity in the resting state data, the \mathbf{C} matrix was adapted to model low frequency fluctuations seen in this state using deterministic inputs as task conditions (see Materials and Methods, below).

Alternate Model Architectures

As pointed out in Stocco et al. (2021), DCM is a strictly top-down, theory-driven method, and cannot be used to infer an architecture from the data. Instead, to evaluate the CMC as an architecture, its predictions were compared against a collection of alternative networks that consist of the same components, but different connection patterns (Stocco et al., 2021). These alternate models are not exact implementations of other cognitive architecture systems, like ACT-R or SPAUN, but instead represent the space of possible theoretical neural architectures.

The alternate architectures fall into two broad categories, or families. In the "Hub-and-Spoke" family (Fig. 1C), a single ROI is designated as the central "Hub", and is bidirectionally connected to all other ROIs. However, none of the "Spoke" ROIs are connected to any other - all activity must travel through the "Hub". Three different Hub-and-Spoke models are considered, based on whether the role of the hub is played by the Prefrontal Cortex, mapped to Working Memory (as proposed by Cole at al., 2012), the basal ganglia, mapped to Procedural Memory (as proposed by Anderson, 2007), or the temporal lobe, mapped to Long Term Memory (as proposed by Visser et al., 2012).

The "Hierarchical" family of models proposes an alternate structure. wherein brain connectivity implements hierarchical levels of processing that initiate with Perception and culminate with Action (Fig. 1D). Networks in this family conceptualize the brain as a feedforward neural network model in which different regions perform progressively greater levels of representational abstraction (Huntenburg et al., 2018). Three different hierarchical architectures are generated based on the relative position of the basal ganglia (mapped to Procedural Memory) in the hierarchy. Specifically, the basal ganglia can be placed between perception and long-term memory (as in models of procedural categorization: Kotz et al., 2009; Seger et al., 2008), between long-term memory and working memory (as in models of memory retrieval: Scimeca & Badre, 2012), or between working memory and action (as in models of action selection: Houk et al., 2007).

Broadly speaking, the CMC can be considered as a "Hub-and-Spoke" structure, using Working Memory (mapped to the Prefrontal Cortex) as the "Hub" ROI, with an additional direct connection between Perception and Action.



Figure 1: (A) The Common Model of Cognition (CMC); (B) Proposed associations between components and anatomical brain regions. (C) Three variations of Hub-and-Spoke (HUB) models, and of (D) Hierarchical (HIER) models. Arrows: dark blue, connections present in both CMC and candidate models; red, connections unique to candidate models; and dotted, connections present in CMC and absent in candidate models.

Materials and Methods

The Human Connectome Project Dataset

The data used in this analysis was drawn from the Human Connectome Project (HCP), a large scale effort to collect neuroimaging data from healthy young adults. This study in particular analyzed a subset (N=168) of rsfMRI data exclusively. For each subject, 14 minutes of rest data (eyes open with fixation) were recorded prior to a run of task data collection. A second rest run was recorded after the task battery, and was not included in this analysis. Between the two collection days, each subject had a total of 28 minutes of data. Each day's data was modeled separately, and then combined in the final analysis.

Data Processing and Analysis

Image Acquisition and Preprocessing. MRI images were acquired and minimally preprocessed according to HCP guidelines (Barch et al., 2013; Van Essen et al., 2013). Scans were taken on a 3T Siemens Skyra using a 32-channel head coil with acquisition parameters set at TR = 720 ms, TE = 33.1 ms, FA = 52° , FOV = 208×180 mm. Each image contained 72 2.0mm oblique slices with an in-plane 2.0 x 2.0 mm resolution. Images were acquired with a multi-band acceleration factor of 8X. These raw images then underwent minimal preprocessing including unwarping, motion realignment, and normalization to the standard MNI template. From there, the images were then smoothed with an isotropic 8.0 mm full-width half maximum Gaussian kernel.

Simulated Task Events Both general linear modeling (GLM) and DCM analysis require a design matrix that specifies the timing of external events that drive brain activity. Traditionally, these events are task related; the onset or absence of some stimuli. Rest data, by contrast, is collected without any specific task structure, and the recorded activity must be driven by internal and unobservable patterns. Following Di and Biswal's method (2014), a series of slow oscillatory waves of different frequencies were created as input "events" that simulate background brain activity (Fig. 2A). Specifically, eight different driving waves were generated as sine and cosine waves with frequencies of 0.01, 0.02, 0.04, and 0.08 Hz, respectively. The frequencies of these oscillations capture the canonical frequency range (0.1 -0.01 Hz) of spontaneous fluctuations in brain activity (Fox et al., 2005). An event is considered to be occurring during the positive cycle of the wave. A second assumption concerns how different events affect the different regions. In task-based DCM analysis, it is possible to make reasonable assumptions about which regions are affected by which events, such as the presentation of visual stimuli affecting a perceptual region. Di and Biswal (2014) explored a subset of possible regressor-by-region combinations to determine the most appropriate. Here, we followed the procedure of Ketola et al (2020) and let each region be potentially affected by each oscillatory regressor (Fig. 2B). Note that, while being the most general approach, this method goes against our hypotheses that spontaneous brain activity would follow a structured architecture, as it gives every region the greatest opportunity to have its time

series modeled by external inputs rather than by the network effects of other regions.



Figure 2: (A) Oscillatory waves (dotted lines) translated into "box-car" plots of events (solid lines). (B) Each event used as drivers for activity in all ROIs.

GLM. A GLM analysis was carried out to define the event matrix \mathbf{x} that is used in the DCM equation (Eq. 1). In task based DCM, events were differentiated by type and served as input to specific regions of interest. Since the resting state does not have any tasks, and the artificial events were used to capture background activity patterns, all "events" were used as direct inputs to all regions of interest (Fig. 2).

Regions of Interest Definition. Previous DCM analyses relied on task-based activity to define specific regions for each model component, but in the absence of a task structure for rest data, an alternate method was needed to determine regions of expected activity. Initial region masks were created using NeuroSynth (www.neurosynth.org), a platform that combines the results of thousands of published fMRI results and produces meta-analysis images of activity associated with various higher level conceptual category terms. For each of the five model components of the CMC model, a corresponding term was chosen from NeuroSynth's database, and a summary statistical mask was produced for each term, with each voxel having an associated Z value representing the probability that the voxel would show up a study associated with the term. These individual masks, however, were large and produced significant overlap when combined, meaning that activity in a particular voxel could belong to more than one region. To solve this problem, two thresholds were applied to the original masks, one height threshold applied to each individual voxel statistic and a minimal extent threshold applied to each cluster size. Both thresholds were calculated proportionally for each region, i.e. as a proportion of the highest Z-score and of the largest cluster within an image, respectively. The proportional adjustment was done to prevent regions with large clusters and high statistics, like perception, from overtaking regions with comparatively low Z score levels, like procedural memory. The Nelder-Mead (1965) optimization algorithm was then applied to find thresholds in the two-parameter space that would produce the largest possible regions without any overlapping voxels. The final values identified by the Nelder-Mead

algorithm were a proportional height threshold of 0.5359, and a proportional extent threshold of 0.4164. The final masks are shown in Figure 3.



Figure 3: Final regions of interest derived from Neurosynth activity masks. Individual ROIs were selected from most active voxels within these areas for each subject.

Model Fitting. 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 model (Fig. 1). The predicted neural activity for each model was then calculated using Equation 1, and the predicted time course of BOLD signal was then generated by applying а biologically-plausible model of neurovascular coupling to the simulated neural activity of each region. All of the model parameters were estimated through an 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.

Model Comparison. The models were compared on the basis of their likelihood function L(m | x). A model's likelihood is the probability of it producing the observed data x; that is, L(m | x) = P(x | m). Group-level likelihood values for a model m can then be expressed as the product of the likelihood of that model fitting each participant p, i.e., $\prod_p L(m | x_p)$. The log-likelihood is the sum of all of the individual log-likelihoods: $\sum_p \log L(m | x_p)$. Although more sophisticated model comparison procedures have been proposed (e.g., Stephan et al., 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

Regressor Quality Analysis

We first conducted a GLM analysis to ensure that our oscillatory regressors successfully captured brain activity. To do so, we calculated an omnibus ANOVA across all oscillatory regressors at the participant level. This test captures any variance that can be accounted for by any of the oscillatory regressors. The resulting *F*-statistic map was then log-transformed, yielding a measure of the

difference between the variance explained by regressors and the residual variance (i.e., noise). Finally, a performed group-level T-test was on the individual-specific log-transformed F-maps. The result of this analysis is a statistical test of whether the variance captured by the regressors was significantly greater than the variance of the residuals. The results are shown in Figure 4, thresholded at a value of t(160) > 5.212, which corresponds to p < 0.05 when corrected for multiple comparisons through the Family Wise Error correction procedure.

As Figure 4 shows, most of the grey matter voxels exhibit oscillatory activity that was captured by our regressors. Importantly, the significant voxels encompass regions in all of our predefined ROIs, including the medial temporal lobes (long-term memory ROI in Figure 3, visible in the coronal section of Figure 4) and the subcortical basal ganglia (procedural memory ROI in Figure 3, visible in the axial and sagittal sections of Figure 4), which are notoriously affected by lower signal-to-noise ratios in high-density neuroimaging protocols.



Figure 4: *T*-test showing voxels whose brain activity was significantly captured by the oscillatory regressors.

Comparison of Architectures

Each subject had two sessions of rsfMRI data, collected on two separate days. Each session was modeled individually, and then both sessions were combined on a subject level for the comparison analysis. Figure 5 illustrates the group-level log-likelihoods of the models in the rest condition. The figure presents *relative* log-likelihoods: the lowest log-likelihood is subtracted from all the others. As a result, the worst-fitting model always has a relative log-likelihood value of zero, with the best fitting model having the highest positive value.

Across both sessions, the CMC provides the best account of resting state brain activity, when compared against each of the six alternate structures. Because log-likelihood is not sensitive to model complexity, it is common to compute log likelihood in some penalized form. For example, the common Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) penalize likelihood by the number of parameters. Both measures assume, however, that parameter values are independently distributed, which is not the case for DCM models (for example, connectivity values for the same node tend to be correlated). For this reason, it is common to use a different, penalized form of likelihood known as Free Energy (Penny et al., 2012), which accounts for non-independent parameters. The values reported in Figure 5 depict this penalized form of likelihood, and thus already account for varying model complexity.



Figure 5: The log-likelihood of the CMC architecture compared to six alternate architectures across both sessions of rsfMRI data.

Analysis of Bayes Factors

Although the evidence in favor of the CMC is apparent, one might wonder exactly how significant the difference in log-likelihood is. To express log-likelihood in an interpretable form, we will use Bayes Factors (*BF*). The $BF_{1,2}$ between two models m_1 and m_2 is defined as:

$$BF_{1,2} = P(m_1 \mid \mathbf{x}) / P(m_2 \mid \mathbf{x})$$

In other words, the value of $BF_{1,2}$ represents the odds of model 1 fitting the data better than model 2. Given the definition of likelihood as $L(m|\mathbf{x}) = P(\mathbf{x}|m)$, $BF_{1,2}$ can be expressed as

$$BF_{12} = e^{\Delta t}$$

where $\Delta L = \log L(m_1 | x) - \log L(m_2 | x)$ is the difference in log-likelihoods between model 1 and model 2. As a guideline, Kass and Raftery (1995) suggest that values of BF > 20 correspond to a value of p < .05 in a canonical null-hypothesis test and provide "strong" evidence in favor of model 1 over model 2, while values of BF > 150 provide "very strong" evidence. All of the *BF* values for the comparisons of the CMC against all the other models exceeded 10^{250} , indicating that the evidence in favor of the CMC is, in fact, overwhelming.

Random-Effects Analysis

Although the results provide strong evidence in favor of the CMC, it should be noted that they are not directly comparable with the model comparison approach reported by Stocco et al. (2021). In the original paper, the authors compared the different architectures by measuring the relative probabilities that each architecture would fit any given participant (Stephan et al., 2009). This approach is conceptually different from the log-likelihood approach because it is based on relative, rather than absolute, fit to the data and because participants are considered as a random factor, thus giving different architectures the opportunity to fit different subgroups of participants.

To provide a better comparison to the original findings, we replicate the analysis method of Stocco et al. (2021) with the current resting-state data. The results are reported in Figure 6. In the figure, the curves represent the densities of the relative probabilities that each architecture would fit a participant. The superiority of the CMC is shown by the fact that its probability density function lies to the right of all other architectures. Architectures can be quantitatively compared in terms of exceedance probabilities, i.e. the probability that a point randomly sampled from their density distributions would have a higher probability than any other architectures. In this case, the Common Model had an exceedance probability of 96.4%, further confirming its superiority.



Figure 6. Probability densities that each architecture would best fit the data from a participant in our sample.

Discussion

The major finding of this paper is the apparent presence of an underlying structure of brain connectivity that predicts activity even during undirected and task free behavior. The implications of these results are broad.

First, they demonstrate the success in adapting a traditional DCM analysis to resting state through the use of simulated task events and externally generated ROIs, paving the way for future explorations of resting state data. In particular, the use of summarized fMRI data to determine ROIs presents the opportunity to explore increasingly complex model structures involving more specific brain areas. While the CMC provides the best

account of underlying connectivity, it remains only the best model of those that we have tested so far, and is deliberately composed of a few, high level components. The use of more localized ROIs opens the door to examining each component in greater detail; separating visual perception from auditory perception, for example, or decomposing the long term memory component into semantic and episodic memory. The DCM framework also allows models to account for modulatory connections between regions, which, while not used in this paper, provide further opportunities to define and specify a general purpose framework of cognition.

While specific ROIs will always differ slightly across subjects and tasks, the localized ROIs used in the present study represent a much smaller search space than the broad parcellation used to define ROIs in the original CMC study. The data from the original study should be reanalyzed using the more specific maps to ensure that the findings still hold, with the ultimate aim of defining more exact regions that can be used in future analyses.

The goal behind this study was to test large-scale architectures in a task-free paradigm, using only signals originating from spontaneous neural activity that would capture the intrinsic organization of the brain (Fox et al., 2005). Although this procedure has become the accepted standard in neuroimaging research, the extent to which resting state activity is truly spontaneous remains debated: even at rest, participants do typically engage in some form of thought, such as daydreaming or mind wandering. A recent computational model of mind wandering (Taatgen et al., 2021), for instance, argues that mind-wandering is generated by the occasional intrusion of task-unrelated goals and that, when activated, it triggers a cascade of mental processes, such as memory retrieval, that are similar to those required by canonical tasks. The fact that the same architecture that was found to best capture brain activity across multiple cognitive tasks (Stocco et al., 2021) also explains brain activity at rest supports the assumptions of this model and the idea that spontaneous thought follows the same patterns as task-directed thought.

The implications of a general framework for cognition that remains persistent in the resting state will significantly increase its applicability to other domains, such as computational psychiatry and neurology.

Acknowledgements

This research was supported by grant FA9550-19-1-0299 from the Air Force Office of Scientific Research to AS.

References

- Anderson, J. R. (2007). How can the human mind occur in the physical universe? *Oxford University Press*.
- Cole, M. W., Yarkoni, T., Repovš, G., Anticevic, A., & Braver, T. S. (2012). Global connectivity of prefrontal cortex predicts cognitive control and intelligence.

Journal of Neuroscience, 32(26), 8988-8999.

- 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.
- 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.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273–1302.
- Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., ... & Petersen, S. E. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron*, 98(2), 439-452.
- Hohenfeld, C., Werner, C. J., & Reetz, K. (2018). Resting state connectivity in neurodegenerative disorders: Is there potential for an imaging biomarker? *NeuroImage: Clinical*, 18, 849–870.
- Houk, J. C., Bastianen, C., Fansler, D., Fishbach, A., Fraser, D., Reber, P. J., ... & Simo, L. S. (2007). Action selection and refinement in subcortical loops through basal ganglia and cerebellum. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1573-1583.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. Journal of the American Statistical Association, 90(430), 773–795.
- Ketola, M., Thompson, S., Madhyastha, T., Grabowski, T., & Stocco, A. (2020). Applying the Common Model of Cognition to resting-state fMRI Leads to the identification of abnormal functional connectivity in Parkinson's Disease. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society.*
- Kotz, S. A., Schwartze, M., & Schmidt-Kassow, M. (2009). Non-motor basal ganglia functions: A review and proposal for a model of sensory predictability in auditory language perception. *Cortex*, 45(8), 982-990.
- 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.
- Nelder, J.A. & Mead, R. (1965). A simplex method for function minimization. *Comput. J.*, 308–313.
- Penny, W. D. (2012). Comparing dynamic causal models using AIC, BIC and free energy. *Neuroimage*, 59(1), 319-330.
- Scimeca, J. M., & Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron*, 75(3),

380-392.

- Seger, C. A. (2008). How do the basal ganglia contribute to categorization? Their roles in generalization, response selection, and learning via feedback. *Neuroscience & Biobehavioral Reviews*, 32(2), 265-278.
- Shehzad, Z., Kelly, A. C., Reiss, P. T., Gee, D. G., Gotimer, K., Uddin, L. Q., ... & Milham, M. P. (2009). The resting brain: unconstrained yet reliable. *Cerebral cortex*, 19(10), 2209-2229.
- 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., 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
- Taatgen, N. A., van Vugt, M. K., Daamen, J., Katidioti, I., Huijser, S., & Borst, J. P. (2021). The Resource-availability Model of Distraction and Mind-wandering. Cognitive Systems Research.
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K., & WU-Minn HCP Consortium. (2013). The Human Connectome Project: An overview. *NeuroImage*, 80, 62–79.
- Visser, M., Jefferies, E., & Lambon Ralph, M. A. (2010). Semantic processing in the anterior temporal lobes: a meta-analysis of the functional neuroimaging literature. *Journal of cognitive neuroscience*, 22(6), 1083-1094.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Smoller, J. W., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H., & Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 106(3), 1125–1165.