Deriving an Architecture from Brain Data

Andrea Stocco[1], Zoe Steine-Hanson[2], Natalie Koh[3], John Laird[4], Christian Lebiere[5], Paul Rosenbloom[6]
[1] University of Washington, [2] Oregon State University
[3] Northwestern University, [4] University of Michigan
[5] Carnegie Mellon University, [6] University of Southern California

Testing architectures with brain data

Task-based approach:

1. Find a task that looks different in two architectures
2. Find a way to derive **neural** predictions
3. **Compare** observed vs. predicted data

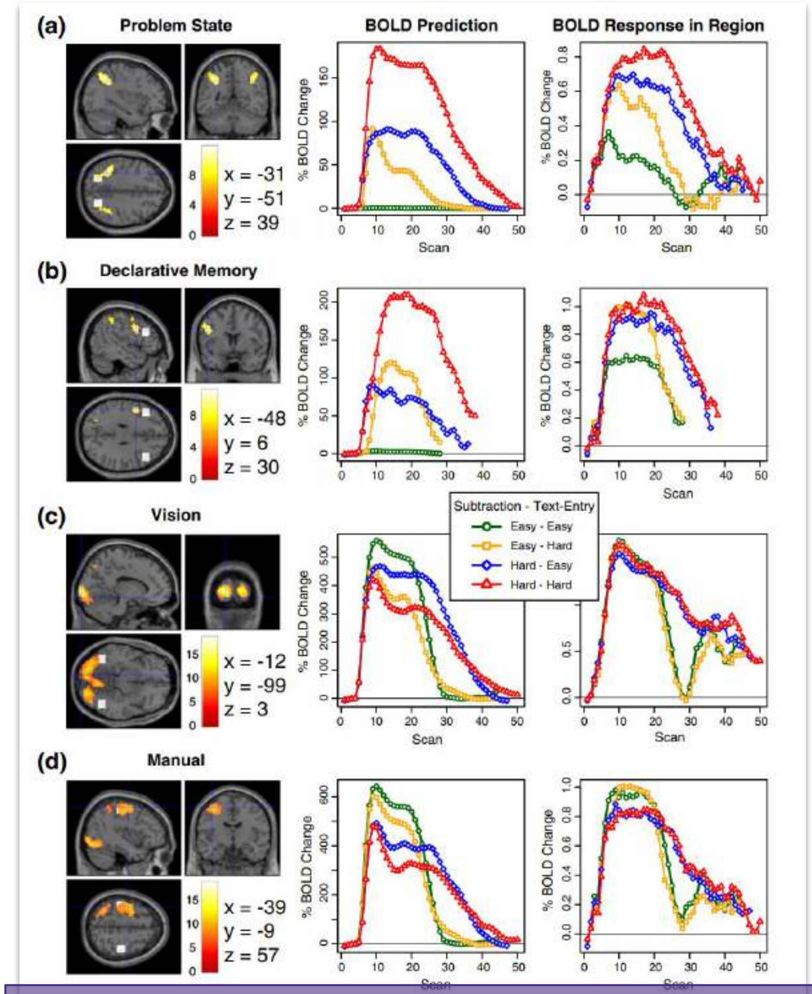

Borst, Van Rijn, Taatgen, *Neuroimage*, 2011

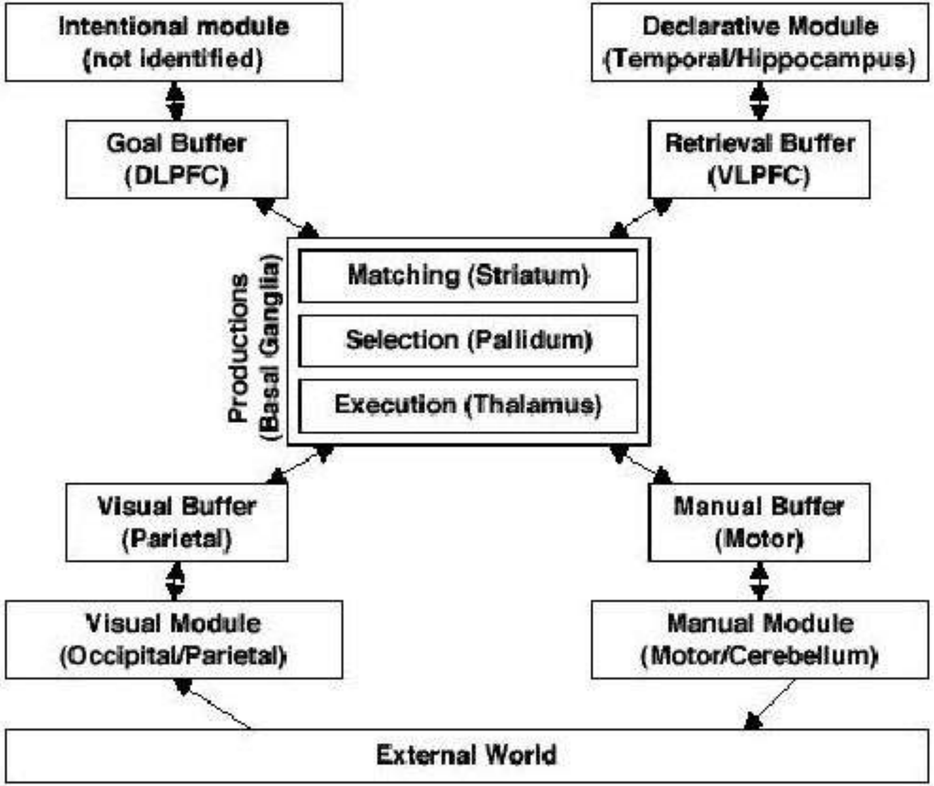
Intentional module
(not identified)
Declarative Module
(Temporal/Hippocampus)
Goal Buffer
(DLPFC)
Retrieval Buffer
(VLPFC)
Productions
(Basal Ganglia)
Matching (Striatum)
Selection (Pallidum)
Execution (Thalamus)
Visual Buffer
(Parietal)
Manual Buffer
(Motor)
Visual Module
(Occipital/Parietal)
Manual Module
(Motor/Cerebellum)
External World

Architecture: We should be able to see it

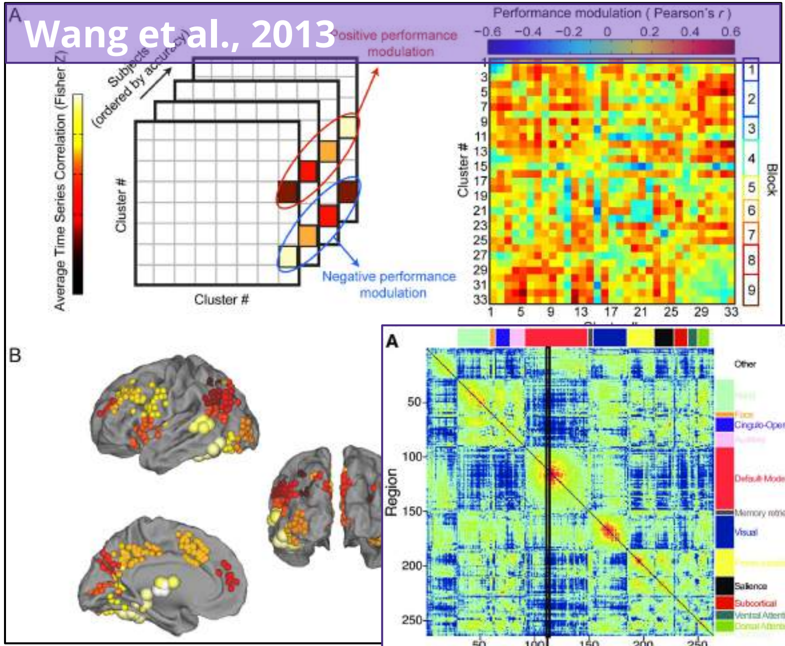

Wang et al., 2013

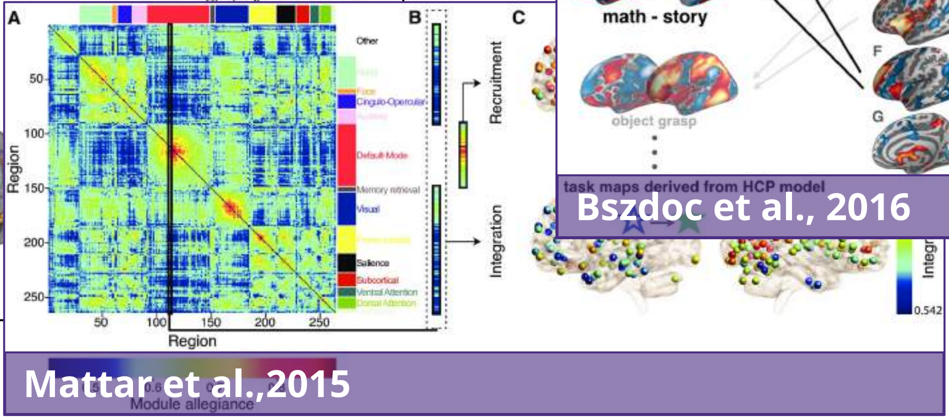

Mattar et al.,2015

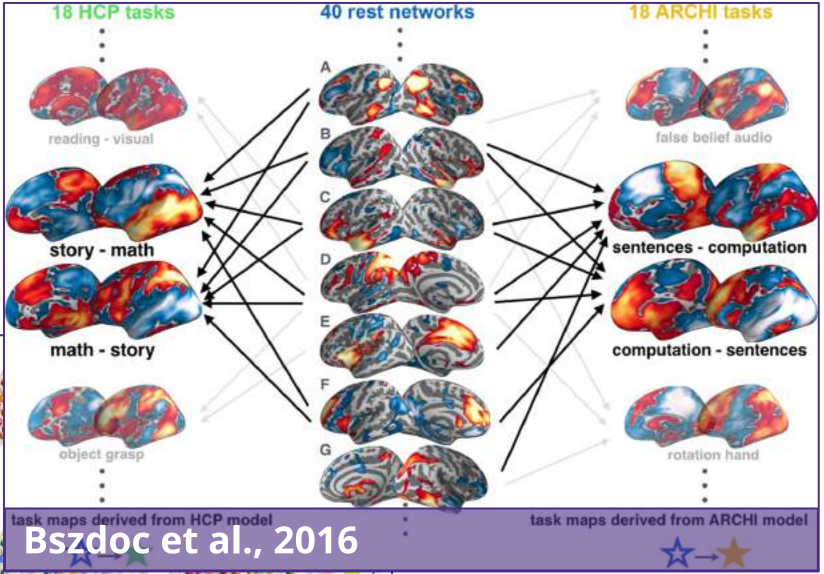

Bszdoc et al., 2016

A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics

John E. Laird, Christian Lebiere, Paul S. Rosenbloom

■ *A standard model captures a community consensus over a coherent region of science, serving as a cumulative reference point for the field that can provide guidance for both research and applications, while also focusing efforts to extend or revise it. Here we propose developing such a model for humanlike minds, computational entities whose structures and processes are substantially similar to those found in human cognition. Our hypothesis is that cognitive architectures provide the appropriate computational abstraction for defining a standard model, although the*

A mind is a functional entity that can think, and thus support intelligent behavior. Humans possess minds, as do many other animals. In natural systems such as these, minds are implemented through brains, one particular class of physical device. However, a key foundational hypothesis in artificial intelligence is that minds are computational entities of a special sort — that is, cognitive systems — that can be implemented through a diversity of physical devices (a concept lately reframed as substrate independence [Bostrom 2003]), whether natural brains, traditional general-purpose computers, or other sufficiently functional forms of hardware or wetware.

Laid, Lebiere, & Rosenbloom, 2017, *AI Magazine*

Laird, Lebiere, & Rosenbloom (2017)

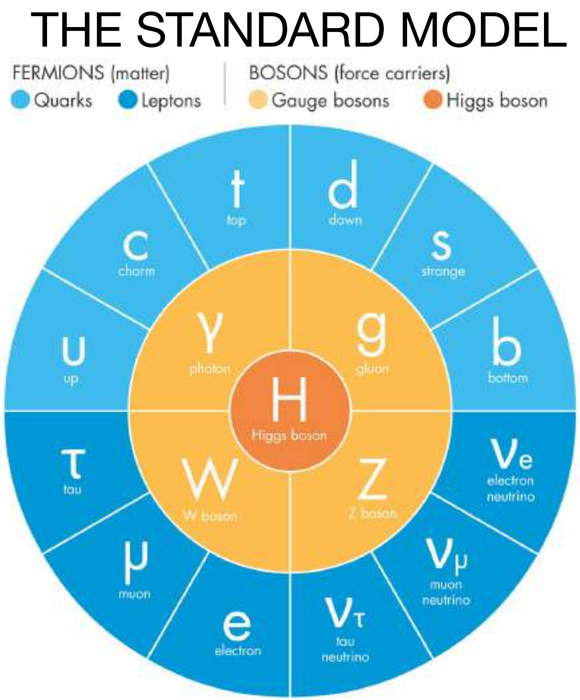

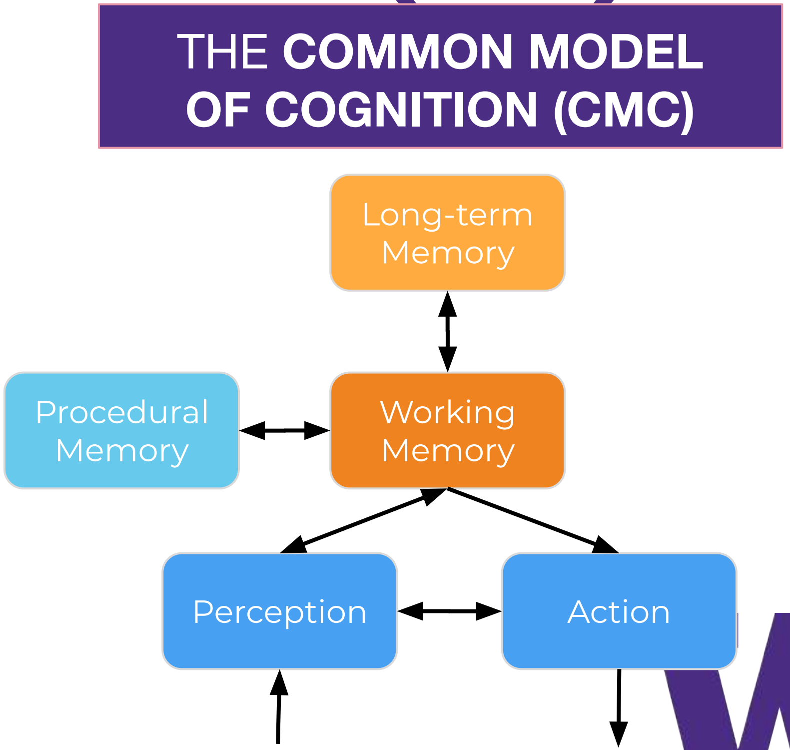

The CMC: Structural components

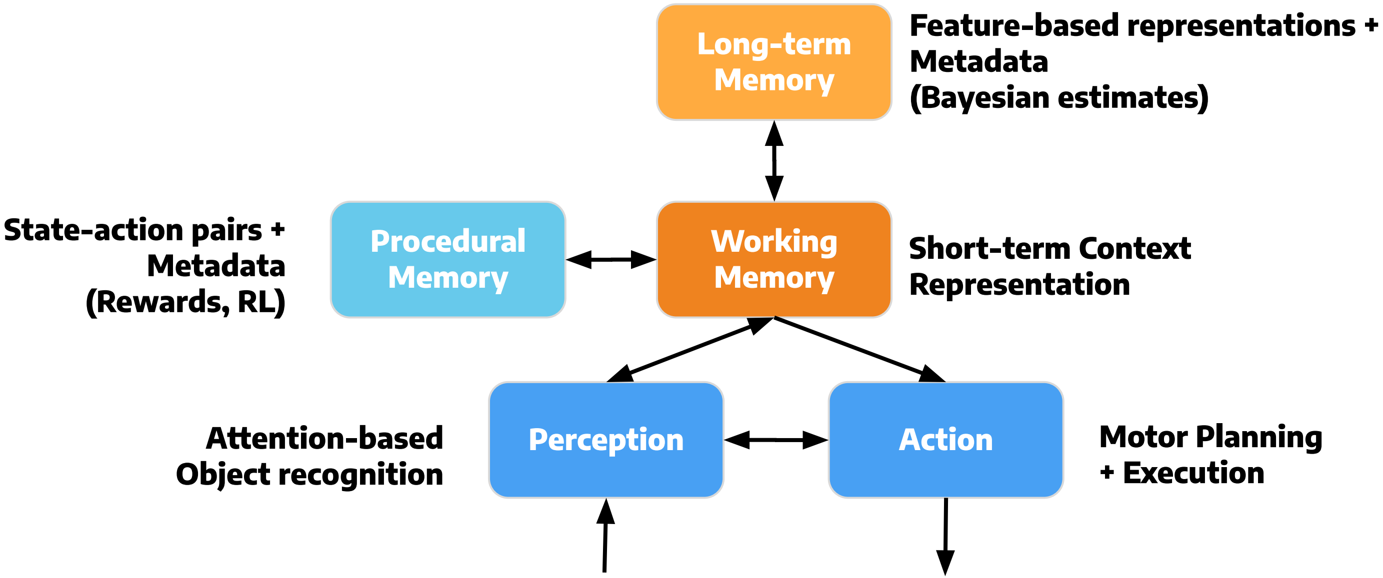

Can we see the CMC in neural data?

- The CMC makes predictions about how components are connected
- This should be reflected in patterns of functional connectivity between regions
- It should be independent of the task

The Methods

What do we need?

- Need a method to:
  - Identify architecture components with brain structures
  - Derive neural predictions from the architecture
  - Compare between alternative architectures

What do we need?

- Need a method to:
  - **Identify architecture components with brain structures**
  - Derive neural predictions from the architecture
  - Compare between alternative architectures

Large-scale component identification

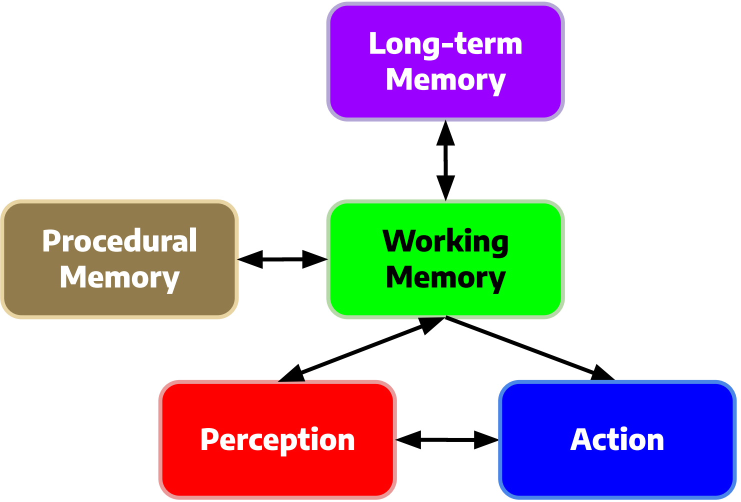

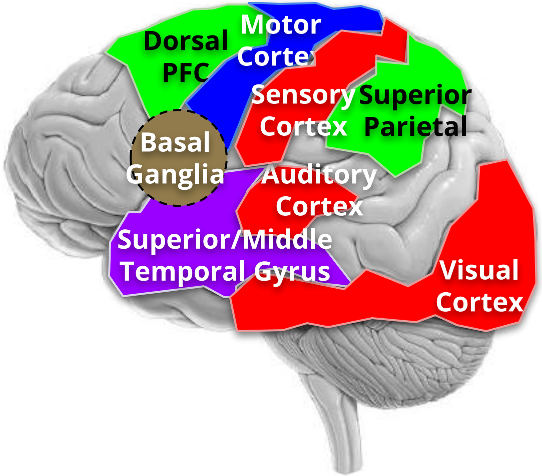

What do we need?

- Need a method to:
  - Identify architecture components with brain structures
  - **Derive neural predictions from the architecture**
  - Compare between alternative architectures

Creating a network model

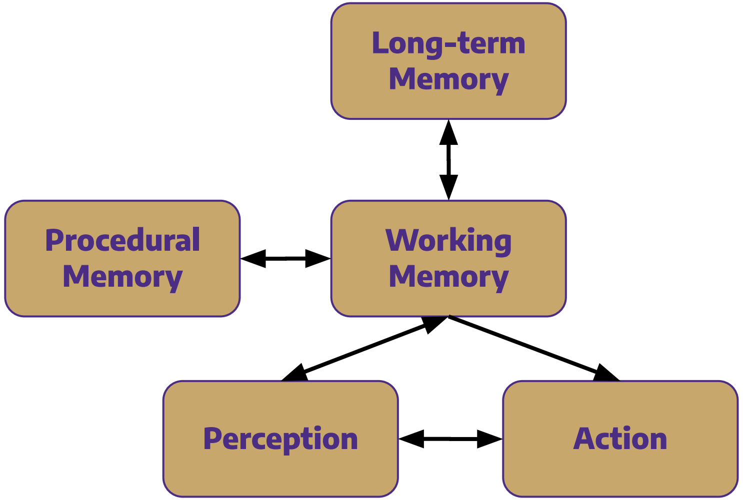

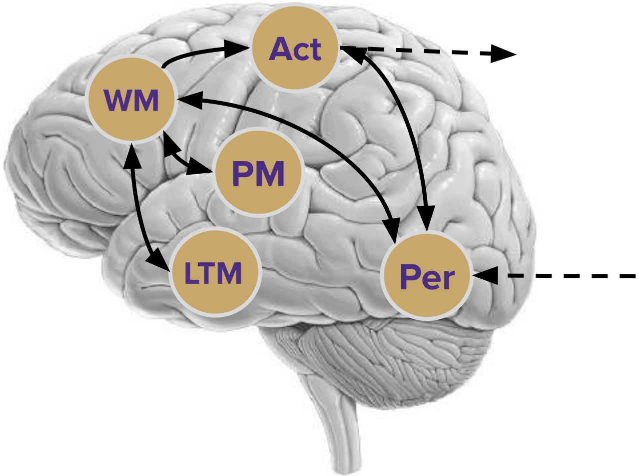

Implementing the network model: Dynamic Causal Modeling

Traditional GLM

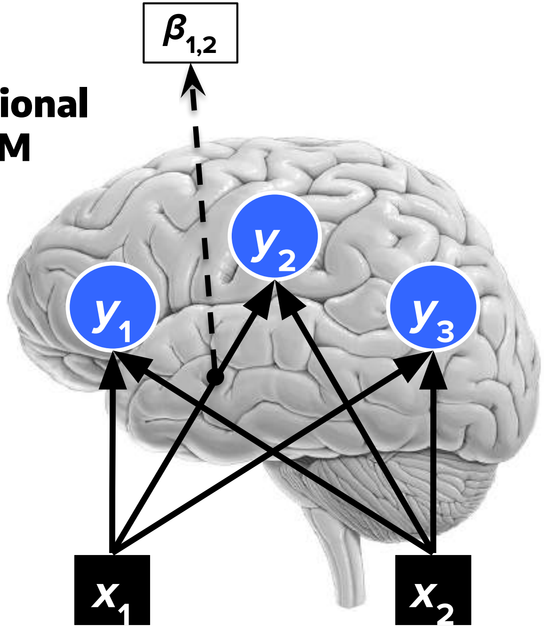

$\mathbf{y} = \Sigma_i \boldsymbol{\beta}_i{}^* \mathbf{x}_i$

Dynamic Causal Modeling

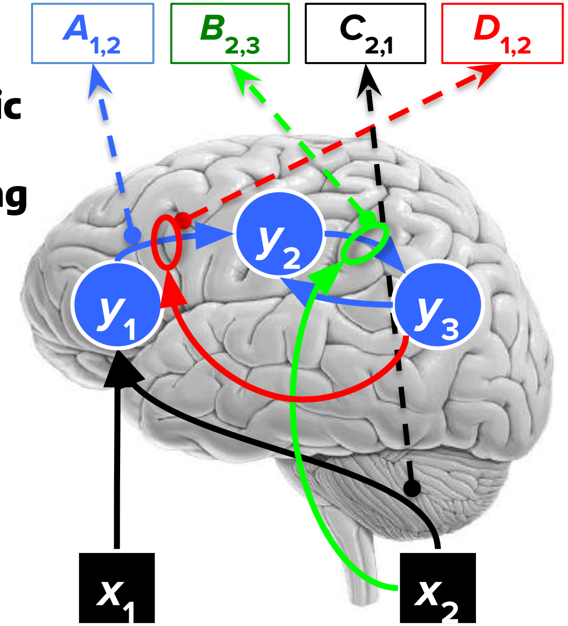

$d\mathbf{y}/dt = \mathbf{A}y + \Sigma_i x_i \mathbf{B}(i)y + \mathbf{C}x + \Sigma_j y_j \mathbf{D}(j)y$

Estimating the dynamic model

- Removing **modulatory** terms yields a **linear model**
  - $dy/dt = \mathbf{A}y + \cancel{\boxed{\Sigma_i x_i \mathbf{B}[i,,]y}} + \mathbf{C}x + \cancel{\boxed{\Sigma_j y_j \mathbf{D}[j,,]y}}$
  - $dy/dt = \mathbf{A}y + \mathbf{C}x$
- **A** and **C** estimated **through** Expectation/Maximization
- **Variational Bayes** (older but faster) instead of MCMC (newer but slower) to calculate PDFs

What do we need to test the CMC?

- Need a method to:
  - Identify architecture components with brain structures
  - Derive neural predictions from the architecture
  - **Compare between alternative architectures**

Alternatives 1 and 2: Hierarchical, recursive architecture

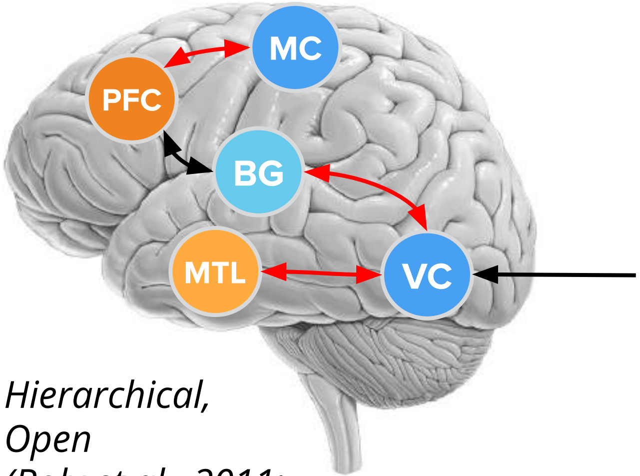

Hierarchical, Open (Boly et al., 2011; Margulis et al., 2018)

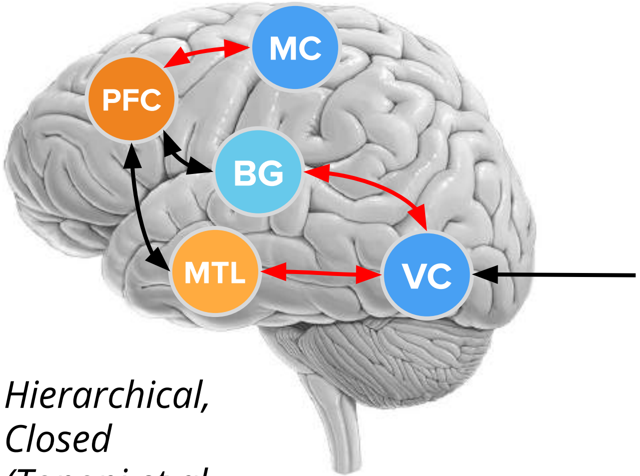

Hierarchical, Closed (Tononi et al., 2011)

Alternatives 3 and 4: Hub & Spoke brain architectures

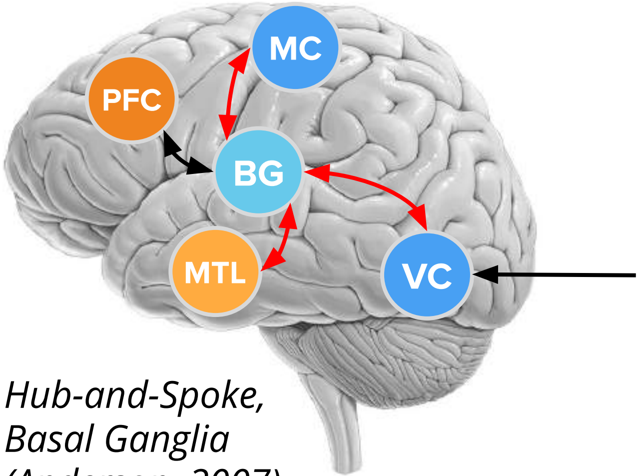

Hub-and-Spoke,
Basal Ganglia
(Anderson, 2007)

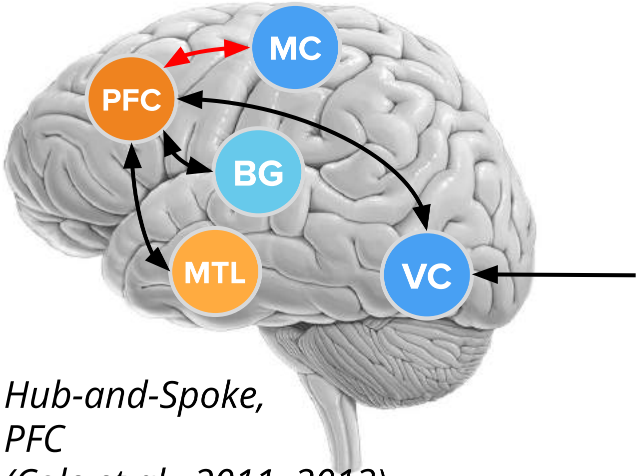

Hub-and-Spoke,
PFC
(Cole et al., 2011, 2013)

Comparing architectures

- Many criteria exists
  - AIC, BIC, Log-likelihood…
- Here, **Bayesian approach**: Posterior probability that a model is **true**, given the **data**
  - Each architectures' PDF is modeled as a **Dirichlet** distribution ~ Dir(α)
  - Two metrics: **expected** and **exceedance** probabilities.

Comparing architecture

Calculate a distribution of probabilities *q* that a model would fit participants' data

Given *q*, calculate the probabilities that each model 1, 2… *k* would **best** fit the data

Given *m*, estimate the probability that every model *k* would fit each subject 1, 2, … *n*

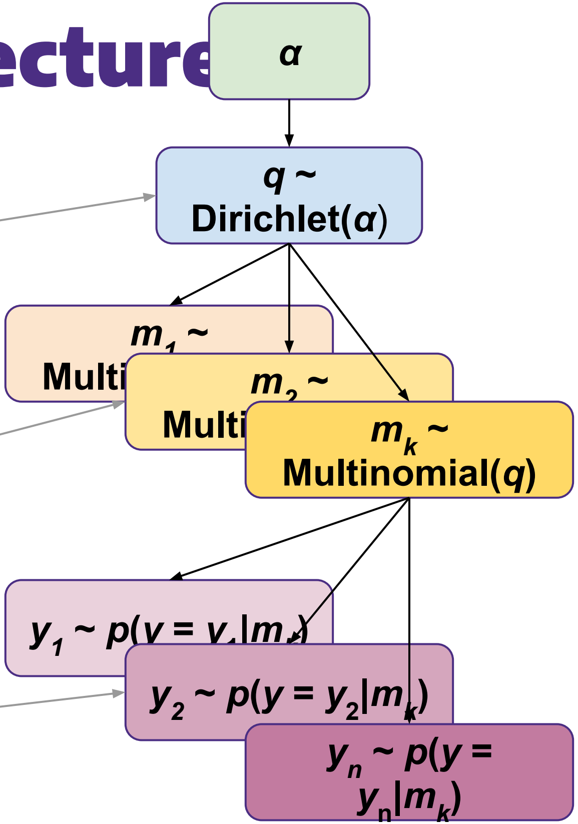

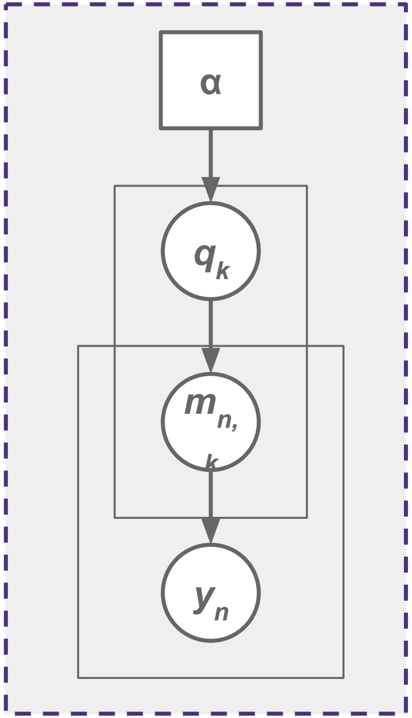

Comparing models

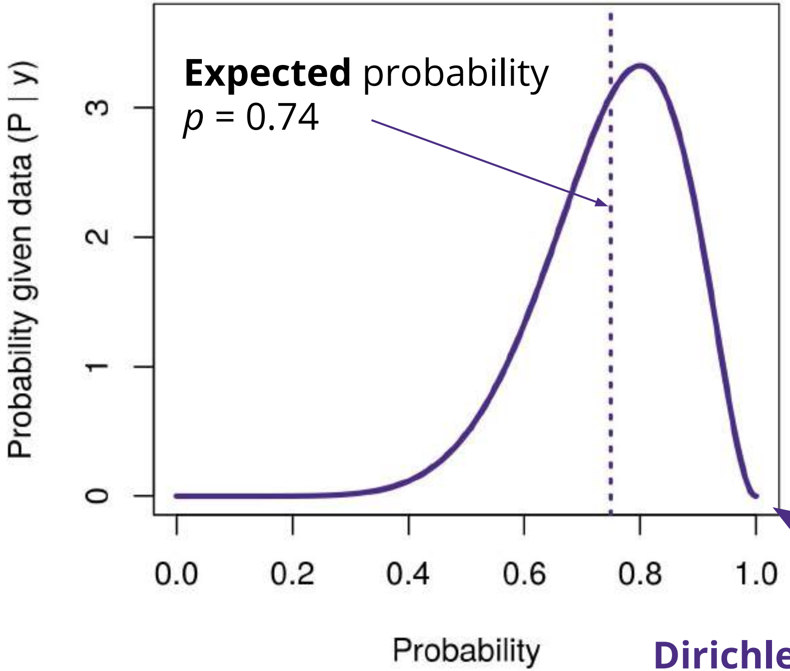

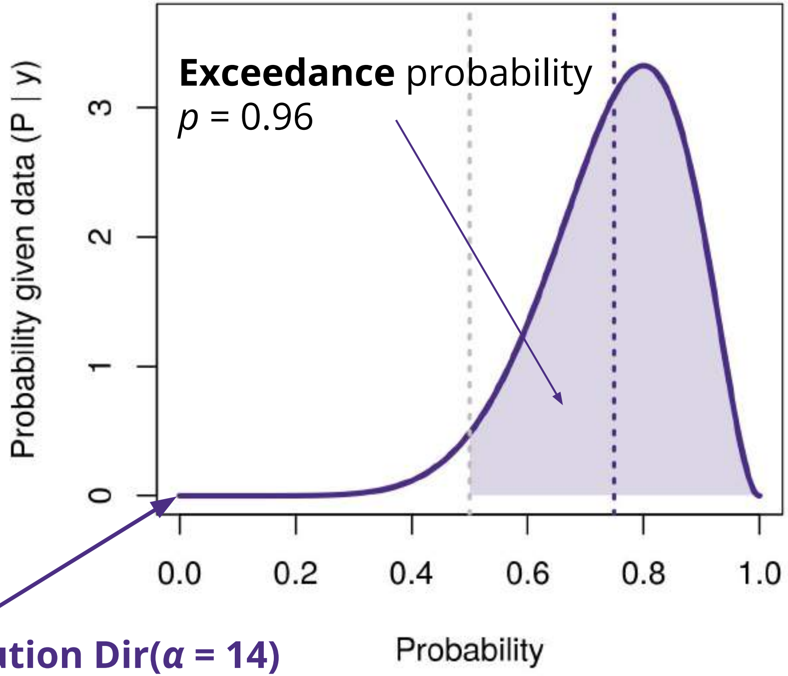

Dirichlet distribution Dir(α = 14)

The Data

The Human Connectome Project

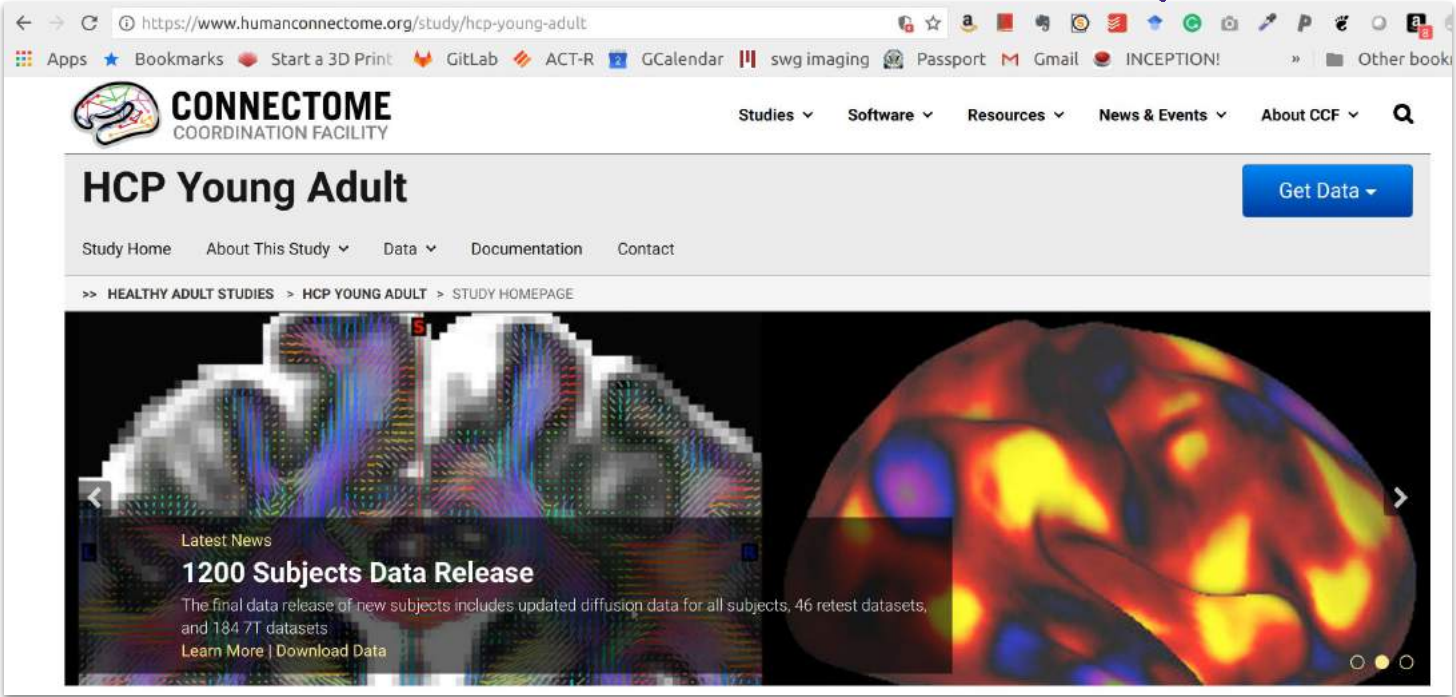

The Human Connectome Project

> Contains high-quality neuroimaging data:
  - 1,200 Adult Participants (July 2018)
  - 7 Different Tasks
  - 4 Resting State Sessions
  - fMRI + MEG data (subset)

The Human Connectome Project

- Contains high-quality neuroimaging data:
  - **200 out of** 1,200 Adult Participants
  - **6 out of** 7 Different Tasks
  - 4 Resting State Sessions
  - **fMRI** + MEG data

Siemens Skyra, Multiband					
TR	720 ms	**MB factor**	8x	***N* Slices**	72
TE	33.1 ms	**FOV**	208 x 180 mm	**Slice Gap**	0mm
FA	52°	**In-plane res**	2 x 2 mm	**Slice thick**	2mm

Tasks

Task	Reference	Description
Motor	Buckner et al. (2011)	**Hand, arm, foot, leg, voice** responses
Emotional	Hariri et al. (2002)	**Fearful faces** vs. Neutral Shapes
Gambling	Delgado et al. (2000)	**“Losing” blocks** vs. “Winning” blocks of choices
Language	Binder et al. (2011)	**Language blocks** vs. Math blocks
Relational	Smith et al. (2007)	**Relational arrays** vs. Control arrays
Social	Whitley et al. (2007)	**Interacting shapes** vs. Randomly moving
Working Memory	Dobryshevsky et al. (2006)	**2-back** vs 0-back blocks

Tasks

Task	Reference	Description
~~Motor~~	~~Buckner et al. (2011)~~	~~Hand, arm, foot, leg, voice responses~~
Emotional	Hariri et al. (2002)	**Fearful faces** vs. Neutral Shapes
Incentive	Delgado et al. (2000)	**"Losing" blocks** vs. "Winning" blocks of choices
Language	Binder et al. (2011)	**Language blocks** vs. Math blocks
Relational	Smith et al. (2007)	**Relational arrays** vs. Control arrays
Social	Whitley et al. (2007)	**Interacting shapes** vs. Randomly moving
Working Memory	Dobryshevsky et al. (2006)	**2-back** vs 0-back blocks

Analysis pipeline

HCP Data

- Emotion
- Incentive
- Language
- Relational
- Social
- WM

GLM analysis, 1st level (canonical)

Subject-level quality Control

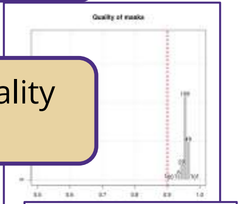

GLM analysis, 2nd Level

Group-level VOI coordinates

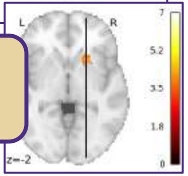

GLM analysis, 1st level (for DCM)

Individual-level VOIs coordinates

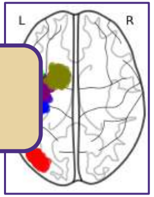

VOI data extraction (PCA)

DCM Model Estimation

- Common Model
- Hierarchical 1
- Hierarchical 2
- Hub & Spoke 1
- Hub & Spoke 2

Bayesian model comparison

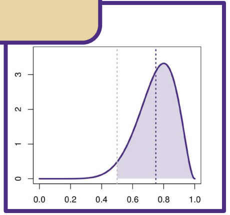

Task-Specific Regions of Interest

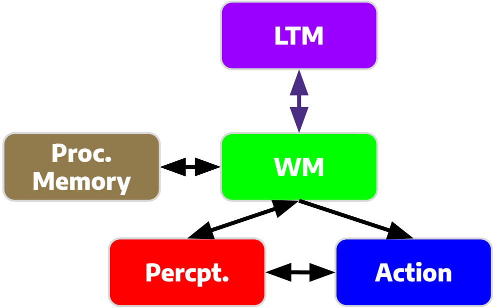

Emotion (N=187)

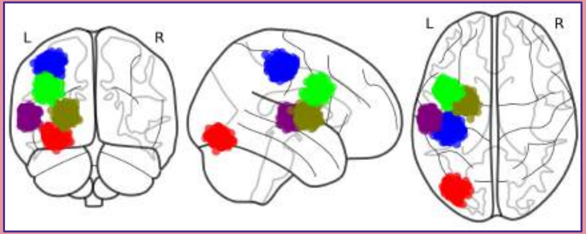

Incentive (N=199)

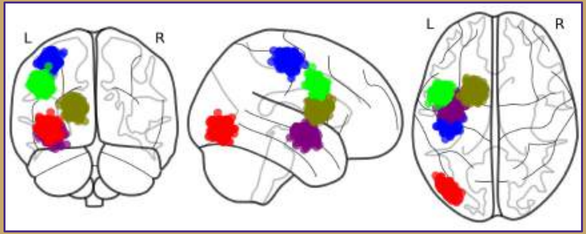

Language/Math (N=187)

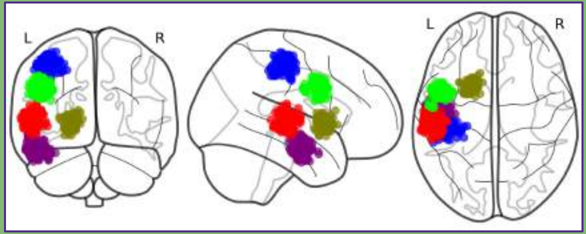

Relational (N=185)

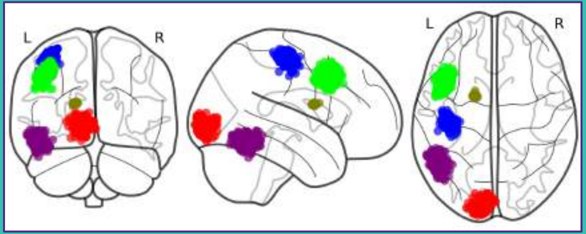

Social (N=188)

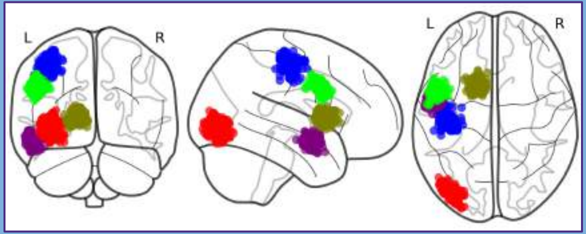

Working Memory (N=188)

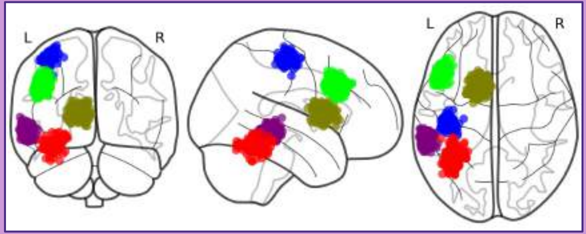

The Results

Canonical GLM analysis (Group-level)

Emotion (*N*=187)

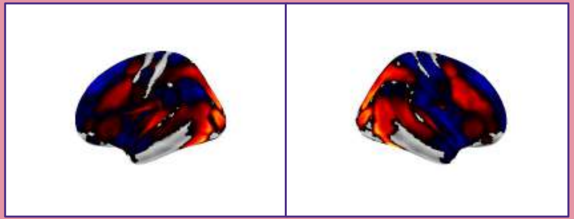

Incentive (*N*=199)

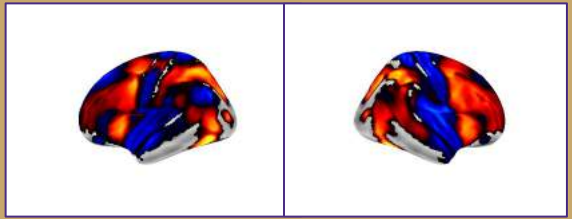

Language/Math (*N*=187)

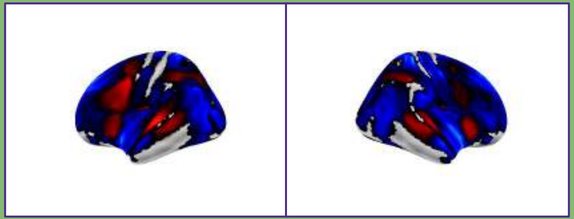

Relational (*N*=185)

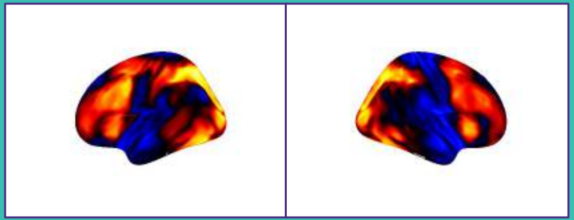

Social (*N*=188)

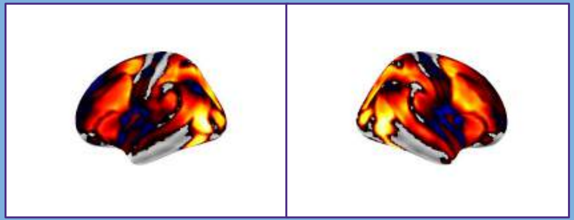

Working Memory (*N*=188)

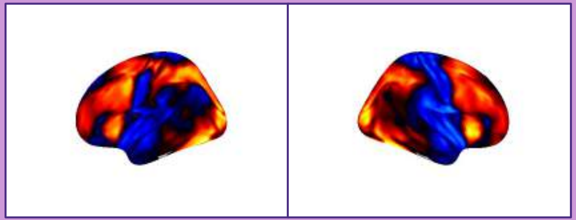

Bayesian model comparison

- Reminder
  - Estimating **posterior probability** $P(M \mid Y)$ of a model M given data Y.
- Two measures
  - Expected probability
  - Exceedance probabilities

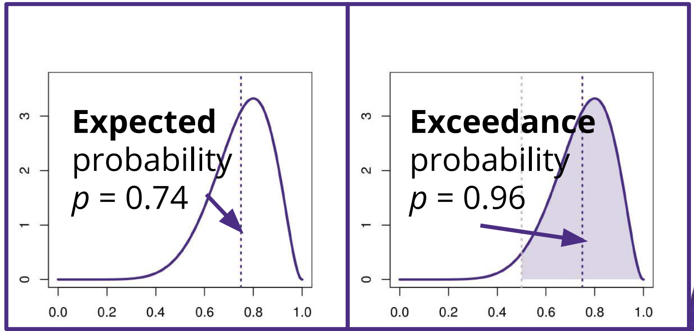

Results: Probability densities by task

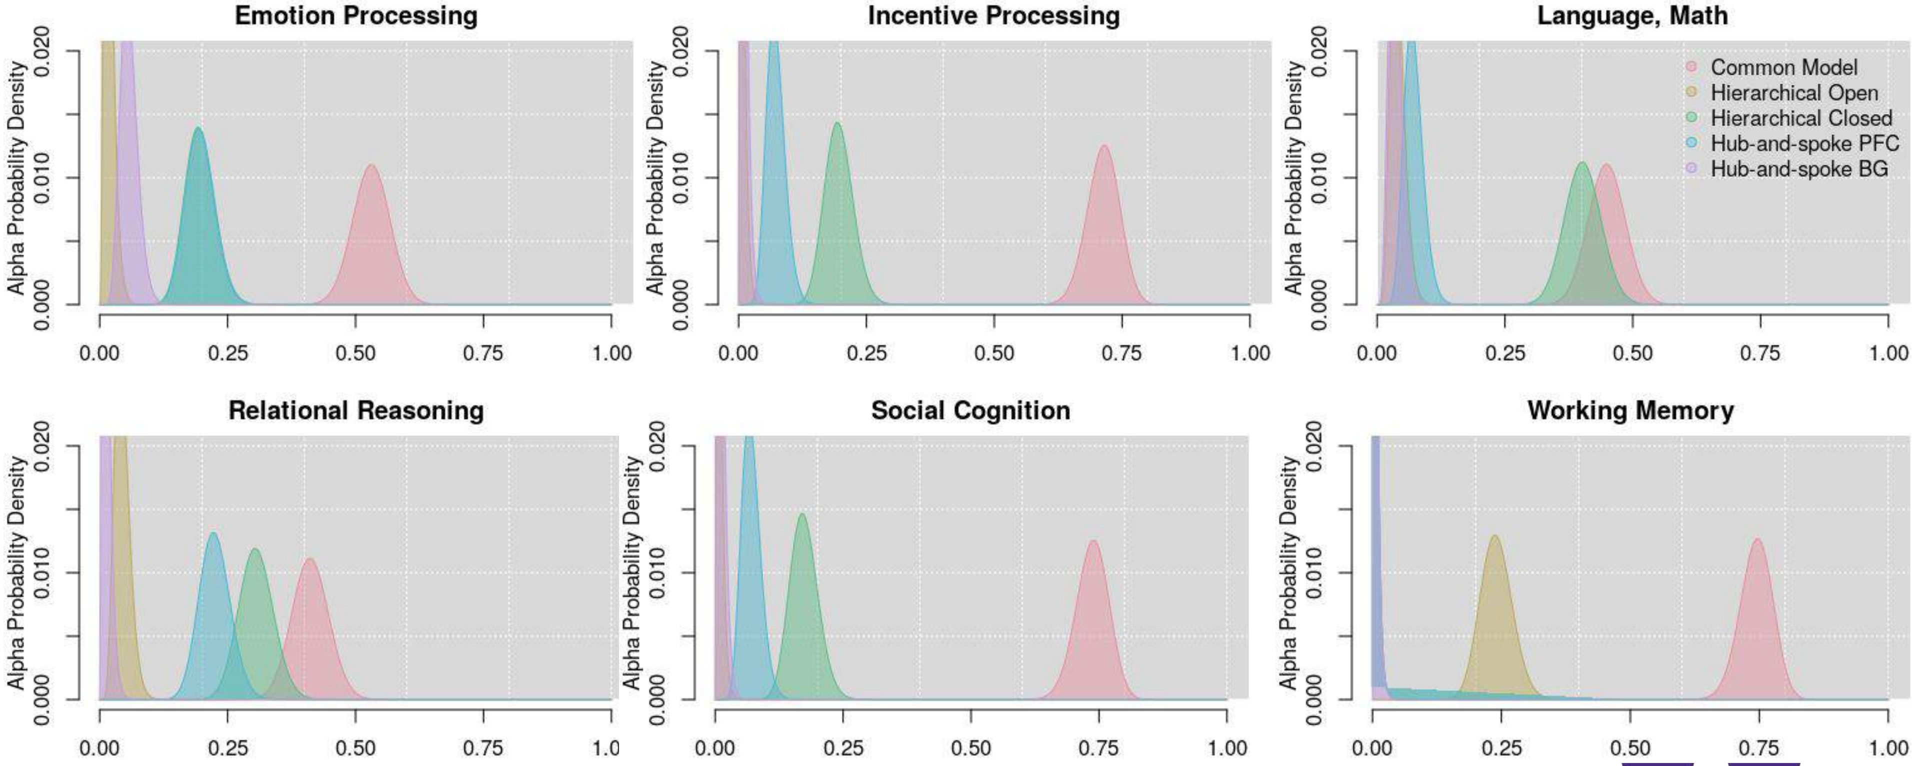

Results: Expected Probability by Task

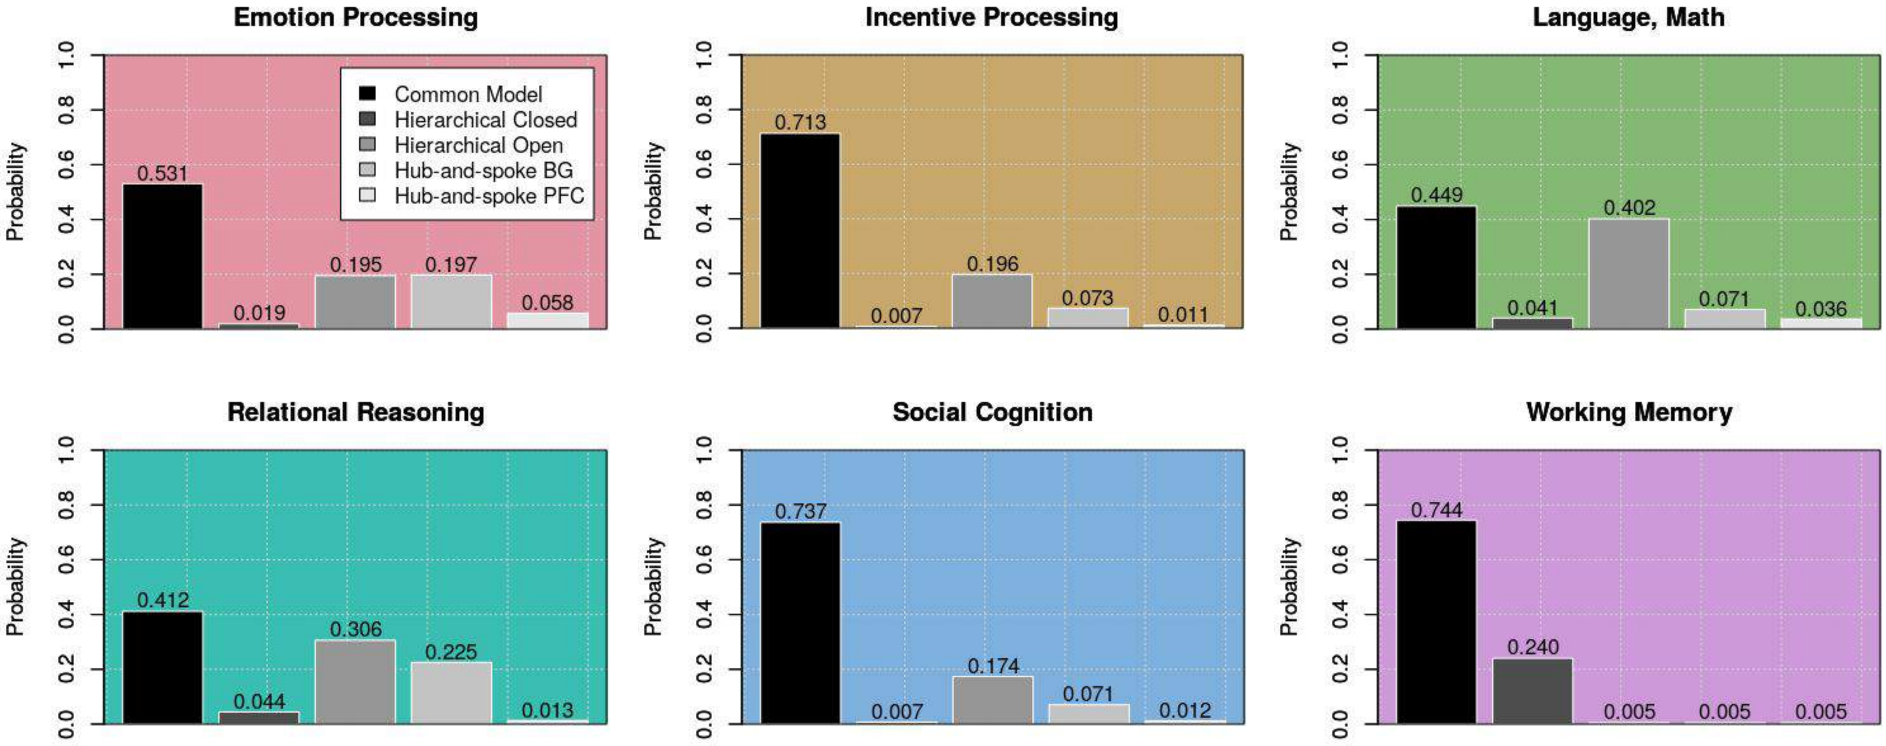

Results: Exceedance probability by task

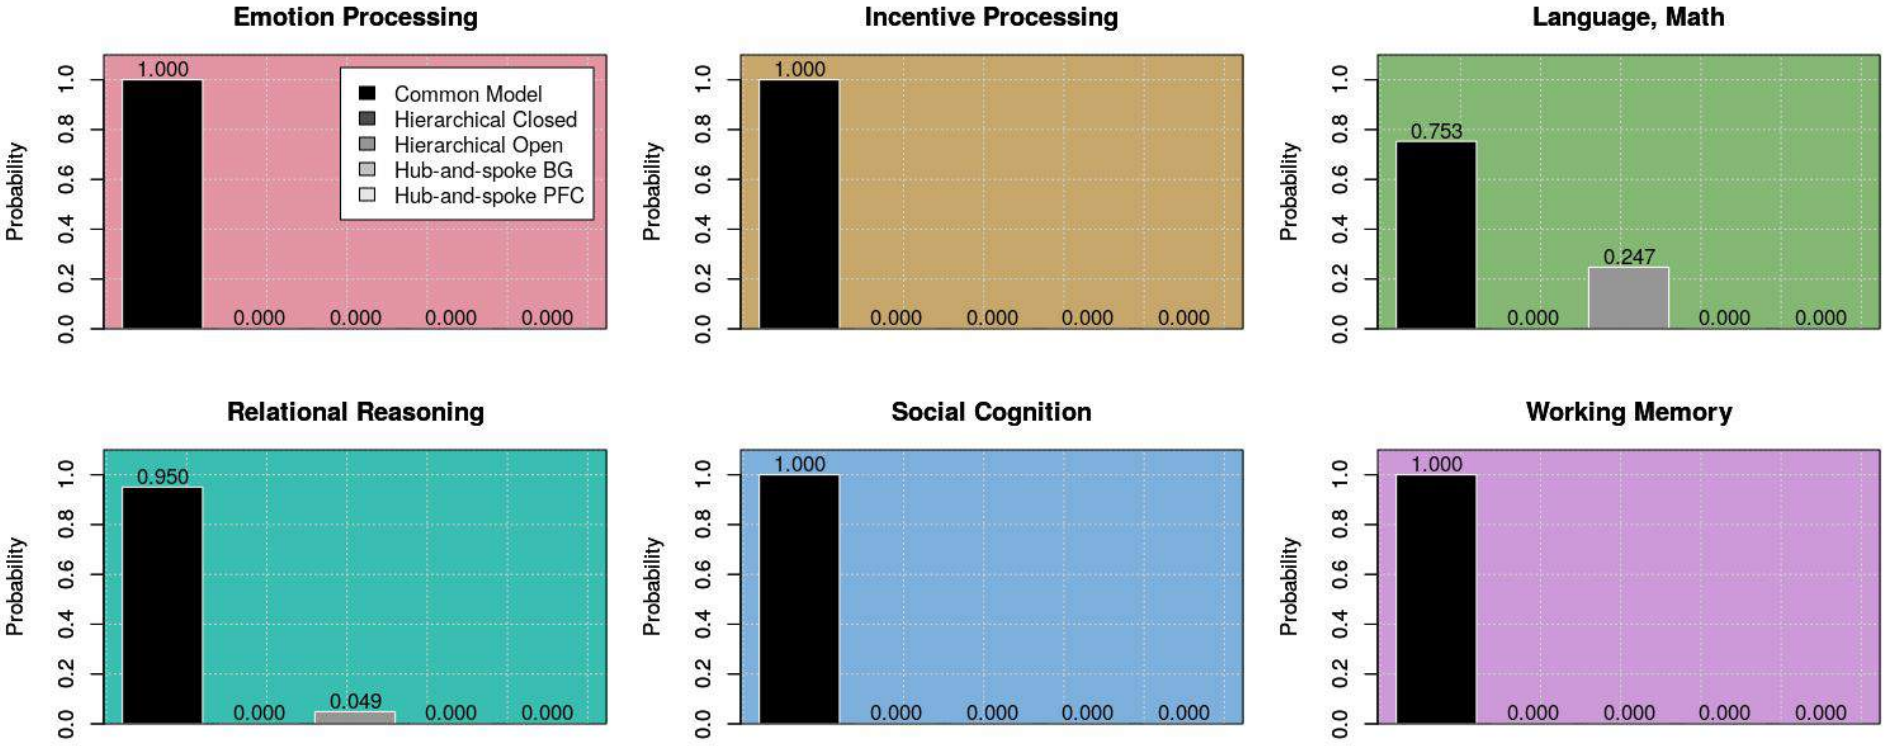

All tasks combined (repeated measures)

Expected Probability, All Tasks (*N*=168)

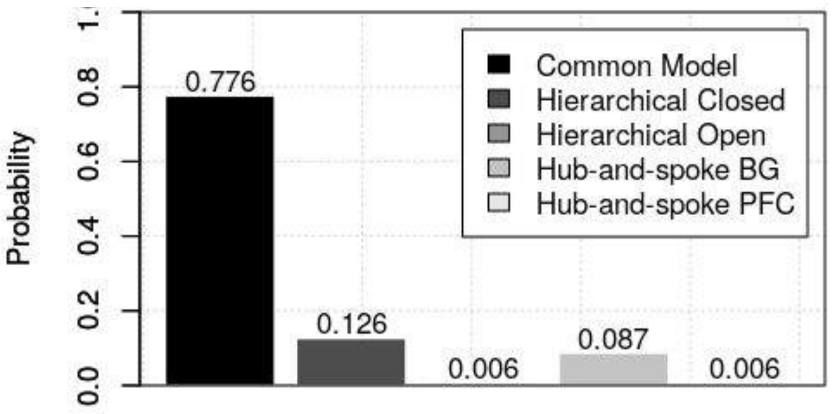

Exceedance Probability, All Tasks (*N*=168)

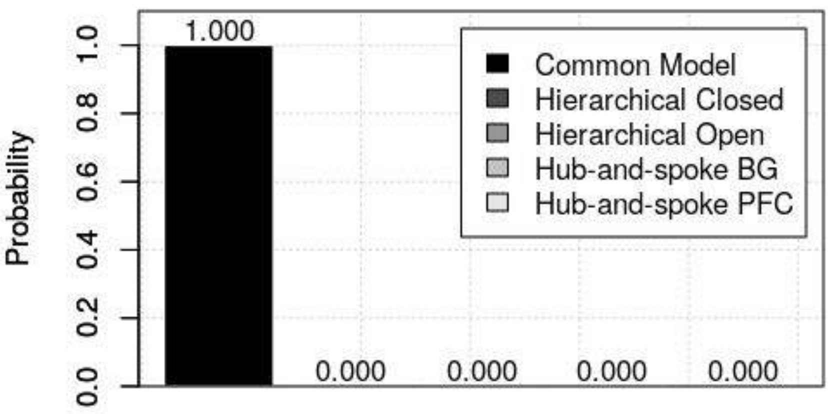

Is the CMC reasonable?

- Check connectivity parameters
  - Single model, all tasks
- All parameters are **positive**
  - Except self-connections
- All parameters are **likely**
- **Values** change by task

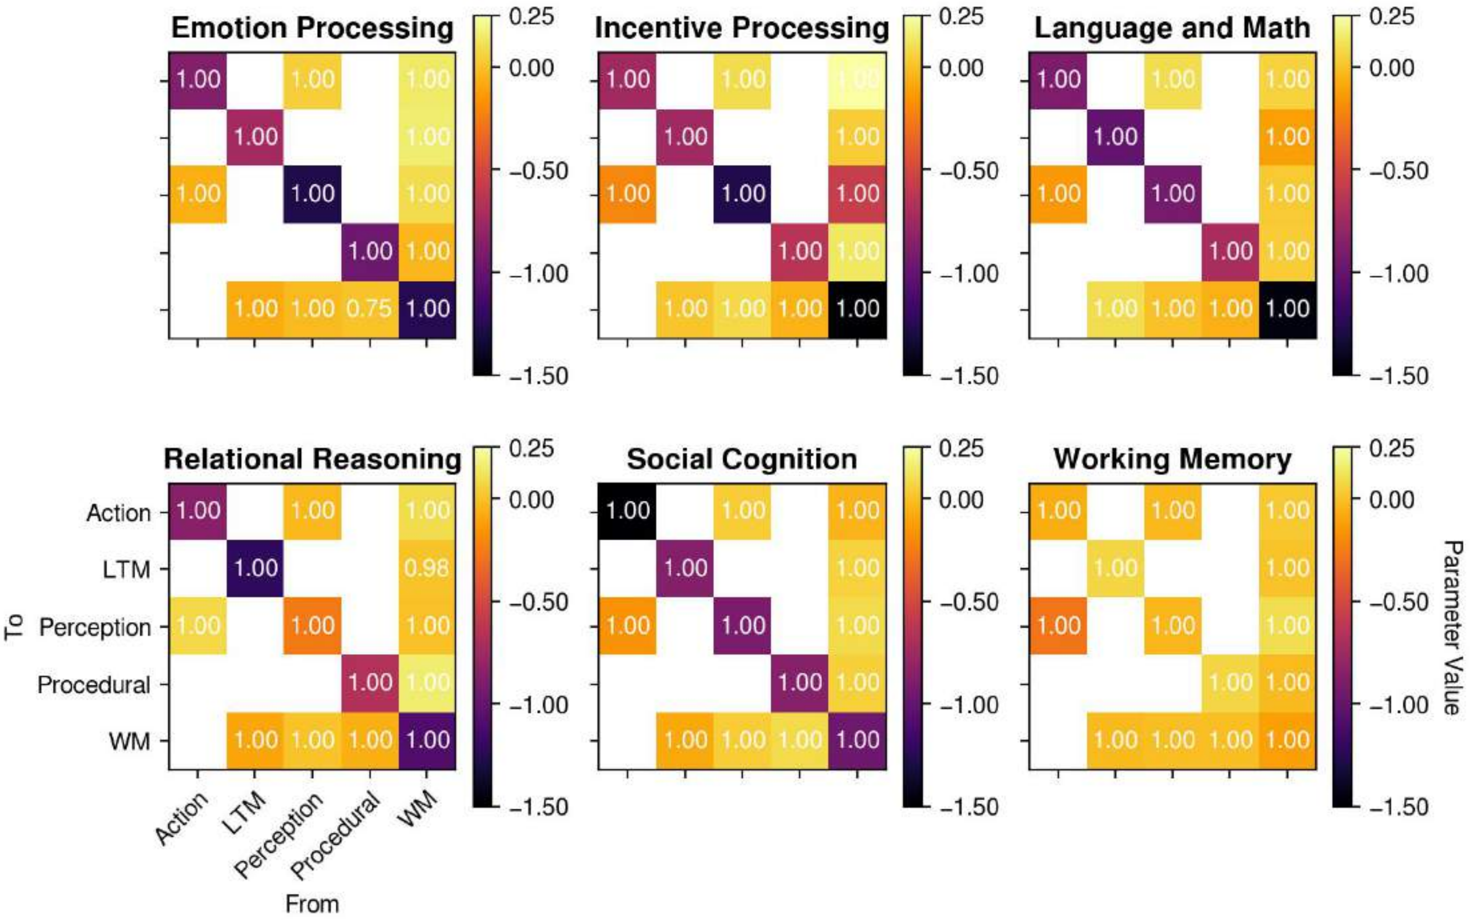
Emotion Processing
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00
1.00 1.00
1.00 1.00 0.75 1.00
Incentive Processing
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00 1.00
Language and Math
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00 1.00
Relational Reasoning
1.00 1.00 1.00
1.00 0.98
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00 1.00
Social Cognition
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00 1.00
Working Memory
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00
1.00 1.00
1.00 1.00 1.00 1.00
0.25
0.00
−0.50
−1.00
−1.50
Parameter Value
To
Action
LTM
Perception
Procedural
WM
From
Action
LTM
Perception
Procedural
WM

Same approach can be used to compare models within architectures

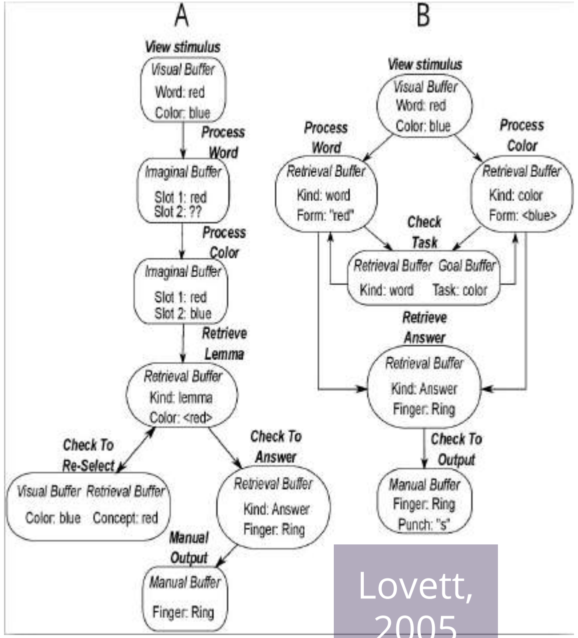

Lovett, 2005

Altmann, 2001

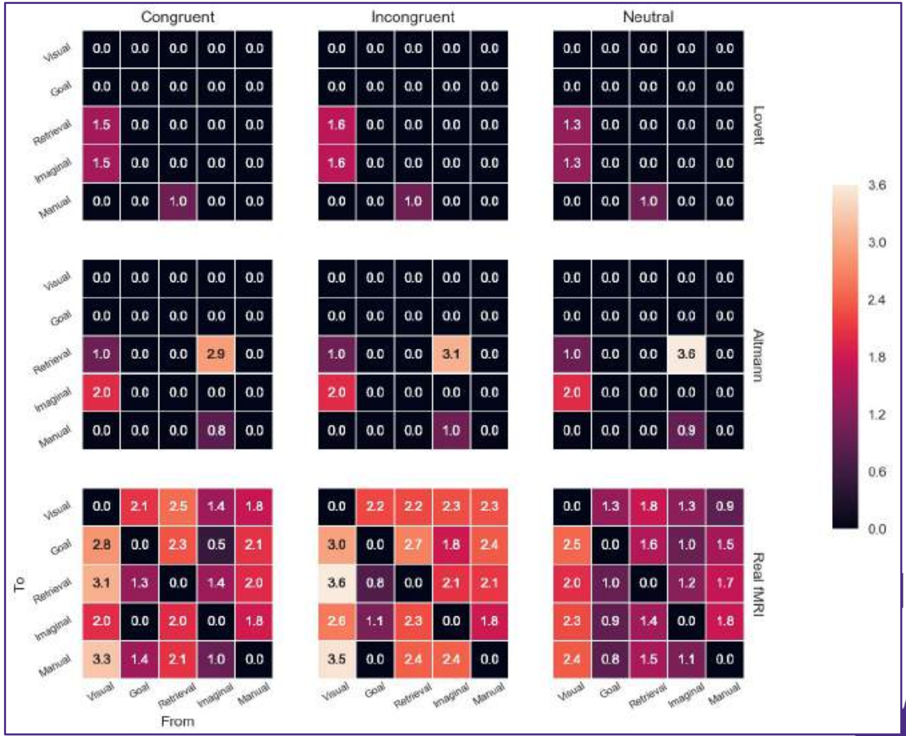

Same approach can be used to compare models within architectures

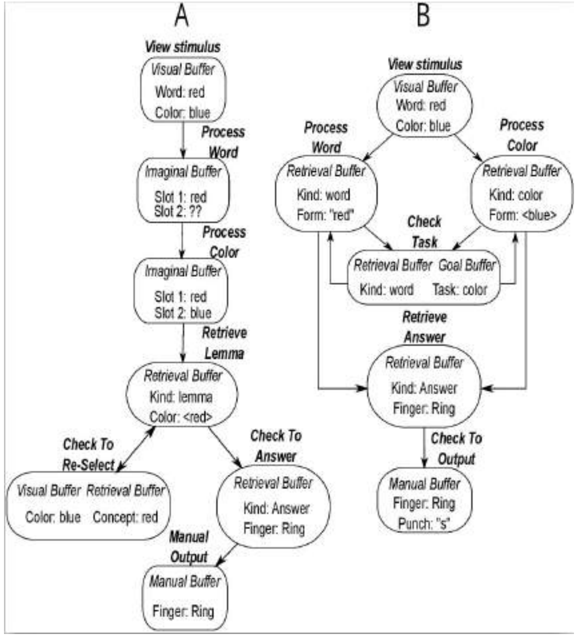

Ketola, Jiang, & Stocco, *Comparing Models with Effective Connectivity.*

Thursday 7/25 @ CogSci, Paper session 18, 2:30-4pm

Thank you!

… Questions?

Credits

John Laird

Christian Lebiere

Paul Rosenbloom

Natalie Koh

Zoe
Steine-Hanson

PSA: We are looking for a post-doc to work on this! Email stocco@uw.edu