A lens into cognition: Topology and geometry of neural systems

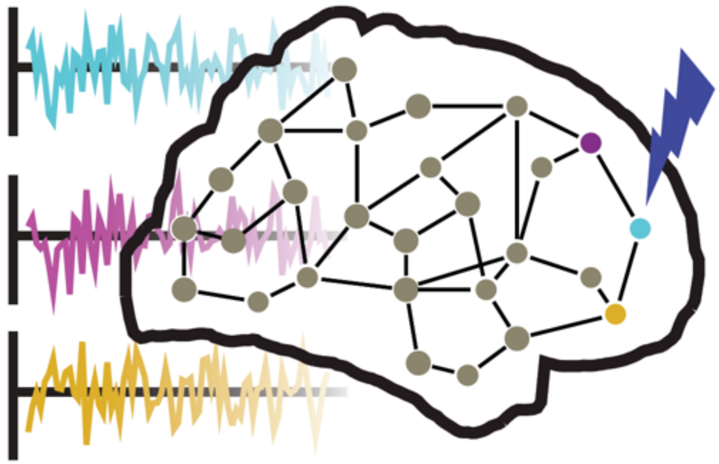

Evelyn Tang
Max Planck Institute
for Dynamics and Self-Organization

KITP Active20
May 14, 2020

Emergent phenomena across scales

Space

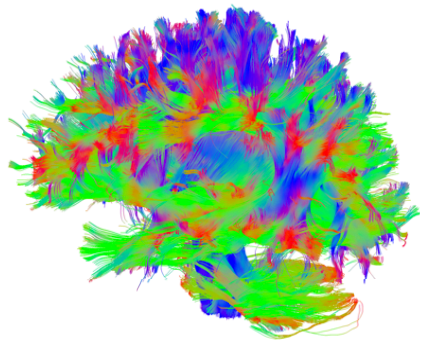

Brain networks and control

Tang et al., *Nature Comm* 2017

Tang & Bassett, *Rev Mod Phys* 2018

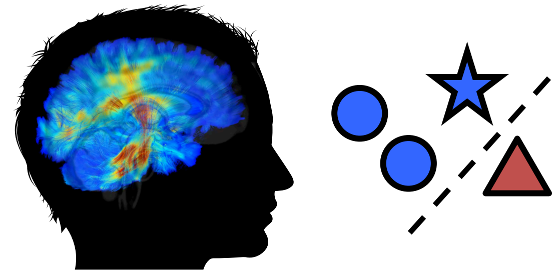

Effective learning

Tang et al., *Nature Neuro* 2019

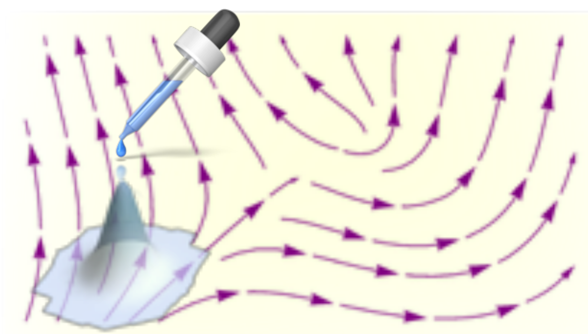

Information in fluid flows

Tang & Golestanian, *arXiv* 2019

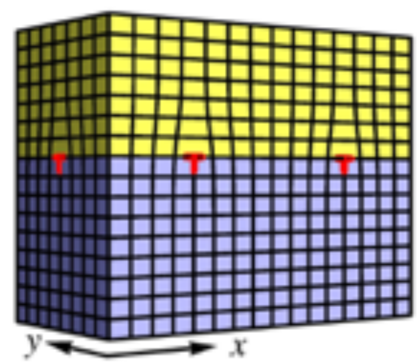

Topological phases of matter

Tang and Fu, *Nature Phys* 2015

Tang et al., *Phys Rev Lett* 2012

Time

Learning: an out-of-equilibrium process

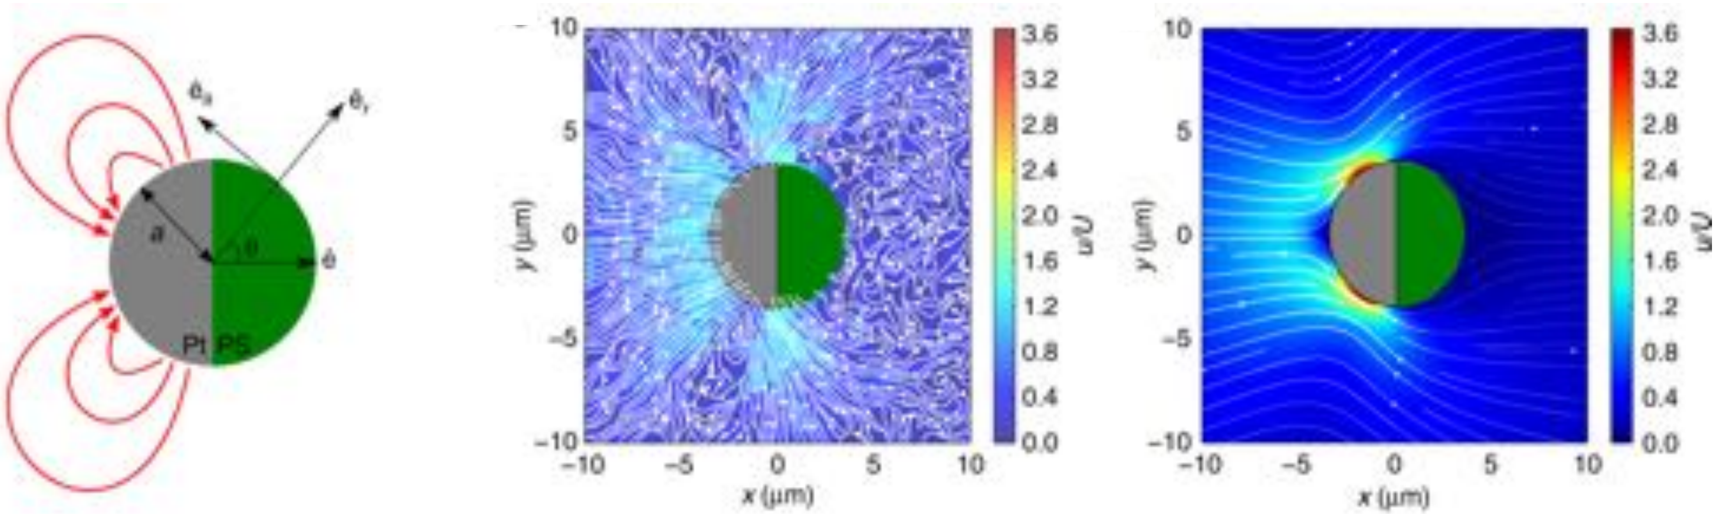

Janus sphere with controllable orientation

Golestanian & Ebbens groups, *Nat Comm* 2015

As we gain understanding and control of active systems, can we “teach” them what to do?

Machine learning is successful but opaque and expensive

Cost to train a new model

Strubell, Ganesh & McCallum, *Proc. 57th Comp. Ling.* 2019

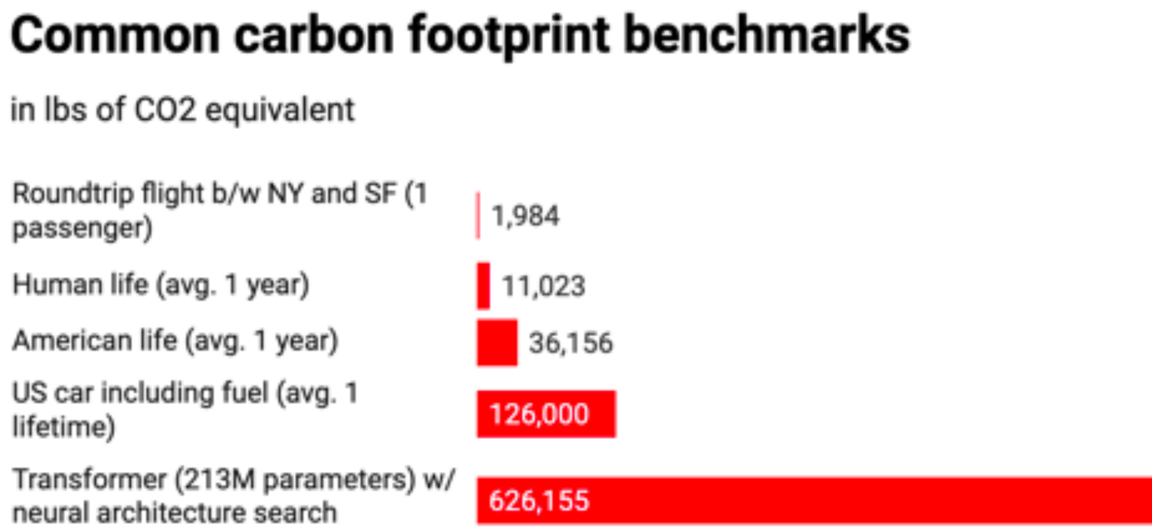

Huge number of parameters

Biological learning is quick and efficient

A baby who sees
their parent
use cell phones

What are underlying principles of learning?

Probing learning in humans is difficult

- No ground up theory for cognition: we can model neuron dynamics, but coarse-graining methods lacking
- No controlled experiments
- Sample sizes are small
- Data is noisy; has side effects from other physiological processes

What features in neural data can distinguish between cognitive states?

Coarse-grained feature of a multi-dimensional dataset

Learning engages complex dynamics:
coordination over different modalities
including sensory, attentional, memory

Given noisy data, won't study specific dynamics

Hypothesis: there exists a suitable dimension
for computational complexity

Neural data can be separated along a dimension

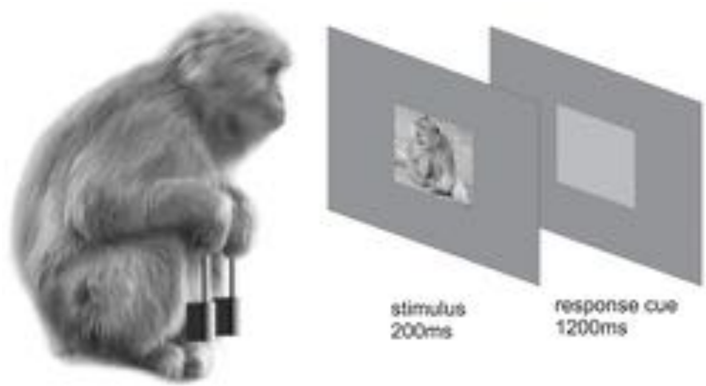

Trained to categorize cats and dogs

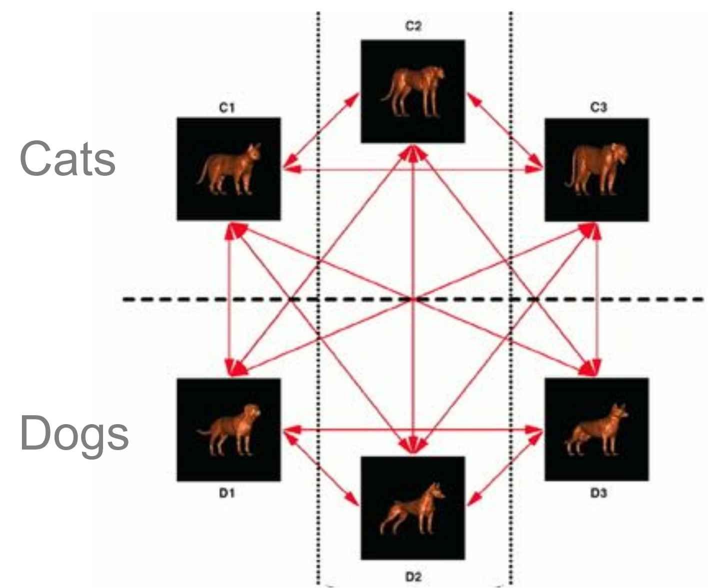

Freedman et al., *Science* 2001

Activity in lateral prefrontal cortex of monkeys could be classified according to animal type

Effective dimension can be lower than that of measurement space

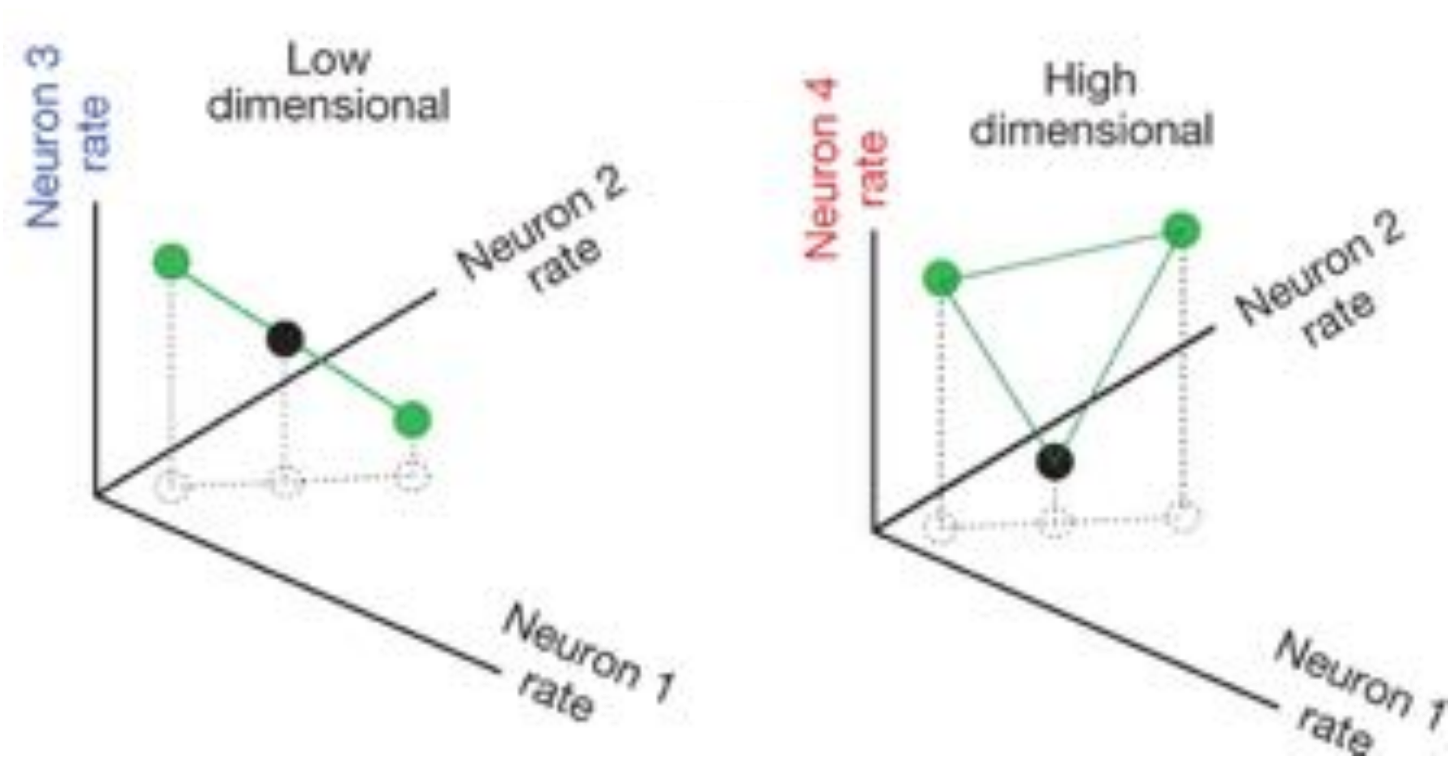

Rigotti et al., *Nature* 2013

Probe the appropriate dimension for successful learning

Combinatorial approach to estimate dimension for noisy data

Given n types of data (shapes): 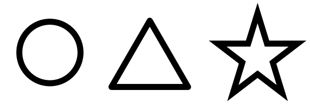 n categories

Assign binary labels (blue or red): 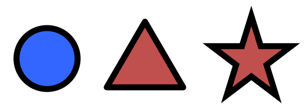 2^n ways

1D:

2D:

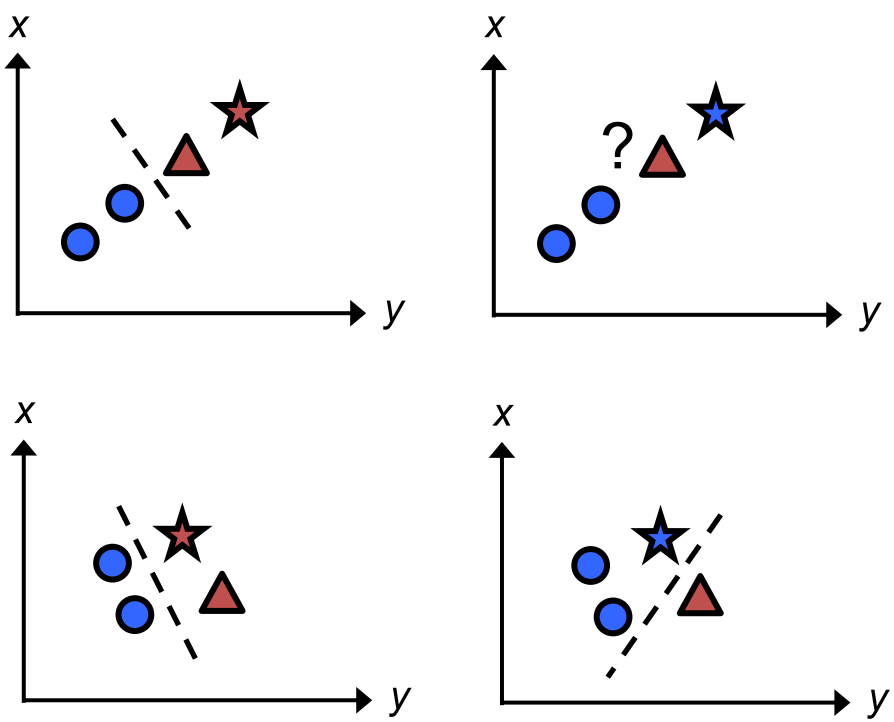

Their linear separability (over different assignments) estimates the dimension

Unlike spectral analysis: does not depend on a metric

Rigotti et al., *Nature* 2013

Experiment with complex cognitive stimuli

Computer-generated shapes with similar statistical properties

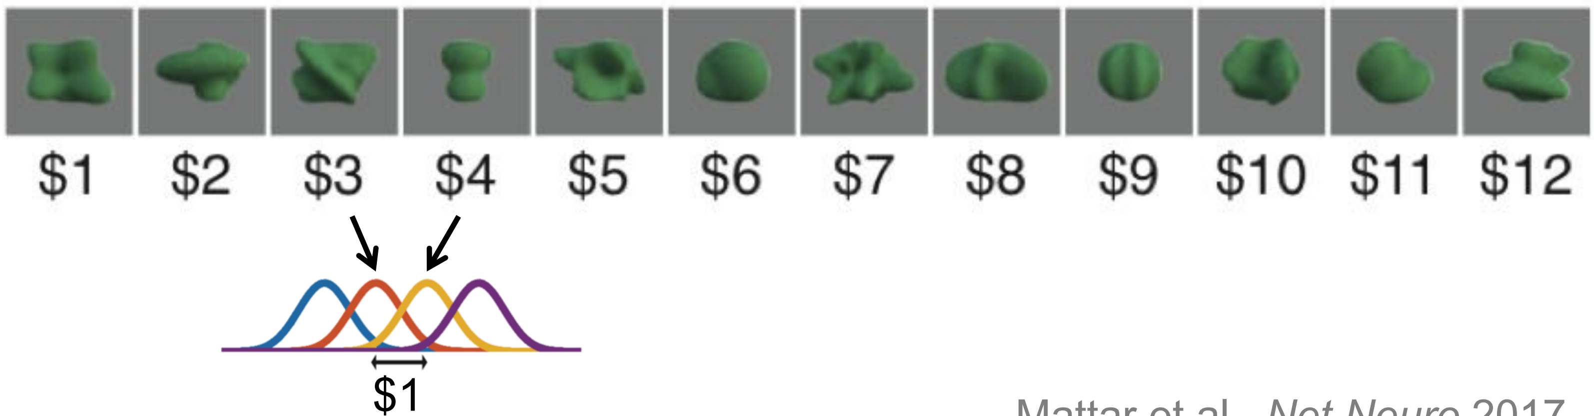

Mattar et al., *Net Neuro* 2017

Shapes have value drawn from a Gaussian with fixed mean

Participants had to associate dollar values to each new shape

Adult participants learned the values of these shapes through feedback

3 training sessions a day, over 4 days

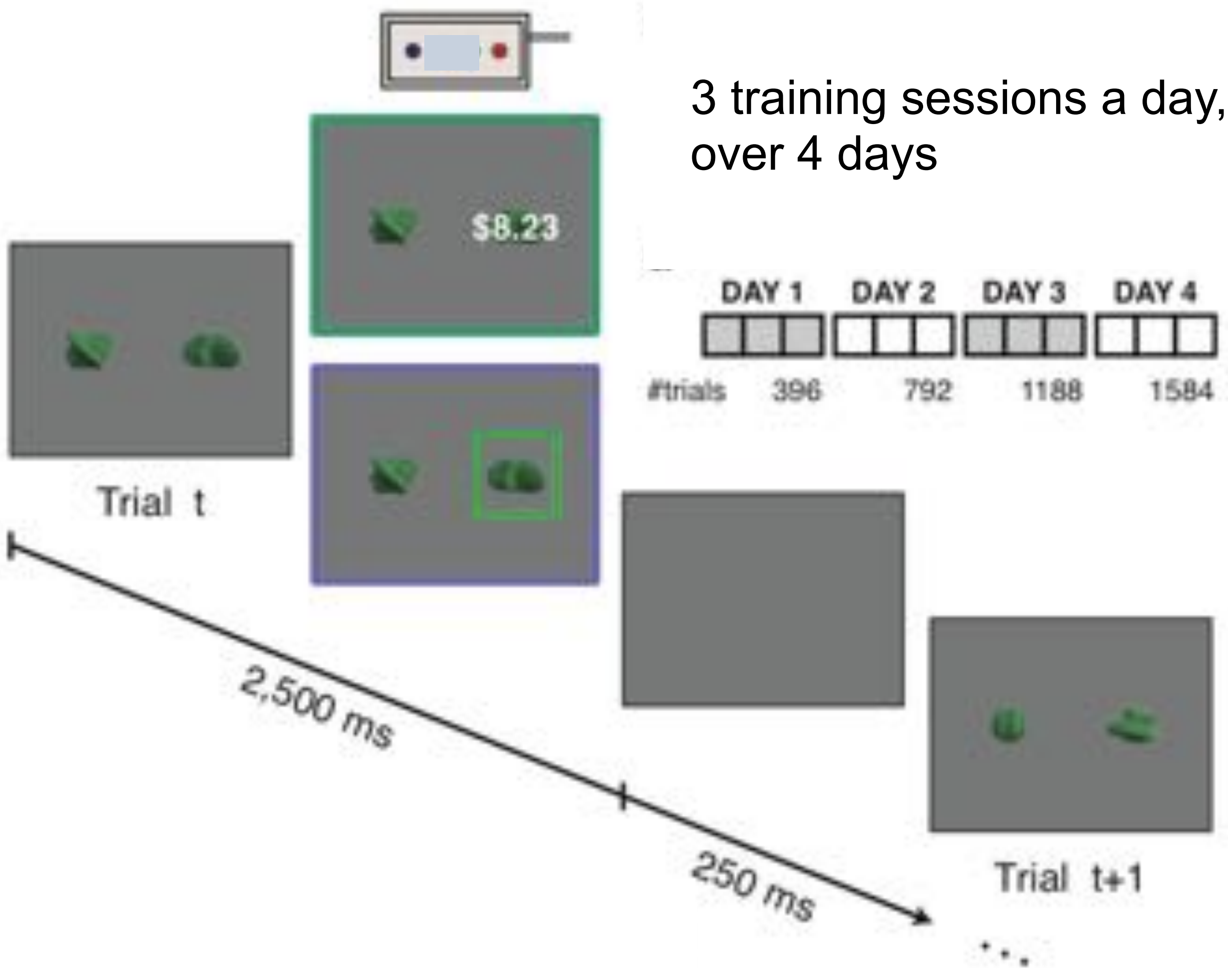

Their neural patterns scanned using fMRI throughout the experiment

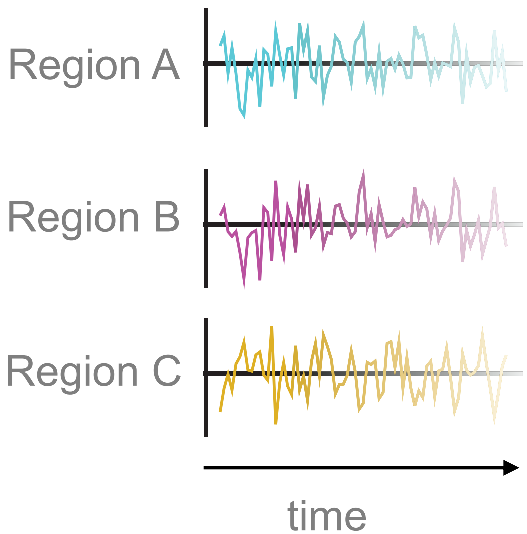

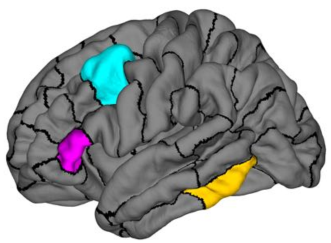

Each session: 140 pairs shown

Blood-oxygen-level dependent activation measured

Coarse-grained approach:

83 regions parcellation of whole-brain

Neural responses across all regions form a geometric representation

140 shapes each contribute a point in data cloud

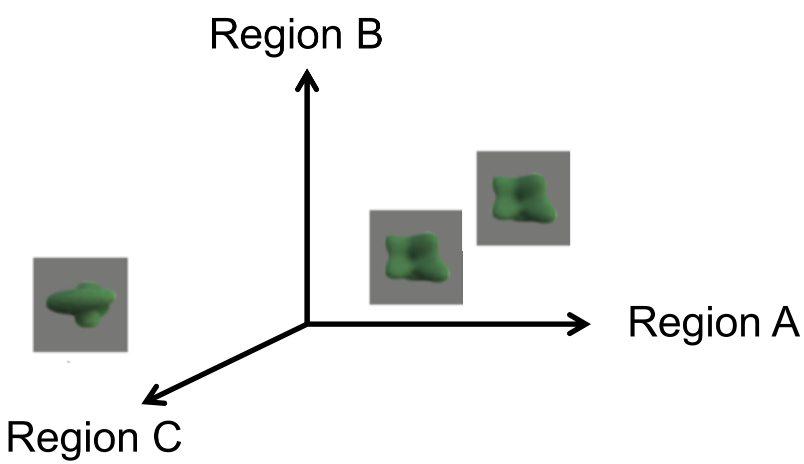

For n categories:
2^n ways to assign binary labels

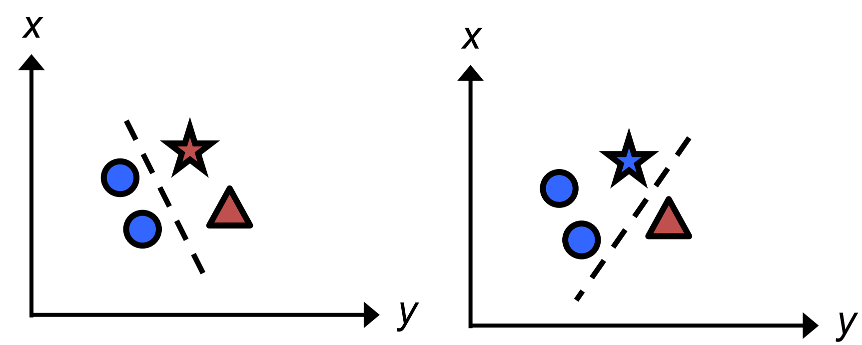

Average separability over hyperplanes is a proxy for dimension – large combinatorics allows method to be robust to noise

Fast learners have a higher dimensional representation of neural data

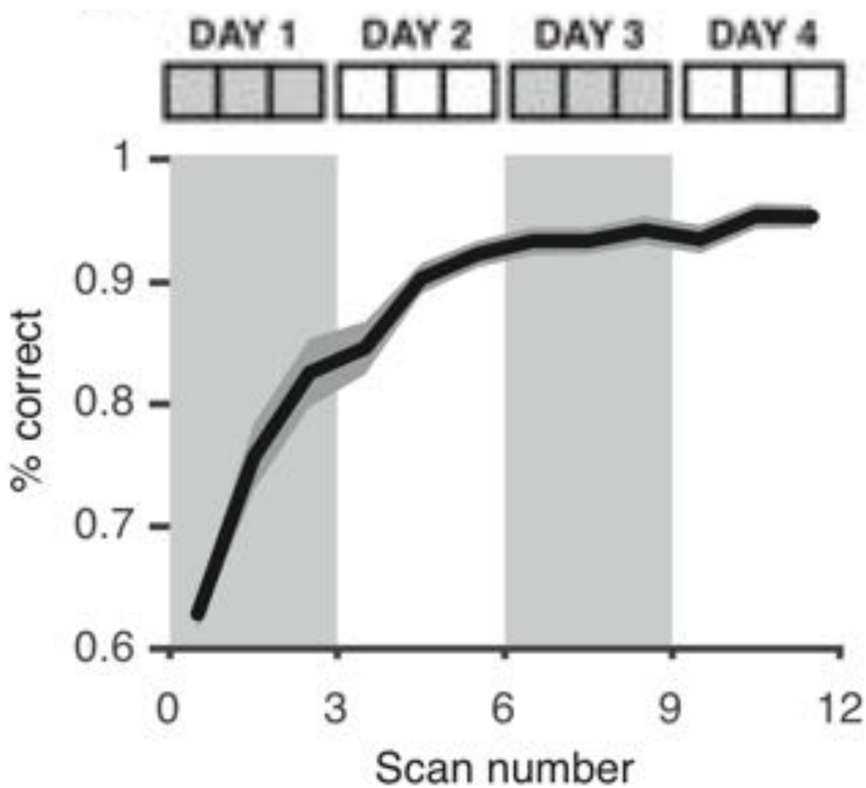

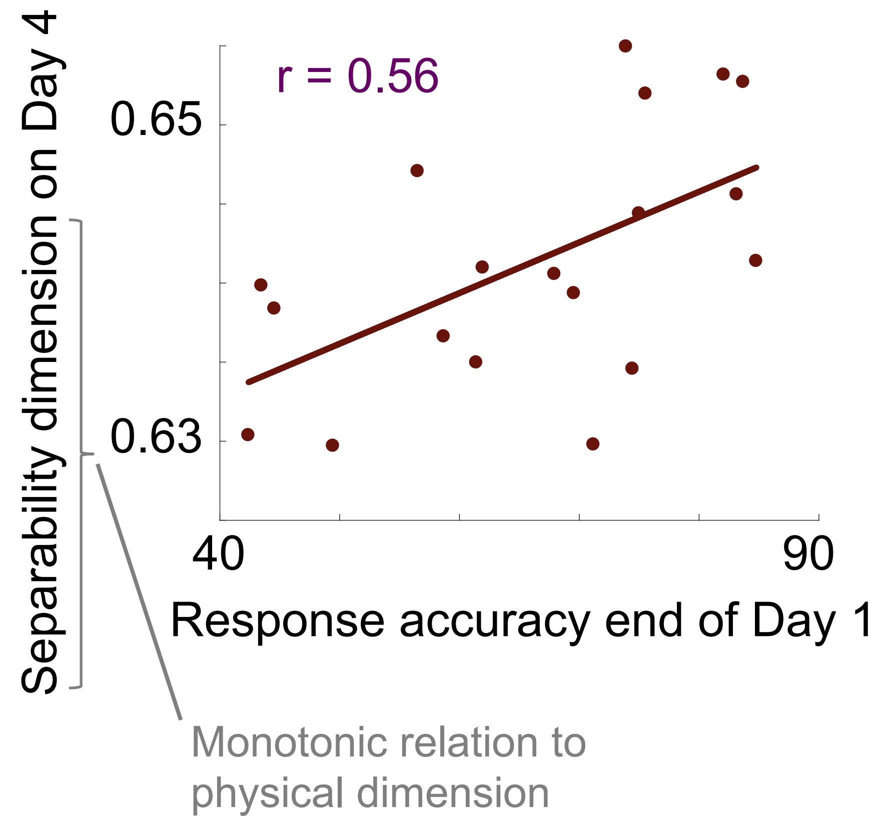

Higher dimensional neural representations are associated with effective learning on this task

Test result reliability using a null model

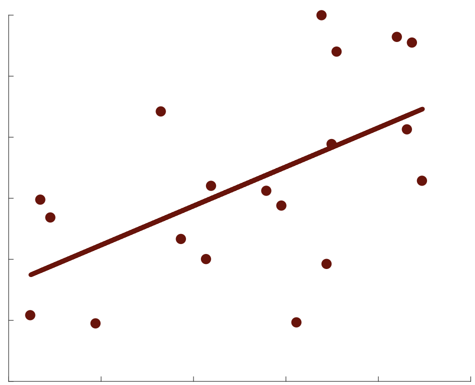

Without theory, and without repeated experiments to fit – need to identify "no result"

Obtain a baseline comparison from similar data without task information

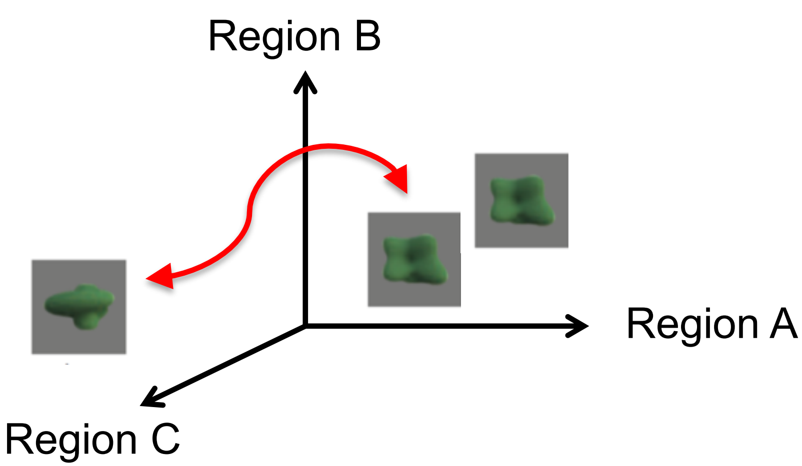

1. Shuffle task labels of data points
2. Repeat analysis

Null model shows result is significant

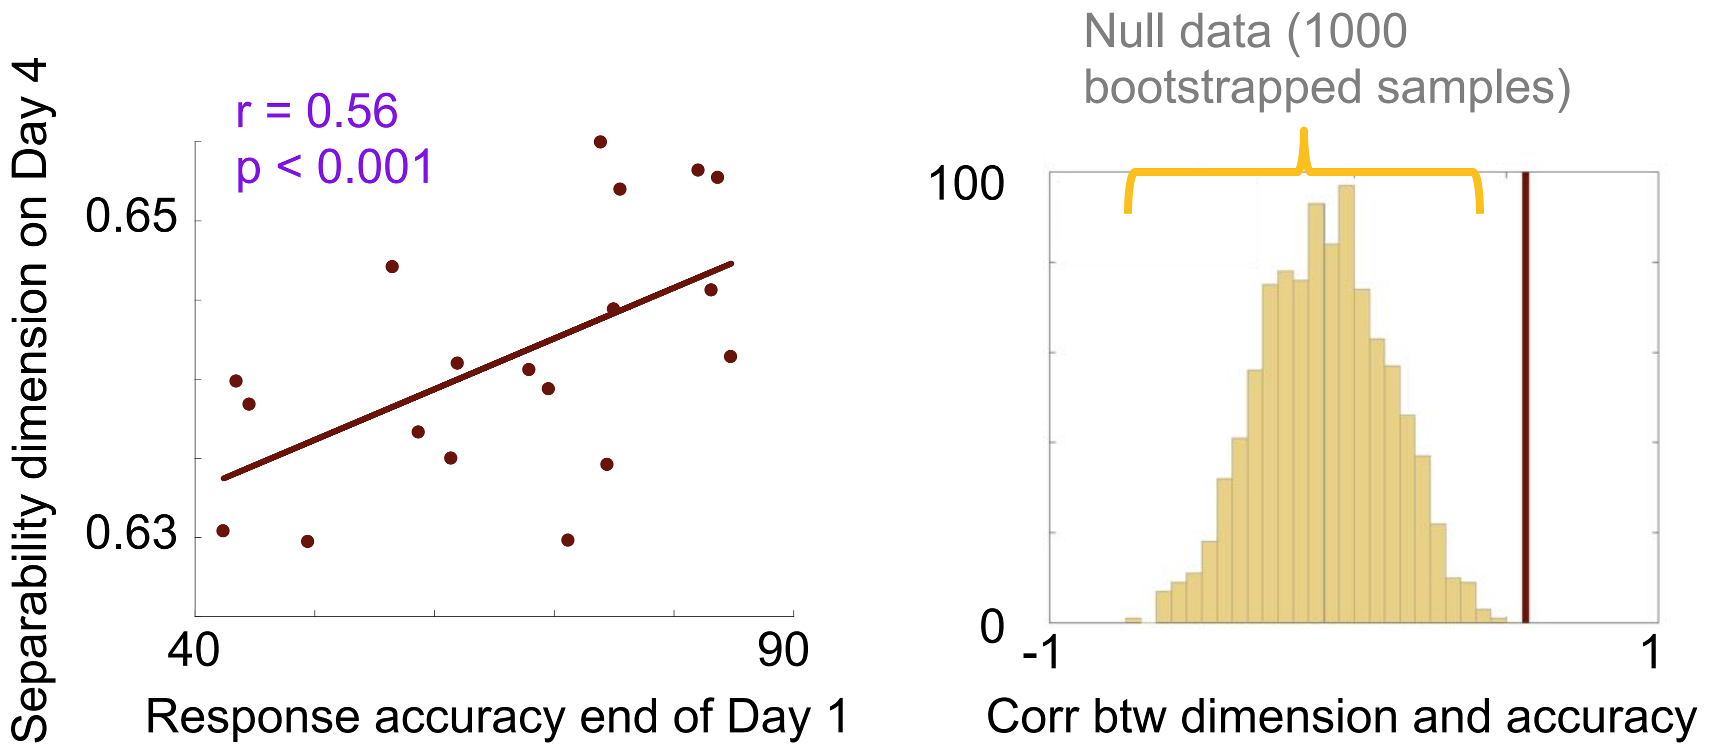

Null data also shows smaller dimension for fast learners

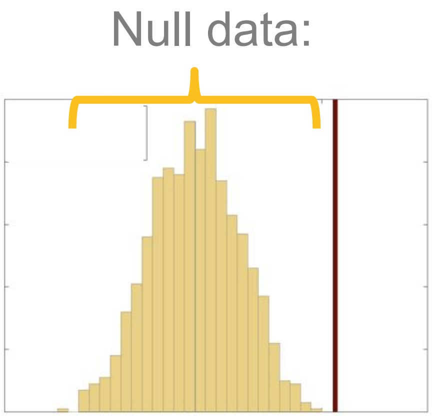

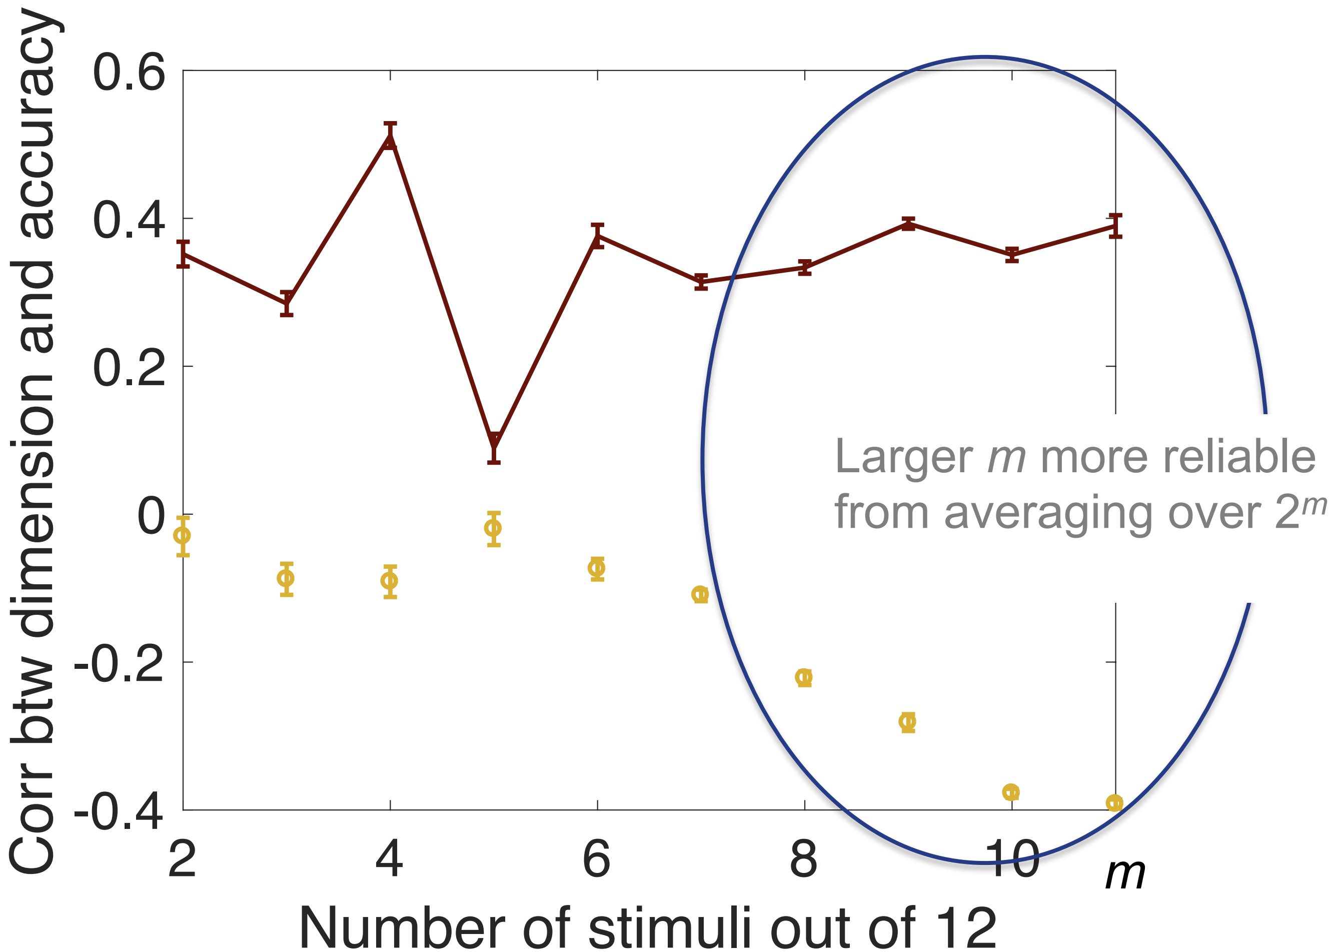

Negative correlation between null data dimension and learning accuracy

Fast learners have higher task-based dimension and lower embedding dimension

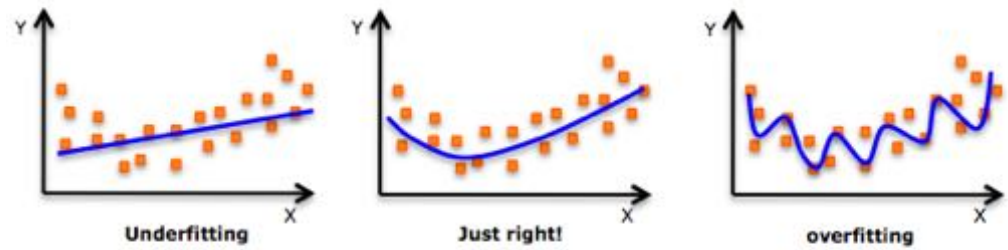

Efficient **Inefficient**

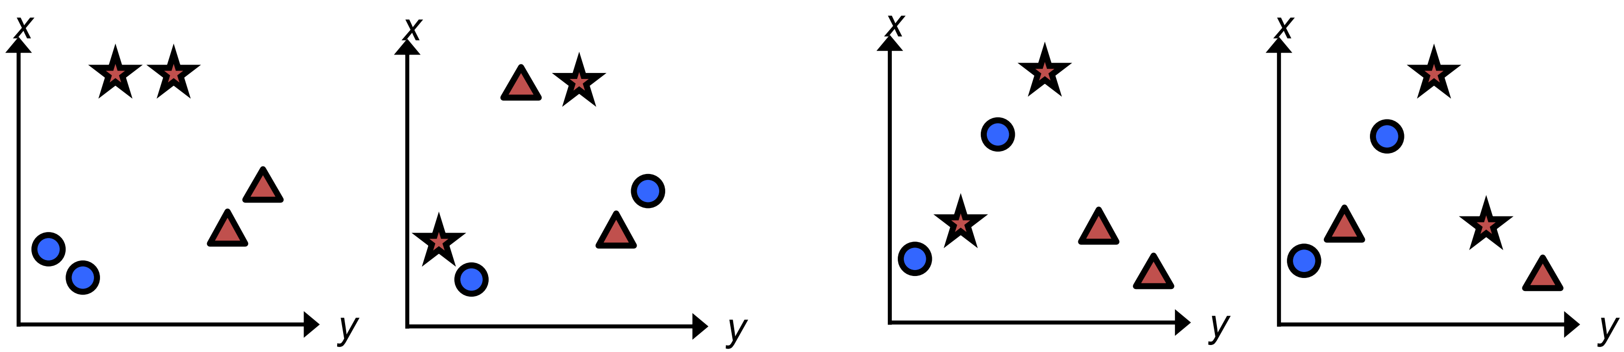

Task data dim ~ 3 Shuffled data dim < 2 Task data dim < 3 Shuffled data dim ~ 2

Fast learners have an efficient representation:
high ratio of information-coding to resources used

Tang, Mattar, Giusti, Lydon-Staley, Thompson-Schill & Bassett, *Nature Neuroscience* 2019

Analogous results seen in data manifolds of neural networks

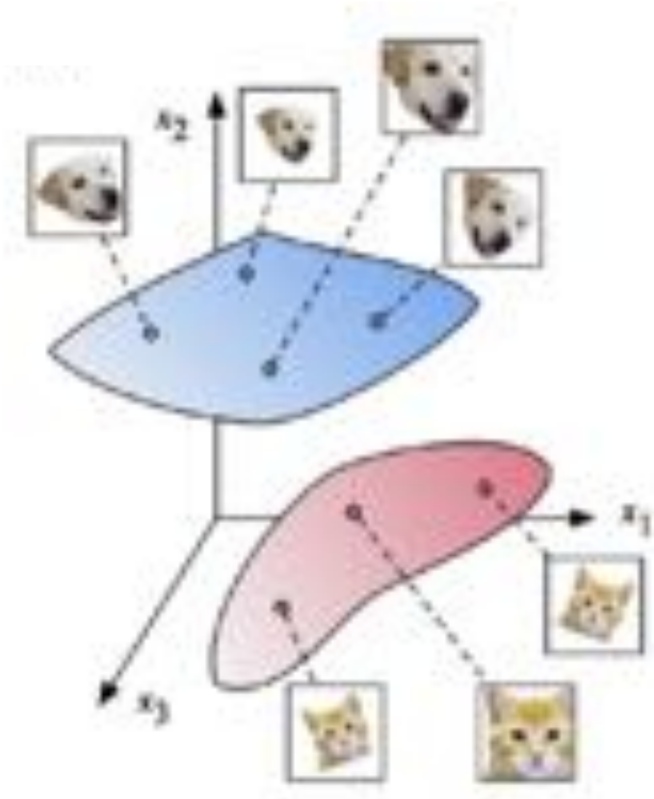

Analysis of the shape of manifold representations in neural networks

Two kinds of dimensionality that can behave in different ways with training

Chung, Lee & Sompolinsky, *PRX* 2018

Virtual lesioning: data-driven approach to identify which brain regions contribute most

Brain regions are removed one at a time: result recalculated
Largest change (in correlation of accuracy with dimension) due to:

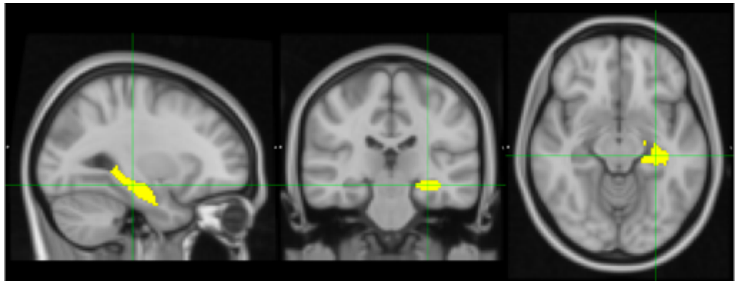

Left hippocampus

Associated with rapid learning of stimulus associations

Squire, *Psych Rev* 1992

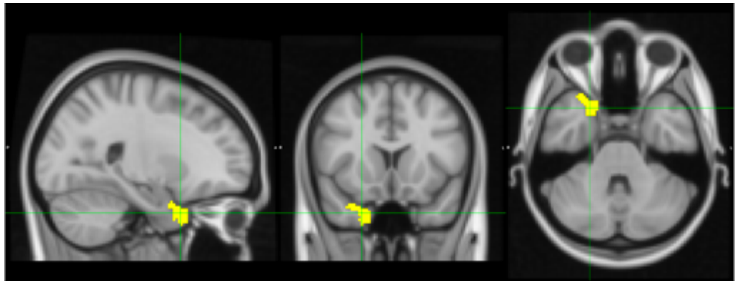

Right temporal pole

Represents information about abstract conceptual properties (such as value)

Peelen & Caramazza,
J Neuroscience 2012

Recapitulation of effect on smaller voxel-level in some regions

Study 5 regions in each hemisphere with 300 (of fewer) voxels

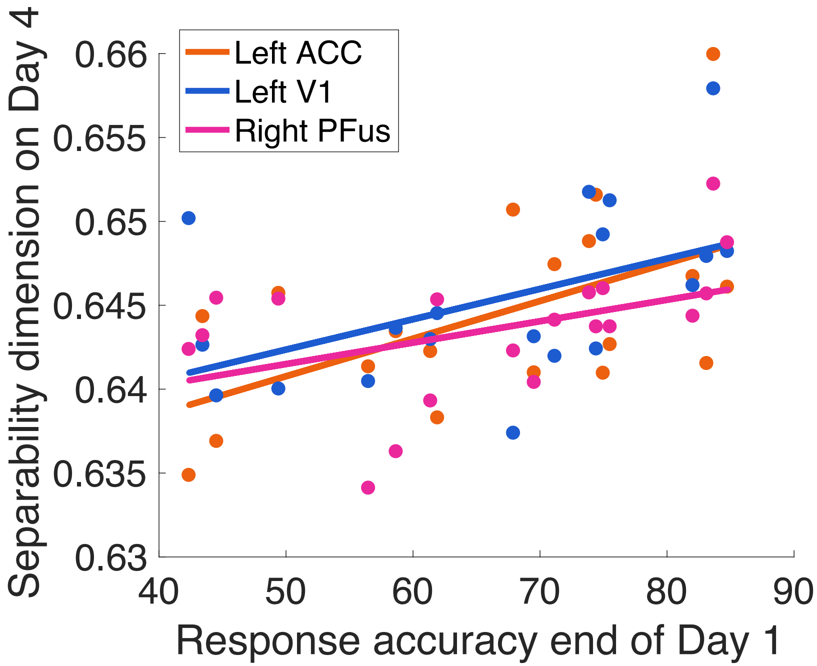

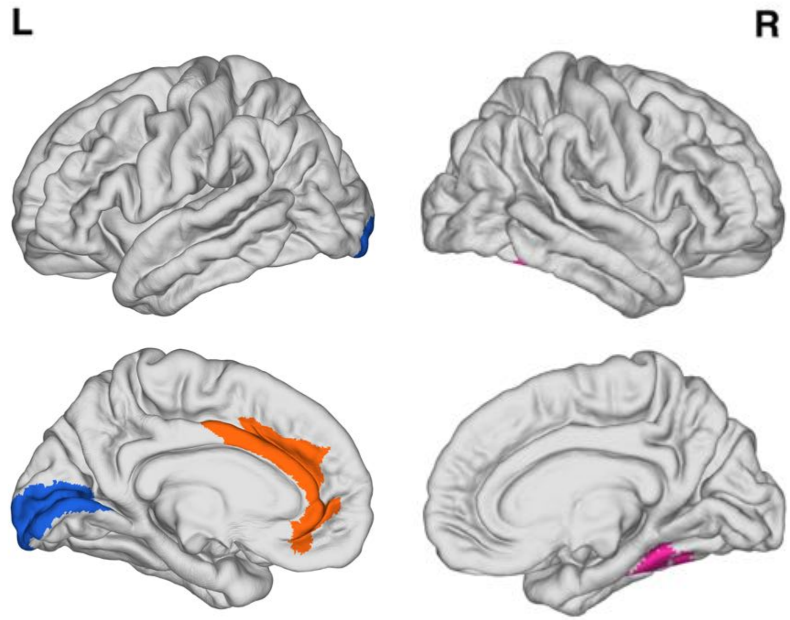

Left anterior cingulate cortex has strongest result and known role in reward-based learning

Bush et al., *PNAS* 2002

Followed by left V1 and right posterior fusiform

The geometry of neural activity reflects cognitive performance

1. Fast learners have higher dimensional representations of neural activity.

2. This allows objects of different value to be more easily distinguished.

3. Fast learners also have lower embedding dimension: hence they have more efficient representations with a high ratio of information-coding to resources used.

Emergent phenomena across scales

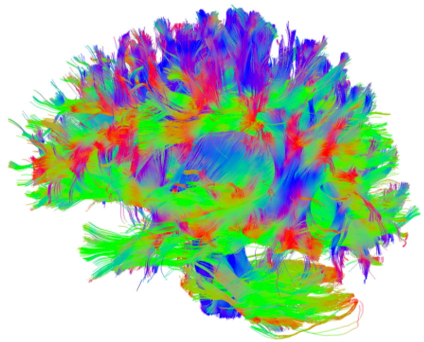

Brain networks and control

Tang et al., *Nature Comm* 2017

Tang & Bassett, *Rev Mod Phys* 2018

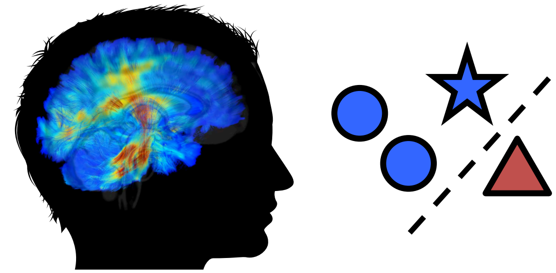

Effective learning

Tang et al., *Nature Neuro* 2019

Space

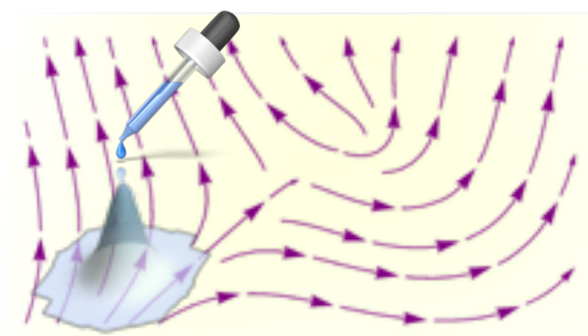

Information in fluid flows

Tang & Golestanian, *arXiv* 2019

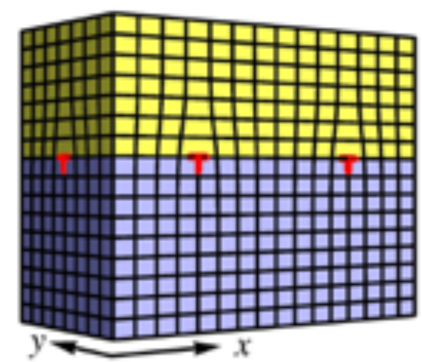

Topological phases of matter

Tang and Fu, *Nature Phys* 2015

Tang et al., *Phys Rev Lett* 2012

Time

How does brain structure subserve dynamics and function?

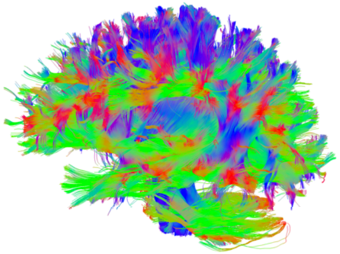

Control theory and dynamical models to probe the role of connectivity;

and in changes across development

Towards understanding function: e.g. children are more spontaneous while adults are better at cognitive control

Network control theory models dynamics in heterogeneous real-world systems

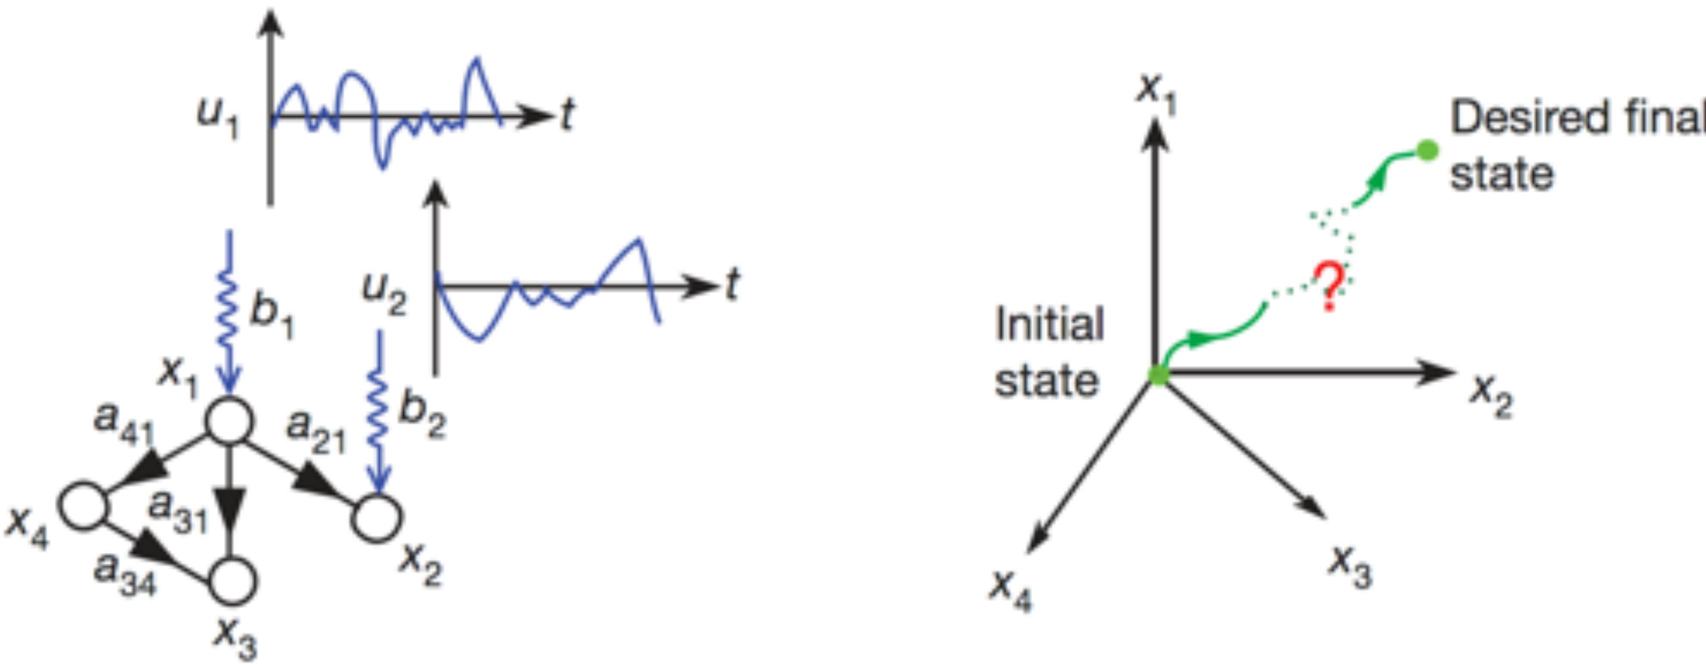

Liu et al., *Nature* 2011

Models the driving of dynamical changes across neurons or neural systems

External Input

Stimulation

Neurofeedback

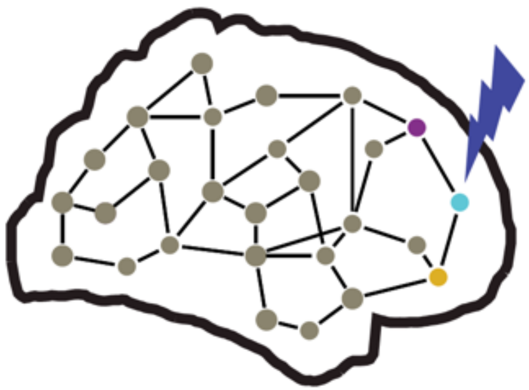

Internal Control

Cognitive Control

Homeostasis

Bassett et al., *Ann Rev Biomed Eng* 2017

Topology of brain connectivity mapped with non-invasive neuroimaging

Identify brain regions on the mesoscale

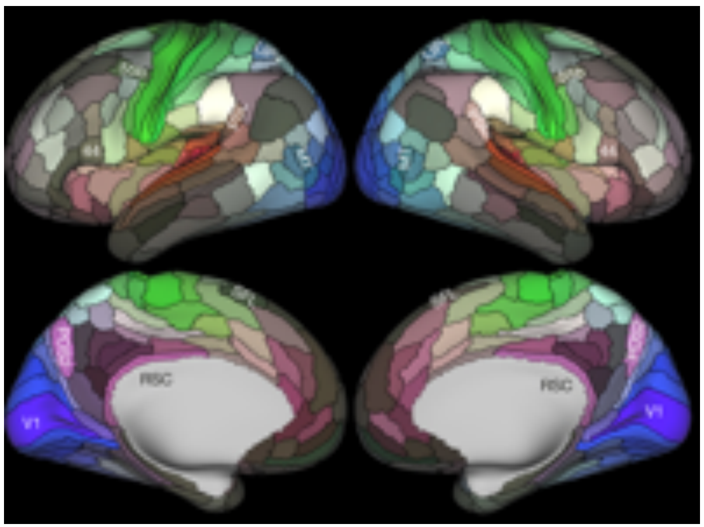

Glasser et al., *Nature* 2016

White matter pathways inferred from movement of water molecules diffusing along tracts

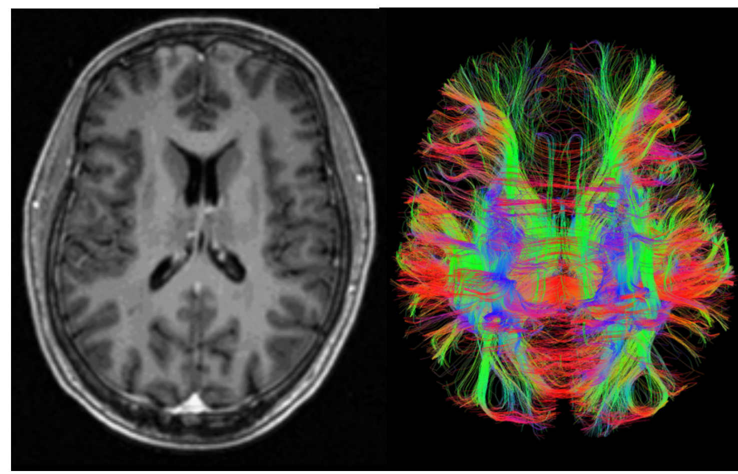

Tuch et al., *Neuroimage* 1997

Build brain network which estimates strength of connections between regions

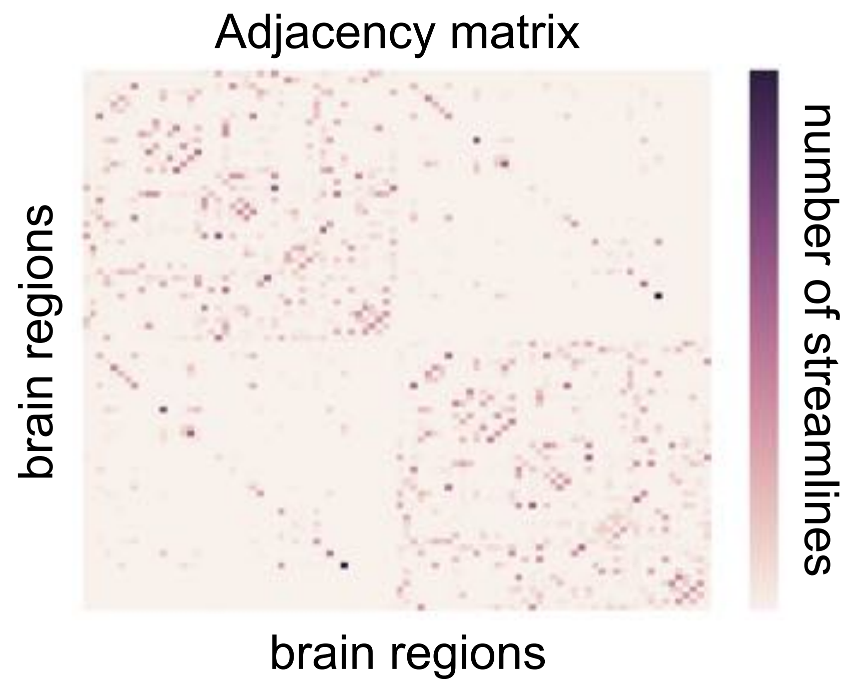

We represent the pattern of white matter tracts between brain regions as an undirected, weighted adjacency matrix

Bullmore & Sporns, *Nat Rev Neurosci* 2009

Linear dynamical model
+ input into system

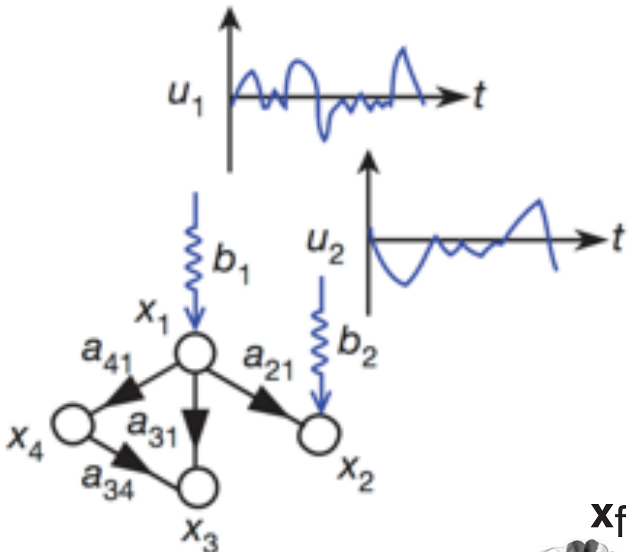

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$$

Regions of input

Input into system

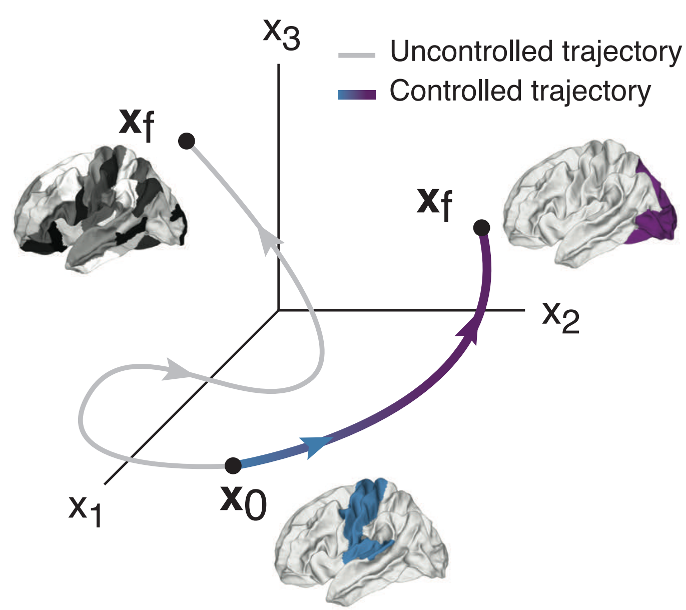

Input into systems defines an energy landscape

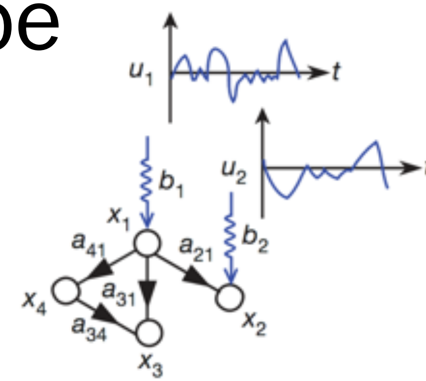

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$$

After T steps,

$$\mathbf{x}(T) = \mathbf{C}_T \underbrace{\begin{bmatrix} \mathbf{u}(T-1) \\ \vdots \\ \mathbf{u}(0) \end{bmatrix}}_{\text{Input energy}}; \qquad \mathbf{C}_T := \underbrace{\begin{bmatrix} \mathbf{B} & \mathbf{AB} & \cdots & \mathbf{A}^{T-1}\mathbf{B} \end{bmatrix}}_{\text{Depends on brain network and input regions (structural)}}.$$

$$E(\mathbf{u}, T) := \sum_{t=0}^{T-1} ||\mathbf{u}(t)||^2$$

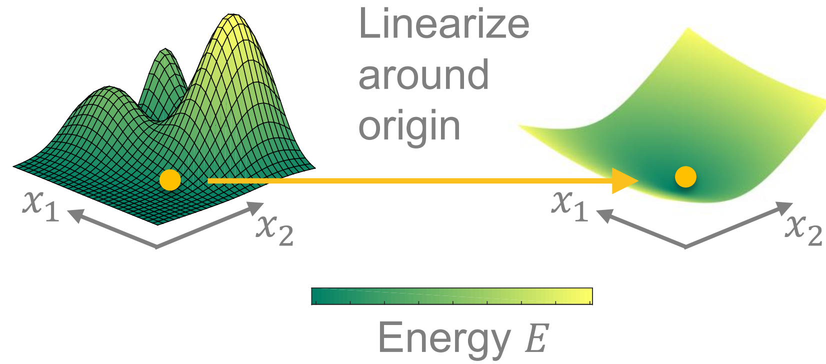

Structural network properties determine the minimum input energy

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$$

Minimum input energy

$$\mathbf{x}(T) = \mathbf{C}_T \begin{bmatrix} \mathbf{u}(T-1) \\ \vdots \\ \mathbf{u}(0) \end{bmatrix}; \qquad \begin{bmatrix} \mathbf{u}^*(T-1) \\ \vdots \\ \mathbf{u}^*(0) \end{bmatrix} = \mathbf{C}_T^\top (\mathbf{C}_T \mathbf{C}_T^\top)^{-1} \mathbf{x}_f$$

Kailath,
Linear Systems 1980

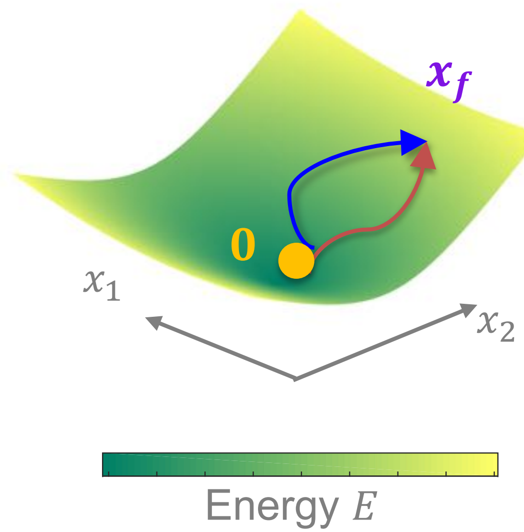

$$E^*(T) = \sum_{t=0}^{T-1} ||\mathbf{u}^*(t)||^2 = \mathbf{x}_f^\top (\mathbf{C}_T \mathbf{C}_T^\top)^{-1} \mathbf{x}_f$$

Gramian $\mathbf{W}_T := \mathbf{C}_T \mathbf{C}_T^\top = \sum_{t=0}^{T-1} \mathbf{A}^t \mathbf{B}\mathbf{B}^\top (\mathbf{A}^\top)^t$

Network connectivity and strength determine possible dynamical transitions

Minimum input energy $E^*(T) = \mathbf{x}_f^{\mathsf{T}} \mathbf{W}_T^{-1} \mathbf{x}_f; \quad \mathbf{W}_T := \sum_{t=0}^{T-1} \mathbf{A}^t \mathbf{B} \mathbf{B}^{\mathsf{T}} (\mathbf{A}^{\mathsf{T}})^t$

When $\mathbf{x}_f$ is an eigenvector of $\mathbf{W}_T$ with eigenvalue λ, $E^*(T) = \lambda^{-1}$

Average input energy (over all $\{\mathbf{x}_f\}$): $\mathrm{Tr}(\mathbf{W}^{-1})$
Ability to control with least average energy: $\mathrm{Tr}(\mathbf{W})$

$T \to \infty$
$$\mathrm{Tr}(\mathbf{W}_T^{-1}) \geq \frac{N^2}{\mathrm{Tr}(\mathbf{W}_T)}$$

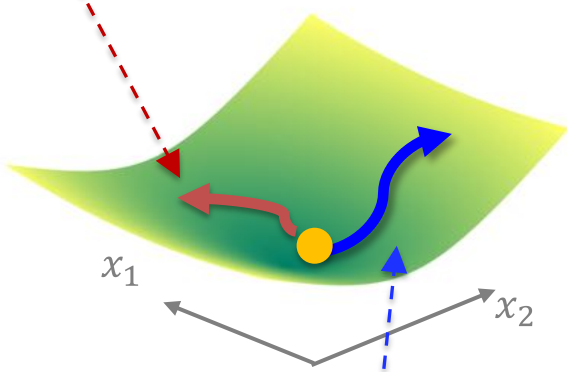

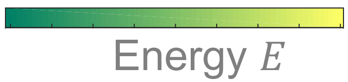

Pasqualetti et al., *IEEE TCNS* 2014
Gu et al., *Nat Comm* 2015

v_j: j^{th} eigenvector of A with eigenvalue ξ_j.
If v_{ij} is small, then j^{th} mode is poorly controllable from i (extension of PBH test).

"Worst-case" energy from $E^*(T) \leq \lambda_{\min}^{-1}(\mathbf{W}_T)$
Ability for modal control of most costly transition from i: $\sum_j (1 - \xi_j^2(\mathbf{A})) v_{ij}^2$

Use network control metrics on cohort of 882 youth from 8 to 22 years

Diffusion data from Philadelphia Neurodevelopmental Cohort

Roalf et al., *Neuroimage* 2016

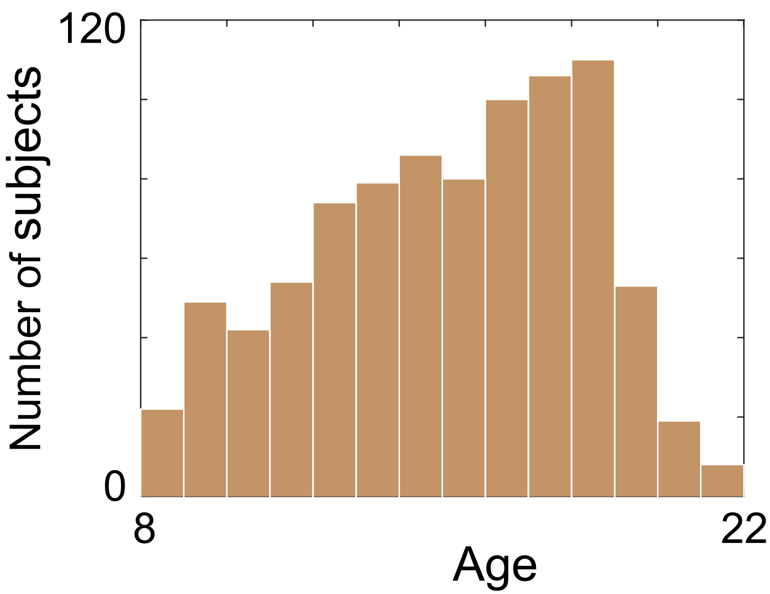

Theodore Satterthwaite

Do brain networks show increasing control with age?

Topology of brain networks supports more dynamical transitions

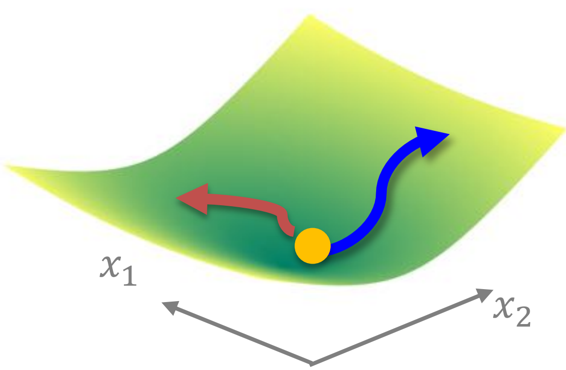

Coarse metrics across whole-brain

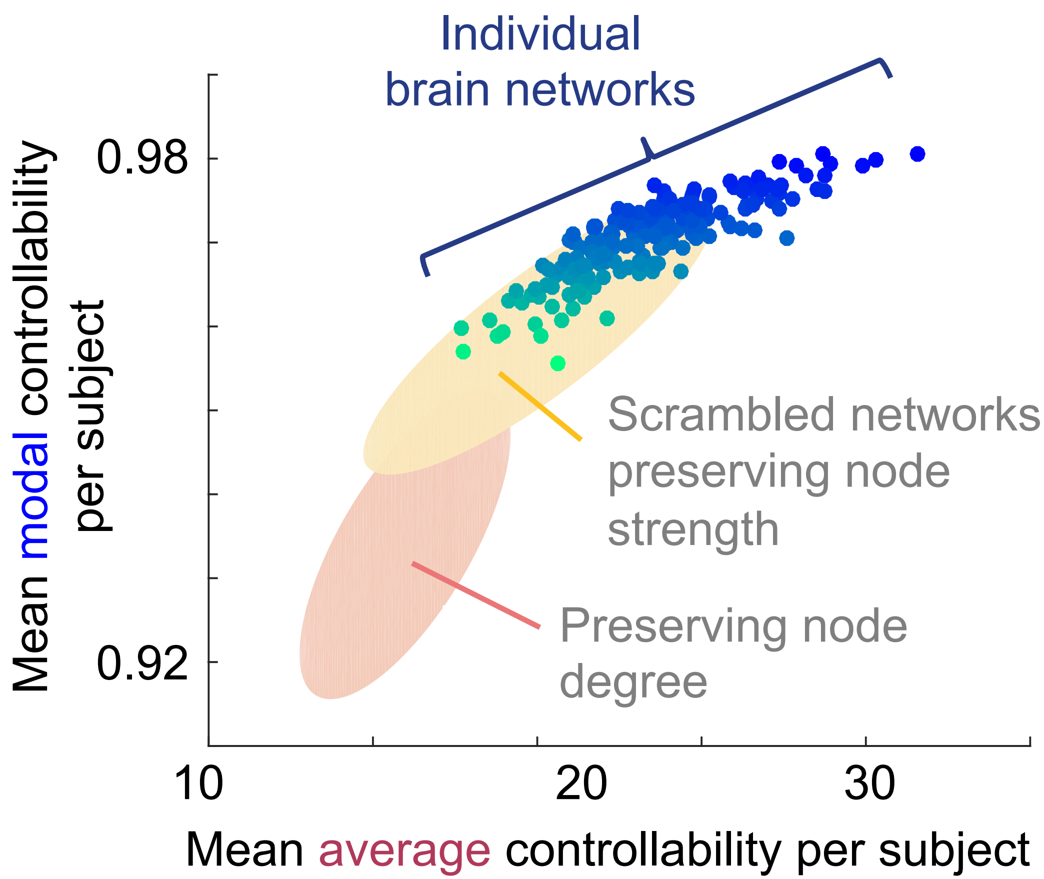

Brain networks increasingly support diverse dynamics with age

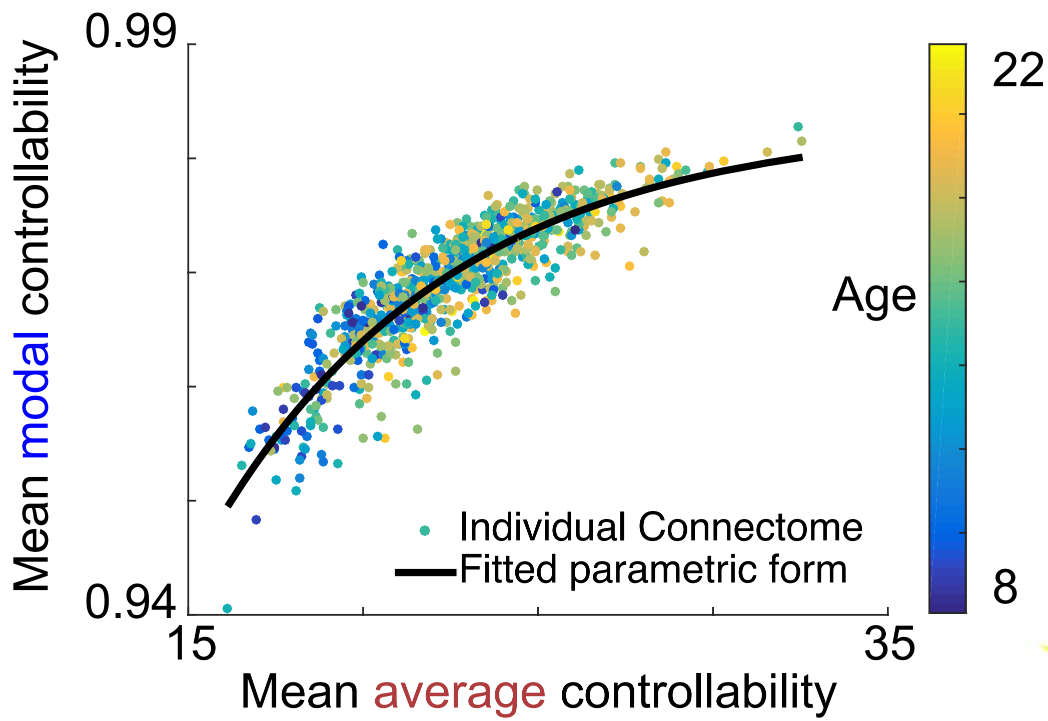

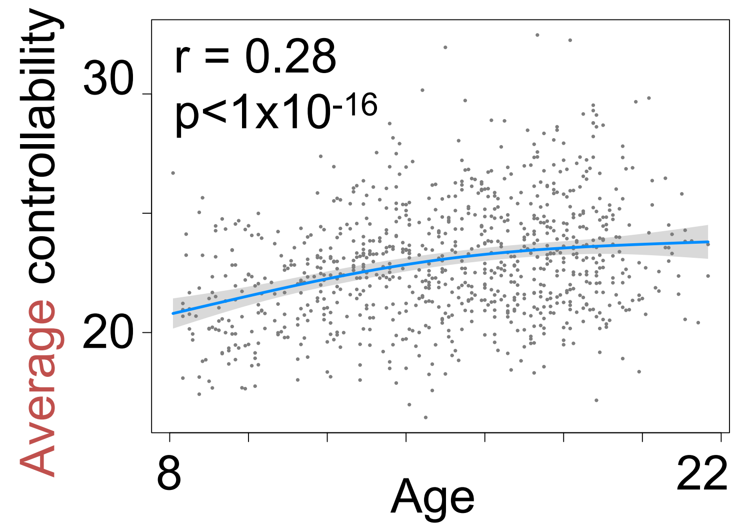

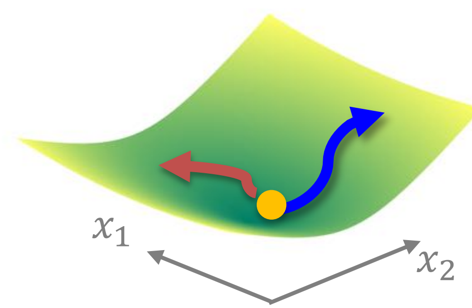

Older subjects have a larger range of possible dynamics from low to high energy transitions: increased specialization

What can a finer look tell us?

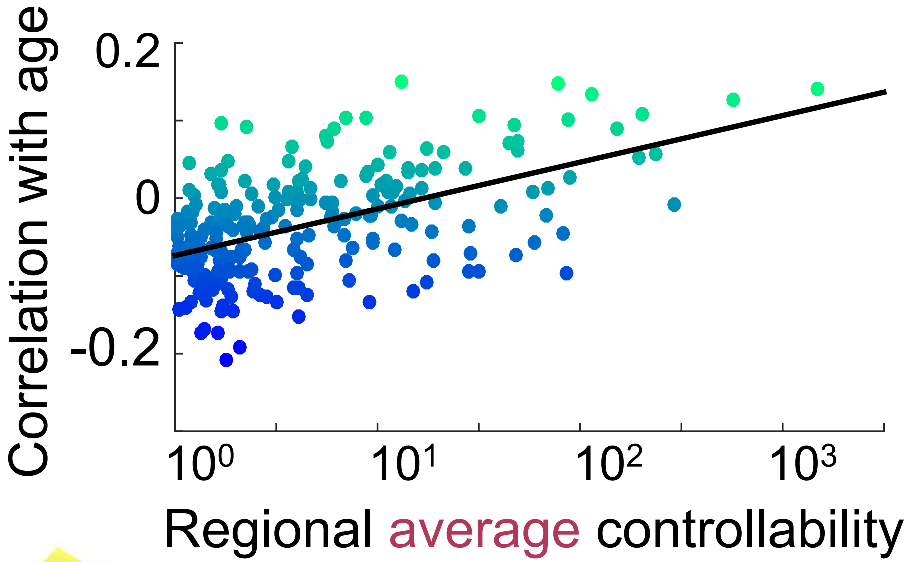

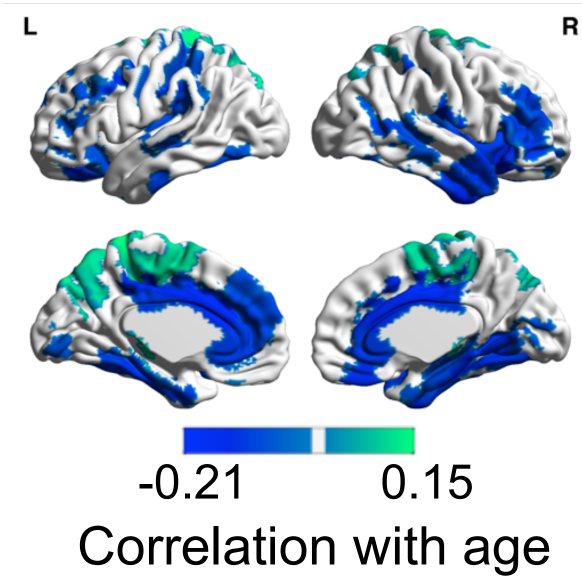

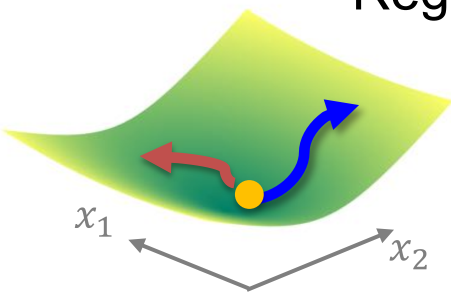

Regions high in average controllability increase in controllability with age

Regional specialization of control with age or "super-controllers"

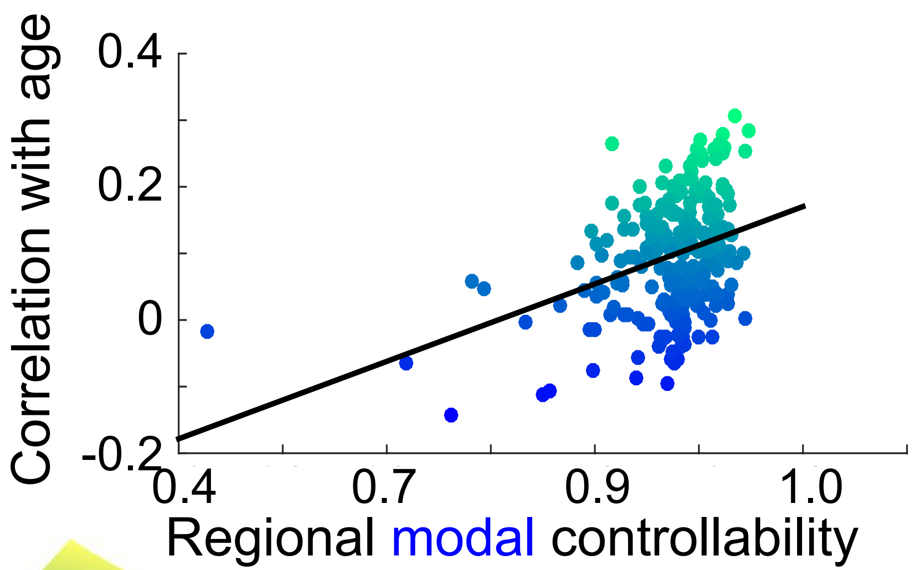

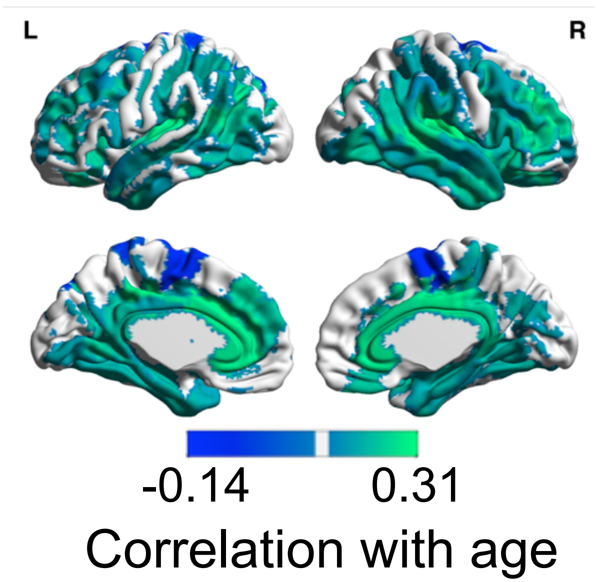

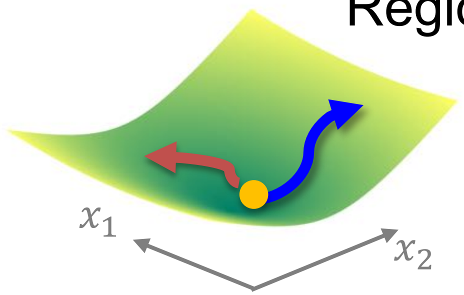

Seen in both average and modal control regions

Which brain regions have high control for effective cognition?

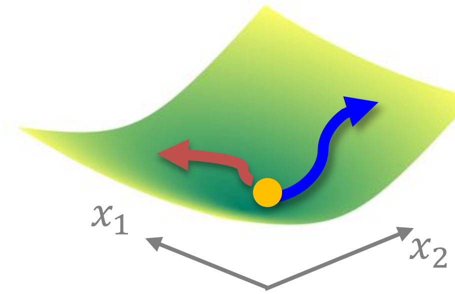

Subjects with high subcortical controllability exhibit poorer cognitive performance

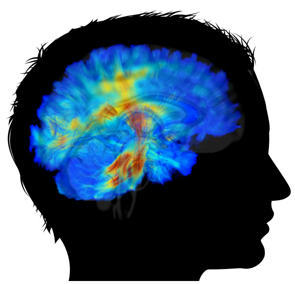

Subjects participated in battery of cognitive tests

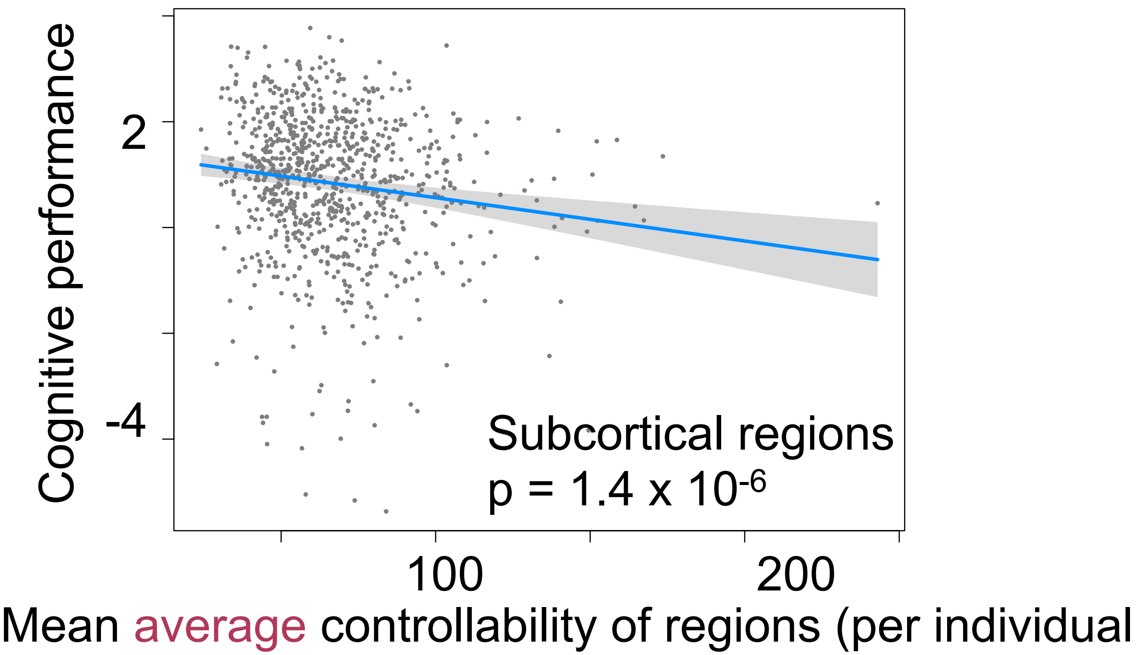

Consistent with evidence that segregation between neural systems is associated with improved cognitive ability

Wig, TICS 2017

Synchronous neural activity is often associated with pathology

Recording from a scalp electrode during a seizure

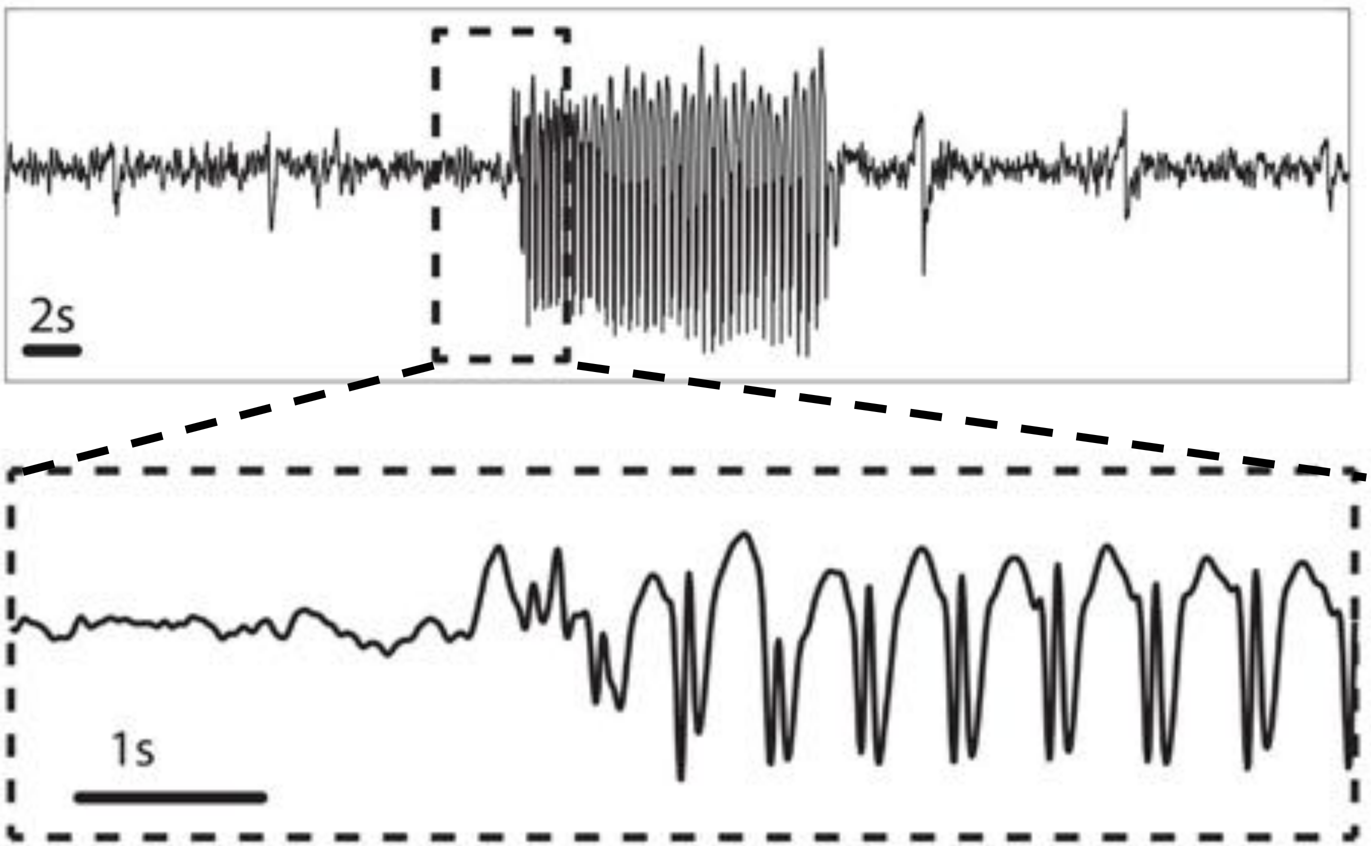

Taylor et al., *Front Neurosci*, 2015

Synchronizability measures network ability to sustain globally similar dynamics

Synchronous state described by Laplacian connectivity matrix and its eigenvalues $\{\lambda_j\}$

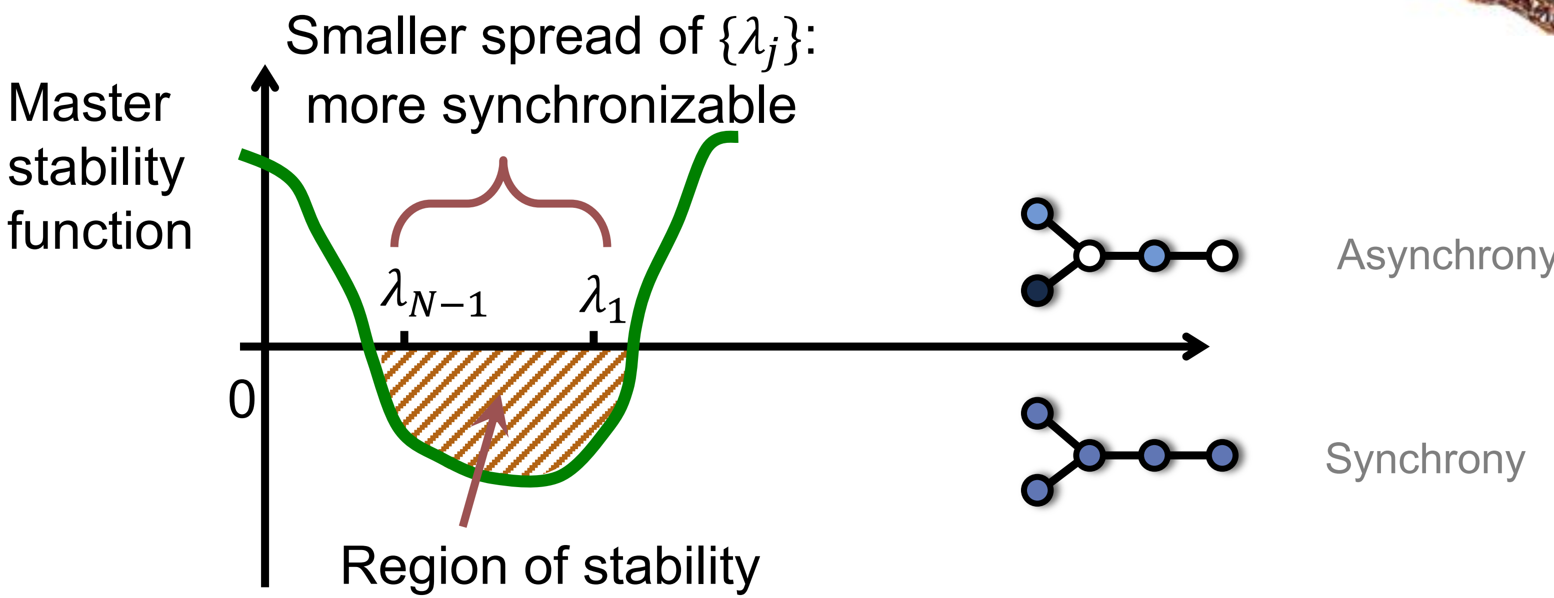

Pecora and Carroll, *Phys Rev Lett* 1998

Do brain networks have less susceptibility to synchronous (potentially pathological) dynamics with age?

Topology of brain networks is less synchronizable

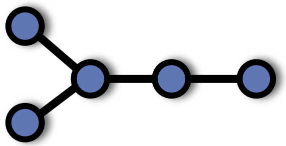

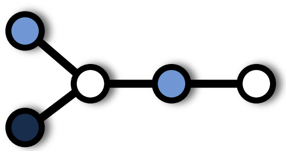

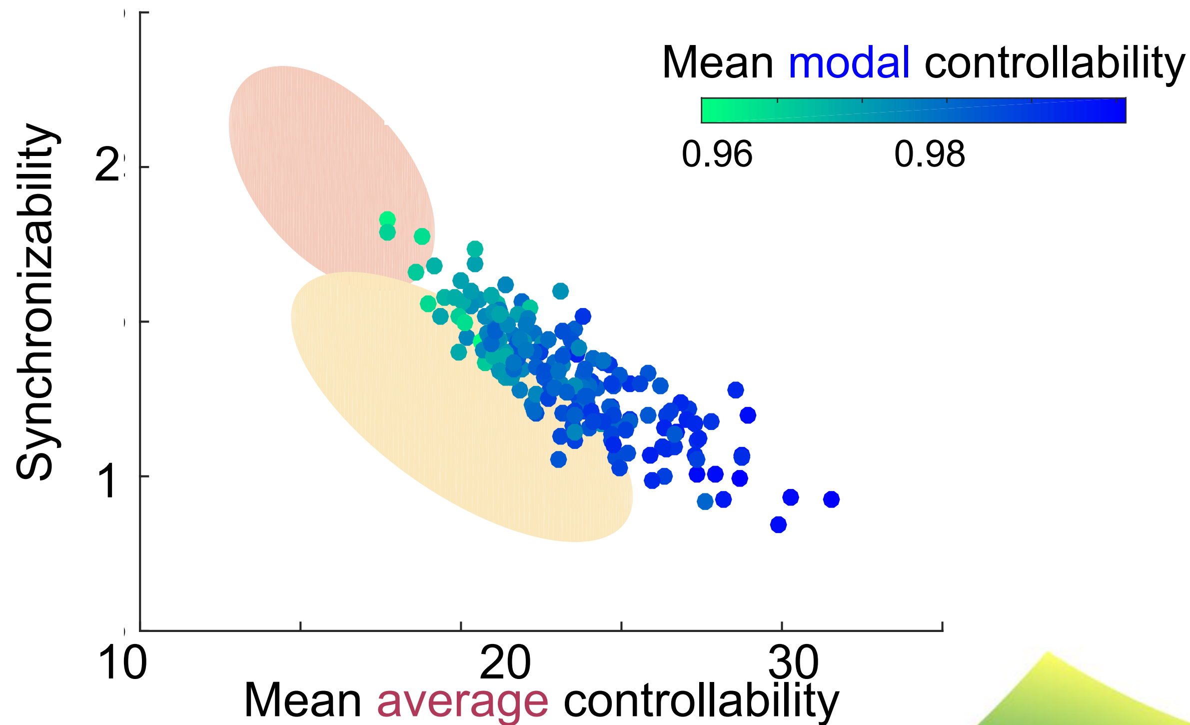

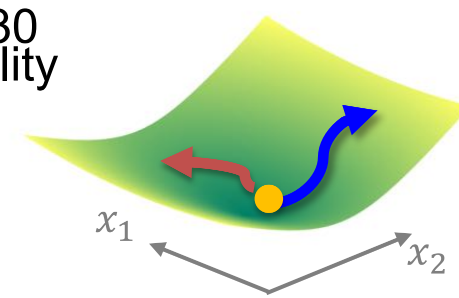

Brain networks less vulnerable to synchrony with development

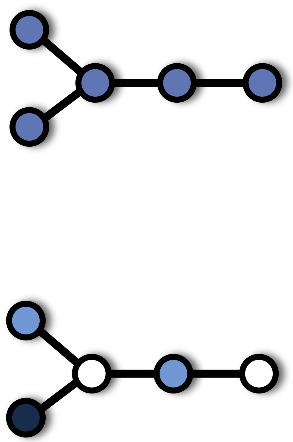

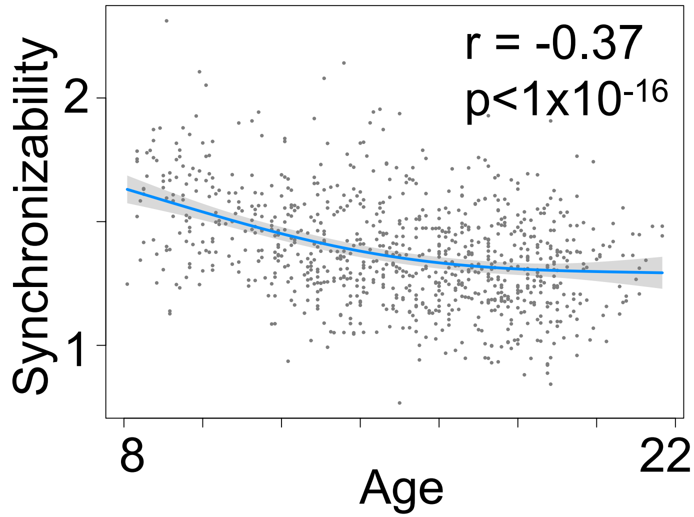

Older subjects show less susceptibility to synchronous dynamics

Quantifies intuitions on emerging control and decreased susceptibility to global inputs with age

A phenomenological wiring rule that promotes healthy development?

If these findings suggest a mechanism for development

Rewire brain networks for higher controllability and lower synchronizability

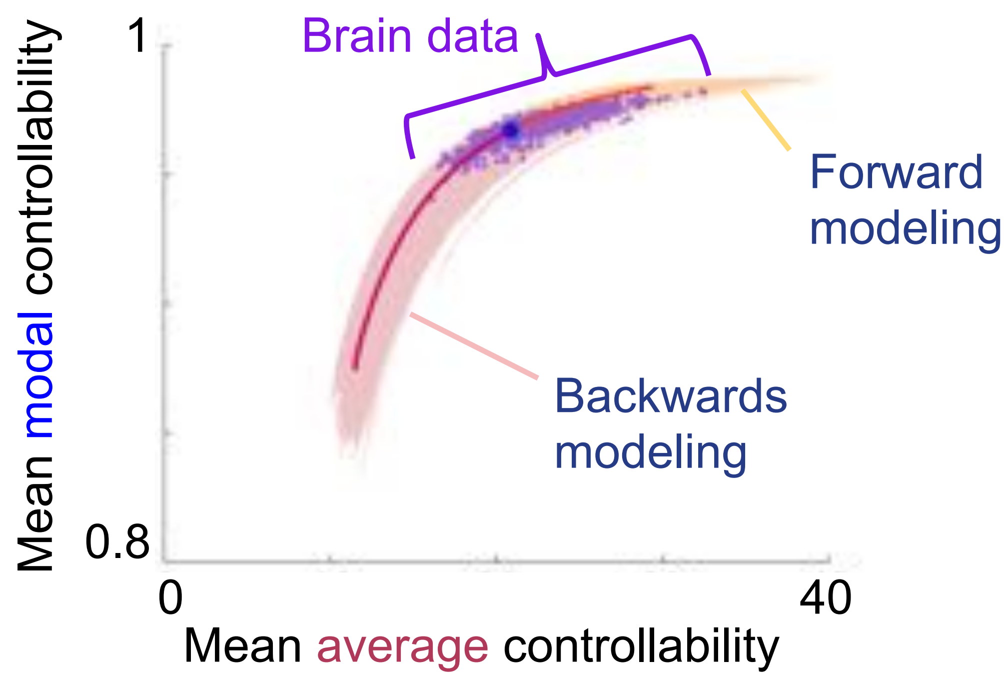

Simulations recapitulate developmental arc

Tang et al., *Nat Comm* 2017

Topology of brain structure and changes across development

1. We develop dynamical models that link changes in white matter to predicted dynamics and function.

2. Our results formalize intuitions about increasing specialization of brain connectivity across development, at the expense of greater flexibility.

3. We identify regional changes and drivers of cognition.

Acknowledgements

Complex Systems Group

Danielle Bassett

Chad Giusti

Marcelo Mattar

Living Matter Physics

Ramin Golestanian

Lorenzo Piro

Laura Collesano

Emergent phenomena across scales

Space

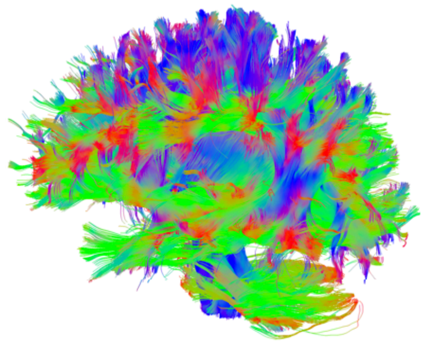

Brain networks and control

Tang et al., *Nature Comm* 2017

Tang & Bassett, *Rev Mod Phys* 2018

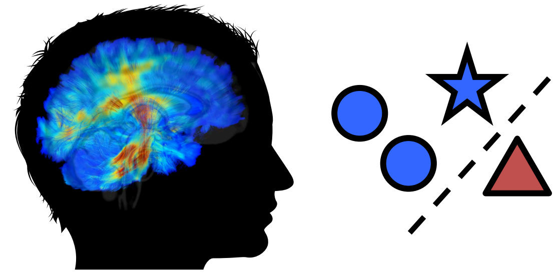

Effective learning

Tang et al., *Nature Neuro* 2019

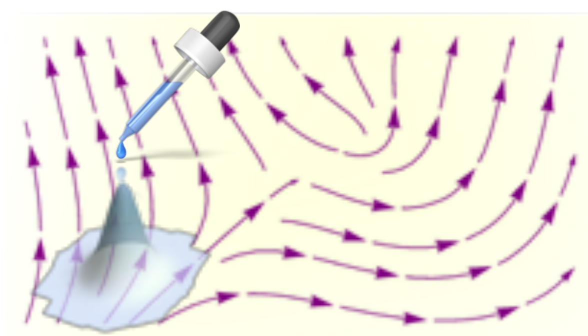

Information in fluid flows

Tang & Golestanian, *arXiv* 2019

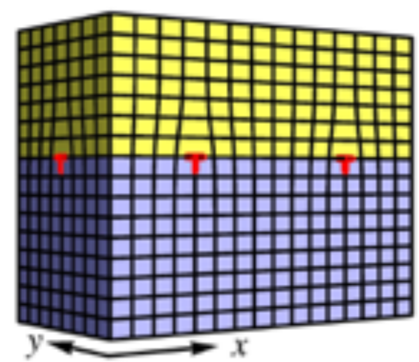

Topological phases of matter

Tang and Fu, *Nature Phys* 2015

Tang et al., *Phys Rev Lett* 2012

Time