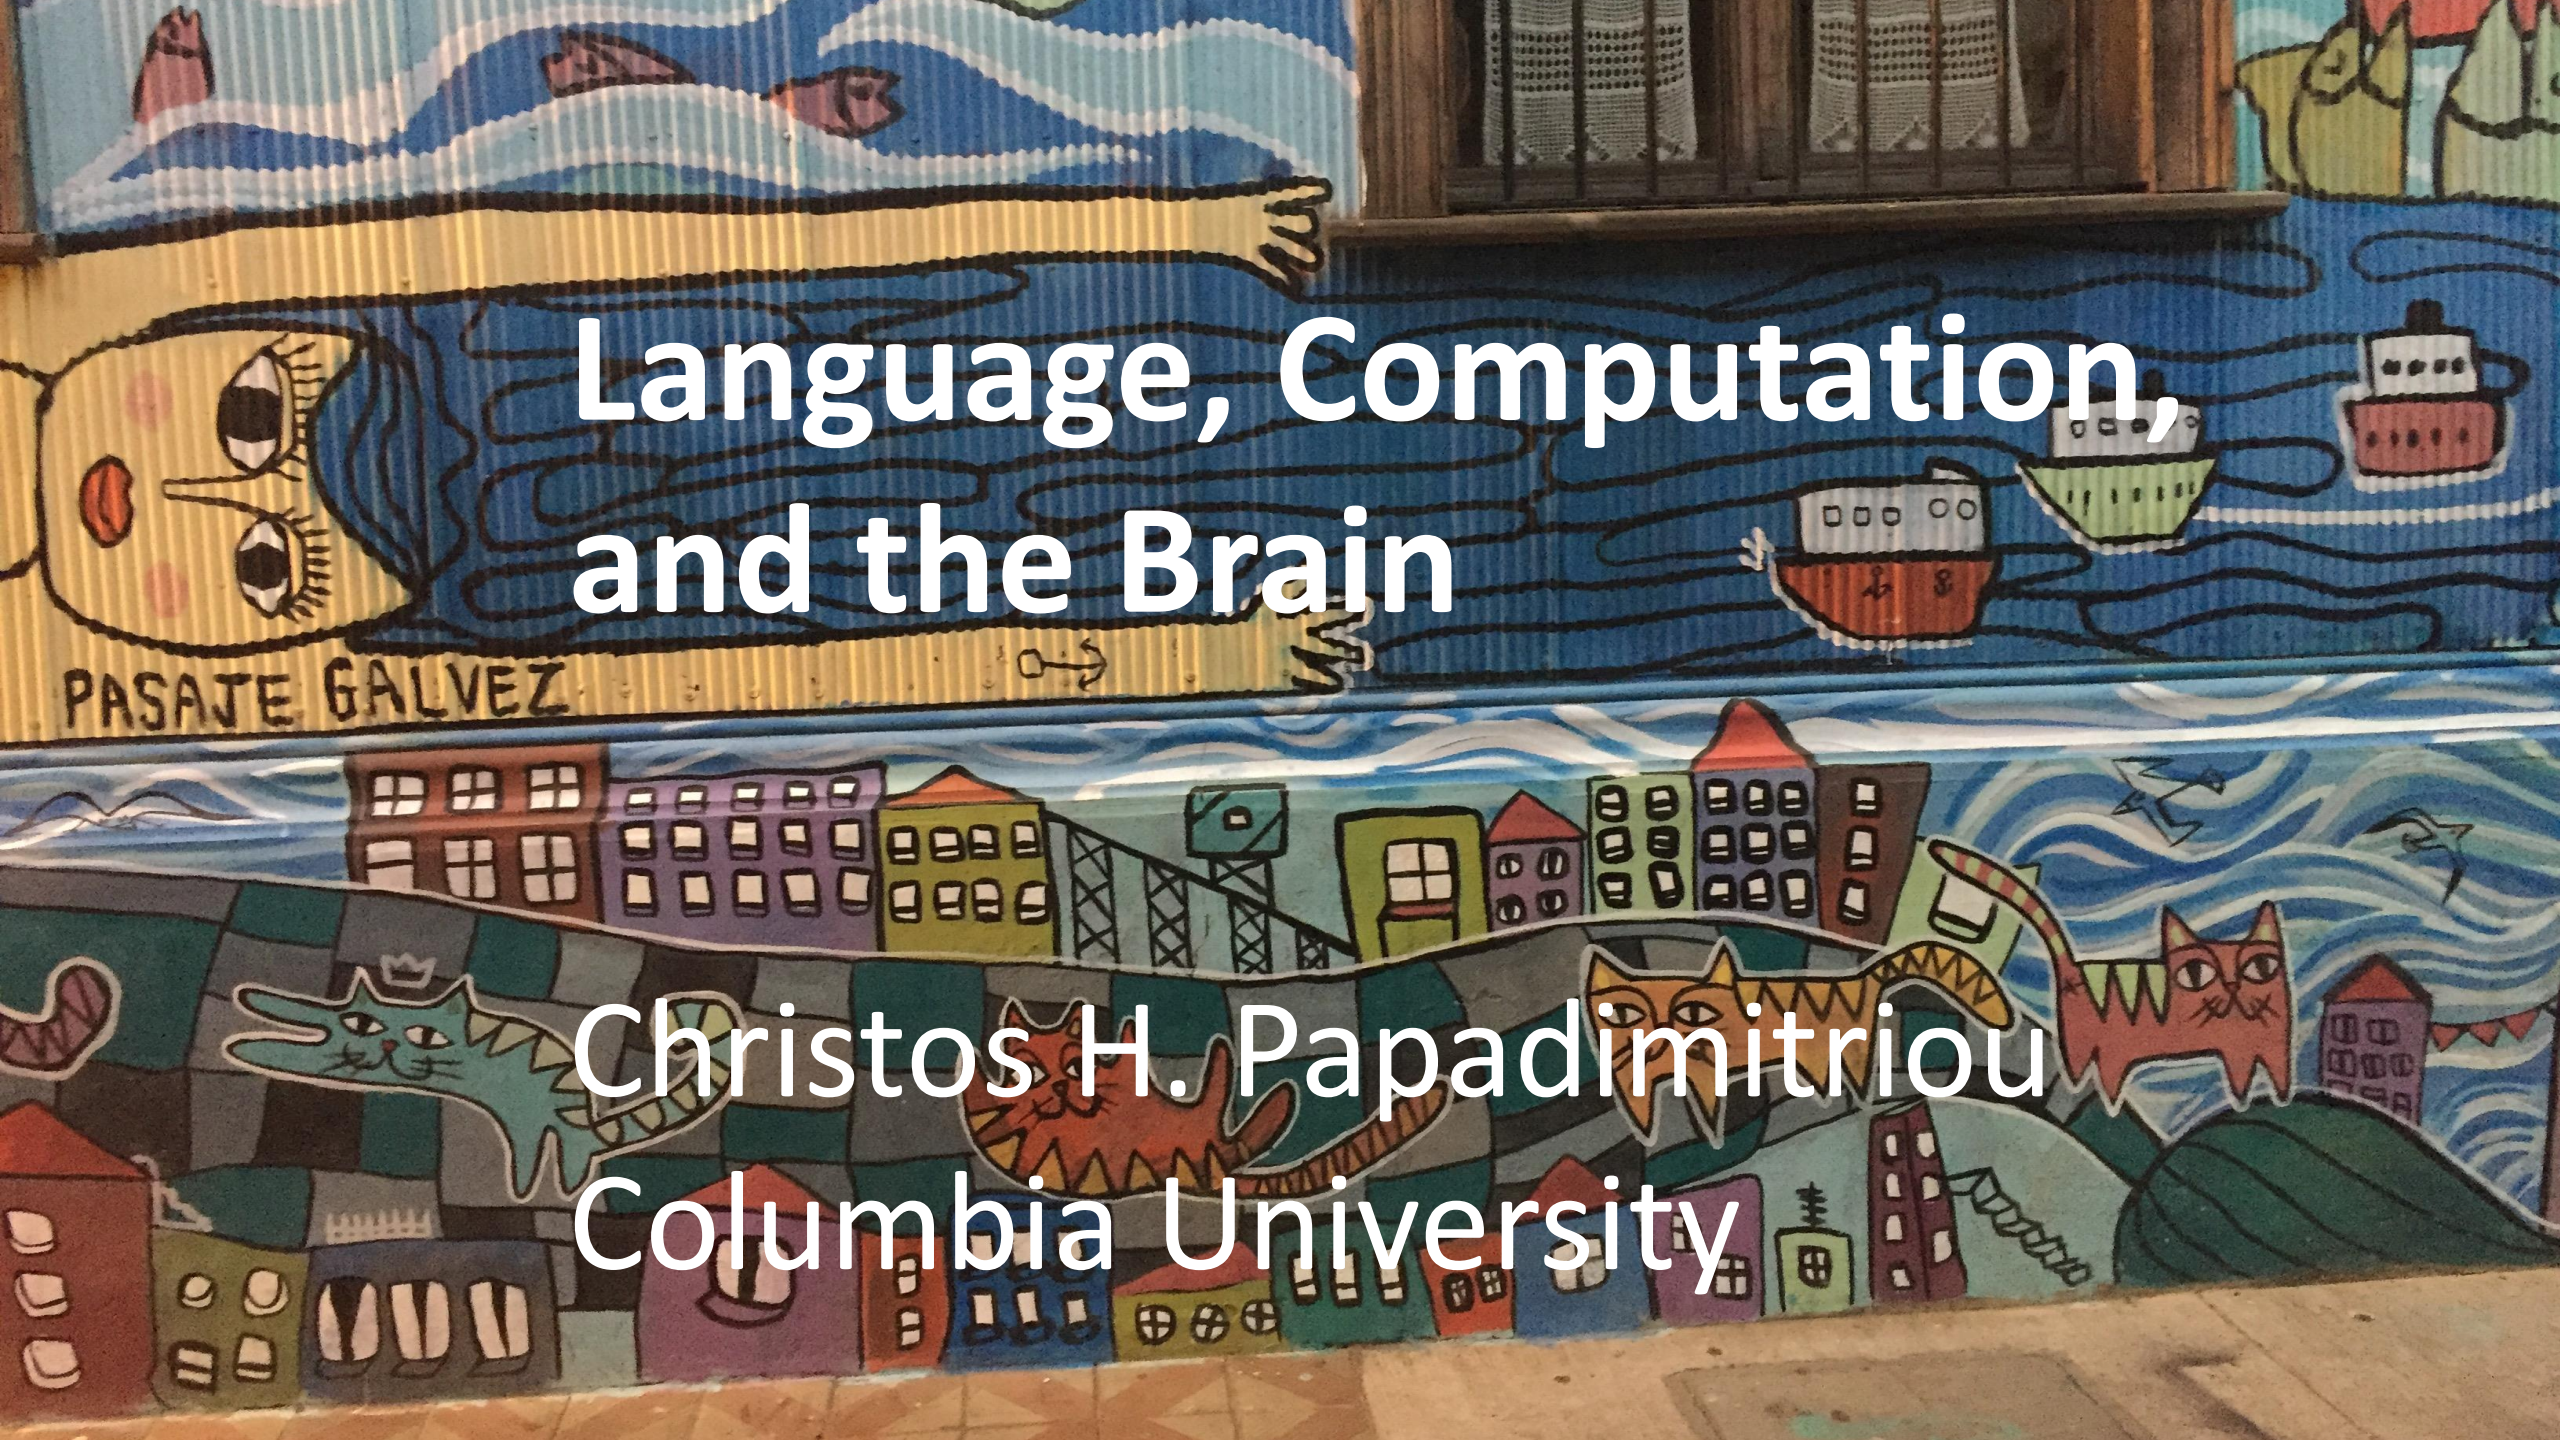

Language, Computation, and the Brain

Christos H. Papadimitriou
Columbia University

How does language happen?

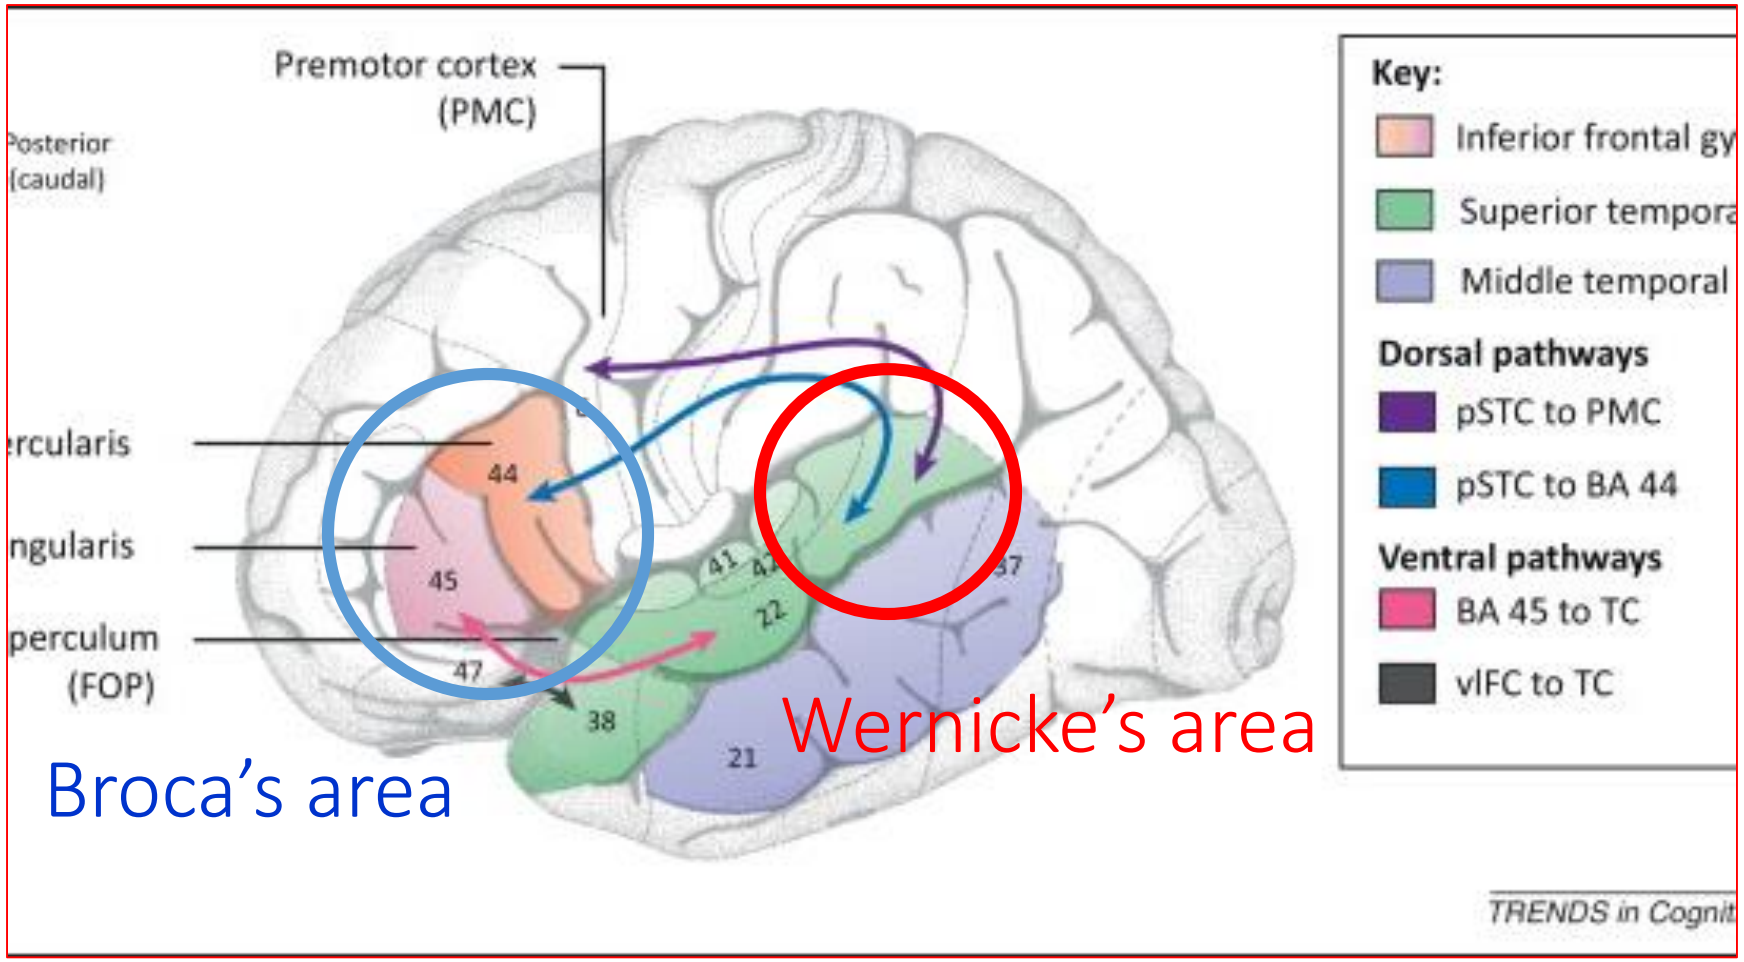

The [Poeppel 2016] experiment

fret ship hill give true melt fans blue guess hits then cats

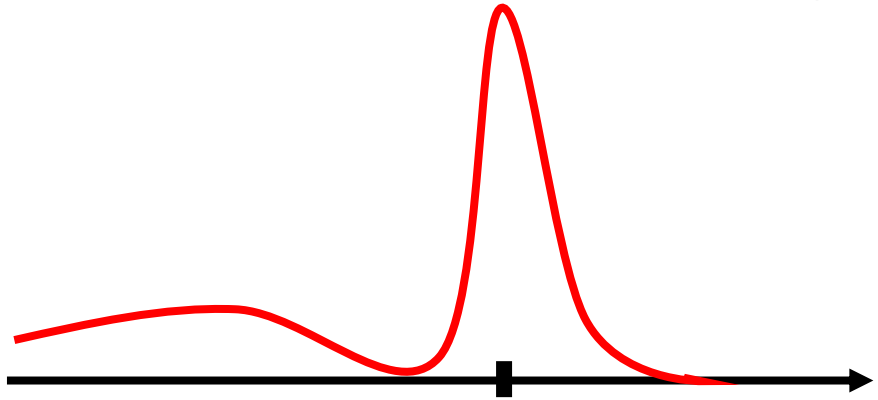

4 hertz

The [Poeppel 2016] experiment, stage II

bad cats eat fish new plan gave joy little boy kicks ball

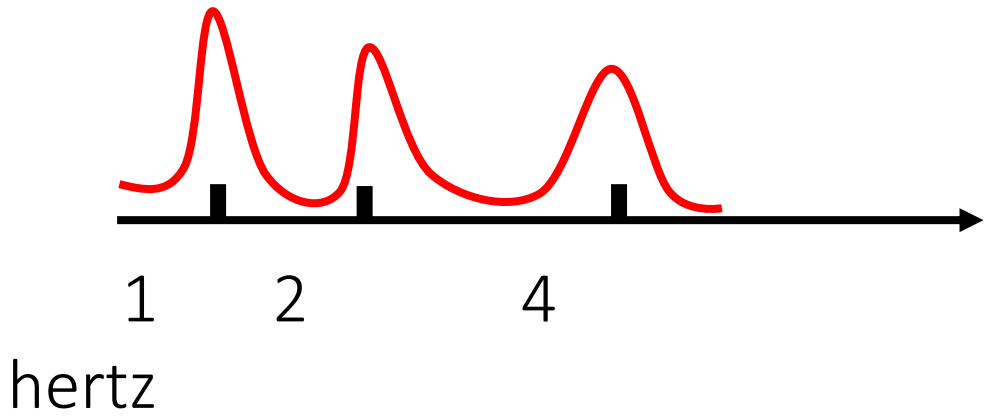

My interpretation

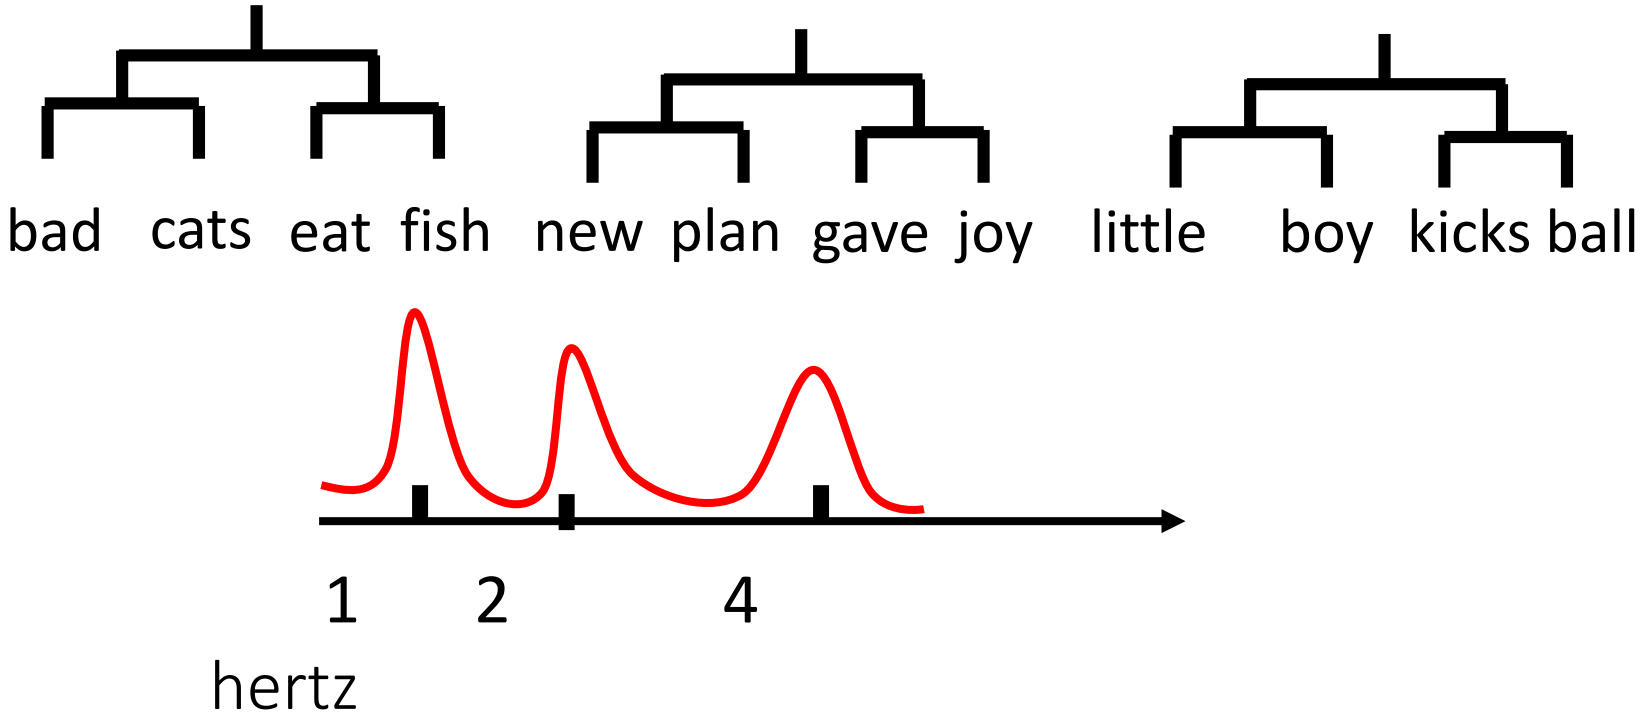

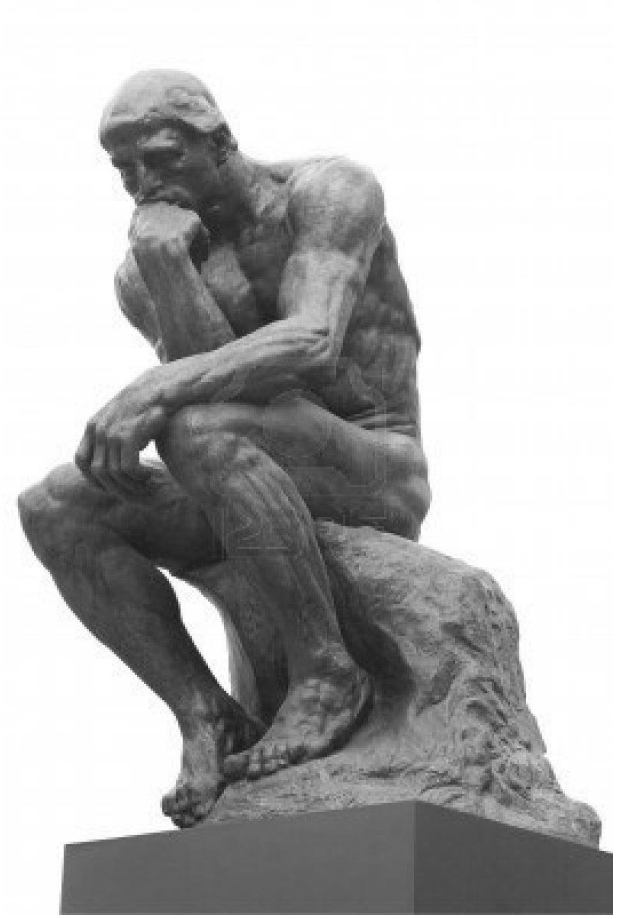

How are these tree representations created?
in about a dozen spikes per step?

(Why a dozen?
neuron spikes ~ 50Hz :
language ~ 4Hz)

[Frankland & Greene PNAS 2015]

"The ball hit the truck"

vs

"The truck hit the ball''

Different areas of the Superior Temporal Gyrus (STG) responded to "truck" in the two sentences

The first area also responded to

"The truck was hit by the ball"

Zaccarella & Friedericci “Merge in the human Brain” *Front. Psych. 2015*

- The completion of phrases, and especially of sentences, activates parts of Broca’s area

The Language Organ

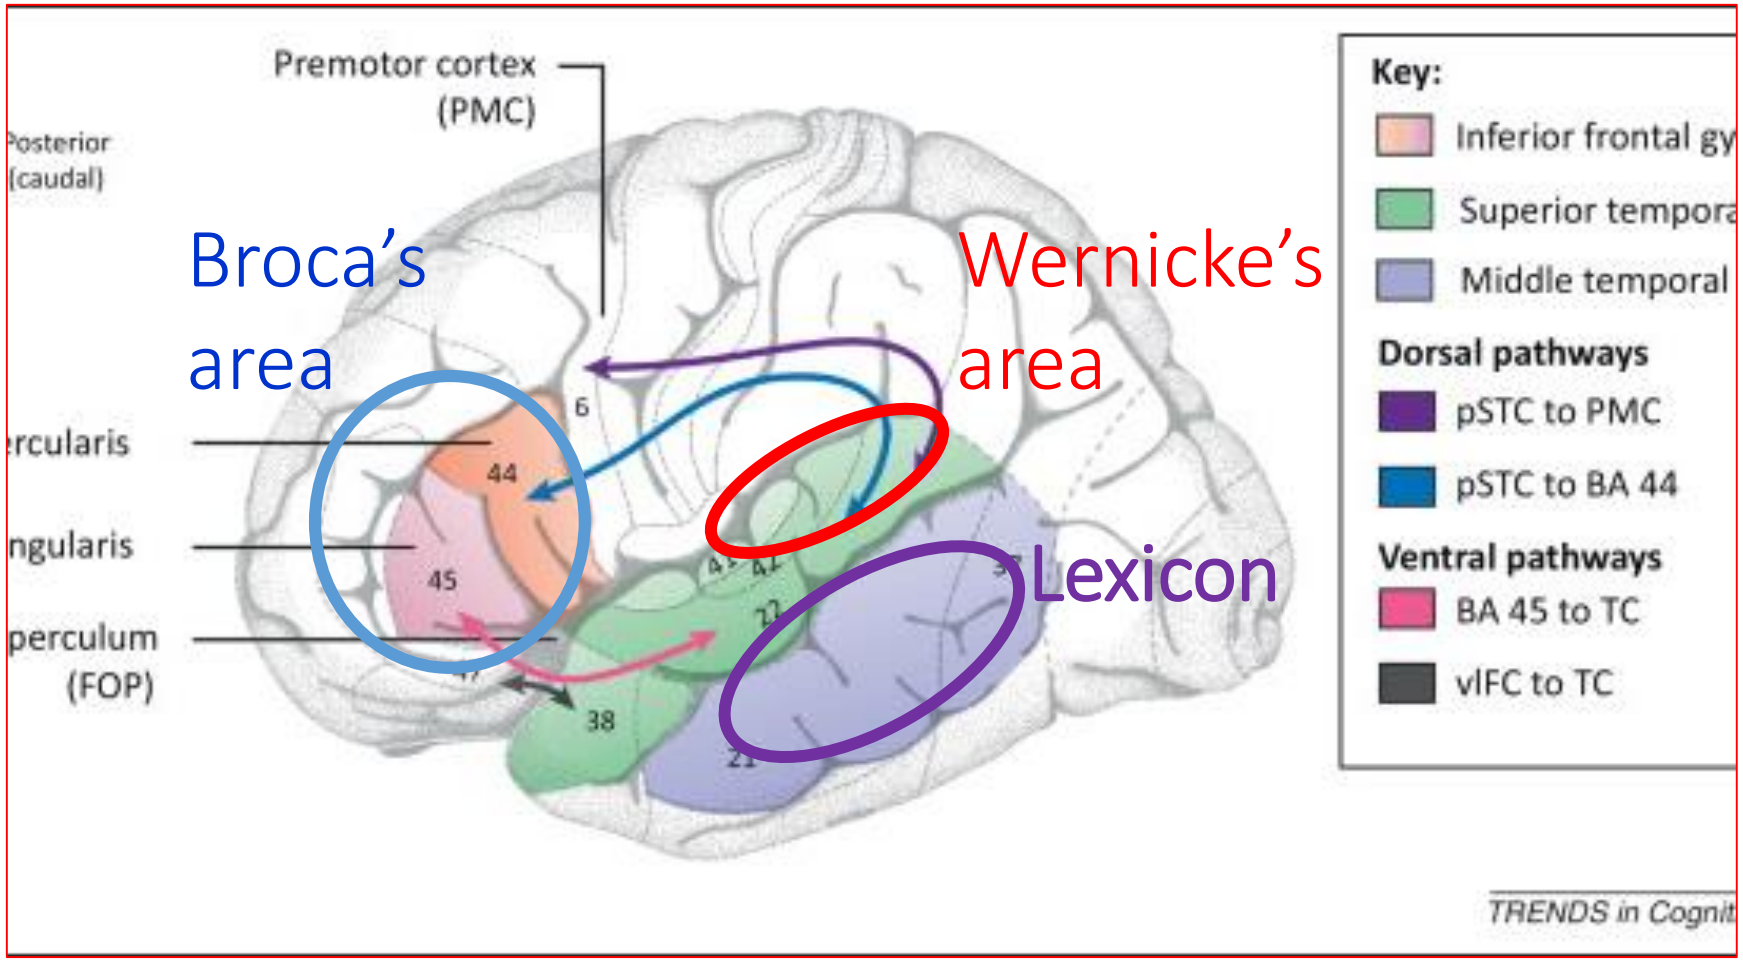

Now let’s go ten miles high

How does the Mind

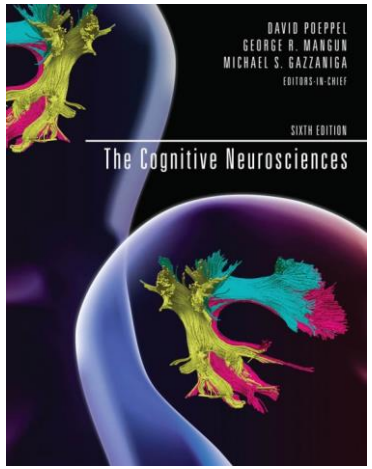

emerge from the Brain?

*"…we do not have a **logic** for the transformation of neural activity into thought … I view discerning [this] logic as the most important future direction of neuroscience…"*

Richard Axel, *Neuron*, Sep 2018

So, what kind of formal theory would qualify as Axel's *logic*?

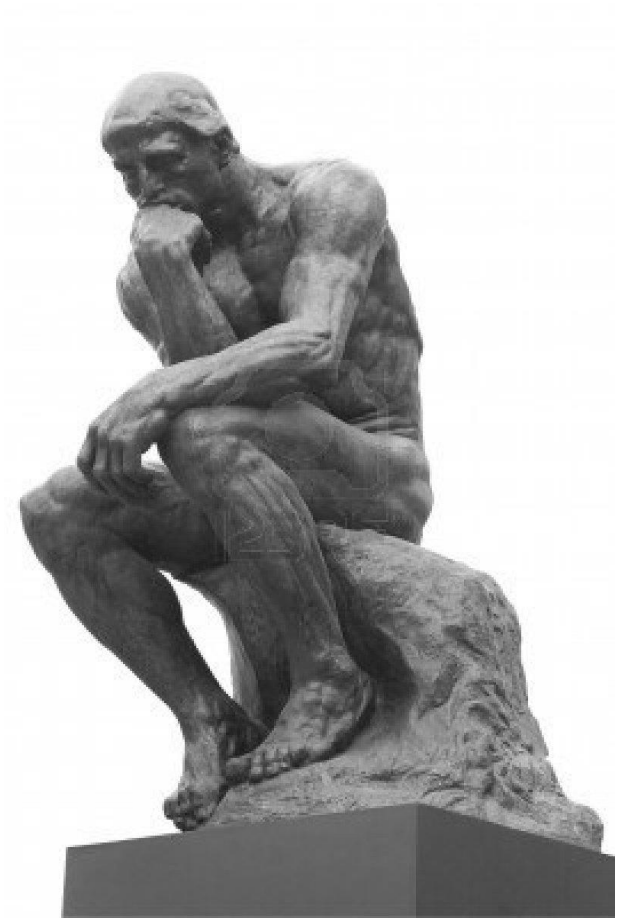

Computation in the brain: What is the right level?

- Whole brain?

?Assemblies of neurons!

- Spiking neurons and synapses?
- Dendrites?
- Molecules?

1980s

cell

1990s

cell**S**

Short history of assemblies (or *ensembles*) of neurons

- [Hebb 1949, Harris 2003, 2005; Buzsaki 2008, 2010, Yuste et al. 2019,…]
- Assembly: A large ***and densely intraconnected*** set of excitatory neurons in a brain area whose firing (in a pattern) is tantamount to our thinking of a particular memory, concept, person, name, word, episode, etc.
- ***G. Buzsaki 2020: "assemblies are the alphabet of the brain"***

CP et al "Computation by assemblies of neurons" *PNAS,* June 2020: **The assembly hypothesis**

- There is an intermediate level of brain computation
- Call it the Assembly Calculus
- It is implicated in higher cognitive functions such as reasoning, planning, language, story-telling, math, science, music, …
- Assemblies of neurons are its basic representation – its "data type"

The assembly hypothesis (cont.)

What are the operations of the Assembly Calculus?

- Projection
- Association
- Merge
- Bind (or reciprocal projection)
- ...plus control commands: Read, Fire, Inhibit, Disinhibit

Projection of an assembly: (RP&C)*

Assembly

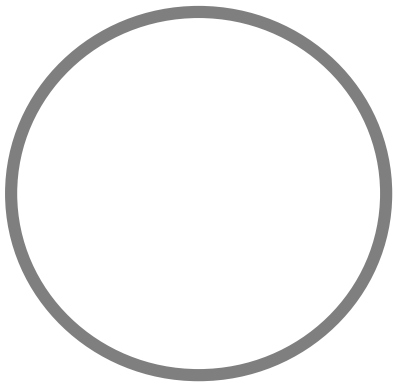

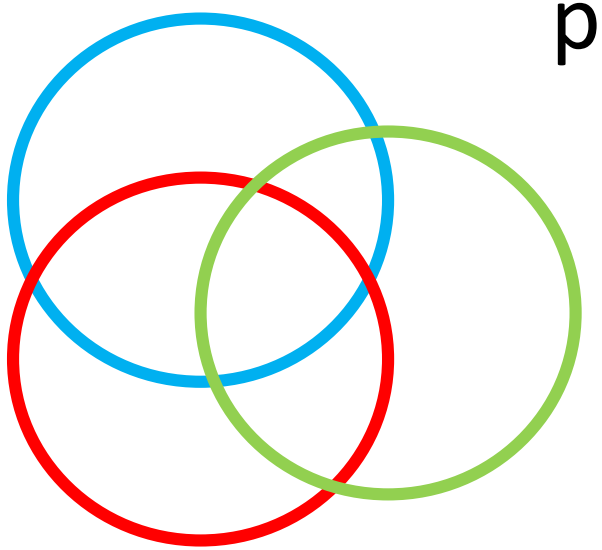

Also, Hebbian plasticity

Q: Does this process converge?
After a dozen or so steps?
And is the result densely connected?

math model?

A mathematical model of the brain

Finite number of brain regions

Each contains n neurons

Inhibition: only k fire (selection)

Some pairs of areas are
connected by directed **$B_{n,p}$**
(= bipartite $G_{n,p}$)

All are recurrently connected
by directed **$G_{n,p}$**

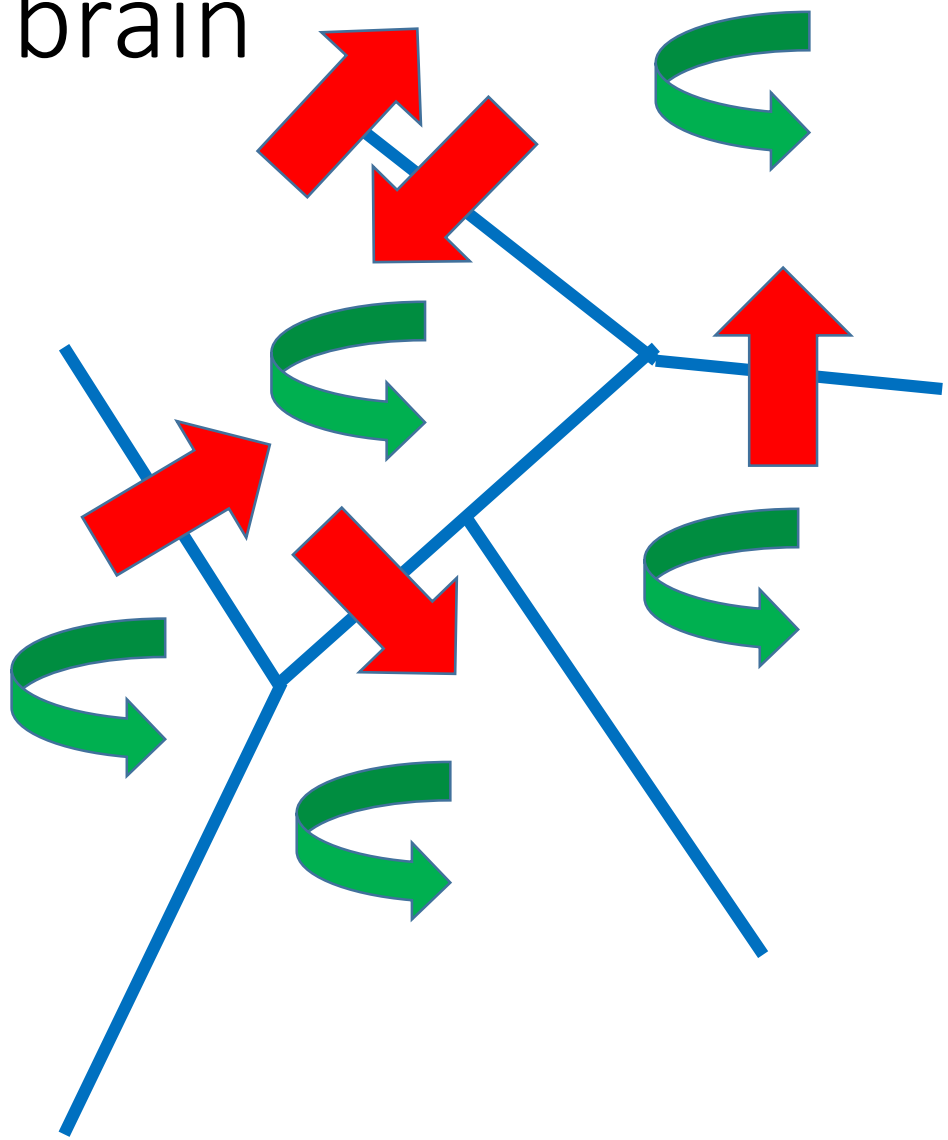

The model (cont.)

Neurons fire in discrete steps
The k with highest input
are selected to fire
this is called
“random project and cap”
RP&C
Important primitive, models
inhibition-excitation equilibrium

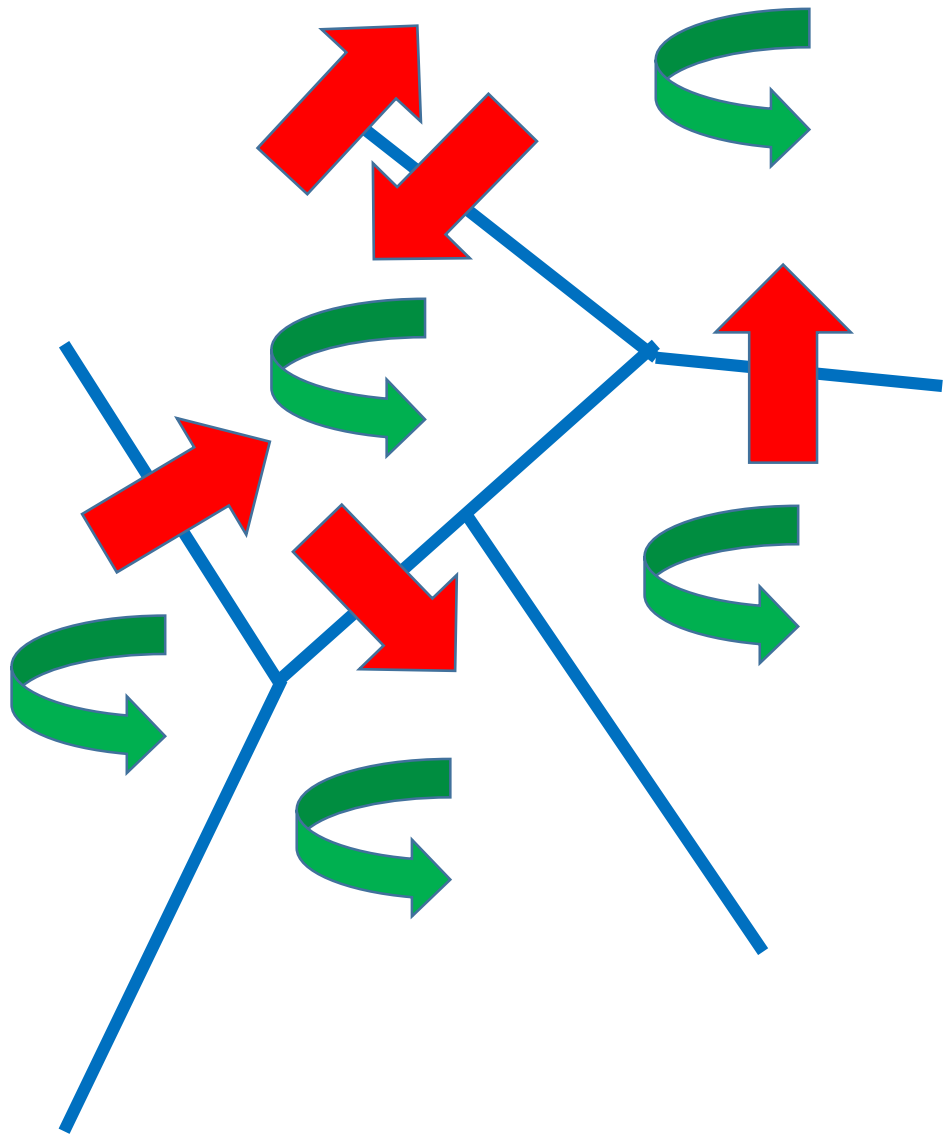

The model (cont.)

Connections between areas can be inhibited/disinhibited

Plasticity: If i→j, i fires and in the next step j fires, the weight of i→j is multiplied by (1 + β)

(also: homeostasis, forgetting…)

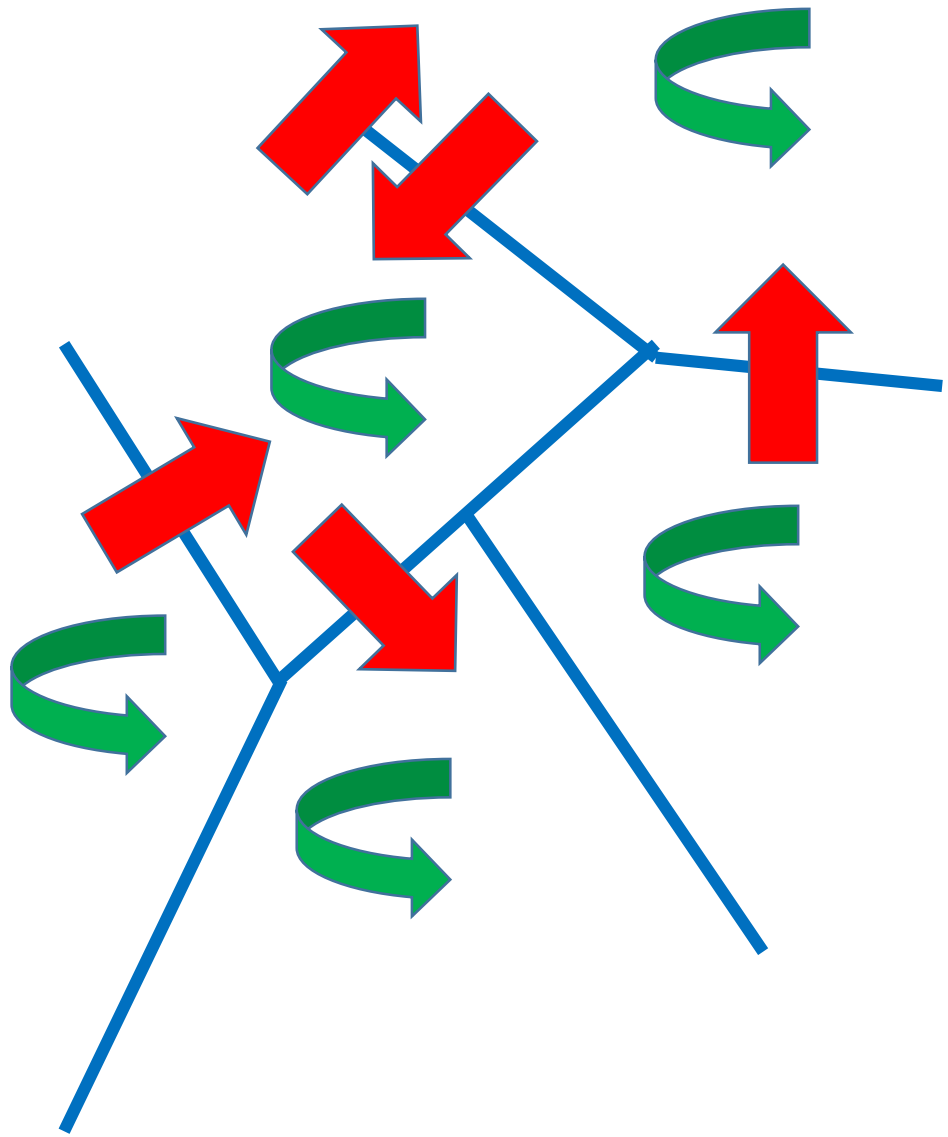

Main parameters of the model, intended values

- $n \sim 10^{7\text{-}8}$
- $k \sim 10^{3\text{ - }4}$
- $p \sim 0.001$
- $\beta \sim 0.10$

Main ideas: randomness, selection, plasticity

Basic operation: RP&C

But, is the model realistic?

- Discrete time steps are unrealistic but not distortive: comparisons with asynchronous models
- The RP&C operation is stylized but natural
- Assemblies and plasticity have not been thought to work at this time scale (~ 2 - 4 Hz) before
- ***All said, as formal models of the brain go, I believe this is quite realistic***

The scope of the model: the association cortex

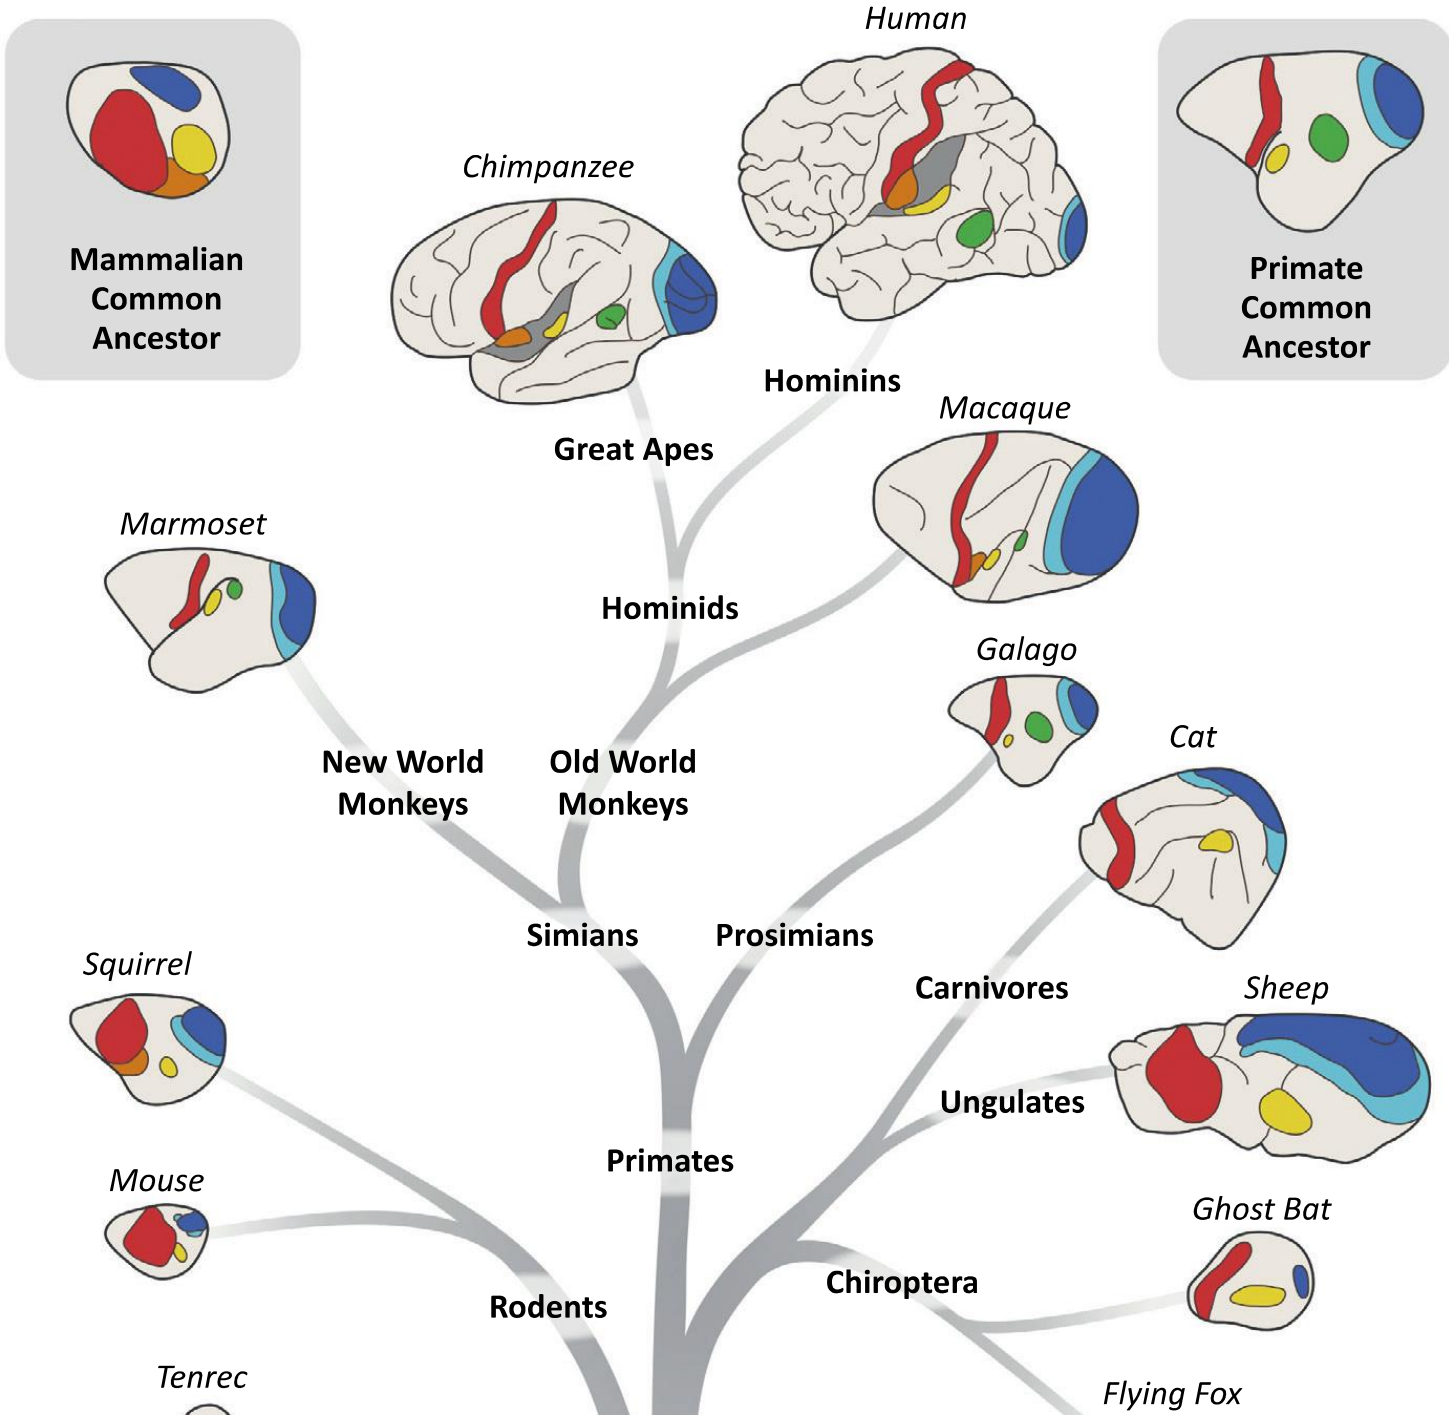

Association of two assemblies

two
stimuli

assembly
representations

The [Ison et al. 2016] experiment

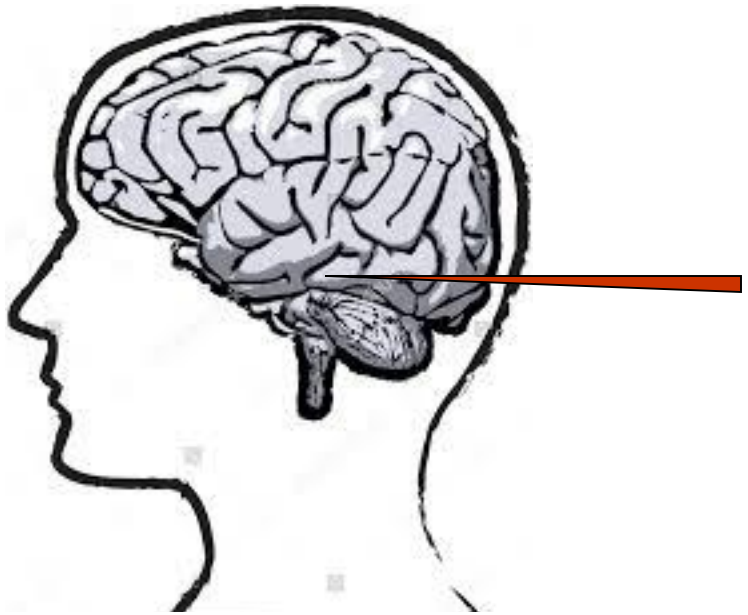

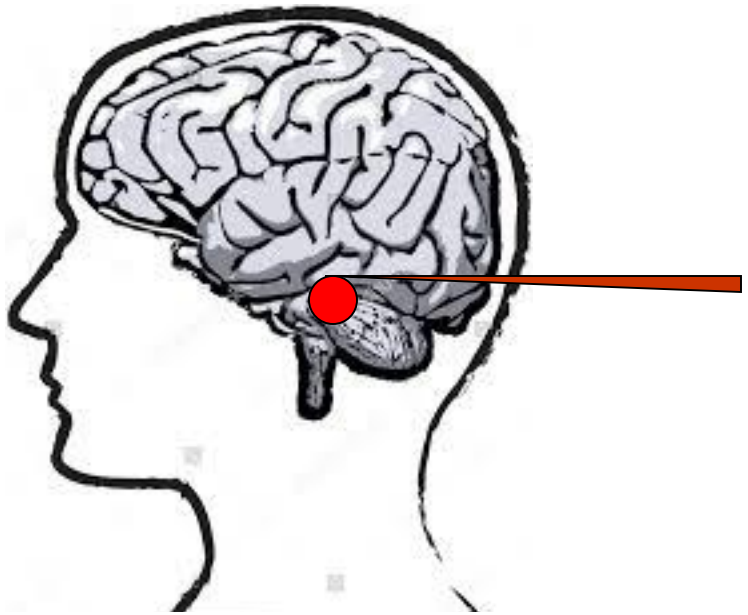

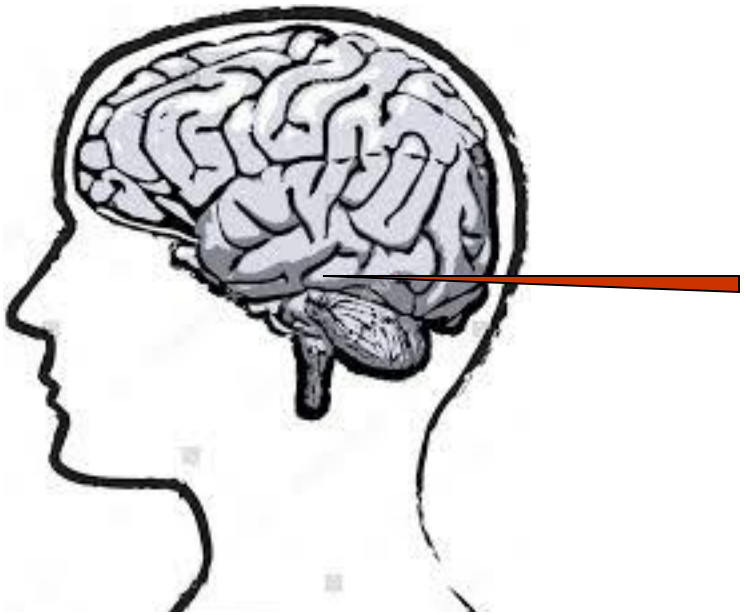

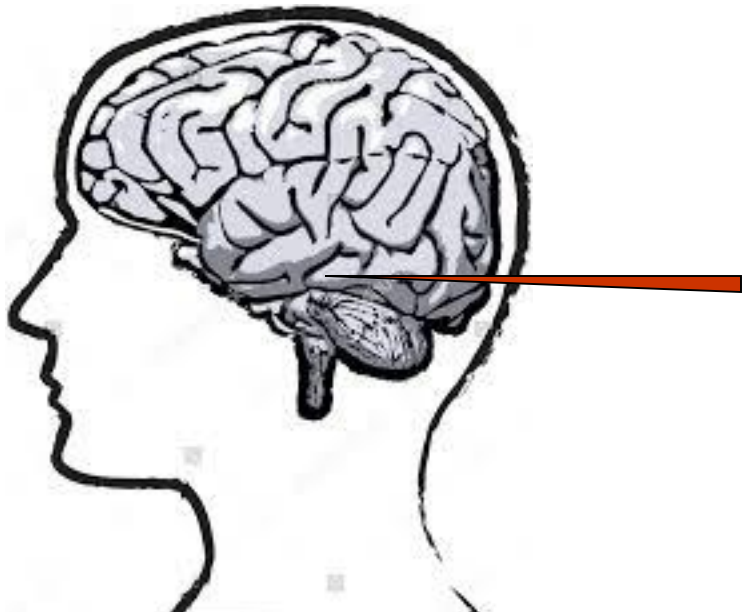

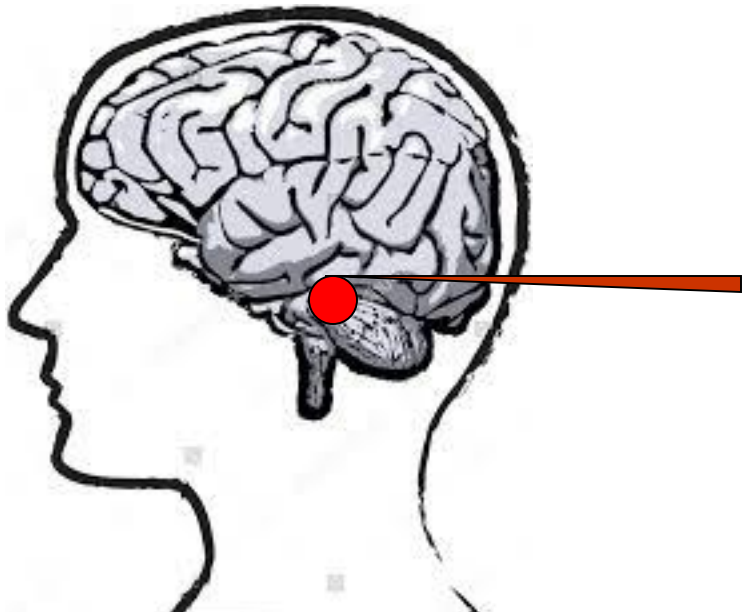

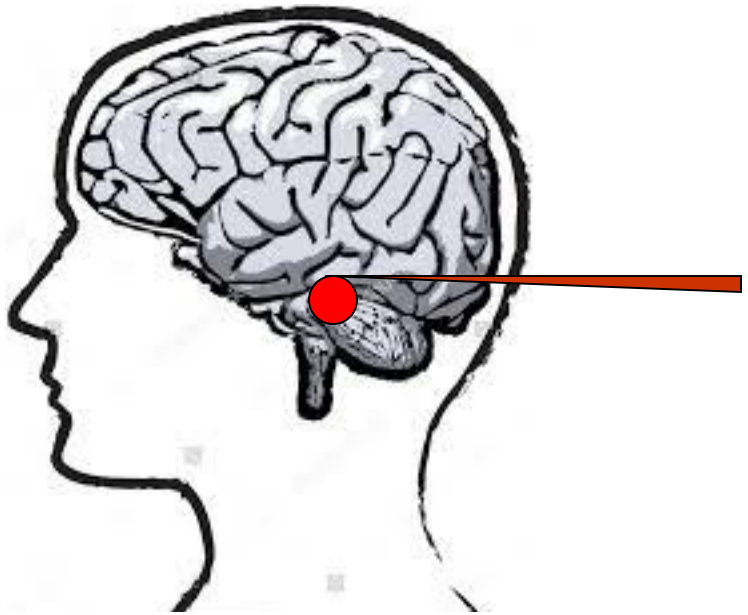

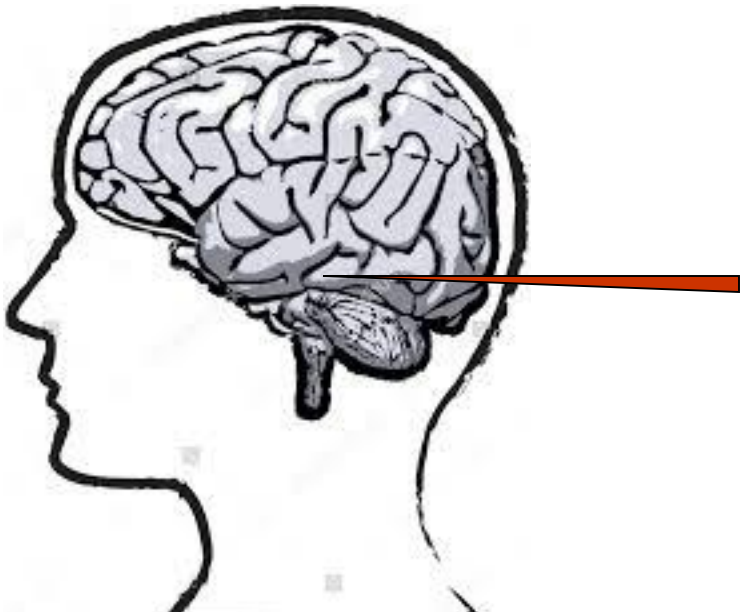

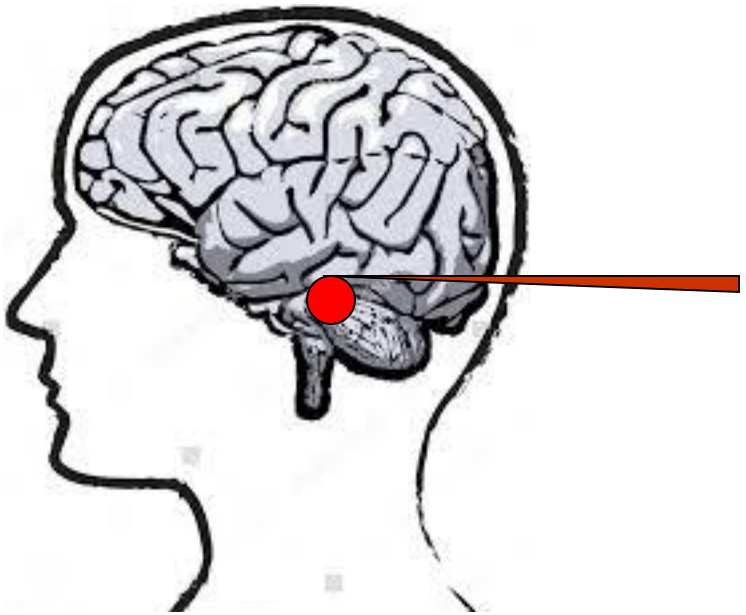

A Web of Associations

- Assemblies in a brain area overlap ***randomly***
- They also overlap as a result of association operations, to ***encode affinities***: co-occurrence, shared attributes, similarity
- Also, ***overlap preserved*** through projection
- Such overlaps may enable ***associative recall***
- And possibly other, more sophisticated modes of randomized computation…

Assembly operations are ***Real***

- ***They correspond to behaviors of assemblies that have been observed in experiments…***
- ***…or are strongly suggested by experiments, cf Merge***
- *They can be **provably compiled down** to the activity of neurons and synapses, through:*
  - *Mathematical proofs and simulations in our model*
  - *Simulations in systems of neuro-realistic spiking neurons*

Assembly Calculus recap

- project(y, B, x)
- associate(x, y)
- pattern_complete(x, y)
- merge(x, y, B, z)
- Plus: fire(x), read(A), disinhibit(A)

Q: ***How powerful is this system?***

A: ***can perform arbitrary √n-space computations***

Merge of assemblies: it’s complicated…

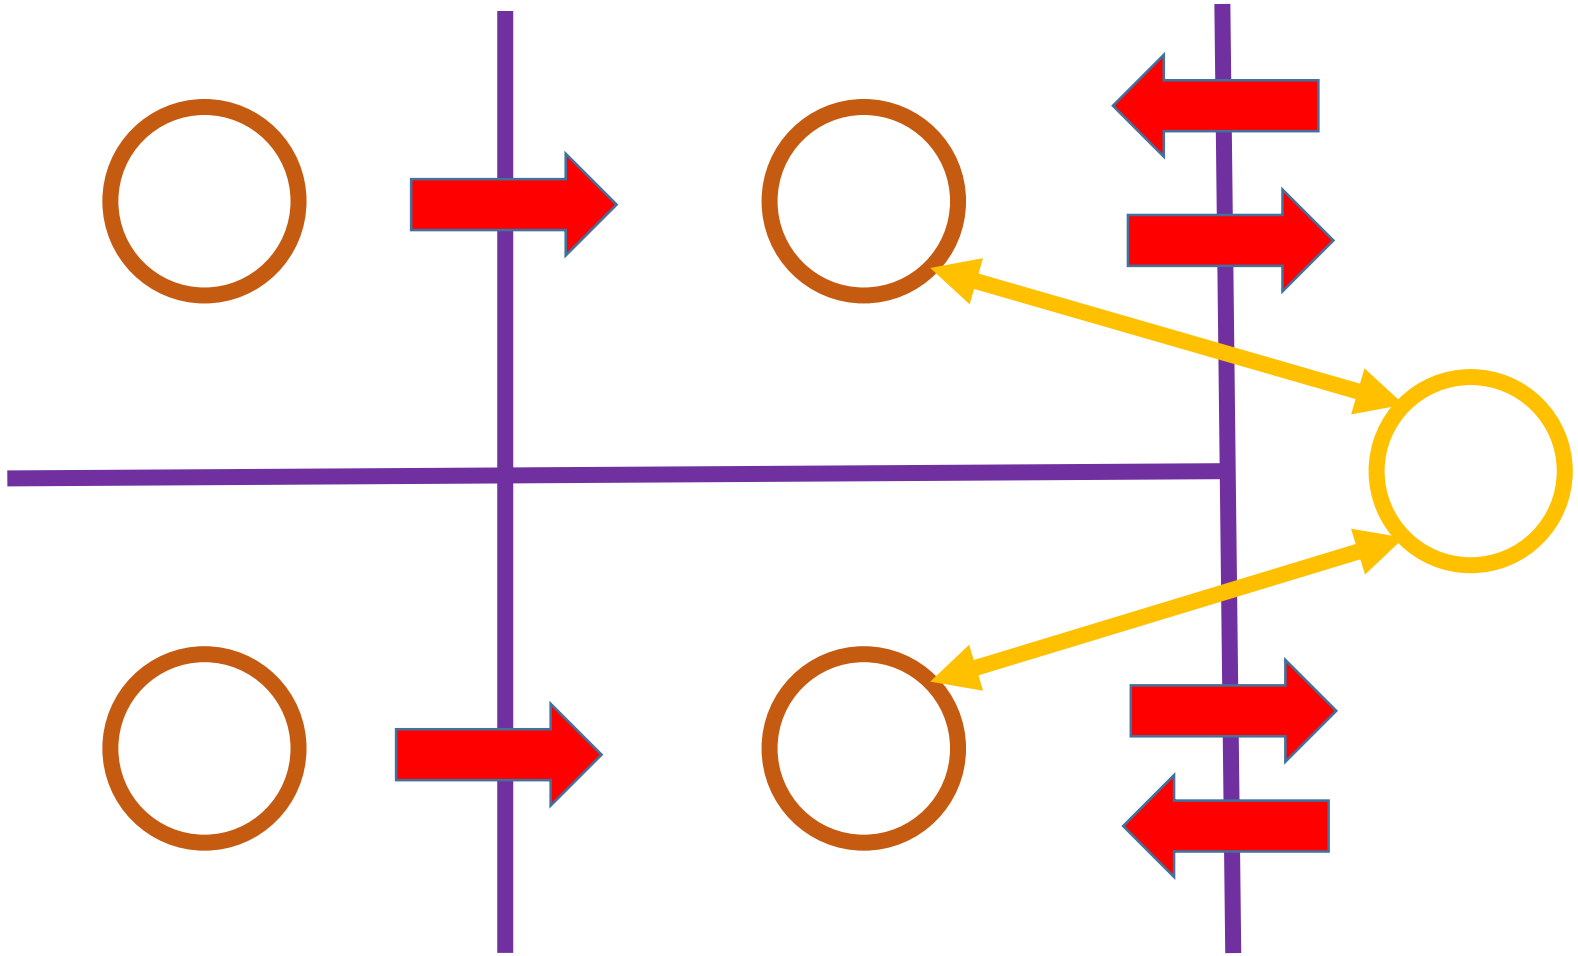

Merge: Does it need enhanced hardware? The mystery of the Arcuate Fasciculus

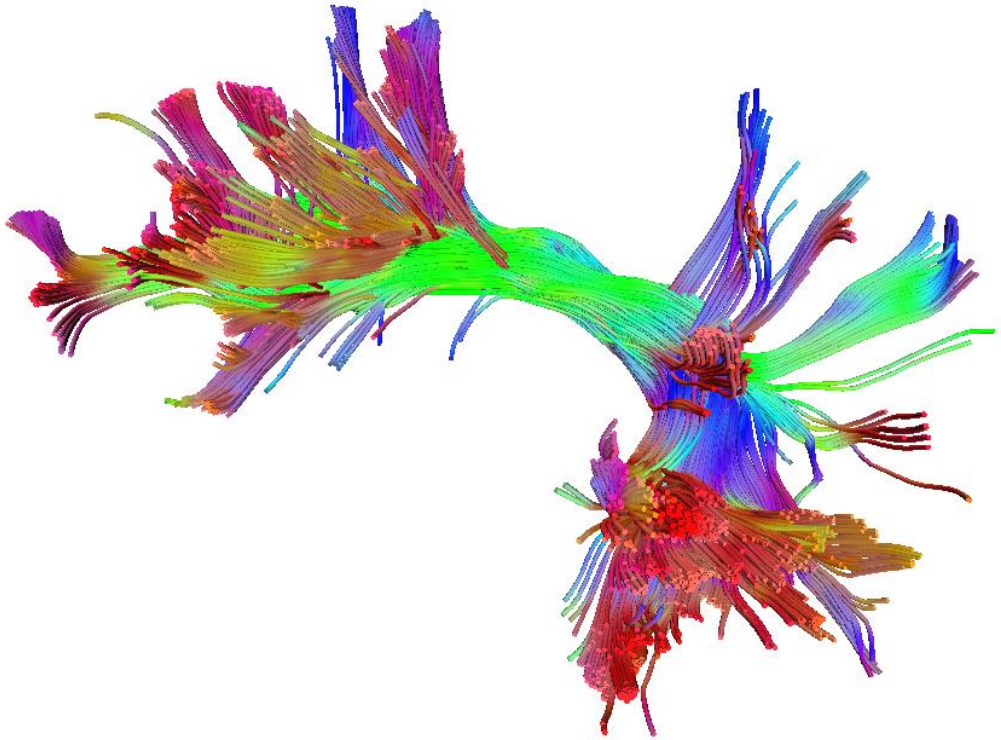

The Language Organ

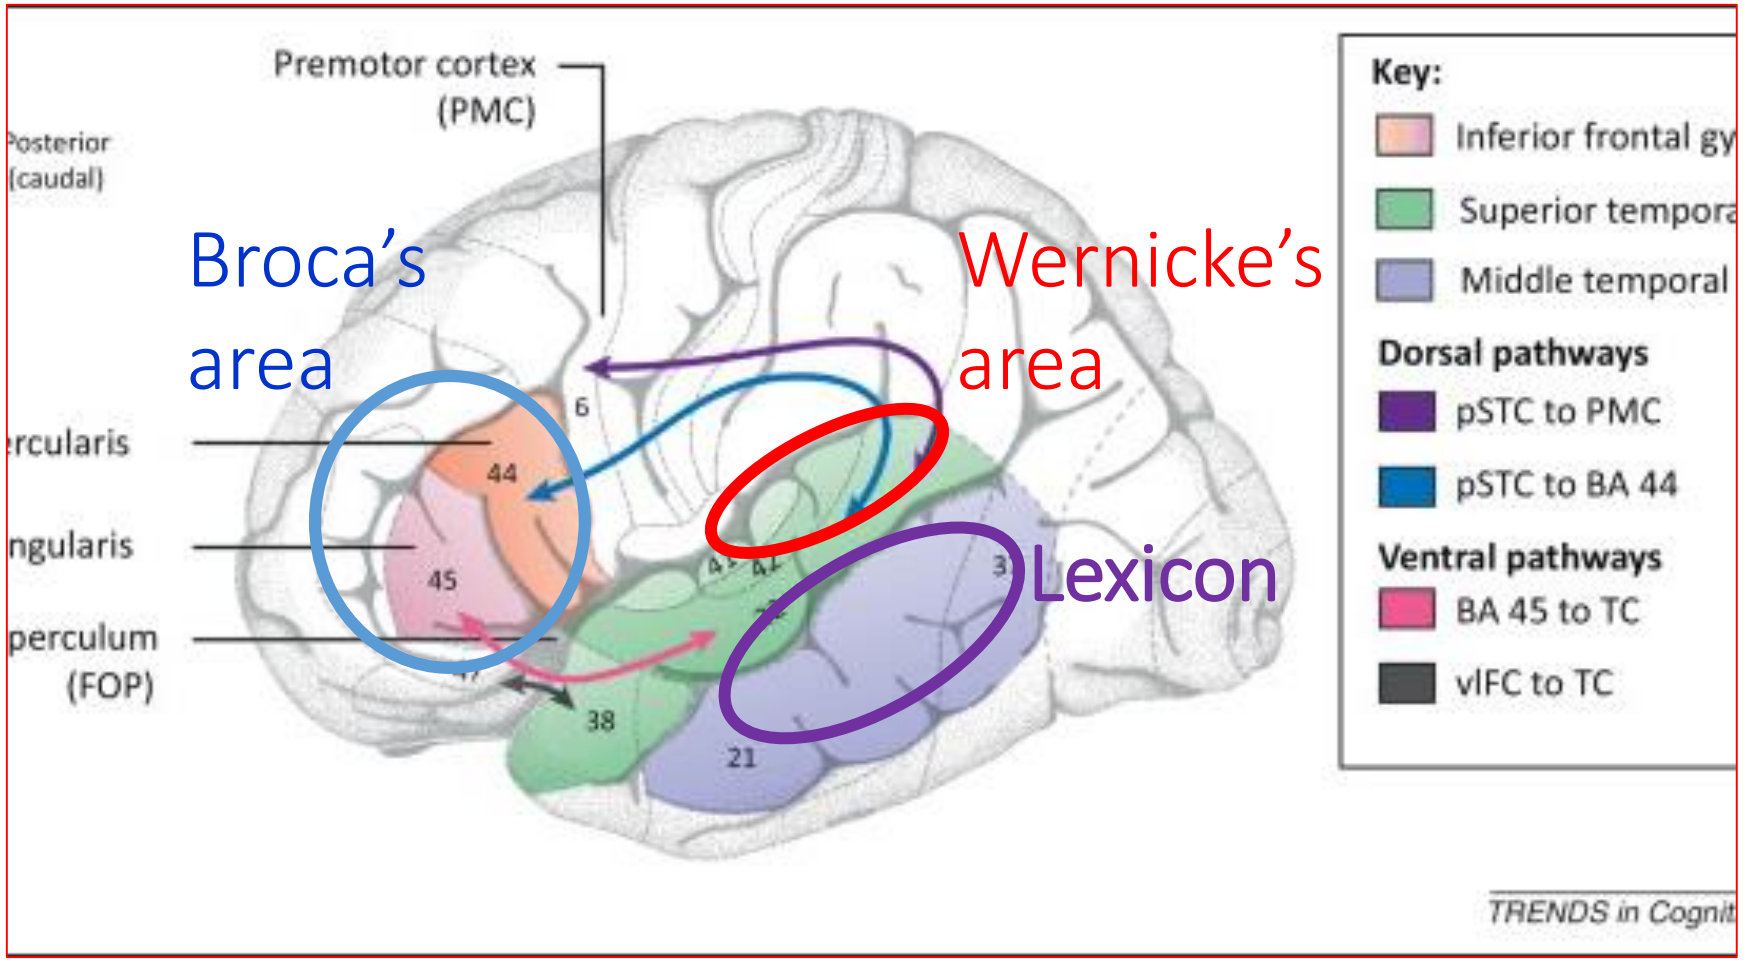

But why study Language?

- A last-minute adaptation
- 3,000 generations ago (~1% of split from chimpanzee)
- It evolved to exploit the brain's strengths
- Invaluable lens for studying the brain
- ***A deluge of recent experiments!***

Language generation through the Assembly Calculus

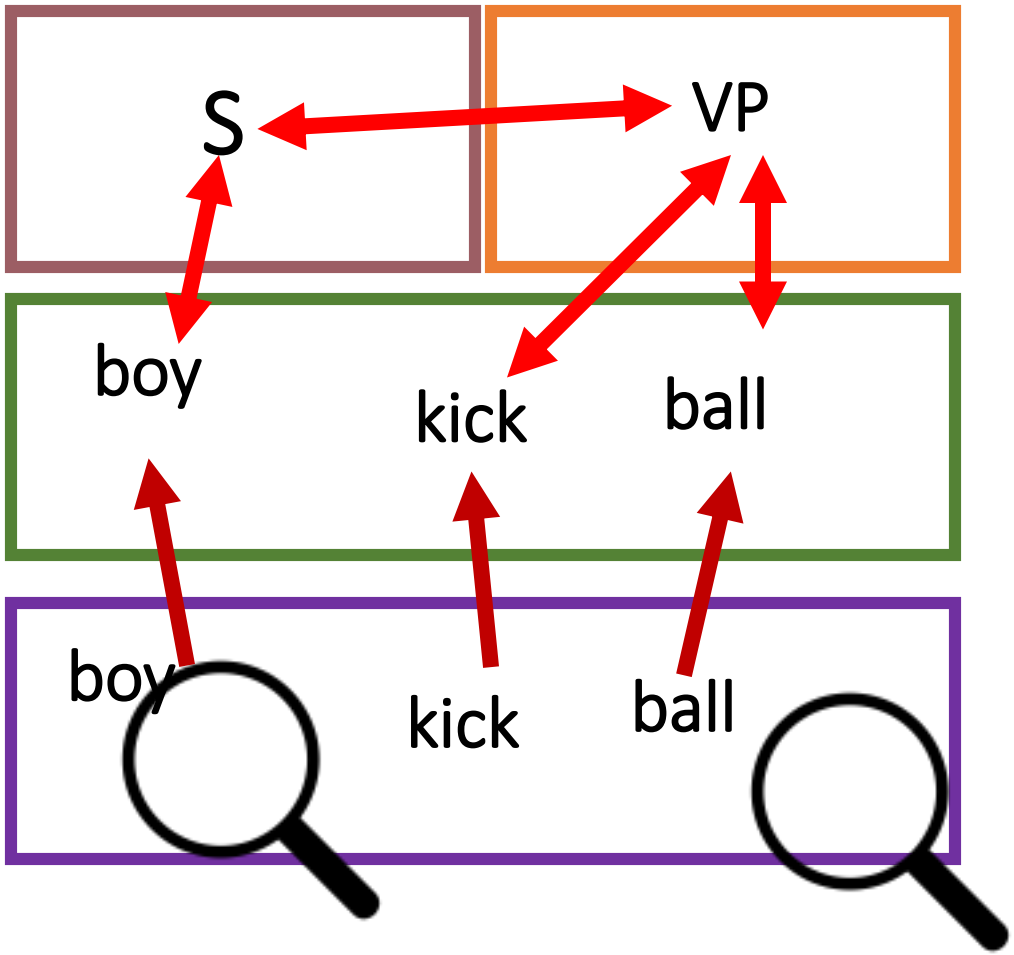

But how about comprehension?

- Some think it came much later than generation
- Reverse engineering of generation?
- Online parsing seems much more complicated than generation
- Must solve problems like ambiguity, polysemy, backtracking, recursion, grammaticality, …

A Bold Experiment: A Parser for English written in the Assembly Calculus

- Architecture: Brain Areas and Fibers
- Can be inhibited/disinhibited individually
- LEX area contains assemblies for words
- Two rival architectures, we implemented one

Parser: The architecture

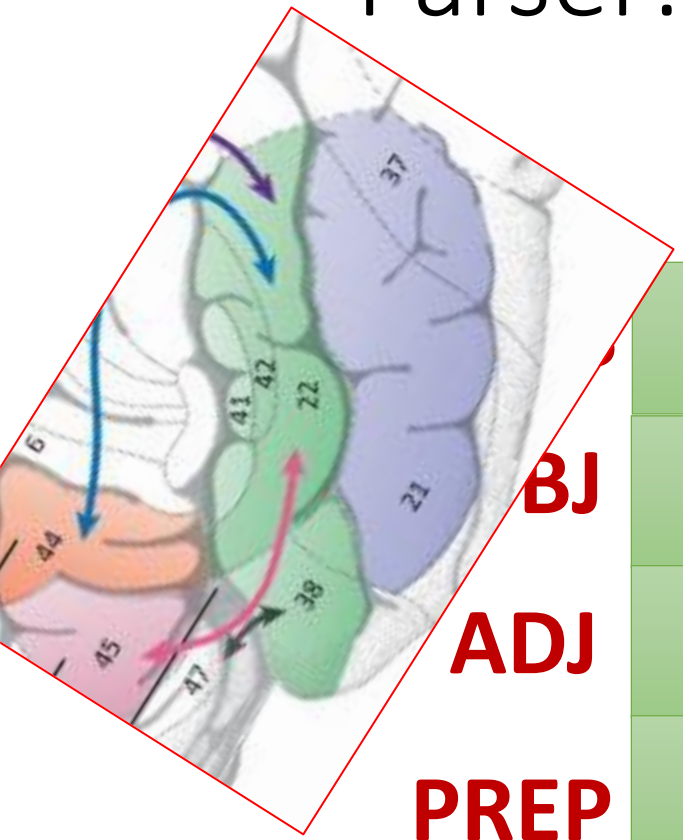

BJ

ADJ

PREP

OBJ

…

fibers

LEX

chased

dog

cat

a

A Parser written in the Assembly Calculus

- Architecture: Brain Areas and Fibers
- Can be inhibited/disinhibited individually
- LEX area contains an assembly for each word
- Each word assembly has an **action set** (inhibit/disinhibit commands)
- Encoding its part of speech, its syntactic role
- When word assembly fires, its action set is ***executed***
- The sum total of word action sets: **≈ *the grammar***

The Parser

- Written in an extension of the Assembly Calculus
- Added: **assemblies with actions**, **project***
- That is, ***the parser is implemented exclusively through the spikes of stylized spiking neurons***
- After processing a sentence, it leaves behind a ***retrievable*** trail of side effects: ***”the parse”***
- ≈ a dependency tree
- ***Prototype is running, paper submitted***
- Joint with Mike Collins and Dan Mitropolsky

The Parser (cont.)

- It parses simple sentences such as:

“the young couple in the next house saw the old little white car of the main suspect quite clearly”

- Speed: about ***20-25 spikes*** (= 0.3-0.5 sec) ***per word***
- More features can be added easily
- With much more work***: recursion and sentence embedding, polysemy, backtracking, …***
- Also running: a simple ***Russian*** parser

is this the neural basis of language?

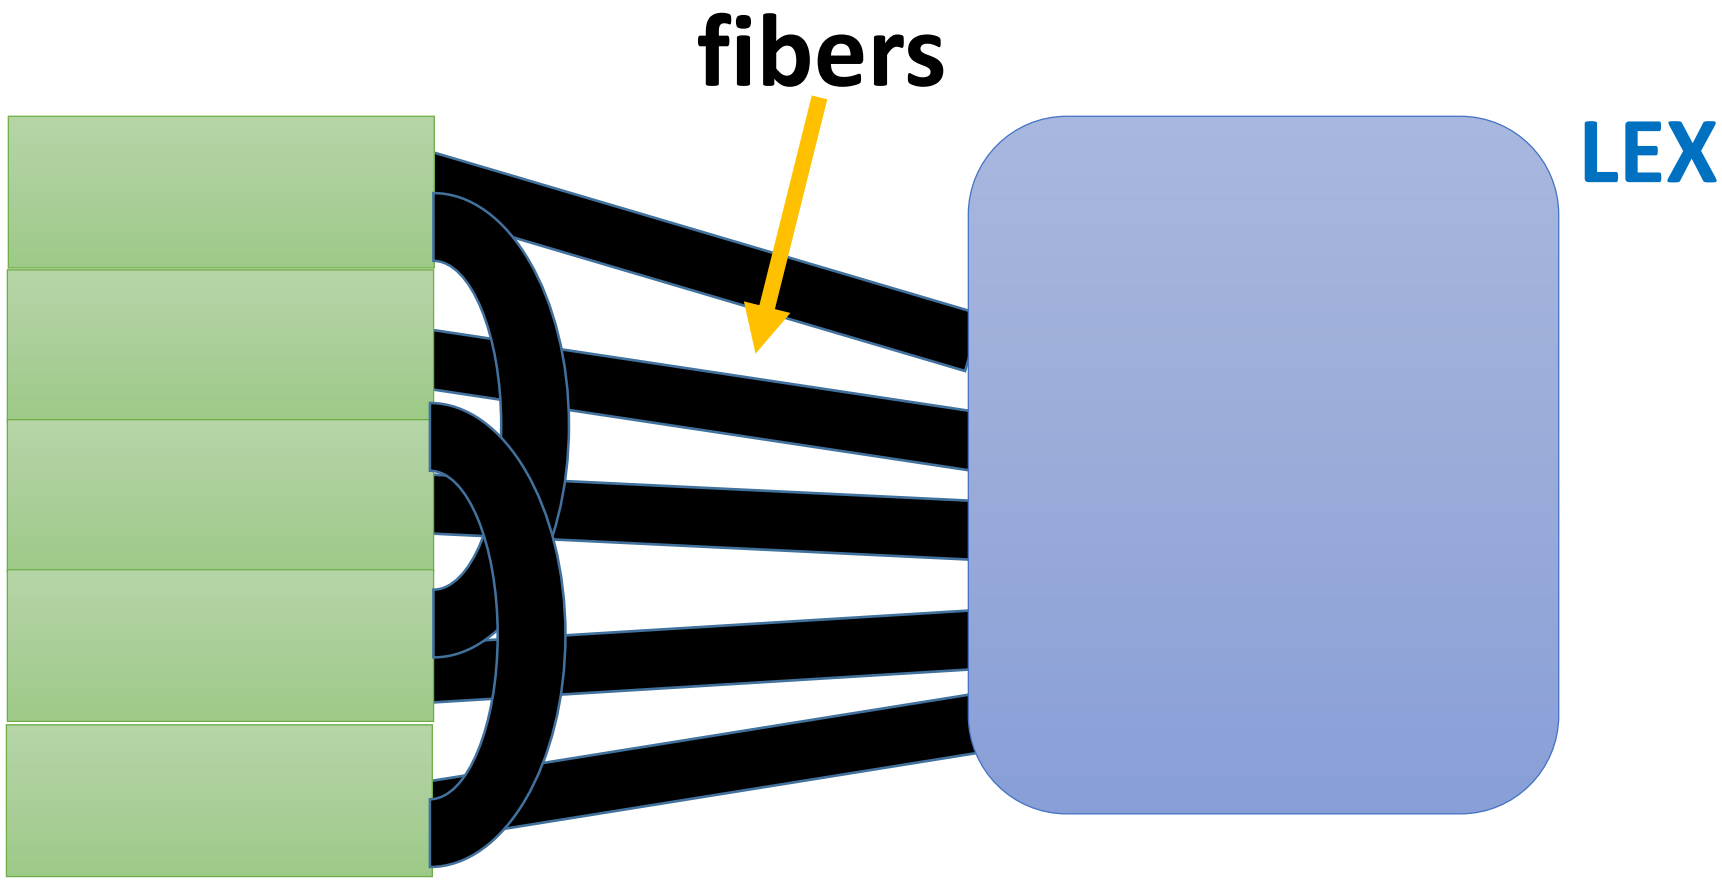

Soooooooo…

- The study of the Brain is fascinating and bottomless
- Assemblies and their operations may be one productive path to computation in the brain – to bridging the gap
- Are assemblies the seat of Axel's "logic"?
- How do assemblies learn and predict?
- How can one test/verify/falsify the Assembly Hypothesis through experiments?
- The parser experiment: engaging with deep questions

my collaborators!

Santosh
Vempala
GaTech

Wolfgang
Maass
TU Graz

Mike Collins
Columbia & Google

Dan
Mitropolsky
Columbia

Thank You!