


Plenary Speaker: Christos H. Papadimitriou

- ▶ He has been awarded the Knuth Prize, IEEE's John von Neumann Medal, the EATCS Award, the IEEE Computer Society Charles Babbage Award, and the Gödel Prize.
- ▶ He is a fellow of the ACM and the National Academy of Engineering, and a member of the National Academy of Sciences and the American Academy of Arts and Sciences.
- ▶ He has written five textbooks and many articles on algorithms and complexity, and their applications to optimization, databases, control, AI, robotics, economics and game theory, the Internet, evolution, and the brain.
- ▶ He co-founded my field, algorithmic game theory, introducing worst-case equilibria and *price of anarchy*.
- ▶ He settled the complexity of Euclidian traveling salesman and Nash equilibrium.
- ▶ He holds a PhD from Princeton (1976), and eight honorary doctorates.
- ▶ He has also written three novels: Turing, Logicomix, and his latest, Independence.

To: Fred Roberts
From: Tom Leighton and Christos Papadimitriou (chair)
Date: Thu, 6 Mar 1997
Subject: DIMACS Research Evaluation

Our two-member research evaluation team was given seven specific questions (below) relating to the strategic and tactical research goals of DIMACS, and





**A Calculus for
Brain Computation**

Christos H. Papadimitriou
Columbia University

Computer Science 1936 – 1995: the Computer



Computer Science 1936 – 1995: the Computer



compilers

operating systems

graphics

databases

intelligent systems


algorithms

chips

networks

P vs NP

Computer Science 1995 – : the Internet



Computer Science 1995 – : the universe



economics
game theory

behavior

dynamics

statistical physics

quantum
physics

networks

Computer Science 1995 – : the universe



Computer Science 1995 – : the universe



behavior


dynamics

statistical physics

quantum
physics

networks

*Computation
as a lens on the Sciences*



Game Theory and Economics




Game Theory and Economics

Finding a
Nash
equilibrium
is an
intractable
problem!



Evolution 160 years later

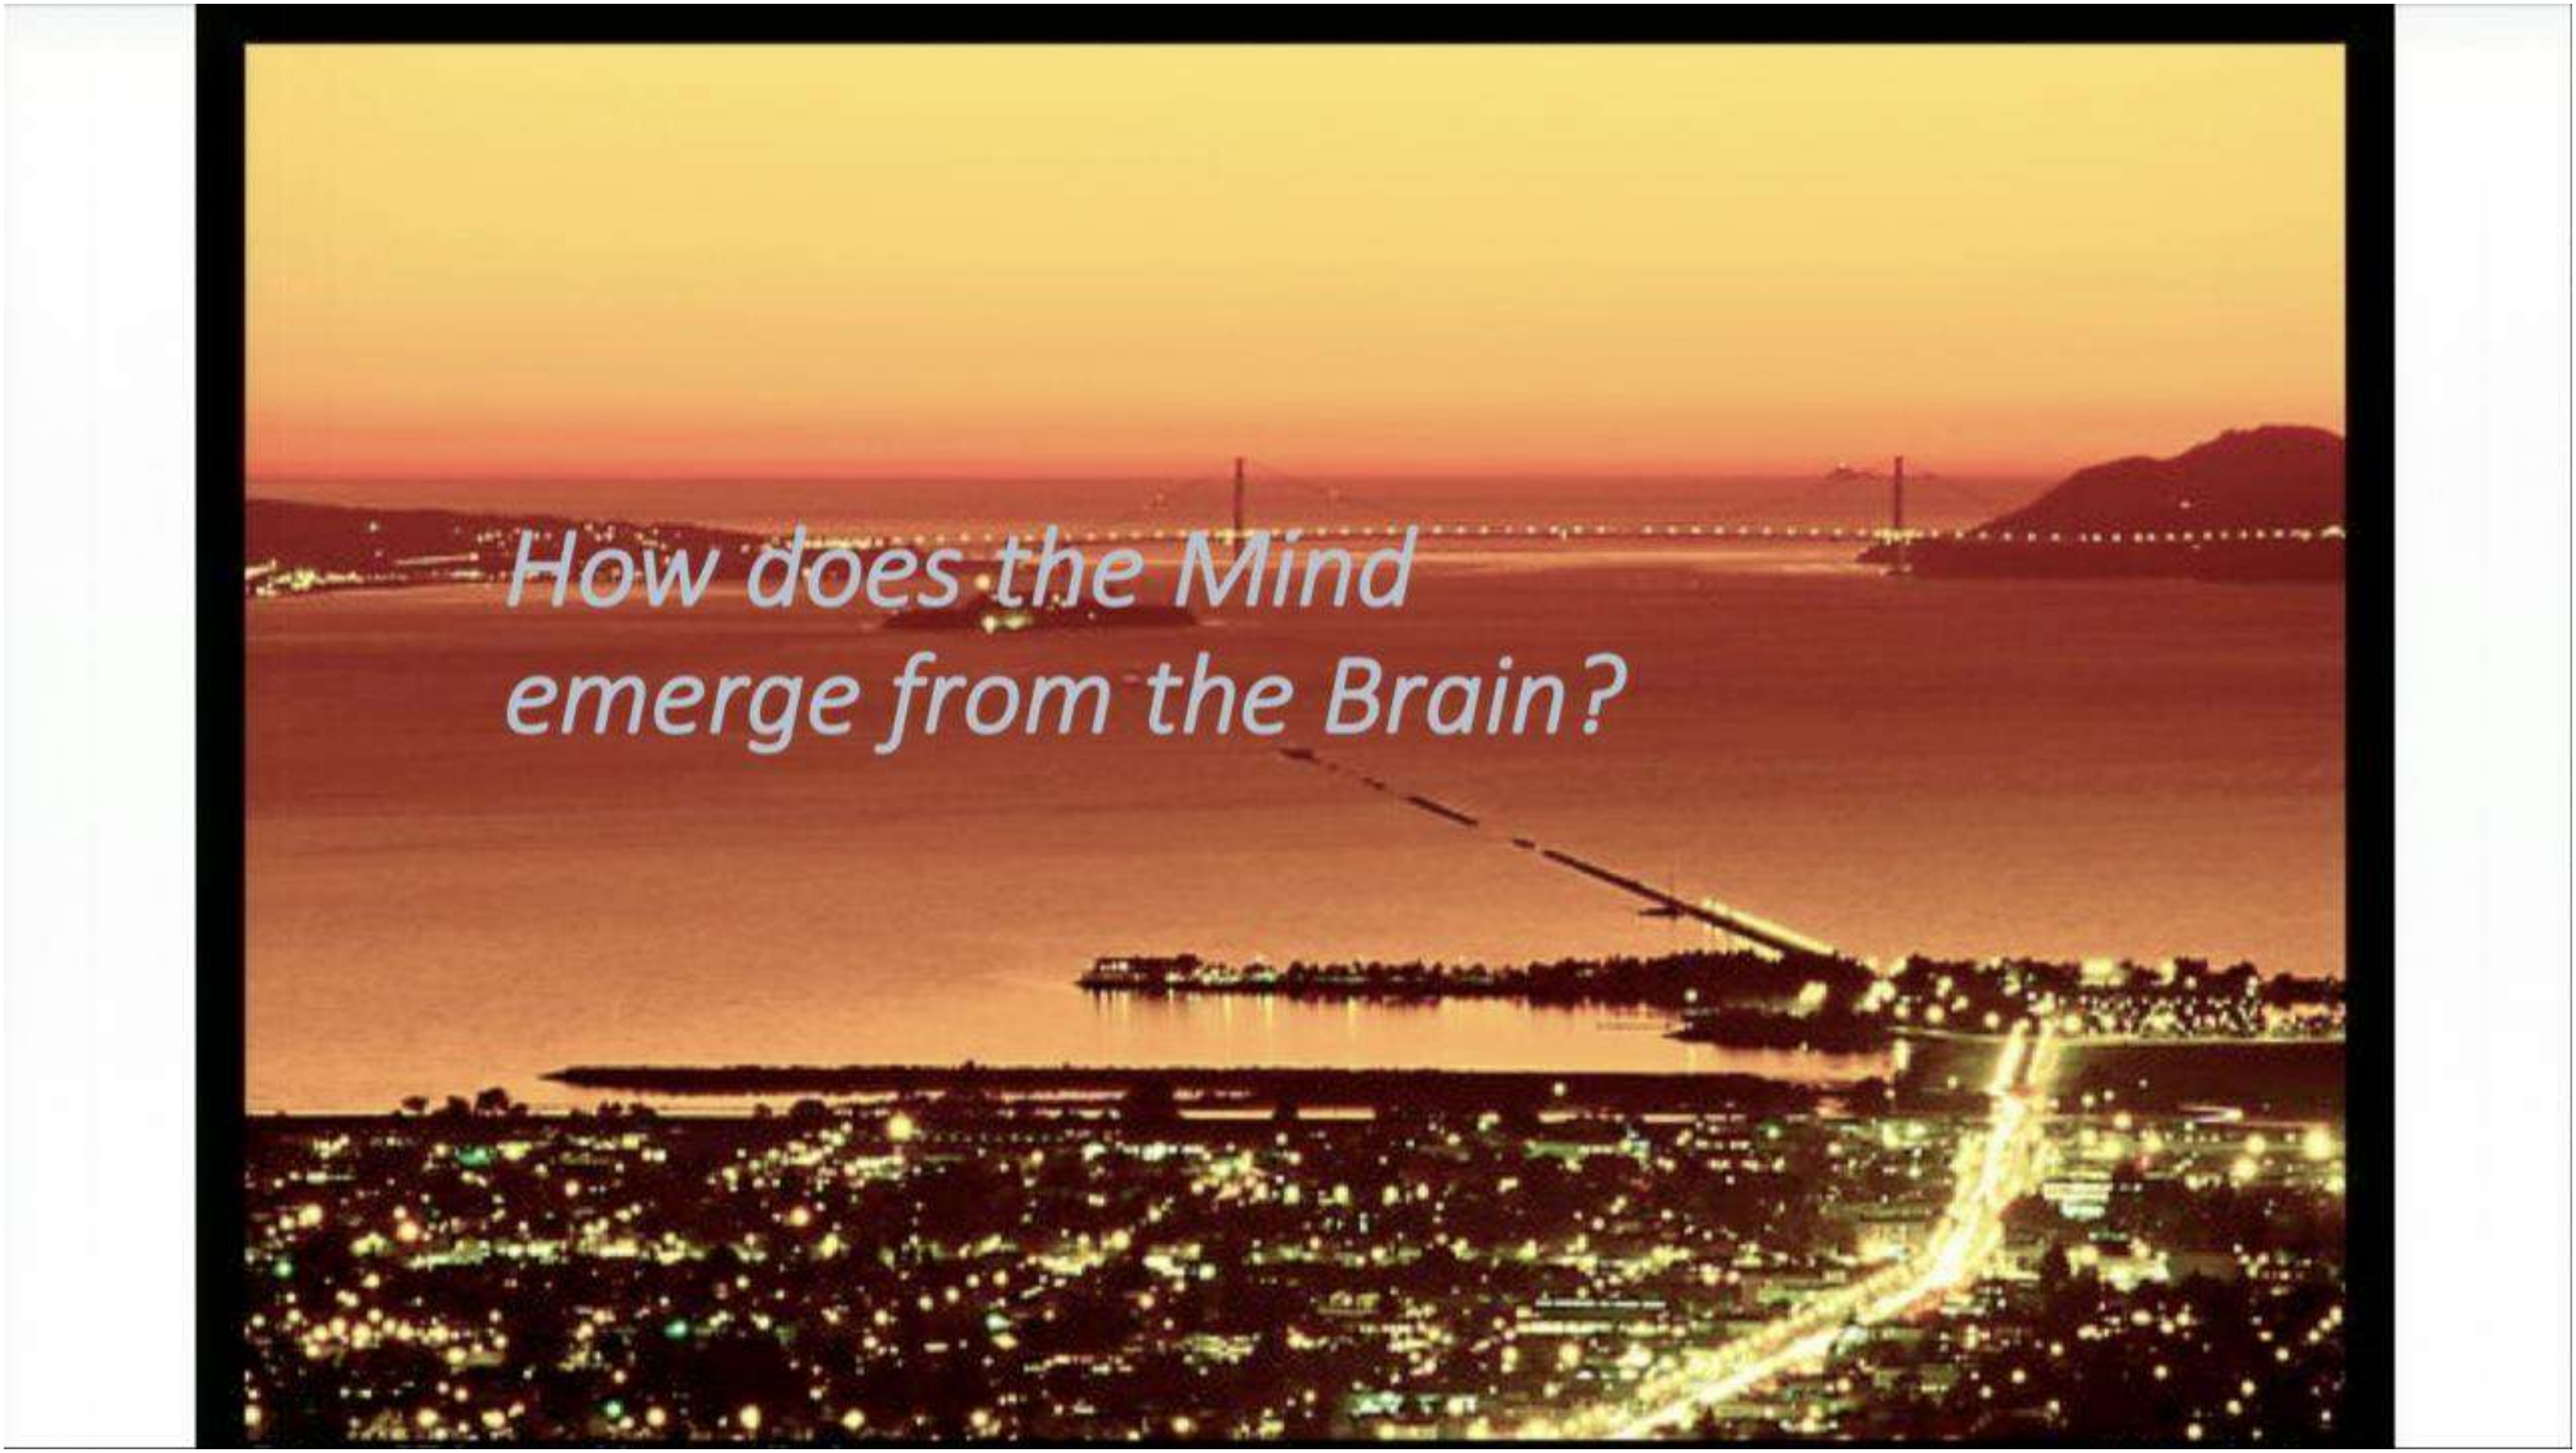


The evolution of a population of genotypes is tantamount to the **genes** playing a **repeated game** through **ADABOOST** (the well known AI algorithm)

Furthermore, every gene in every generation optimizes a **trade-off** between cumulative **fitness** and **allele entropy**

...and next....

*How does the Mind
emerge from the Brain?*

An aerial night photograph of a city, likely Seattle, showing a large body of water, a suspension bridge, and city lights. The text is overlaid on the image.

*How does the Mind
emerge from the Brain?*



PRINCIPLES
OF NEURAL
SCIENCE
Fifth Edition

Eric R. Kandel
James H. Schwartz
Thomas M. Jessell
Steven A. Siegelbaum
A. J. Hudspeth

Photo by
Sando Medical Books

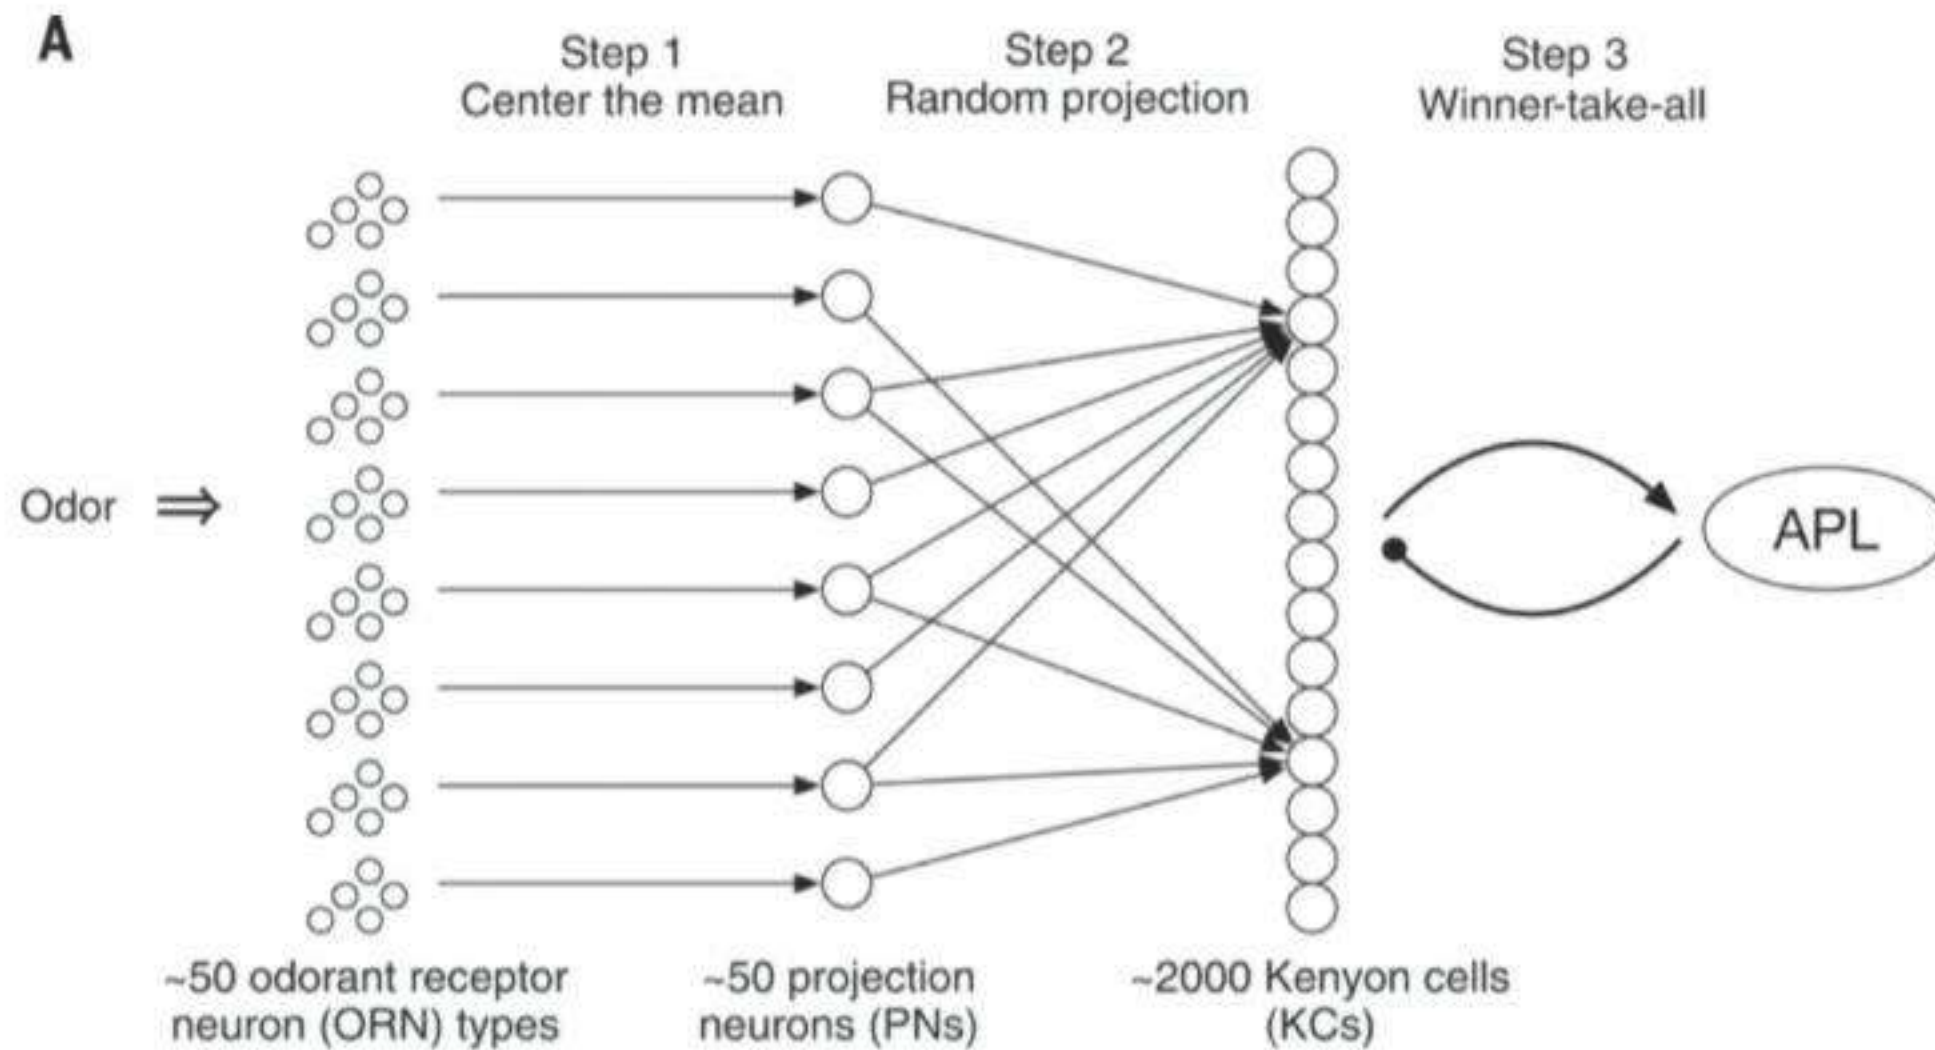


“...we do not have a logic for the transformation of neural activity into thought ... I view discerning [this] logic as the most important future direction of neuroscience...”

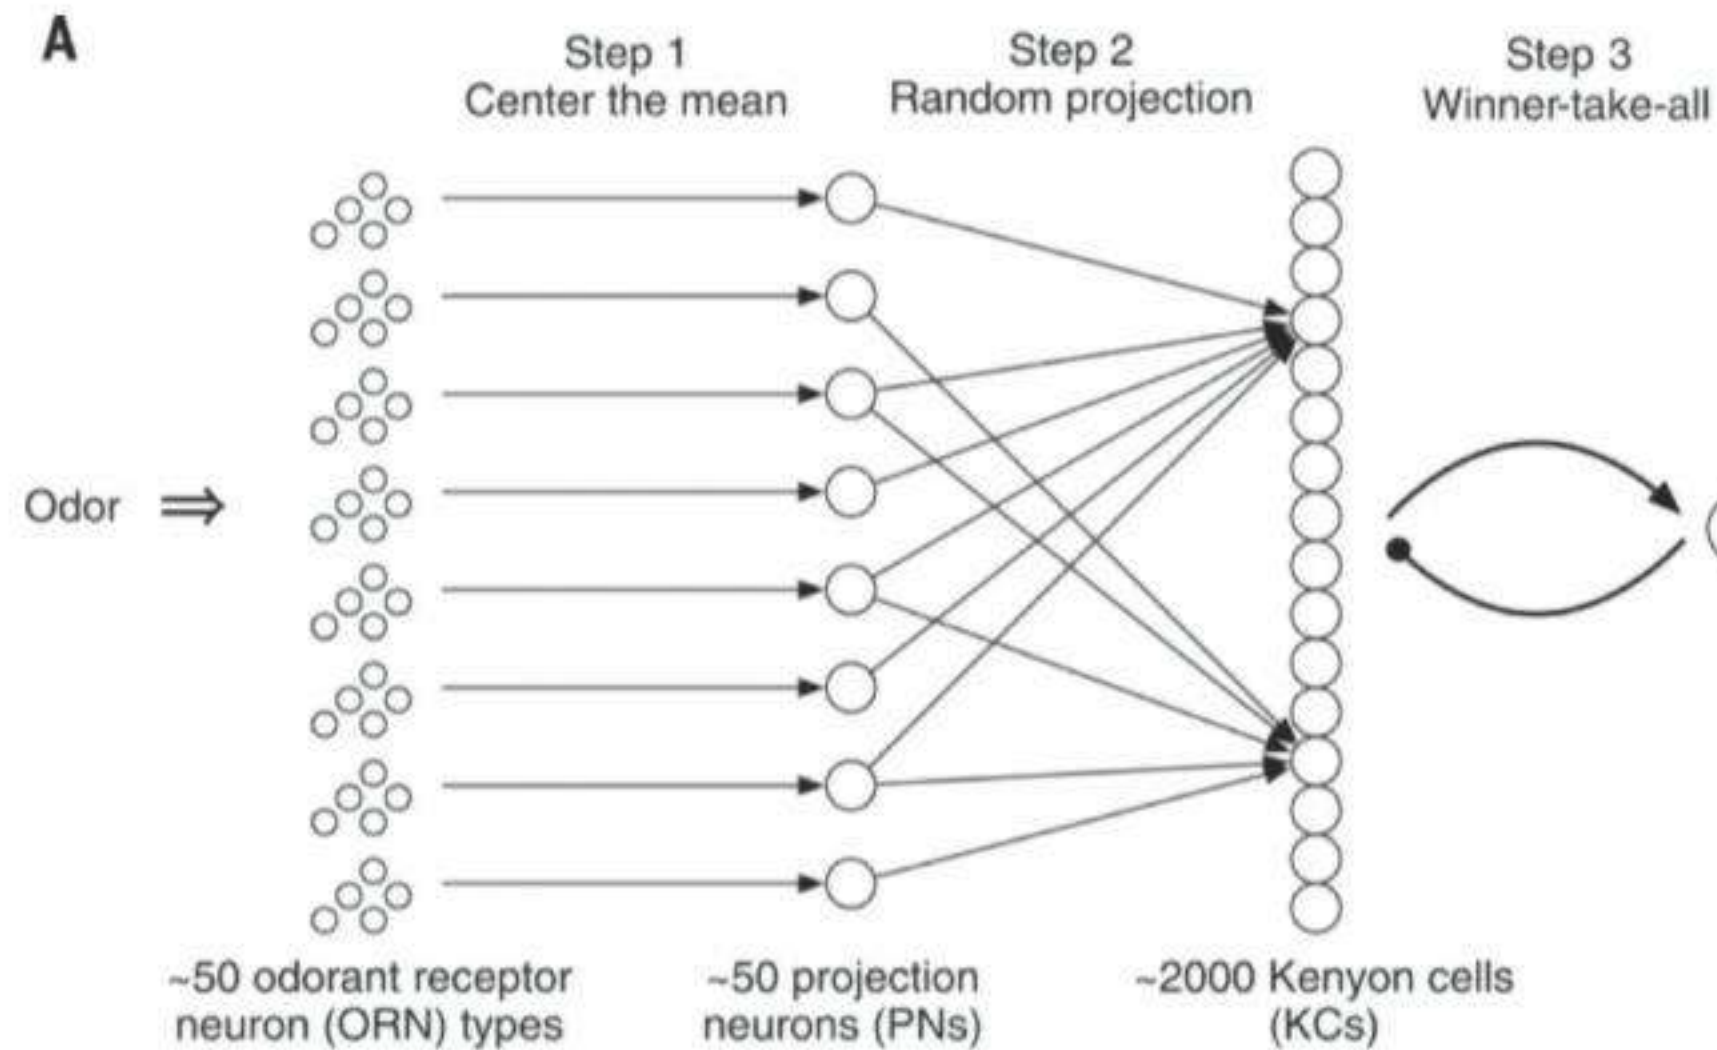
Richard Axel: *Neuron* Sep 2018



How fruit flies remember smells

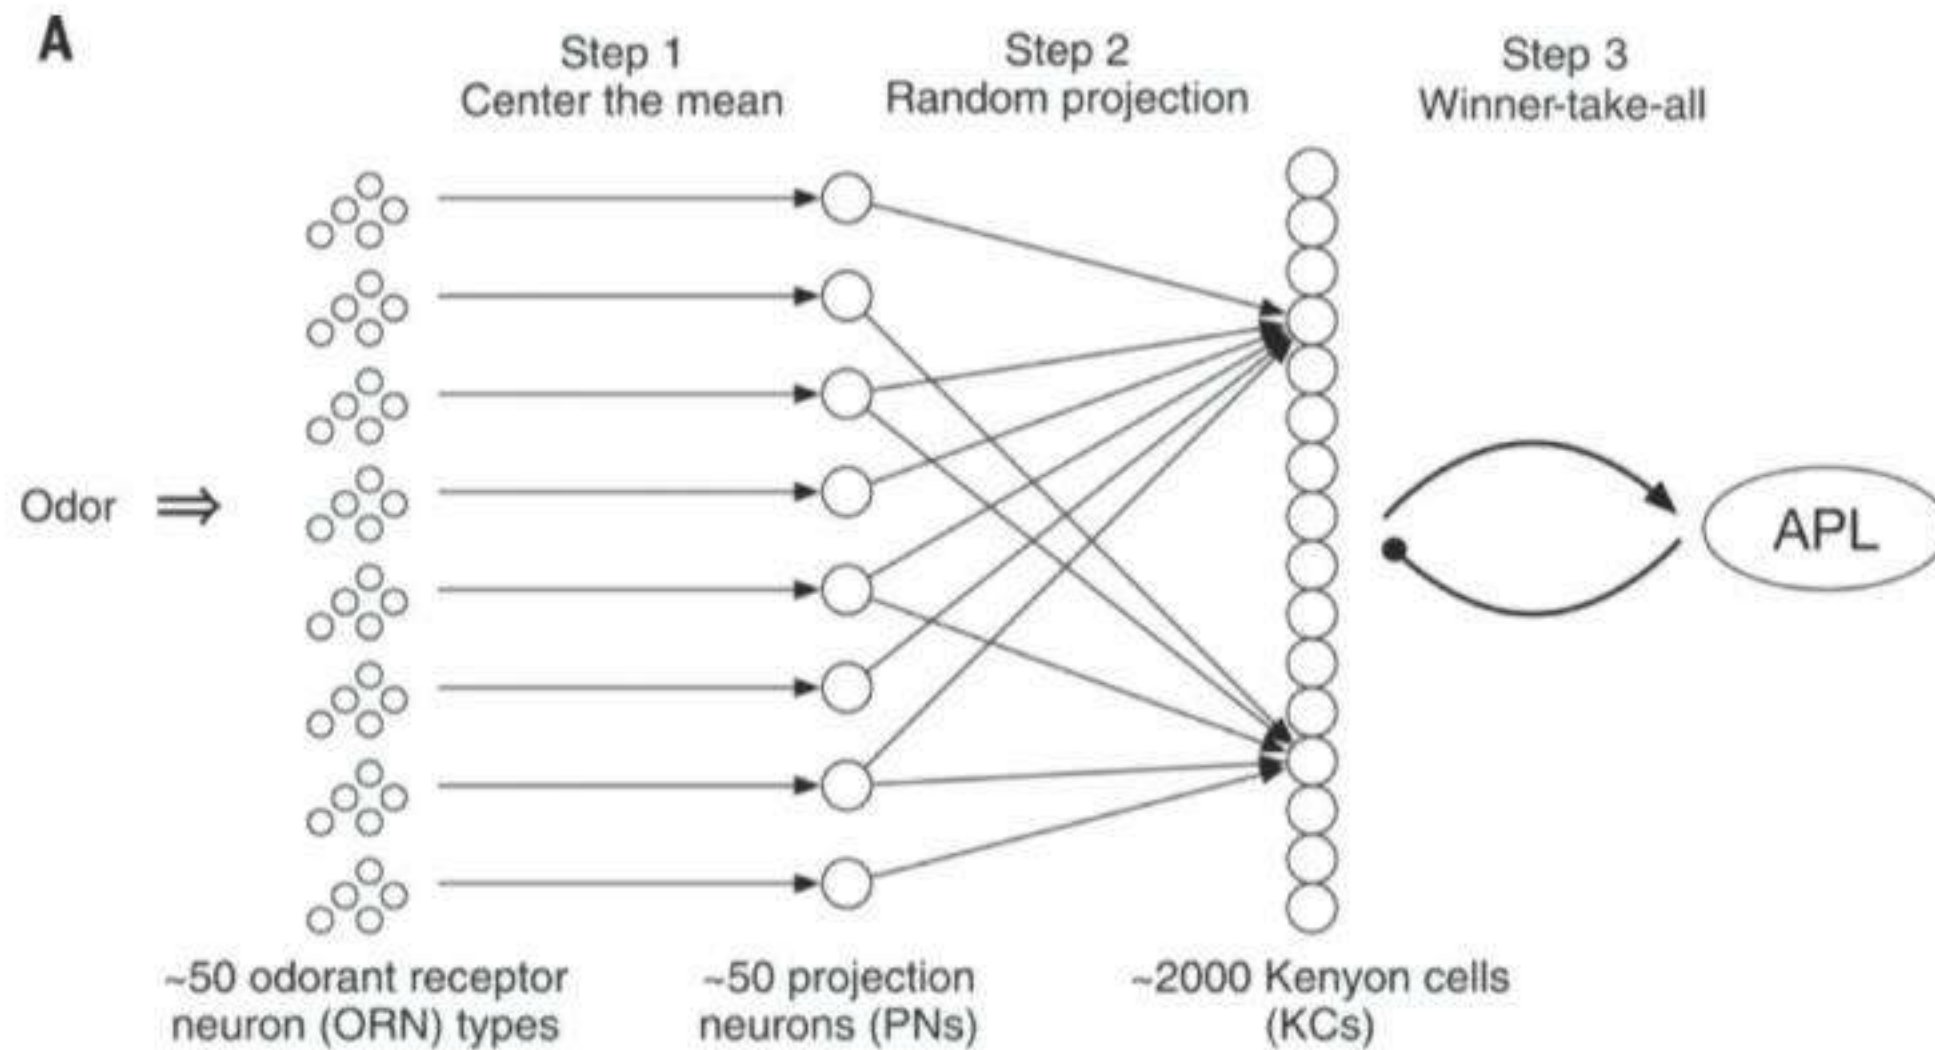


How fruit flies remember smells

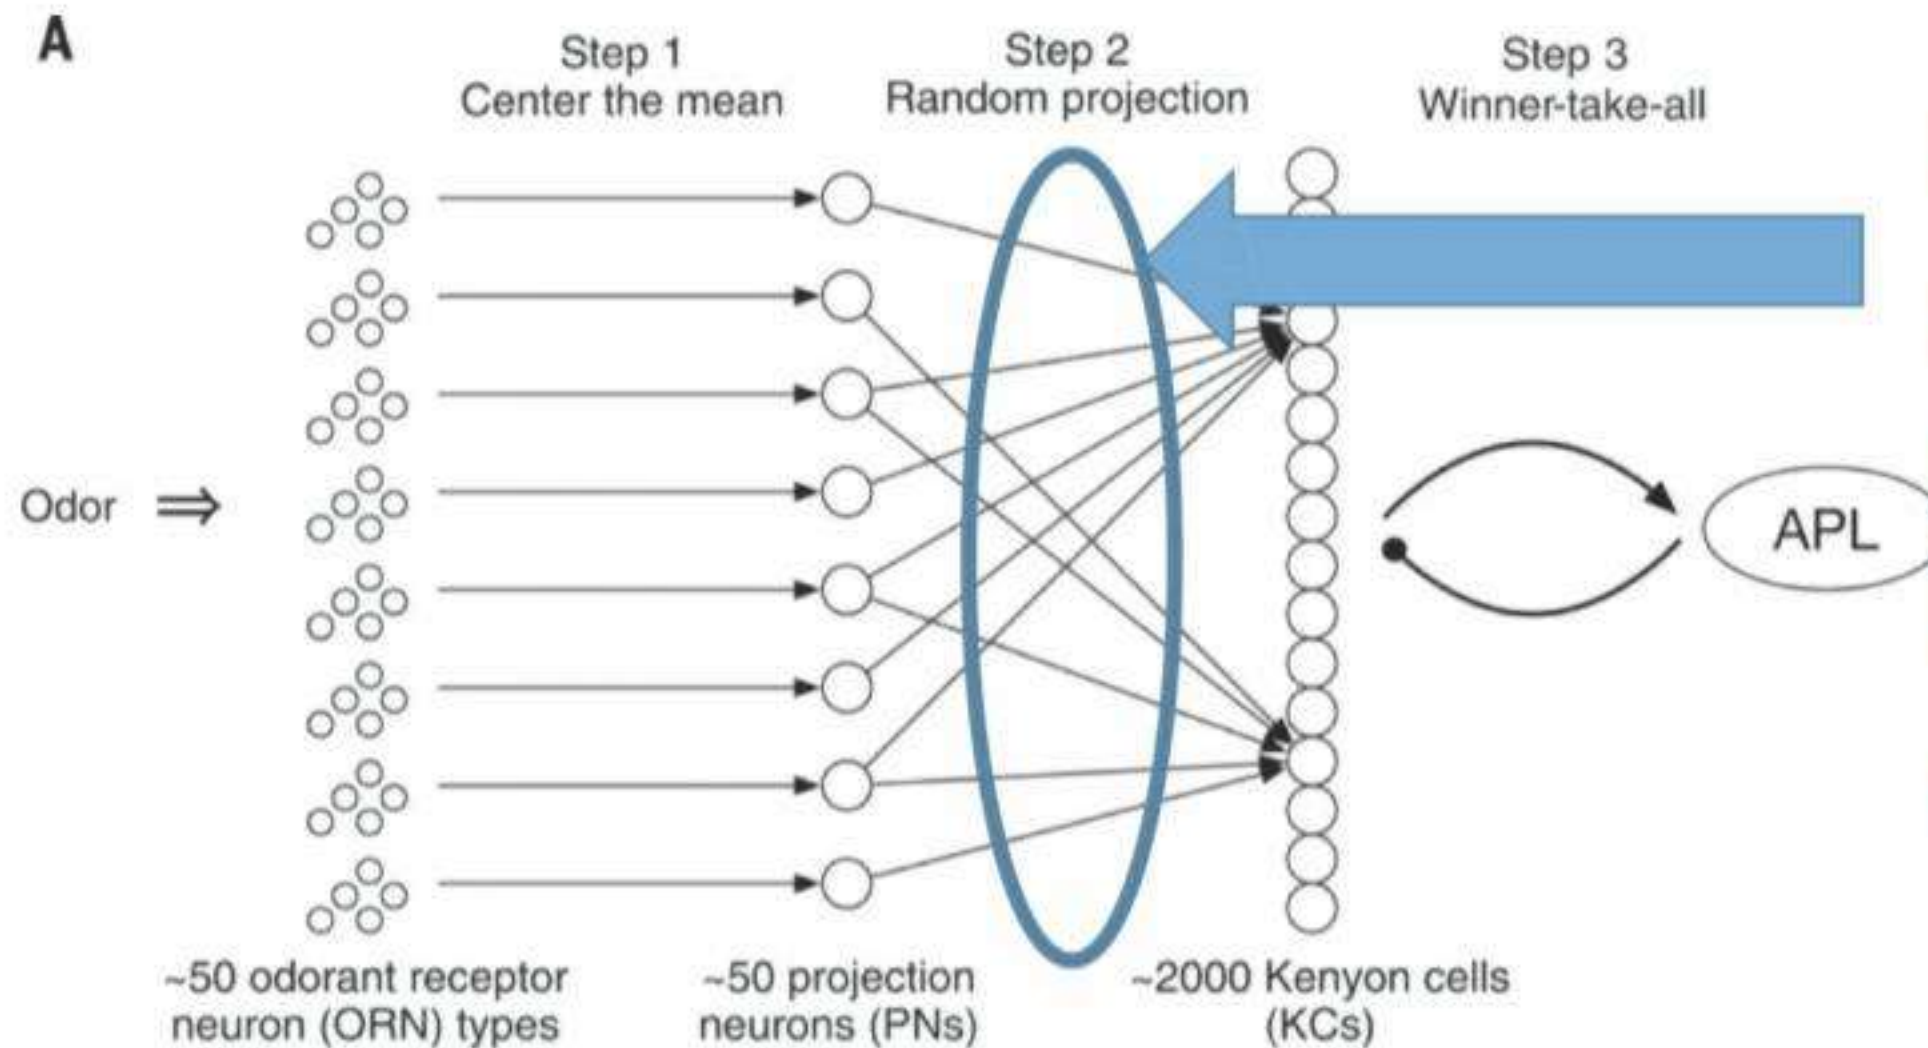


**random
projection
followed by
cap (RP&C):
100 winners
(out of 2000)
take all**

How fruit flies remember smells




How fruit flies remember smells




Q: but wait a minute! Is this a random bipartite graph?

A: *Random convergence of olfactory input in the Drosophila mushroom body* by S. Caron, V. Ruta, L. Abbott, R. Axel, 2013


Bottom line:
looks like a
random
bipartite graph,
except that the
degree
distribution
of the LHS is
not uniform



RP&C




RP&C




RP&C preserves similarity!

- Random $n \times n$ bipartite graph




RP&C preserves similarity!

- Random $n \times n$ bipartite graph
- A set **A** of k out of n nodes of the LHS fire



RP&C preserves similarity!

- Random $n \times n$ bipartite graph
- A set **A** of k out of n nodes of the LHS fire
- A new set **cap(A)** is formed by RP&C
- Repeat now for **B, cap(B)**
- If **A** and **B** overlap in αk nodes, what is the overlap of **cap(A)** and **cap(B)**?




The underlying mathematical reason:

Theorem [P., Vempala, 2018]:

The intersection of **cap(A)** and **cap(B)** will be, with high probability, at least

$$\frac{\left(\frac{k}{n}\right)^{\frac{1-\alpha}{1+\alpha}}}{\left(\ln n/k\right)^{\frac{\alpha}{1+\alpha}}}$$

The
underlying
mathematical
reason:
compare
with
simulations




So much for the fruit fly...




- Q: Does something homologous happen in mammals?

Yes!



K. Franks, M. Russo, S. Sosulki, A. Mulligan, S. Siegelbaum, R. Axel
“Recurrent Circuitry Dynamically Shapes the Activation of
Piriform Cortex,” Neuron October 2011

Yes!



K. Franks, M. Russo, S. Sosulki, A. Mulligan, S. Siegelbaum, R. Axel
“Recurrent Circuitry Dynamically Shapes the Activation of
Piriform Cortex,” Neuron October 2011

From the *Discussion* section of Franks *et al.*

An odorant may [cause] a small subset of [PC] neurons [to fire].


Inhibition triggered by this activity will prevent further firing

This small fraction of ... cells would then generate sufficient recurrent excitation to recruit a larger population of neurons.

In the extreme, some cells could receive enough recurrent input to fire ... without receiving [initial] input...


In pictures...

set of spiking
neurons




In pictures...

set of spiking
neurons



RP&C




In pictures...

set of spiking
neurons



RP&C




third RC&P,
and so on...


new RC&P

In pictures...

set of spiking
neurons



RP&C



new RC&P

Also, Hebbian
plasticity

third RC&P,
and so on...

Does this process converge?

And does it preserve similarity?

The model

Finite number of brain regions

Each contains n neurons

Inhibition: only k fire


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.)

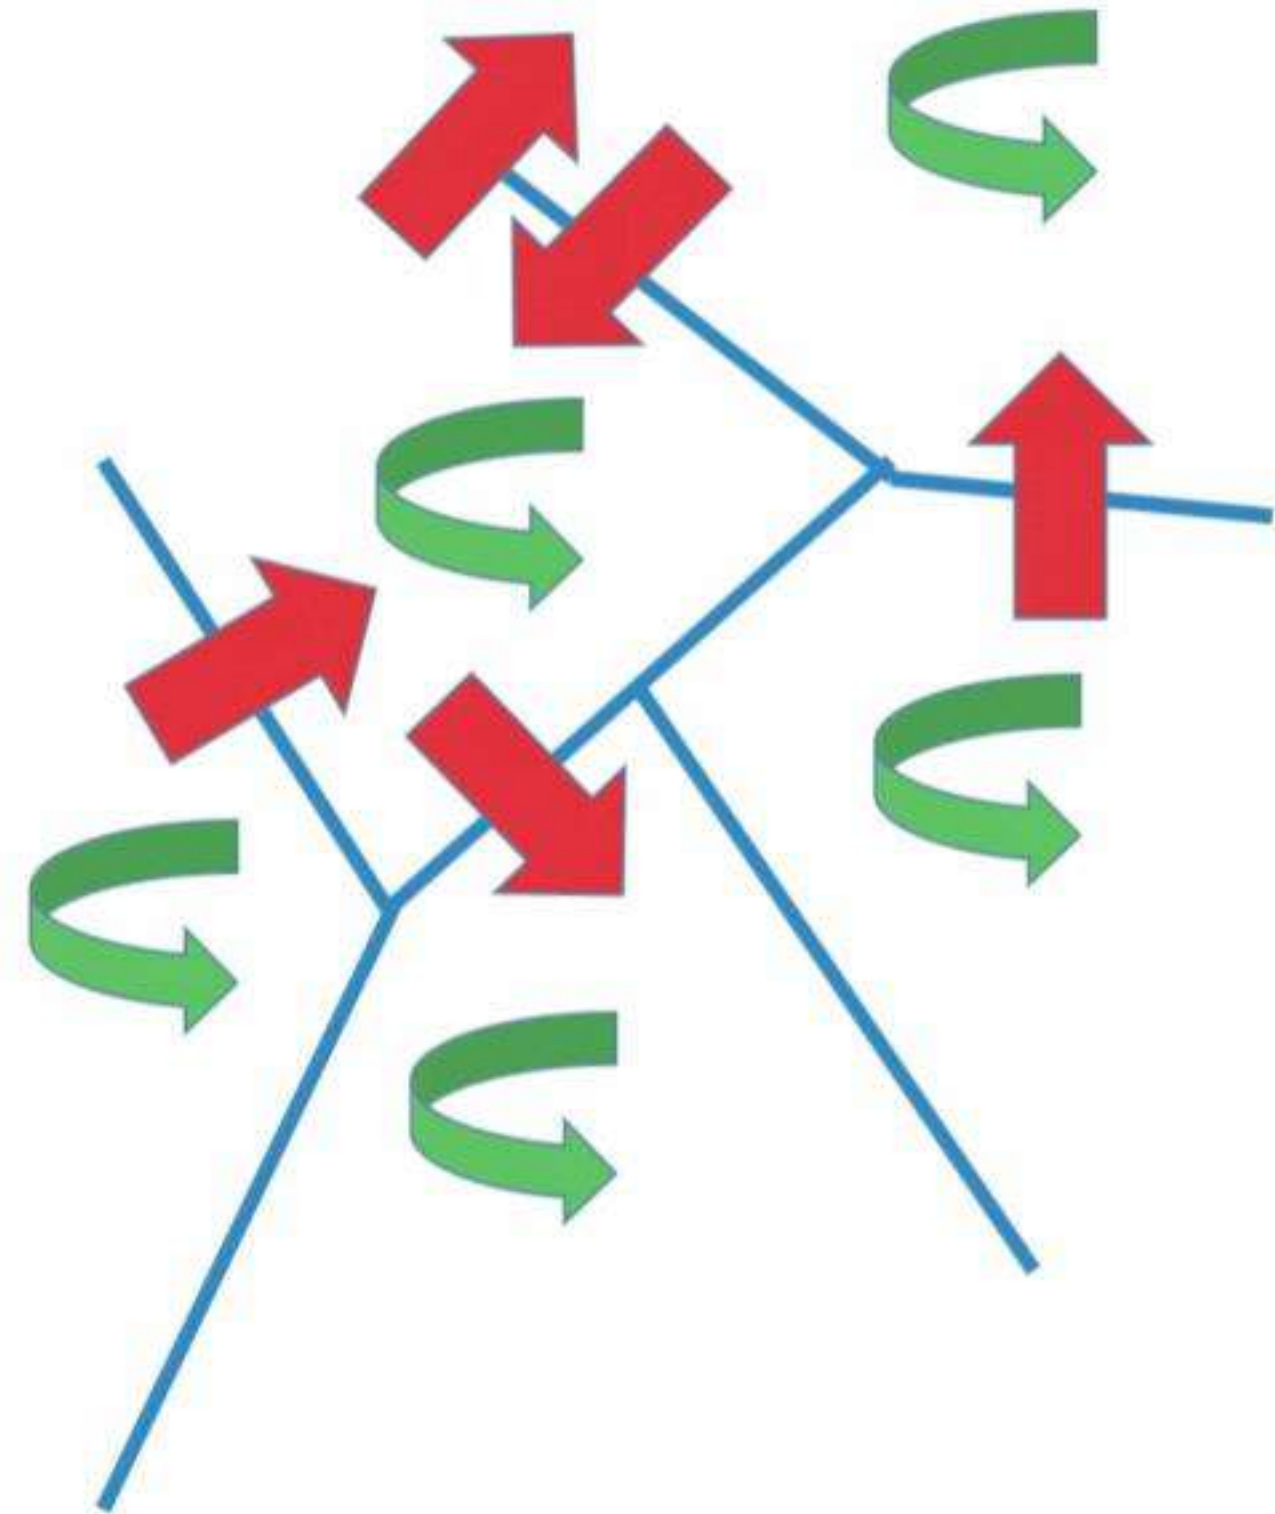
Neurons fire in **discrete steps**

Selected by RP&C

Connections **between** areas can be enabled/disabled

Plasticity: If $i \rightarrow j$, i fires and in the next step j fires, the weight of $i \rightarrow j$ is multiplied by $(1 + \beta)$

Also: homeostasis, forgetting



Main parameters, intended values

- $n \sim 10^7$
- $k \sim 10^3 - 4$
- $p \sim 0.001$
- $\beta \sim 0.10$



Main parameters, intended values

- $n \sim 10^7$
- $k \sim 10^3 - 4$
- $p \sim 0.001$
- $\beta \sim 0.10$

Main ideas: randomness, selection, plasticity

Theorem (P., Vempala 2016-18): The process converges exponentially fast, with high probability, and the *total number of cells involved* is **at most**:

- If $\beta \geq \beta^*$: $k + o(k)$
- If $0 < \beta < \beta^*$: $k \cdot \exp(0.17 \cdot \ln(n/k) / \beta)$
- $\beta^* = (\sqrt{2} - 1) / (1 + \sqrt{pk / \ln n})$



The result of such projection: an *Assembly*

- Set of $\approx k$ neurons in a brain area whose firing (in a pattern) is tantamount to our thinking of a particular memory, concept, name, word, episode, etc.
- [Hebb 1949, Harris 2003, 2005; Buzsaki 2008, 2010, Yuste et al. 2017]
- Simulations of a far more biologically accurate model [Pokorny et al 2018, under submission]



cell



cells (or concept cells)

The Big Picture

Computation in the brain: What is the right level?

The Big Picture

Computation in the brain: What is the right level?

- Molecules?

The Big Picture

Computation in the brain: What is the right level?

- Spiking neurons and synapses?
- Molecules?

The Big Picture

Computation in the brain: What is the right level?

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

The Big Picture

Computation in the brain: What is the right level?

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

The Big Picture

Computation in the brain: What is the right level?

- Whole brain?



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

Recall...

“...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...”

Neuron, Sep 2018



The Big Picture

Computation in the brain: What is the right level?

- Whole brain?



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

Recall...

“...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...”

Neuron, Sep 2018



The assembly hypothesis

The assembly hypothesis

- There is an **intermediate level** of brain computation

The assembly hypothesis

- There is an **intermediate level** of brain computation
- Implicated in carrying out **higher cognitive functions** such as reasoning, planning, language, story-telling, math, music, ...

The assembly hypothesis

- There is an **intermediate level** of brain computation
- Implicated in carrying out **higher cognitive functions** such as reasoning, planning, language, story-telling, math, music, ...
- Assemblies are its basic **representation** – its main “data structure”

The assembly hypothesis

- There is an **intermediate level** of brain computation
- Implicated in carrying out **higher cognitive functions** such as reasoning, planning, language, story-telling, math, music, ...
- Assemblies are its basic **representation** – its main “data structure”
- **What are its fundamental operations?**

The assembly hypothesis

- There is an **intermediate level** of brain computation
- Implicated in carrying out **higher cognitive functions** such as reasoning, planning, language, story-telling, math, music, ...
- Assemblies are its basic **representation** – its main “data structure”
- **What are its fundamental operations?**
- NB: an operation must be **useful** and **plausible**

Useful and plausible...

- **Useful**: must help explain experiments
- **Plausible**: can be “compiled down” to neurons and synapses


The assembly hypothesis: operations

- **Project**(x, A, y)
- $A = \text{area}(y)$, $x = \text{parent}(y)$
- (Plus, this is how an assembly is first created)


The assembly hypothesis: operations

- **Project**(x, A, y)
- $A = \text{area}(y)$, $x = \text{parent}(y)$
- (Plus, this is how an assembly is first created)
- Q: *Other operations?*
- A: Two assemblies may be **associated** by sharing cells
- Association encodes “affinity”, similarity, co-occurrence...


Association of two assemblies



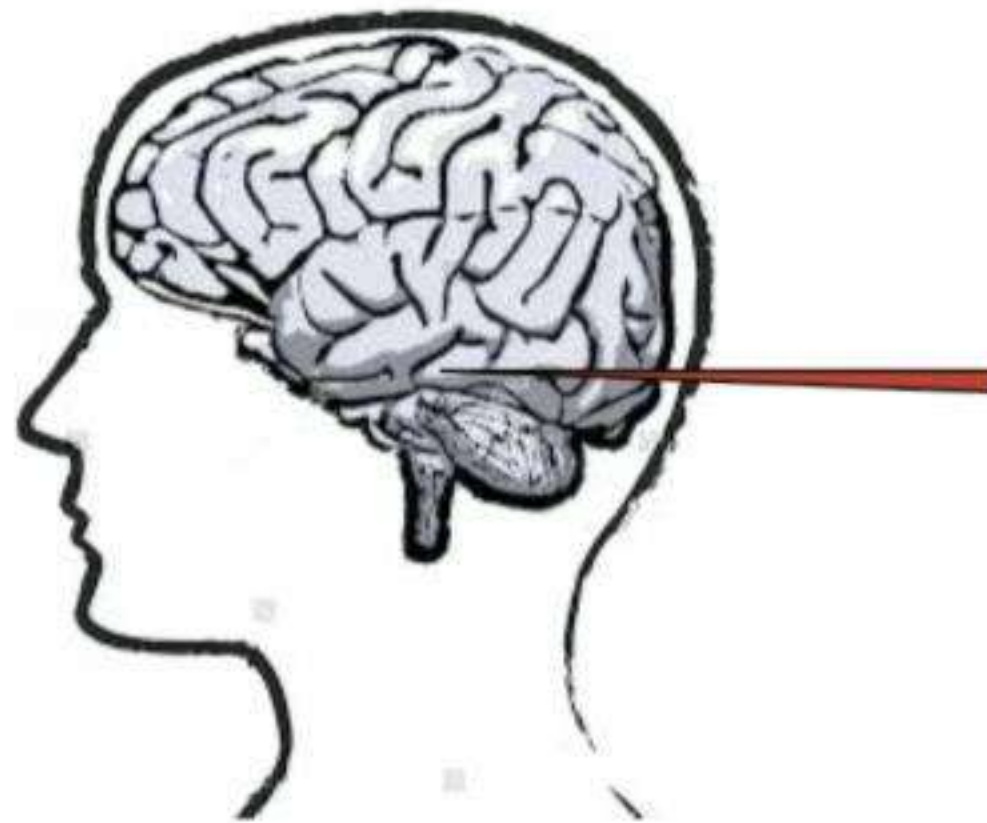
Association of two assemblies






Association of two assemblies





The [Ison et al. 2016] experiment




















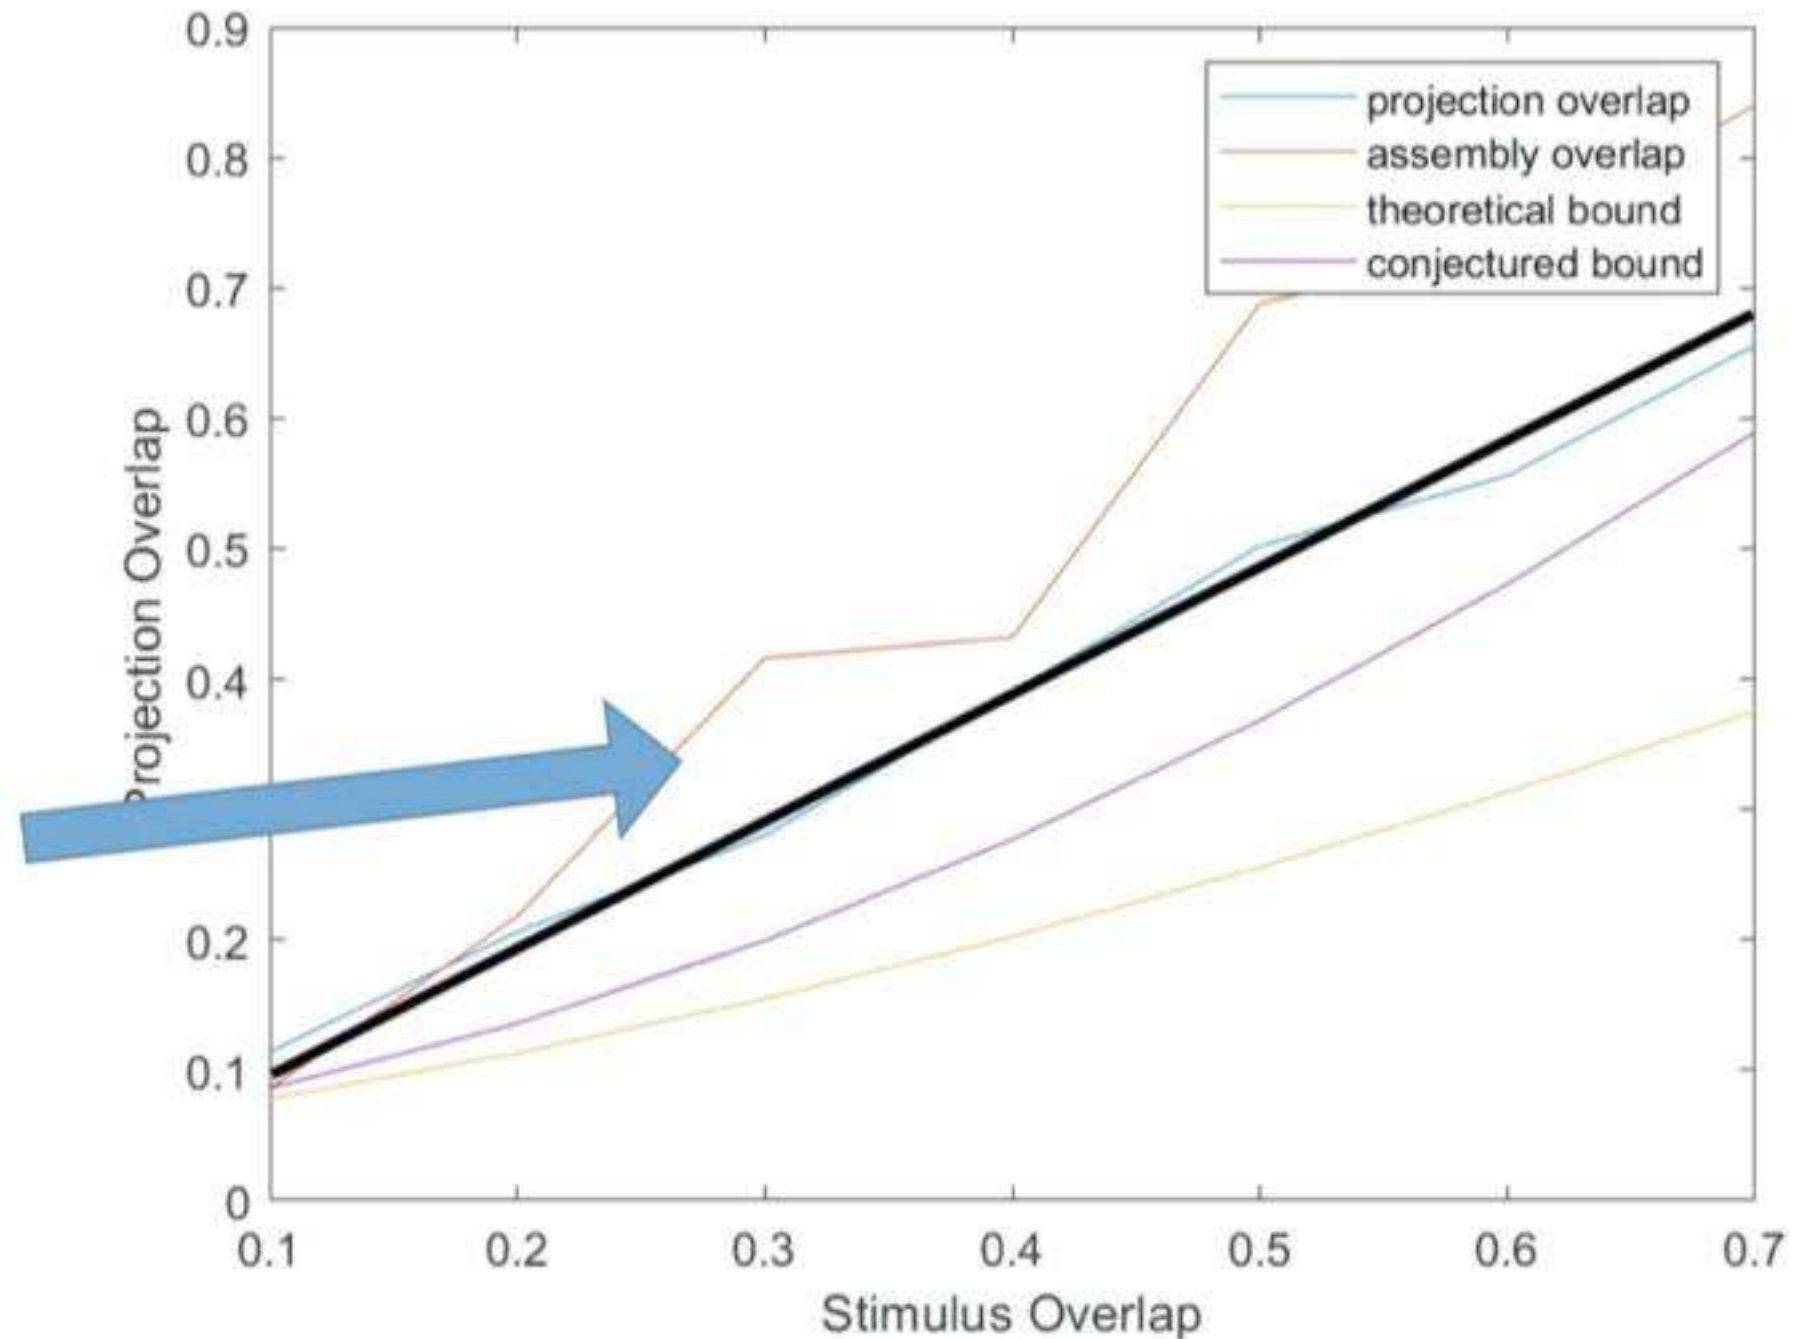




Q: is association preserved under projection?

Recall the fly,
and similarity
preservation:


association of assemblies
seems to be *very
well preserved* under
projection




Other operations?

- **merge**(x, y, A, z)
- (assemblies x, y, from different areas, project to create **one** assembly in area A, call it z)
- Creates hierarchies
- Valuable for implementing **language**

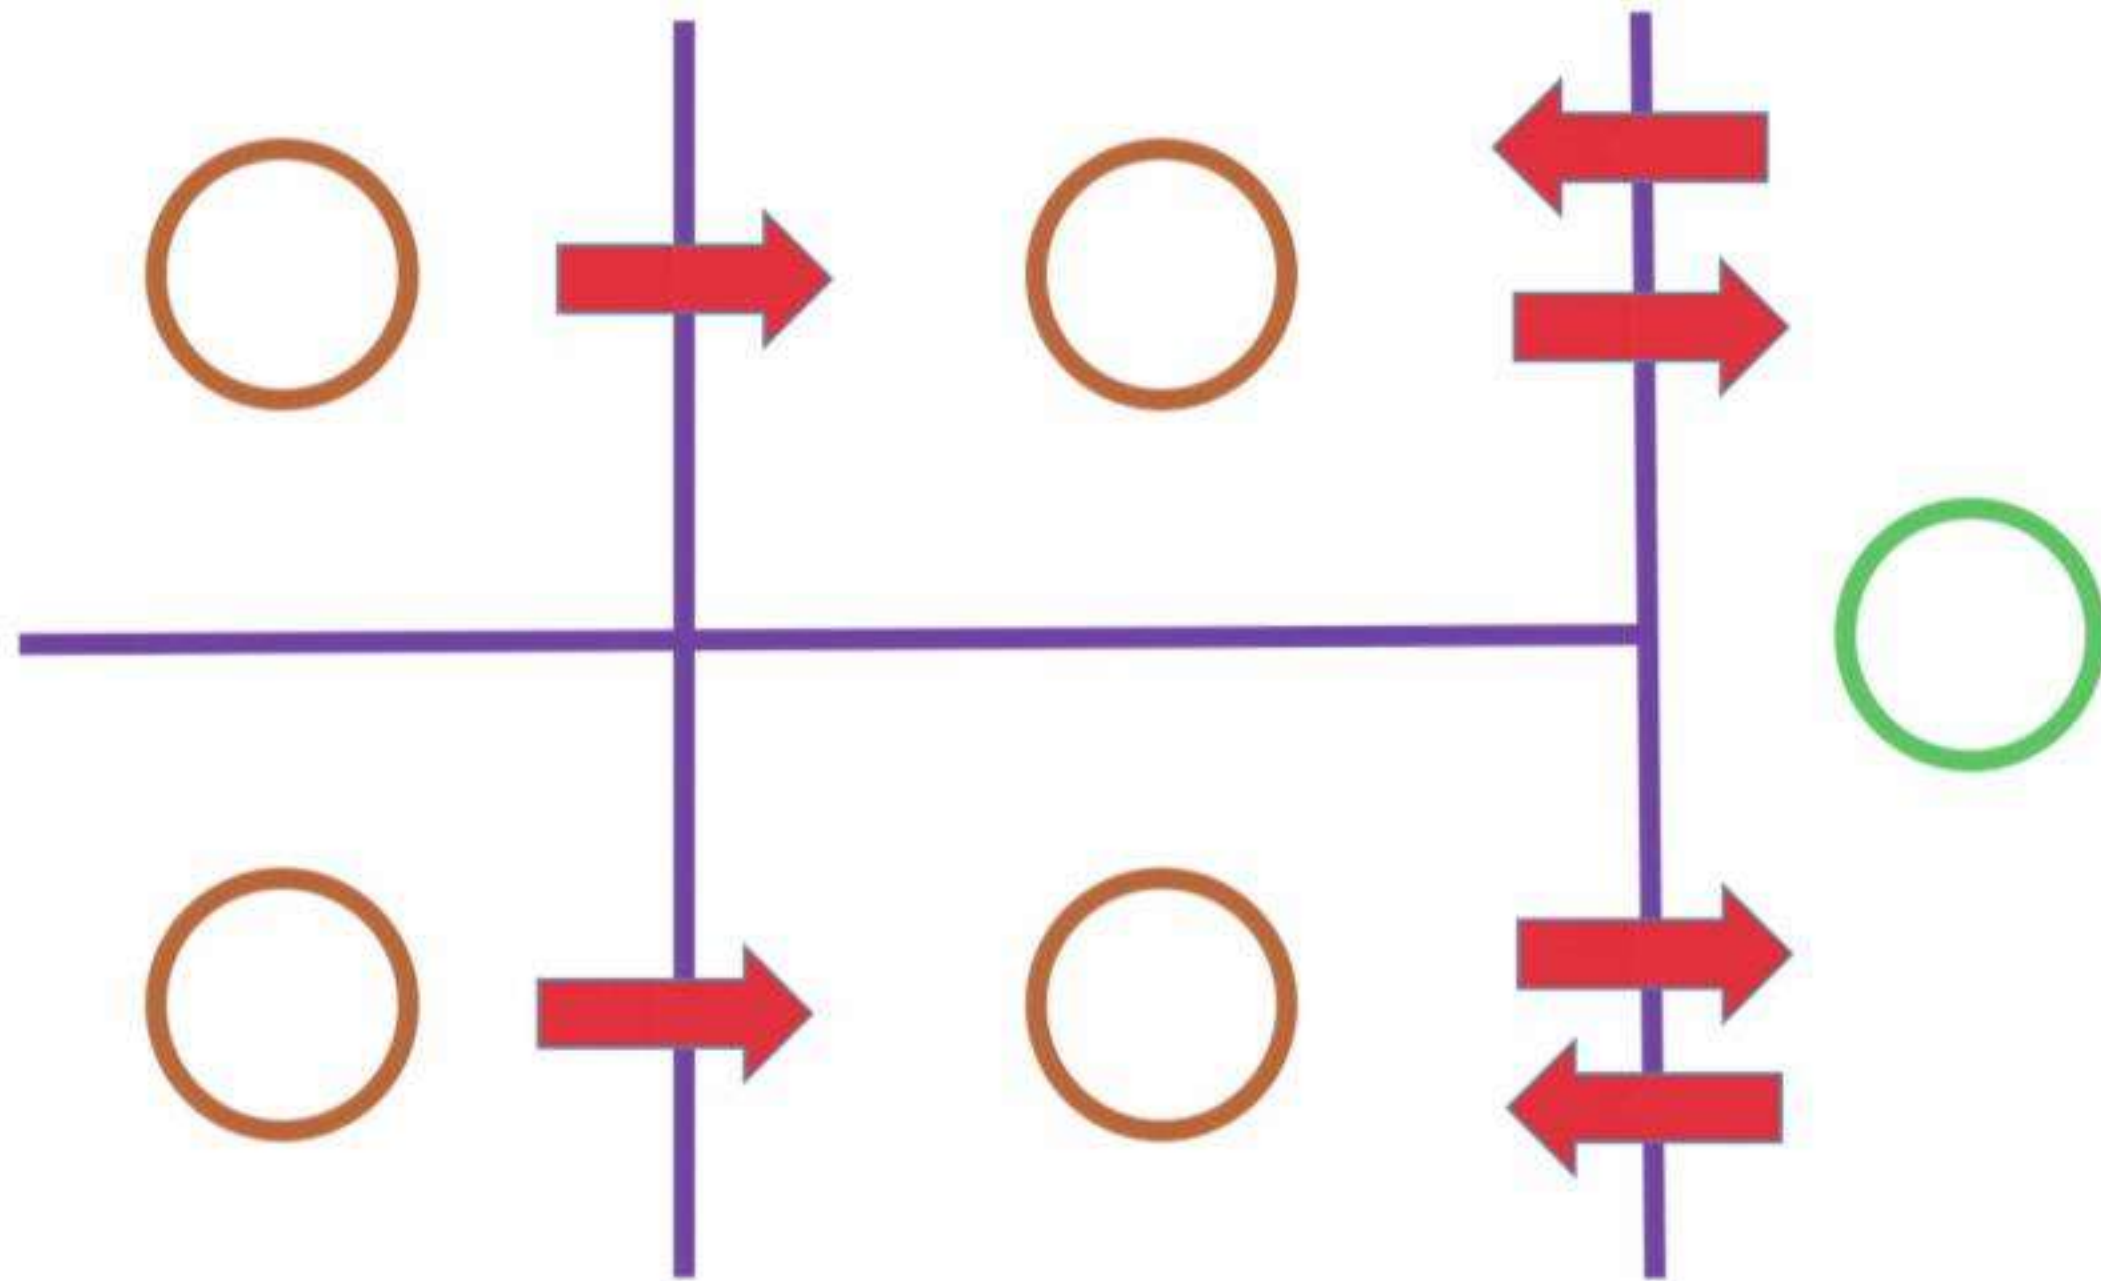
Merge: It's complicated...




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



Merge: It's complicated...



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



Assembly Operations recap

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

Assembly Operations recap

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

Q: *How powerful is this system?*

Assembly Operations recap

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

Q: *How powerful is this system?*

A: *can perform arbitrary $\forall n$ -space computations*

Ultimately: **Language**

- An environment created *by us* a few thousand generations ago
- A “last-minute adaptation”
- Hypothesis: it evolved so as to exploit the Brain’s strengths
- Invaluable lens for studying the Brain
- *A deluge of recent experiments!*

Ultimately: Language

- An environment that provides us a few thousand years of experience
- A “last-mile” problem: how to exploit the brain’s capacity for language
- Hypothesis: the brain is designed to exploit the environment by using language
- Invaluable tool for exploiting the Brain
- *A deluge of new ideas and inventions!*



Ultimately: **Language**

- An environment created *by us* a few thousand generations ago
- A “last-minute adaptation”
- Hypothesis: it evolved so as to exploit the Brain’s strengths
- Invaluable lens for studying the Brain
- *A deluge of recent experiments!*

The [Poeppel 2016] experiment

The [Poeppel 2016] experiment


fret ship hill give true melt fans ɓ!ʌc

The [Poeppel 2016] experiment

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

The [Poeppel 2016] experiment

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



The [Poeppel 2016] experiment, stage II

...

The [Poeppel 2016] experiment, stage II

bad cats eat fish

The [Poeppel 2016] experiment, stage II


bad cats eat fish new plan gave joy

The [Poeppel 2016] experiment, stage II



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

The [Poeppel 2016] experiment, stage II

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



My interpretation



[Frankland & Greene PNAS 2015]

“The ball hit the truck”

VS

“The truck hit the ball”

Different areas of the STG responded to “truck” in the two sentences [*Recall relations...*]

[Frankland & Greene PNAS 2015]

“The ball hit the truck”

VS

“The truck hit the ball”

Different areas of the STG responded to “truck” in the two sentences [*Recall relations...*]

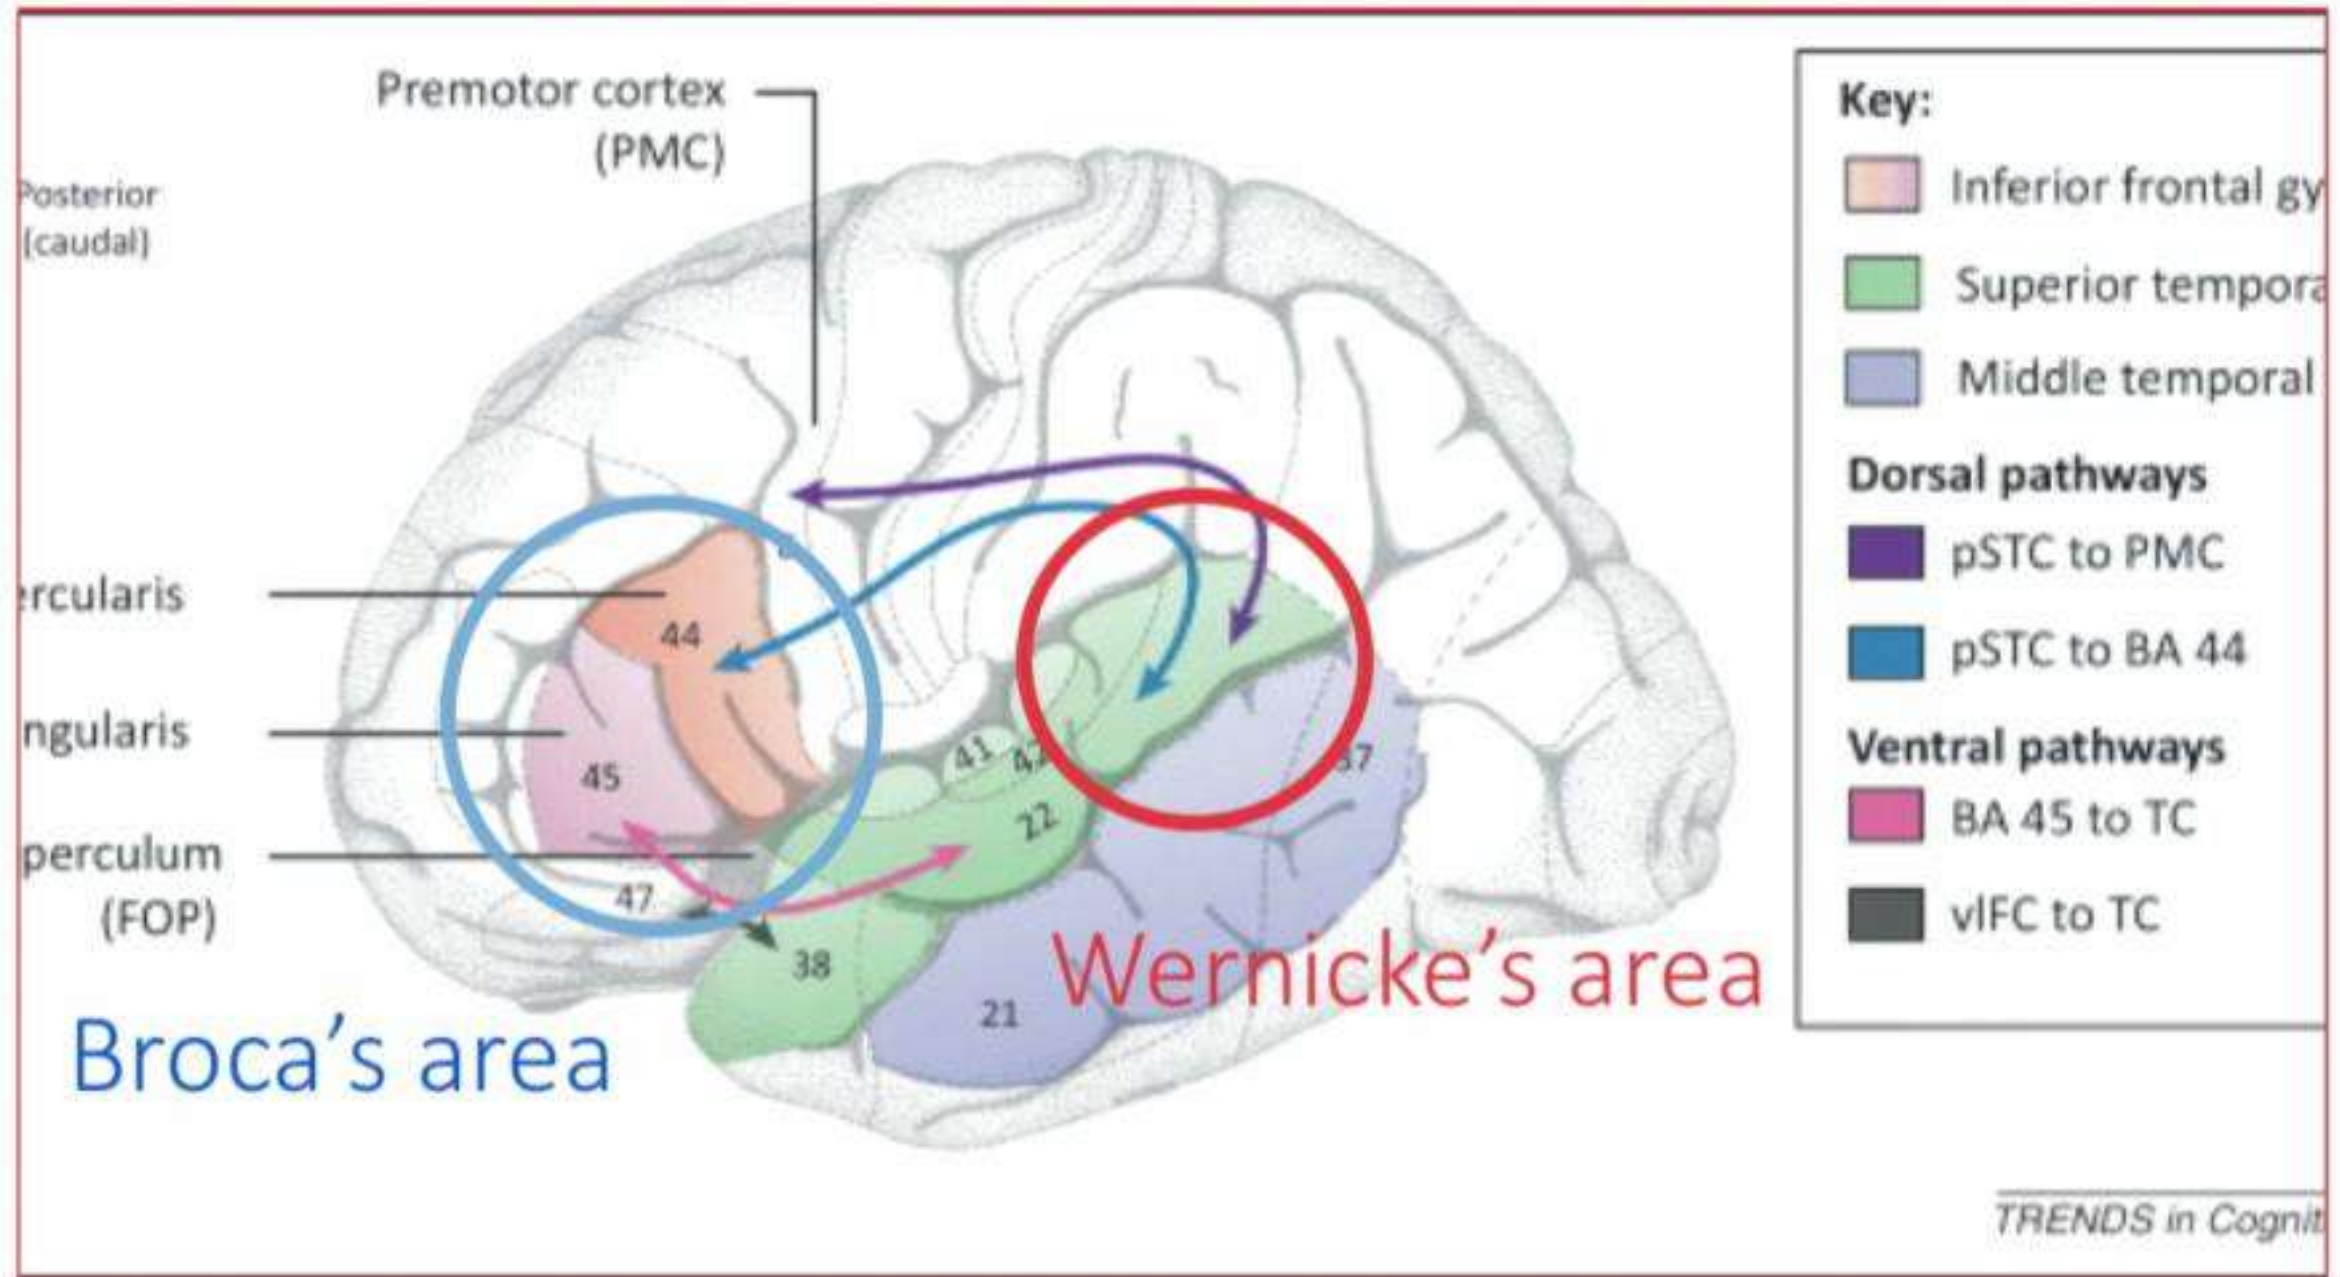
The first area also responded to

“The truck was hit by the ball”

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

- The completion of phrases, and especially of sentences, **lights up parts of Broca's area**


[ZF 2010]: Neural pathways for syntax?





a brain architecture for syntax



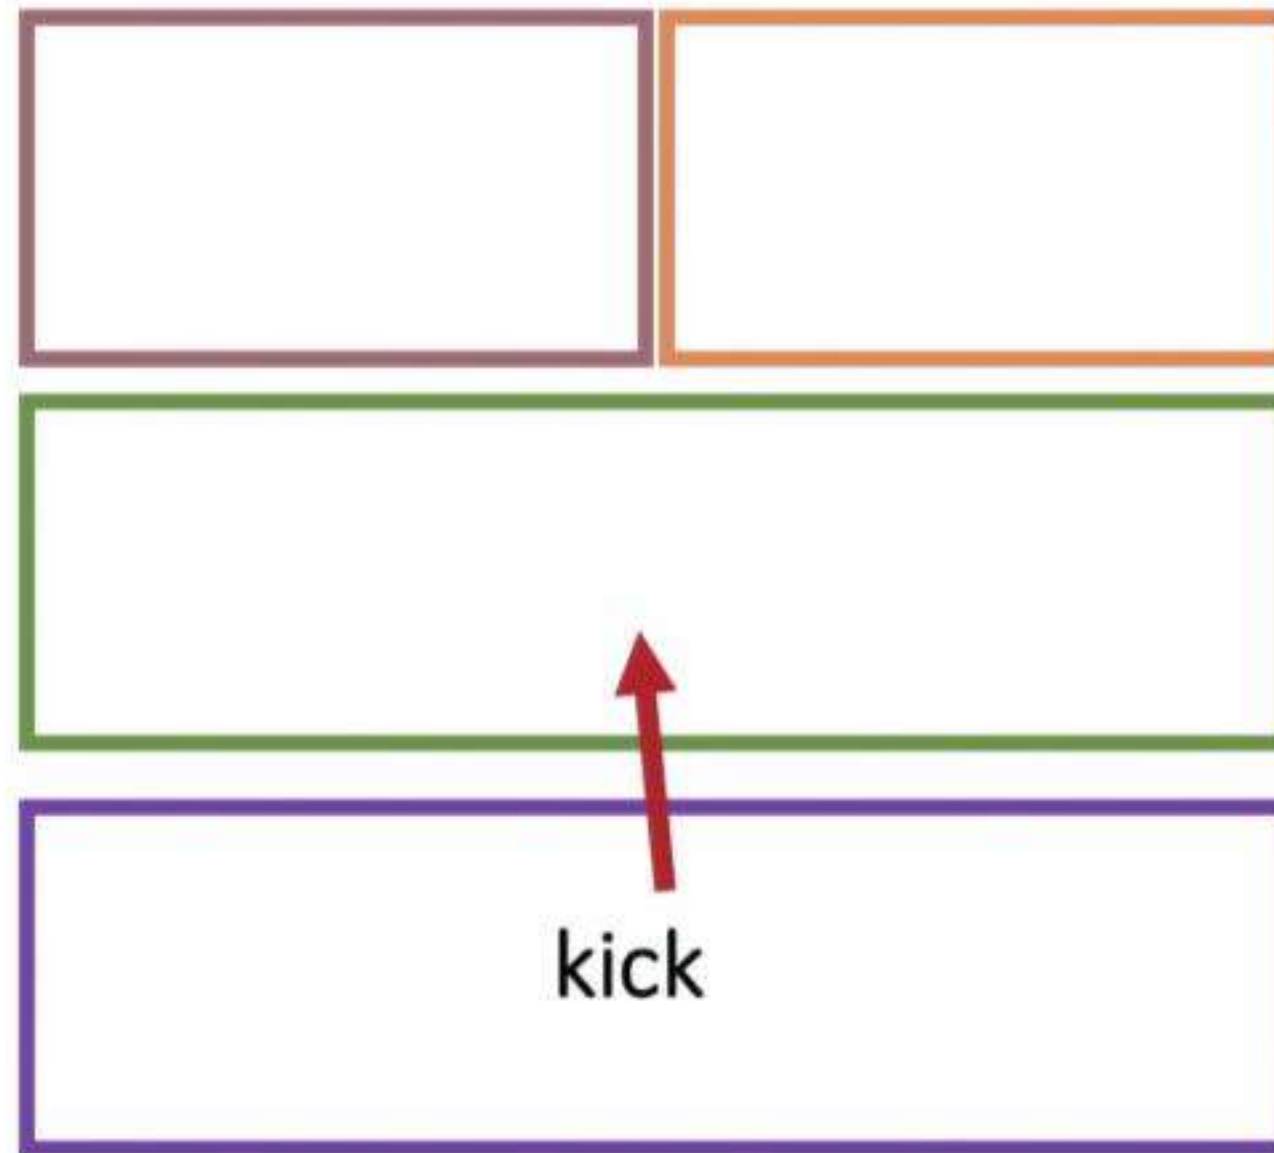
a brain architecture for syntax




a brain architecture for syntax




a brain architecture for syntax



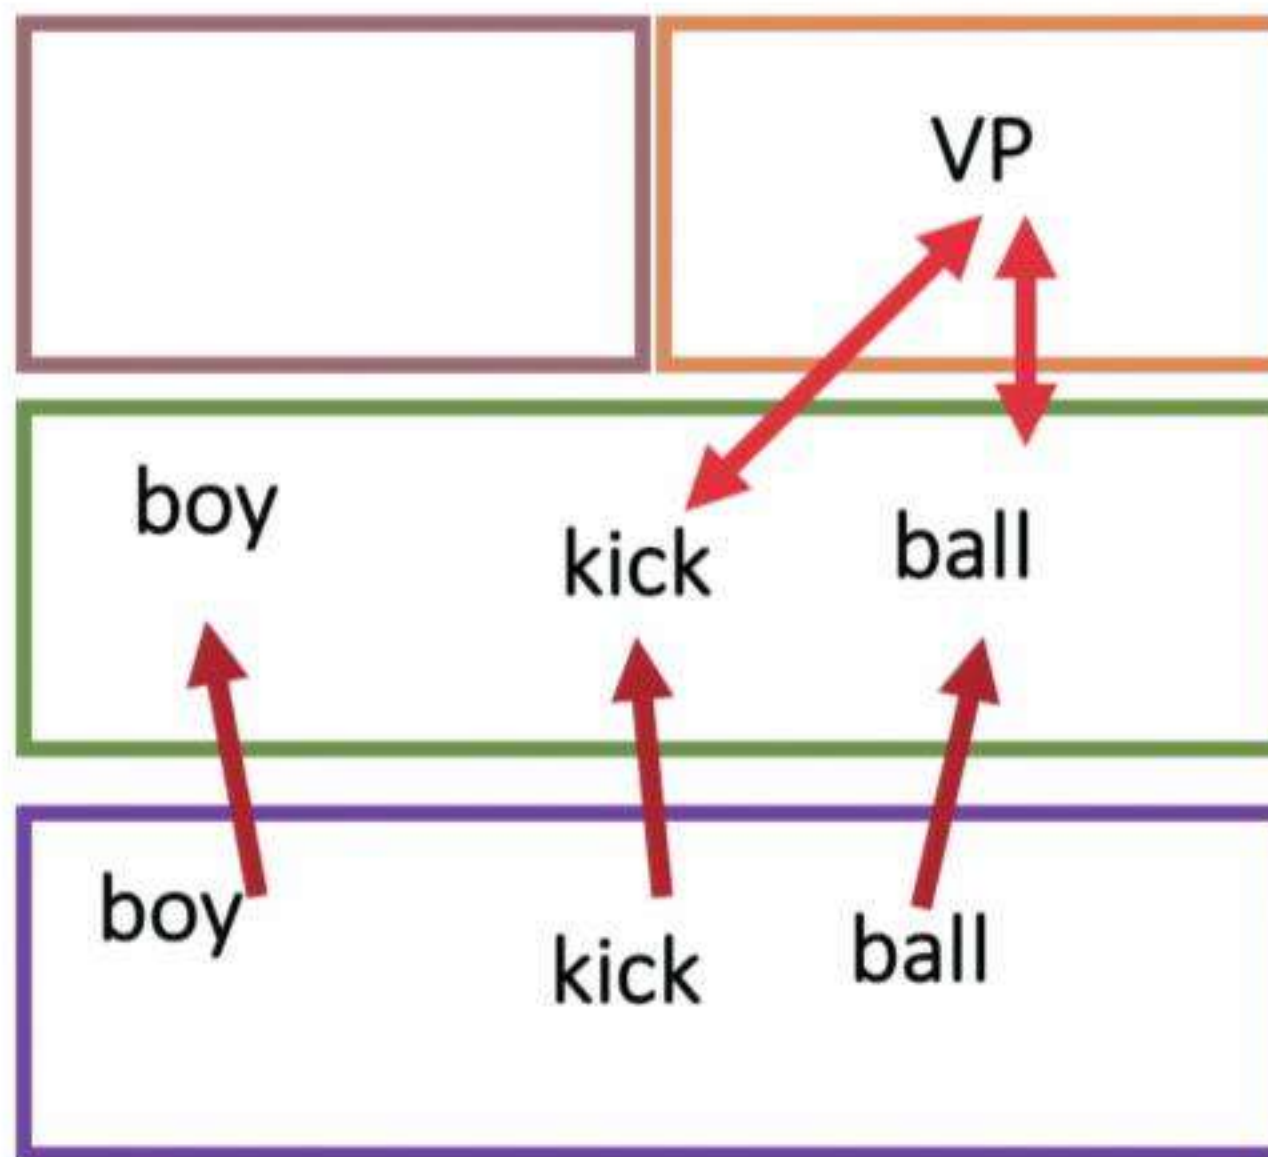
a brain architecture for syntax




a brain architecture for syntax



a brain architecture for syntax




a brain architecture for syntax




In conclusion:

What **C Golgi** and **S Ramon y Cajal** saw



The assembly hypothesis



Soooooo...

- The study of the Brain is fascinating and bottomless
- Assemblies and their operations may be one productive path to thinking about computation in the brain
- Are they the seat of Axels's logic?
- How do assemblies learn and predict?
- How can one test/verify/falsify the Assembly Hypothesis through experiments?



Santosh Vempala
GaTech



Wolfgang Maass
TUGraz

Larry Abbott
Columbia



Dan Mitropolsky
Columbia



Mike Collins
Columbia

Thanks!

PASAJE GALVEZ

