Programming in the Brain

Randall C. O'Reilly
University of Colorado Boulder
eCortex, Inc

Parallel vs. Serial

Turing machines are universal *because* they are serial (and it doesn't take much hardware)

Conversely, some (large class of) problems *cannot* be solved in parallel (dependencies..)

To achieve universal flexibility, inherently parallel neural processing must be come serial

More Serial Advantages

Solving novel tasks requires novel combinations of existing subroutines (or new subroutines, which is much harder)

Serialization allows generic recombination of subroutines.

Parallel processing requires separate hardware for each routine, and connecting them is tricky..

Parallel is great too..

Fast, high-dimensional constraint satisfaction

Parallel gradient search learning: pursue many different solutions in parallel (serial search takes until the end of the universe..)

The *human* brain is the best known combination of parallel and serial processing! (animals are too parallel)

The Biological Architecture

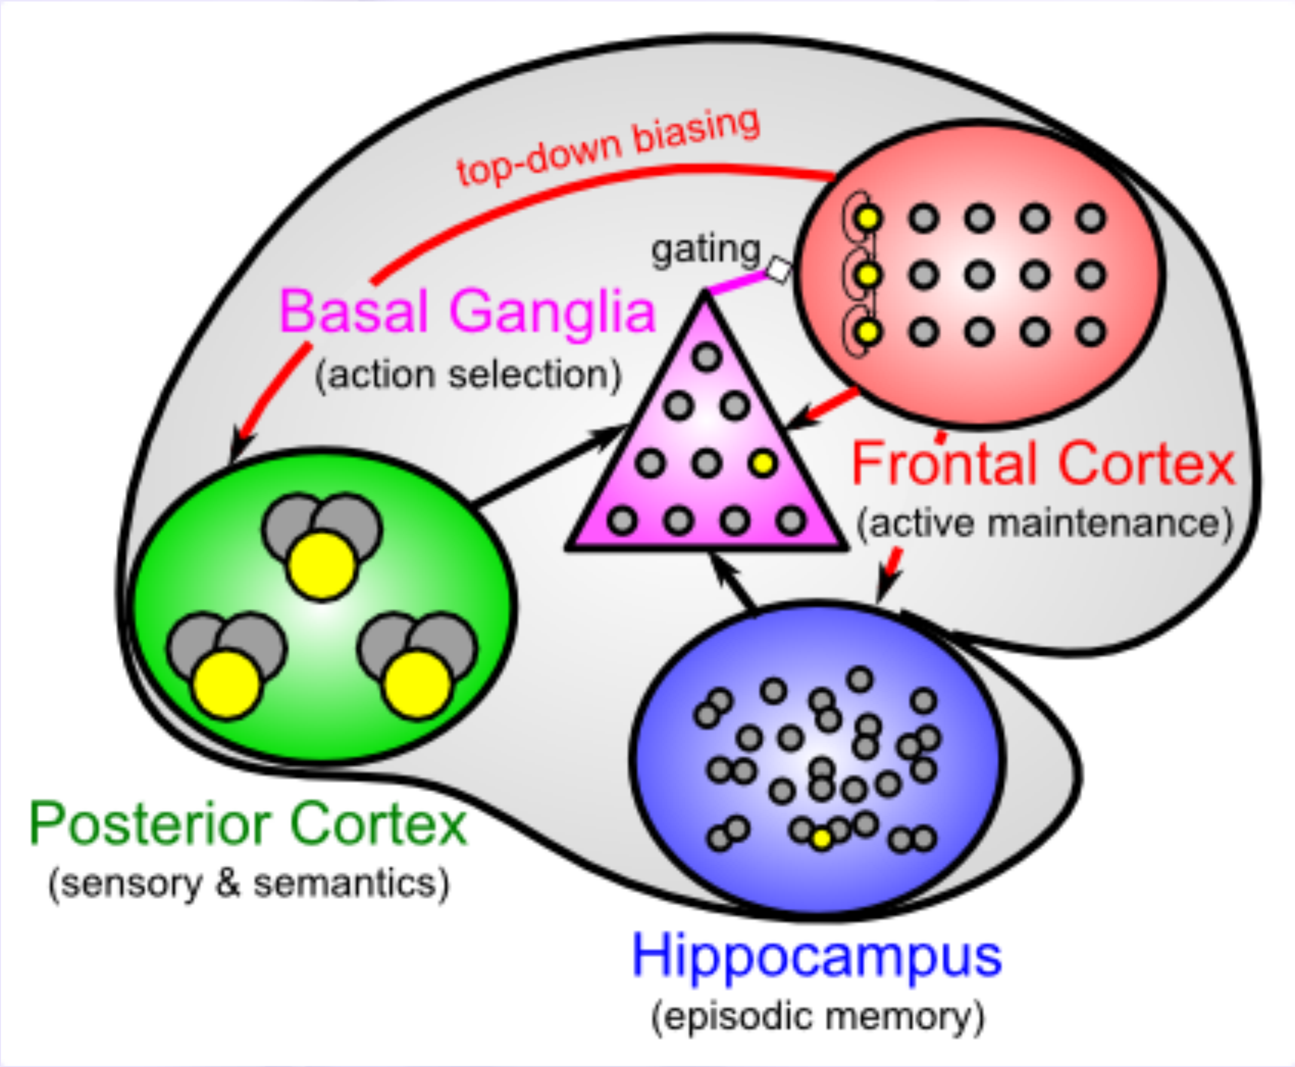

PC is parallel

BG / FC enables serial via *gating* and *active task maintenance*

Hippo is fast, high-capacity memory cache

Proof of Concept: ACT-R

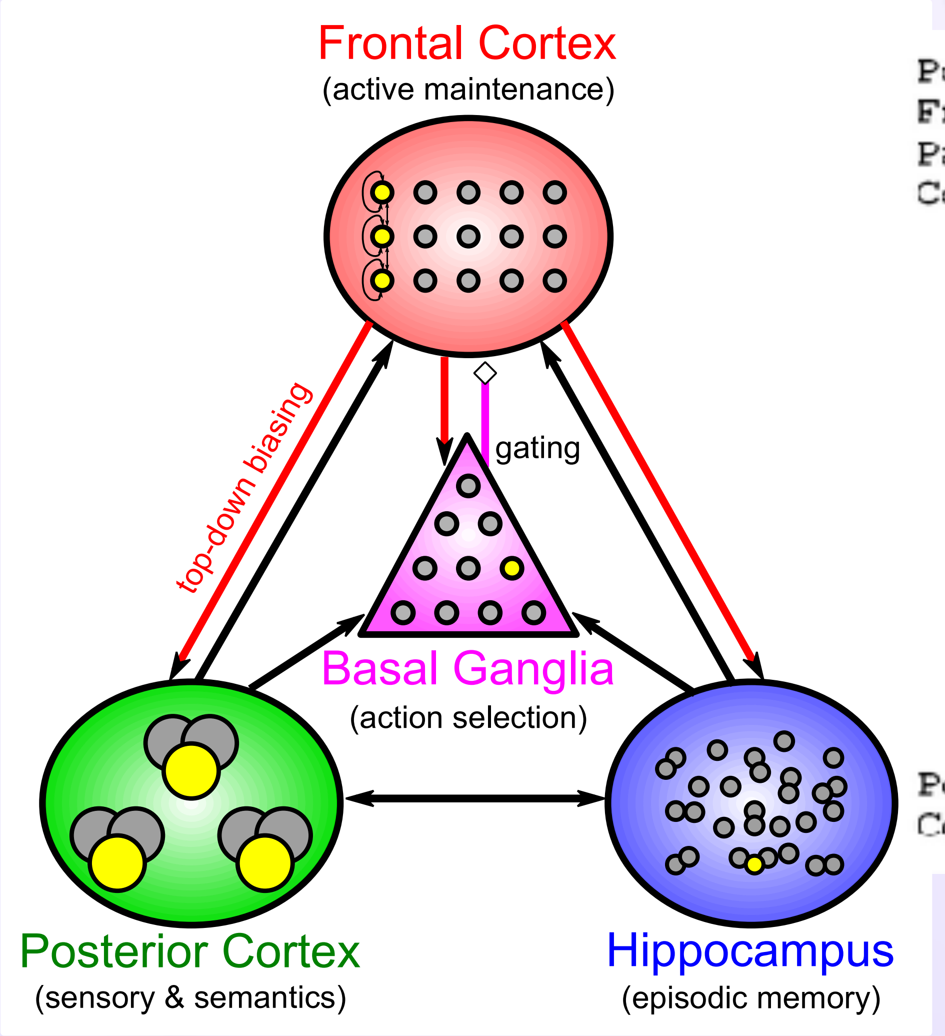

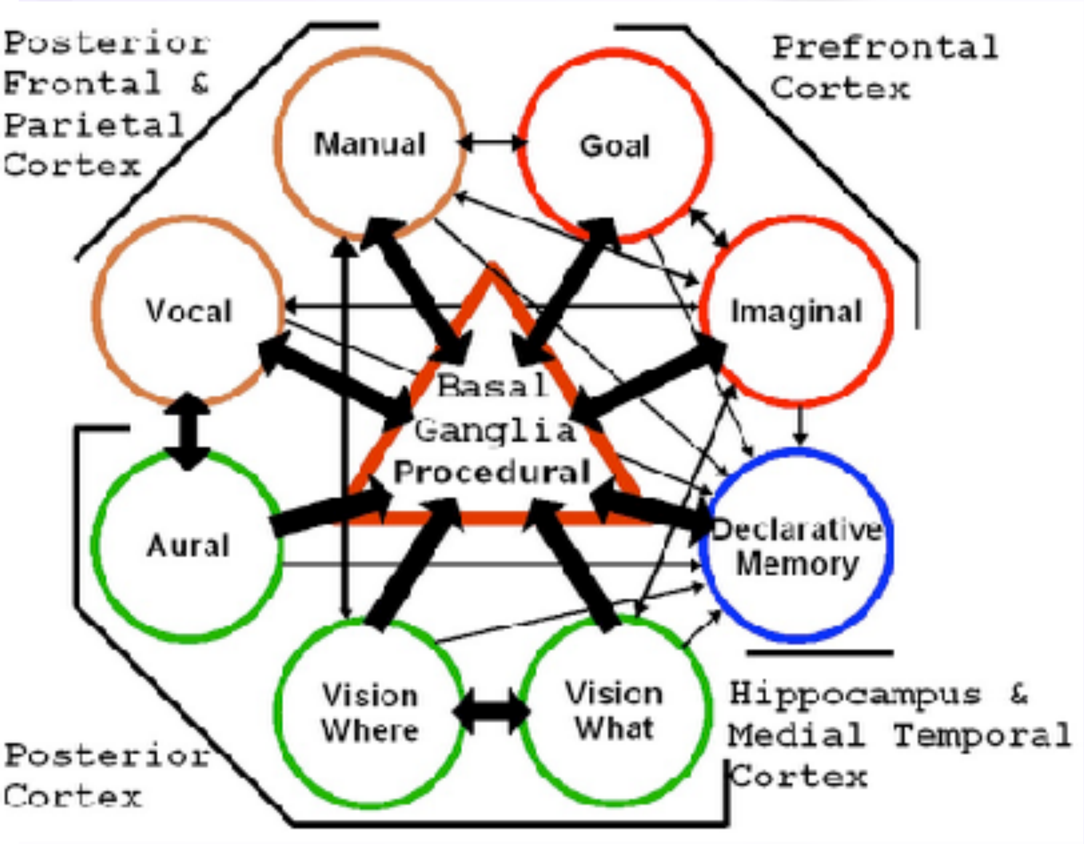

BG production system forces serialization -> flexible combination of productions. Goal buffer & declarative memory coordinate.

Leabra Biologically-based Cognitive Architecture

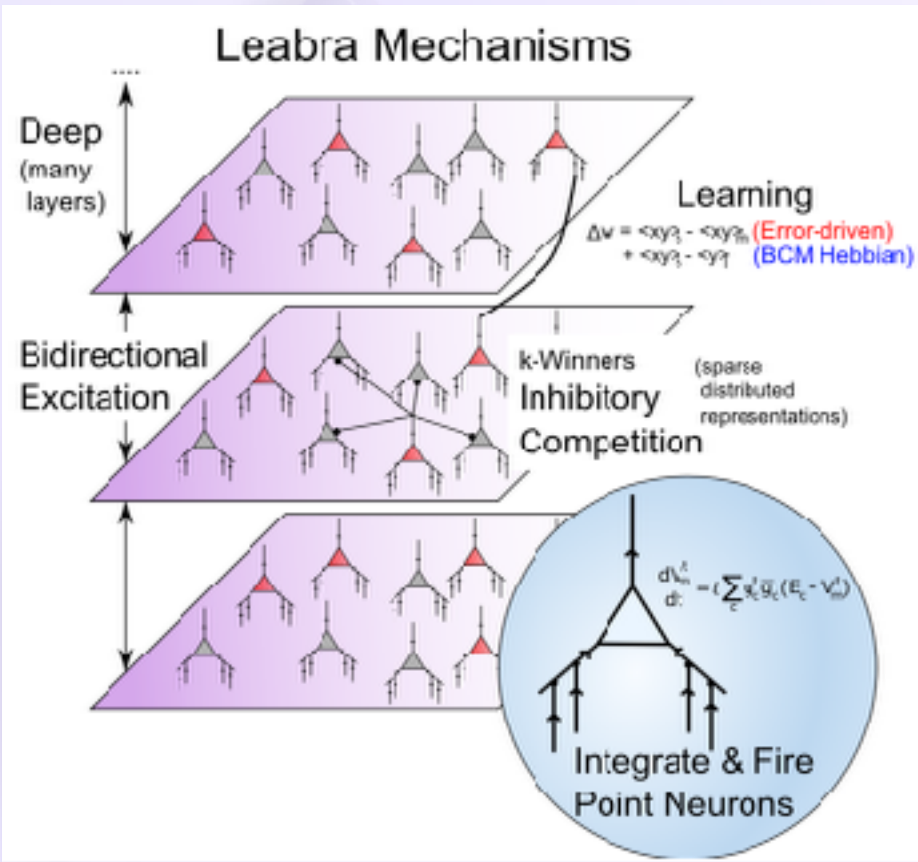

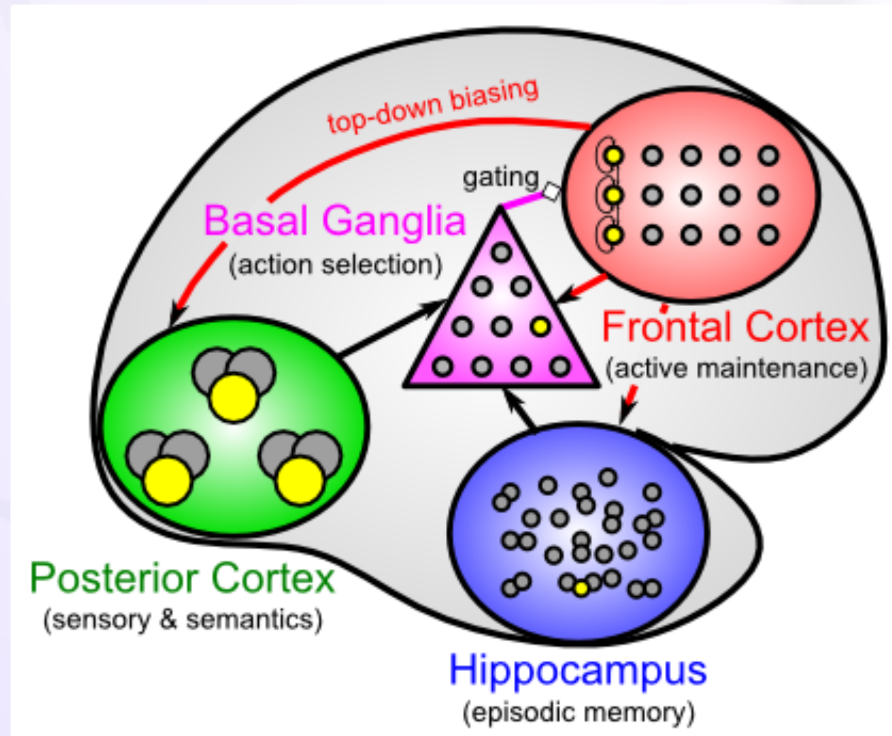

Same framework accounts for wide range of cognitive neuroscience phenomena: perception, attention, motor control and action selection, learning & memory, language, executive function – ***all built out of the same neurons.***
http://ccnbook.colorado.edu

PBWM System (bio LSTM)

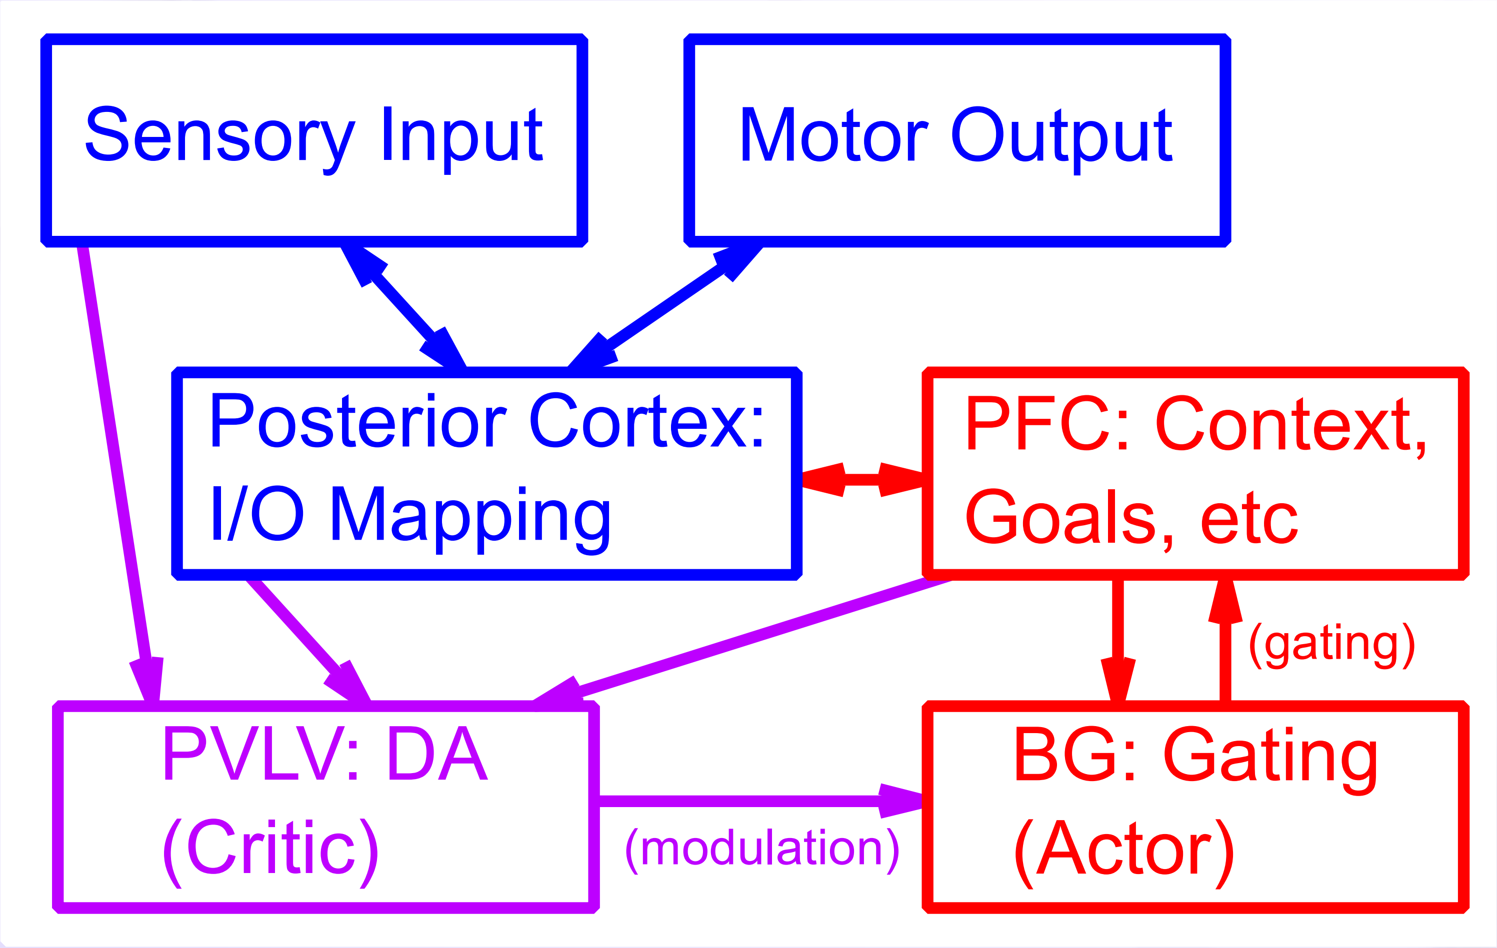

Three levels of modulation to get anything done..

Dynamic, Adaptive Gating (BG)

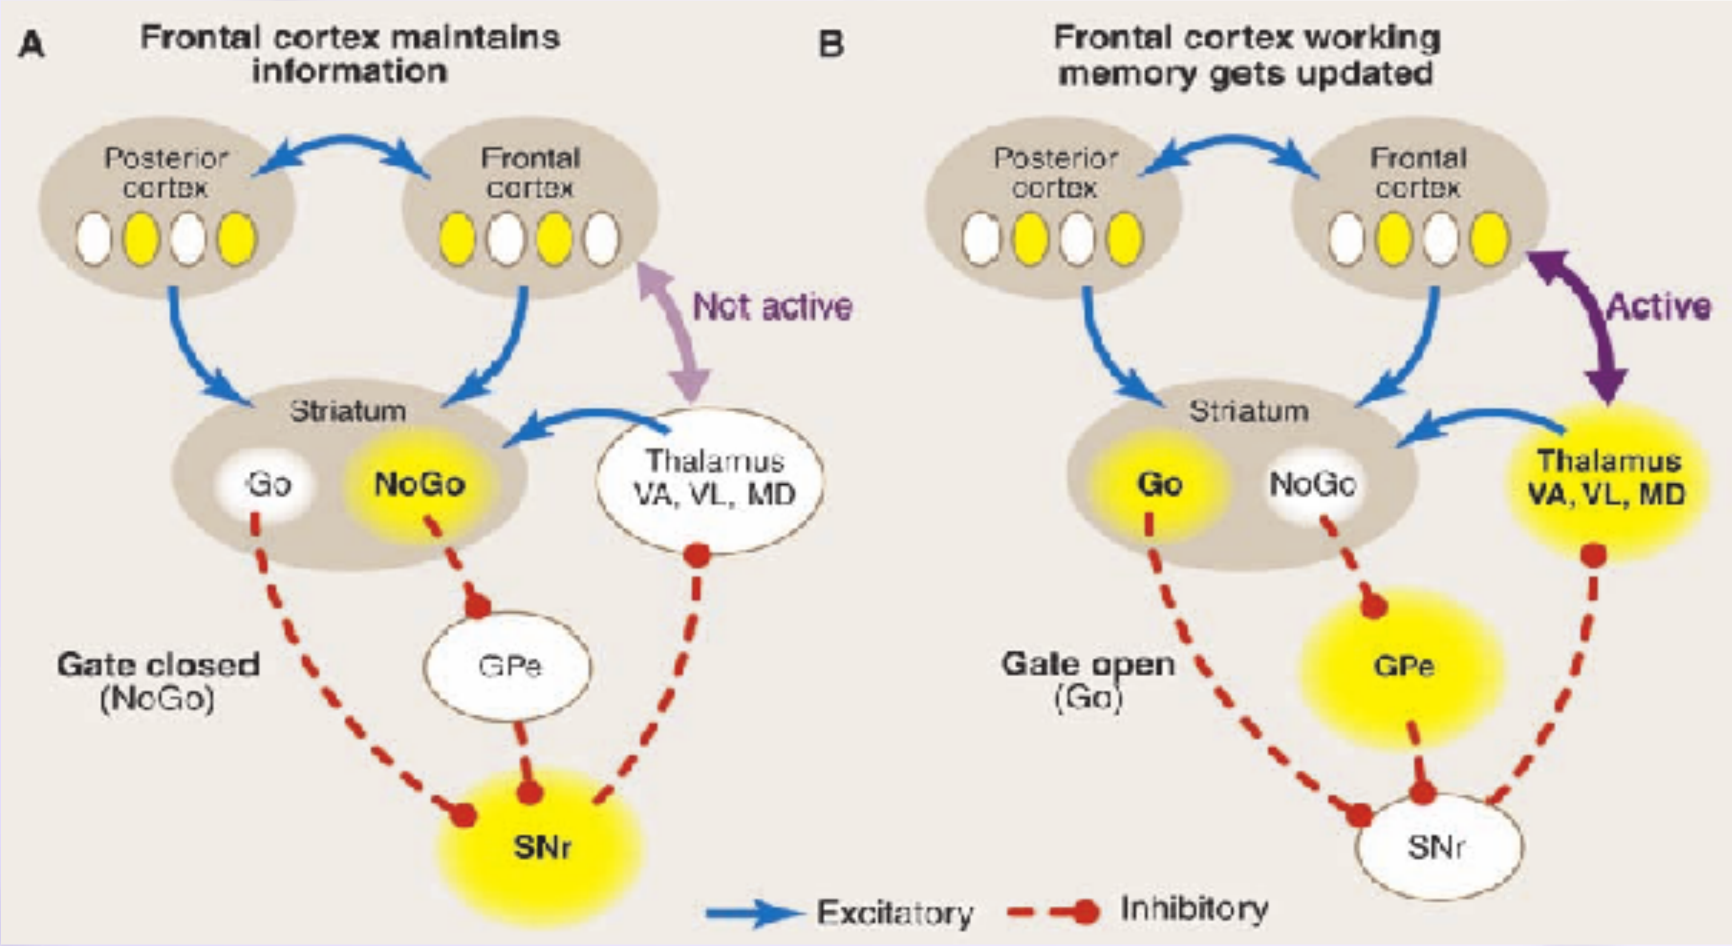

- BG Toggles PFC bistable states

Basal Ganglia Reward Learning

(Frank, 2005…; O'Reilly & Frank 2006)

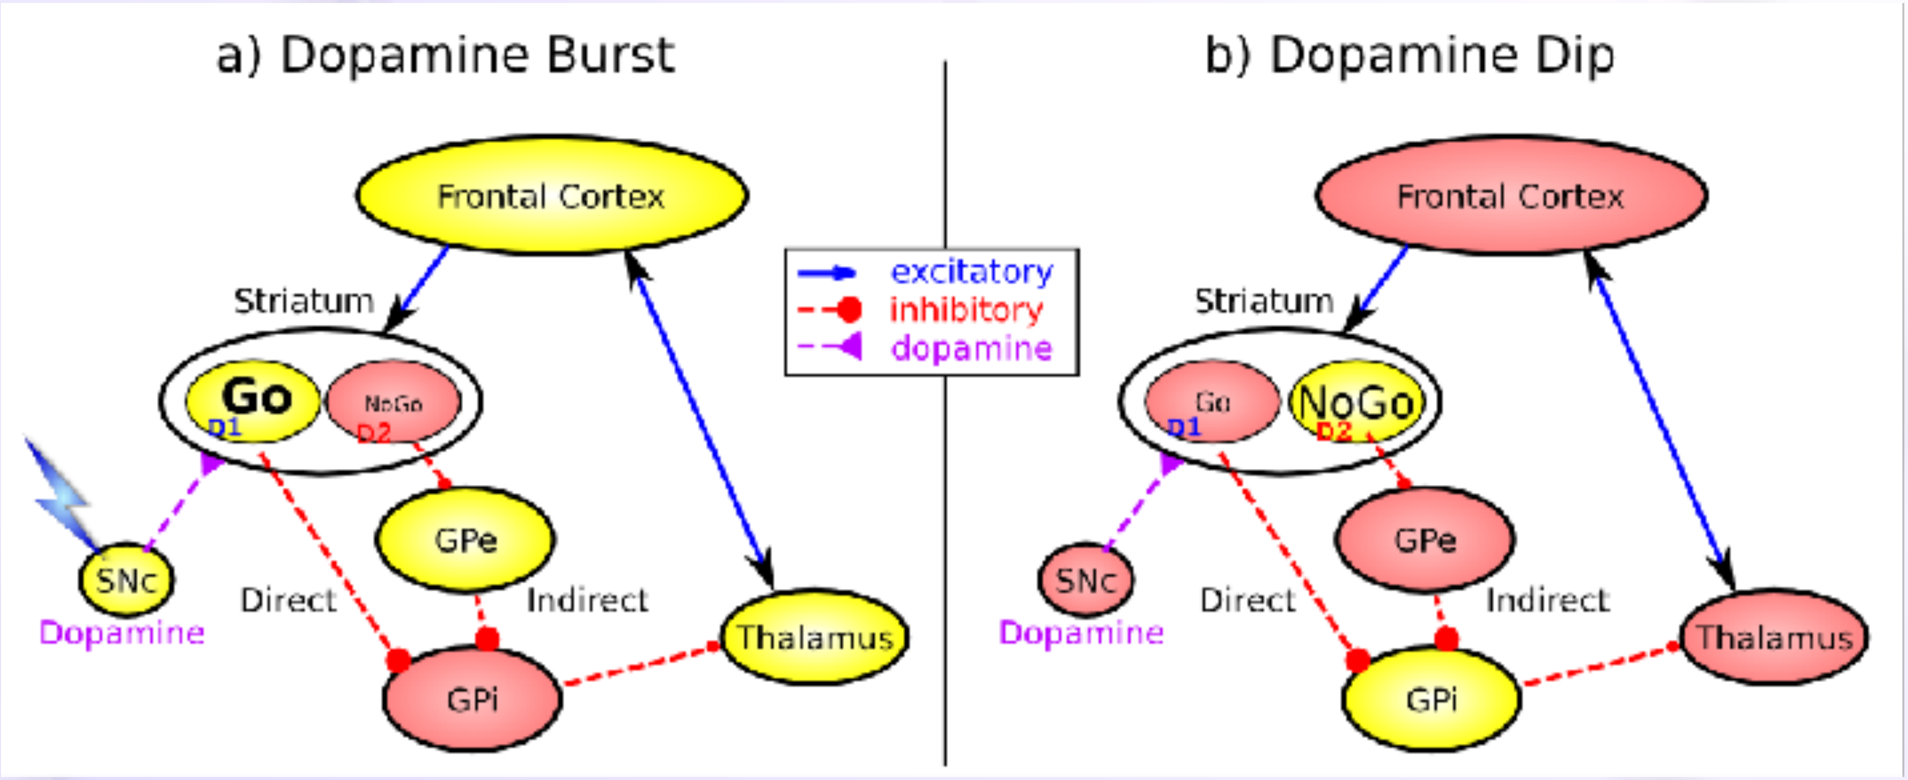

· Trial and error learning: action of updating PFC is evaluated in terms of previous success associated with PFC state

BG Gates Super -> Deep

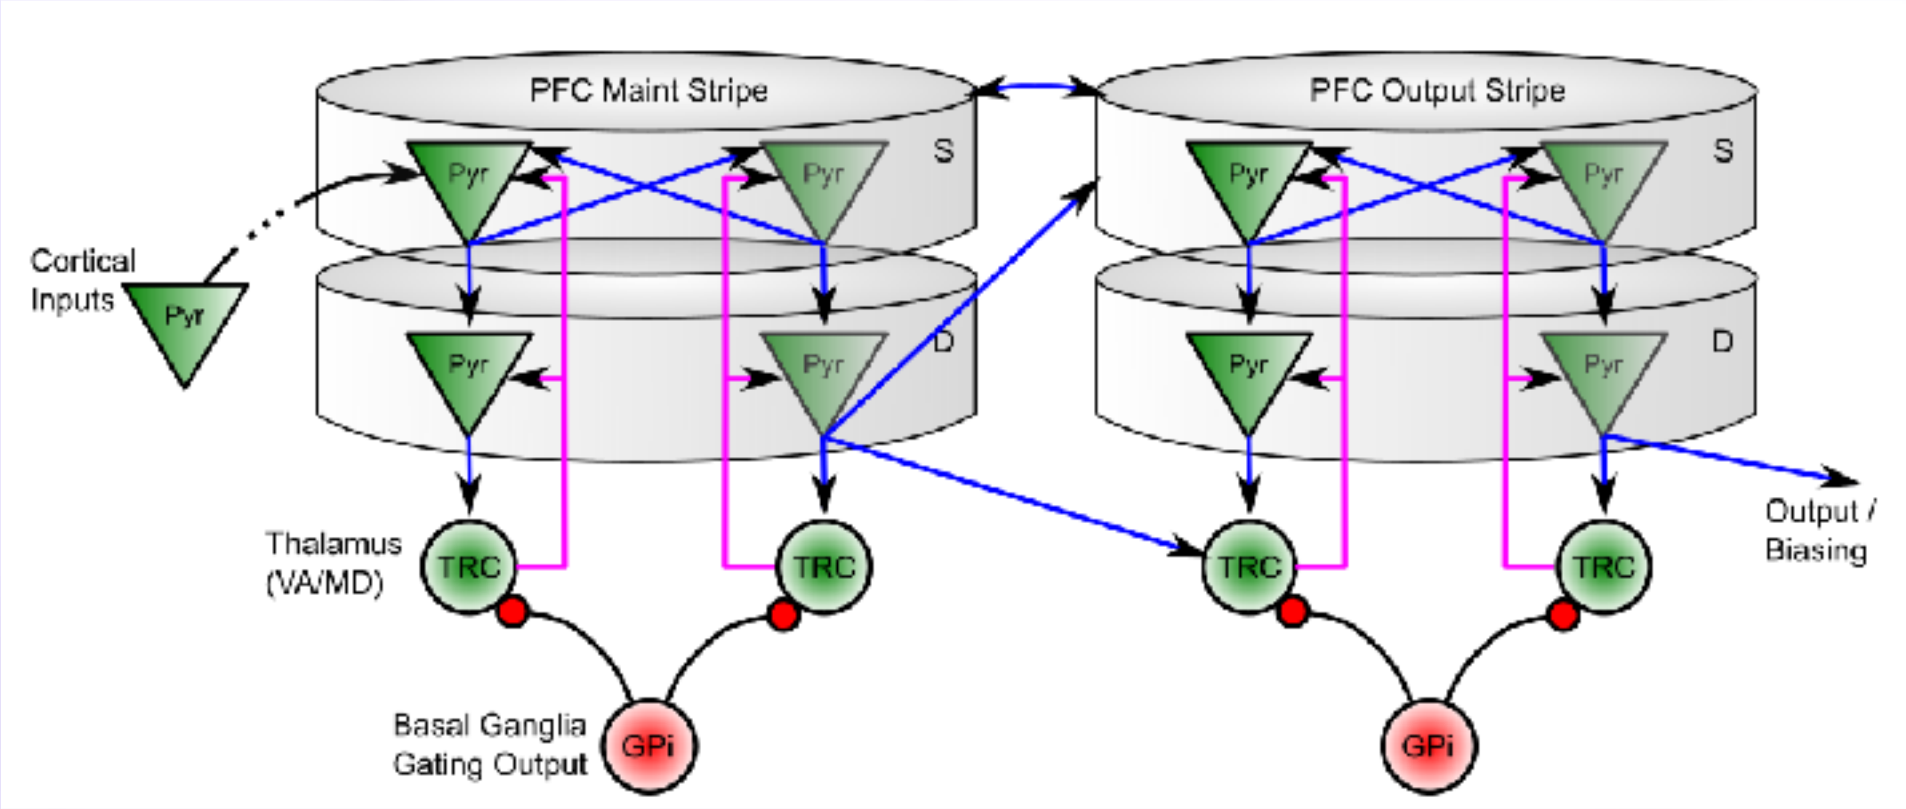

- Maintenance via Thalamocortical loops (TRC <-> Deep), BG disinhibits
- Superficial reflects inputs *and* maintenance
- Separate Maintenance vs. Output PFC / BG *stripes*

BG Gates Super -> Deep

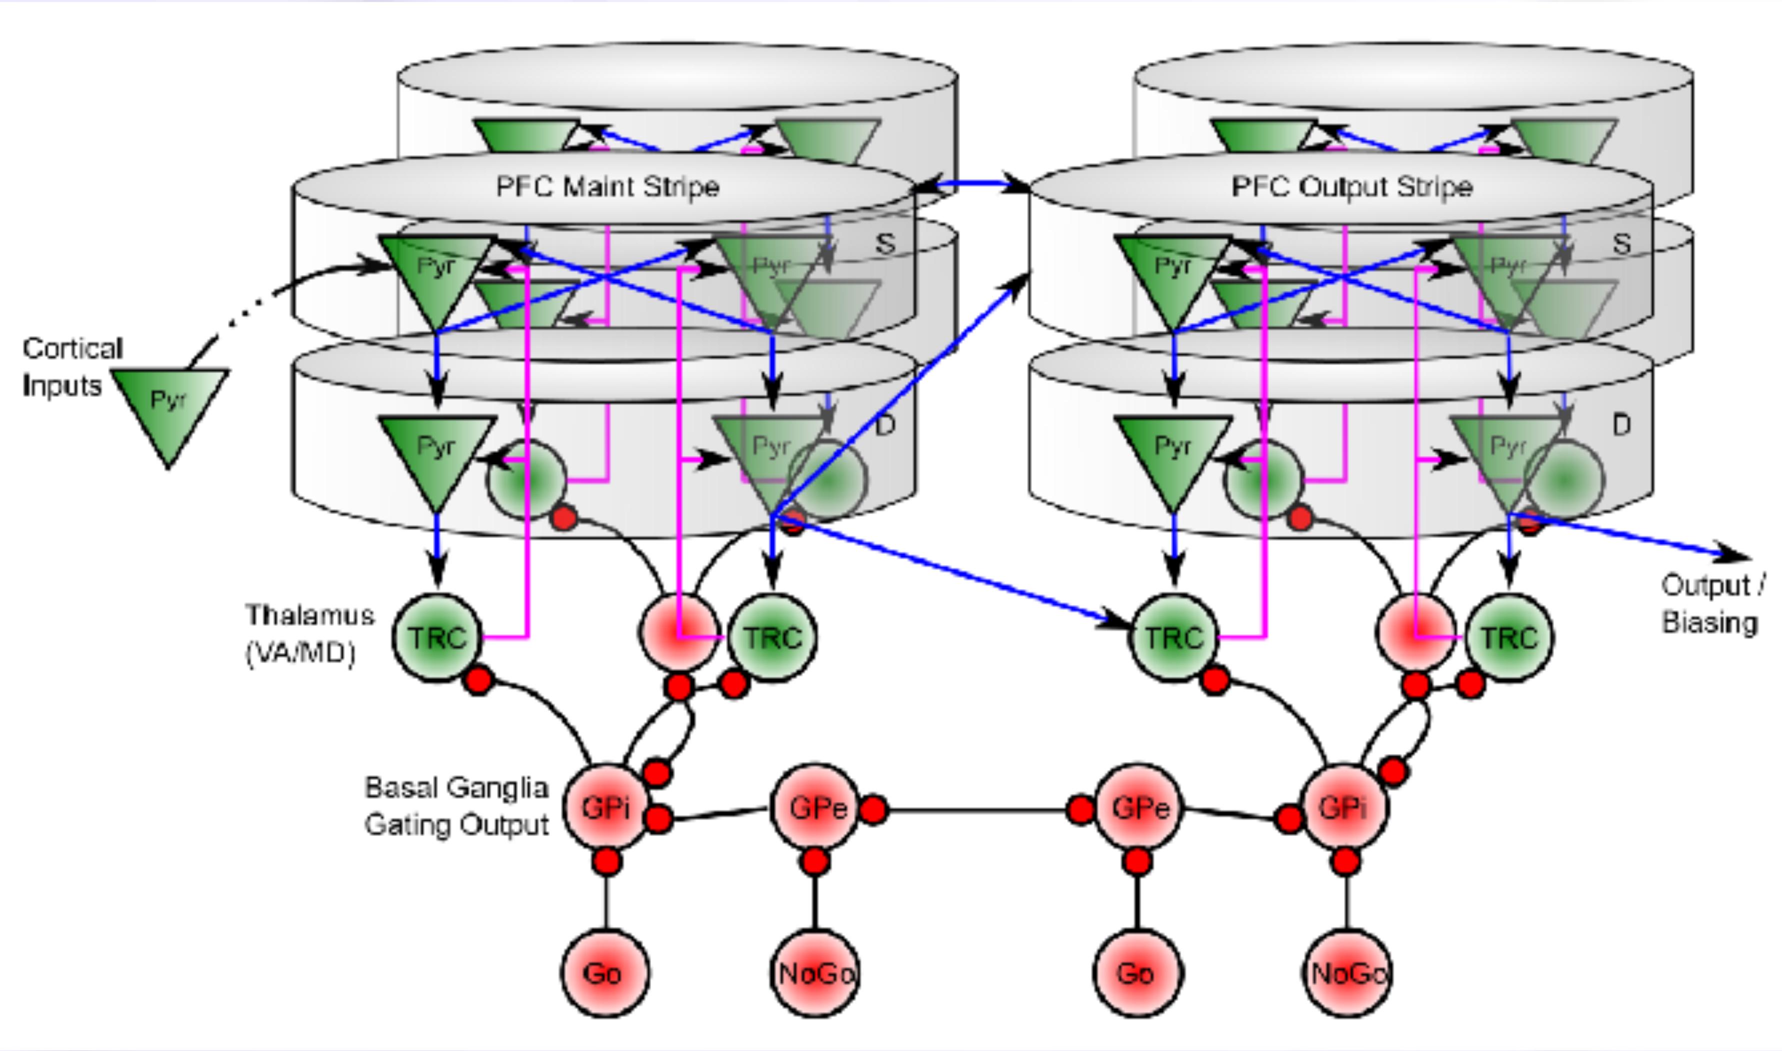

- Competition in GPi, GPe (and striatum) between Maint vs. Output, and diff stripes
- Pre-competition in GPe for NoGo veto power (cannot veto everything!)
- Asymmetric learning rate for dopamine dips vs. bursts on NoGo (D2) – hypercritical..

PBWM Applications

- Learns from raw trial & error experience to perform complex working memory tasks, including N-Back (Chatham et al, 2011), 1-2-AX CPT (O'Reilly & Frank, 2006), Keep Track (Friedman et al, in prep)
- Role of BG and DA on working memory tested in a number of expts (e.g., Frank & O'Reilly, 2006)

Demo of SIR model

PFC is Exec, driven by bottom line..

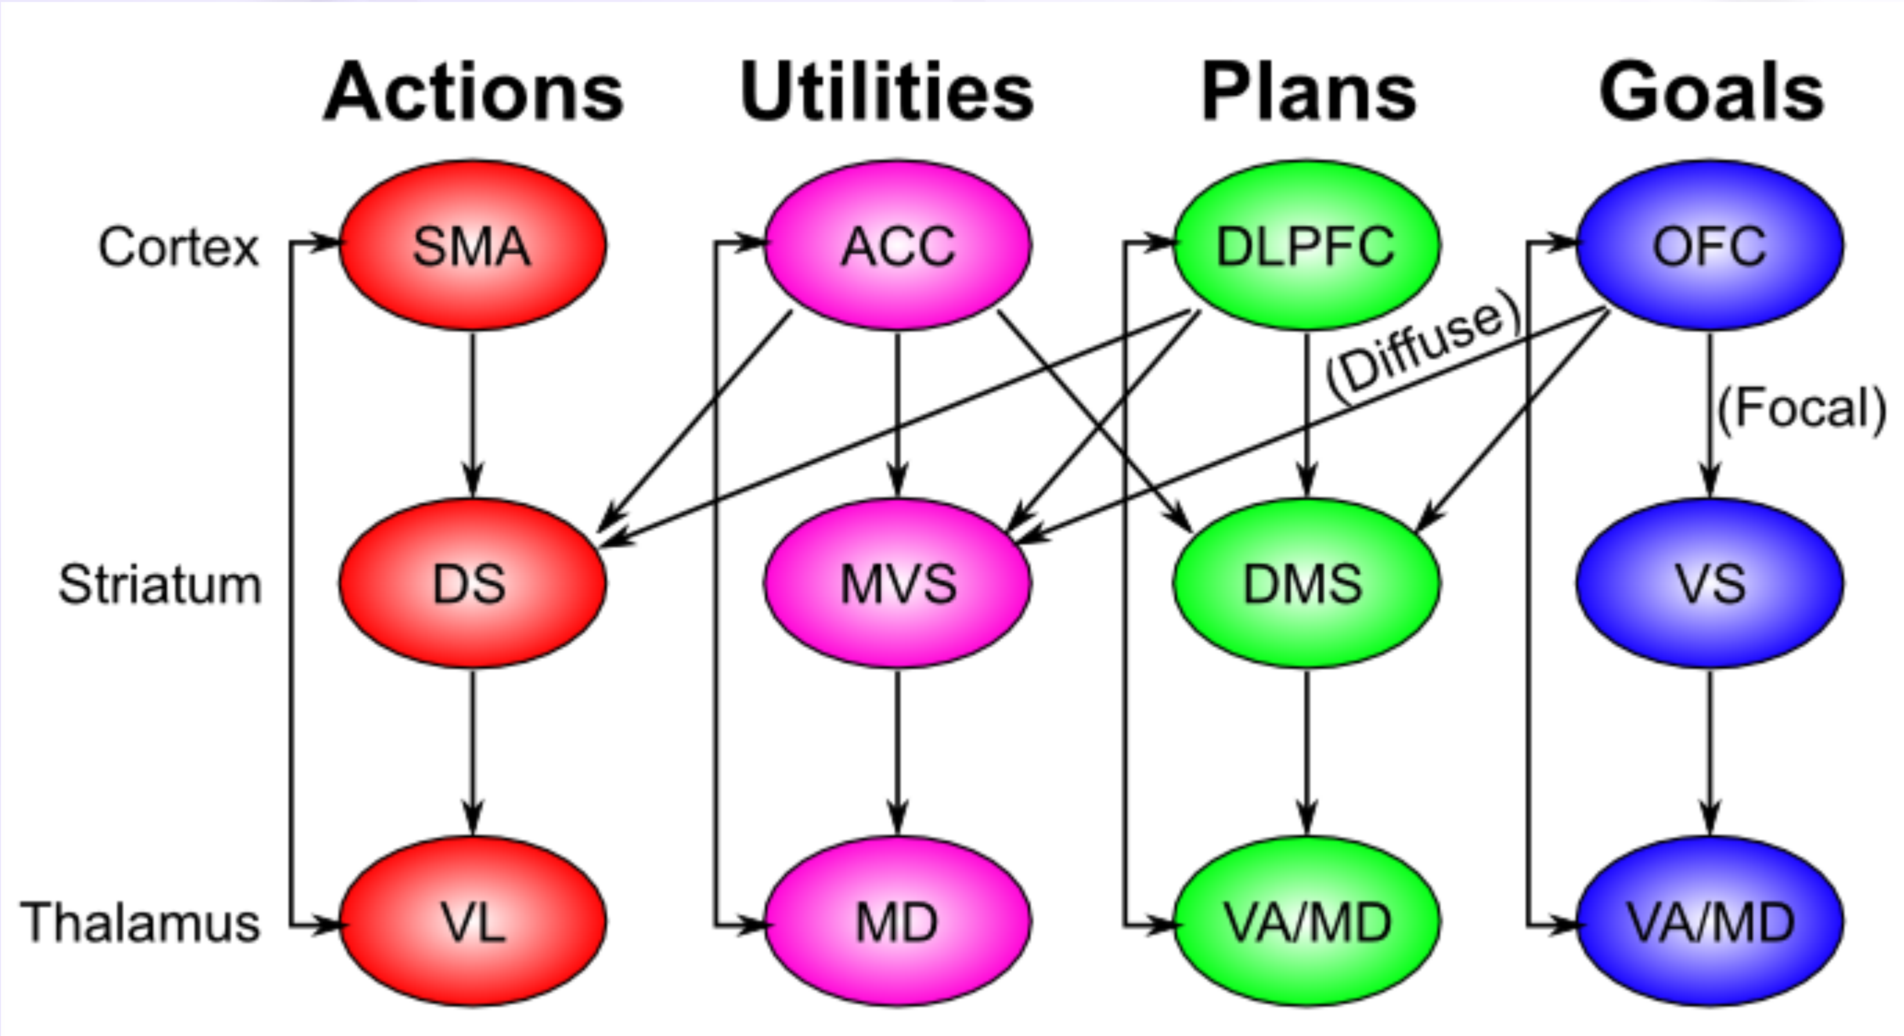

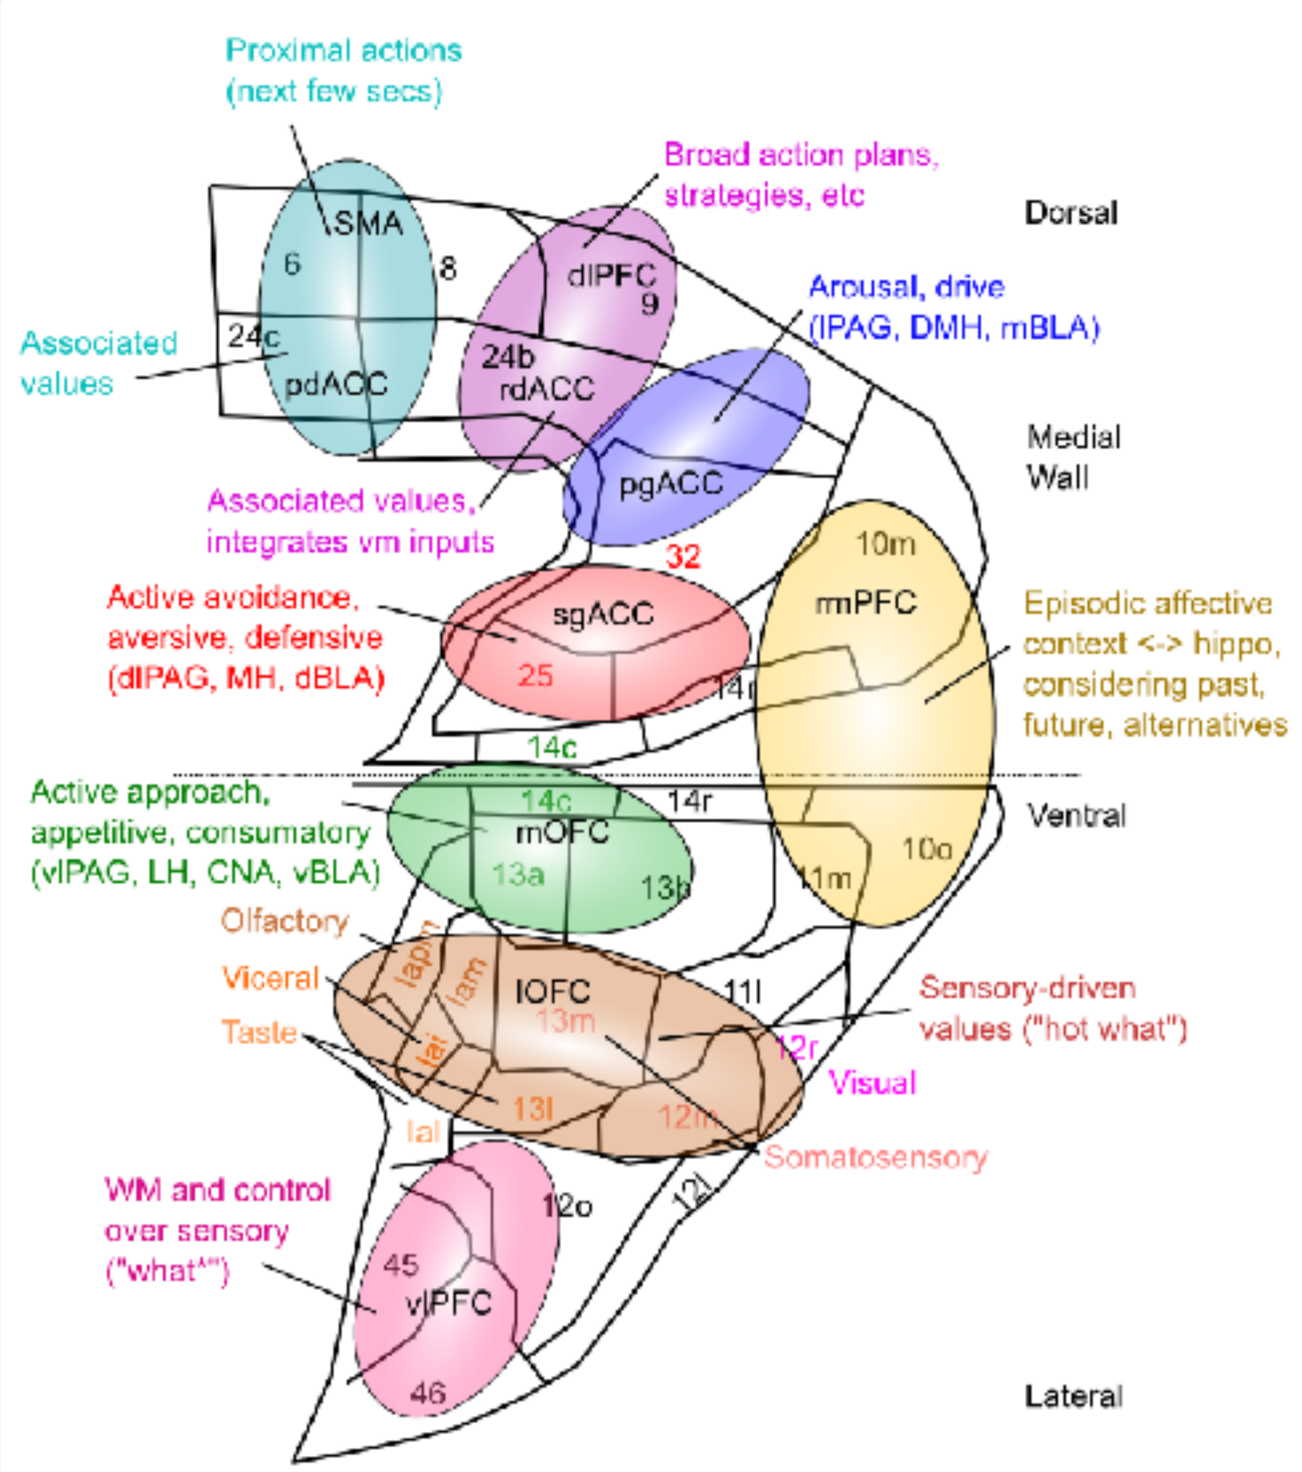

Medial Frontal Map of Values

This is your emotional life

Executive Function

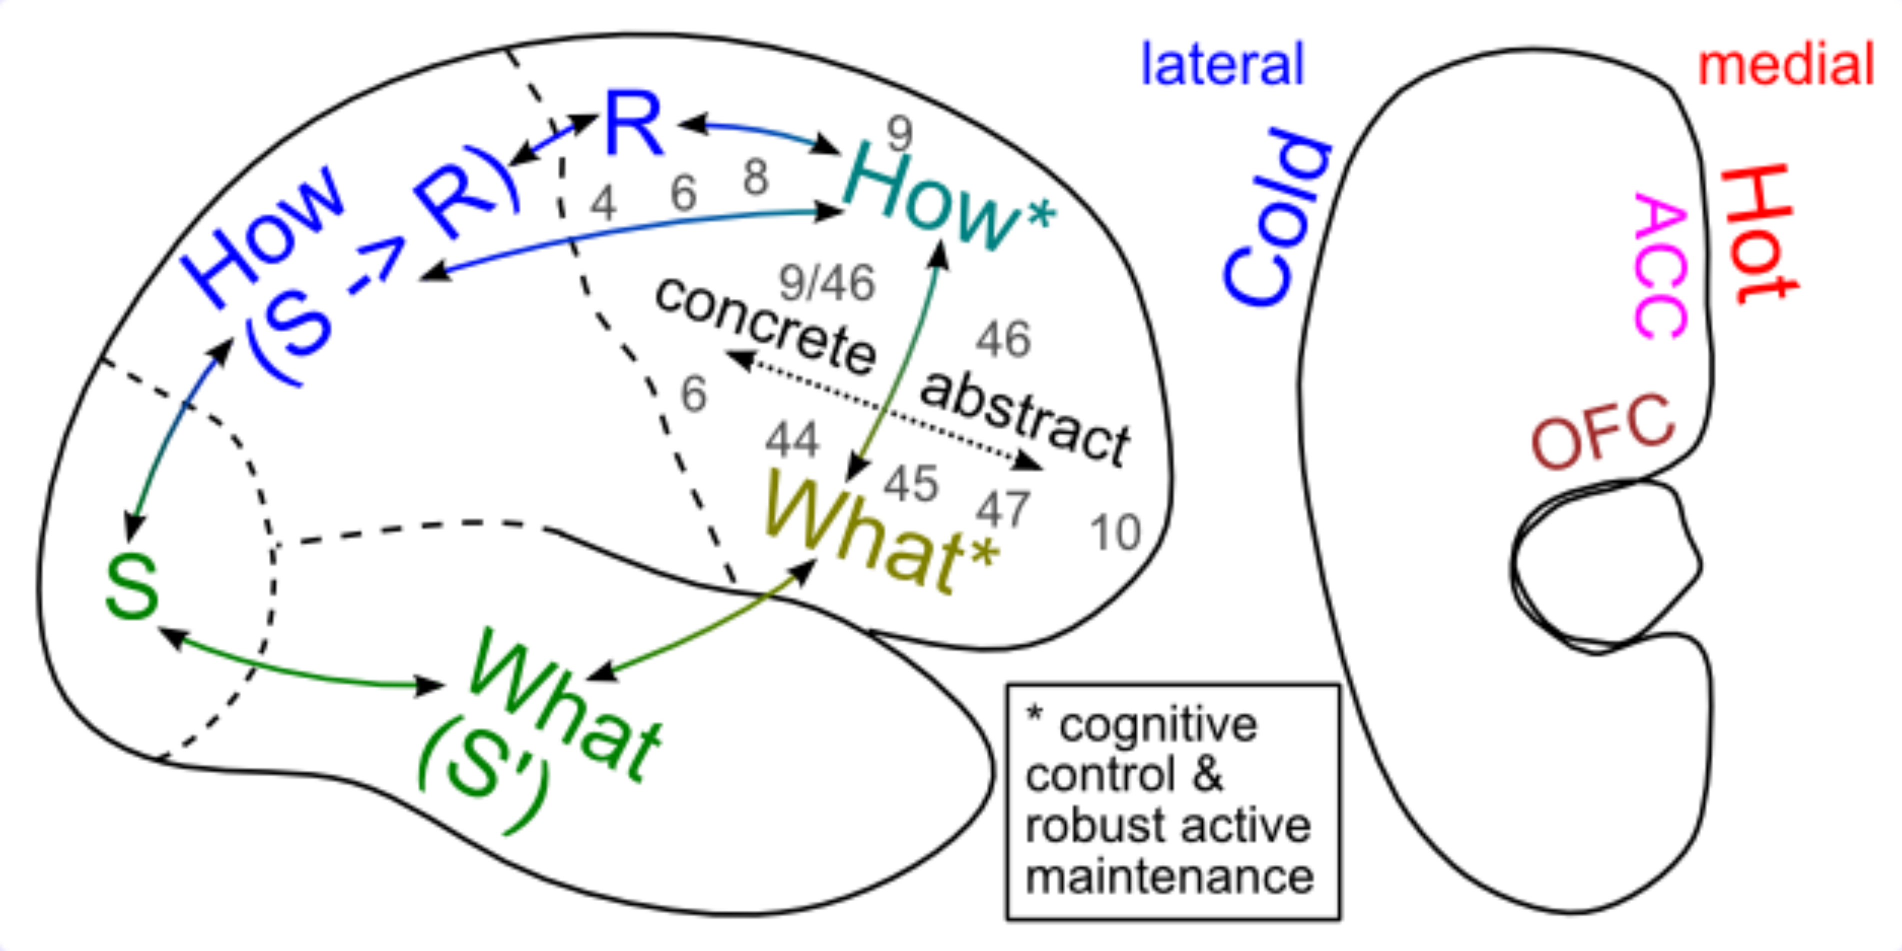

Hippocampal System

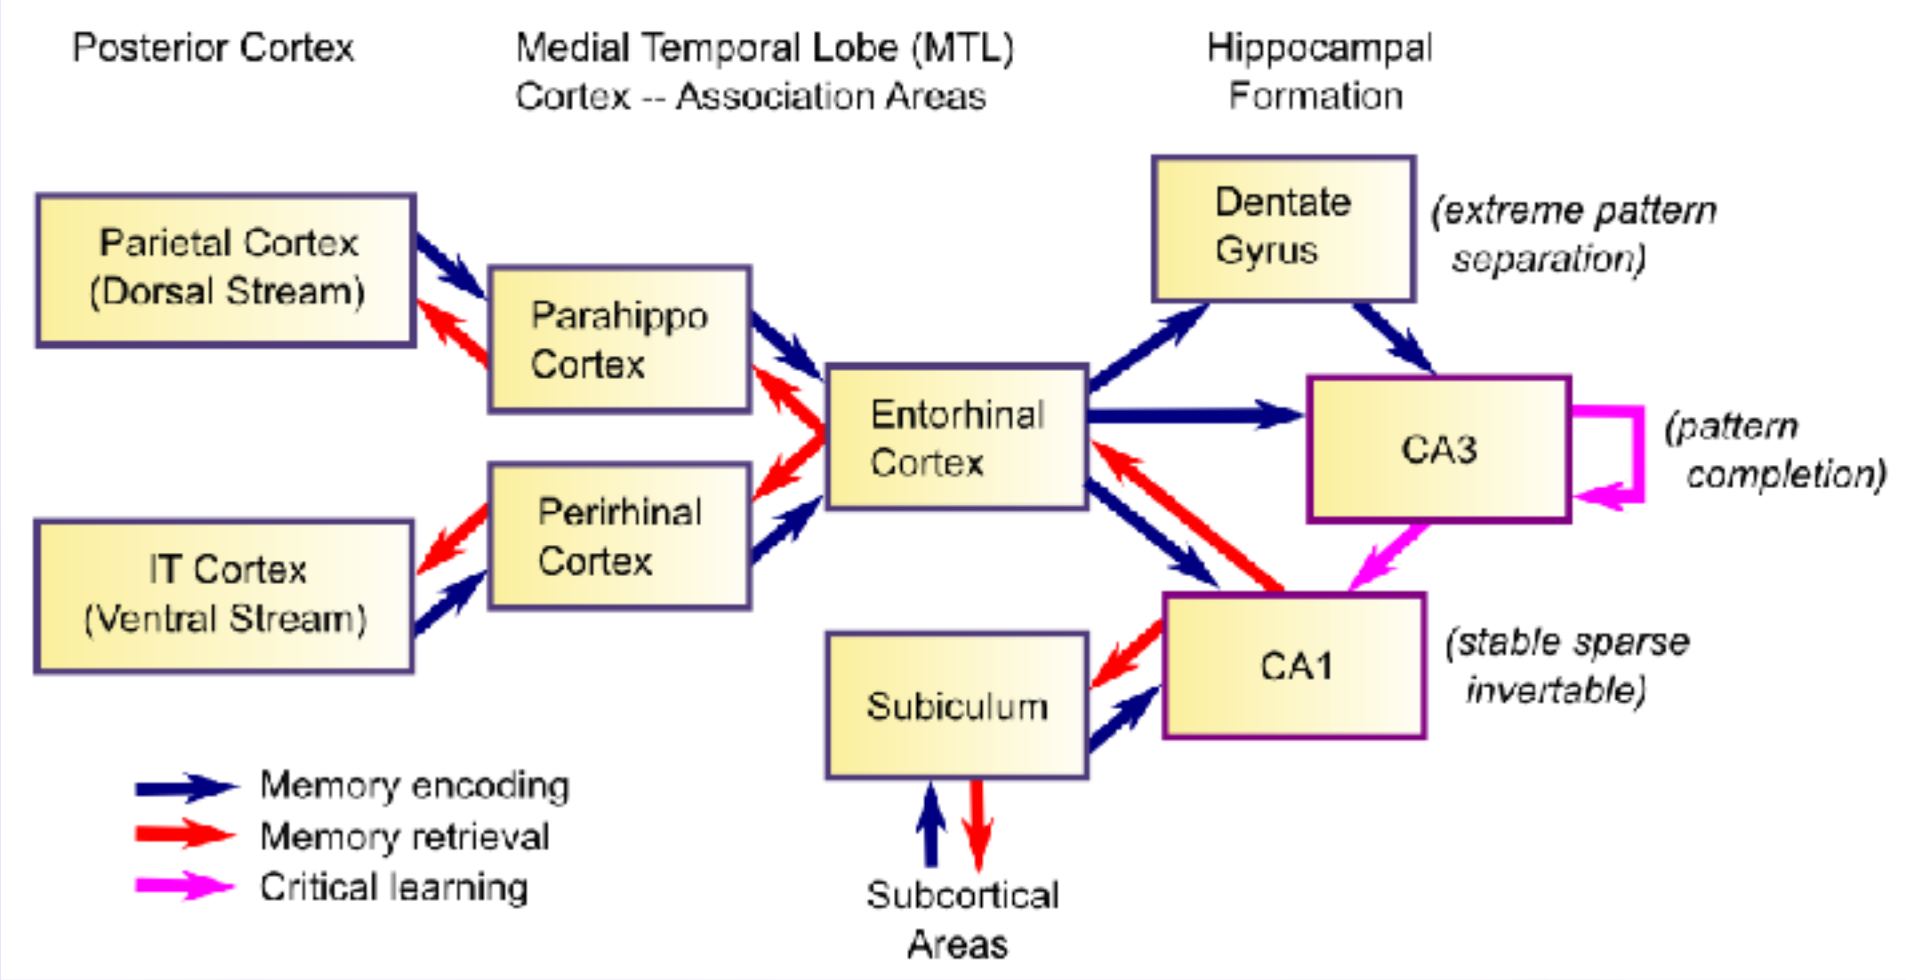

Sparse = Pattern Separation = rapid binding of arbitrary information

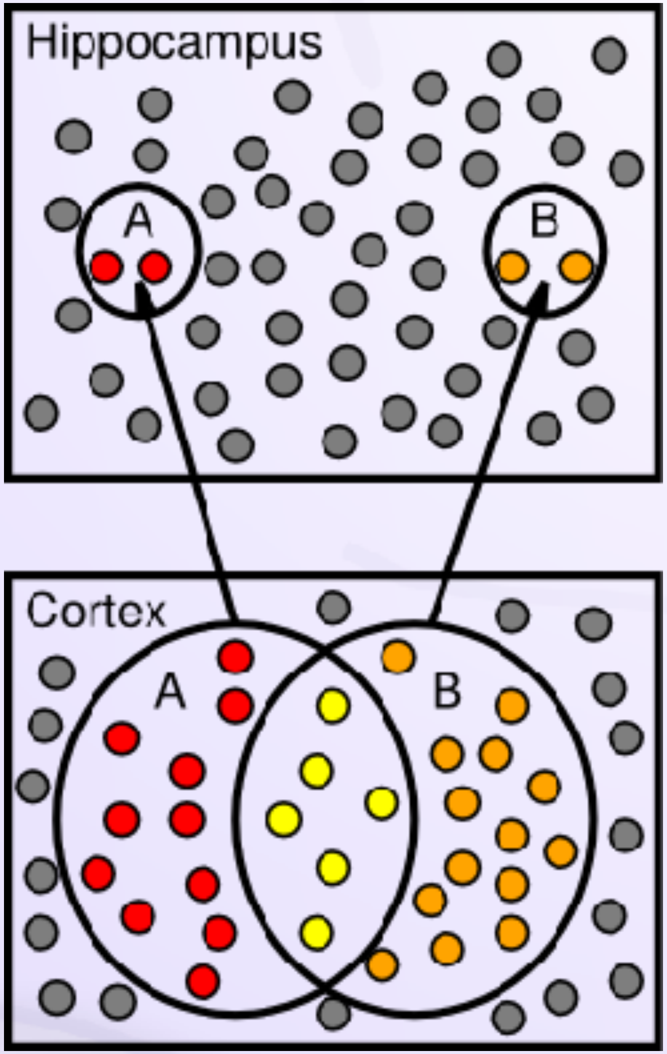

Combinatorial Instruction Following

(Huang, Hazy, Herd & O'Reilly, 2013)

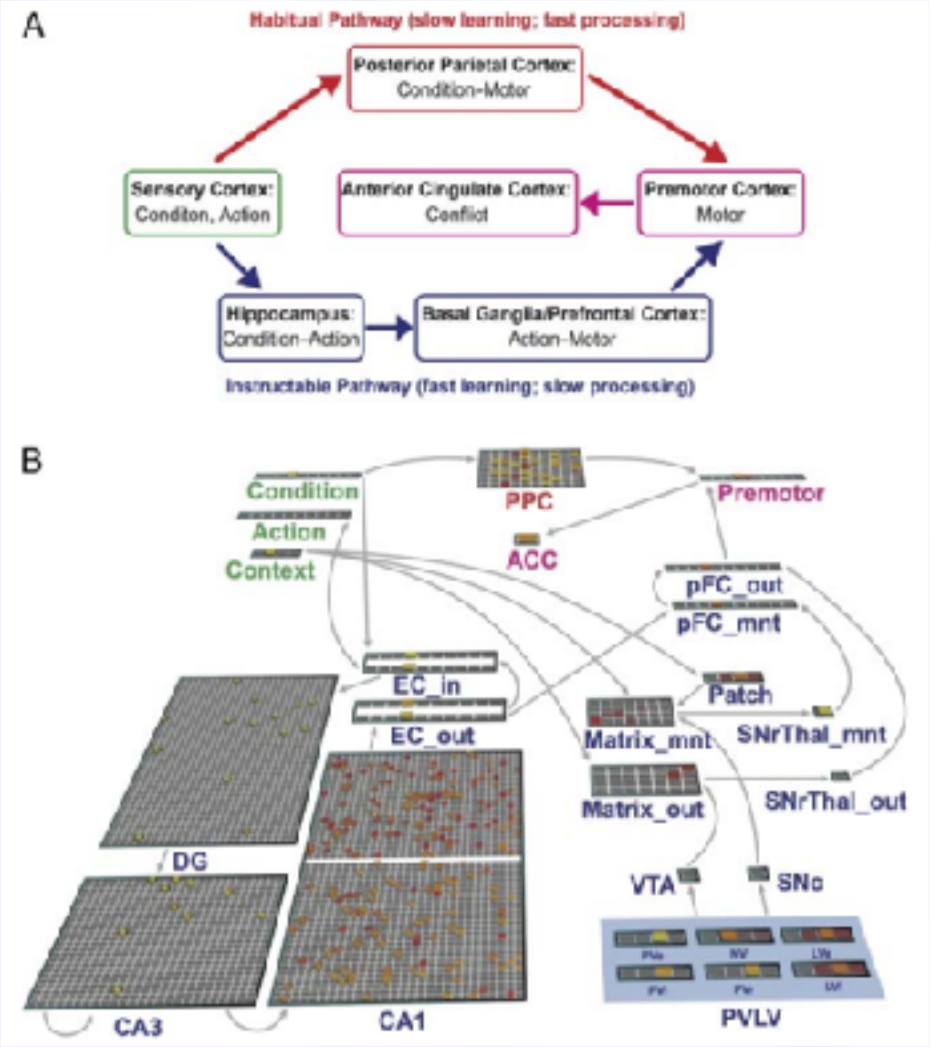

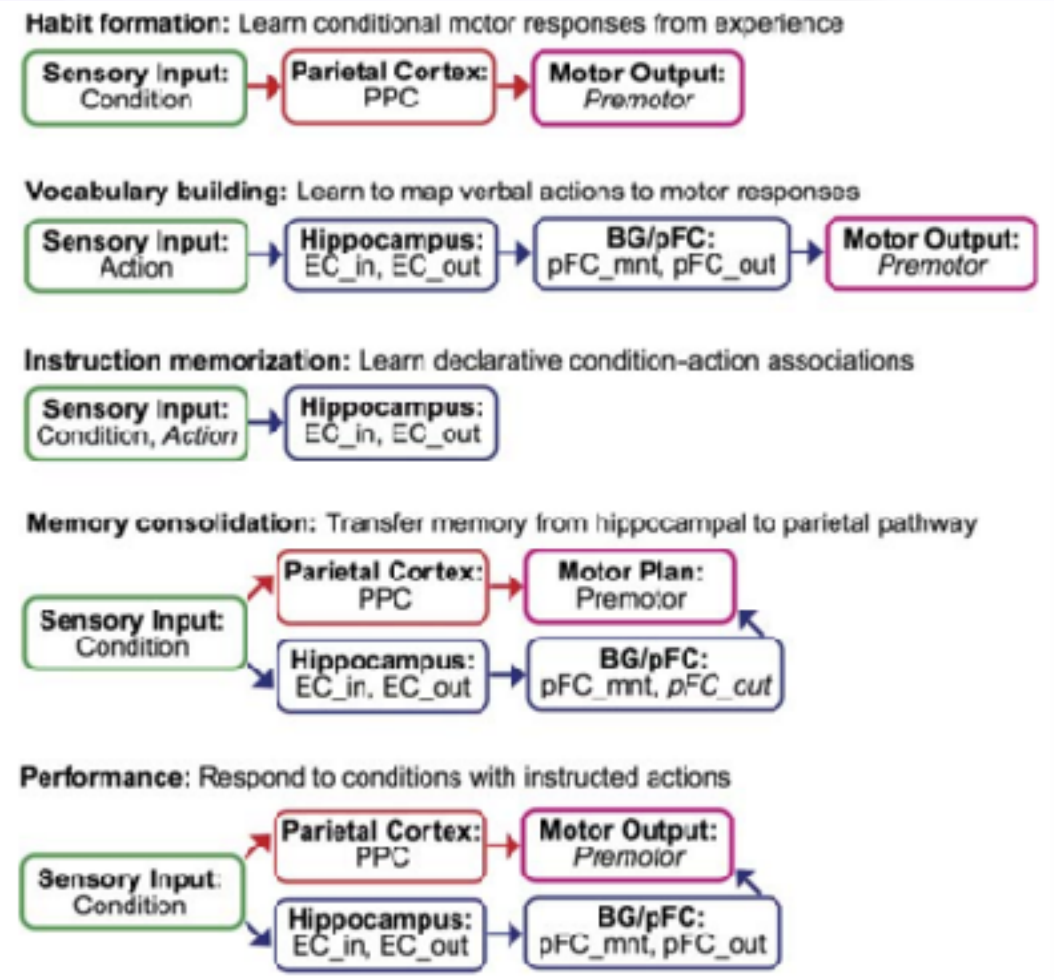

Serial vs. Parallel Summary

People can approximate a Turing machine by using LSTM-like BG/PFC gating & working memory (cf. "Neural Turing Machines", Graves..)

Essential for combinatorial flexibility: recombining existing subroutines to do novel tasks

While still leveraging huge advantages of parallel learning and processing..

The 3R's of Serial Processing

Reduce binding errors by serial processing: spatial attention spotlight – eg DCNN's that focus on object BBox – less adversarial image issues..

Reuse same neural tissue across many different situations – improved generalization (e.g., RN)

Recycle activity throughout network to coordinate all areas on one thing at a time: consciousness

Relational Network (Santoro et al)

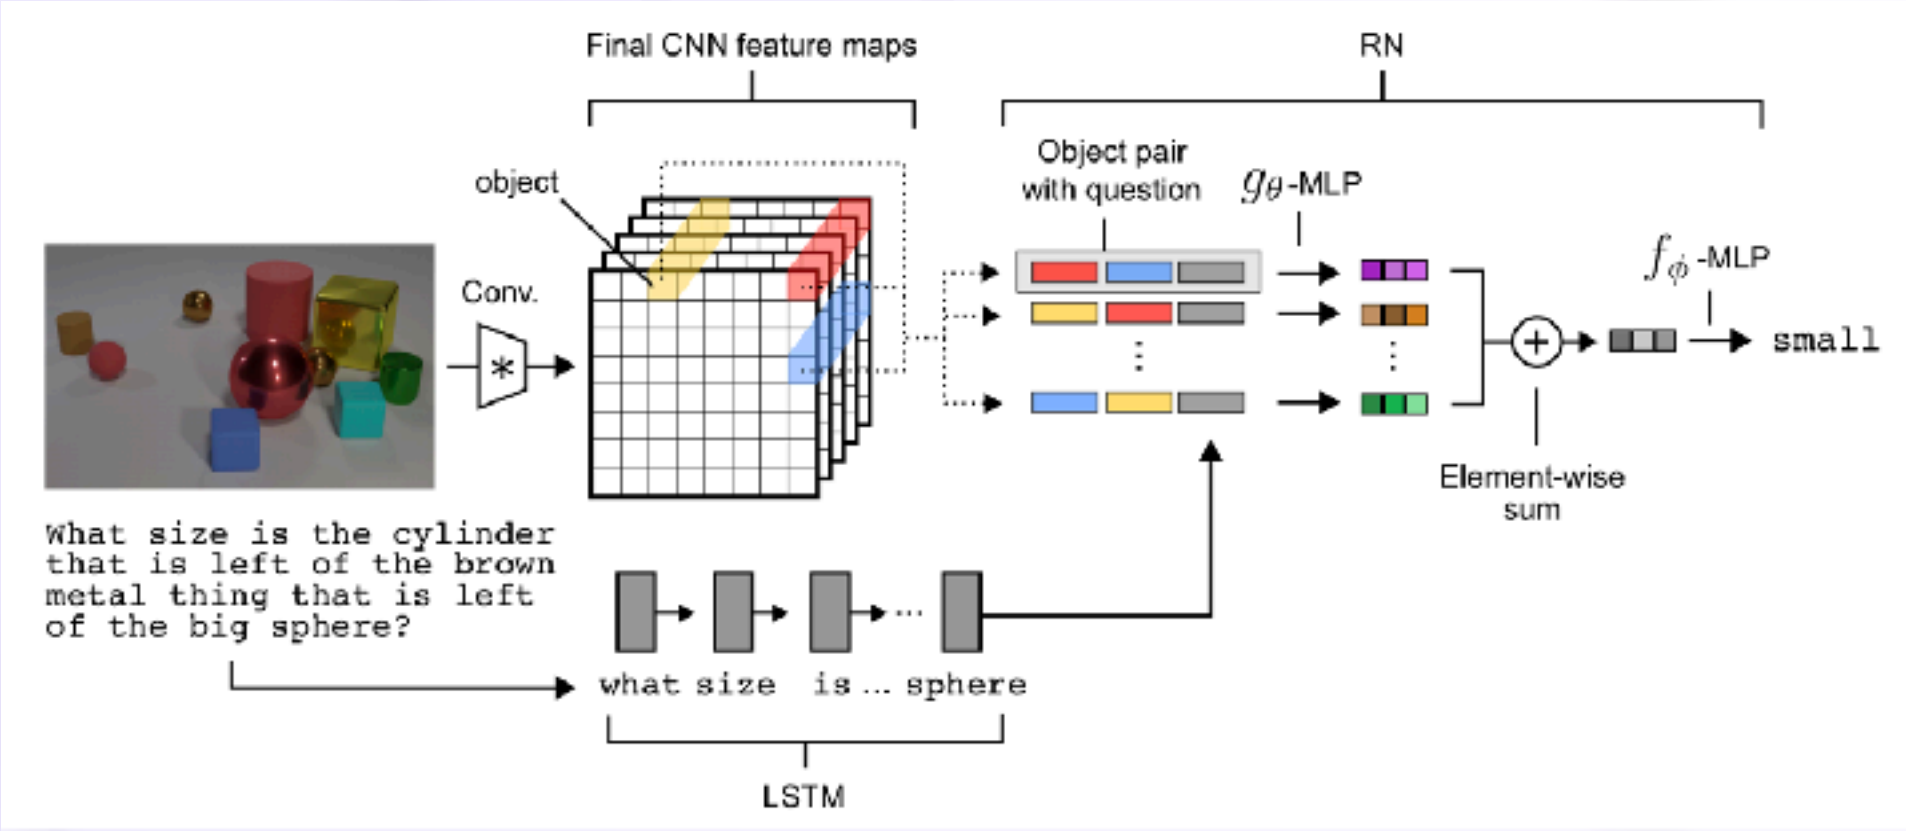

RN reuses same low-dimensional (pairwise) weights for all possible comparisons (using convolutional shared weights, in parallel).

But cannot generalize outside of training set (no combinatorial gen..)

Recurrent Processing -> Consciousness
(Lamme, 2006; cf. Bengio 2017)

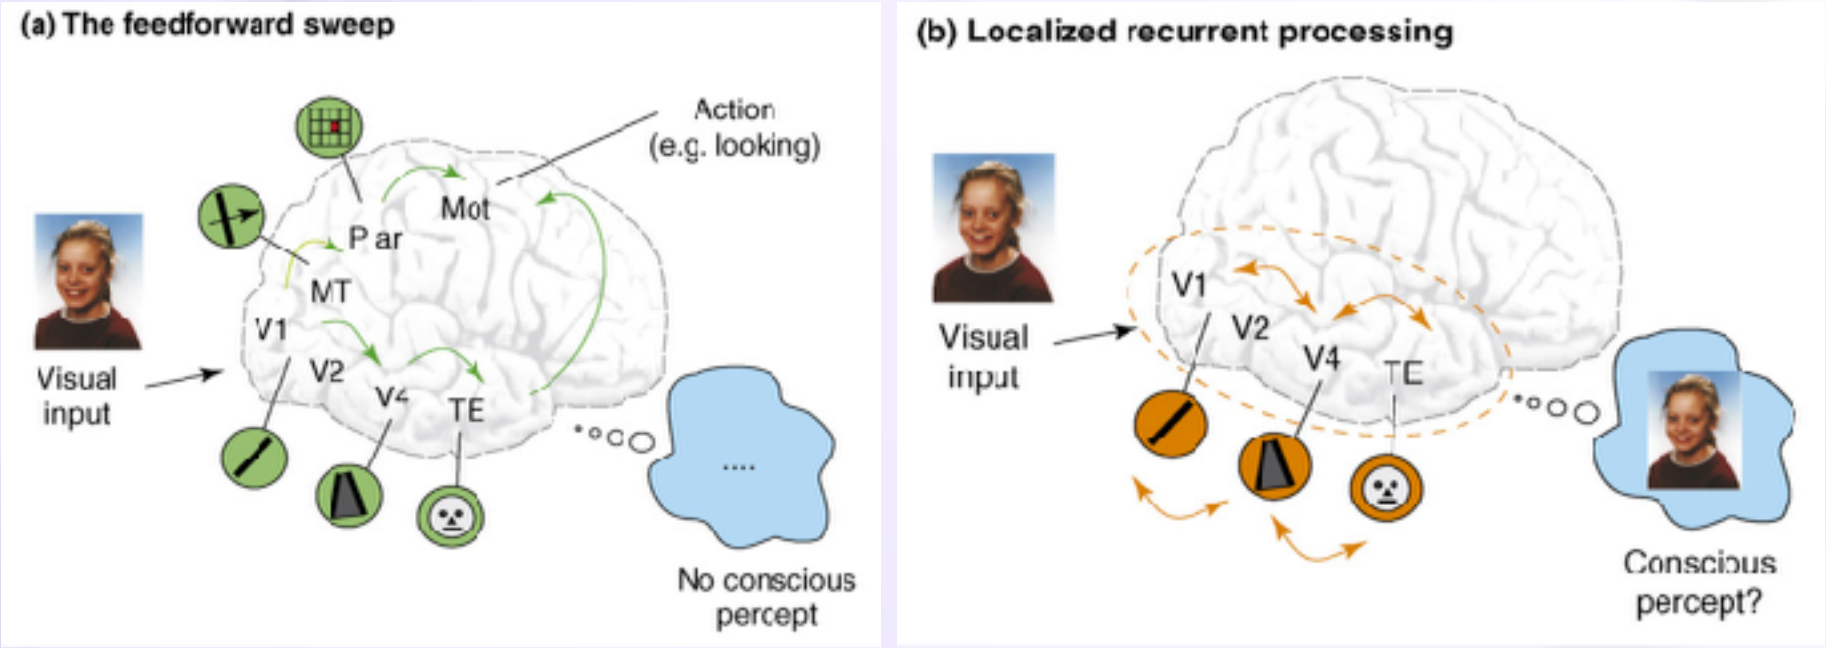

Consciousness is *Unitary*

Recurrence coordinates all areas on one thing (emerges via popular vote)

Consciousness is *Functional*

Helps organize, prioritize behavior – “focusing” key for difficult problems

Consciousness *Flows*

Temporal dynamics and information processing (multi-step cognition)

Recurrent Processing

Current DCNN's are almost exclusively feedforward

Cortex is massively recurrent (bidirectional excitatory connections)

Leabra model uses bidirectional excitatory connections to drive error-driven learning (O'Reilly, 1996) *and* constraint satisfaction (ala Hopfield)

Remaining Mysteries

How do we learn a full combinatorial vocabulary of productive subroutines?

What is the API? How do they communicate? Language, spatial attention spotlight..

Cognitive sequencing, planning: how do we write programs on the fly in our brains??

Thanks To

CCN Lab

- Tom Hazy
- Seth Herd
- Tren Huang
- Dave Jilk (eCortex)
- Nick Ketz
- Trent Kriete
- Kai Krueger
- Scott Mackie
- Brian Mingus
- Jessica Mollick
- Wolfgang Pauli
- John Rohrlich
- Sergio Verduzco-Flores
- Dean Wyatte

Funding

- **ONR** – Hawkins

Extras