Circuits and Mechanistic Neuroscience

Adam Marblestone

Submit questions here:

https://app.sli.do/event/92gy6nuo

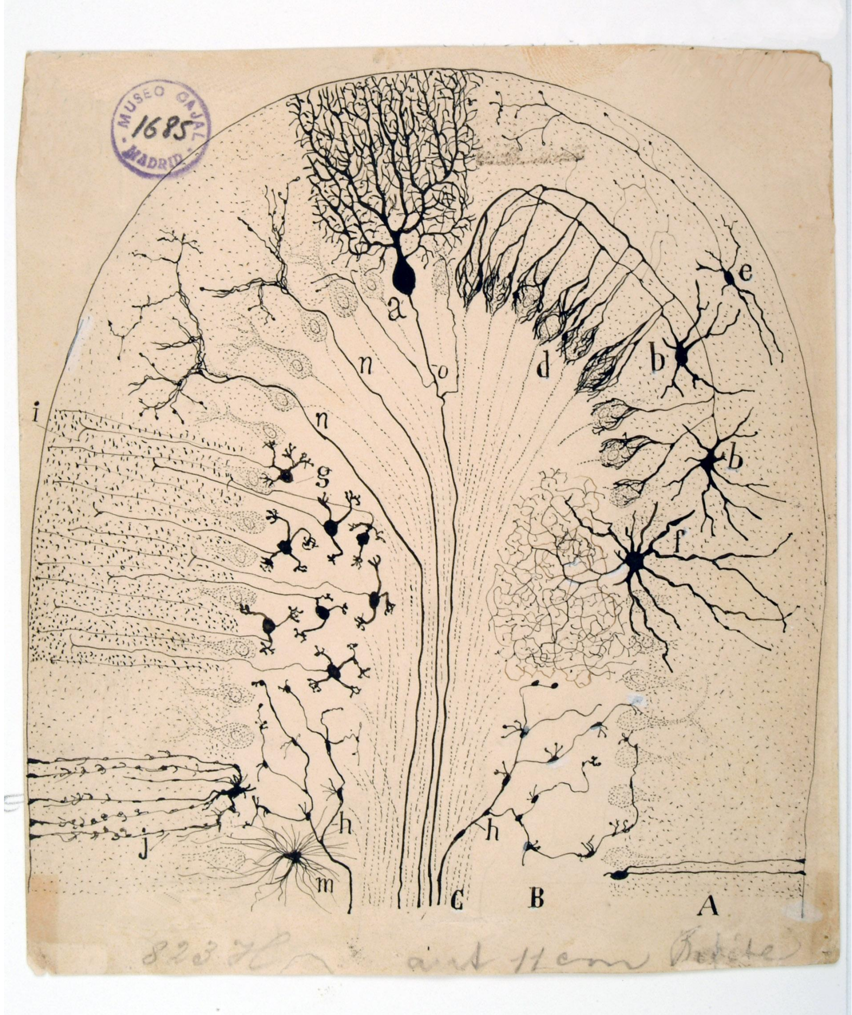
MUSEO CAJAL
1685
MADRID

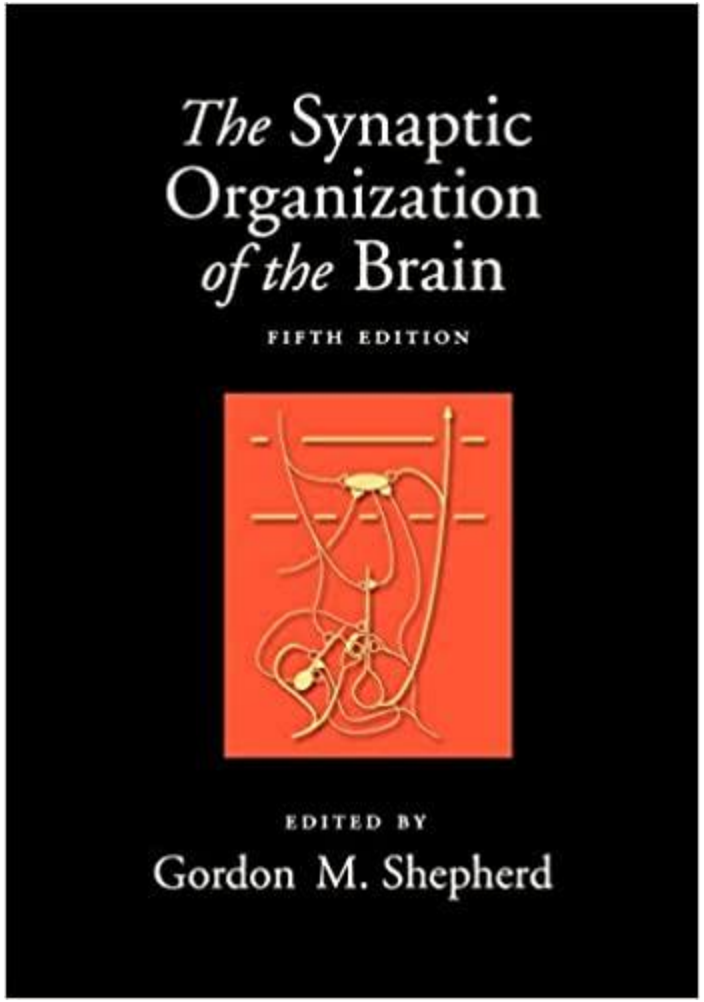
The Synaptic Organization of the Brain
FIFTH EDITION
EDITED BY
Gordon M. Shepherd

Do we know *in circuit-level detail* how brain implements at least one ML algorithm?

Do we know *in circuit-level detail* how brain implements at least one ML algorithm?

Yes: RL

Do we know *in circuit-level detail* how brain implements at least one ML algorithm?

Yes: **RL (or at least a simple instance)**

Circuit-level understand of RL

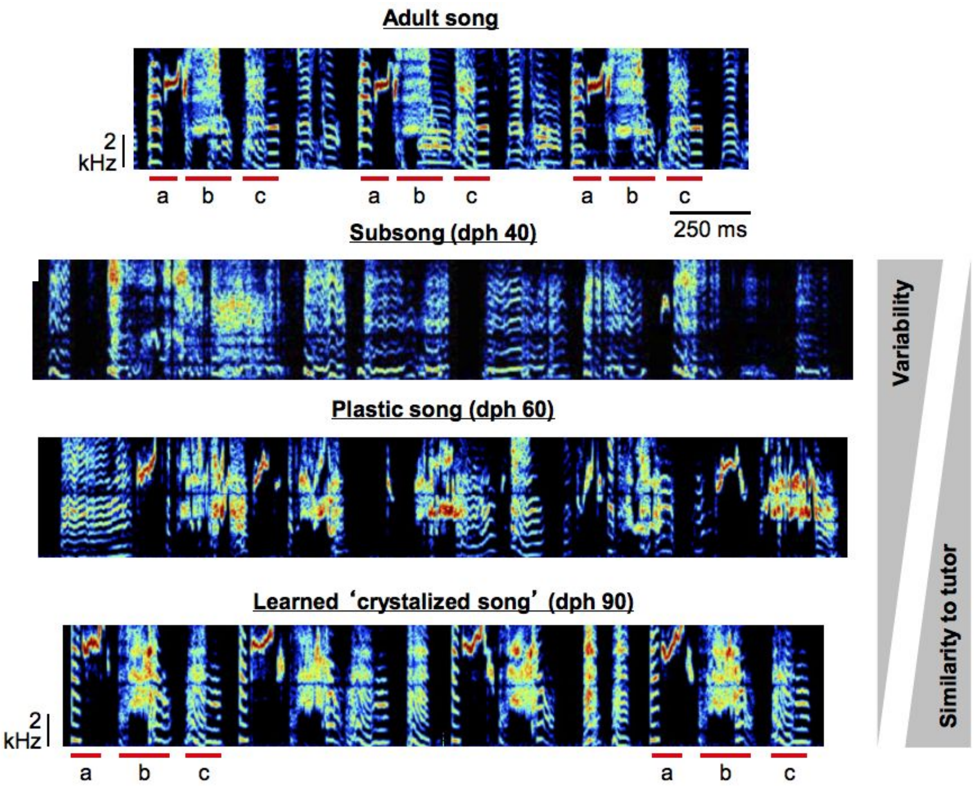

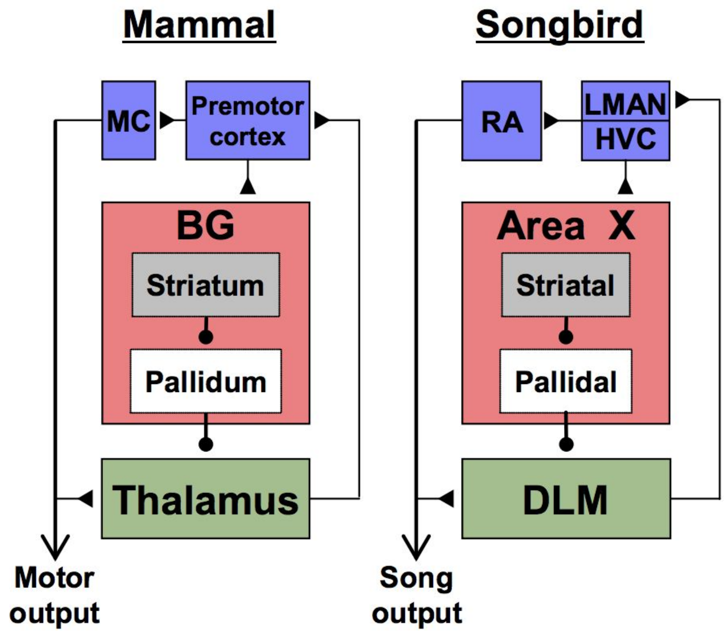

A hypothesis for basal ganglia-dependent reinforcement learning in the songbird

M. Fee, J. Goldberg · Published 2011 · Biology, Medicine · Neuroscience

Circuit-level understand of RL

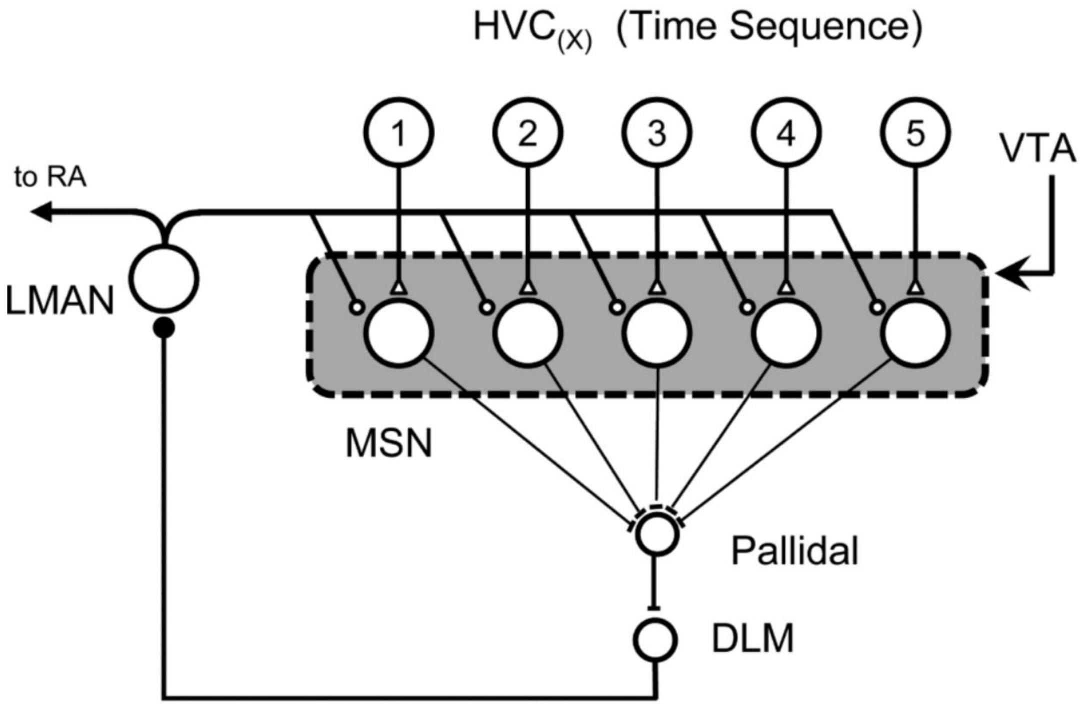

"Node perturbation" in each of many parallel "channels" through striatum

Ultimately "consolidated" into cortico-cortical connections

A hypothesis for basal ganglia-dependent reinforcement learning in the songbird

M. Fee, J. Goldberg · Published 2011 · Biology, Medicine · Neuroscience

Circuit-level understand of RL

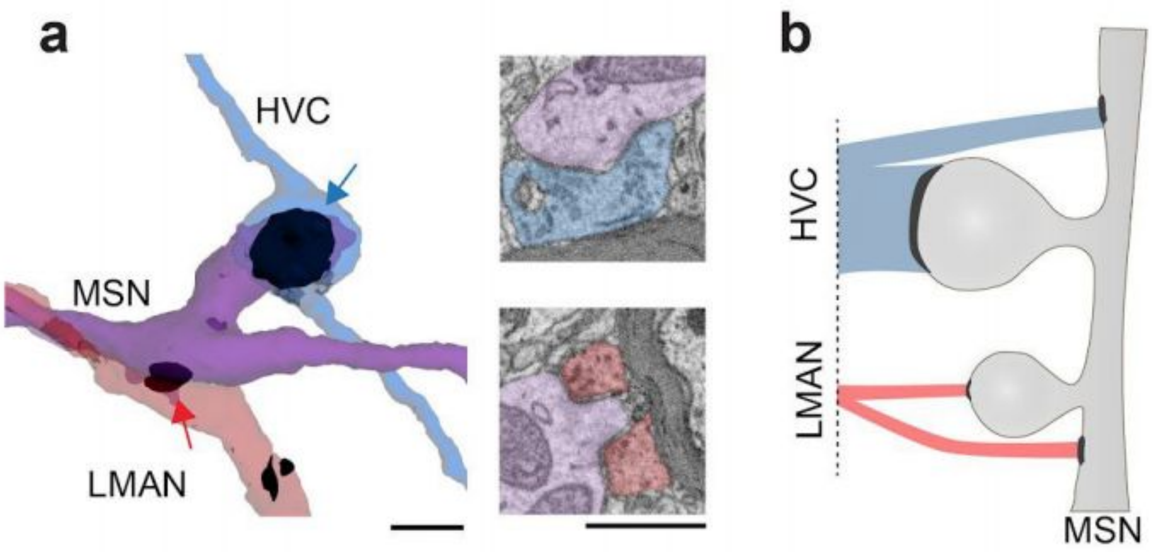

Deep synapse-level circuit reconstructions validate predictions of this theory

An anatomical substrate of credit assignment in reinforcement learning

J Kornfeld, M Januszewski, P Schubert, V Jain, W Denk, MS Fee

RL in the (songbird) brain does not use one monolithic reward

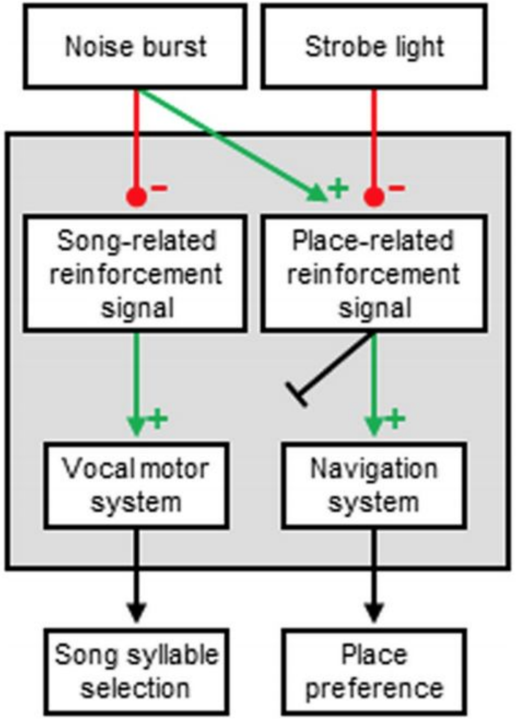

Spatially segregated circuits see distinct reinforcers

Reinforcers are internally-generated
(e.g., based on comparison to stored tutor song template)

Place preference and vocal learning rely on distinct reinforcers in songbirds

Don Murdoch, Ruidong Chen & Jesse H. Goldberg

What about in mammals?

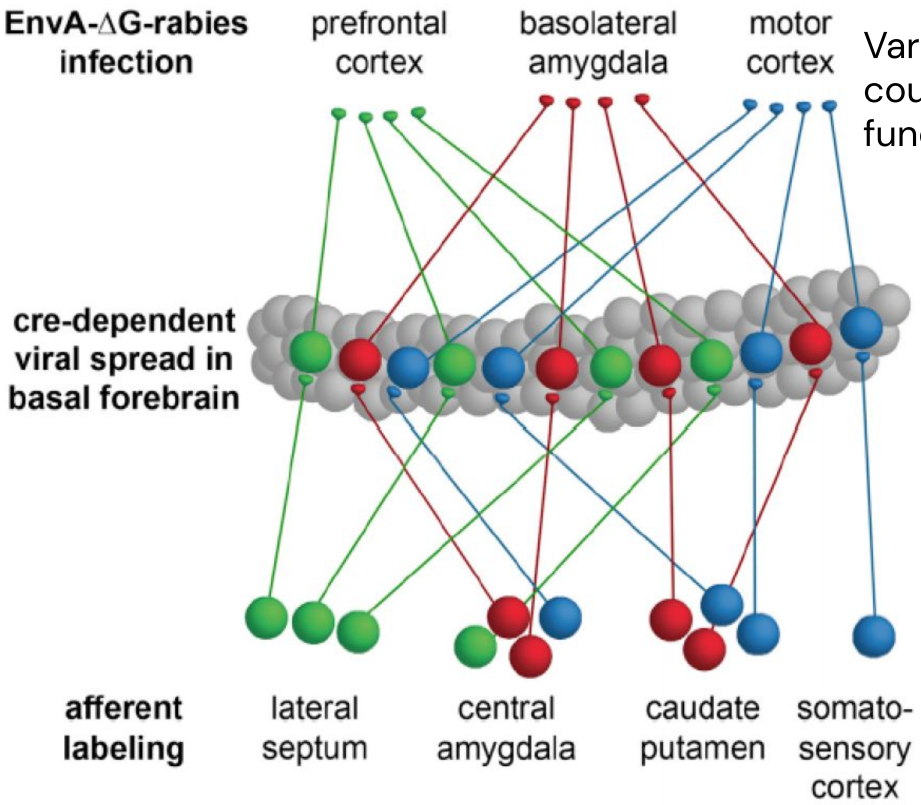

Various neural pathways such as the acetylcholine system could support something like multiple cost or reward functions applied to different sub-circuits.

A Cholinergic Mechanism for Reward Timing within Primary Visual Cortex

Alexander A. Chubykin[3], Emma B. Roach[3], Mark F. Bear, Marshall G. Hussain Shuler

[3] These authors contributed equally to this work

Nucleus basalis-enabled stimulus-specific plasticity in the visual cortex is mediated by astrocytes (i.e., glia not neurons)

Naiyan Chen[a,1], Hiroki Sugihara[a,1], Jitendra Sharma[a,b], Gertrudis Perea[a], Jeremy Petravicz[a], Chuong Le[a], and Mriganka Sur[a,2]

The Input-Output Relationship of the Cholinergic Basal Forebrain

Matthew R. Gielow and Laszlo Zaborszky

What about in mammals?

Complex cortical networks feeding into something like a "value function"?

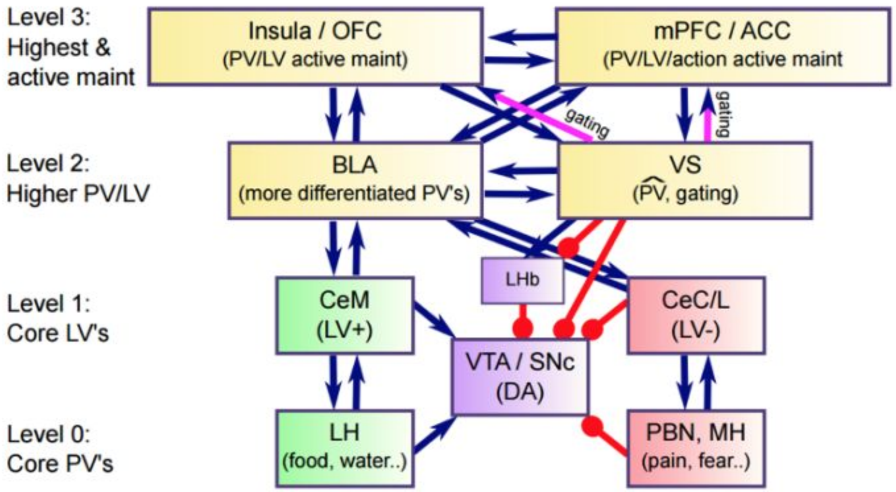

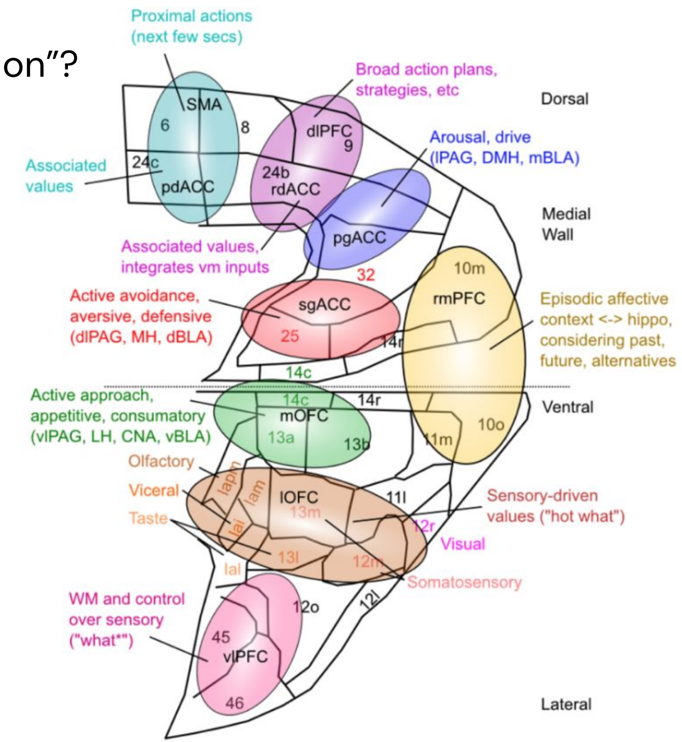

Randall O'Reilly

Do we know *in circuit-level detail* how brain implements a multi-layer credit assignment mechanism of similar power to backprop?

Do we know *in circuit-level detail* how brain implements a multi-layer credit assignment mechanism of similar power to backprop?

Trick question: not sure it does at all

Do we know *in circuit-level detail* how brain implements a multi-layer credit assignment mechanism of similar power to backprop?

Trick question: but there are ideas…

Idea: *Use multiple dendritic compartments to store both "activations" and "errors"*

soma voltage ~ activation

dendritic voltage ~ error derivative

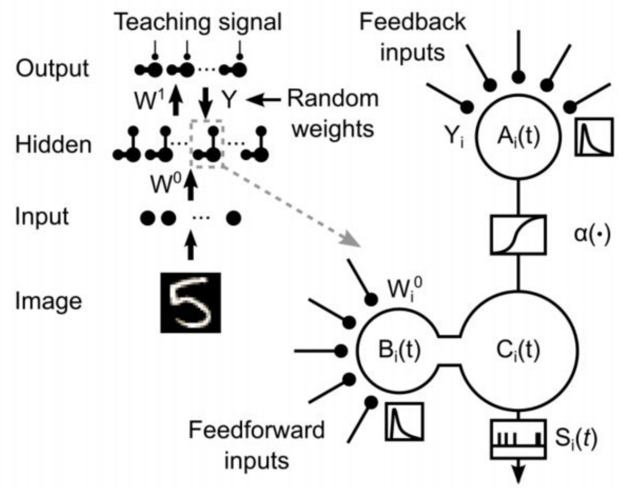

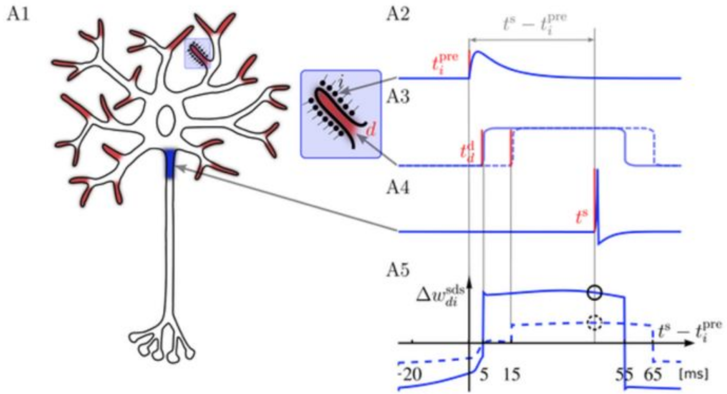

Supervised and Unsupervised Learning with Two Sites of Synaptic Integration

KONRAD P. KÖRDING AND PETER KÖNIG

Institute of Neuroinformatics, ETH/UNI Zürich, Winterthurerstr. 190, 8057 Zürich, Switzerland

Deep learning with segregated dendrites

Jordan Guergiuev[1,2], Timothy P. Lillicrap[4], Blake A. Richards[1,2,3,*]

Somato-dendritic Synaptic Plasticity and Error-backpropagation in Active Dendrites

Mathieu Schiess, Robert Urbanczik, Walter Senn

Does the need to implement a powerful learning algorithm underly the cellular complexity of cortical micro-circuits?

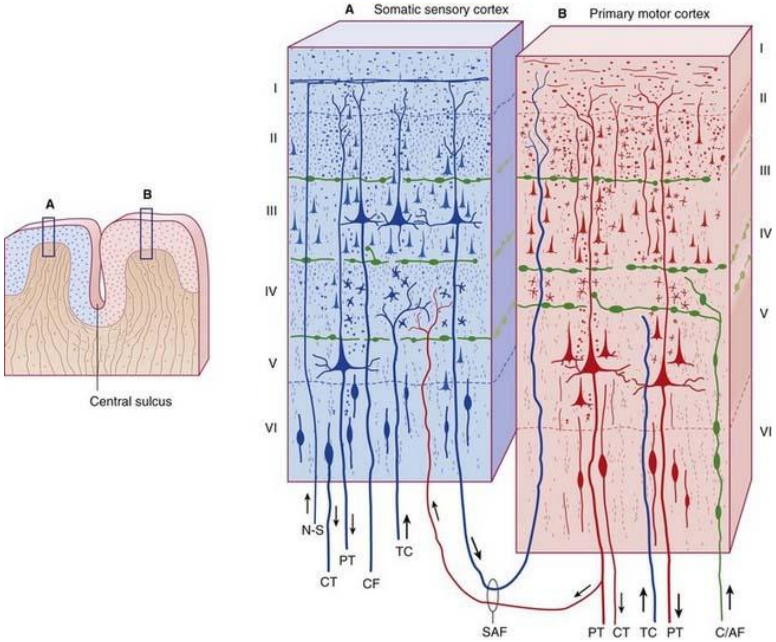

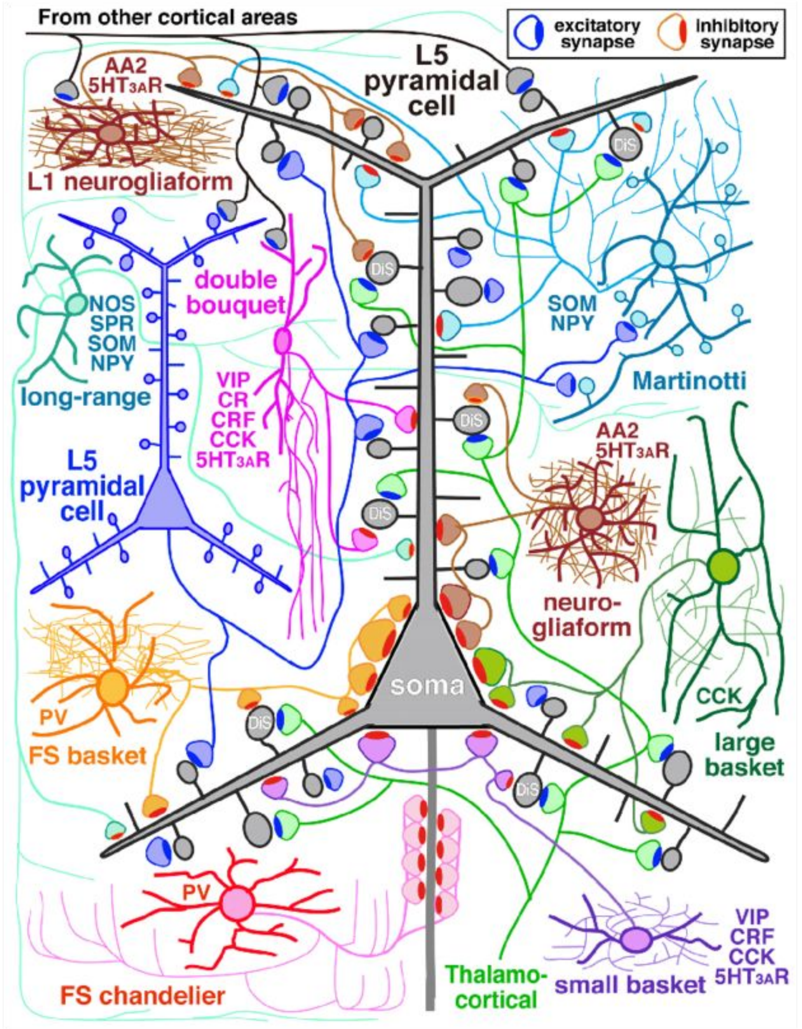

Does the need to implement a powerful learning algorithm underly the cellular complexity of cortical micro-circuits?

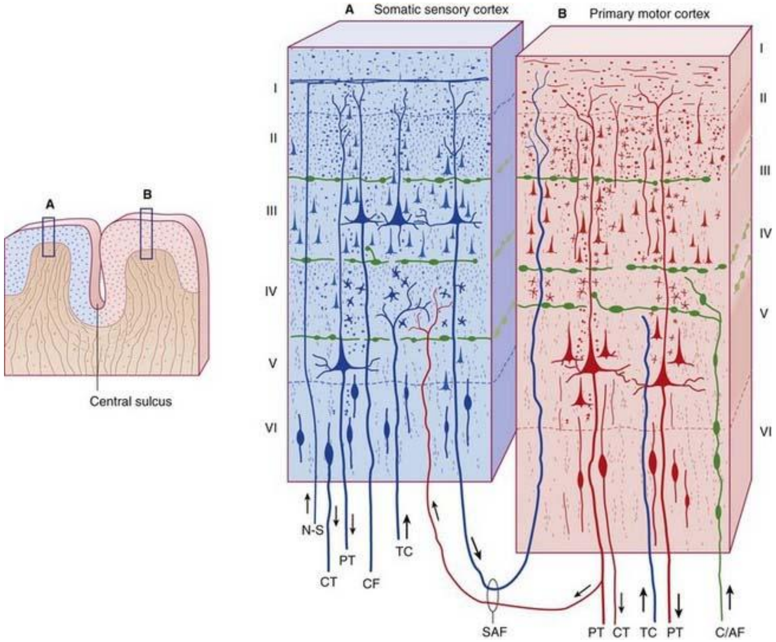

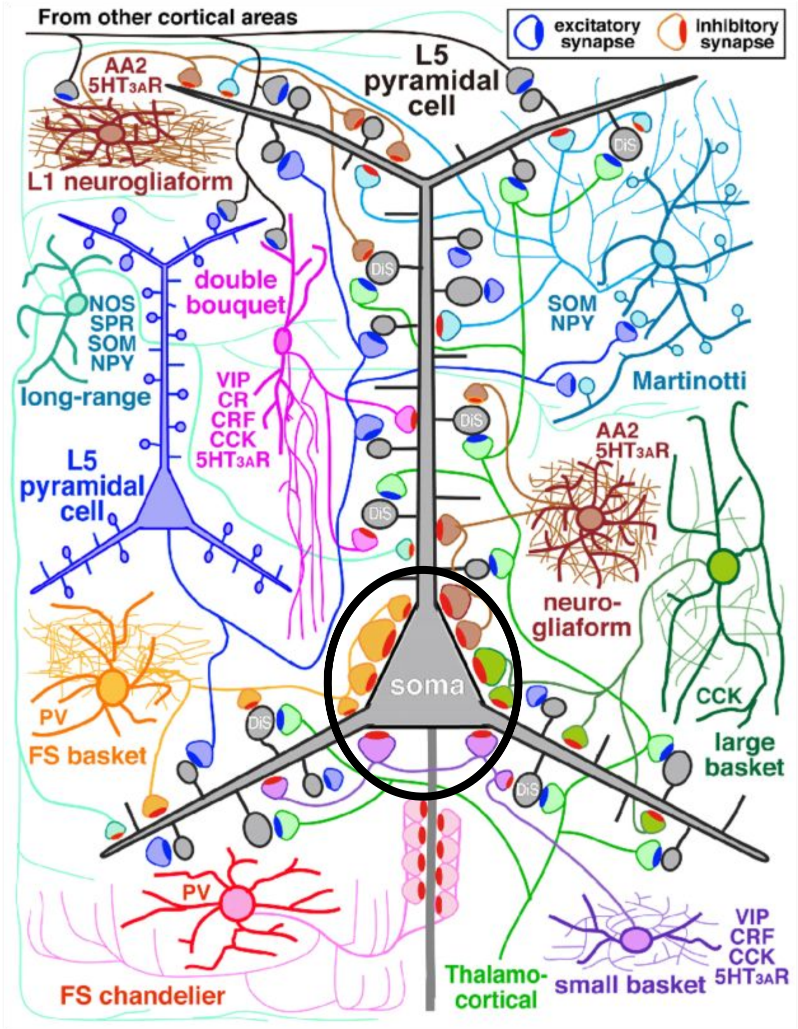

Table 1. Comparison of Models

		Temporal-error model		Explicit-error model	
		Contrastive learning	Continuous update	Predictive coding	Dendritic error
Properties[a]	Control signal	Required	Required	Not required	Not required
	Connectivity	Unconstrained	Unconstrained	Constrained	Constrained
	Propagation time	L-1	L-1	2L-1	L-1
	Pre-training	Not required	Not required	Not required	Required
Error encoded in		Difference in activity between separate phases	Rate of change of activity	Activity of specialised neurons	Apical dendrites of pyramidal neurons
Data accounted for		Neural responses and behaviour in a variety of tasks	Typical spike-time-dependent plasticity	Increased neural activity to unpredicted stimuli	Properties of pyramidal neurons
MNIST performance[b]		~2–3	–	~1.7	~1.96

Review

Theories of Error Back-Propagation in the Brain

James C.R. Whittington[1,2] and Rafal Bogacz[1,*]

Can such ideas be validated?

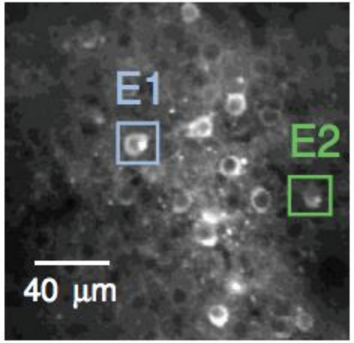

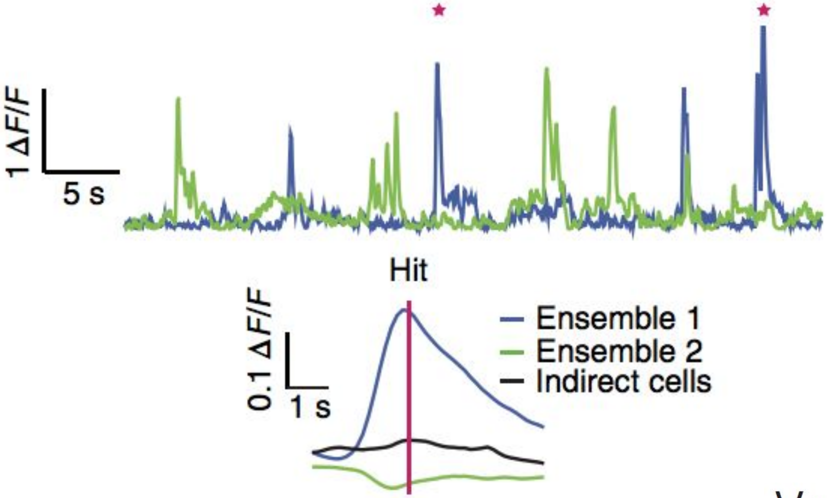

Clever experimental design with “brain computer interface” experiment in animals might allow us to tease out credit assignment mechanisms from *movies of neurons firing (calcium imaging)*

Rewards are made contingent on the firing of particular neurons on the circuit

Volitional modulation of optically recorded calcium signals during neuroprosthetic learning

Kelly B Clancy[1,7], Aaron C Koralek[2,7], Rui M Costa[3], Daniel E Feldman[2,4] & Jose M Carmena[2,5,6]

Other interpretations of cortical circuit

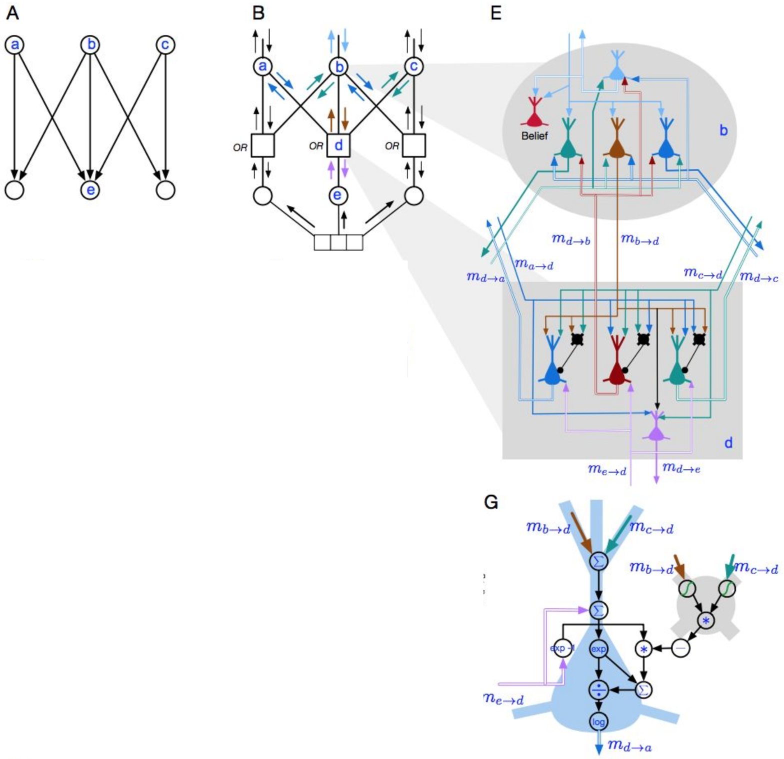

In other interpretations, cortical circuit architecture is determined more by the need for complex inference (e.g., as in message passing algorithms for graphical models of particular structures)

Rather than by the need for weight transport for deep learning-like multilayer credit assignment

A detailed mathematical theory of thalamic and cortical microcircuits based on inference in a generative vision model

Dileep George, Miguel Lázaro-Gredilla, Wolfgang Lehrach, Antoine Dedieu, Guangyao Zhou

Other interpretations of cortical circuit

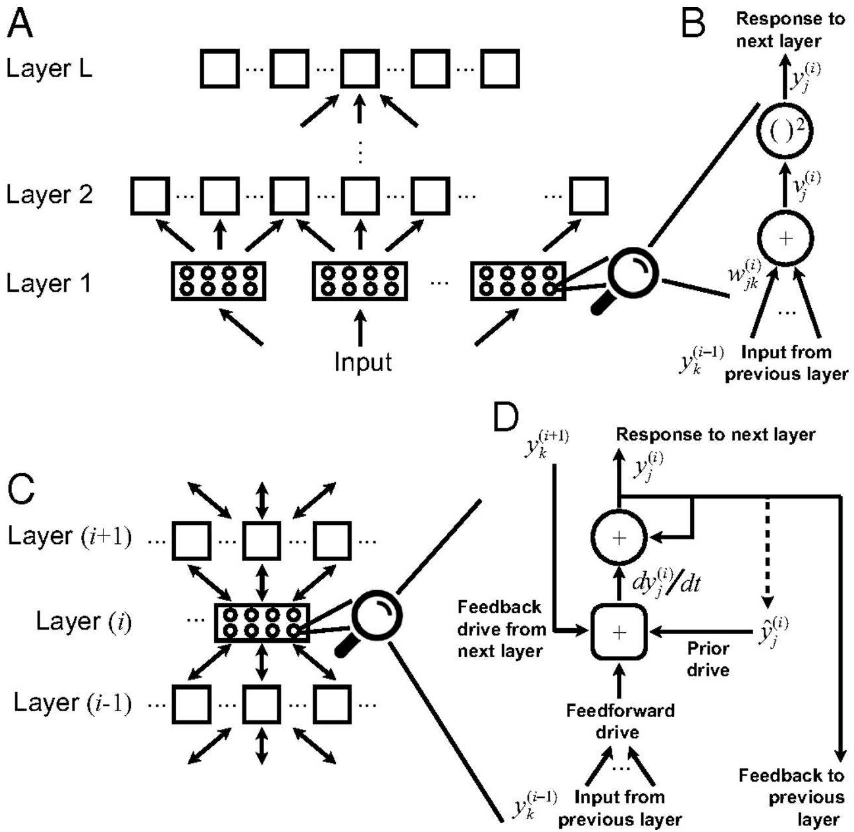

In yet other interpretations, cortical circuits reflect energy-based models that incorporate both top-down prior information and bottom-up sensory information into the energy function

Such models are meant to explain perceptual phenomena in a "perception as inference" framework

Theory of cortical function

David J. Heeger

Content-addressable memory in brain

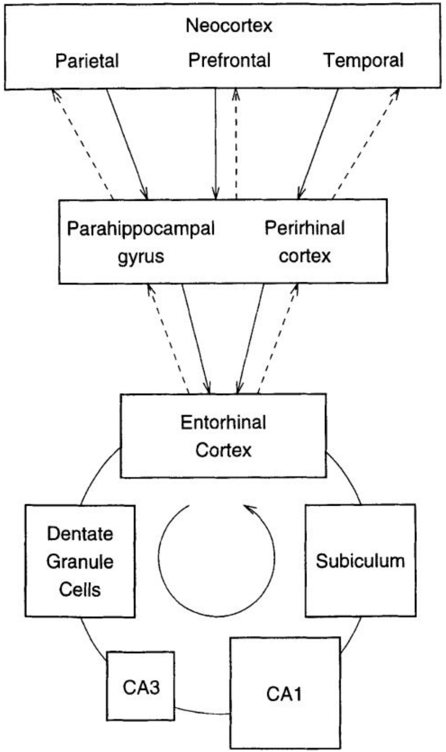

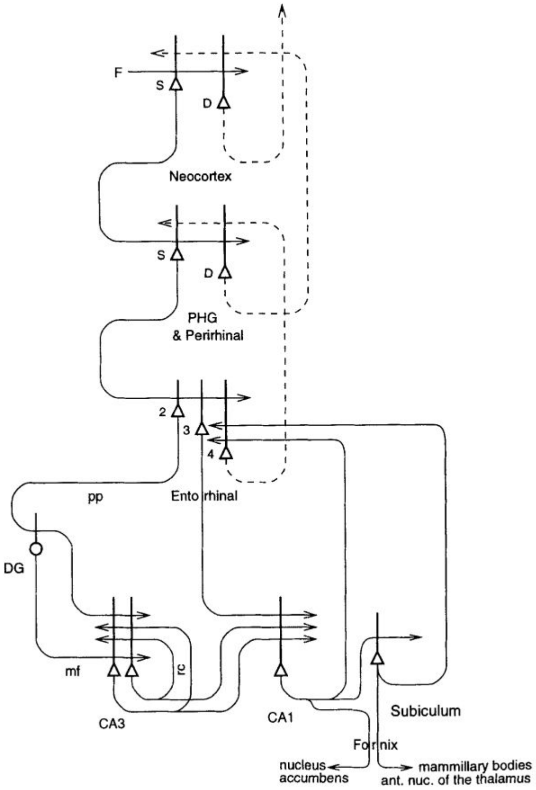

The hippocampus looks like a content-addressable memory system taking input from many higher cortical areas

Pattern separation (dentate gyrus), *pattern completion* (CA3), and *reconstruction* (CA1)

Not just for spatial navigation…

A Theory of Hippocampal Function in Memory

Edmund T. Rolls

Content-addressable memory in brain

Abstract

Neuronal circuits produce self-sustaining sequences of activity patterns, but the precise mechanisms remain unknown. Here we provide evidence for autoassociative dynamics in sequence generation. During sharp-wave ripple (SWR) events, hippocampal neurons express sequenced reactivations, which we show are composed of discrete attractors. Each attractor corresponds to a single location, the representation of which sharpens over the course of several milliseconds, as the reactivation focuses at that location. Subsequently, the reactivation transitions rapidly to a spatially discontiguous location. This alternation between sharpening and transition occurs repeatedly within individual SWRs and is locked to the slow-gamma (25 to 50 hertz) rhythm. These findings support theoretical notions of neural network function and reveal a fundamental discretization in the retrieval of memory in the hippocampus, together with a function for gamma oscillations in the control of attractor dynamics.

To a degree one can even see "pattern completion" happening in activity recordings...

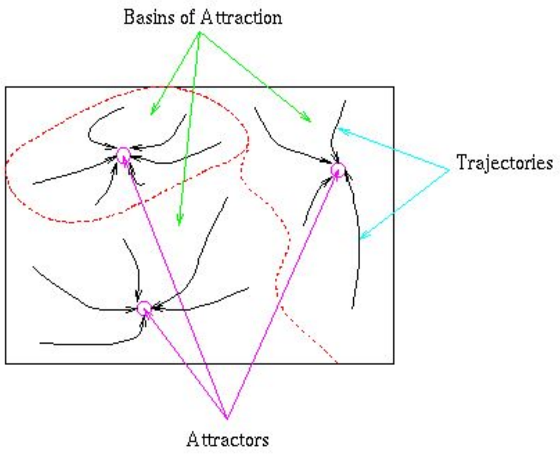

Autoassociative dynamics in the generation of sequences of hippocampal place cells

Brad E. Pfeiffer*, David J. Foster†

http://staff.itee.uq.edu.au/janetw/cmc/chapters/Hopfield/

Content-addressable memory in brain

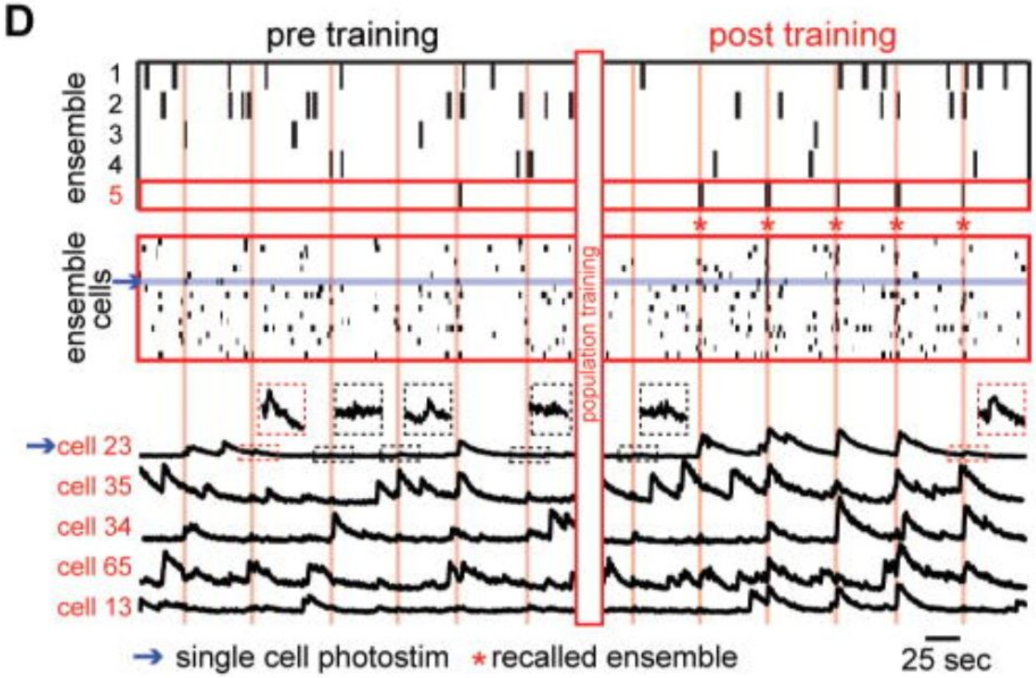

To a degree one can even see "pattern completion" happening in activity recordings...

Imprinting and Recalling Cortical Ensembles

Luis Carrillo-Reid, Weijian Yang, Yuki Bando, Darcy S. Peterka, and Rafael Yuste

Using such mechanisms for symbolic processing?

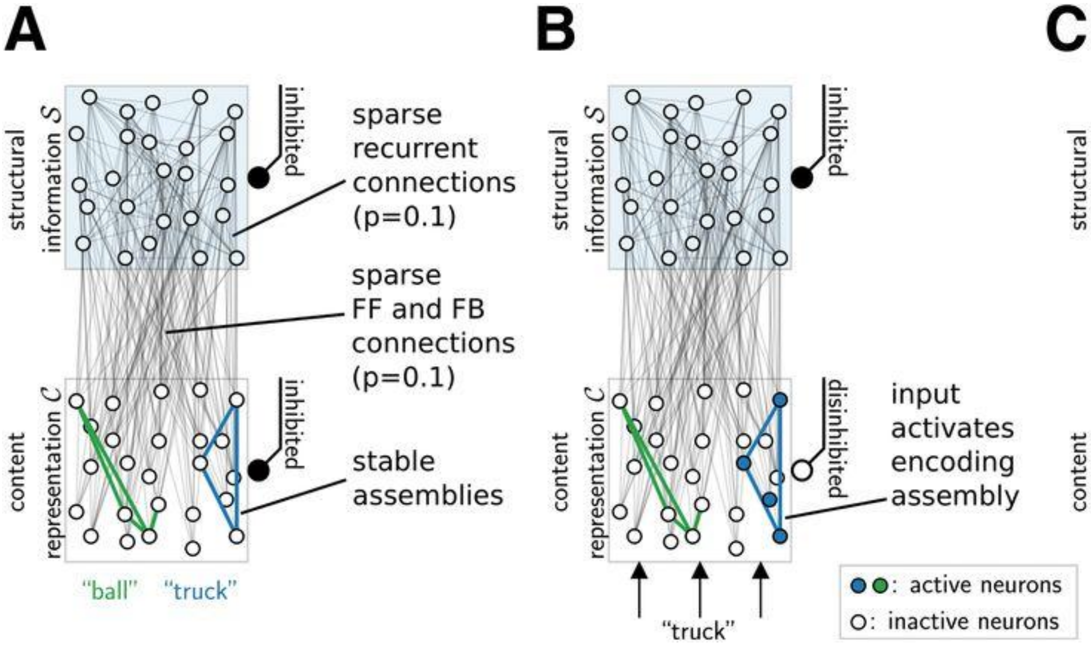

A Model for Structured Information Representation in Neural Networks of the Brain

Michael G. Müller, Christos H. Papadimitriou, Wolfgang Maass, and Robert Legenstein

Looking to the future: better circuit level data

IARPA MICRONS (2015-2020) 1mm^3 scale circuits (mouse brain ~500 mm^3, human ~500,000 mm^3)

MICrONS will combine neuroscience and data science to advance machine learning by uncovering how the cortex performs computations at the mesoscale

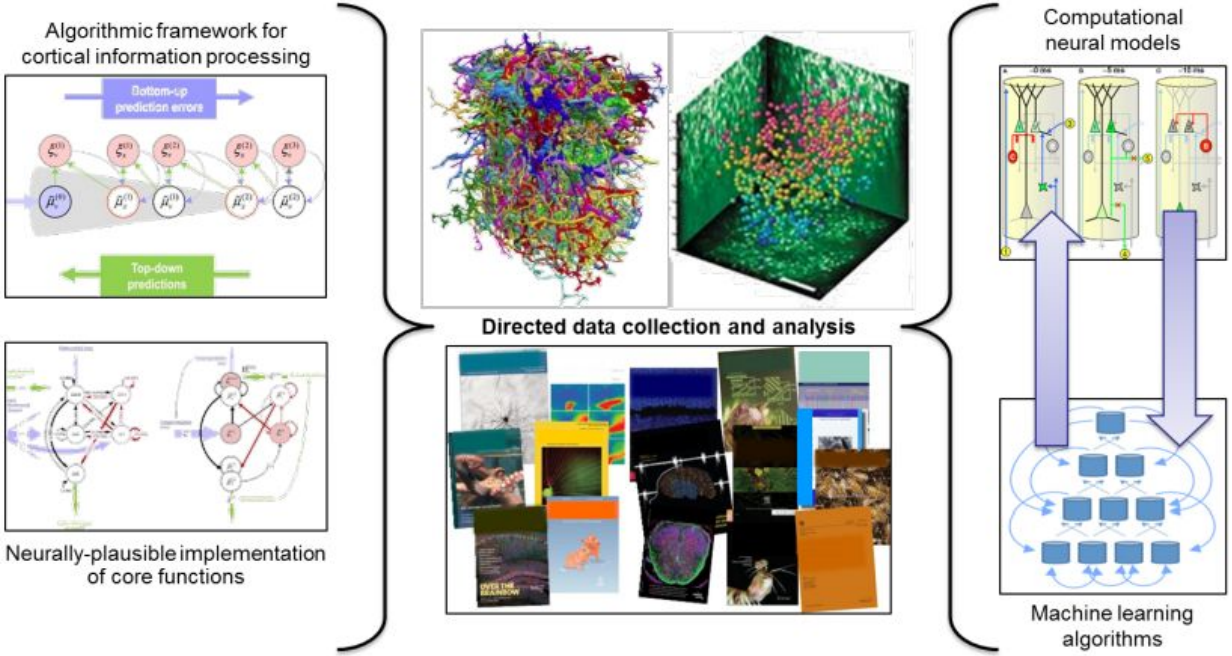

Looking to the future: better circuit level data

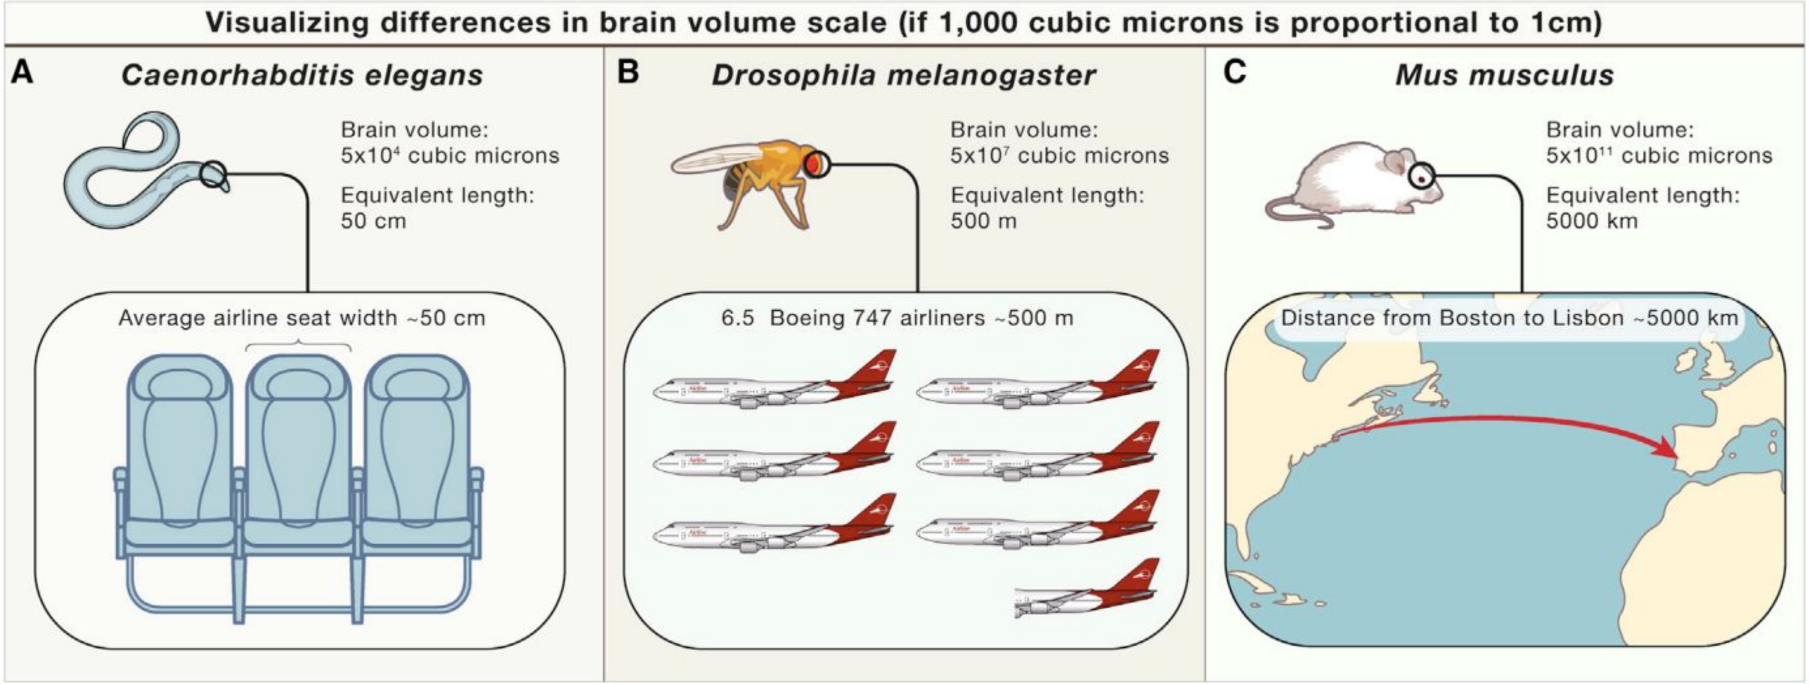

1980s
Connectome
~2020
Full activity map

~2020
Connectome

The Mind of a Mouse ~2025-2030?

Larry F. Abbott [1], Davi D. Bock [2], Edward M. Callaway [3], Winfried Denk [4, 25], Catherine Dulac [5], Adrienne L. Fairhall [6], Ila Fiete [7], Kristen M. Harris [8], Moritz Helmstaedter [9], Viren Jain [10, 25], Narayanan Kasthuri [11], Yann LeCun [12], Jeff W. Lichtman [13, 25], Peter B. Littlewood [14], Liqun Luo [15], John H.R. Maunsell [16], R. Clay Reid [17, 25], Bruce R. Rosen [18] … David C. Van Essen [24]

Take-home messages: circuits part

Some algorithms being used in ML (and we've only covered a subset here) have biologically plausible hypothesized mappings onto neural circuits, but much remains to be understood and validated

For some ML concepts, like reinforcement learning or pattern completion, there are detailed mapping to particular types and circuits of cells in the brains of birds or mammals

There are many basic things we still don't know, like how uniform the cortex is, whether some approximation of backprop is used by the brain, how the brain represents probabilities or does something like Bayesian inference, whether it uses symbolic mechanisms, and so on

Technologies have scaled activity recording from a handful of cells to thousands and can start to be more useful to address these issues, while connectomics has scaled to the level of a fly brain and there is a major push to scale it to a mammal brain

Still early days, but we see hints that circuit neuroscience can inform AI more deeply in next 10-15 years…