

Comparing the progress of Neuroscience and AI

Ni Lao

4.21.2026

Plan




- ***Introduction***
 - Neuroscience & algorithms
 - Retinal Waves & SOM
 - Motor learning & Minimax/ELM/PRA
- ***Spatial Representation & Reasoning***
 - Location cell and grid cells
 - agentic vision
- ***Memory & Reinforcement Learning***
 - World models
 - Dream & experience replays
- ***Knowledge Representation***
 - the ideal distribution of knowledge?

Neuroscience & Algorithm

- Neuroscience development
 - Neuroprosthetics -- understanding and repairing disabilities
 - Neurodegenerative Diseases -- Alzheimer's, Parkinson's, and Huntington's diseases
- Neuroscience provides mechanisms that serve certain purposes
 - e.g., retinotopic mapping
 - e.g., motor learning
- Neuroscience has a long history of inspiring the development of AI (the other being physics)

Let's compare notes

- Where is the current state of AI in terms of neuroscience?
we can exam chapter by chapter



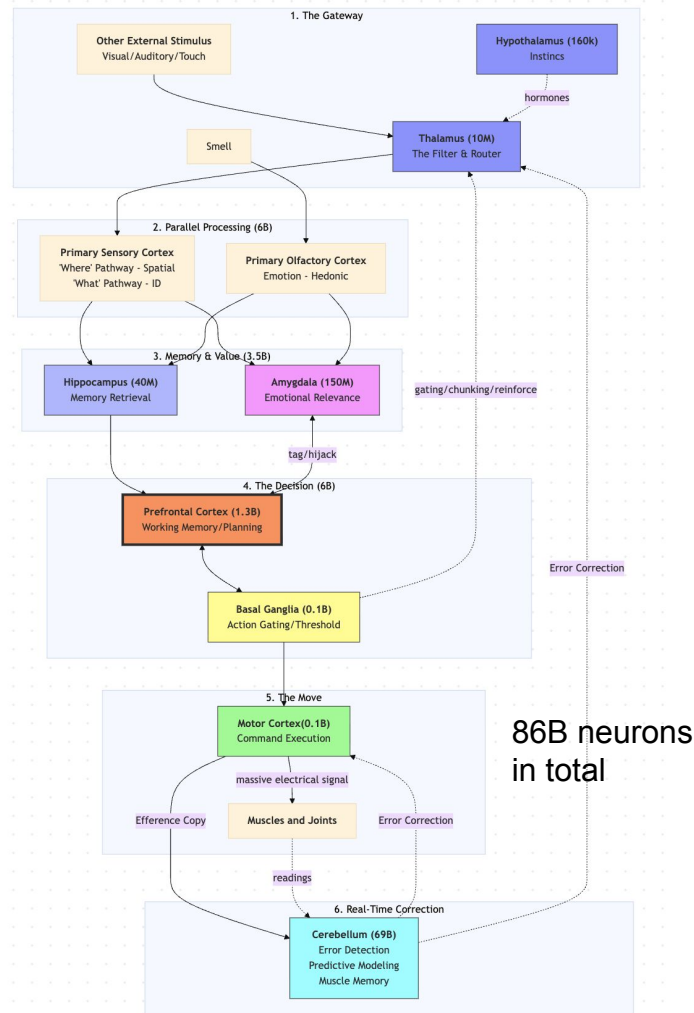
EDITORS

Dale Purves • George J. Augustine
David Fitzpatrick • William C. Hall • Anthony-Samuel LaMantia
Richard D. Mooney • Michael L. Platt • Leonard E. White

UNIT I	Neural Signaling	31
2	Electrical Signals of Nerve Cells	33
3	Voltage-Dependent Membrane Permeability	49
4	Ion Channels and Transporters	65
5	Synaptic Transmission	85
6	Neurotransmitters and Their Receptors	113
7	Molecular Signaling within Neurons	145
8	Synaptic Plasticity	169
UNIT II	Sensation and Sensory Processing	191
9	The Somatosensory System: Touch and Proprioception	193
10	Pain	213
11	Vision: The Eye	233
12	Central Visual Pathways	261
13	The Auditory System	281
14	The Vestibular System	305
15	The Chemical Senses	323
UNIT III	Movement and Its Central Control	355
16	Lower Motor Neuron Circuits and Motor Control and Spinal Cord	357
17	Upper Motor Neuron Control of the Brainstem	381
18	Modulation of Movement by the Basal Ganglia	407
19	Modulation of Movement by the Cerebellum	427
20	Eye Movements and Sensorimotor Integration	447
21	The Visceral Motor System	465
UNIT IV	The Changing Brain	489
22	Early Brain Development	491
23	Construction of Neural Circuits	521
24	Circuit Development: Intrinsic Factors and Sex Differences	551
25	Experience-Dependent Plasticity in the Developing Brain	571
26	Repair and Regeneration in the Nervous System	595
UNIT V	Complex Brain Functions and Cognitive Neuroscience	625
27	Cognitive Functions and the Organization of the Cerebral Cortex	627
28	Cortical States	643
29	Attention	667
30	Memory	681
31	Emotion	703
32	Thinking, Planning, and Deciding	725
33	Speech and Language	745
34	Development and Evolution of Cognitive Functions	767

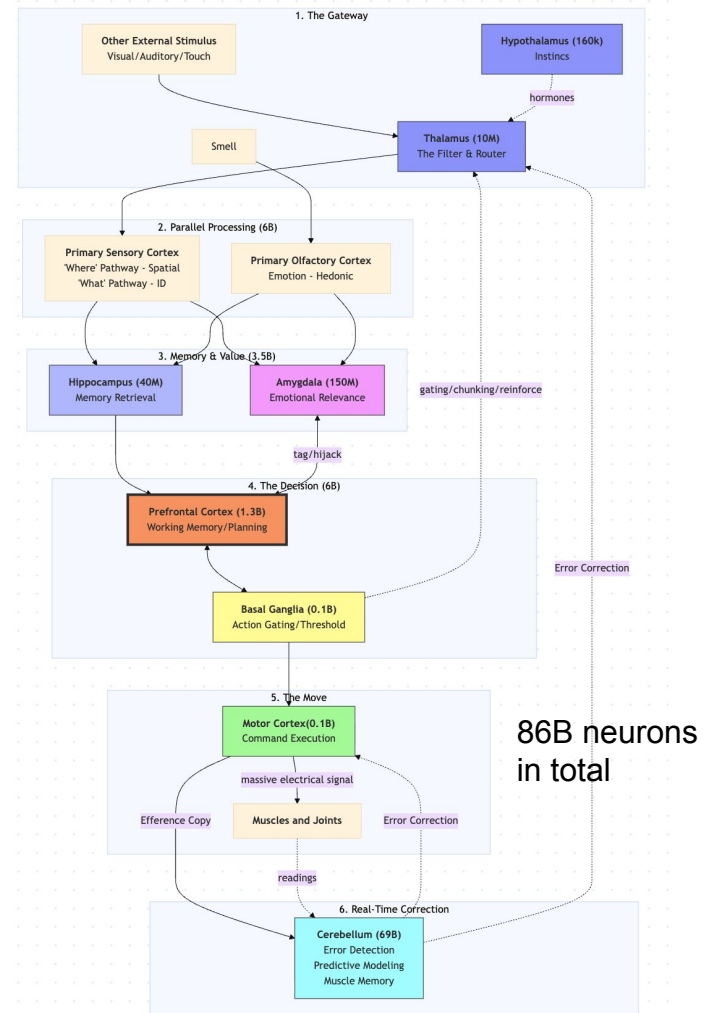
Brain information flow

1. **The Filter (Thalamus):** To focused on a task it suppresses "noise" and passes only relevant signals.
2. **Parallel Processing (Sensory Cortex):** "What is it?" "Where is it?" "What emotion?" **[Semantic Knowledge]**
3. **Memory Lookup (Hippocampus) and value tagging (Amygdala):** checks if the object is dangerous, and if there is past memory of it **[Experience]**
4. **Reasoning (PFC) and decisions (Basal Ganglia):** PFC simulates different outcomes. BG triggers actions above its threshold
5. **The execution (Motor Cortex):**
6. **The rule store / Correction Loop (Cerebellum):** As MC fires, it sends a copy of the plan to the Cerebellum, which detects motion errors and sends lightning-fast correction back to MC. **[Procedural Knowledge]**



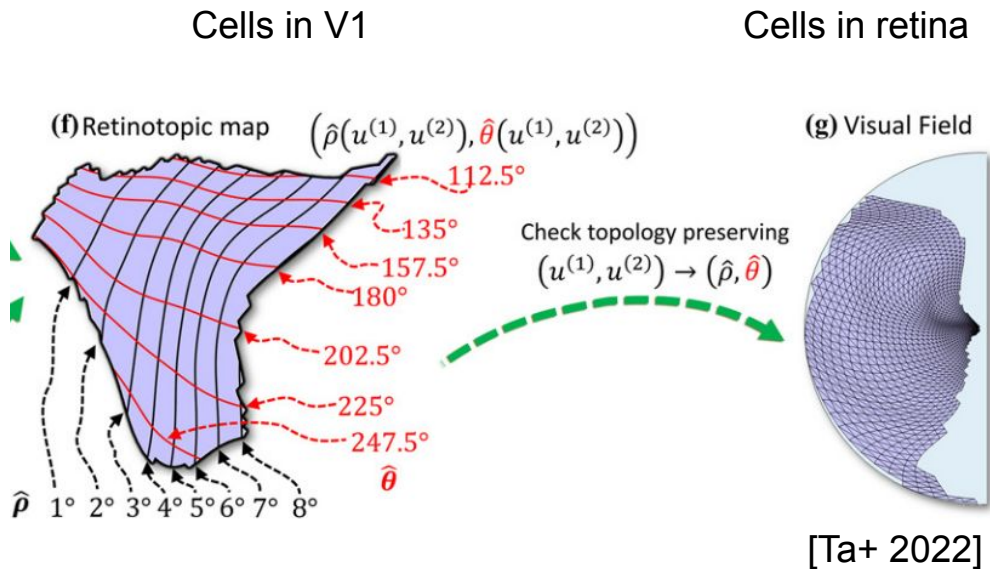
The related concepts in AI

1. **The Filter** (Attention):
2. **Parallel Processing** (Modality specific encoders):
3. **Memory Lookup and value tagging** (Experience replay and value function):
4. **Making decisions** (sequence model capable of reasoning/planning and reinforcement learning):
5. **The execution** (function calls, e.g., write files):
6. **The rule store / Correction Loop** (system prompts, skill, and logical rules): no need to make decisions



Micro-organization of V1


- **V1**: the 1st layer of visual cortex
- V1 retinotopic mapping is **Topology Preserving** -- cells closer in retina are projected to targets closer in V1
- How was it achieved during development?
- It can be simulated by **Self-Organizing Maps (SOM)** (Schulten+, 1990, 1995)





Retinal Waves & SOM

- **Self-Organizing Maps** was invented to model the development of visual and somatosensory maps (Kohonen, 1982, 1990, 2001)
- **Application:** e.g. visualize a dataset as a 2D map

(A) A Spontaneous Retinal Wave



(C)



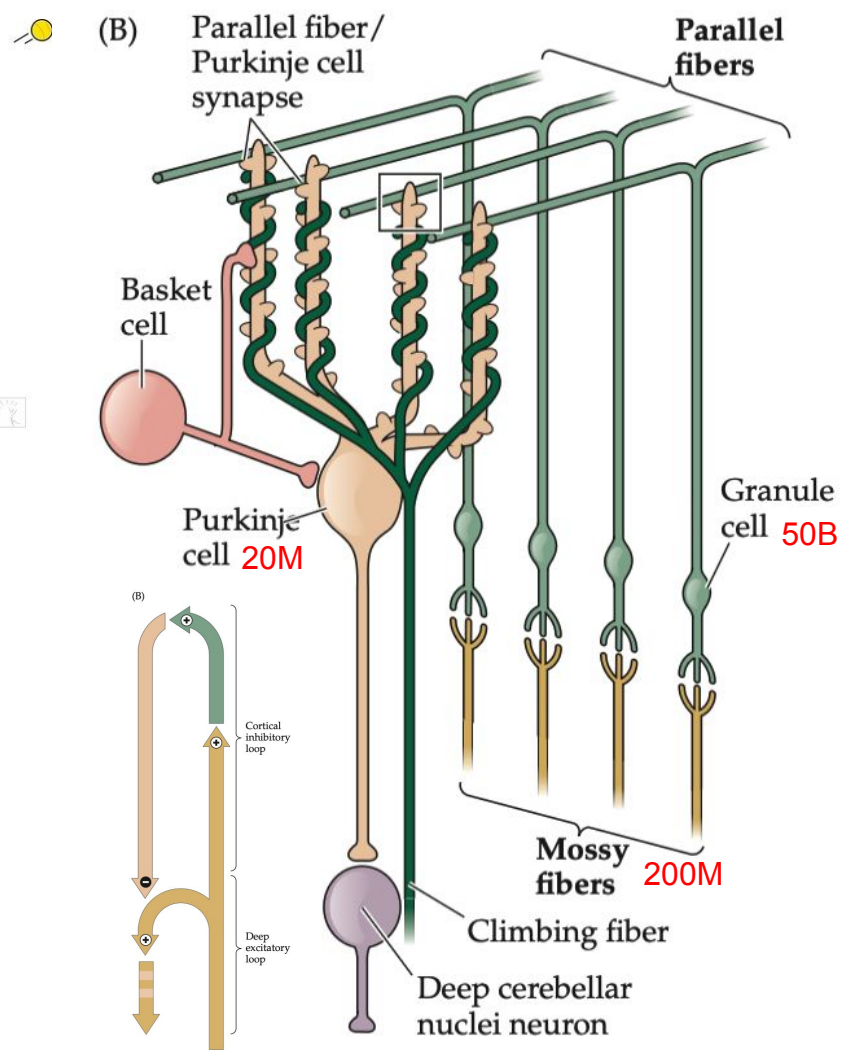
	Retinotopic Map	SOM
Initialization	chemical markers	random or PCA
Data	spontaneous waves create spatiotemporal patterns	similarity scores of data points
Finding a winner and its neighborhood	a small cluster of neurons "wins" by lateral inhibition	the node closest to the data point and nodes in its Gaussian neighborhood

[Background]

Cerebellum



- **Context (Mossy Fibers):** "what the motor cortex **wants** to do", "what the organs are **actually** doing"
- **Sparse Coding (Granule Cells \w Parallel Fibers):** **every microscopic combination of body-state and motor-intent gets a unique context.** Each synapse is a rule in the form [break certain action P when certain context G is achieved]
- **Pattern Recognition (Purkinje Cells, inhibitory):** each a linear classification with input from **200k** Parallel Fibers




[Background]

Cerebellum




- **Before Learning:** raw, clumsy, beginner movement (heavily inhibited and stiff)
- **Error Correction (Climbing Fibers):** any detected error cause a massive electrical spike to weaken recently involved synapses -- **Long-Term Depression (LTD)**
- **Perfection (Disinhibition)** after a few tries, most wrong plans are suppressed leaving only the perfect plans.

Now we achieve fluid action without random suppressions



Cerebellum & Minimax

- **Minimax algorithm** chooses the guess that minimizes the maximum number of possible solutions that could remain.
- #surviving solutions is reduced exponentially



	Cerebellum	Minimax Algorithm
Dimension expansion	Billions of Granule Cells as context	All possible 4 digits
Decision	Purkinje Cells triggered by highest weighted context	Pick the guess with the highest vote from remaining rules
Model Update	Climbing Fibers causing LTD to all triggered synapses	Remove all unfit 4 digits


Bulls & Cows

- **Secret number:** 4271
- **Opponent's try:** 1234
- **Answer:** 1 bull and 2 cows.
(The bull is "2", the cows are "4" and "1".)

Cerebellum & ELM

- **Extreme Learning Machine (ELM)**
 - **FFN** (very shallow network for fast inference)
 - **Kernel trick** (dimensionality expansion)

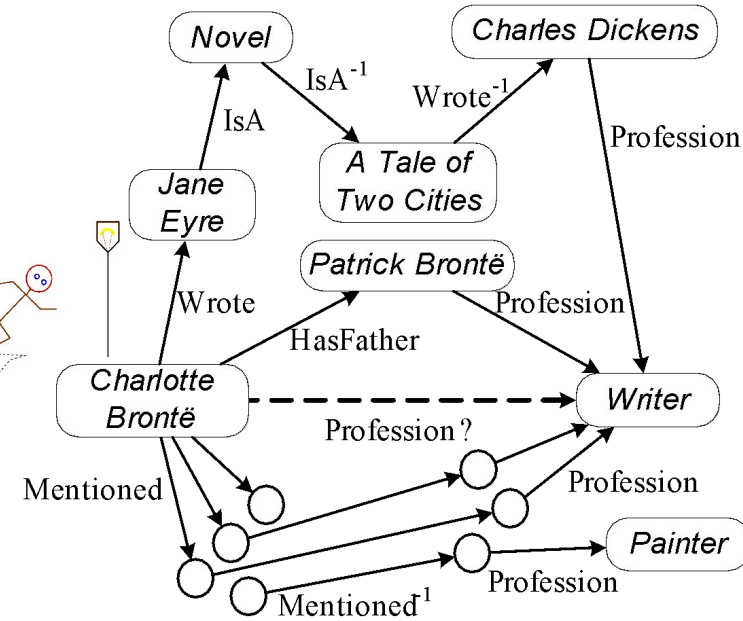
	Cerebellum	ELM
Dimension expansion	Billions of Granule Cells	Kernel Tricks (e.g. random hash, RBF)
Output	Purkinje Cells (perceptron)	Linear Output Layer
Model Update	Climbing Fibers causing LTD	gradient descent for squared error loss



Cerebellum & PRA

- Path Ranking Algorithm (Lao & Cohen 2010)**

PRA uses path-constrained random walks to discover predictive patterns in a graph, and then weigh the importance of these patterns



	Cerebellum	PRA
Dimension expansion	Billions of Granule Cells as context	Millions of sampled logic forms
Model Update	Climbing Fibers causing LTD to all triggered synapses	Filtering logic forms by precisions/recalls. Logistic regression
Decision	Purkinje Cells triggered by highest weighted context	Linear combination of random walk probabilities

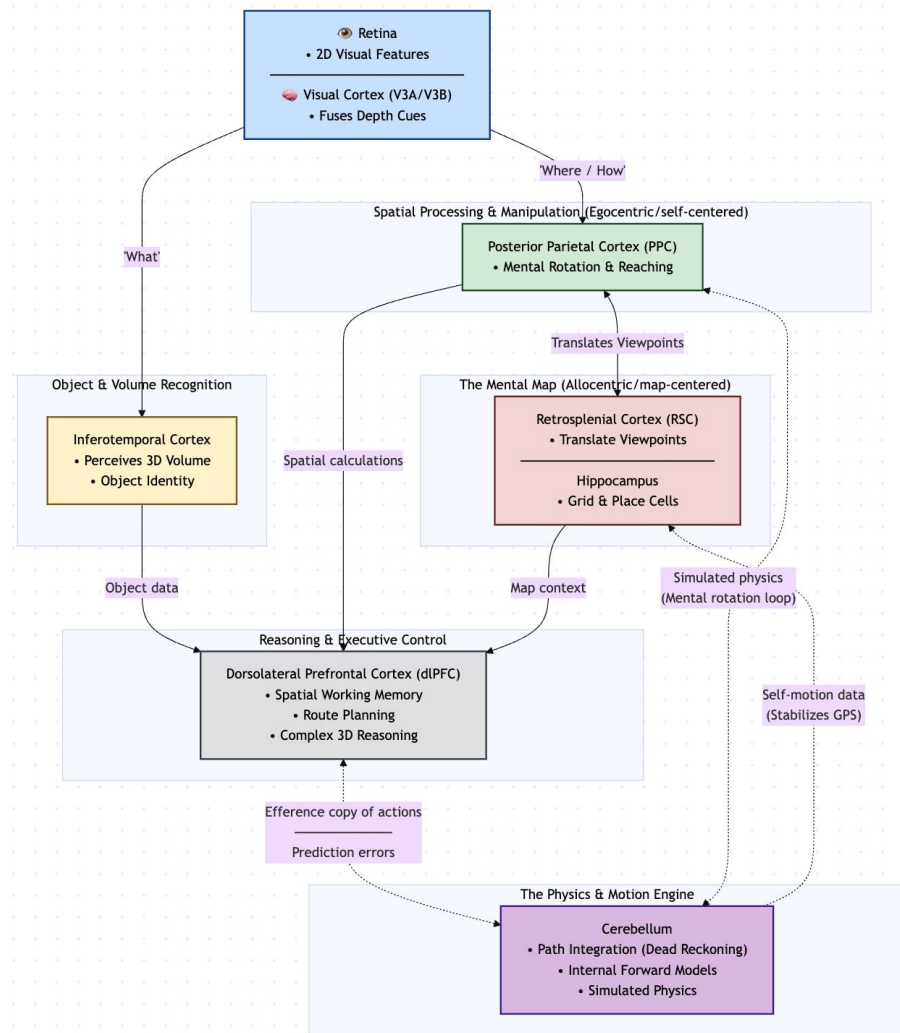
$P(Charlotte \Rightarrow Writer; \langle HasFather, IsA \rangle)$
 $P(Charlotte \Rightarrow Writer; \langle Mention, Mention^{-1}, IsA \rangle)$
 ...
 $P(Charlotte \Rightarrow Painter; \langle HasFather, IsA \rangle)$
 $P(Charlotte \Rightarrow Painter; \langle Mention, Mention^{-1}, IsA \rangle)$
 ...

Plan

- ***Introduction***
 - Neuroscience & algorithms
 - Retinal Waves & SOM
 - Motor learning & Minimax/ELM/PRA
- ● ***Spatial Representation & Reasoning***
 - Location cell and grid cells
 - agentic vision
- ***Memory & Reinforcement Learning***
 - World models
 - Dream & experience replays
- ***Knowledge Representation***
 - the ideal distribution of knowledge?

Spatial information flow


- Input & Early Fusion:** 2D image+ depth
- "What" stream** *what* the 3D shapes are
- "Where" stream:** *where* things are and how to interact with them (mental rotation). [Egocentric]
- The World Map:** translates *where* into a world-map [Allocentric]
- The Player/Policy:** holds all above context in short term memory to make a reasoned decision (e.g., planning the best route).
- The Physics & Motion Engine**
 - report "prediction errors"
 - "forward models" predict the sensory consequences of motor commands



How does the brain represent positions?

- **Morris water navigation task**



- required skills:
 - 1) representation of self position (position encoding)
 - 2) update the position according to movements (path integration)





Grid Cells & Place cells

- Represente places with a form of sparse coding -- **a simple Fourier model of the hexagonal lattice** (Blair et al. 2007)

(A)



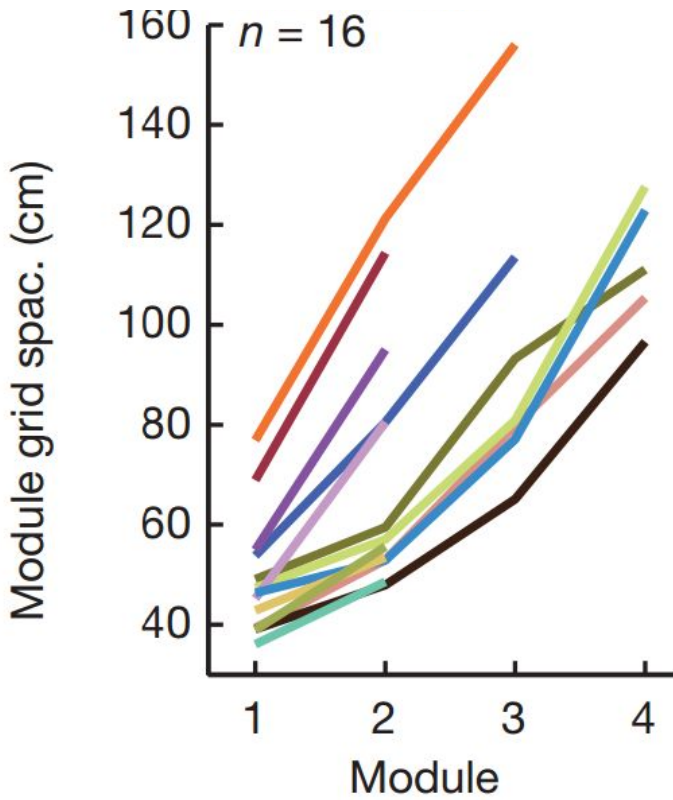
(B)



Multi-Scale Location Encoding

Grid cells in mammals provide a **multi-scale periodic representation** (Banino+ 2018), which is critical for **path integration** and **spatial localization**

$$G(\mathbf{r}) = g\left(\sum_{k=1}^3 \cos(\boldsymbol{\omega}_k \cdot (\mathbf{r} - \mathbf{c}))\right),$$

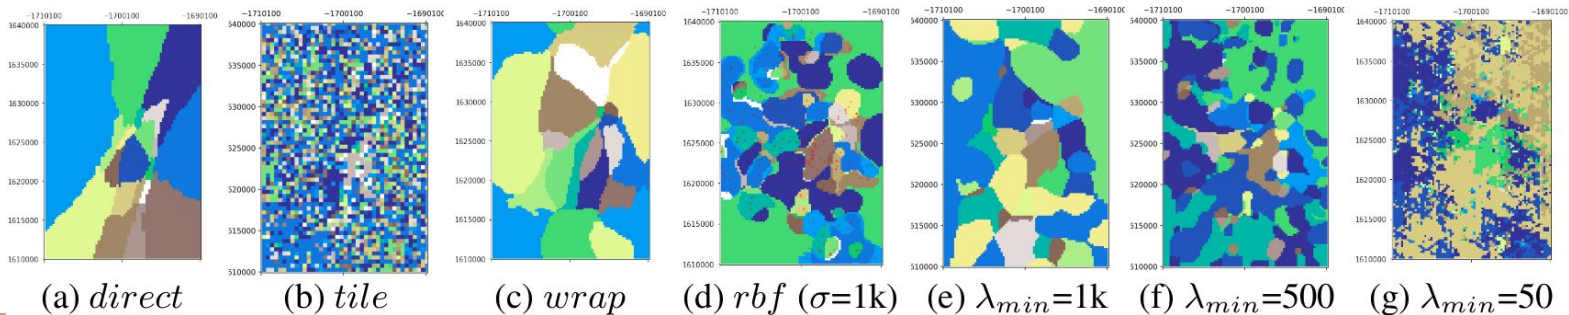


Mean grid spacing for all modules (M1–M4) in all animals (colour-coded)

Multi-scale Representation in Space2Vec

POI classification result for different distribution characteristics:

POI Groups	Clustered ($r \leq 100m$)	Middle ($100m < r < 200m$)	Even ($r \geq 200m$)
<i>direct</i>	0.080 (-0.047)	0.108 (-0.030)	0.084 (-0.047)
<i>wrap</i>	0.106 (-0.021)	0.126 (-0.012)	0.122 (-0.009)
<i>tile</i>	0.108 (-0.019)	0.135 (-0.003)	0.111 (-0.020)
<i>rbf</i>	0.112 (-0.015)	0.136 (-0.002)	0.119 (-0.012)
<i>theory</i>	0.127 (-)	0.138 (-)	0.131 (-)
# POI	16,016	7,443	3,915
Root Types	Restaurants; Shopping; Food; Nightlife; Automotive; Active Life; Arts & Entertainment; Financial Services	Beauty & Spas; Health & Medical; Local Services; Hotels & Travel; Professional Services; Public Services & Government	Home Services; Event Planning & Services; Pets; Education



Multi-step reasoning for vision

- Prompt

Read the gauge in (bar). Detect hub and needle tip and nearby labeled ticks, then calculate the result with math (never estimate or eyeballing). Give the result on the last line with number and unit only.



Multi-step reasoning for vision

- Both PFC and LLM play the orchestrator role in reasoning

Function	PFC	LLM
Short term memory	what, where, goal	user prompt and conversation history
Planning	simulations using Hippocampus and Cerebellum	simulate with thoughts and function calls
Making Decisions	decide attention and the next action	decide the next token or function call

Zoom in with agentic vision

```
import PIL.Image
import PIL.ImageDraw

# Define locations
gauge_center = (415, 420) # Roughly
# Looking closely at the gauge
# Hub is at ~ (420, 425) based on needle hinge
# needle tip ~ (250, 260)
```

Detect key points and calculate

```
import math

def get_angle(p, center):
    return math.degrees(
        math.atan2(p[1] - center[1], p[0] - center[0]))

...
hub = (155, 168)
tip = (114, 71)
t10 = (102, 75)
t15 = (202, 70)
a10 = get_angle(t10, hub)
atip = get_angle(tip, hub)
a15 = get_angle(t15, hub)
print(f"a10: {a10}, atip: {atip}, a15: {a15}")

# bar/degree = (15-10)/(a15 - a10)
bar_per_deg = 5.0 / (a15 - a10)
val = 10 + (atip - a10) * bar_per_deg
print(f"Calculated value: {val}")
```

Result

```
a10: -119.67848992513521, atip: -112.91282839677166, a15:
-64.37800405539302 Calculated value: 10.611718090895238
```

Final answer

10.6 bar

Plan

- ***Introduction***
 - Neuroscience & algorithms
 - Retinal Waves & SOM
 - Motor learning & Minimax/ELM/PRA
- ***Spatial Representation & Reasoning***
 - Location cell and grid cells
 - agentic vision
-  ● ***Memory & Reinforcement Learning***
 - World models
 - Dream & experience replays
- ***Knowledge Representation***
 - the ideal distribution of knowledge?

BG/PFC Learning & RL

- Both are optimizing an expected value (in the future) and rely on a set of similar strategies

	Basal Ganglia (BG) / Prefrontal Cortex (PFC)	Model-Free RL / Model-Based RL
Learning Objective	Future Dopamine	Future reward $R()$
Planning & World Simulation	habitual responses / simulated by Hippocampus & Cerebellum	none / hard-coded or learned transition dynamics
Credit Assignment	Working memory (sustained firing bridged by dopamine).	Backpropagation Through Time
Offline Learning	Dream	Experience Replay

Directly Optimizing The Expected Reward with RL

- **ML** optimizes the log likelihood of target sequences

$$J^{ML}(\theta) = \sum_q \log P(a_{0:T}^{best}(q)|q, \theta)$$


- **RL** optimizes the expected reward under a stochastic policy

$$J^{RL}(\theta) = \sum_q \mathbb{E}_{P(a_{0:T}|q, \theta)} [R(q, a_{0:T})]$$




[Sutton & Barto 1998, 2018]

RL models generate their training data on the fly

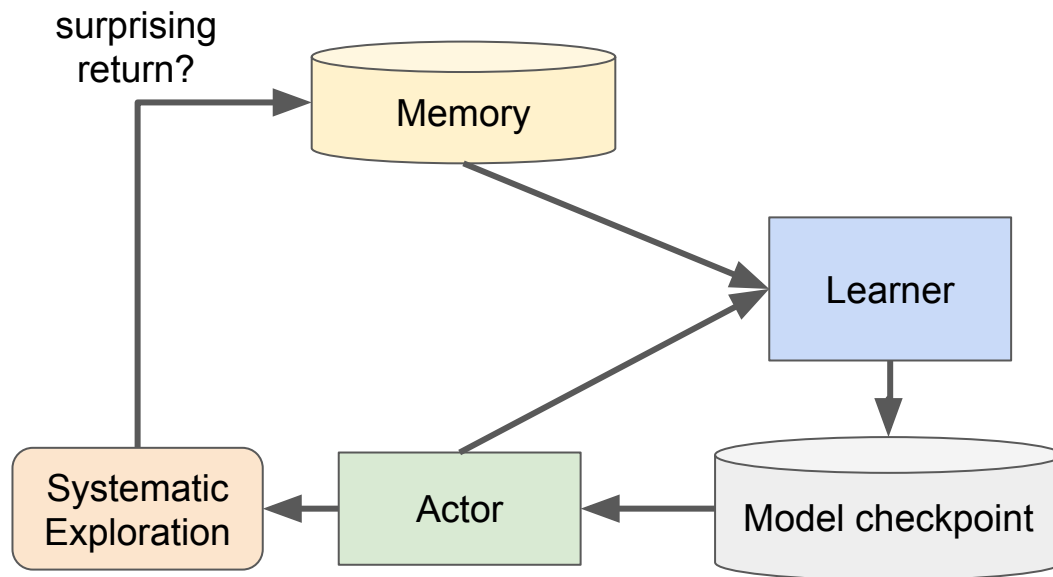


- Training sample management issue
 - Large search space & sparse reward lead to slow and unstable training
 - Spurious reward lead to biased solutions

Memory Augmented Policy Optimization (MAPO)




Memory Augmented Policy Optimization (MAPO)



Inferred mental states from place cells

- **Reversed Replay**: during wakeful rest immediately following an experience
 - -- mark important experience for replay later
- **Deep Sleep (N3)**: Replay of interesting experience at up to 20x speed
 - -- compress experience for long term storage using Sensory Cortex & PFC
 - -- extract patterns (building world models, e.g. Glowing metal = heat = pain)
- **REM (Rapid Eye Movement)**: Preplay of yet-to-explore situations at x1 speed
 - -- explore solutions to unsolved problems



Unbiased gradient estimation w low variances

Given a memory buffer of high return sequences $\mathcal{B} \equiv \left\{ (y^{(i)}, r^{(i)}) \right\}_{i=1}^n$,
re-express expected return as,

$$p(\mathcal{B}) \underbrace{\mathbb{E}_{p(\tilde{y})|\tilde{y} \in \mathcal{B}} R(\tilde{y})}_{\text{inside the buffer}} + (1 - p(\mathcal{B})) \underbrace{\mathbb{E}_{p(\tilde{y})|\tilde{y} \notin \mathcal{B}} R(\tilde{y})}_{\text{outside the buffer}}$$

- For each query
 - Sampling 1 solution from **inside** the buffer according to model
 - Rejection sampling 1 solution from **outside** the buffer according to model

Most of the past experience are not helpful for improving the current model. We pick the ones with high emotional values at sleep



Optimal Sample Allocation

- Assume 0-1 reward
- The optimal strategy (low variance in gradient estimations) is to allocate the **same number of samples** to **reward** vs **no reward** experiences
- This is independent of the model's **current performance**


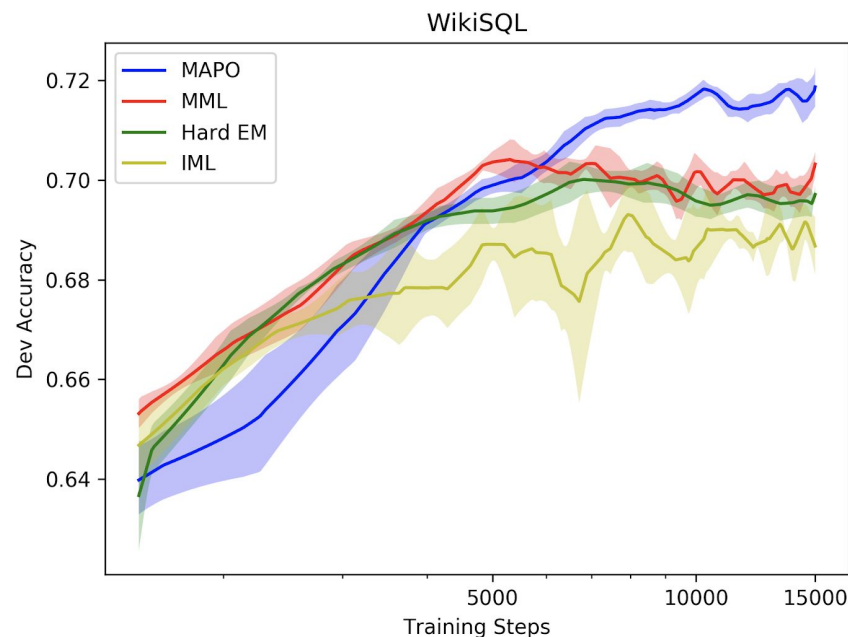
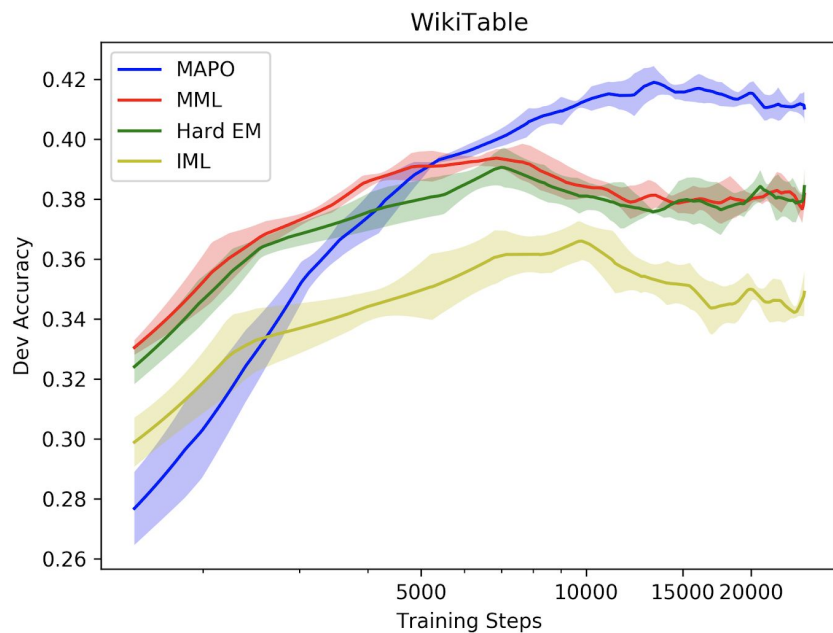


Image source: Guy Harris, 2018
How to Give Feedback in a Non-Threatening Way

Comparison

- MAPO is both stable and converges to more optimal policy



- The shaded area represents the standard deviation of the dev accuracy

Plan

- ***Introduction***
 - Neuroscience & algorithms
 - Retinal Waves & SOM
 - Motor learning & Minimax/ELM/PRA
- ***Spatial Representation & Reasoning***
 - Location cell and grid cells
 - agentic vision
- ***Memory & Reinforcement Learning***
 - World models
 - Dream & experience replays
- ***Knowledge Representation***
 - the ideal distribution of knowledge?

Comparing the Distribution of knowledge


- majority of the neurons are in the "Pattern Store" (procedural knowledge)

	Human Brain (#Neurons)	Gemma 4 31B (#Params)
Total Units	86B	31B
Filtering / Reasoning	The Filter (Thalamus): 10M Reasoning (PFC) and decisions (BG): 6B	Attention Layer ~9B*
Representation	Representation (Sensory Cortex): 6B	Embedding Layer ~2B ViT Encoder ~0.5B
Pattern Store	The machine (Cerebellum): 69B	Feed-Forward layers ~20B*
Training	Memory Lookup (Hippocampus) and value tagging (Amygdala): 3.5B	N/A

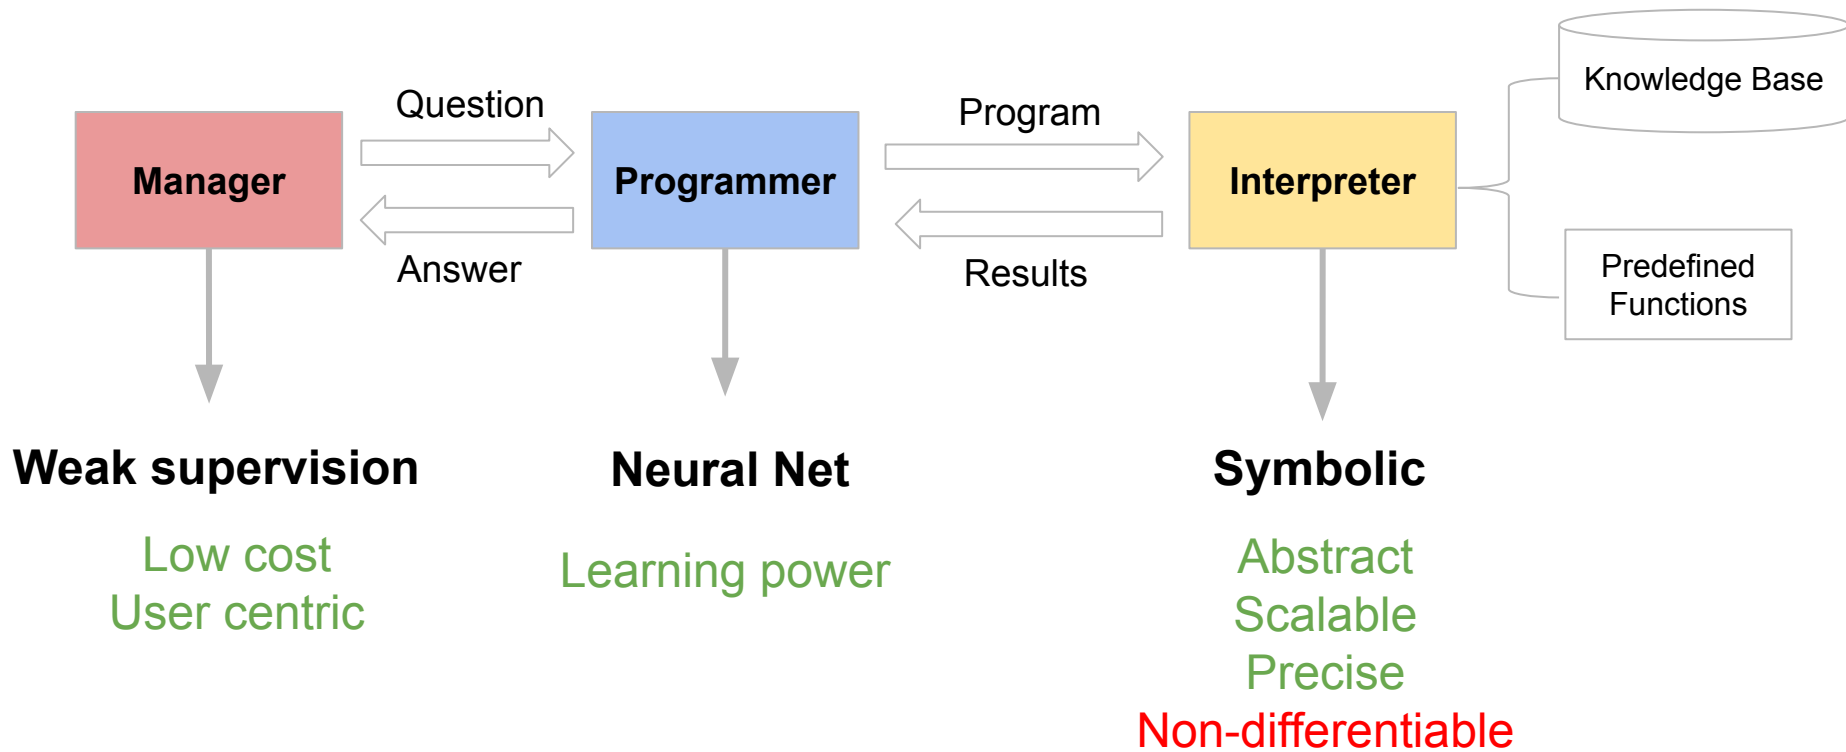
*The standard 2:1 (FFN to Attention) ratio can increase if certain techniques like Grouped-Query Attention (GQA) or Mixture of Experts (MoE) are applied

Symbolic Rules & The Policy

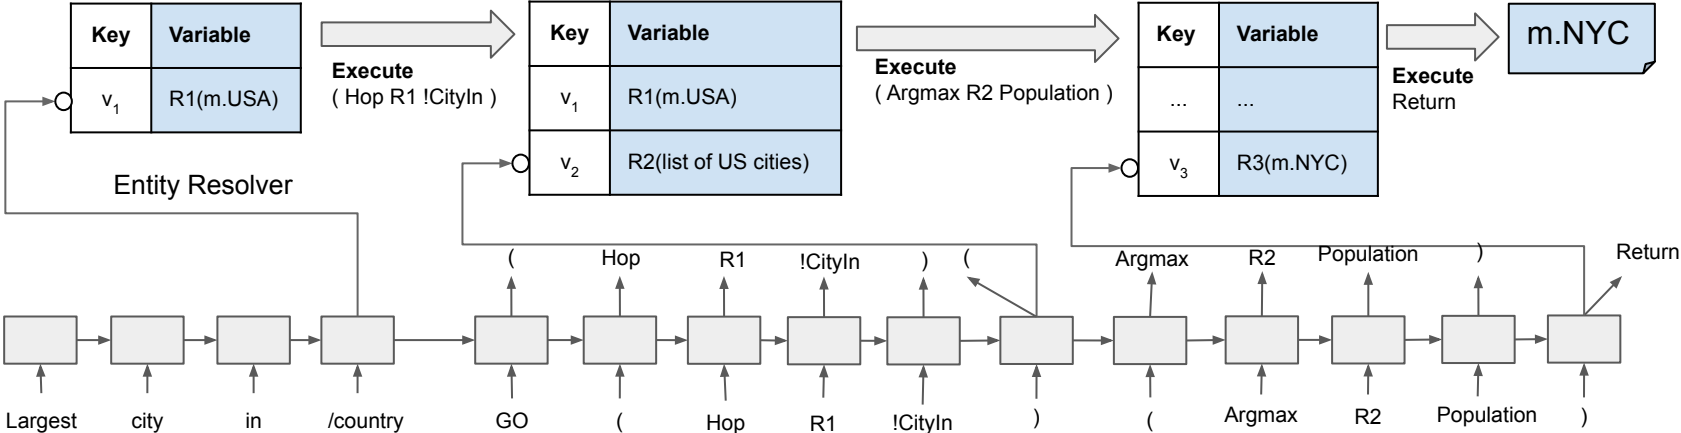
- Combine the **learnability** of neural nets with the **efficiency** of symbolic reasoning




Neural Symbolic Machines



Seq2Seq with Variables for Compositionality



- Equivalent to a linearised bottom-up derivation of the recursive program
- Aggressive pruning by code assists



How is procedural knowledge stored in today's AI

- Model parameters
 - typically $O(10B) \sim O(100B)$
- The Web
 - typically $\sim 10B$ commonly viewed pages (e.g., Bing)
- Knowledge graph
 - typically $\sim 1B$ nodes (e.g., Wikidata, Satori)
- LLM Agent Skills
 - typically $< 1M$ lines (e.g., Claude Code)

 Thanks!