Beyond Backpropagation of Errors

Predictive Coding and Biomimetic Intelligence

AGI 2024 Keynote
August 15, 2024

Alexander Ororbia
The Neural Adaptive Computing (NAC) Laboratory
Rochester Institute of Technology

RIT | Rochester Institute of Technology

NAC
Neural Adaptive Computing Laboratory

Getting Rid of ~~Beyond~~ Backpropagation of Errors

Predictive Coding and Biomimetic Intelligence

AGI 2024 Keynote
August 15, 2024

Alexander Ororbia
The Neural Adaptive Computing (NAC) Laboratory
Rochester Institute of Technology

RIT | Rochester Institute of Technology

NAC
Neural Adaptive Computing Laboratory

Problem Solved! AGI is Upon Us...Or is it?

- The machine learning community, for instance, is producing breakthroughs on a monthly basis
- Many (over)promises that AGI is upon us
  - Generative pre-trained transformers (GPTs) and backpropagation of errors ***are all we need***…

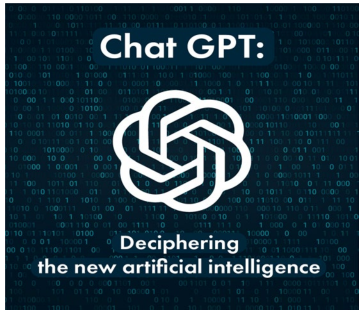

Sparks of Artificial General Intelligence:
Early experiments with GPT-4

Sébastien Bubeck Varun Chandrasekaran Ronen Eldan Johannes Gehrke
Eric Horvitz Ece Kamar Peter Lee Yin Tat Lee Yuanzhi Li Scott Lundberg
Harsha Nori Hamid Palangi Marco Tulio Ribeiro Yi Zhang

Microsoft Research

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

History Tends to Repeat Itself…

- It was once thought that conscious symbolic reasoning / formal rules were ***all you need***…
    - Criticizing this view was met with ridicule and hostility

- Later, it was then thought that kernel machines were ***all you need***…
    - Criticizing this view was met with ridicule and hostility

Geoffrey Hinton spent 30 years hammering away at an idea most other scientists dismissed as nonsense. Then, one day in 2012, he was proven

WHAT COMPUTERS CAN'T DO

OF ARTIFICIAL REASON

By Hubert L. Dreyfus

ALCHEMY AND ARTIFICIAL INTELLIGENCE

Hubert L. Dreyfus

December 1965

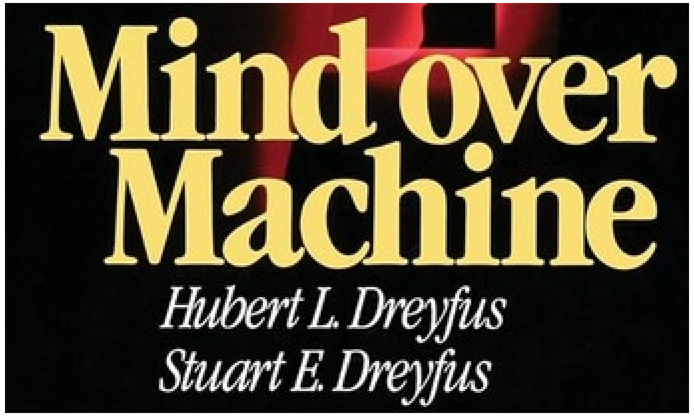

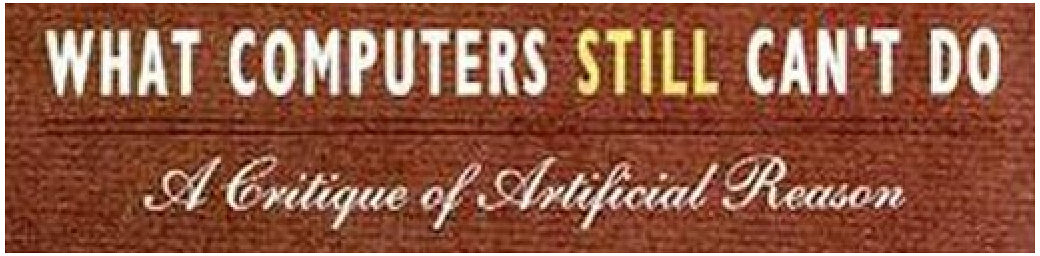

NATURE VOL. 337 12 JANUARY 1989 COMMENTARY 129

The recent excitement about neural networks

Francis Crick

Backprop and Deep GPTs: ALL We Really Need?

- **Red AI** = machine intelligence w/ massive carbon footprints
  - Requires hundreds of technicians for GPUs
  - Addressing this is fundamental to saving money/energy and democratize AI

- Generalization issues
  - Constraints on architecture (must be “backprop-pleasing”)
  - Weak out-of-distribution generalization
  - Catastrophic forgetting is still largely ***an unsolved problem***!

Model training can cost millions of dollars, pollution equal to 1000s of planes!

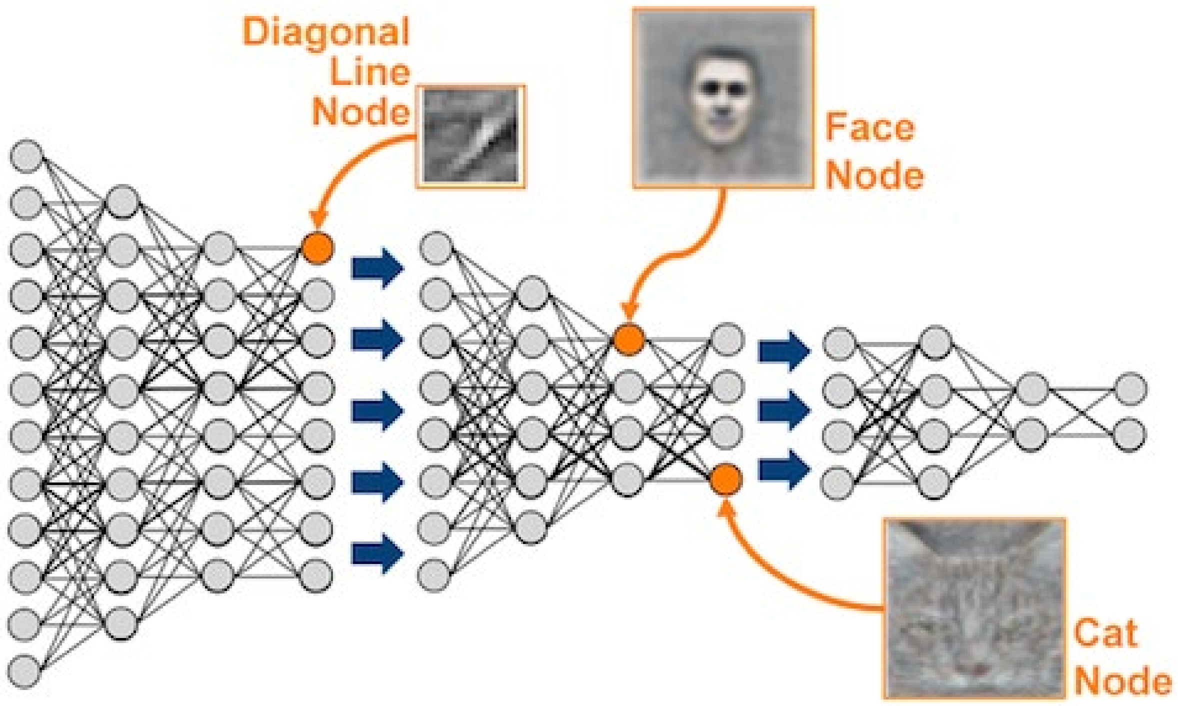

So, what could we turn to once we realize backprop-based deep learning might not be all we need?

Maybe we might consider and look to
biomimetics, bionics, and **mortal computation**

Perspective | Published: 17 April 2020

Backpropagation and the brain

Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman & Geoffrey Hinton

So, what could we turn to once we realize backprop-based deep learning might not be all we need?

Maybe we might consider and look to
biomimetics, bionics, and **mortal computation**

Perspective | Published: 17 April 2020

Backpropagation and the brain

Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman & Geoffrey Hinton

Toward an Integration of Deep Learning and Neuroscience

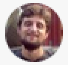 Adam H. Marblestone[1*]  Greg Wayne[2]  Konrad P. Kording[3]

So, what could we turn to once we realize backprop-based deep learning might not be all we need?

Maybe we might consider and look to
biomimetics, bionics, and **mortal computation**

Perspective | Published: 17 April 2020

Backpropagation and the brain

Marris, Colin J. Akerman & Geoffrey Hinton

Catalyzing next-generation Artificial Intelligence through NeuroAI

Anthony Zador, Sean Escola, Blake Richards, Bence Ölveczky, Yoshua Bengio, Kwabena Boahen, Matthew Botvinick, Dmitri Chklovskii, Anne Churchland, Claudia Clopath, James DiCarlo, Surya Ganguli, Jeff Hawkins, Konrad Körding, Alexei Koulakov, Yann LeCun, Timothy Lillicrap, Adam Marblestone, Bruno Olshausen, Alexandre Pouget, Cristina Savin, Terrence Sejnowski, Eero Simoncelli, Sara Solla, ... Doris Tsao

Tow
and

Adam H. Marblestone[1*] Greg Wayne[2] Konrad P. Kording[3]

So, what could we turn to once we realize backprop-based deep learning might not be all we need?

Maybe we might consider and look to
biomimetics, bionics, and **mortal computation**

A Review of Neuroscience-Inspired Machine Learning

Alexander Ororbia
Rochester Institute of Technology
Rochester, NY 14623, USA
ago@cs.rit.edu

Ankur Mali
University of South Florida
Tampa, FL 33620, USA
ankurarjunmali@usf.edu

Adam Kohan
University of Massachusetts Amherst
Amherst, MA, USA
akohan@umass.edu

Beren Millidge
Zyphra, Palo Alto, CA
University of Oxford, Oxford, UK
beren@millidge.name

Tommaso Salvatori
VERSES AI Research Lab, Los Angeles, USA
TU Wien, Vienna, Austria
tommaso.salvatori@verses.ai

Catalyzing next-gener
through NeuroAI

Anthony Zador, Sean Escola, Blake Richar
Matthew Botvinick, Dmitri Chklovskii, Anne
Hawkins, Konrad Körding, Alexei Koulakov,
Olshausen, Alexandre Pouget, Cristina Savin

Tow
and

Adam H. Marblestone[1*] Greg Wayne[2] Konrad P. Kording[3]

So, what could we turn to once we realize backprop-based deep learning might not be all we need?

Maybe we might consider and look to
biomimetics, bionics, and **mortal computation**

A REVIEW OF NEUROSCIENCE-INSPIRED
...NE LEARNING

MORTAL COMPUTATION: A FOUNDATION FOR BIOMIMETIC INTELLIGENCE

Alexander Ororbia
Rochester Institute of Technology
Rochester, NY 14623
ago@cs.rit.edu

Karl Friston
VERSES AI Research Lab
Los Angeles, CA 90016, USA
karl.friston@verses.ai

...kur Mali
... of South Florida
...FL 33620, USA
...nmali@usf.edu

Adam Kohan
University of Massachusetts Amherst
Amherst, MA, USA
akohan@umass.edu

University of Oxford, Oxford, UK
beren@millidge.name

Tommaso Salvatori
VERSES AI Research Lab, Los Angeles, USA
TU Wien, Vienna, Austria
tommaso.salvatori@verses.ai

Matthew Botvinick, Dmitri Chklovskii, Anne ...
Hawkins, Konrad Körding, Alexei Koulakov, Yann LeCun, Timothy Lillicrap, Adam Marblestone, Bruno
Olshausen, Alexandre Pouget, Cristina Savin, Terrence Sejnowski, Eero Simoncelli, Sara Solla, ... Doris Tsao

The Journey to Biomimetic Intelligence

- Bio-motivated self-organization, structural selection, adaptivity
- Credit assignment ***beyond*** backprop
- Bio-plausibility is not a niche property of interest to neuroscientists:
  - Vital for implementation on energy-efficient neuromorphic chips

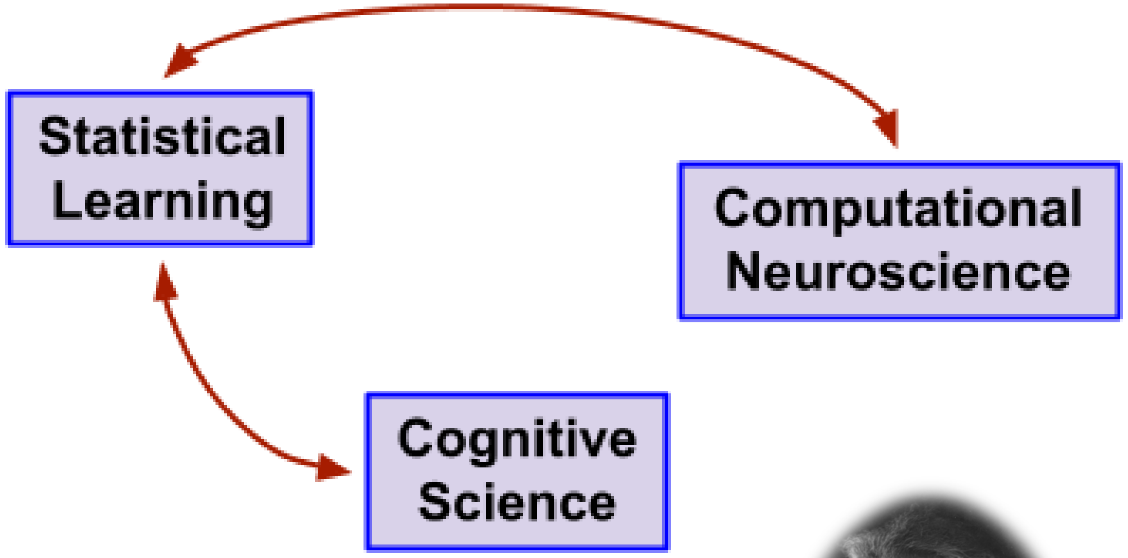

- The devil is in the details
  - How much neurobiological detail is needed?

    "Everything should be made as simple as possible, but no simpler."

    – Albert Einstein

What’s Wrong with Deep Learning and Backprop?

Challenging the way deep learning works today!

The Neural Processing Unit

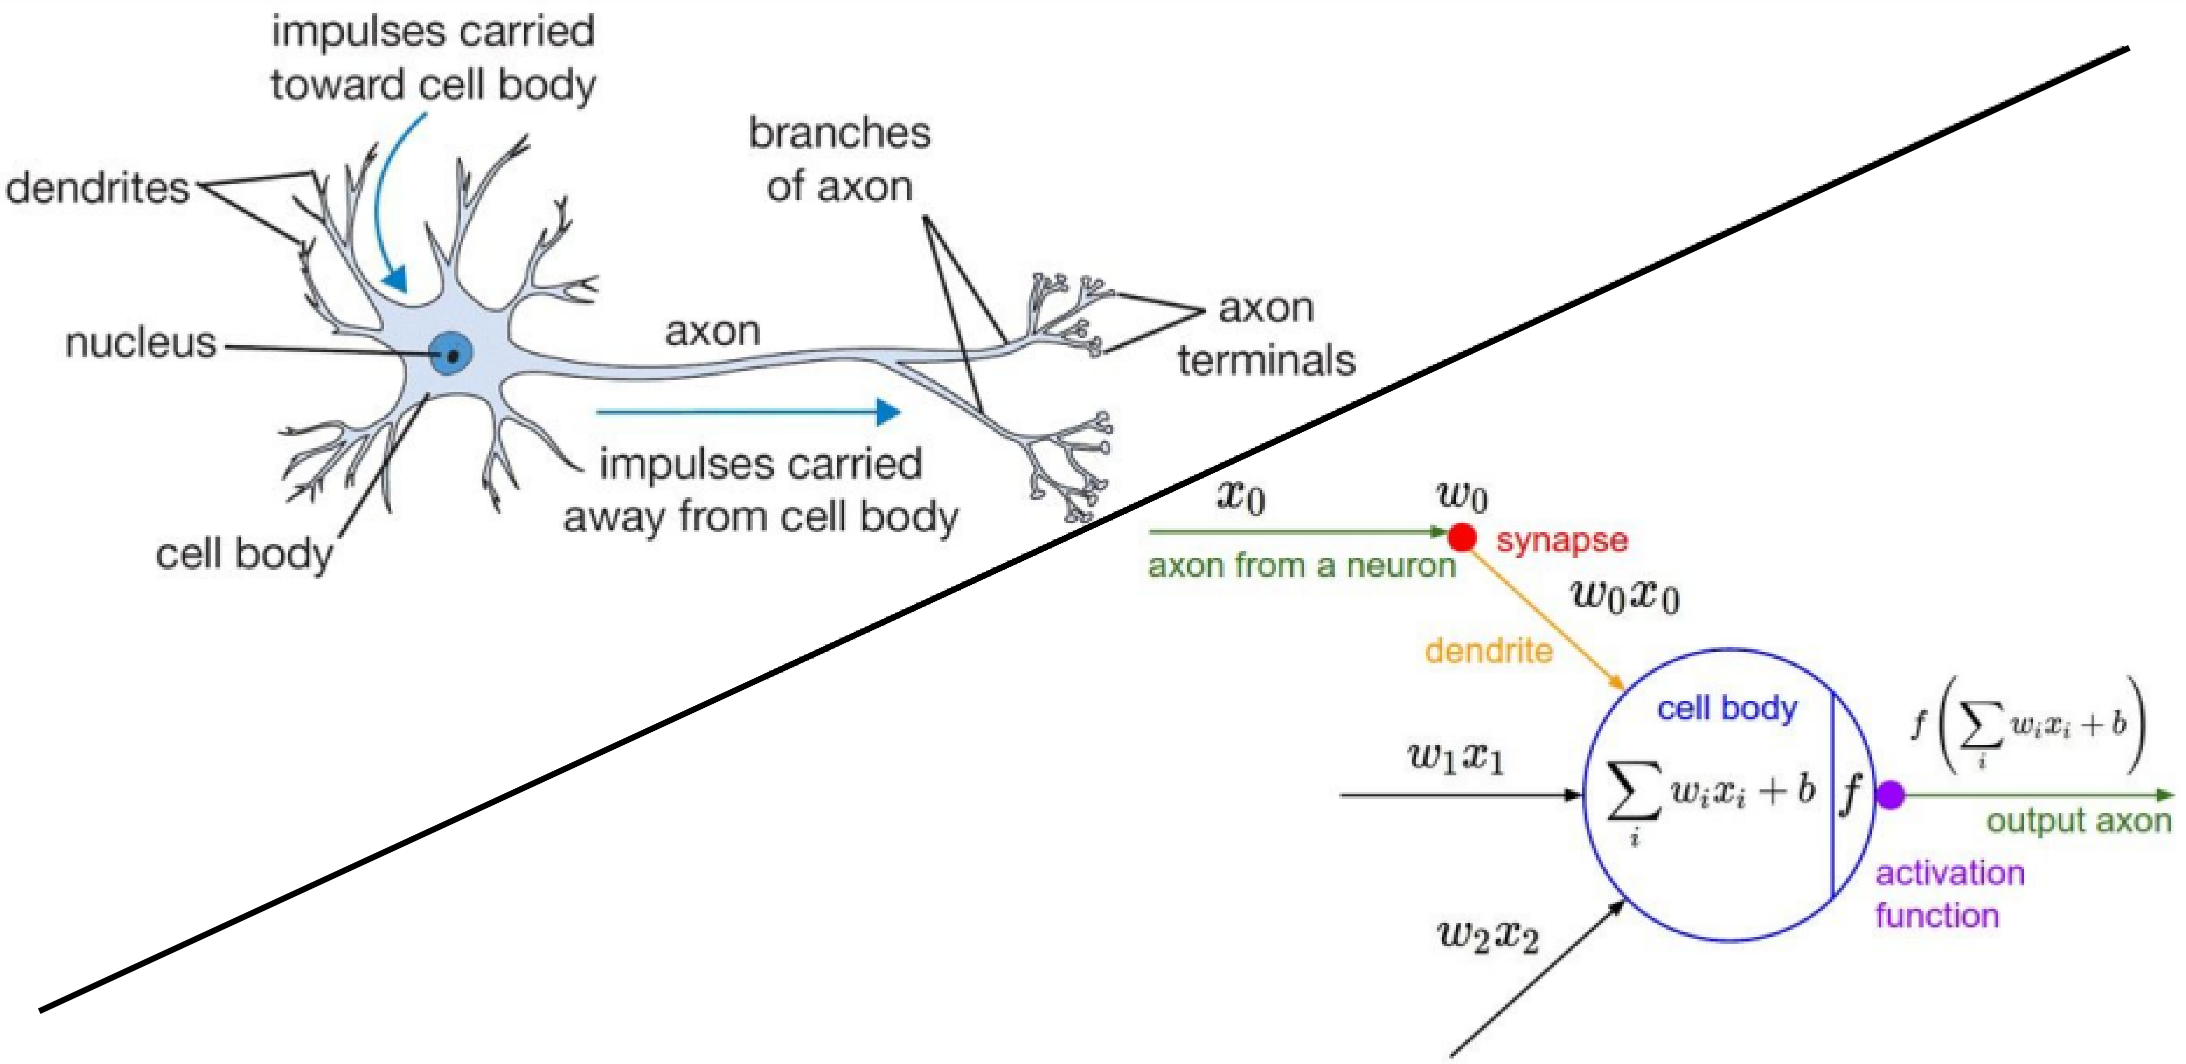

The Neural Processing Unit

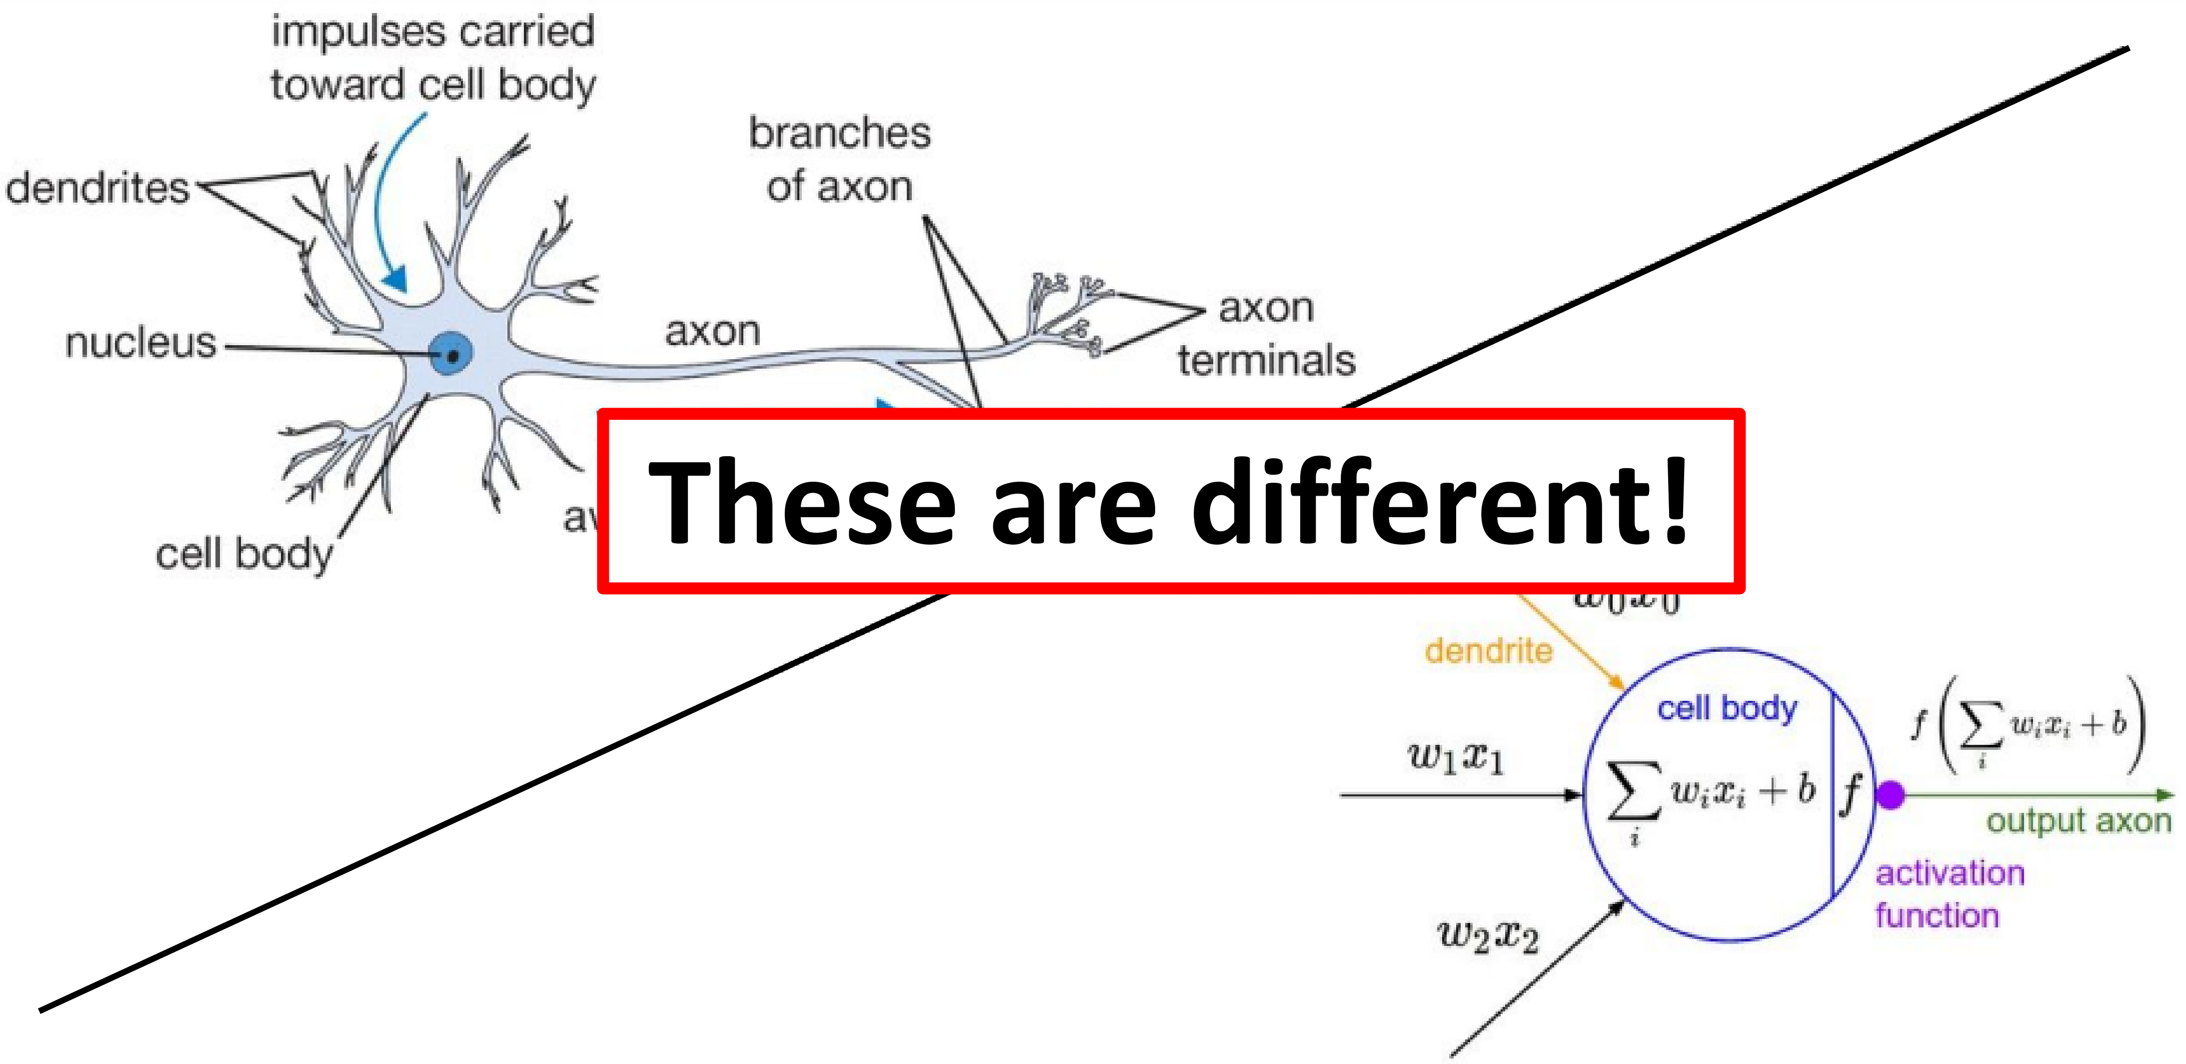

These are different!

Backpropagation of Errors (Backprop)

Used in an update/optimization step (e.g., SGD):

$$\mathbf{W}^\ell \leftarrow \mathbf{W}^\ell - \eta \Delta \mathbf{W}^\ell$$

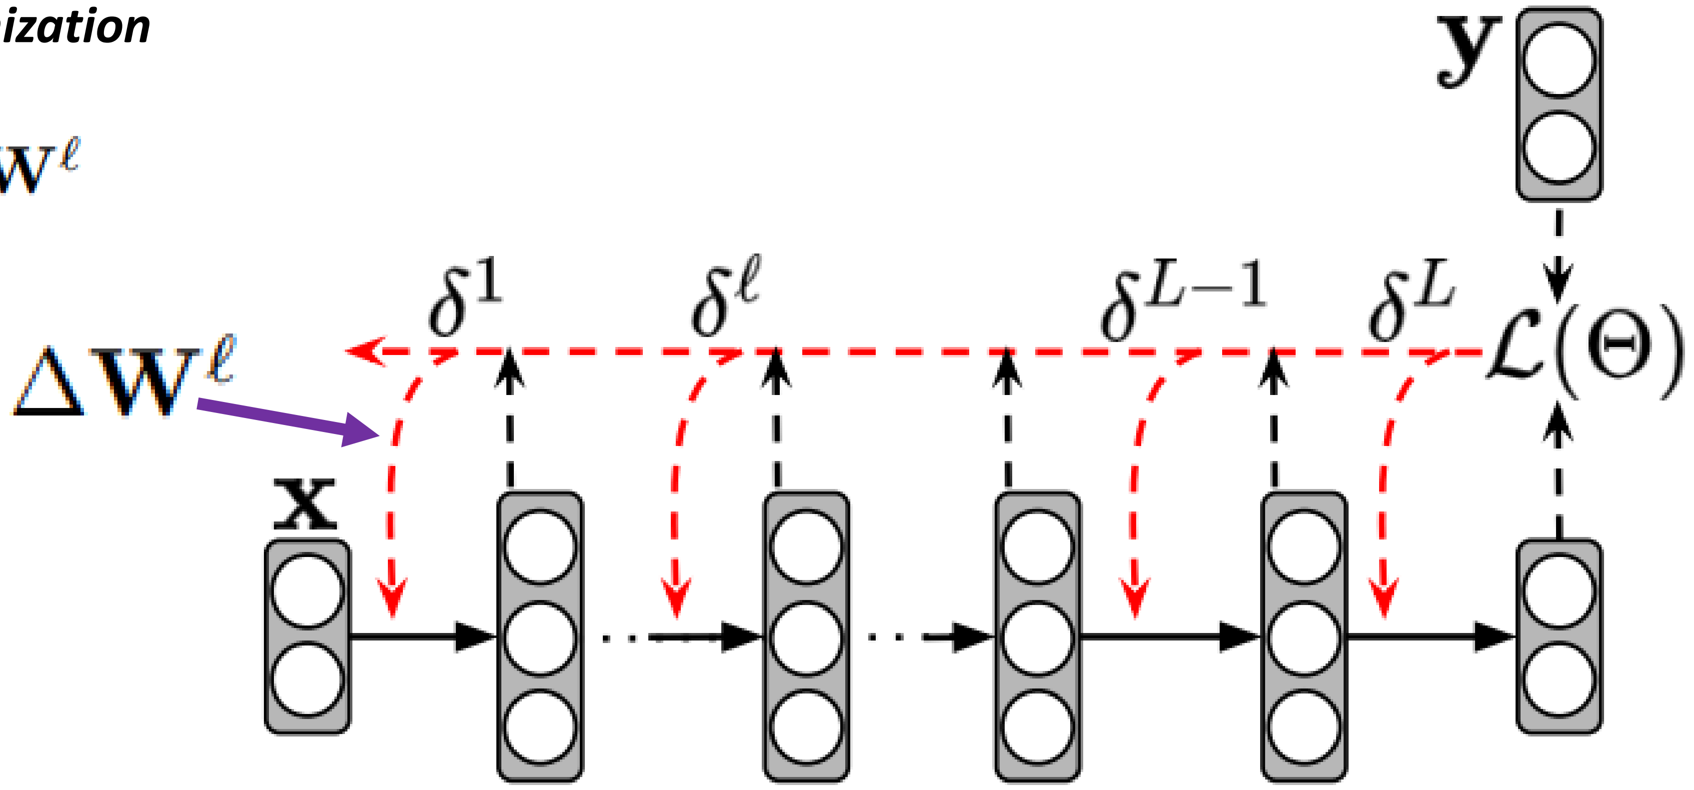

Rumelhart, DE, et al., "Learning representations by back-propagating errors." 1986.

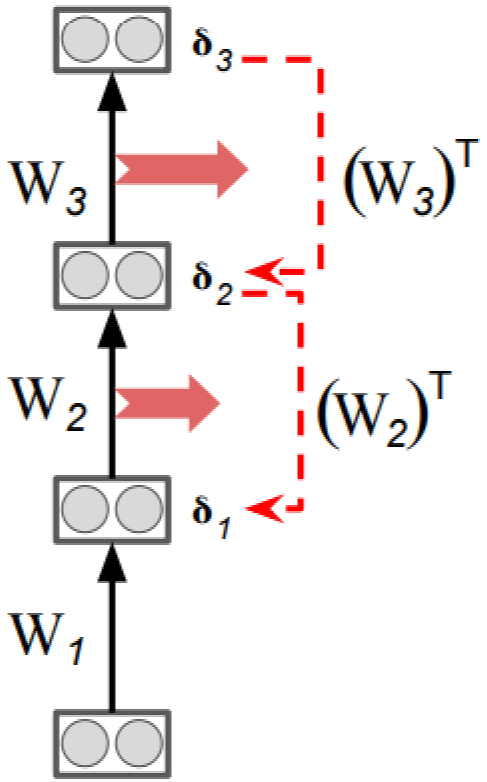

(a) The weight transport problem.

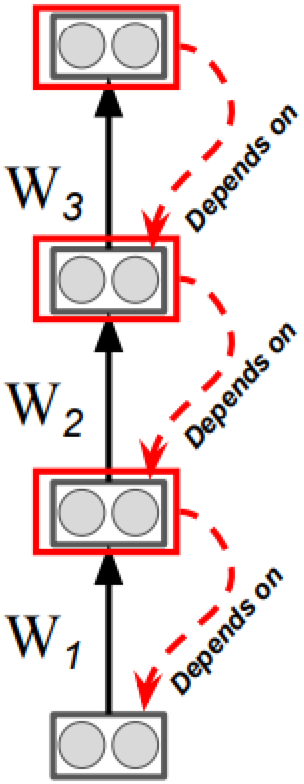

(b) The forward-locking problem.

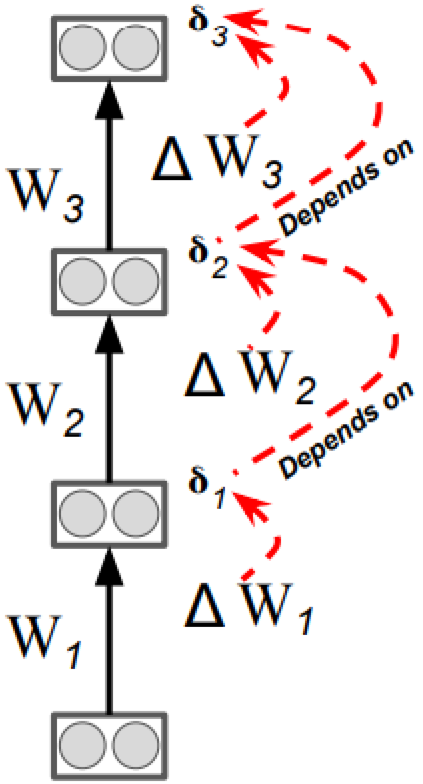

(c) The update-locking problem.

- ***Other Problems with Backprop:***
  - Global Feedback Pathways
  - Inference-Learning Dependency
  - Constraint and Sensitivity (e.g., requires derivatives)
  - Short-term plasticity, dynamics

Credit assignment = the credit/blame game

Grossberg, S. "Competitive learning: From interactive activation to adaptive resonance." 1987.

Crick, F. "The recent excitement about neural networks." 1989.

Ororbia, AG. "Brain-Inspired Machine Intelligence: A Survey of Neurobiologically-Plausible Credit Assignment. 2023.

Backprop through time and recurrent networks:

The issues get worse!!

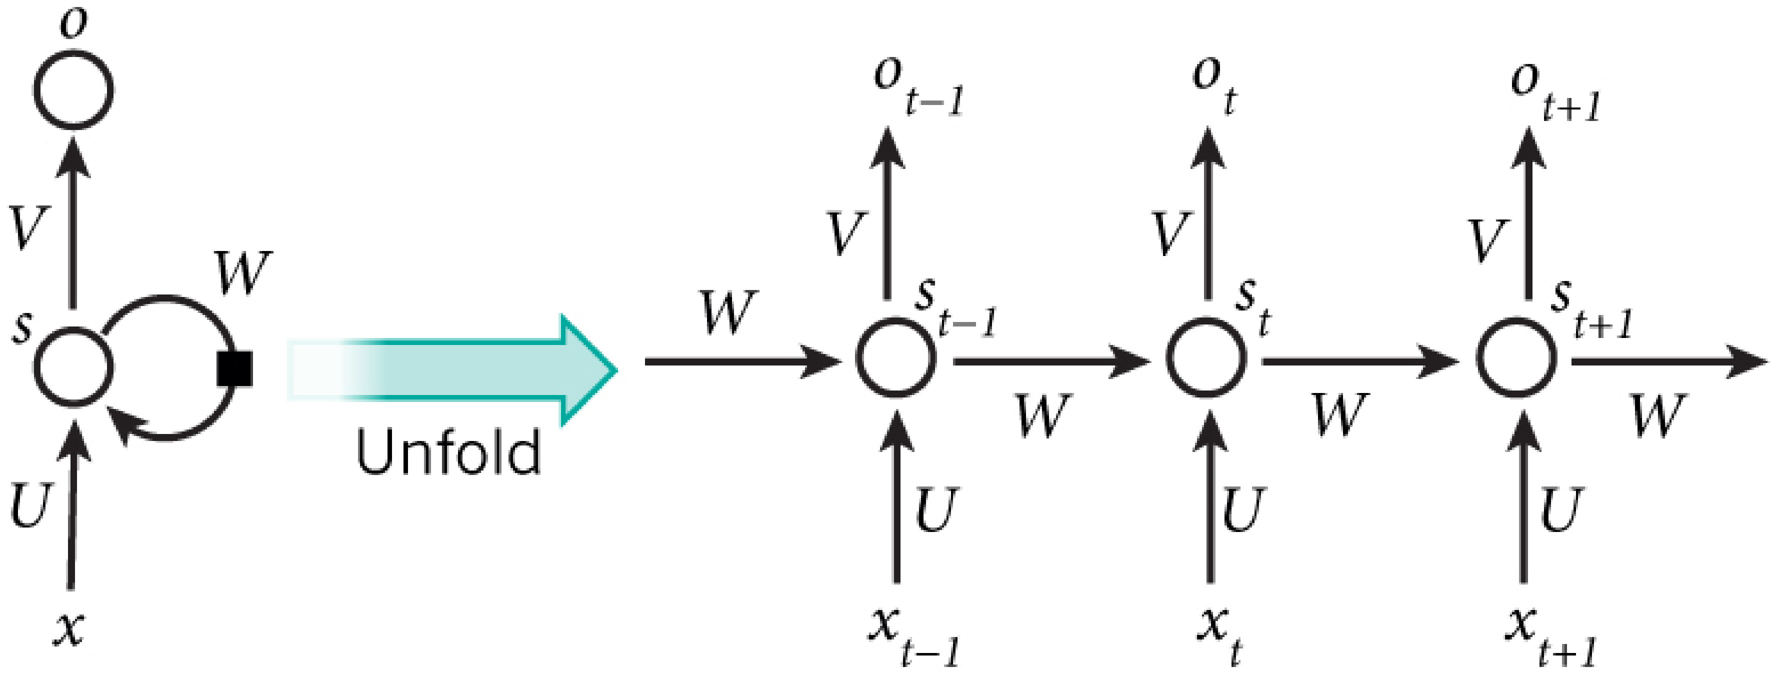

Setting the Stage for Biomimetics

- We are not working w/ the “right” building blocks
  - Backprop is easy and rather fast to use
  - Recurrence is easier to eschew (“time windows are all you need”)
  - Pointwise neurons are easy and rather fast to use
  - Lots of nice innovations / applications

- There are (possibly insurmountable) roadblocks
  - Energy inefficiency, sample inefficiency
  - Catastrophic forgetting
  - Online / real-time learning
  - Sparse reward/signal reinforcement learning

- We have not tapped into full value of what nature and biological neural computation / credit assignment bring to the table

Neurobiological Credit Assignment: A Taxonomy of Frameworks

One way to view and organize biological credit assignment

A Galaxy of Neural Credit Assignment Processes

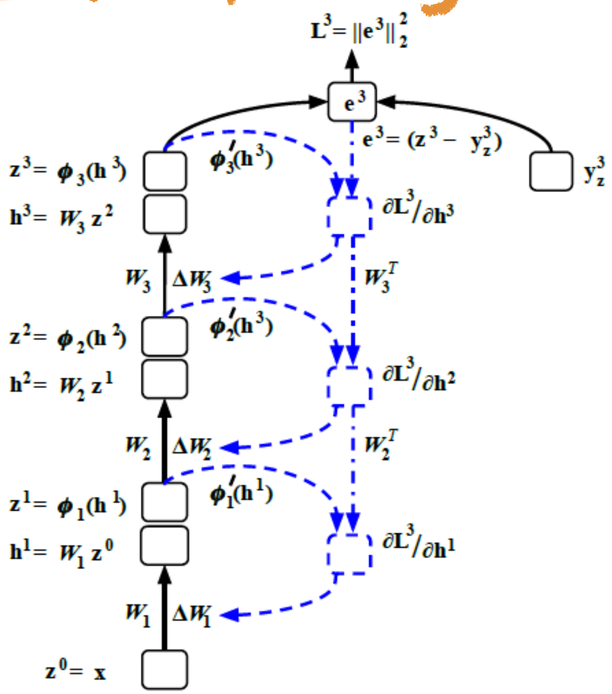

Backprop

A Galaxy of Neural Credit Assignment Processes

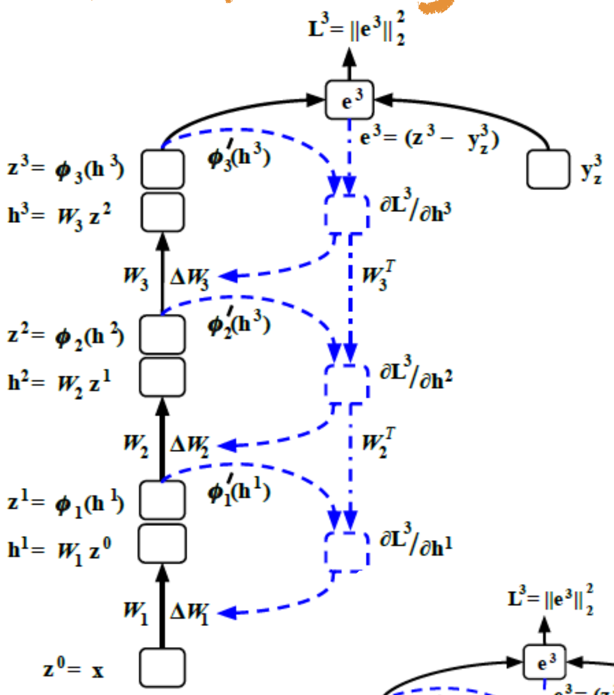

Backprop

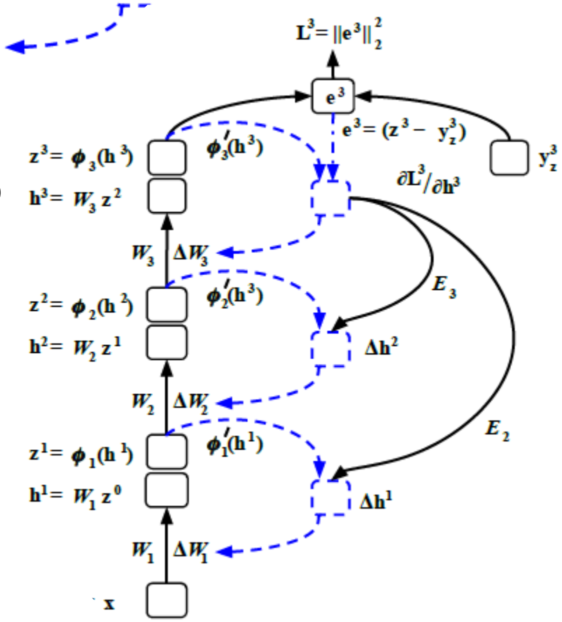

Feedback alignment

(Lillicrap, 2016; Nøkland, 2016)

A Galaxy of Neural Credit Assignment Processes

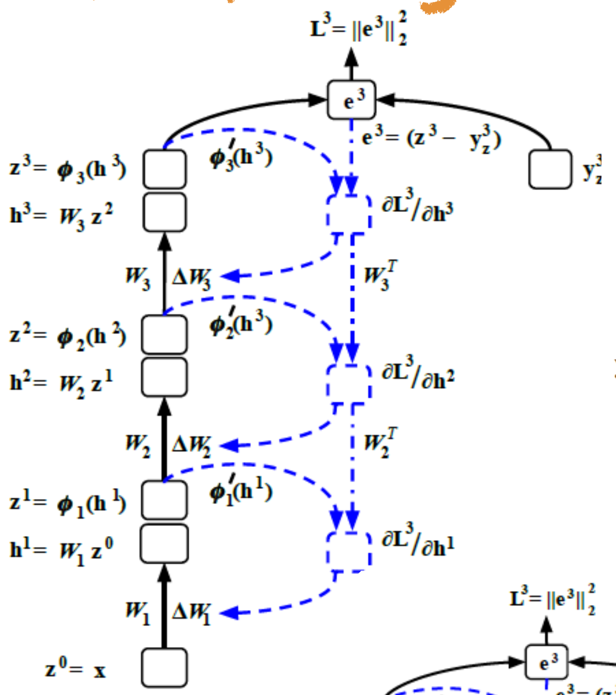

Backprop

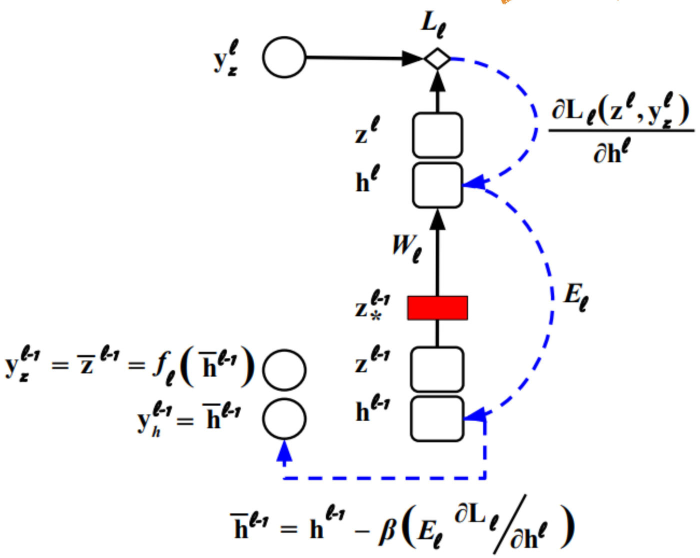

Local Representation Alignment
(Ororbia & Mali, 2019; Ororbia et al., 2023)

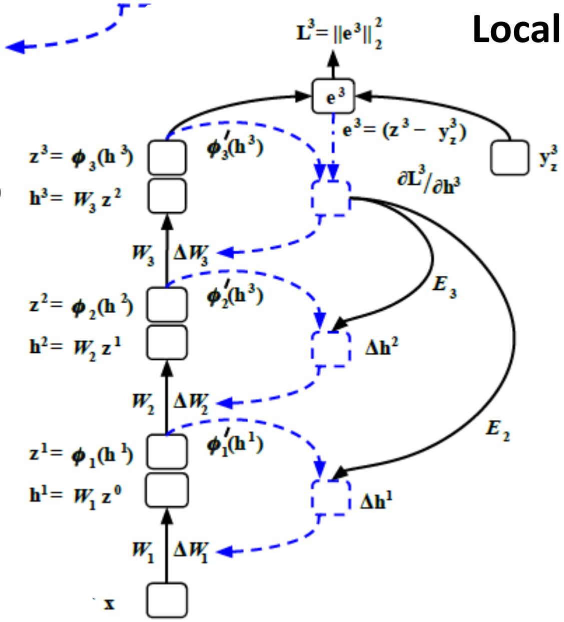

Feedback alignment
(Lillicrap, 2016; Nøkland, 2016)

A Galaxy of Neural Credit Assignment Processes

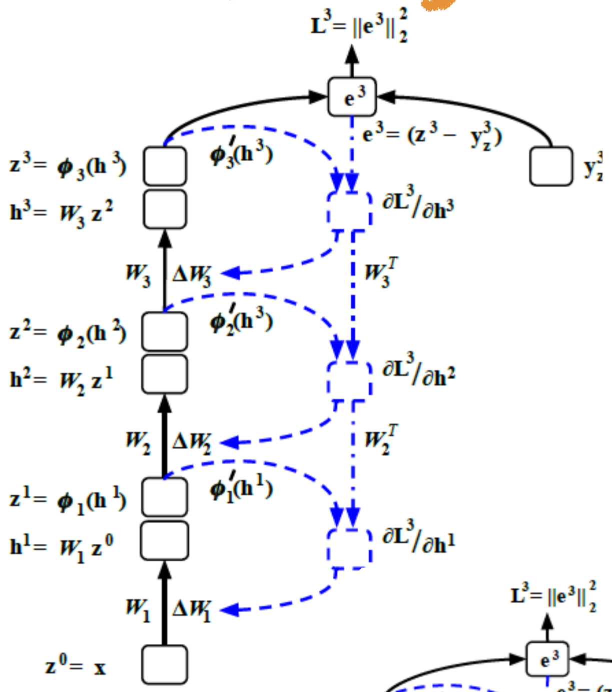

Backprop

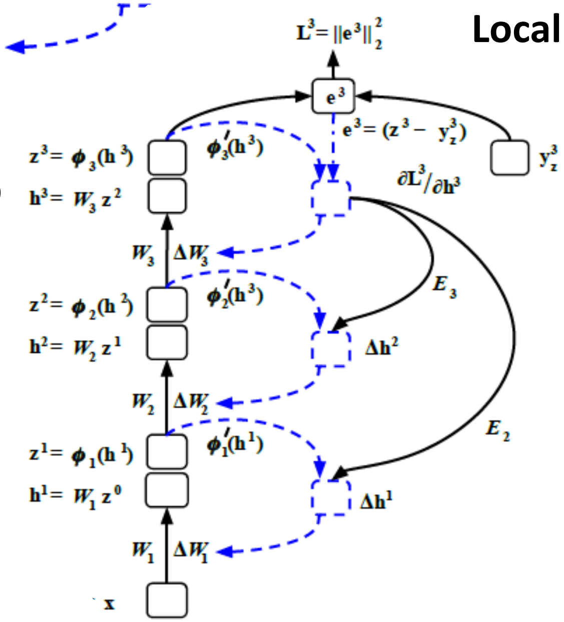

Feedback alignment

(Lillicrap, 2016; Nøkland, 2016)

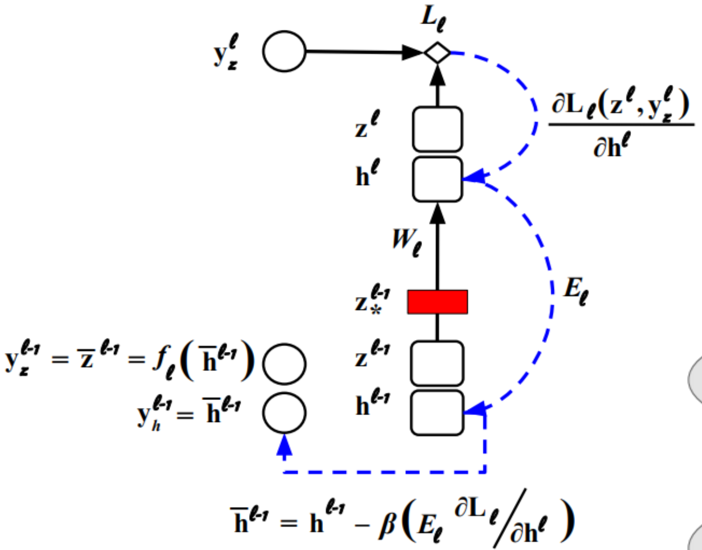

Local Representation Alignment

(Ororbia & Mali, 2019; Ororbia et al., 2023)

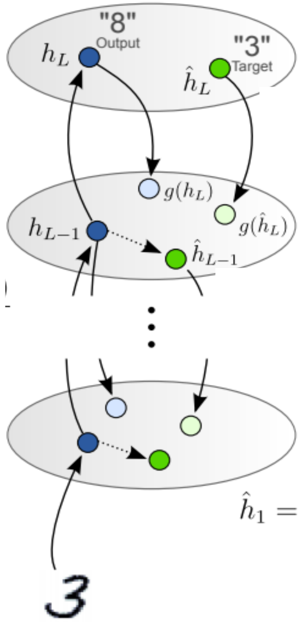

Target Propagation

(Bengio, 1993)

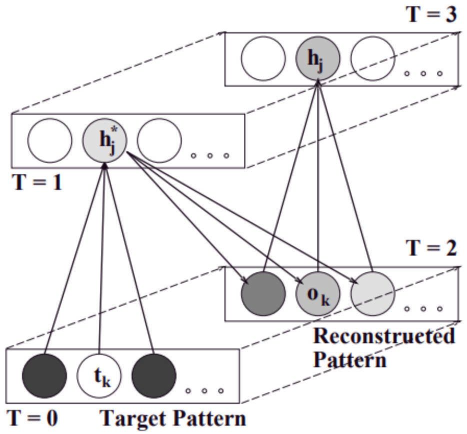

Recirculation

(Hinton, 1987; O'Reilly, 1996)

A Galaxy of Neural Credit Assignment Processes

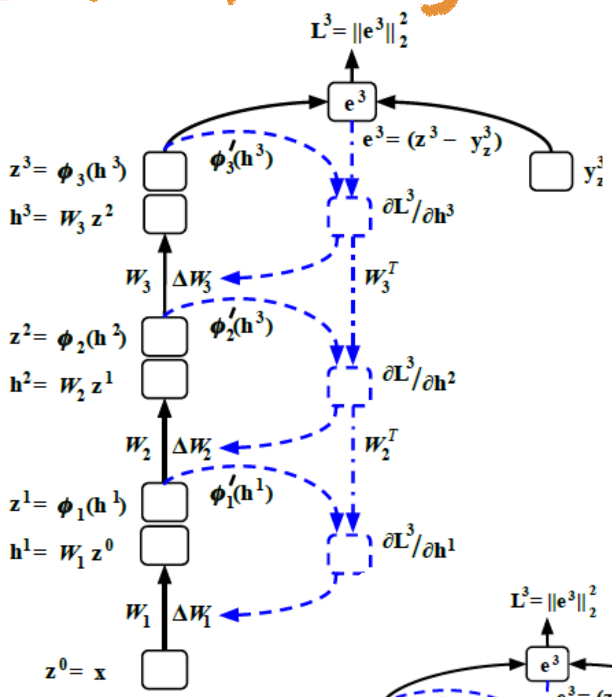

Backprop

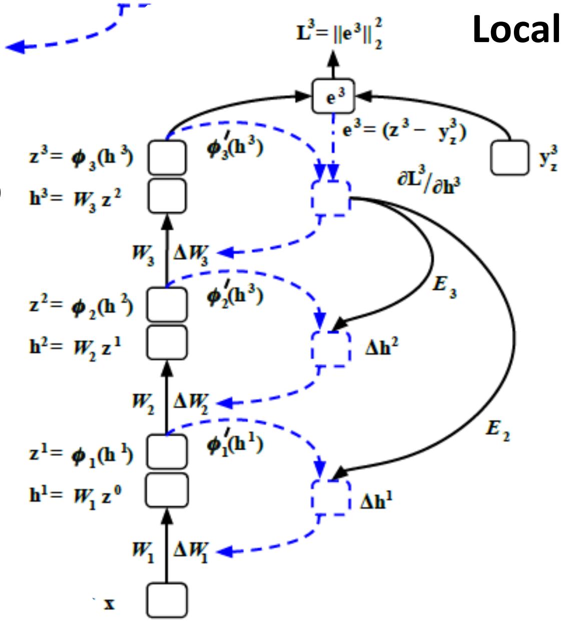

Feedback alignment

(Lillicrap, 2016; Nøkland, 2016)

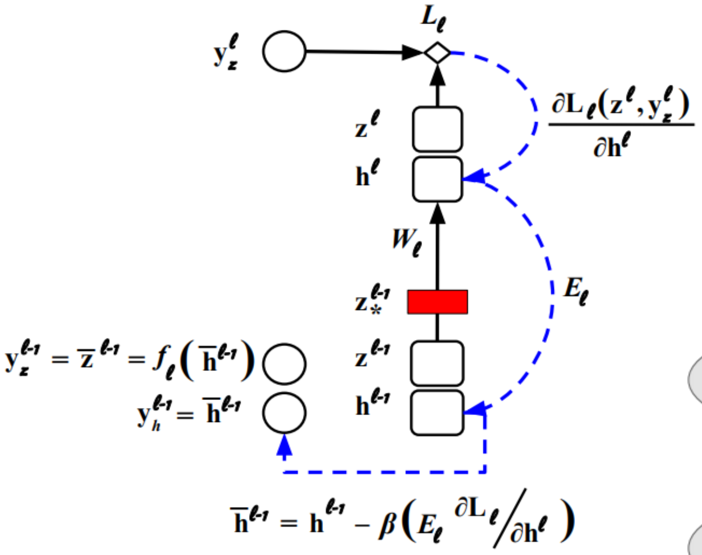

Local Representation Alignment

(Ororbia & Mali, 2019; Ororbia et al., 2023)

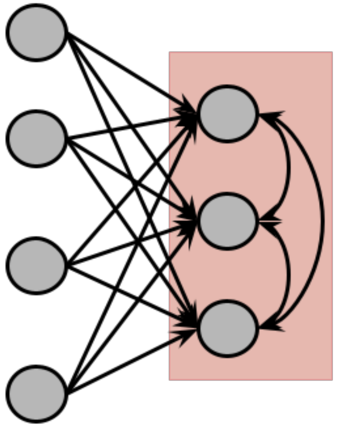

Competitive Hebbian Learning

(Kohonen, 1982; Grossberg 1987; Martinetz, 1993)

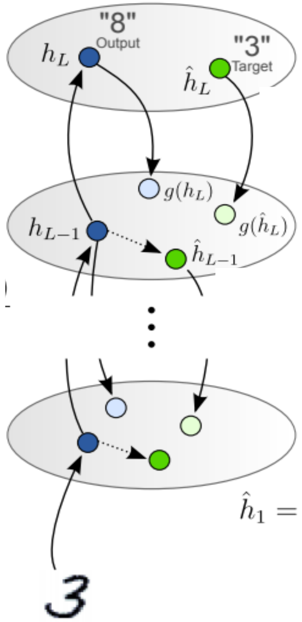

Target Propagation

(Bengio, 1993)

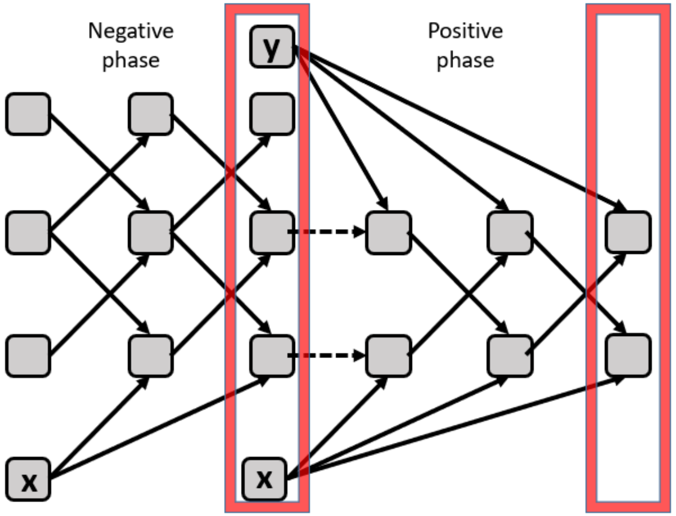

Contrastive Hebbian Learning / Divergence, Wake-Sleep, Equilibrium Propagation

(Movellan, 1991; Hinton, 1995 & 2002; Scellier, et al., 2017)

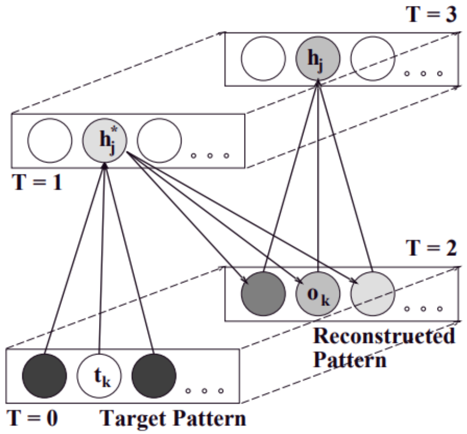

Recirculation

(Hinton, 1987; O'Reilly, 1996)

A Galaxy of Neural Credit Assignment Processes

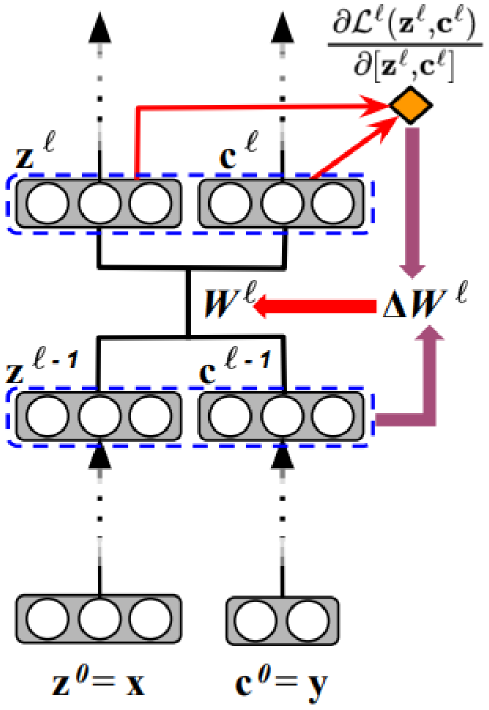

Forward-Only Learning

(Kohan et al., 2022; Hinton 2022; Ororbia & Mali, 2022; Ororbia 2023)

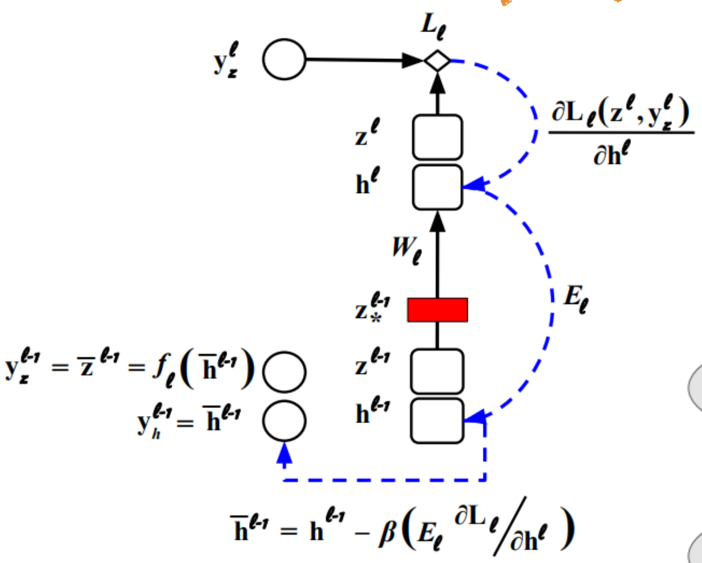

Local Representation Alignment

(Ororbia & Mali, 2019; Ororbia et al., 2023)

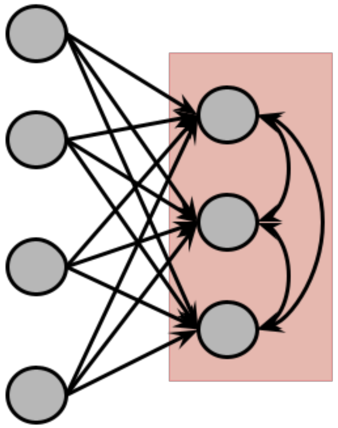

Competitive Hebbian Learning

(Kohonen, 1982; Grossberg 1987; Martinetz, 1993)

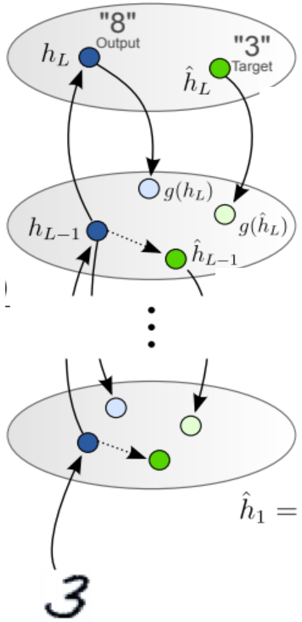

Target Propagation

(Bengio, 1993)

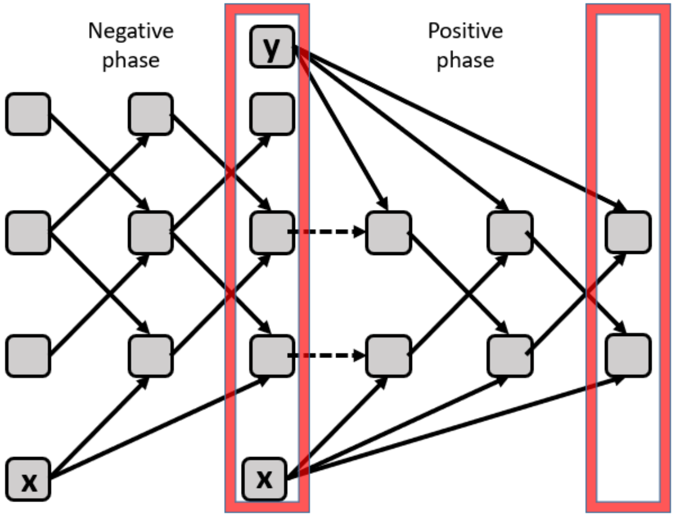

Contrastive Hebbian Learning / Divergence, Wake-Sleep, Equilibrium Propagation

(Movellan, 1991; Hinton, 1995 & 2002; Scellier, et al., 2017)

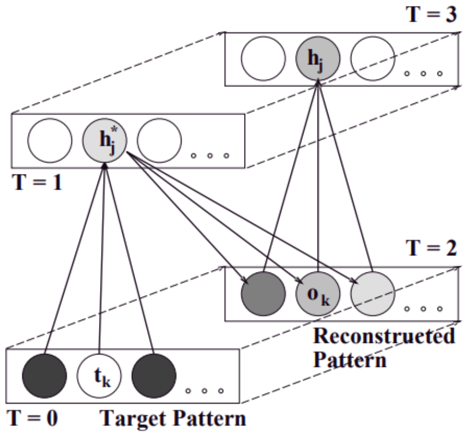

Recirculation

(Hinton, 1987; O'Reilly, 1996)

Okay, let's take a step back and organize this a bit…

Ororbia, A. G. "Brain-Inspired Machine Intelligence: A Survey of Neurobiologically-Plausible Credit Assignment." (2023).

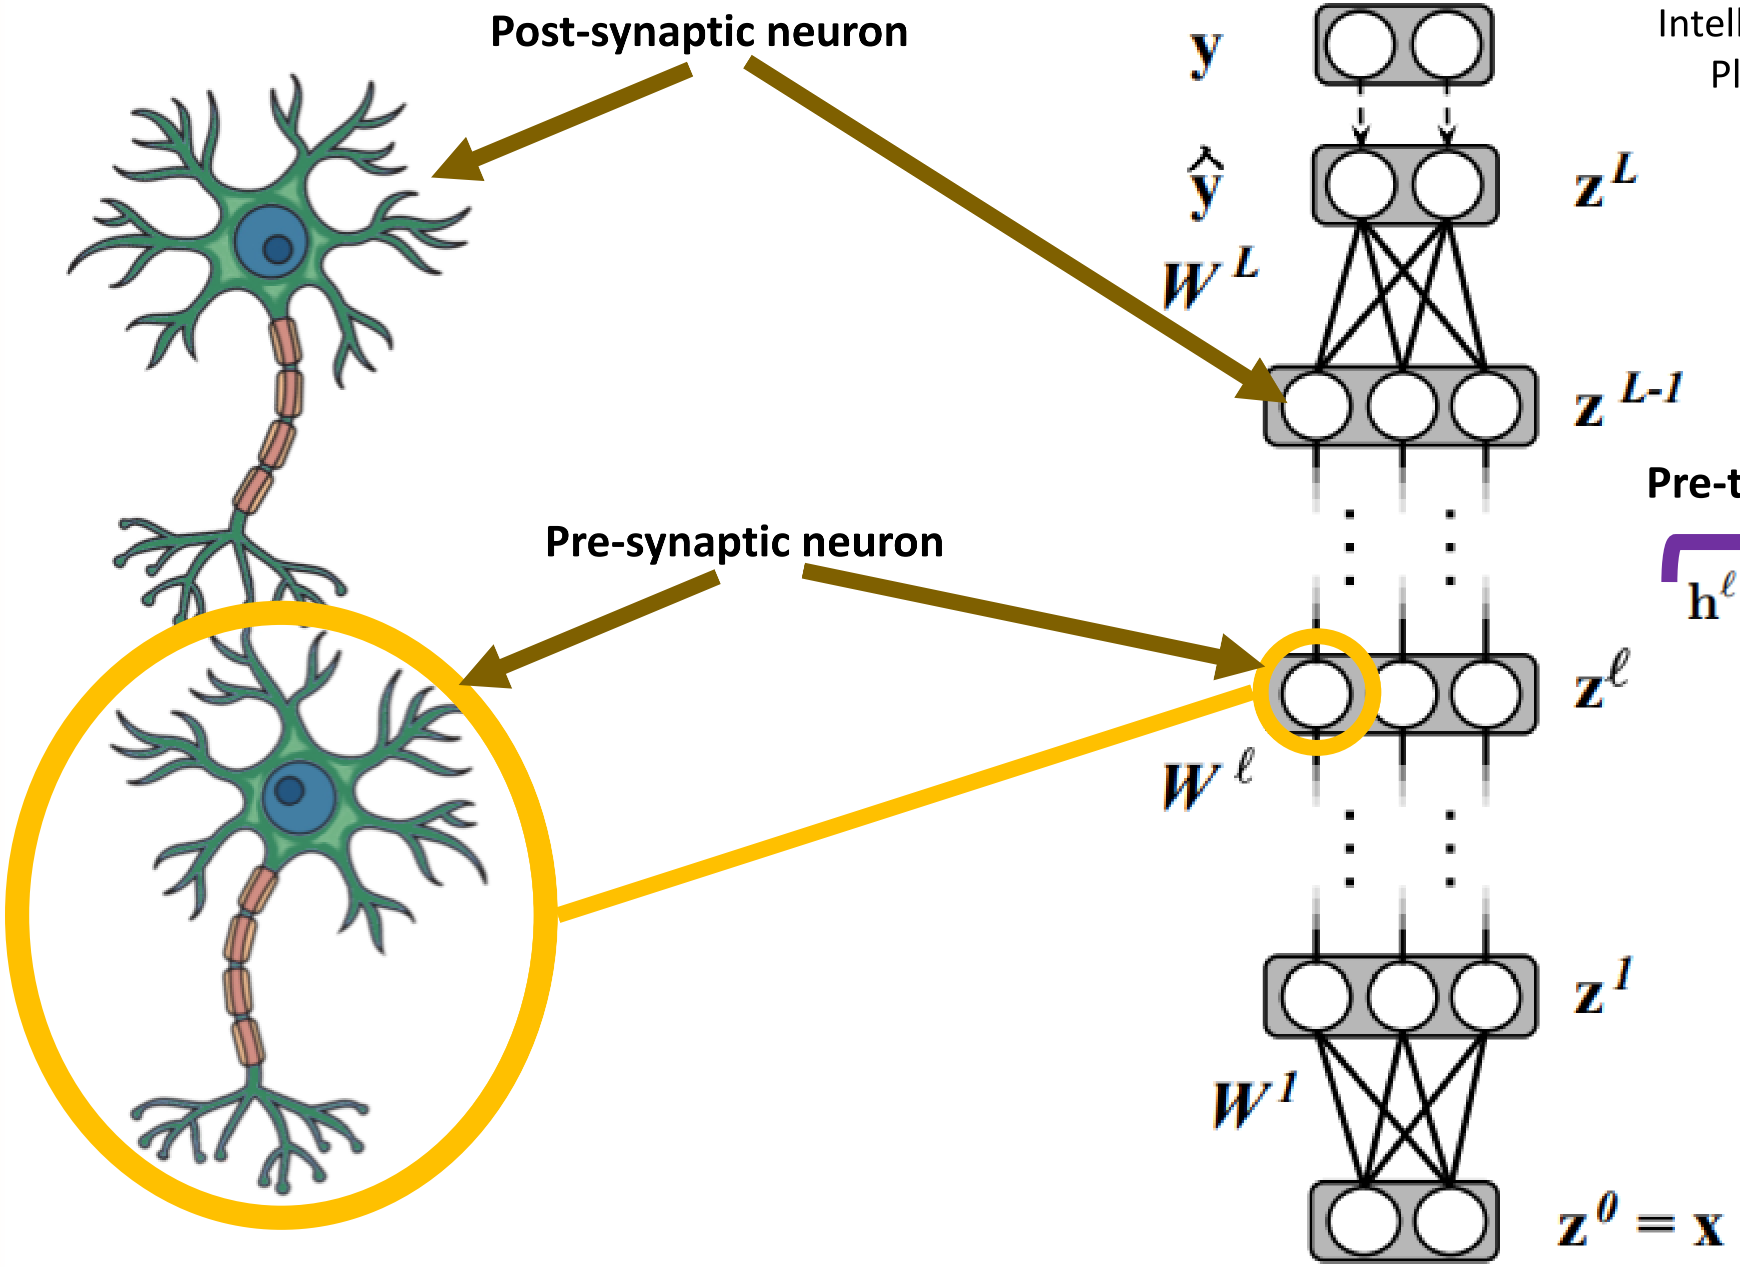

$$\overbrace{\mathbf{h}^\ell = \mathbf{W}^\ell \cdot \mathbf{z}^{\ell-1}}^{\text{Pre-transformation}}, \quad \underbrace{\mathbf{z}^\ell = \phi^\ell(\mathbf{h}^\ell)}_{\text{Post-transformation}}$$

Ororbia, A. G., et al. "A review of neuroscience-inspired machine learning." (2024).

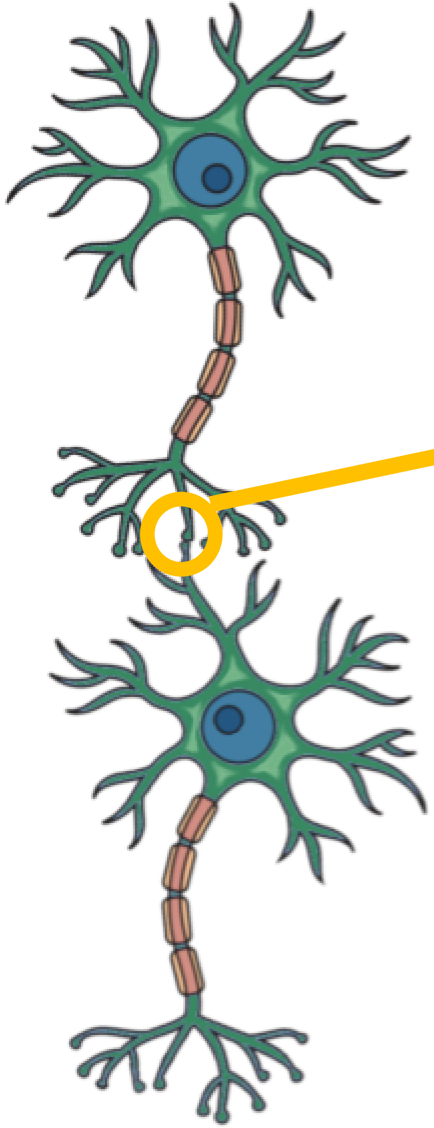

w_{ij}

Synapse (or synaptic juncture) relates transmission between pre-synaptic neuron *j* to post-synaptic neuron *i*

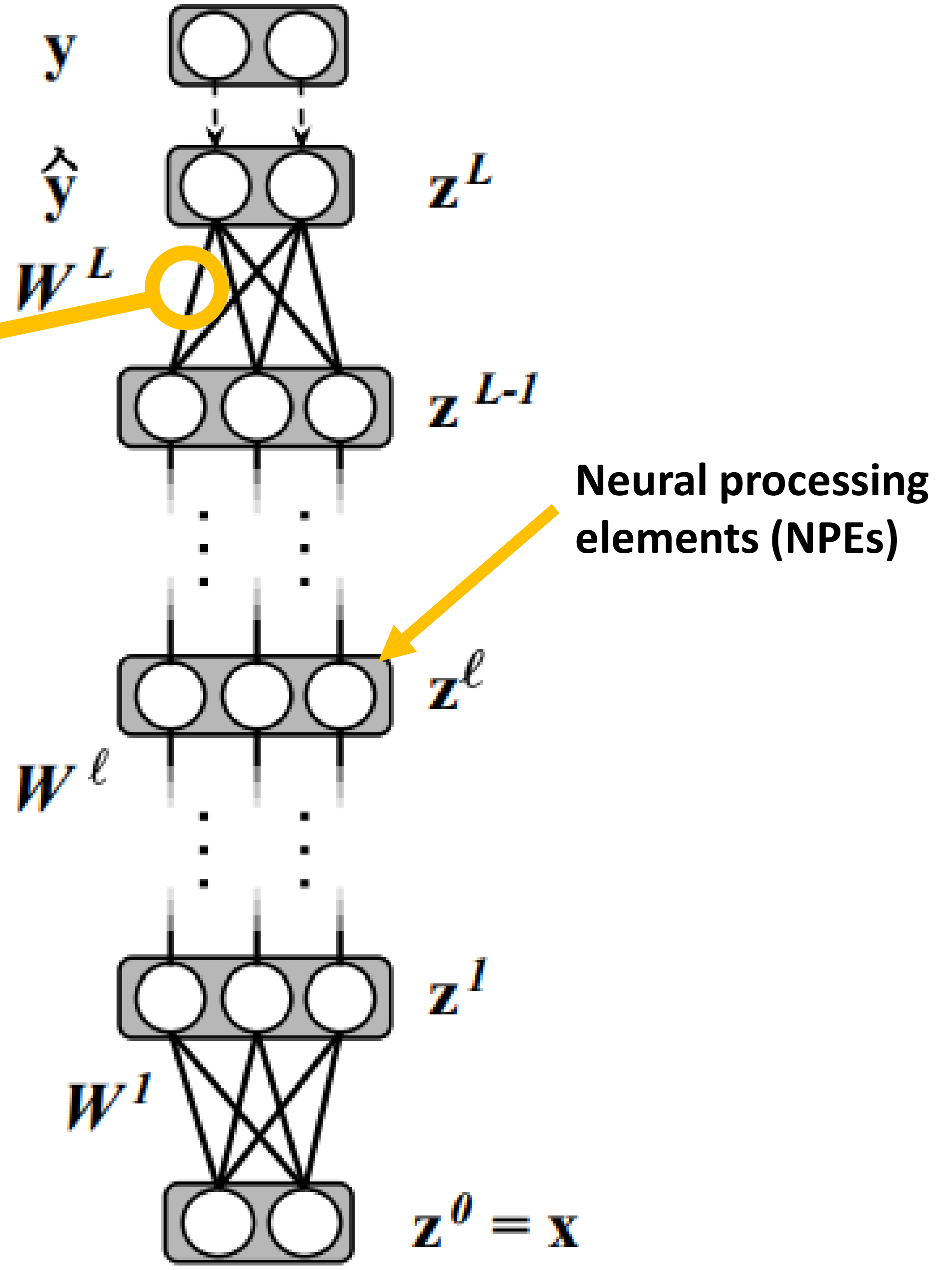

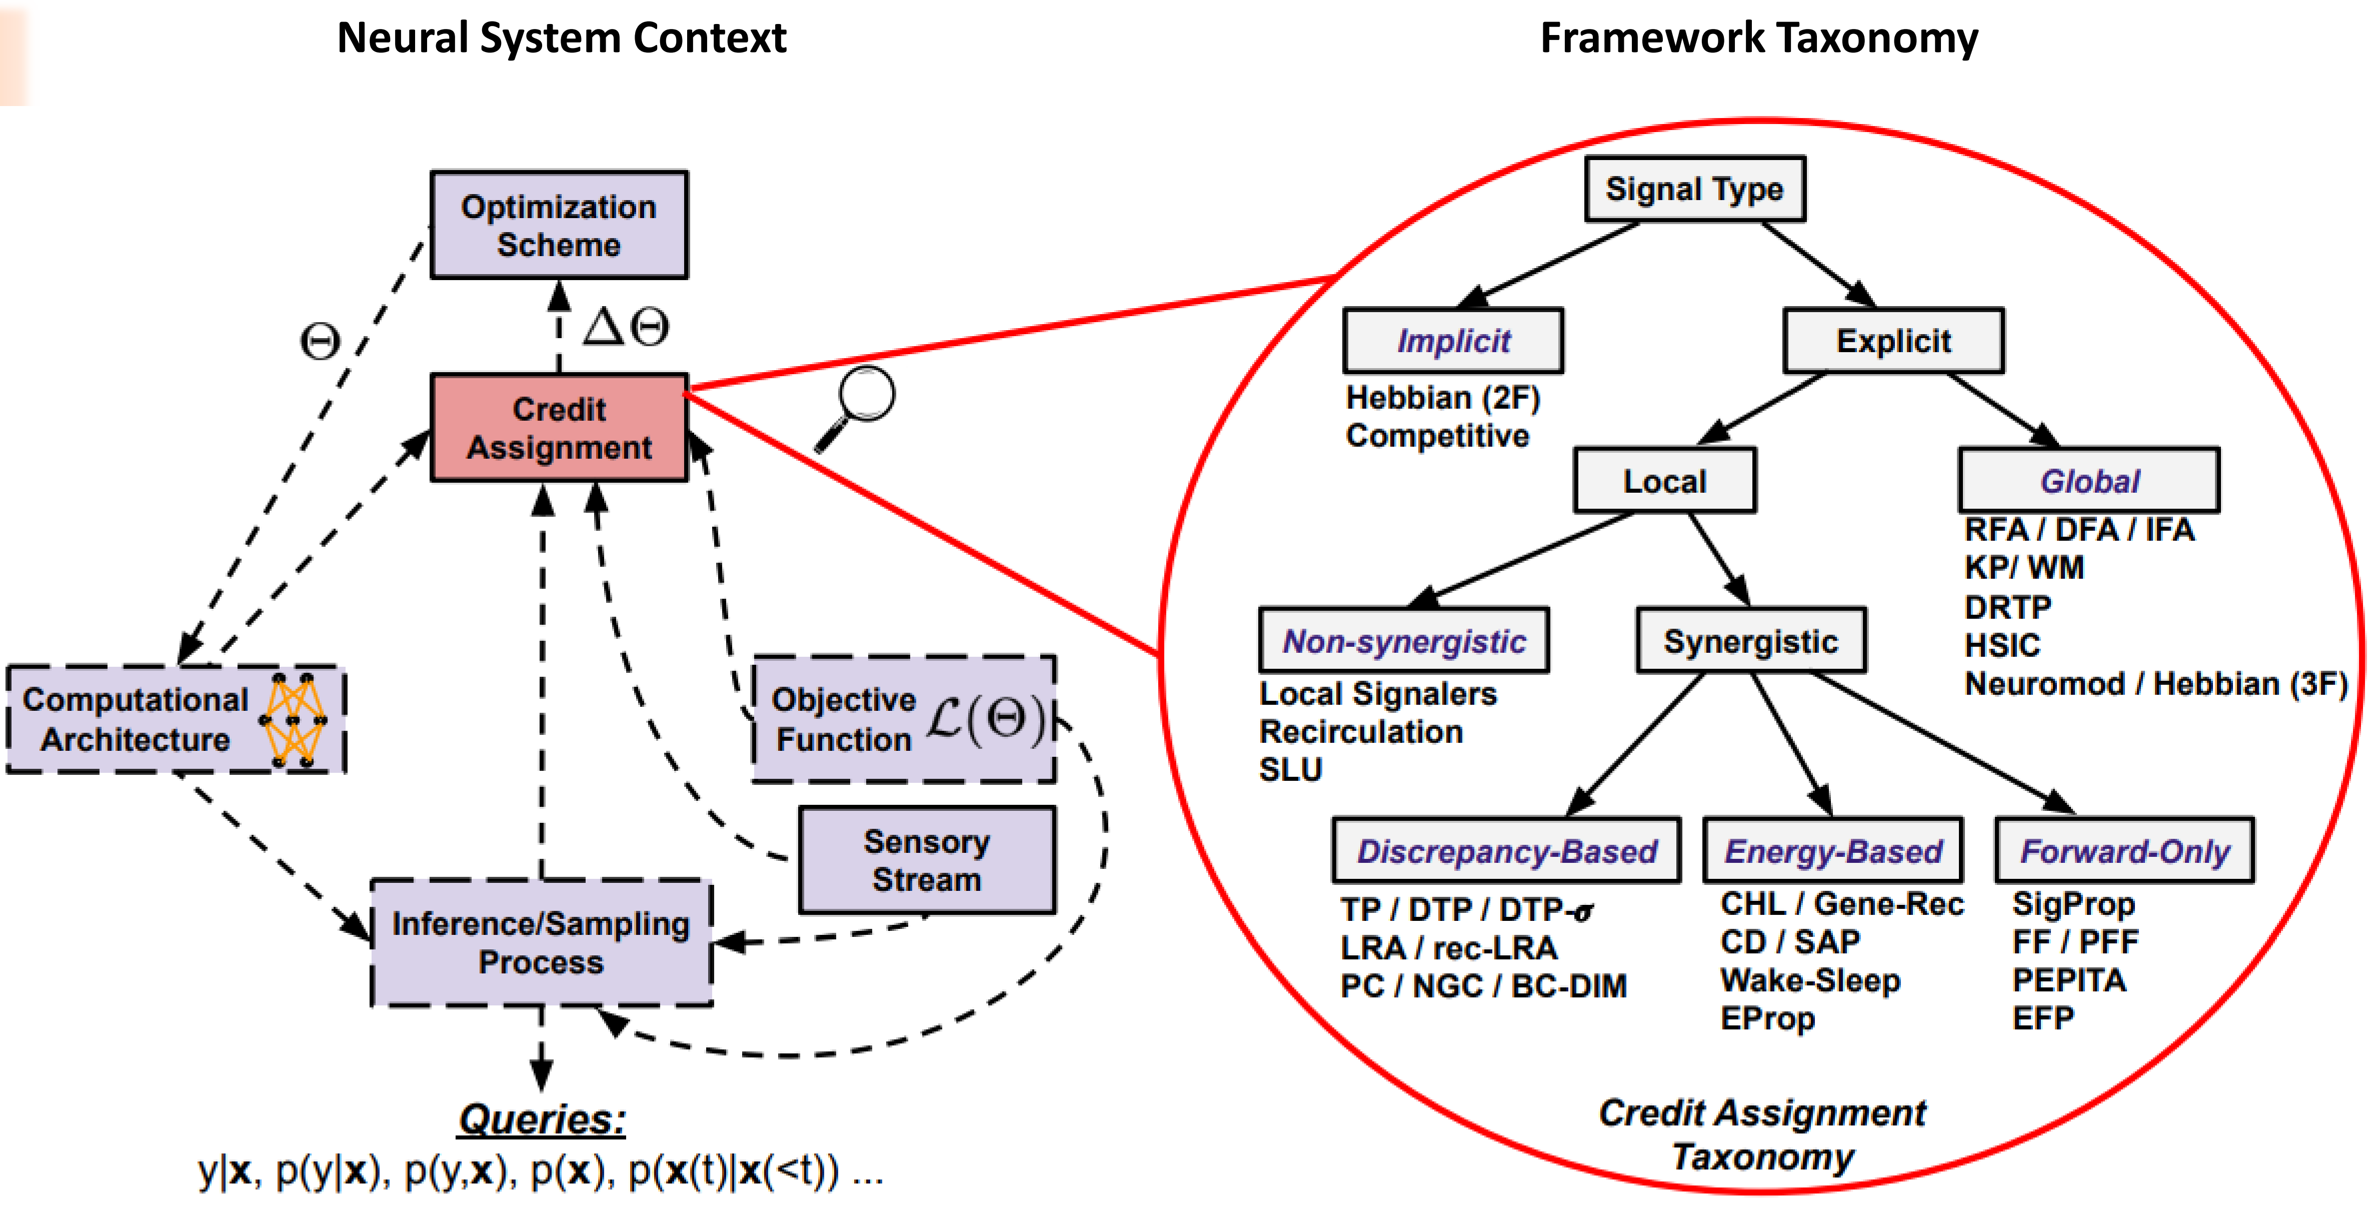

Ororbia, A. G. "Brain-Inspired Machine Intelligence: A Survey of Neurobiologically-Plausible Credit Assignment." (2023).

Let's Zoom into the One of These Families: Discrepancy-Based Learning

This is where predictive coding lives!

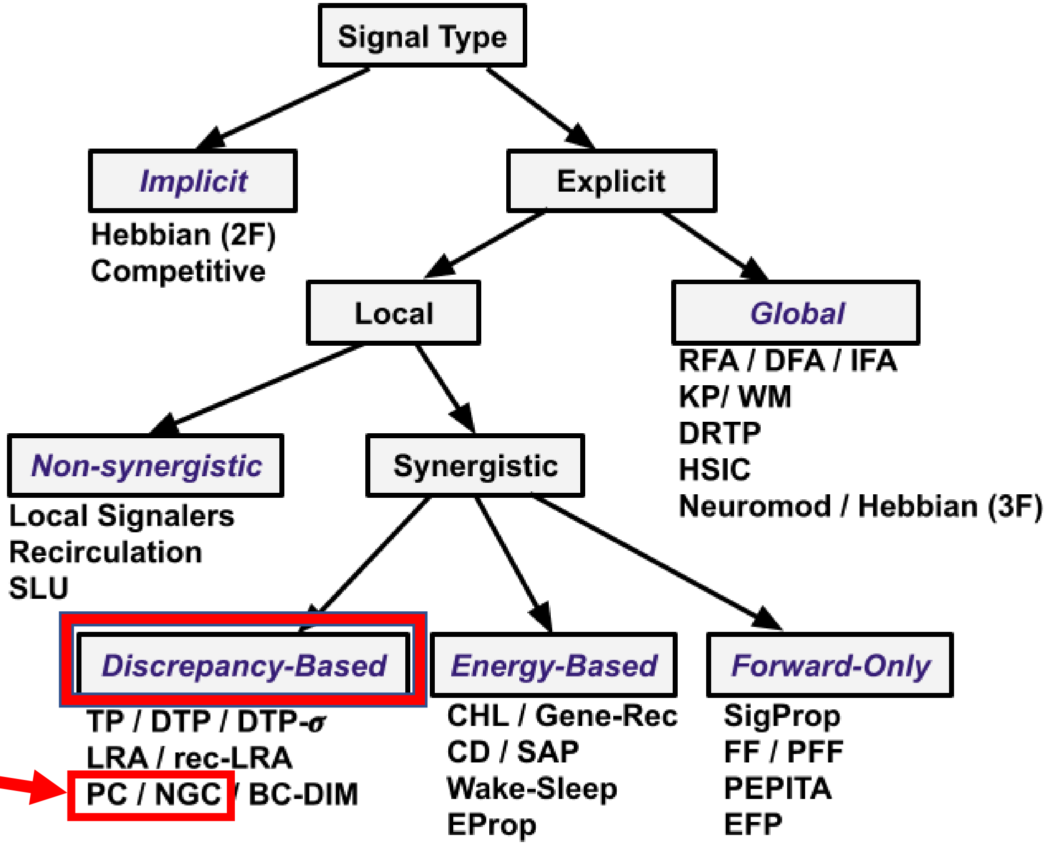

Predictive coding (= a "Synergistic Local Explicit Algorithm")

Neurobiological Motivations

- *Predictive coding*: brain generates hypotheses, adjusts/corrects based on data
  - Sparsity through *sparse coding*
  - Brain = probabilistic, hierarchical / heterarchical

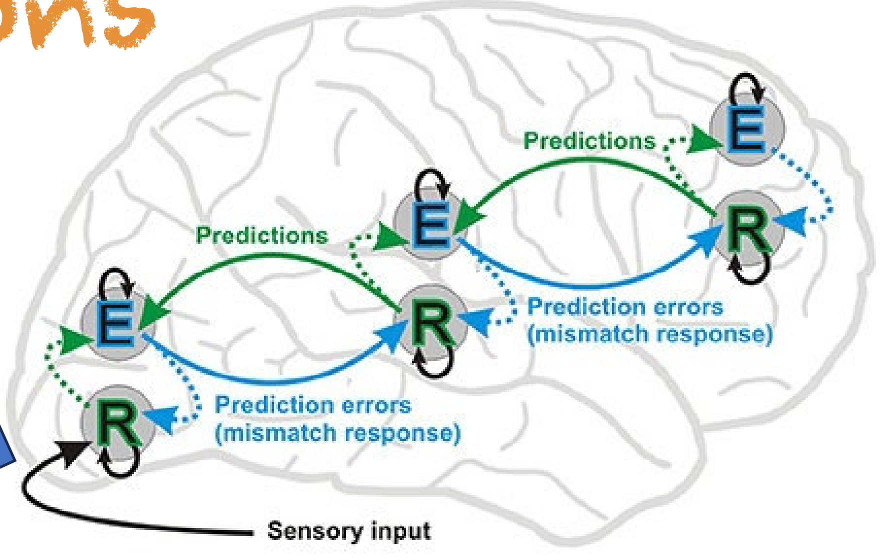

(Image taken from Stefanics et al., 2014.)

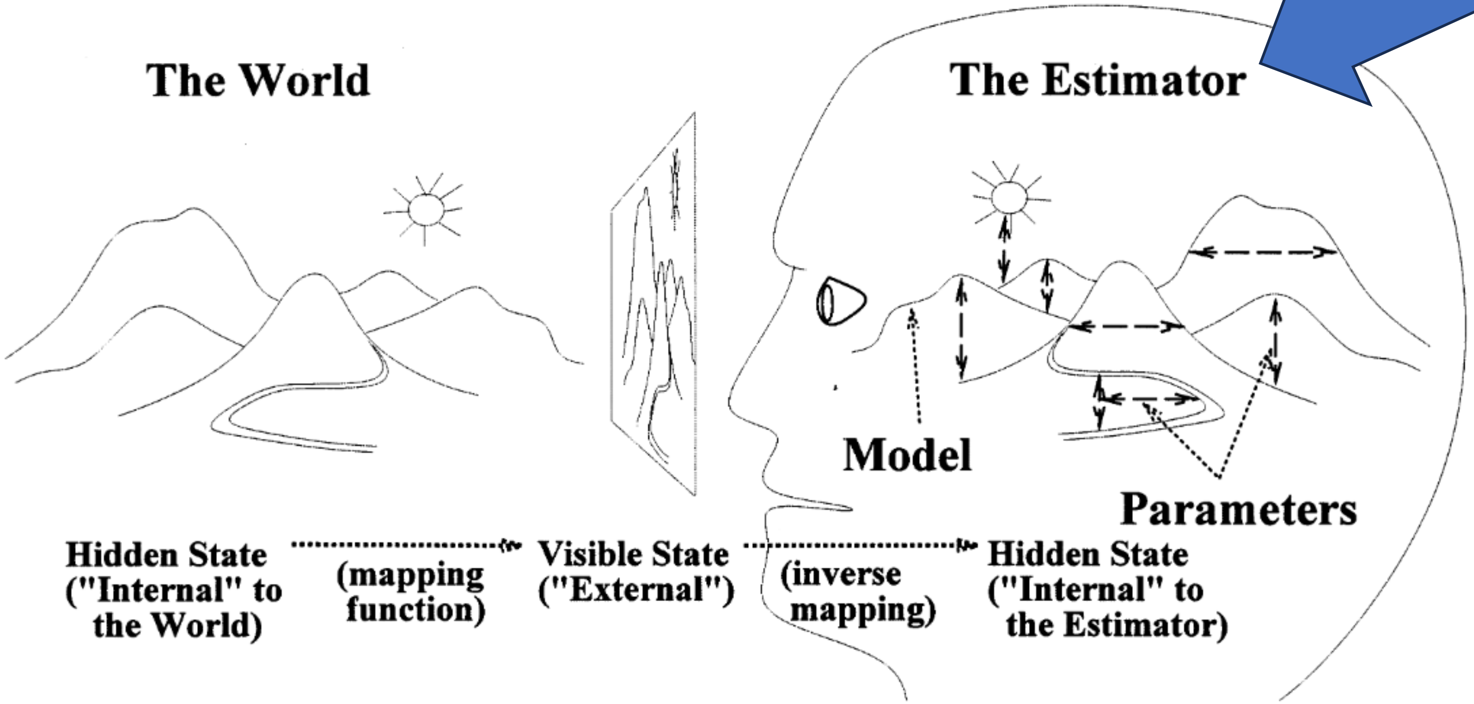

Rao, RPN, Ballard, DH. "Dynamic model of visual recognition predicts neural response properties in the visual cortex." 1997.

Olshausen, BA., Field, DJ. "Emergence of simple-cell receptive field properties by learning a sparse code for natural images." 1996.

Von Helmholtz, Hermann. *Treatise on physiological optics*. 1867.

Neurobiological Motivations

- *Predictive coding*: brain generates hypotheses, adjusts/corrects based on data
  - Sparsity through *sparse coding*
  - Brain = probabilistic, hierarchical / heterarchical

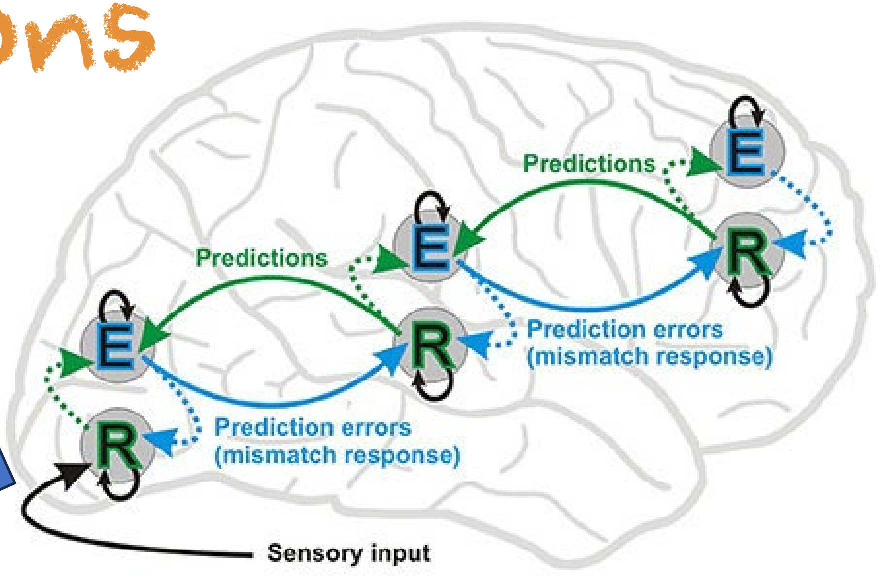

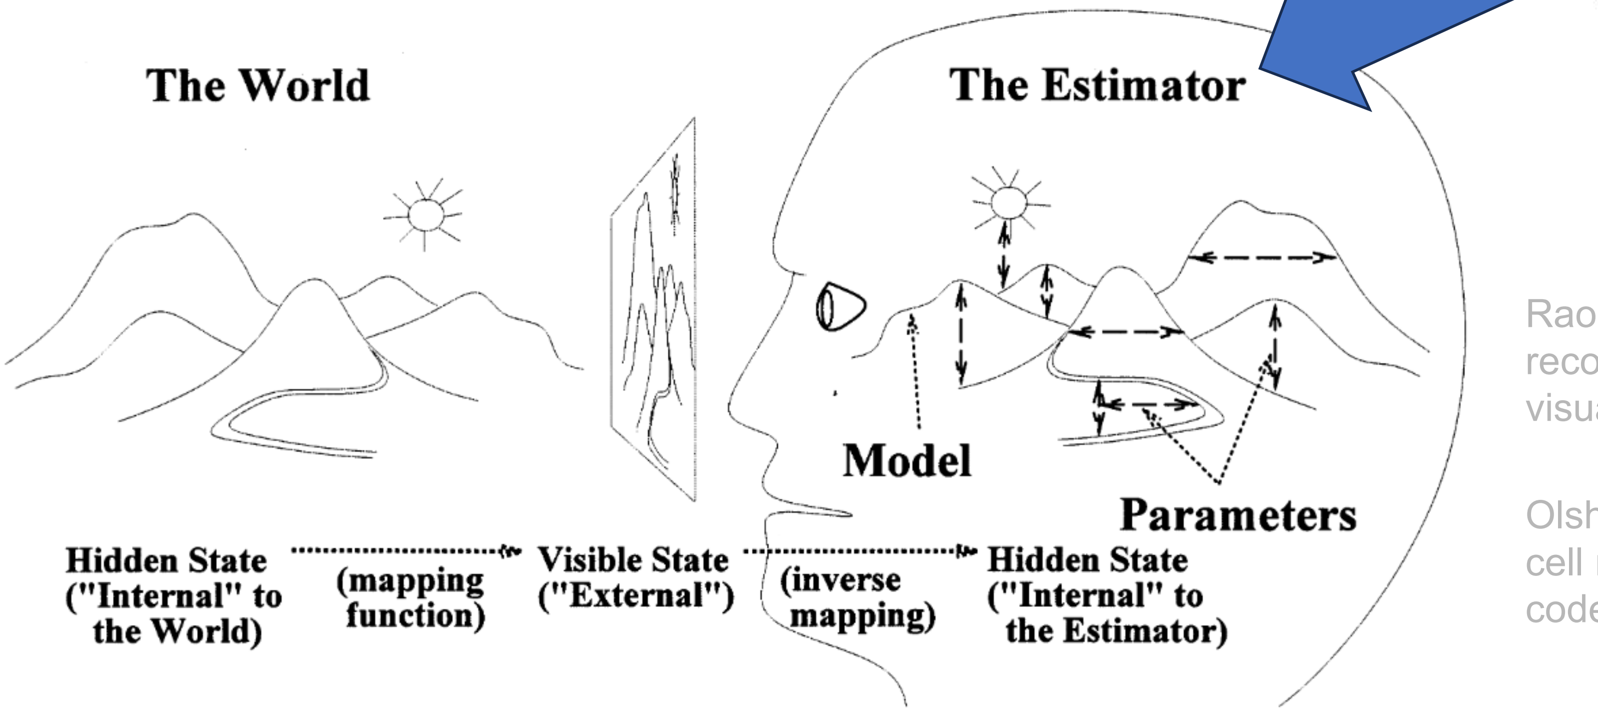

Rao, RPN, Ballard, DH. "Dynamic model of visual recognition predicts neural response properties in the visual cortex." 1997.

Olshausen, BA., Field, DJ. "Emergence of simple-cell receptive field properties by learning a sparse code for natural images." 1996.

Von Helmholtz, Hermann. *Treatise on physiological optics*. 1867.

Mechanics of Predictive Coding

- We take a neuronal dynamics approach to inference and learning
  - Everything is inherently and naturally temporal
  - Passes/aggregates bottom-up & top-down signals
  - Constantly generating hypotheses & adjusting based on data samples

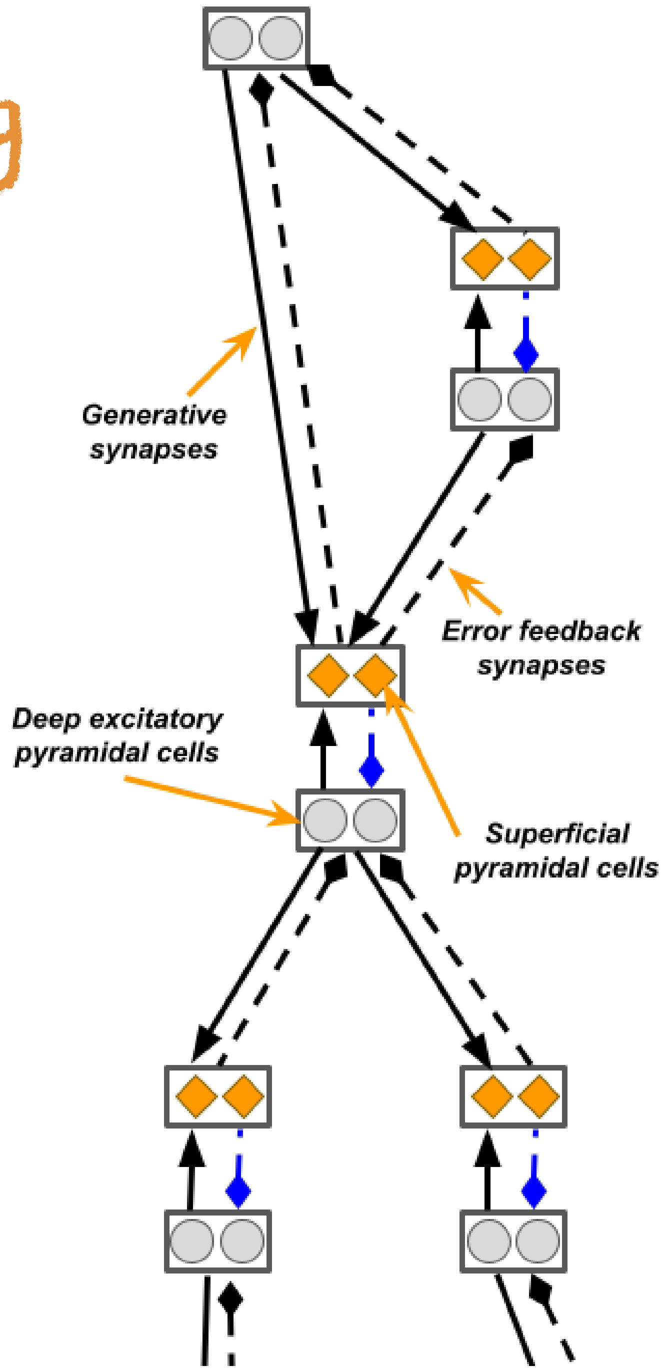

Rao, RPN, Ballard D. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. 1999.

Friston, K. "A theory of cortical responses." 2005

Ororbia, AG, Kifer D. "The neural coding framework for learning generative models.“ 2022.

Step 1:
Hypothesis Generation

A (Variational) Free Energy:

$$\mathcal{F}(\Theta) = \sum_{\ell=0}^{L} \frac{1}{2\Sigma^{\ell}} \sum_{i=1}^{\mathcal{J}_{\ell}} (\mathbf{z}_i^{\ell}(t) - \bar{\mathbf{z}}_i^{\ell})^2$$

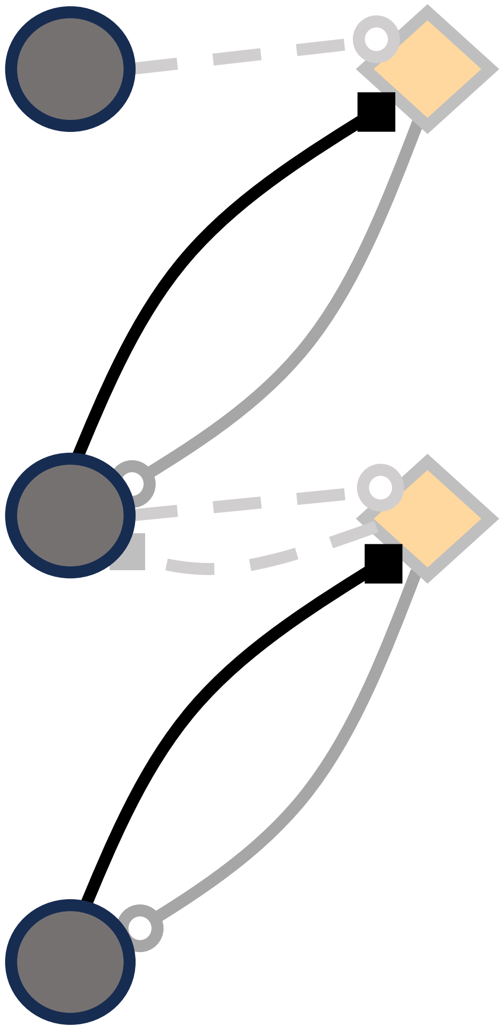

Inhibitory carry-through synapse

Inhibitory synapse

Excitatory synapse

Excitatory carry-through synapse

Step 2:
Mismatch/Error Computation

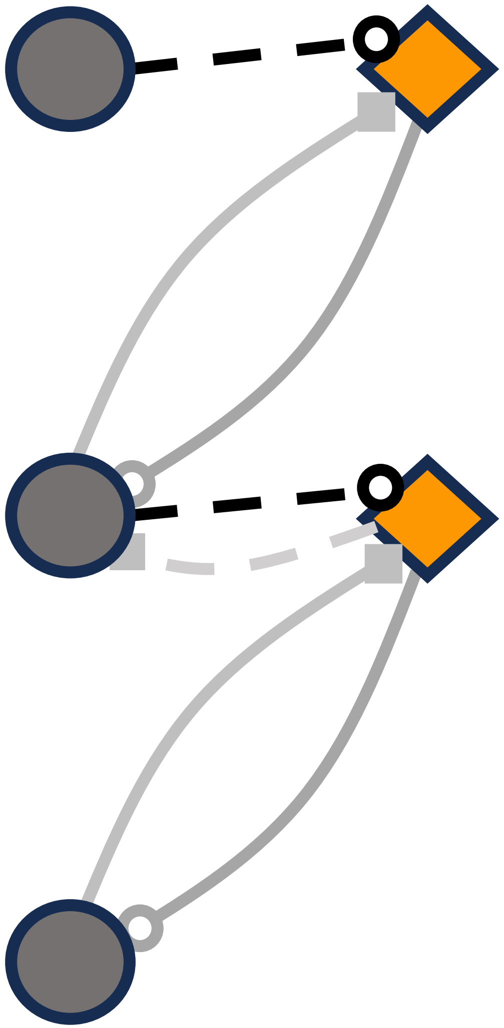

- - - ■ Inhibitory carry-through synapse

——■ Inhibitory synapse

——o Excitatory synapse

- - - o Excitatory carry-through synapse

Step 3:
State
Correction

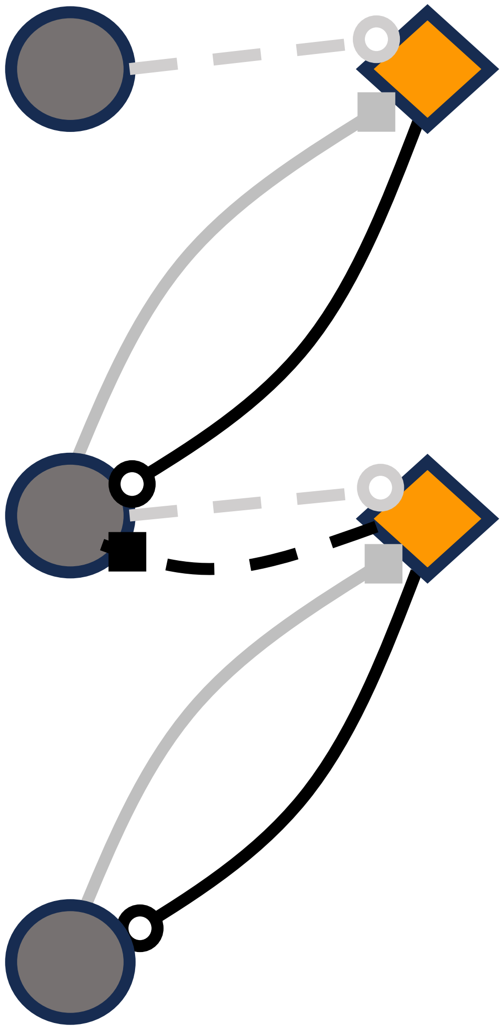

Inhibitory carry-through synapse
Inhibitory synapse
Excitatory synapse
Excitatory carry-through synapse

Go back to *Step 1*

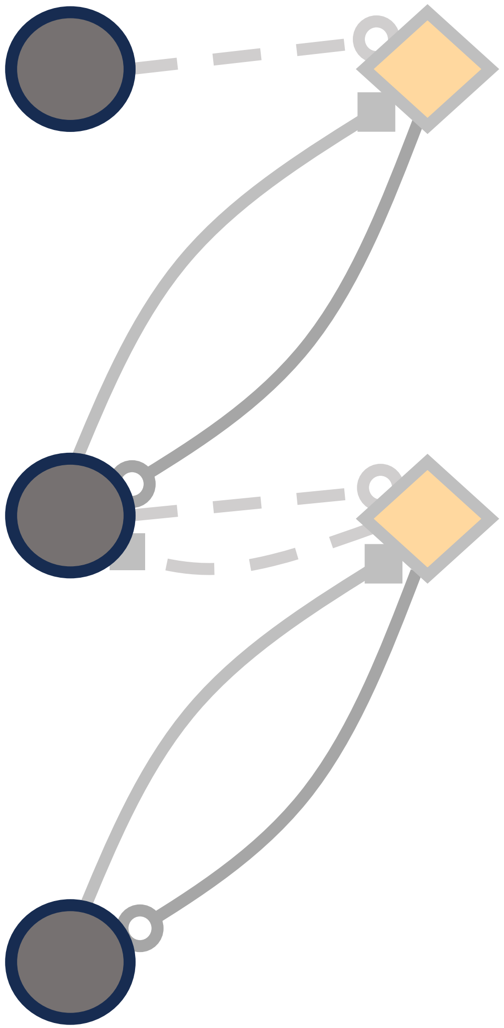

This generate-then-correct process is repeated over a stimulus window of length T

(Conducts bounded iterative inference to converge to stabler state that pleases mapping between input & output signals)

Go back to Step 1

OR

***Update* synaptic efficacies**

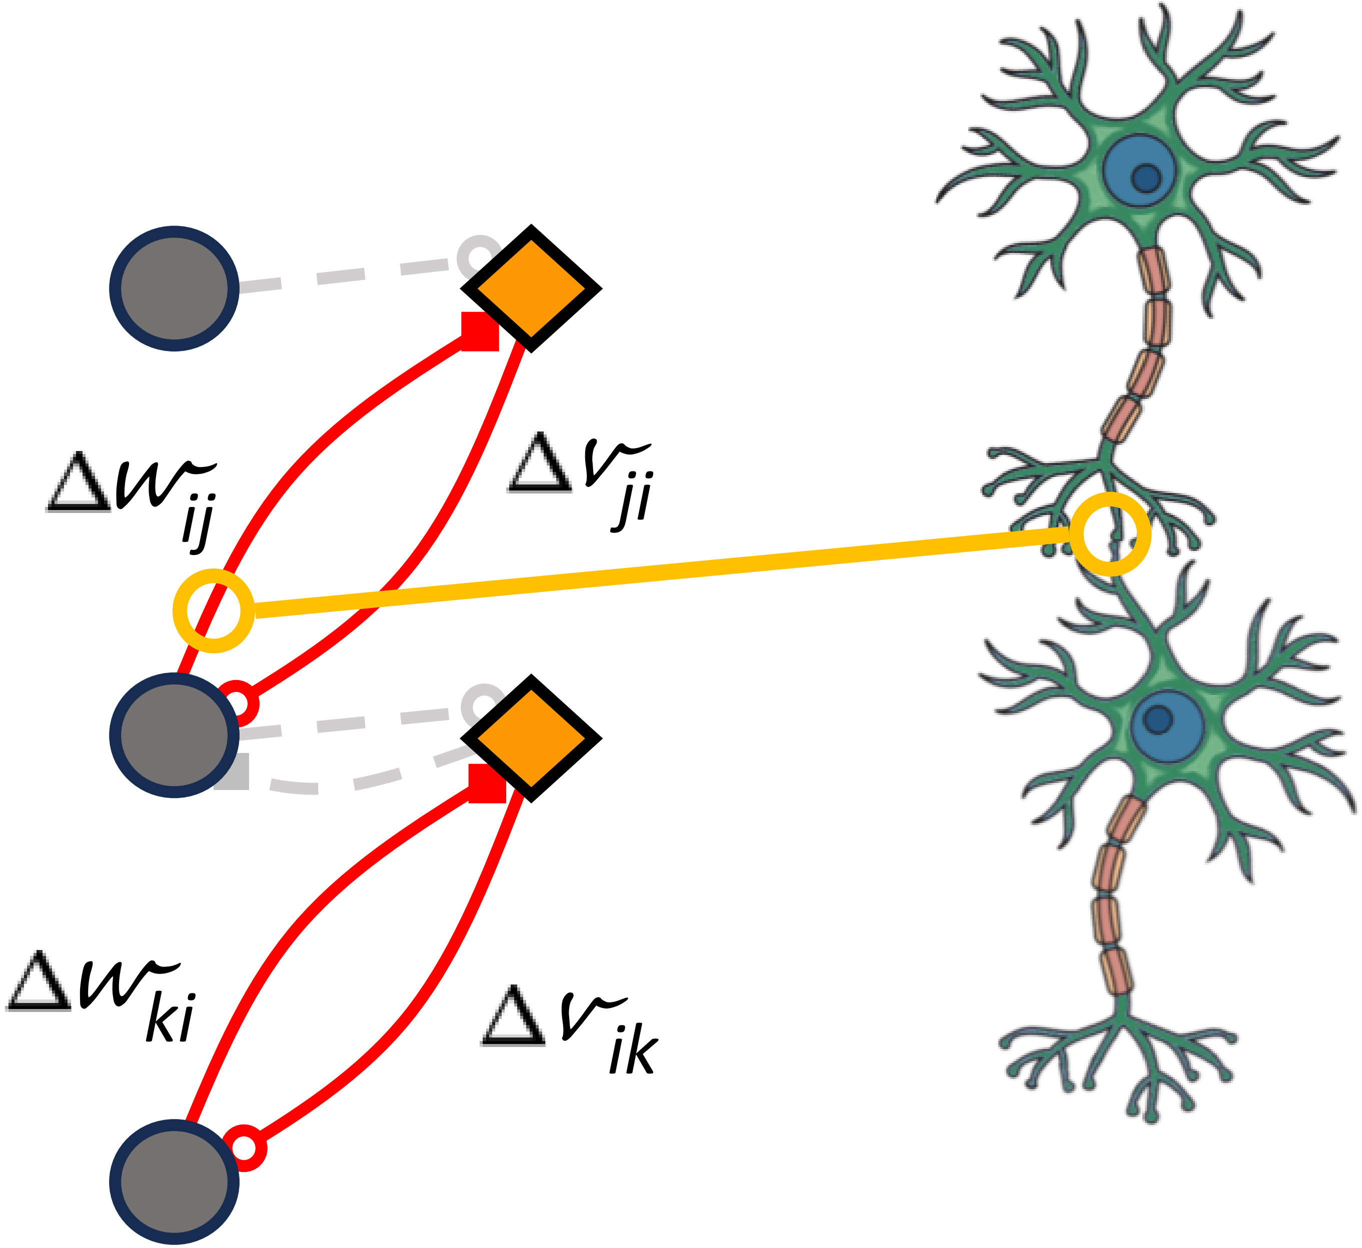

If synaptic adjustments are scheduled every *T* time steps, you get dynamic expectation maximization (*DEM*)

Friston, K.J. "DEM: a variational treatment of dynamic systems." 2008.

What does this framing buy you?

- Optimize lower bound of model evidence (marginal likelihood) = *variational free energy*
- It generalizes to more complex, flexible architectures, e.g., directed graphical models and networks with cycles that resemble ***brain regions***

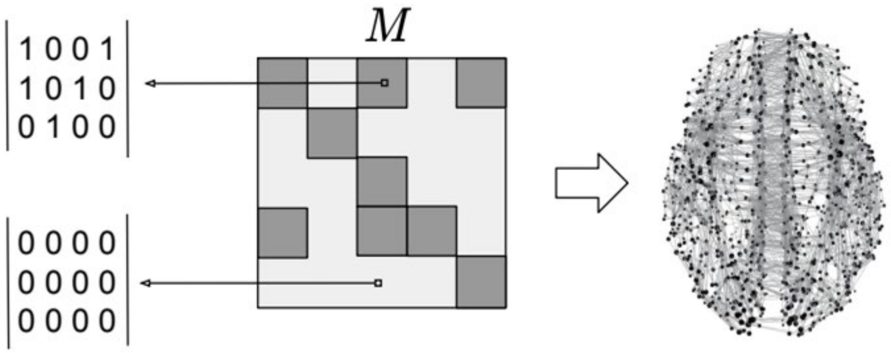

Salvatori, T., et al. "Learning on arbitrary graph topologies via predictive coding." 2022

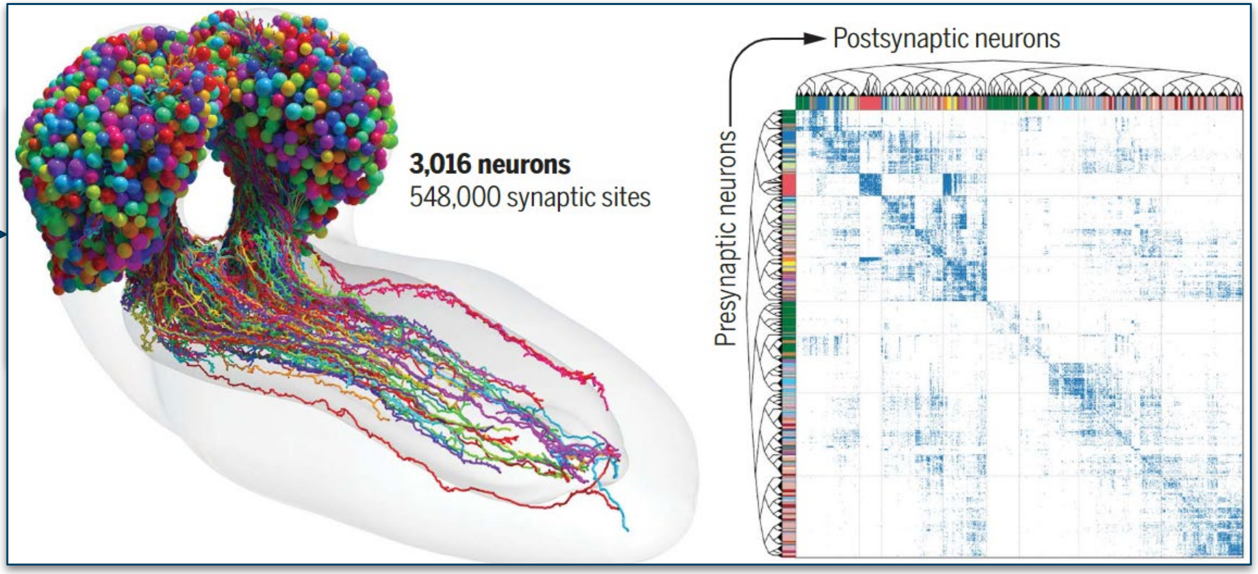

Winding, M., et al. "The connectome of an insect brain." 2023

What else does this framing buy you?

- It has been proven that PC is more robust than standard backprop-trained deep networks
  - Yields advantages in: online learning, training on small datasets, continual learning
- It shares interesting similarities w/ BP
  - It approximates backprop when output error is small
- Perfectly replicates backprop's weight update when adding a temporal scheduling on parameters

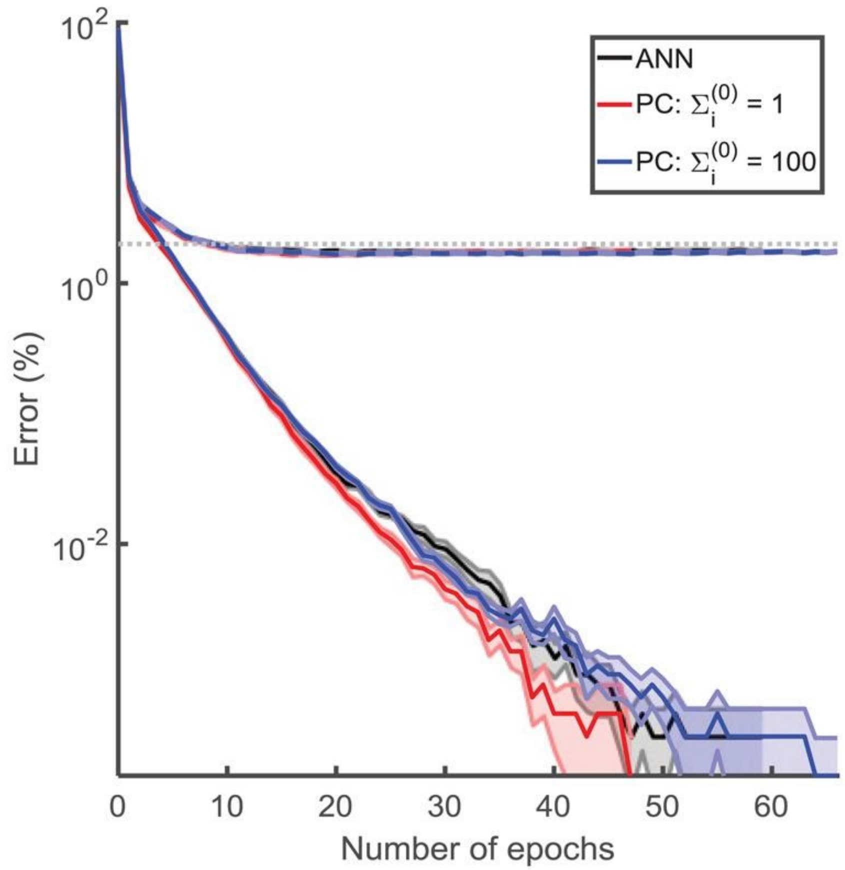

Whittington, JCR, Bogacz, R. "An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity." 2017.

Alonso, Nick, et al. "A Theoretical Framework for Inference Learning." 2022.

Millidge, B. et al. "Predictive coding approximates backprop along arbitrary computation graphs." 2022.

A natural generative model of data!

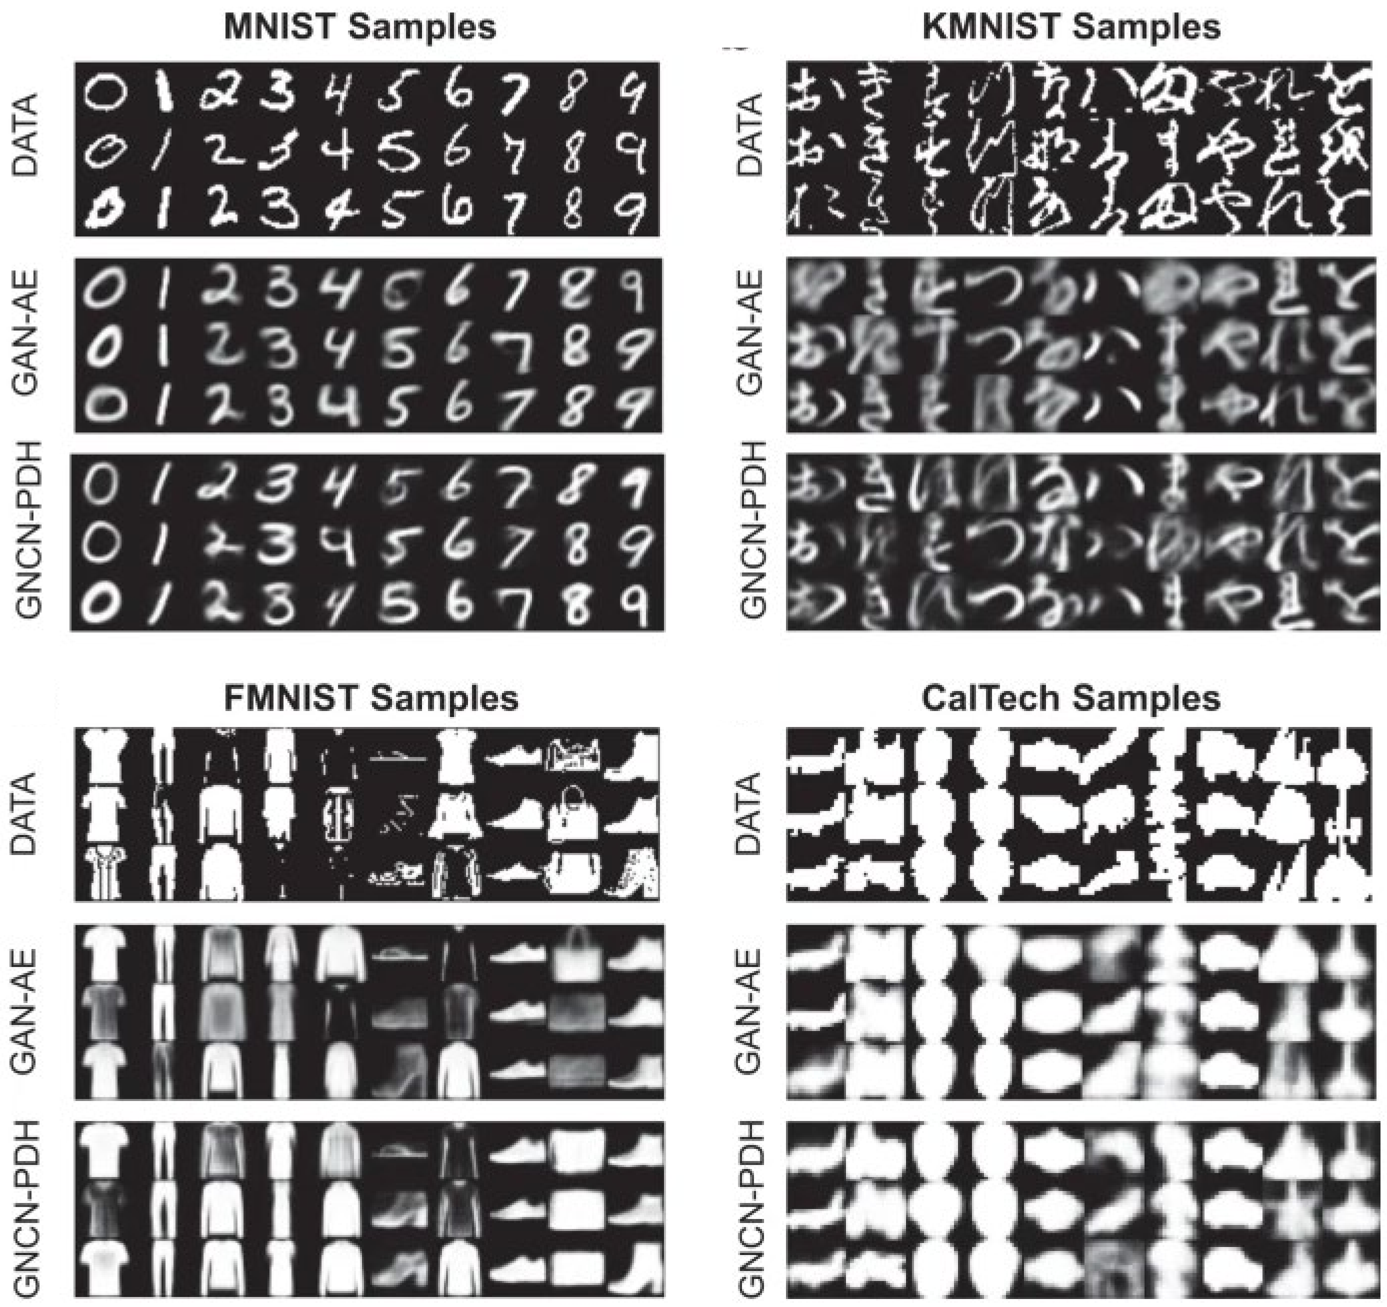

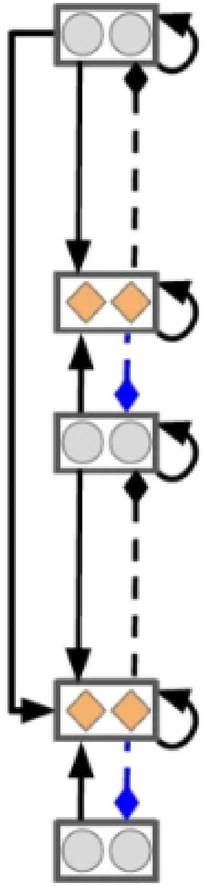

Ororbia, AG, Kifer D. "The neural coding framework for learning generative models." 2022.

A closer look at some NGC confabulations!

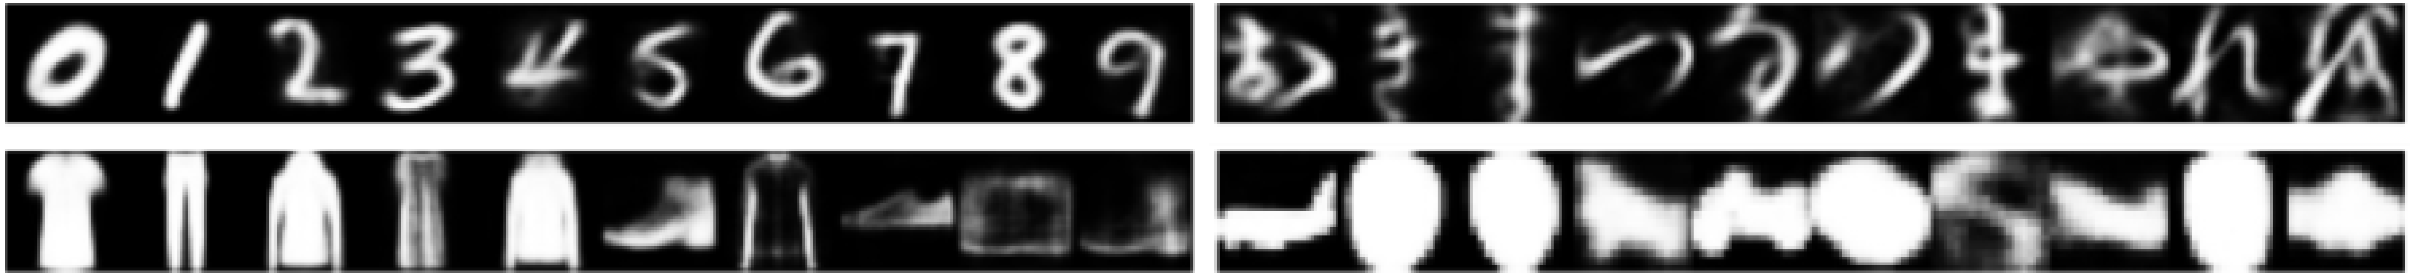

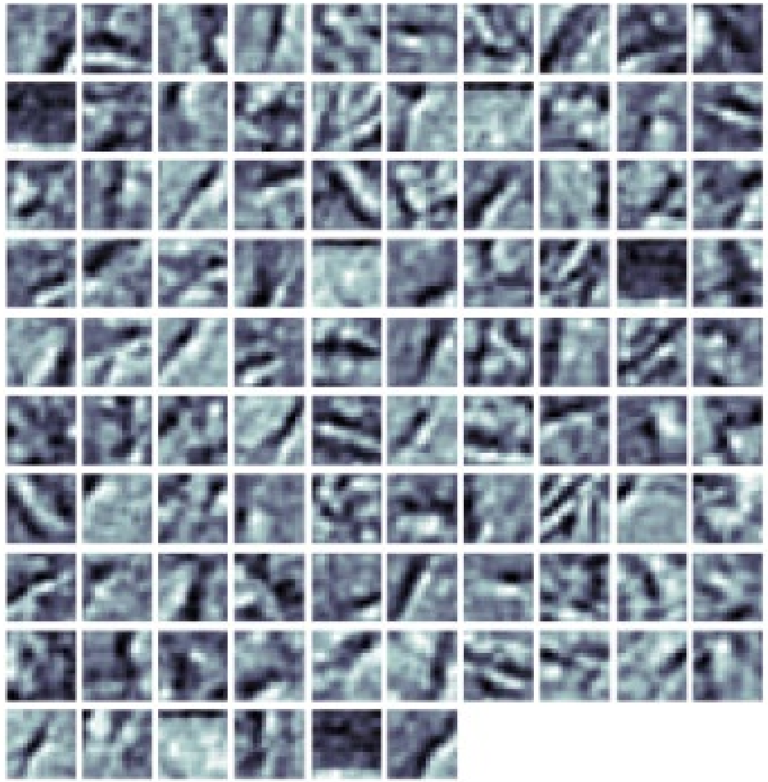

Our careful reproduction of some Rao's sensory level receptive fields on natural image!
(Habibi, F, Ororbia AG, 2024, upcoming)

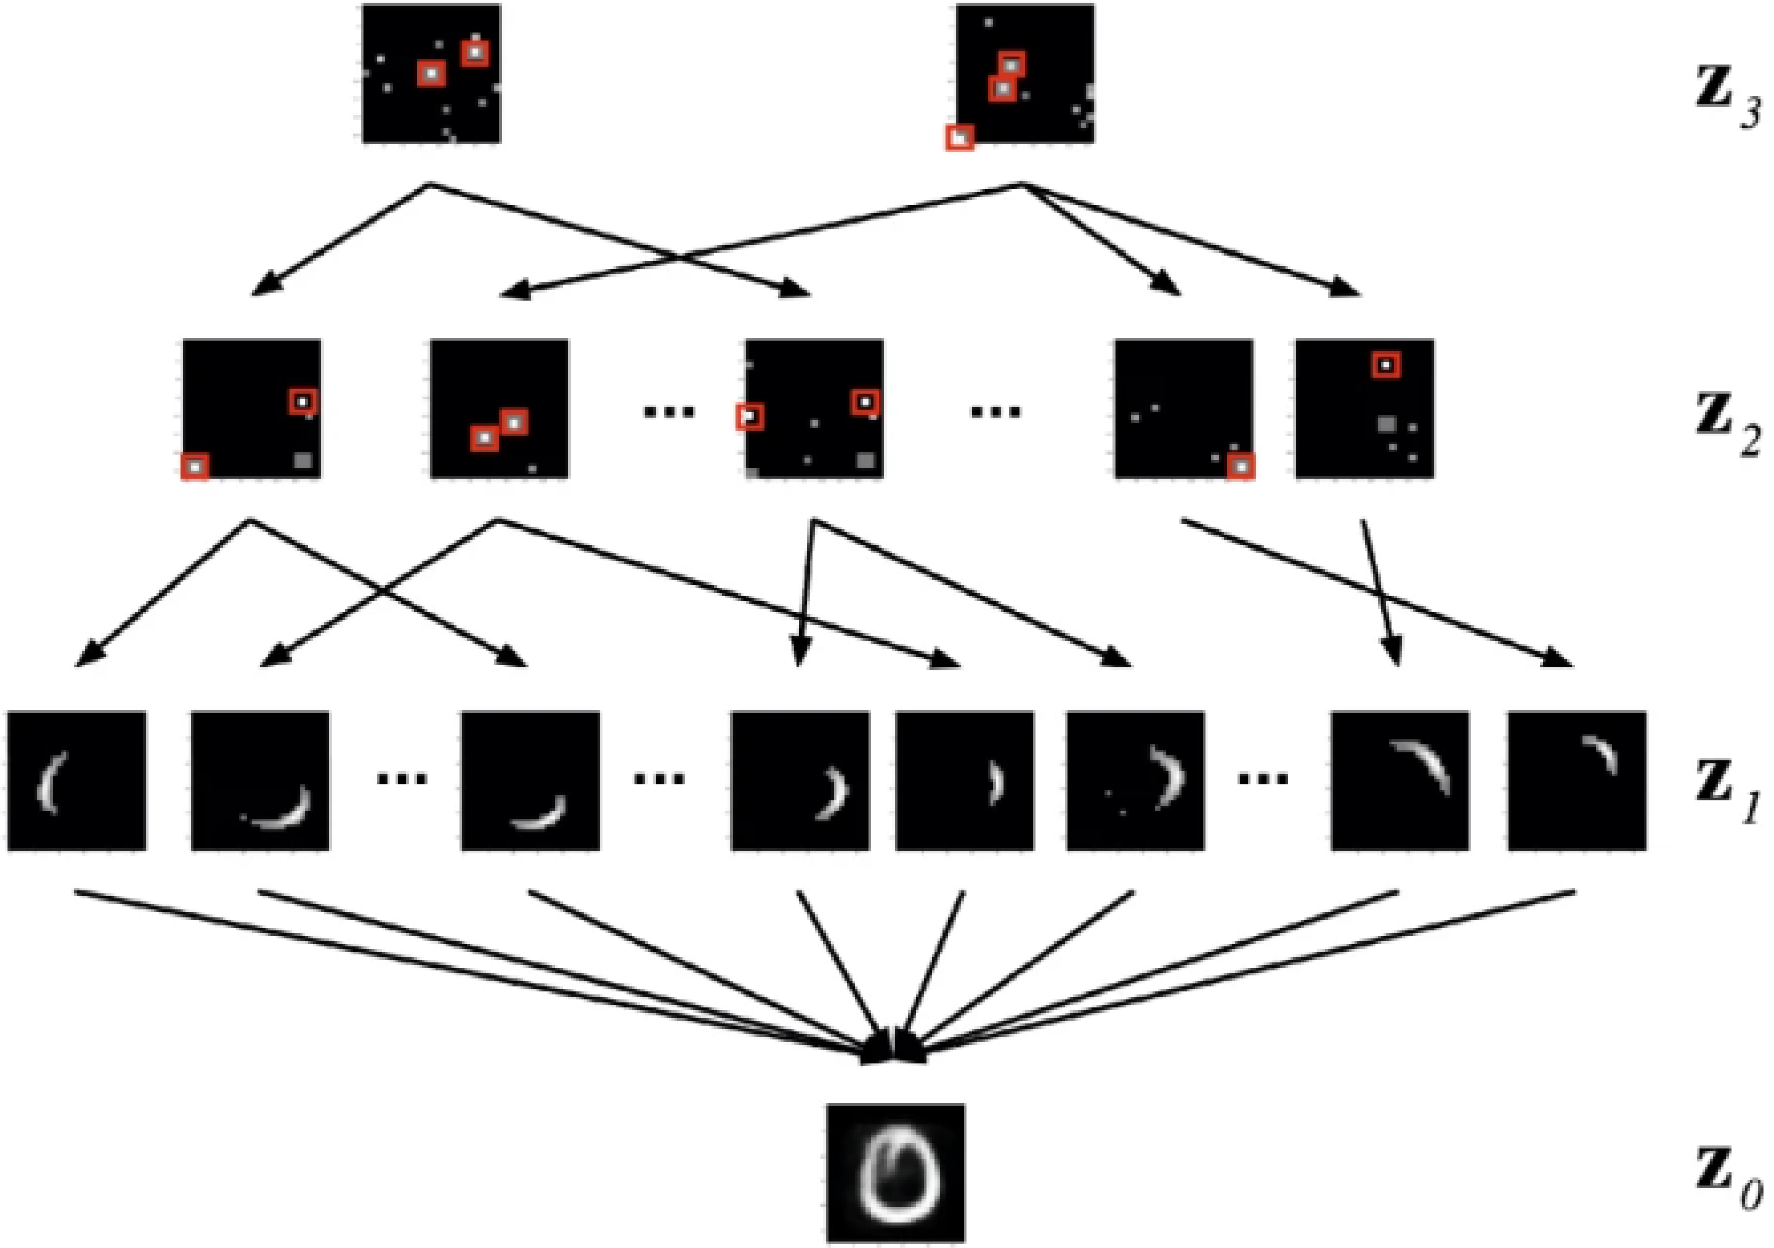

...
...
...
...
...
z_3
z_2
z_1
z_0

An Ever-Growing Body of Work and Results

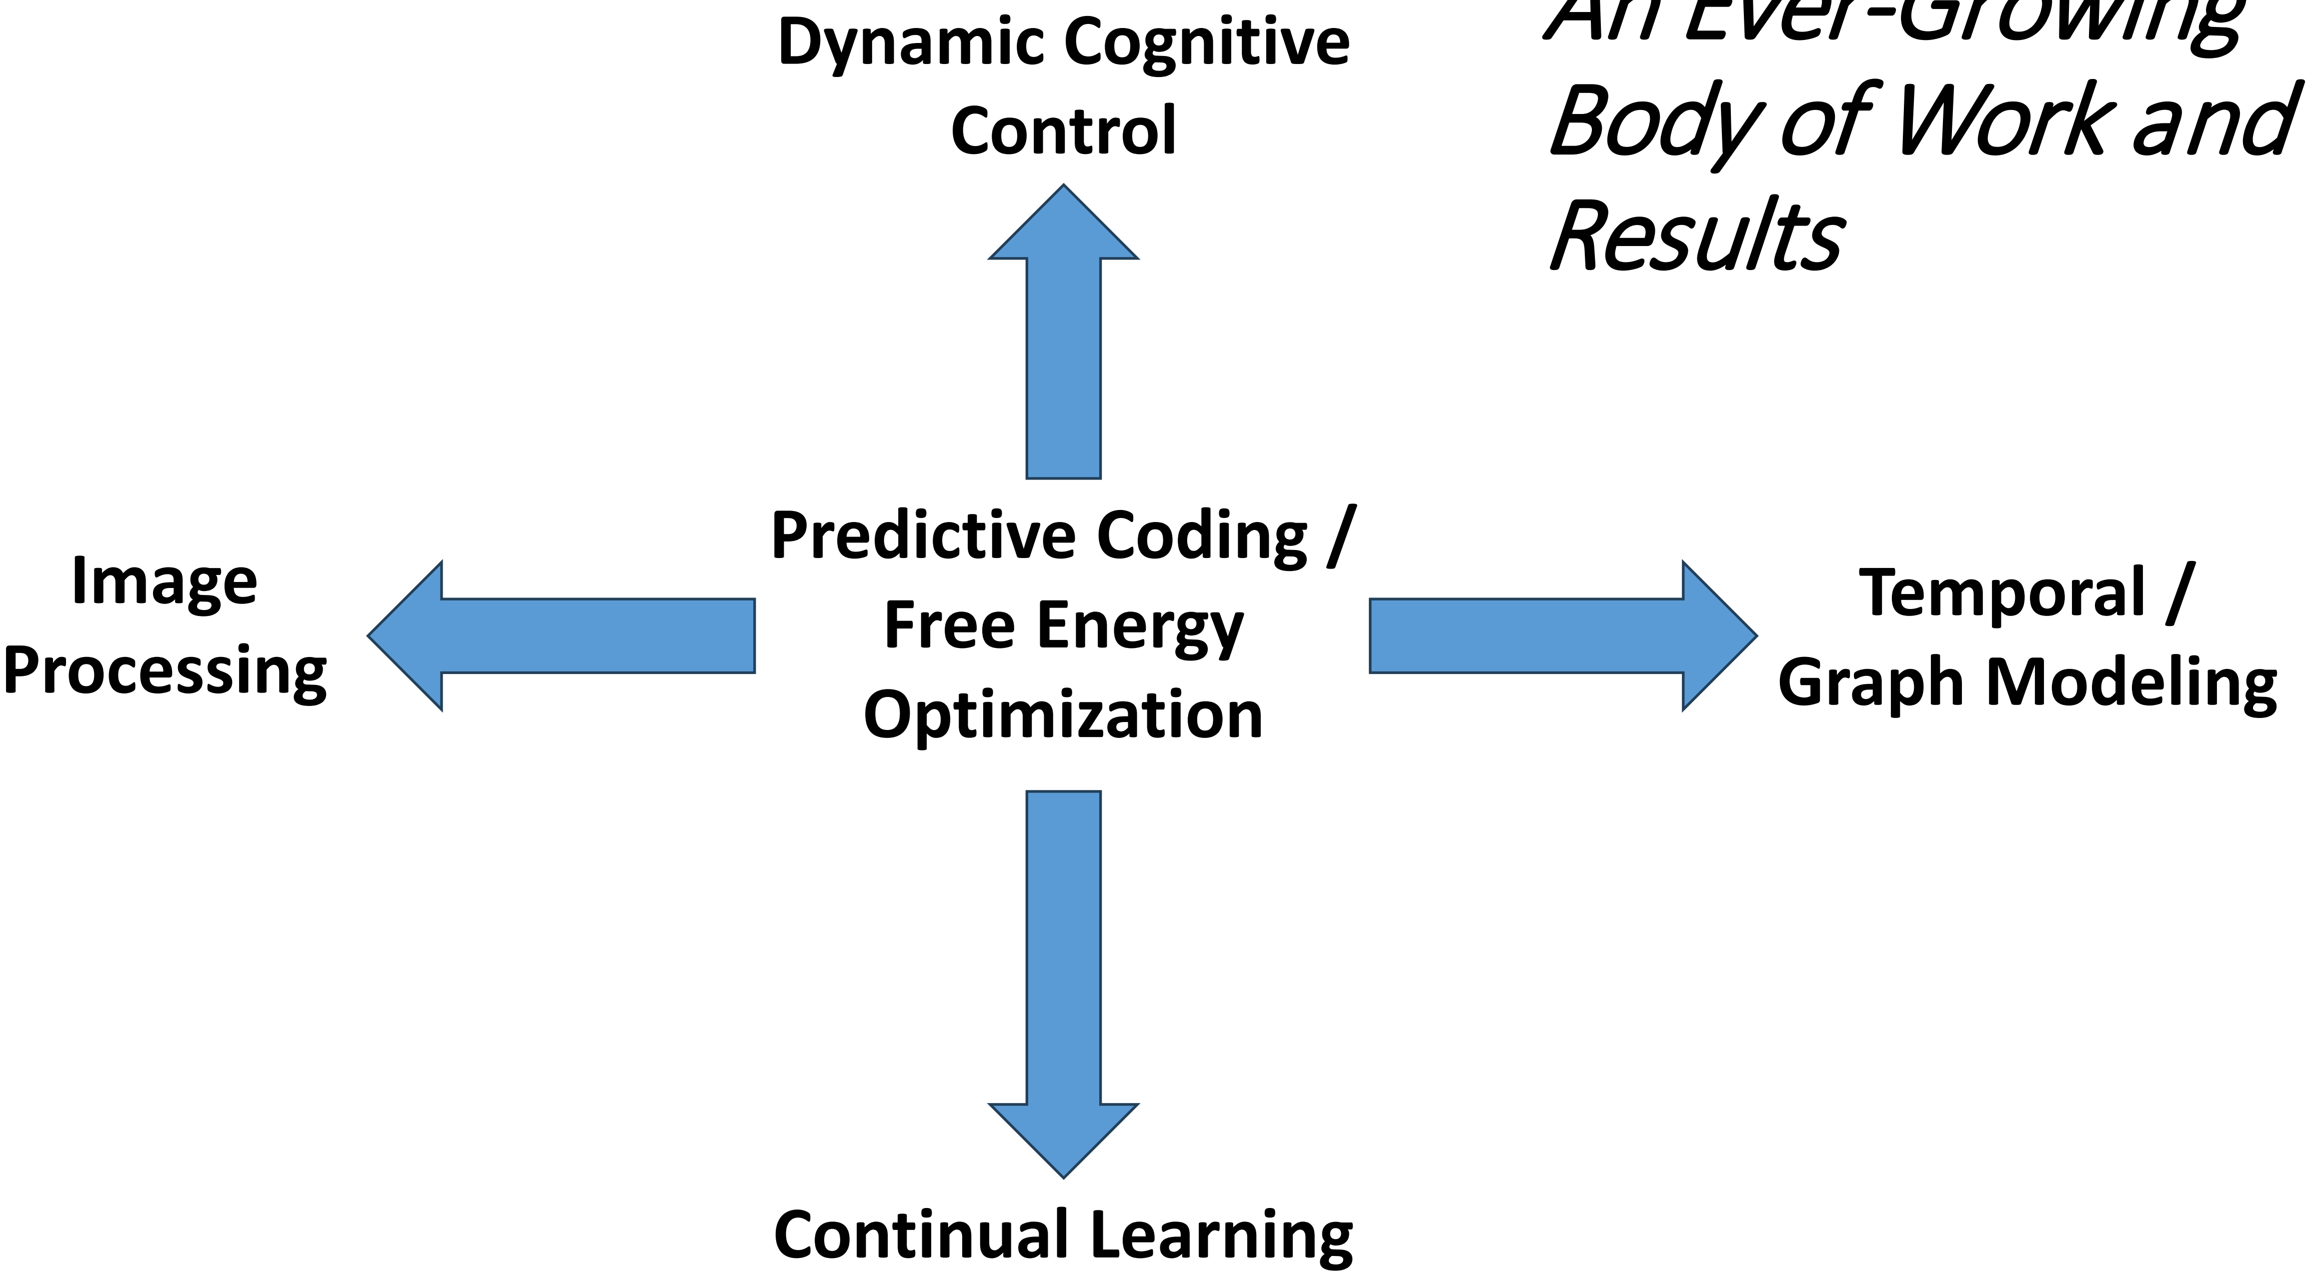

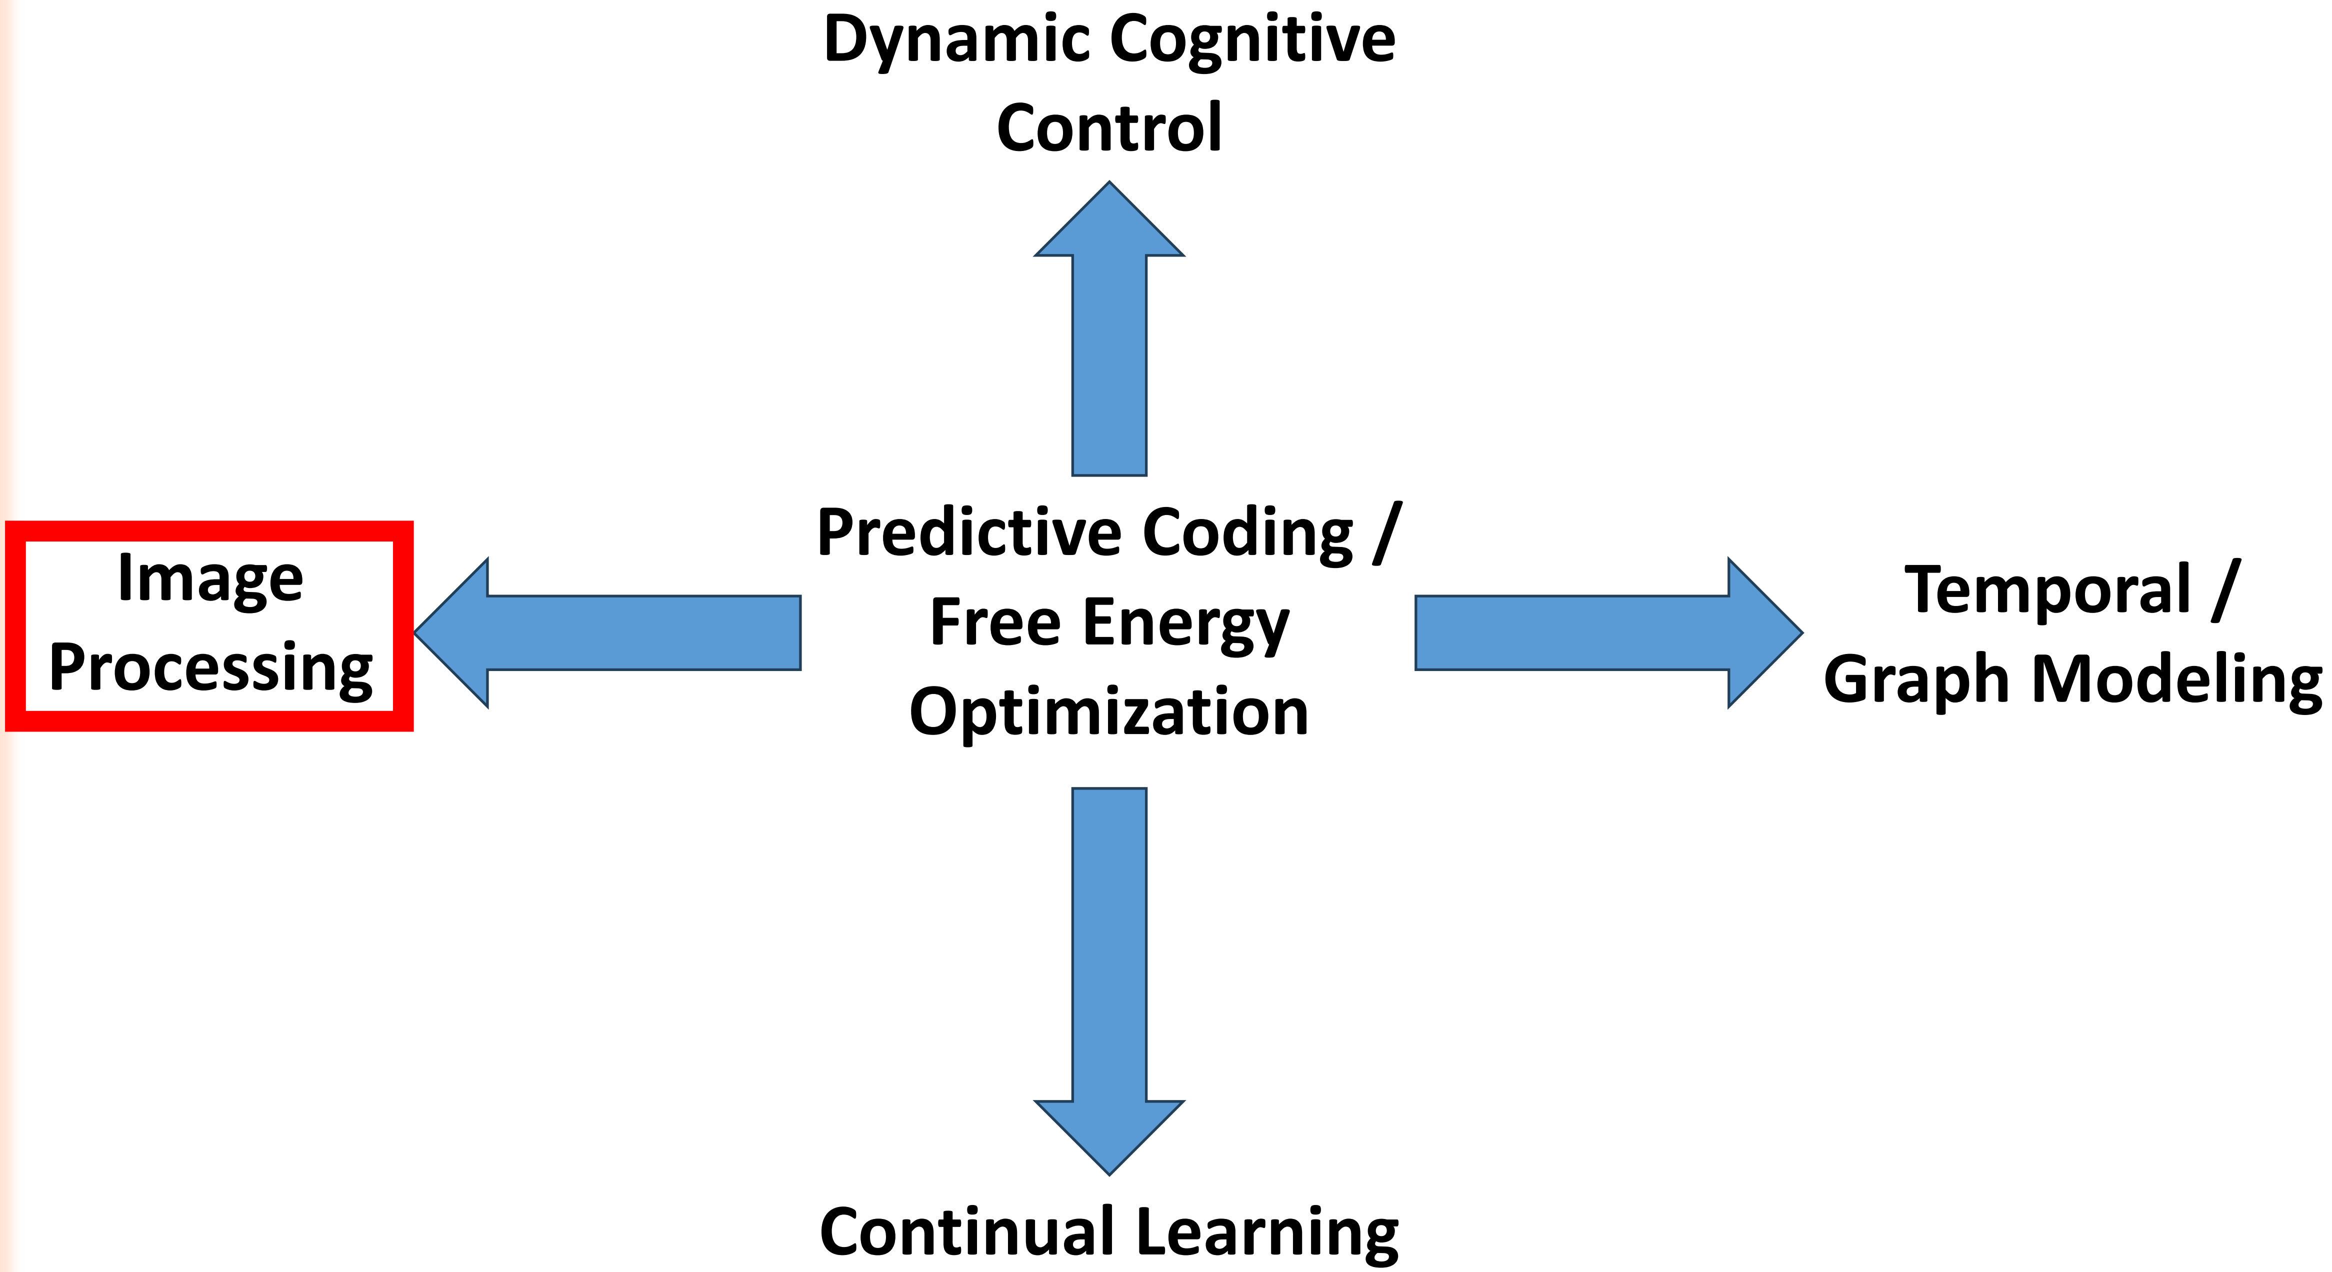
Dynamic Cognitive Control
Image Processing
Predictive Coding / Free Energy Optimization
Temporal / Graph Modeling
Continual Learning

PC able to perform basic denoising

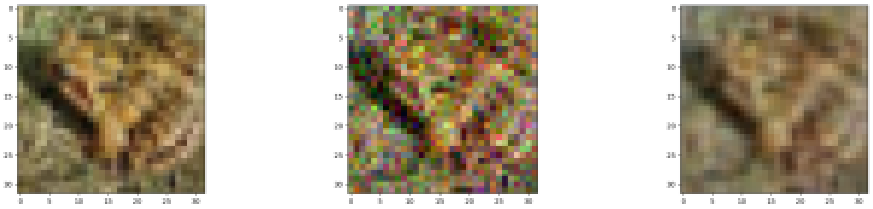

(a) Original image. (b) Corrupted image. (c) Denoised image.

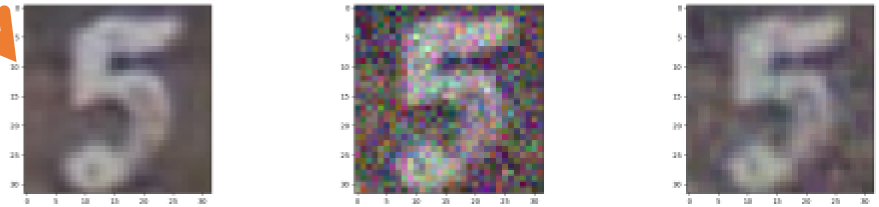

(d) Original image. (e) Corrupted image. (f) Denoised image.

Figure 3: Example image randomly sampled from a dataset test set (Left), the same image corrupted with noise $\sim \mathcal{N}(0, 0.1)$ (Middle), and the Conv-NGC denoising of the corrupted pattern (Right). Top row shows a sample taken from CIFAR-10 while the bottom row shows one taken from SVHN.

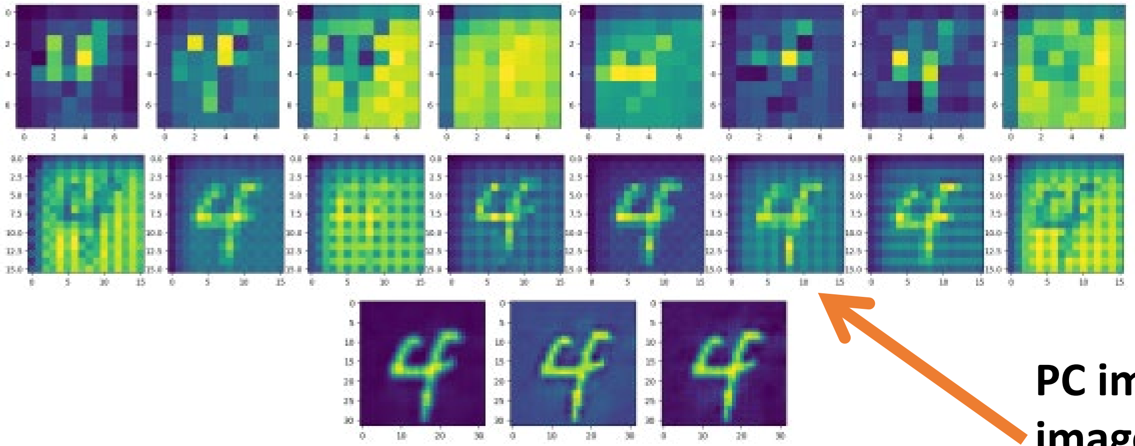

PC implicitly learns an image pyramid in its distributed representations

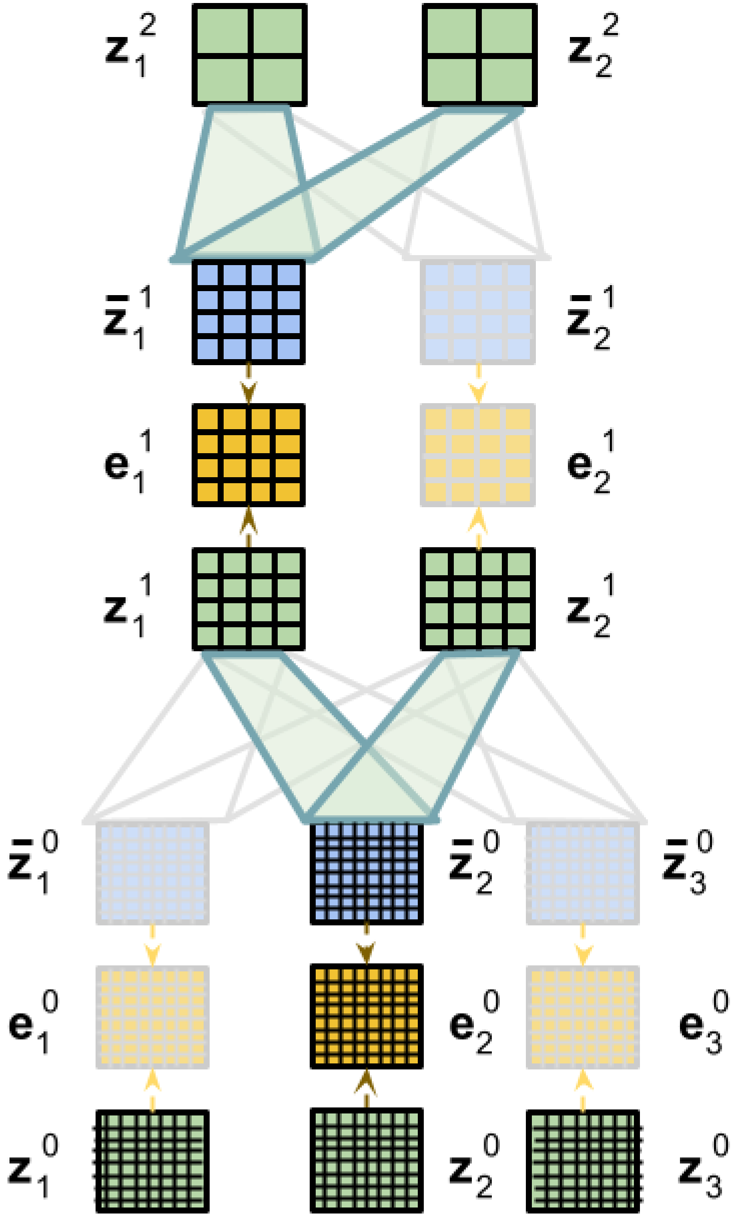

Ororbia, Alexander, and Ankur Mali. "Convolutional neural generative coding: Scaling predictive coding to natural images." 2022.

Predictive Coding and Convolutional Processing

	BP	PC
MLP on MNIST	$98.26\% \pm 0.12\%$	$98.55\% \pm 0.14\%$
MLP on FashionMNIST	$88.54\% \pm 0.64\%$	$85.12\% \pm 0.75\%$
CNN on SVHN	$95.35\% \pm 1.53\%$	$95.53\% \pm 1.54\%$
CNN on CIFAR-10	$69.34\% \pm 0.54\%$	$70.84\% \pm 0.64\%$

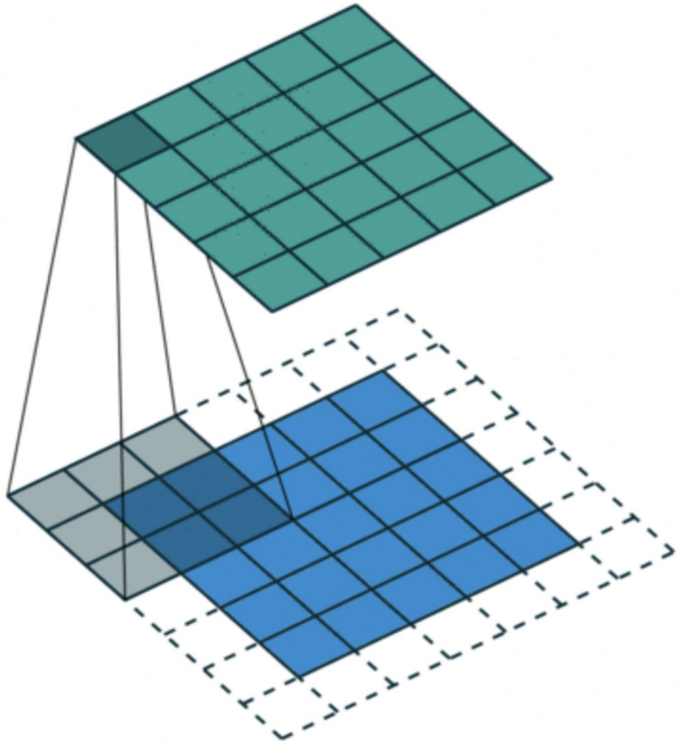

Salvatori, T., et al. "Incremental Predictive Coding: A Parallel and Fully Automatic Learning Algorithm." *arXiv preprint arXiv:2212.00720* (2022).

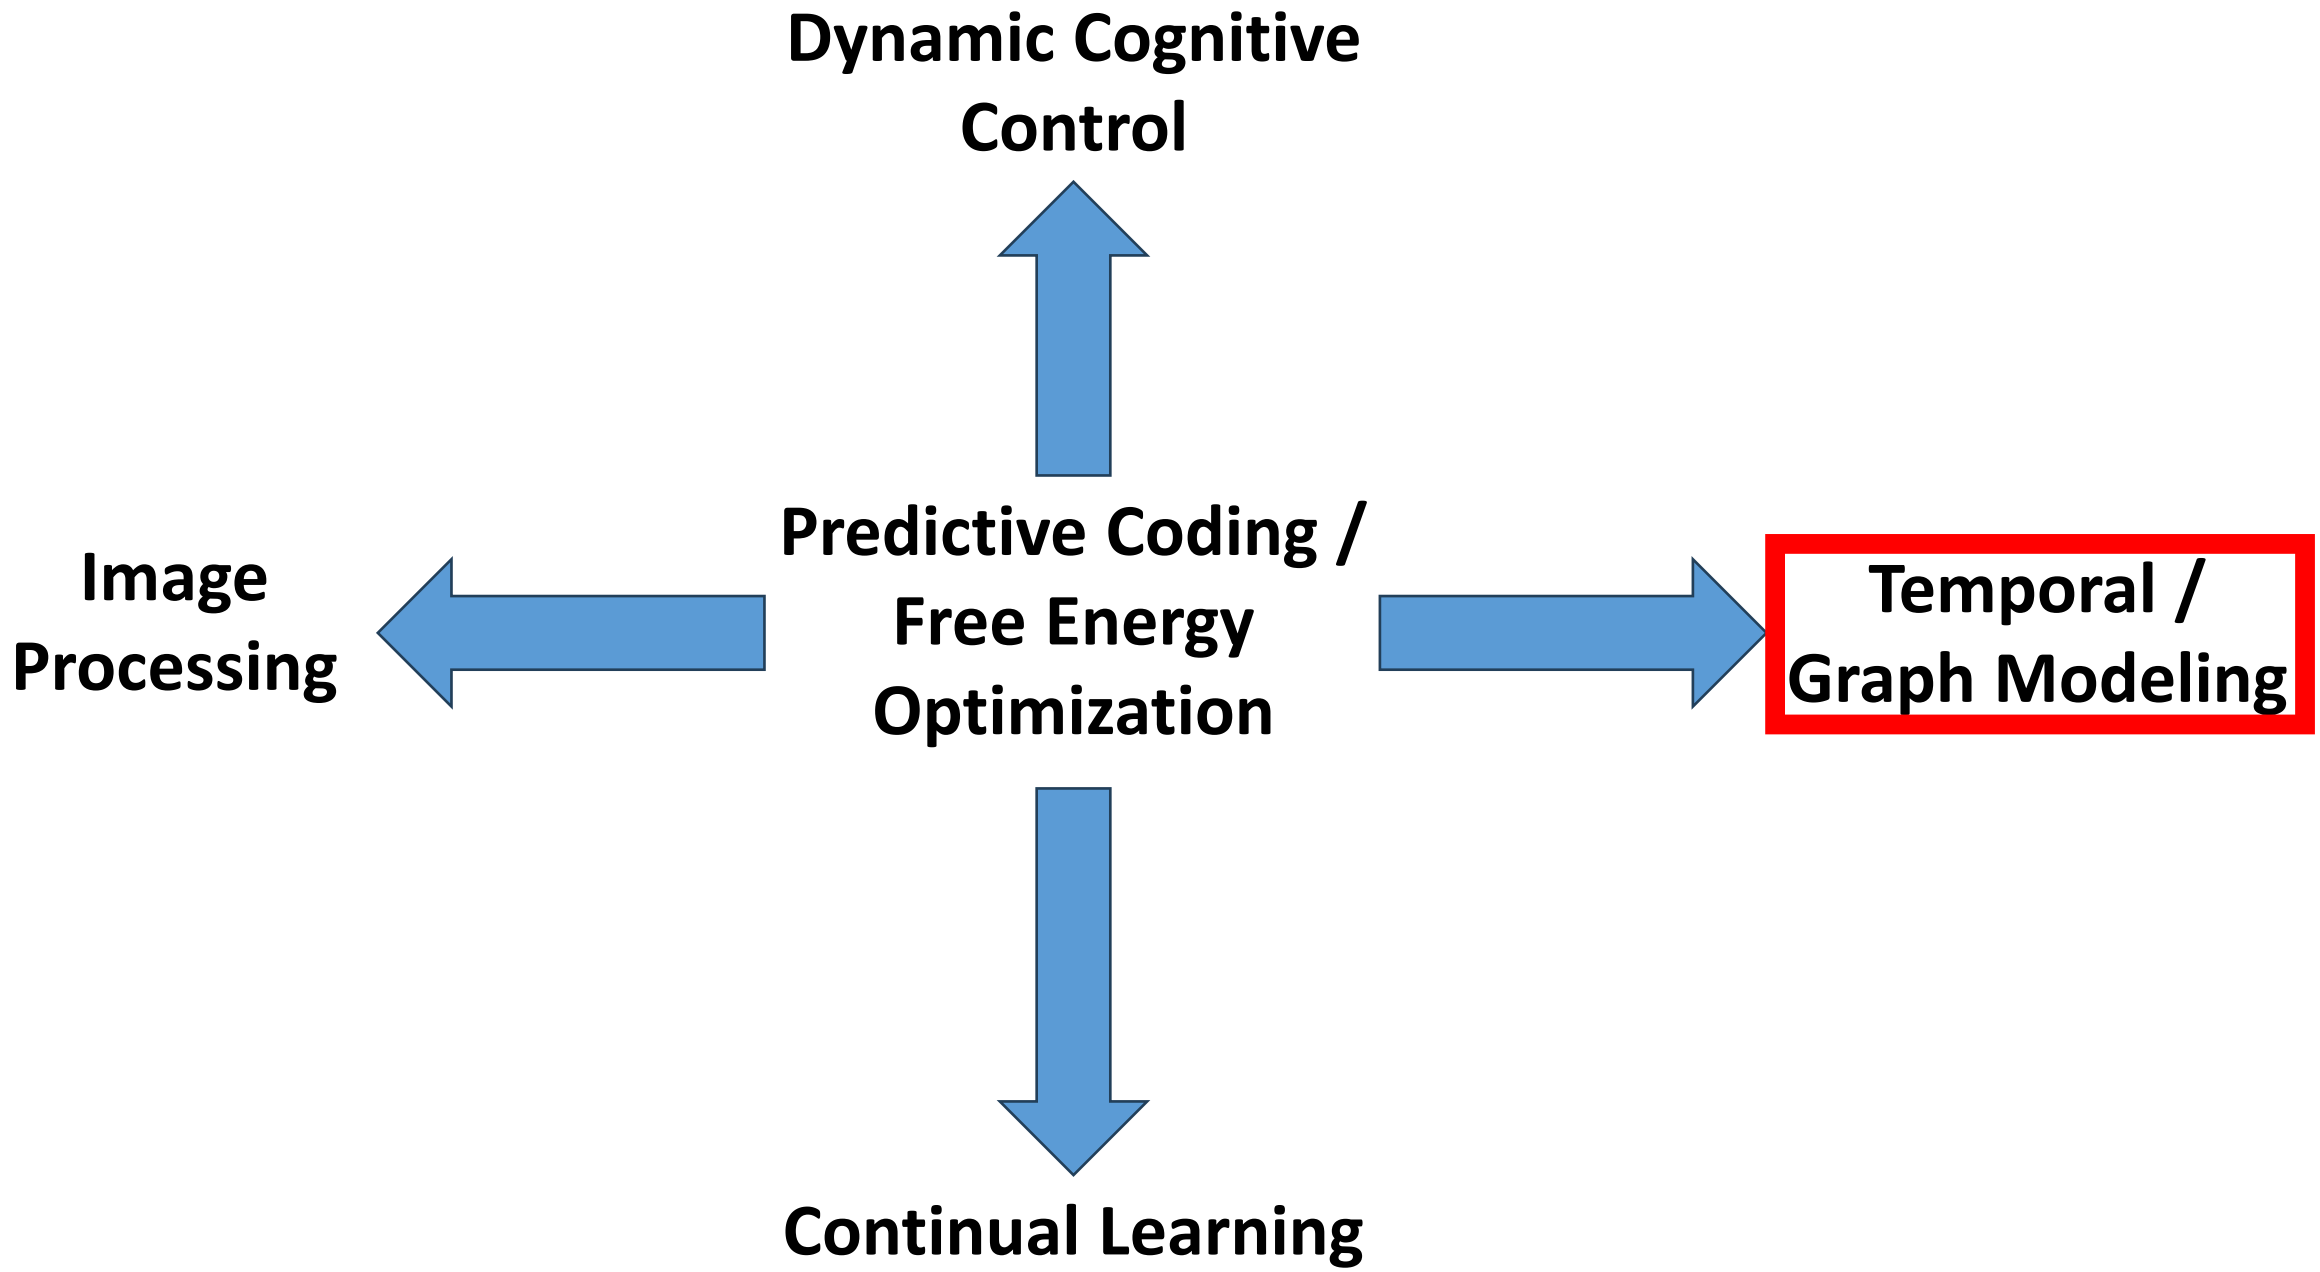
Dynamic Cognitive Control
Image Processing
Predictive Coding / Free Energy Optimization
Temporal / Graph Modeling
Continual Learning

Dynamic, Temporal Predictive Coding

- Can be done with fast-changing weights*
- Or via parallel temporal neural coding

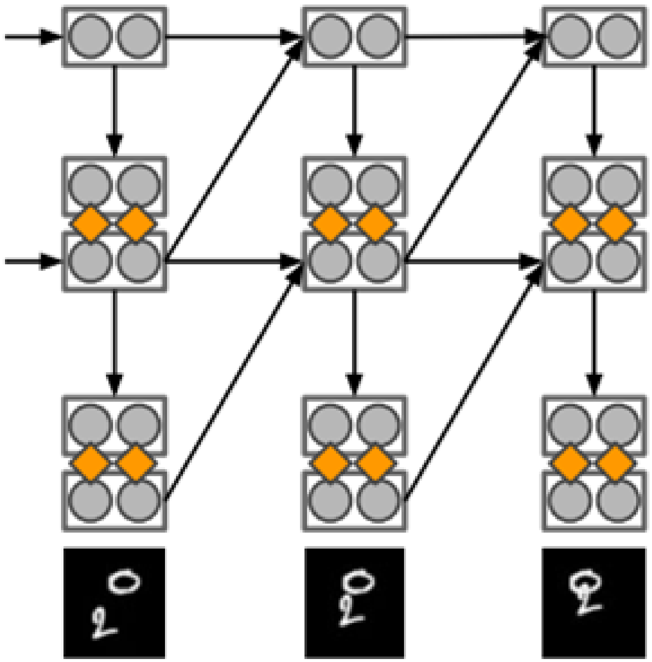

Parallel temporal neural coding

	Bouncing MNIST			Bouncing NotMNIST	
	Model	**Test CE**	**Test SE**	**Test CE**	**Test SE**
w/ unfold	LSTM-FP [56] (BPTT)	350.2	−−	−−	−−
	LSTM-CFP [56] (BPTT)	341.2	−−	−−	−−
	LSTM, BPTT (impl.)	375.42	85.27	787.51	256.66
	LSTM, SAB (impl.)	379.3	86.79	787.59	256.89
	GRU, BPTT (impl.)	375.0	85.18	788.00	257.01
	RNN, BPTT (impl.)	391.4	90.14	795.12	269.29
	RNN, SAB (impl.)	392.7	90.22	794.21	265.21
no unfold	ESN (impl.)	489.2	99.86	812.43	305.57
	LSTM, UORO (impl.)	386.7	89.21	789.48	259.10
	LSTM, RTRL (impl.)	361.2	85.89	778.29	222.08
	P-TNCN (ours)	**338.79**	**79.67**	**713.67**	**176.73**

ZERO-SHOT ADAPTIVE PERFORMANCE OF THE MODELS TRAINED ON NOTMNIST AND TESTED ON MNIST AND VICE VERSA

	NotMNIST → MNIST		MNIST → NotMNIST	
Model, 0-shot	**CE**	**SE**	**CE**	**SE**
LSTM, BPTT	492.21	104.76	1297.26	325.56
LSTM, SAB	499.21	105.87	1299.28	329.59
LSTM, RTRL	447.28	99.89	1211.01	293.56
P-TNCN	**377.30**	**89.39**	**1131.7**	**257.07**

Ororbia, A.. et al. "Learning to Adapt by Minimizing Discrepancy." (2017)

Ororbia, A.. et al. "Continual Learning of Recurrent Neural Networks by Locally Aligning Distributed Representations." (2019)

* Jiang, L. et al. "Dynamic predictive coding: A new model of hierarchical sequence learning and prediction in the cortex." 2022.

Robotic Control Confabulations

Dreaming of electric sheep...er.... eletric blocks!

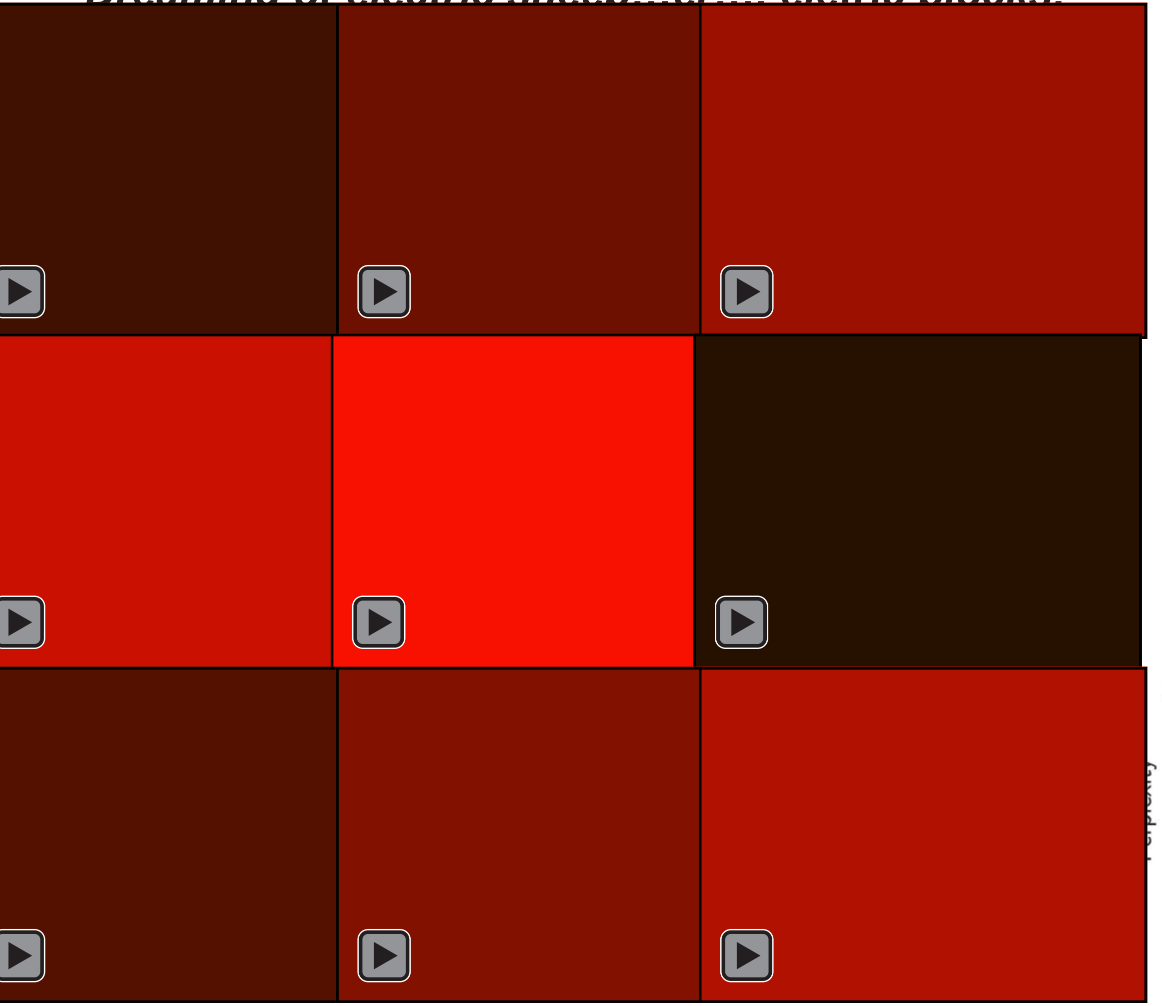

Ororbia et al., 2024; Coming soon!
Keep an eye out on arXiv!

Method	CORA	CiterSeer	PubMed
BP	$\mathbf{80.72 \pm 1.05\%}$	$67.12 \pm 1.53\%$	$\mathbf{77.1 \pm 1.45\%}$
PC	$80.7 \pm 1.09\%$	$\mathbf{67.26 \pm 1.28\%}$	$76.2 \pm 2.44\%$

Graph Neural Network Topologies

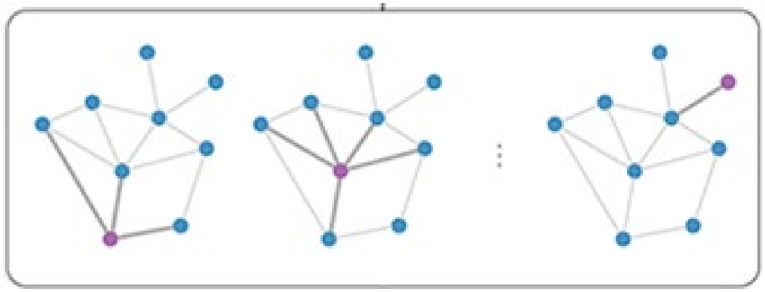

Byiringiro, B., et al. "Robust Graph Representation Learning via Predictive Coding." 2022.

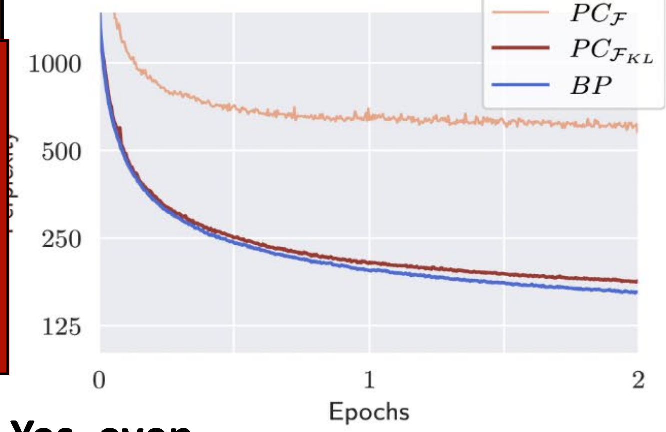

Training Method	Test Perplexity
BP	162.64 ± 0.76
$PC_{\mathcal{F}_{KL}}$	175.90 ± 1.74
$PC_{\mathcal{F}}$	590.08 ± 12.60

Yes, even transformers!

Byiringiro, B., et al. "Robust Graph Representation Learning via Predictive Coding." 2022.

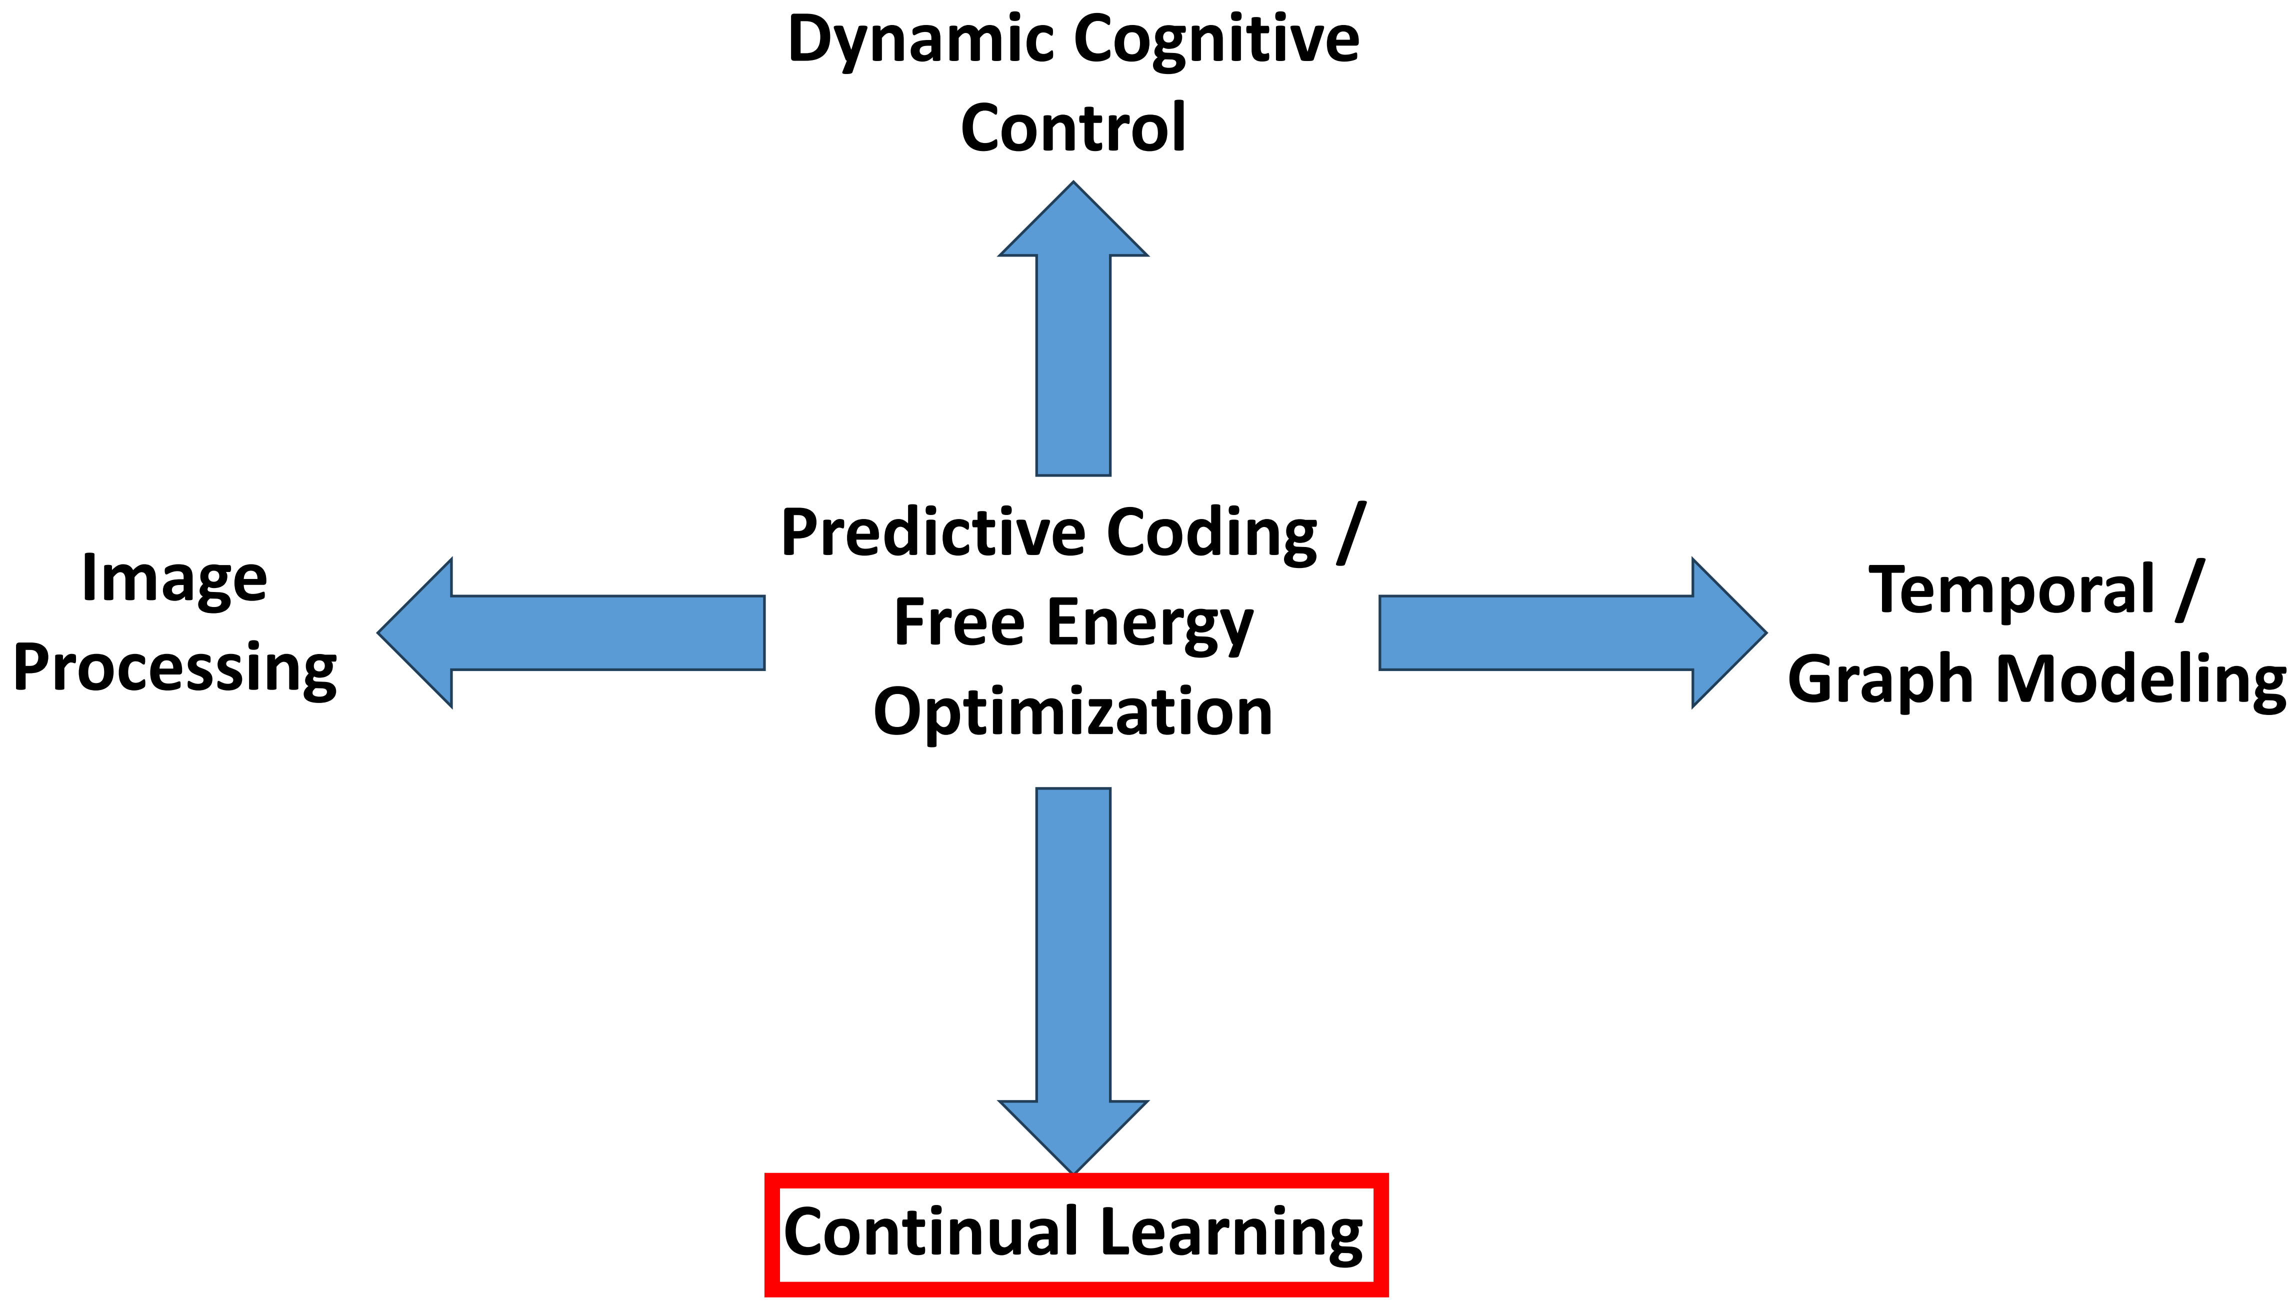
Dynamic Cognitive Control
Image Processing
Predictive Coding / Free Energy Optimization
Temporal / Graph Modeling
Continual Learning

Sequential Neural Coding Network (S-NCN)

Table 3: Generalization metrics (10 trials) for Split MNIST, Split NotMNIST, and Split Fashion MNIST (FMNIST) benchmarks. Note for IMM, we employ the best performing variant, *L2+DT+Md-IMM*.

	MNIST		NotMNIST		FMNIST	
	ACC	BWT	ACC	BWT	ACC	BWT
EWC (SH)	19.04 ± 0.03	-0.3569 ± 0.015	18.55 ± 0.02	-0.3611 ± 0.01	19.88 ± 0.06	-0.3499 ± 0.012
O-EWC (SH)	19.56 ± 0.04	-0.3500 ± 0.01	18.45 ± 0.03	-0.3600 ± 0.012	19.02 ± 0.05	-0.3422 ± 0.01
NR+Mem-1 (SH)	90.58 ± 0.87	-0.05 ± 0.001	89.02 ± 0.030	-0.071 ± 0.004	90.01 ± 0.81	-0.06 ± -0.003
	...			...		
ICarl (SH)	93.99 ± 0.41	-0.100 ± 0.004	88.69 ± 0.102	-0.109 ± 0.007	95.95 ± 0.40	-0.110 ± -0.005
Lucir (SH)	94.02 ± 0.31	-0.103 ± 0.007	93.45 ± 0.093	-0.101 ± 0.006	95.02 ± 0.34	-0.110 ± 0.005
Bic (SH)	90.09 ± 0.86	-0.139 ± -0.009	85.09 ± 0.099	-0.155 ± 0.0091	89.00 ± 0.85	-0.160 ± 0.009
Mnemonics (SH)	96.01 ± 0.32	-0.991 ± 0.005	95.02 ± 0.071	-0.989 ± 0.007	96.75 ± 0.30	0.993 ± 0.006
Lat-S-NCN	0.981 ± 0.003	-0.005 ± 0.004	0.957 ± 0.004	-0.004 ± 0.005	0.982 ± 0.004	-0.003 ± 0.007

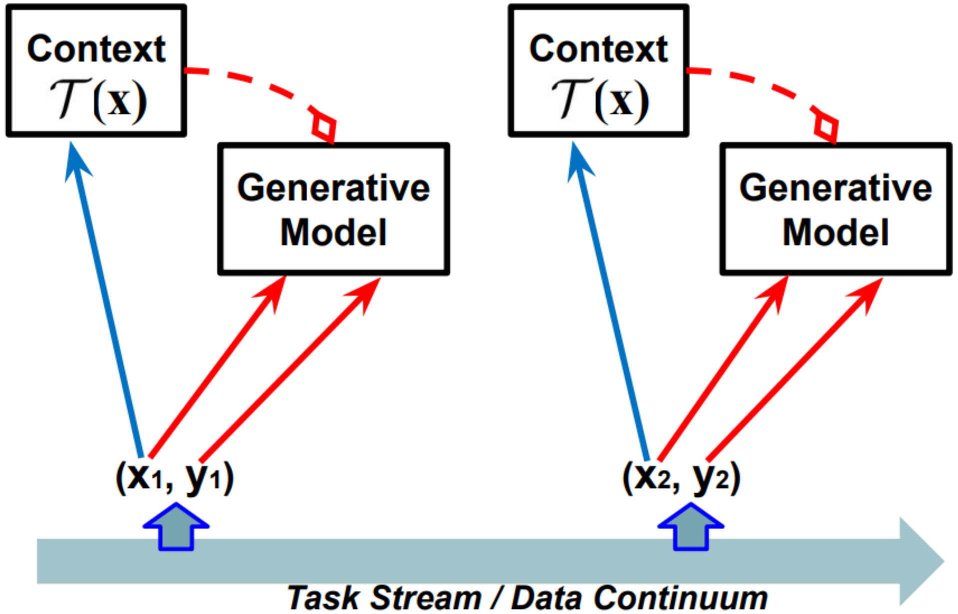

Ororbia, AG, et al. "Lifelong neural predictive coding: Learning cumulatively online without forgetting." 2022.

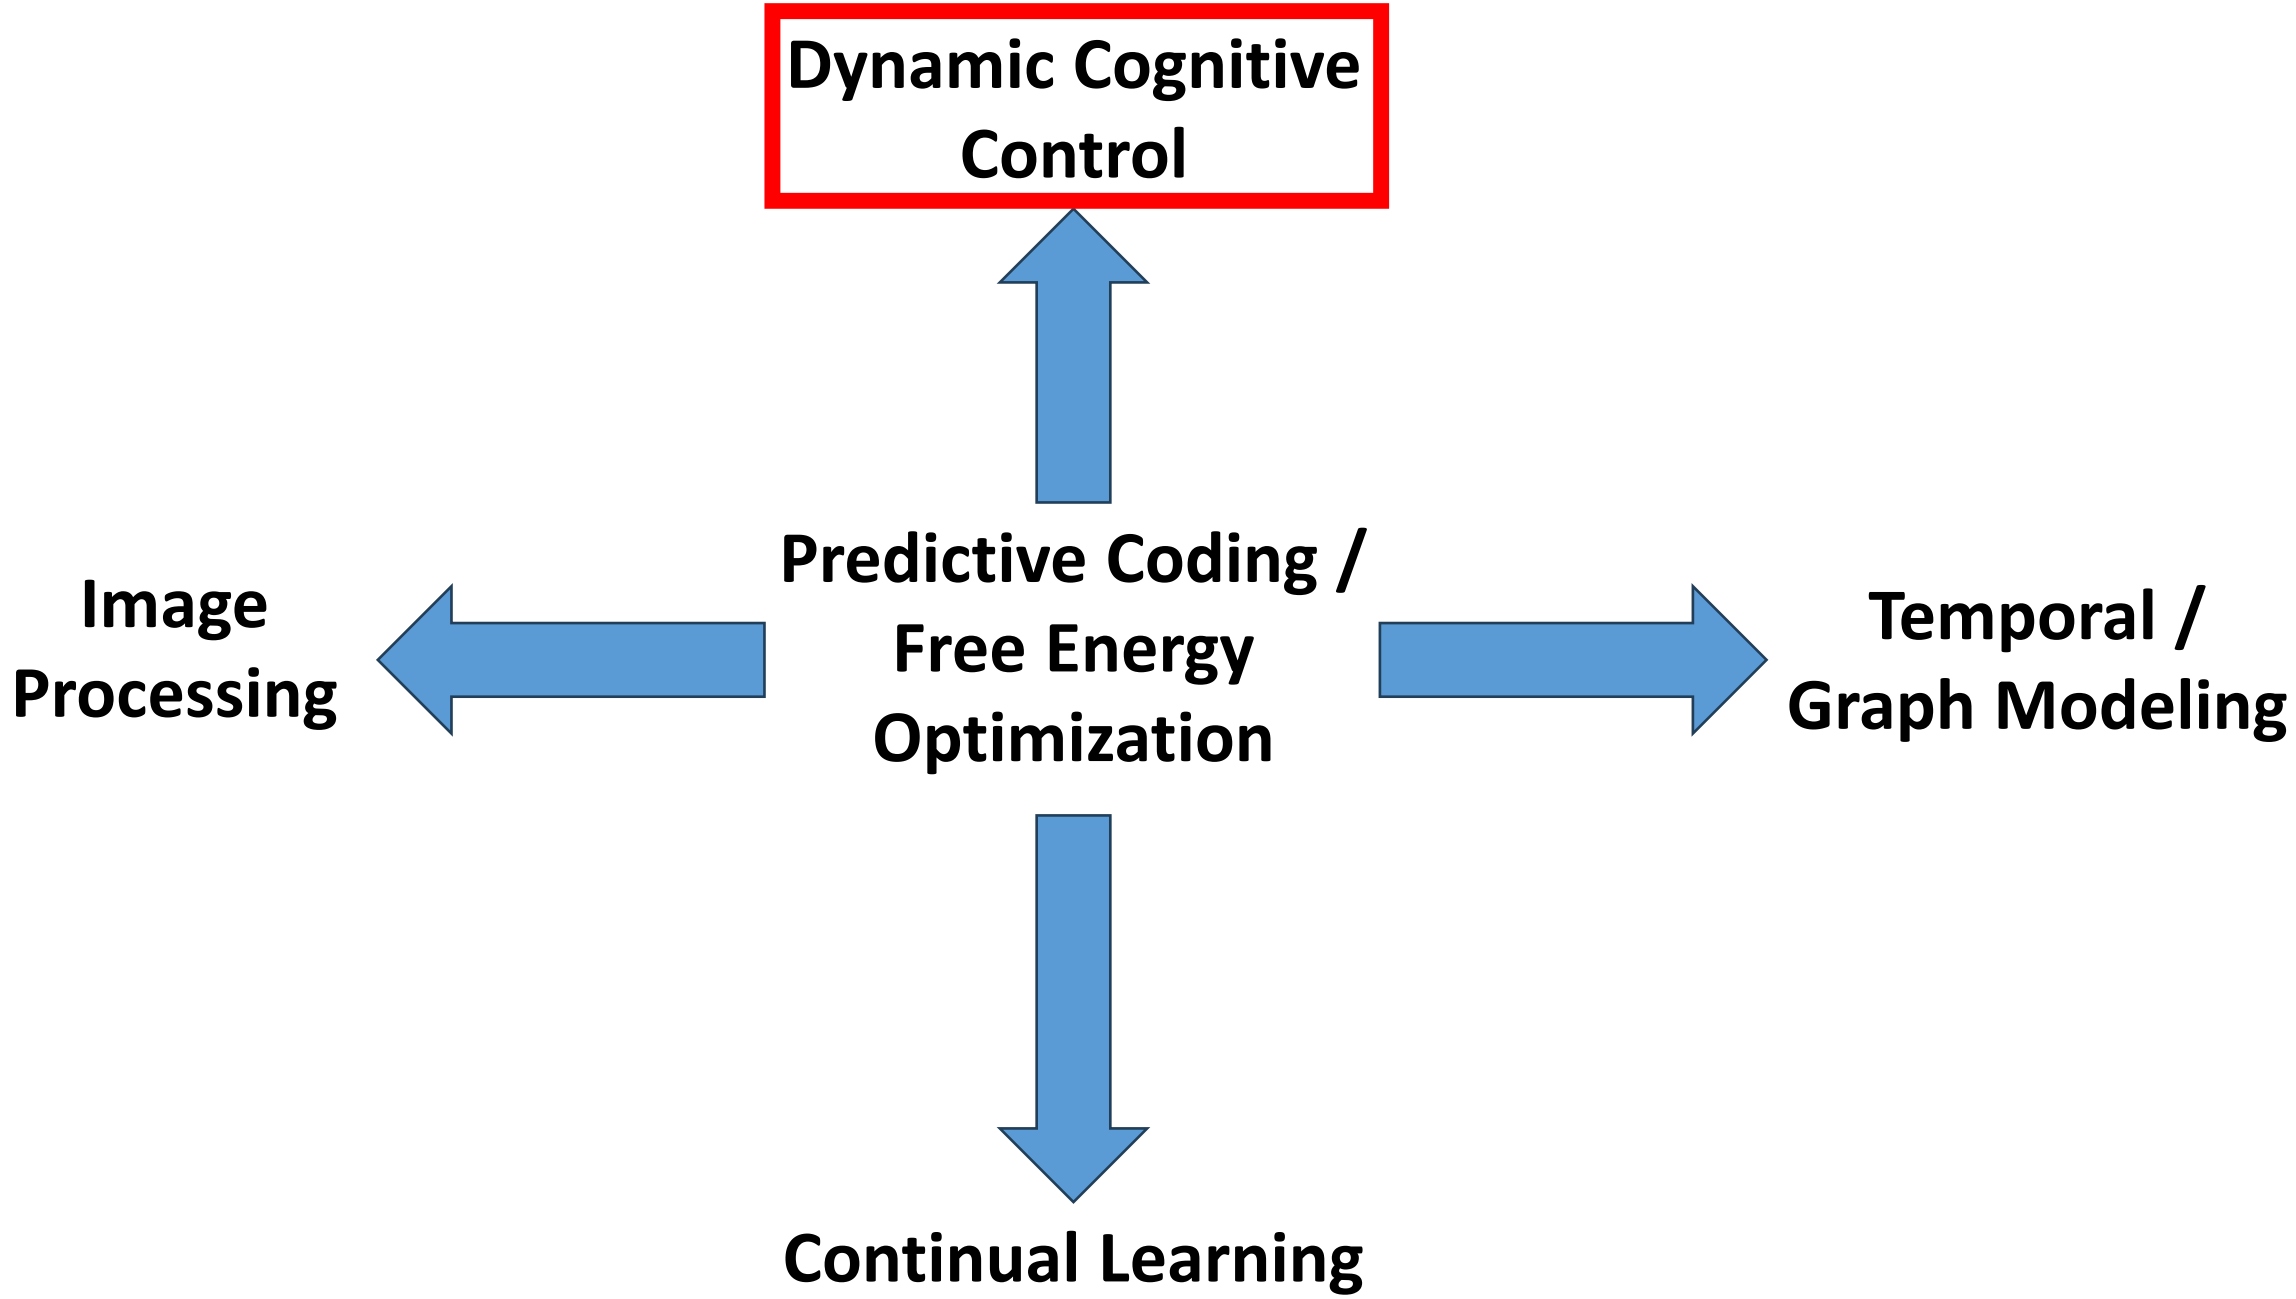
Dynamic Cognitive Control
Image Processing
Predictive Coding / Free Energy Optimization
Temporal / Graph Modeling
Continual Learning

Active Predictive Coding

Fast-changing weights work for actions as well*

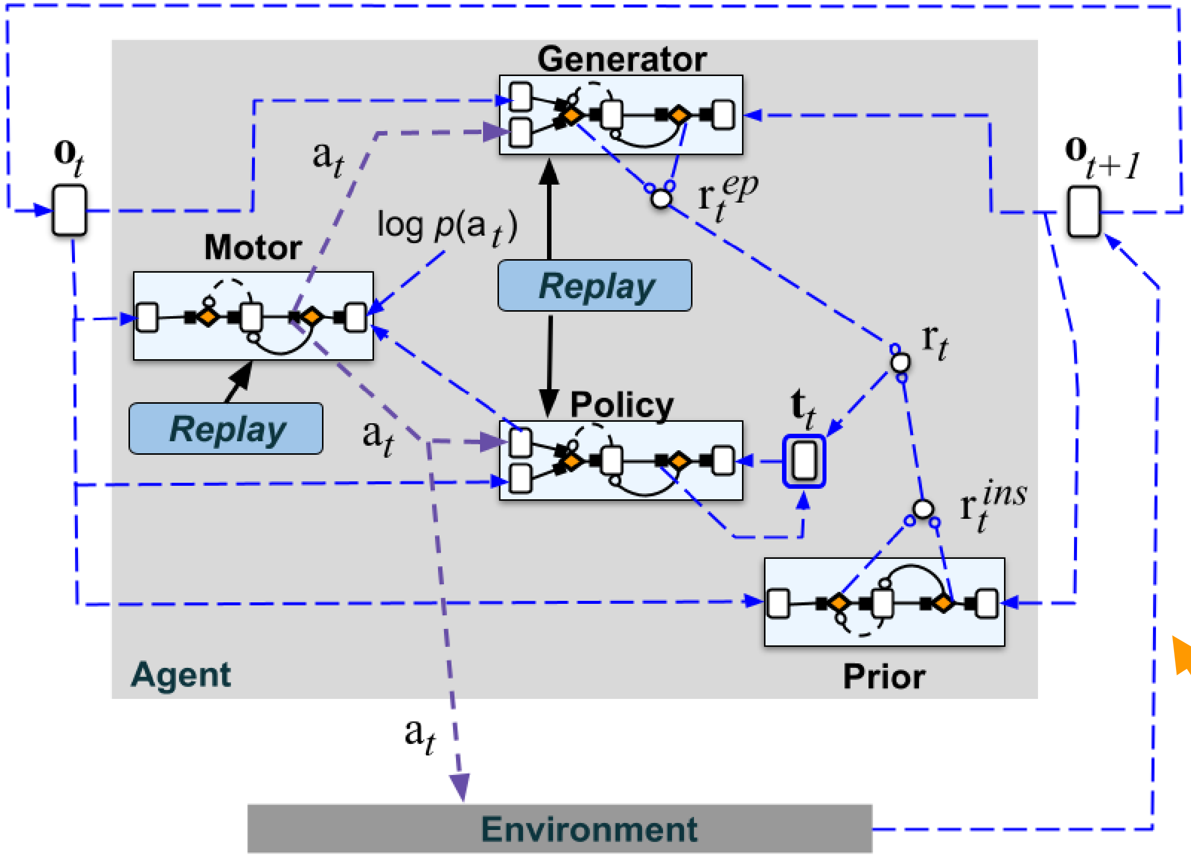

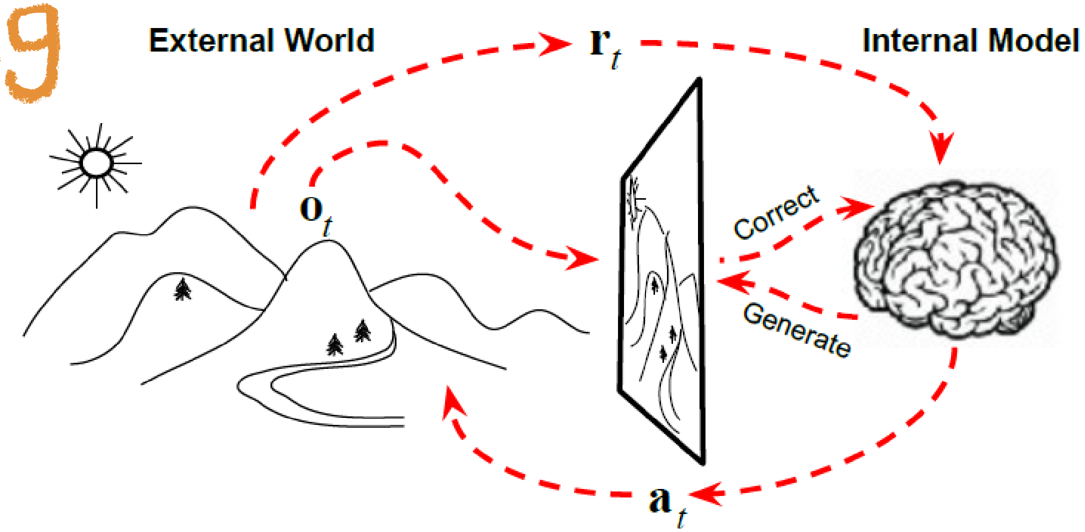

Active predictive coding (*ActPC*) agent for robotic control - one circuit for:

1) prior preference,
2) generative transition dynamics,
3) policy, and
4) motor-action

Ororbia, AG., Mali A.. "Backprop-free reinforcement learning with active neural generative coding." 2022.

Ororbia, AG., Mali A.. "Active predictive coding: Brain-inspired reinforcement learning for sparse reward robotic control problems." 2023.

* Rao, RPN, et al. "Active predictive coding: A unifying neural model for active perception, compositional learning, and hierarchical planning." 2023.

Active Predictive Coding

Fast-changing weights work for actions as well*

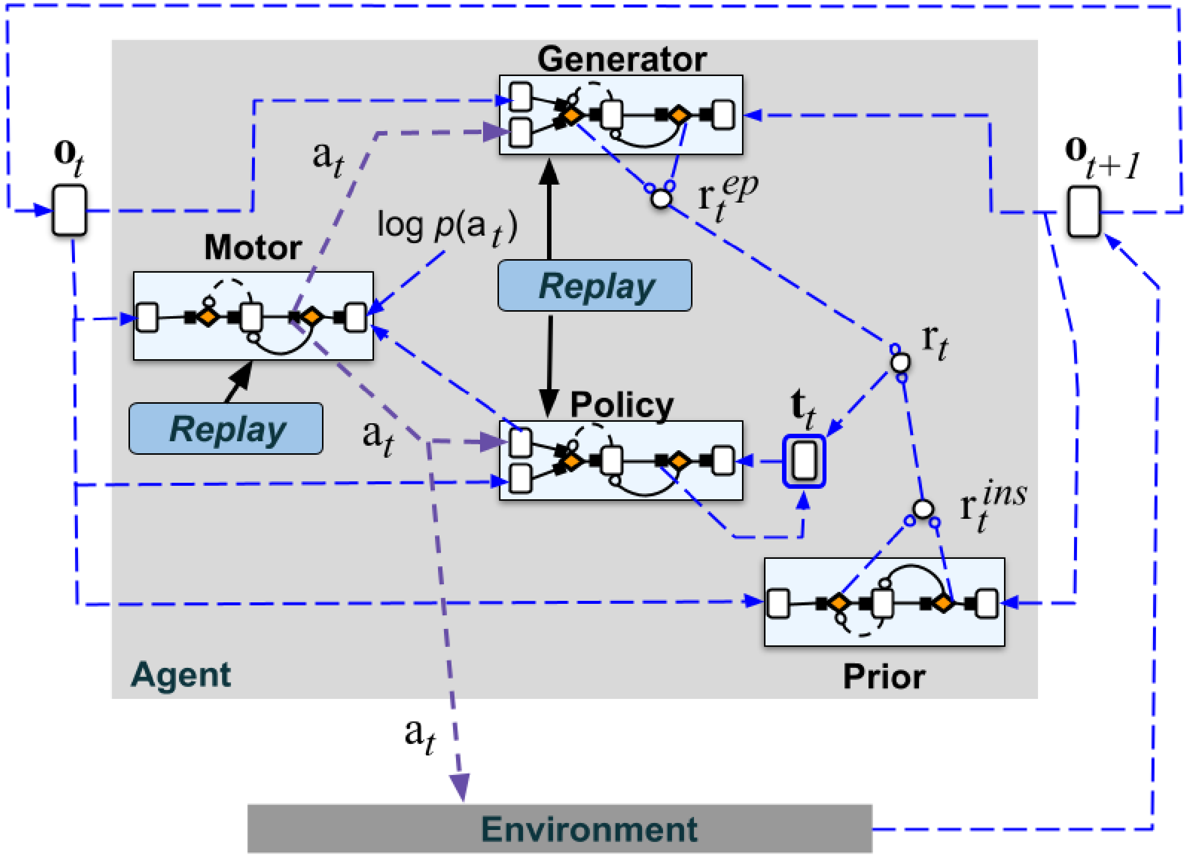

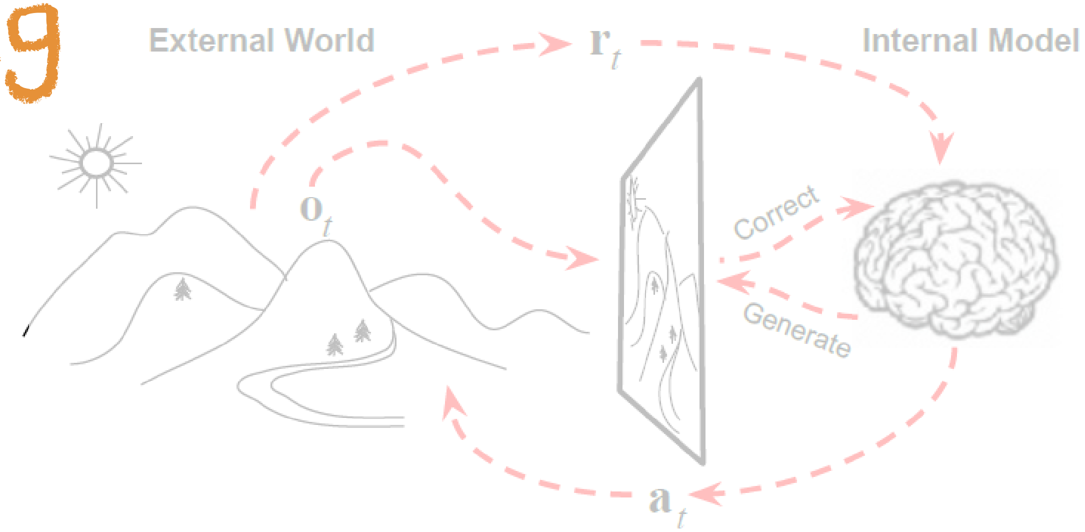

Active predictive coding (*ActPC*) agent for robotic control - one circuit for:

1) prior preference,
2) generative transition dynamics,
3) policy, and
4) motor-action

Ororbia, AG., Mali A.. "Backprop-free reinforcement learning with active neural generative coding." 2022.

Ororbia, AG., Mali A.. "Active predictive coding: Brain-inspired reinforcement learning for sparse reward robotic control problems." 2023.

* Rao, RPN, et al. "Active predictive coding: A unifying neural model for active perception, compositional learning, and hierarchical planning." 2023.

Memory-Augmented Neuronal Dynamics

- Working memory-augmented PC:

$$\bar{\mathbf{z}}^{\ell} = g^{\ell}\left(\mathbf{W}^{\ell+1} \cdot \phi^{\ell+1}(\mathbf{z}^{\ell+1}) + \alpha_m(\mathbf{M}^{\ell+1} \cdot \mathbf{m}_t)\right)$$

$$\mathbf{m}_t = \left[(\mathbf{k}_{t-(H-1)}, ..., \mathbf{k}_{t-i}, ..., \mathbf{k}_{t-1}\right] \text{ and } \mathbf{k}_t = \mathbf{Q} \cdot \mathbf{x}_t$$

H = 7, inspired by:

Miller, GA. "The magical number seven, plus or minus two: Some limits on our capacity for processing information." 1956.

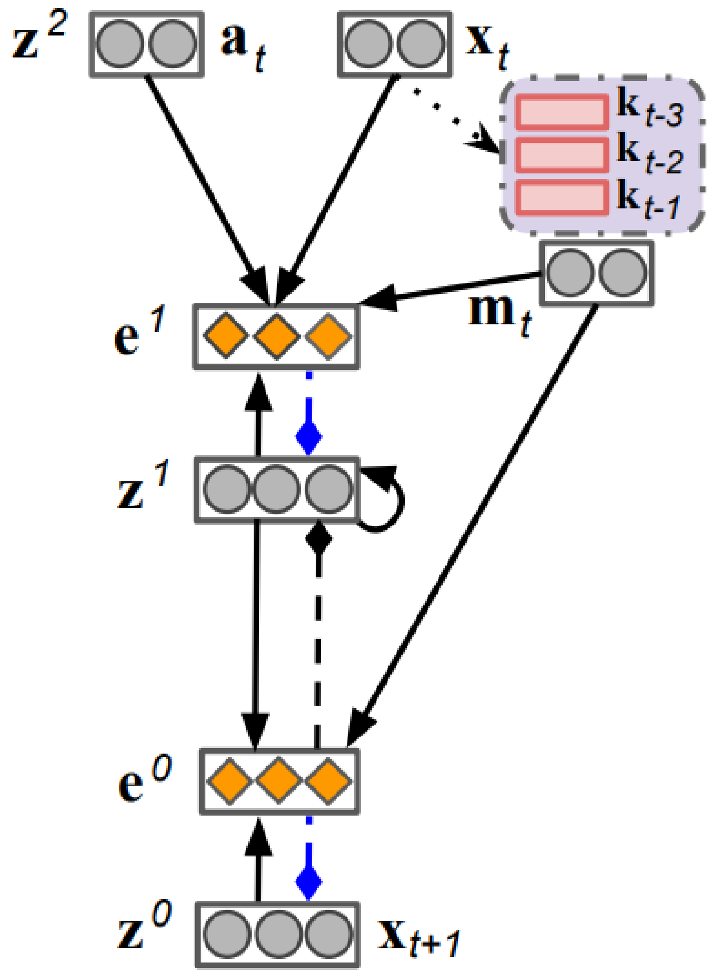

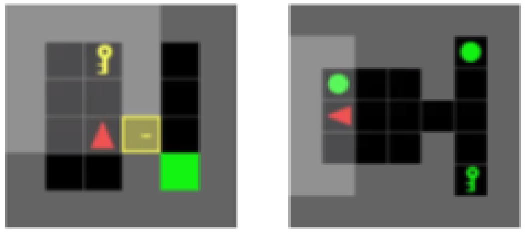

Ororbia, AG, Kelly MA. "Cogngen: Building the kernel for a hyperdimensional predictive processing cognitive architecture." 2022.

Ororbia, AG, Kelly MA. "Maze learning using a hyperdimensional predictive processing cognitive architecture." 2022.

	Avg. Success Rate		Avg. Episode Length	
	DK	Mem	DK	Mem
DQN	0.00	40.0	100.0	41.14
RnD	100.0	48.5	3.71	2.78
BeBold	100.0	48.0	3.93	2.92
CogNGen	100.0	98.5	5.48	2.96

General CogNGen

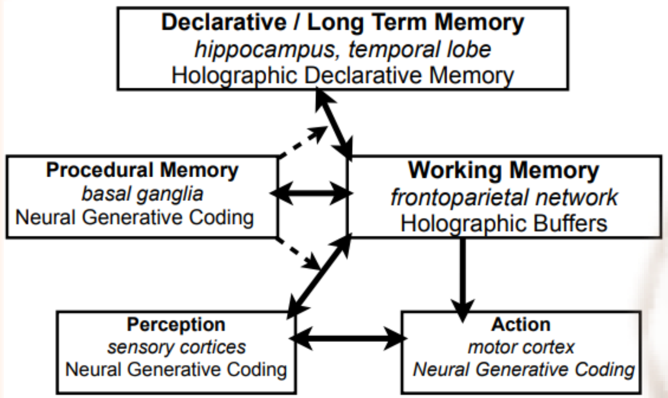

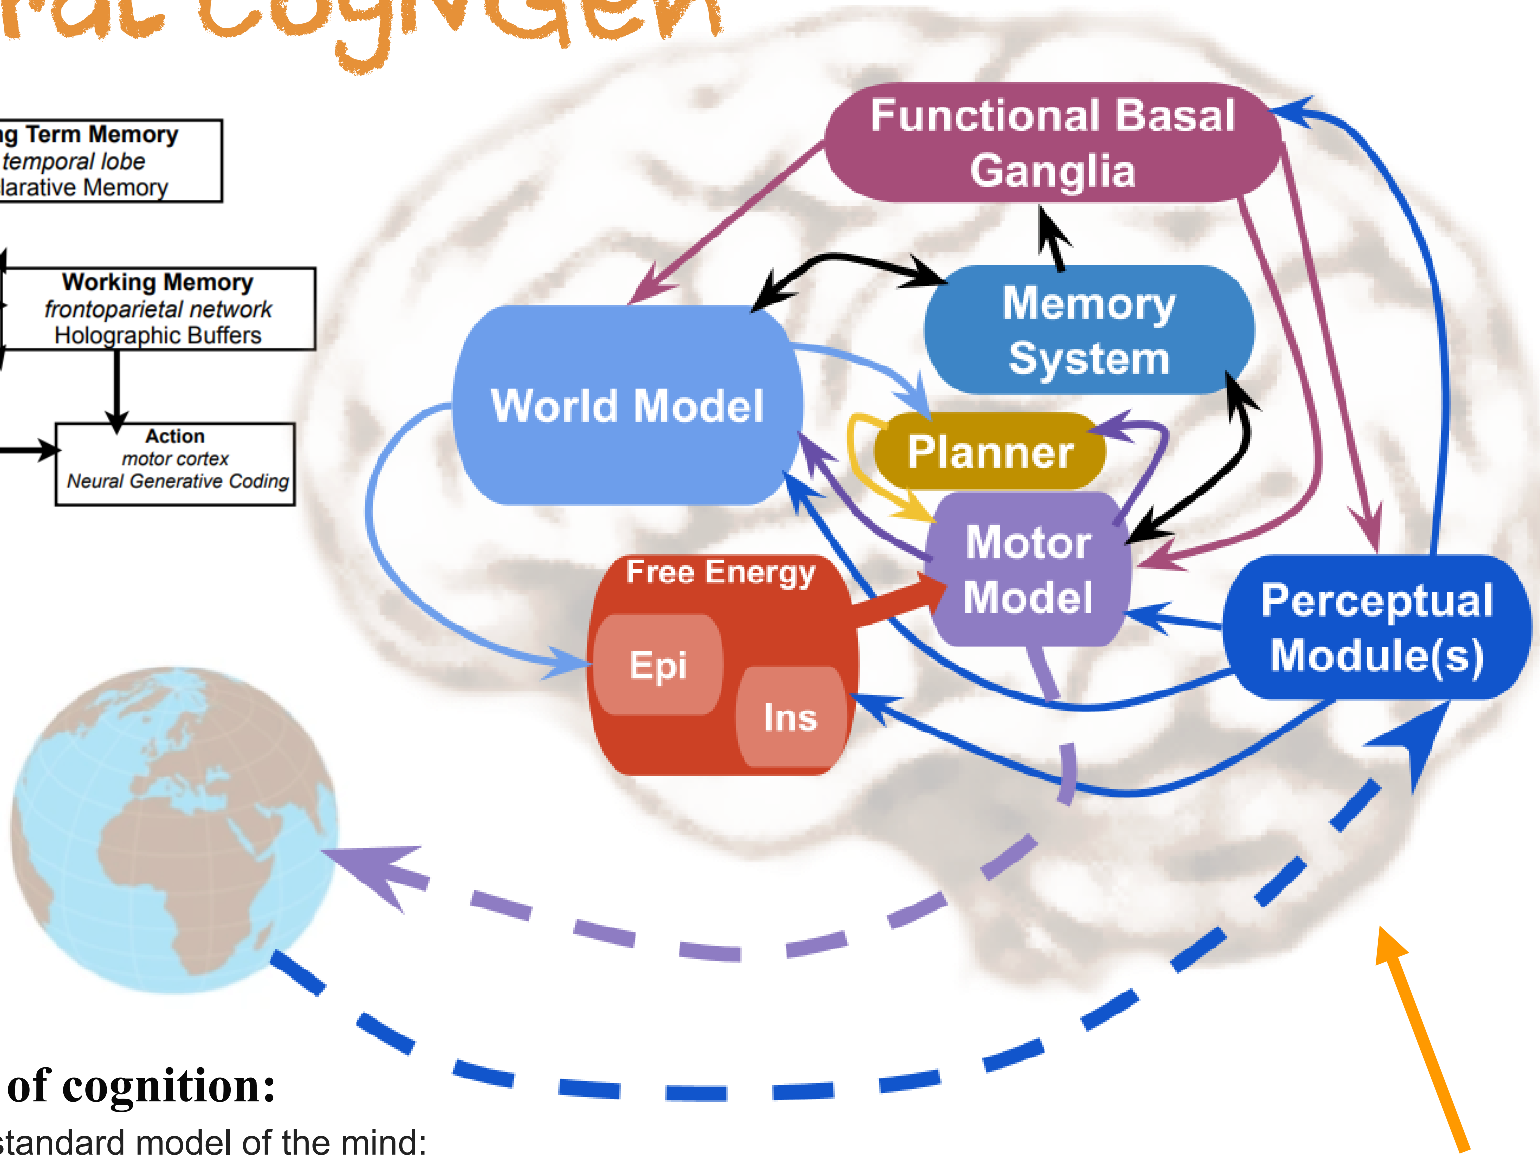

Functional Basal Ganglia

Performs system (task) configuration via excitation, inhibition, & gating

Memory System

Interacts w/ world model to induce episodic recall/replay. Stores triggers & other high-level key activity vectors

Planner

Interacts w/ world & motor models to synthesize action plans

World Model

Learns from perceptual & motor modules an abstract generative model of the agent's environment

Motor Model

Produces the actions (performs action selection) & interacts with the planner.

Free Energy

Computes problem-specific free energy signals (contains the prior model) & exploration-driving signals (via interaction w/ the world model)

Perceptual Modules

Learn to encode/map sensory samples of the environment to a compressed latent vector space.

Common model of cognition:

Laird, JE., et al. "A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics." 2017.

Free Energy Common Model of Cognition:

Ororbia, AG., Kelly, MA. "A neuro-mimetic realization of the common model of cognition via hebbian learning and free energy minimization." 2023.

CogNGen (red arrow agent) navigating a procedurally-generated multi-room environment!

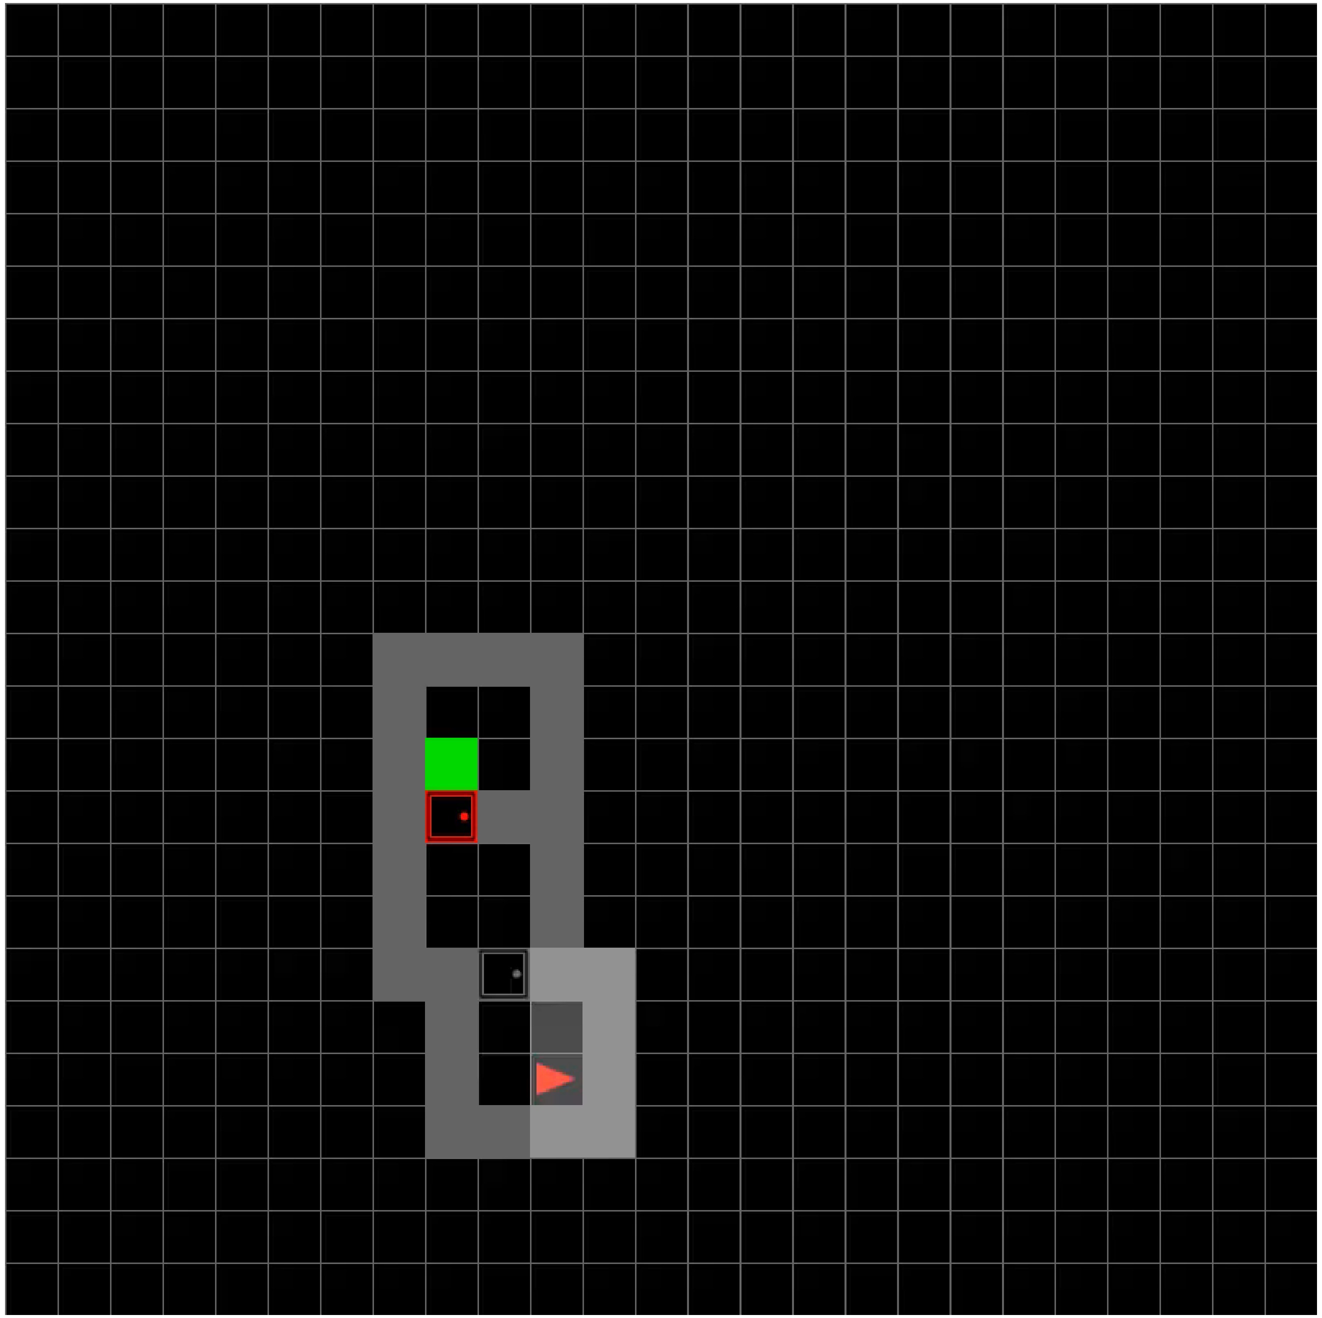

Goal is to get to the green square by moving and learning to open doors on the way!

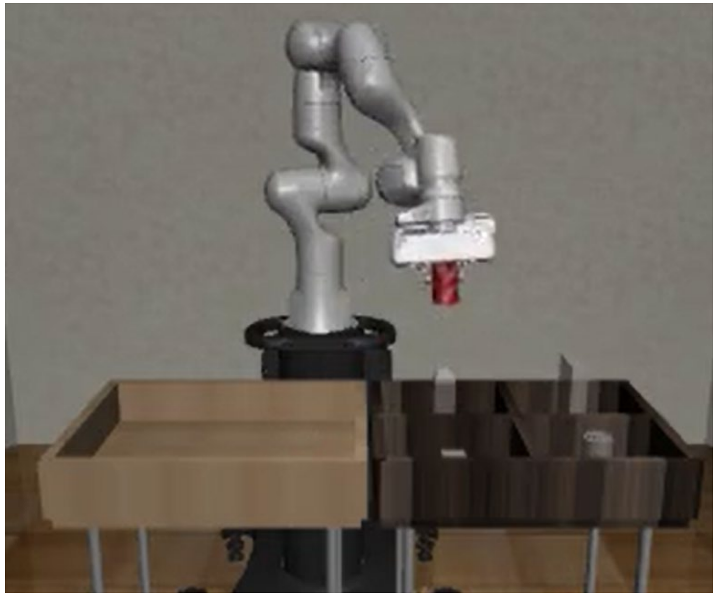

TABLE II: Robosuite results (5-trial mean/std. dev. reported).

Block Lift	**Acc**	**R-Stability**
BC	100.0 ± 0.0	$--$
BC-RNN	100.0 ± 0.0	$--$
BCQ	100.0 ± 0.0	$--$
CQL	56.7 ± 40.3	$--$
HBC	100.0 ± 0.0	$--$
IRIS	100.0 ± 0.0	$--$
DDPG-Demo	63.5 ± 7.8	0.340 ± 0.043
ActPC	96.5 ± 2.1	0.048 ± 0.008
Can Place	**Acc**	**R-Stability**
BC	86.0 ± 4.3	$--$
BC-RNN	100.0 ± 0.0	$--$
BCQ	62.7 ± 8.2	$--$
CQL	22.0 ± 5.7	$--$
HBC	91.3 ± 2.5	$--$
IRIS	92.7 ± 0.9	$--$
DDPG-Demo	51.5 ± 3.5	0.351 ± 0.079
ActPC	94.0 ± 2.1	0.101 ± 0.028

ActPC beats out most strong imitation learning baselines and gets close to best ones (BC-RNN)

ActPC Learns Control Policies

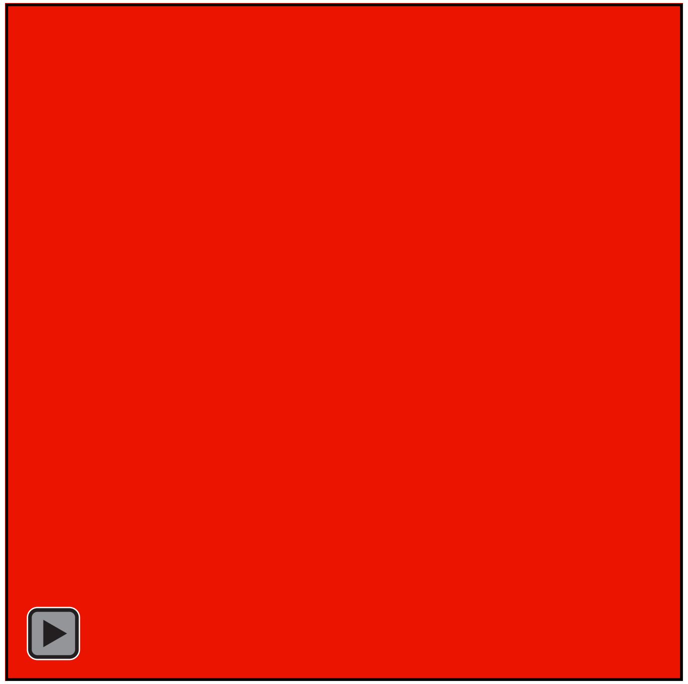

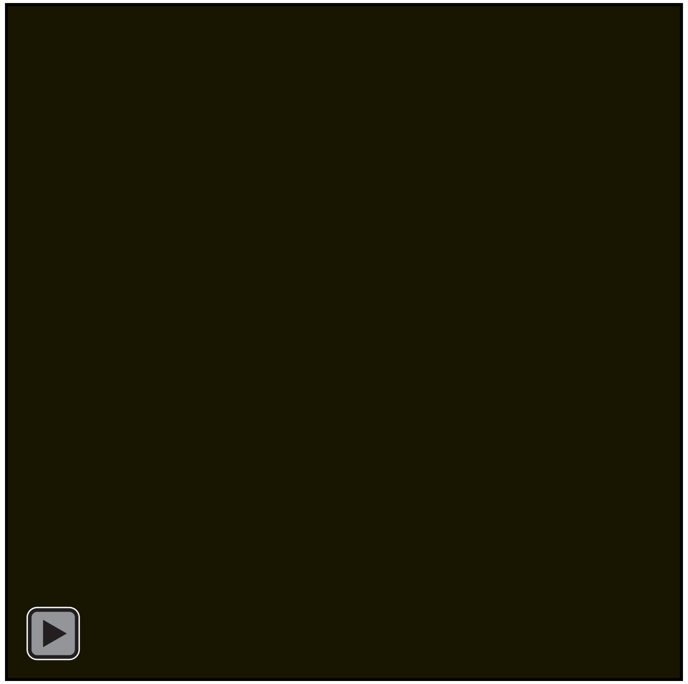

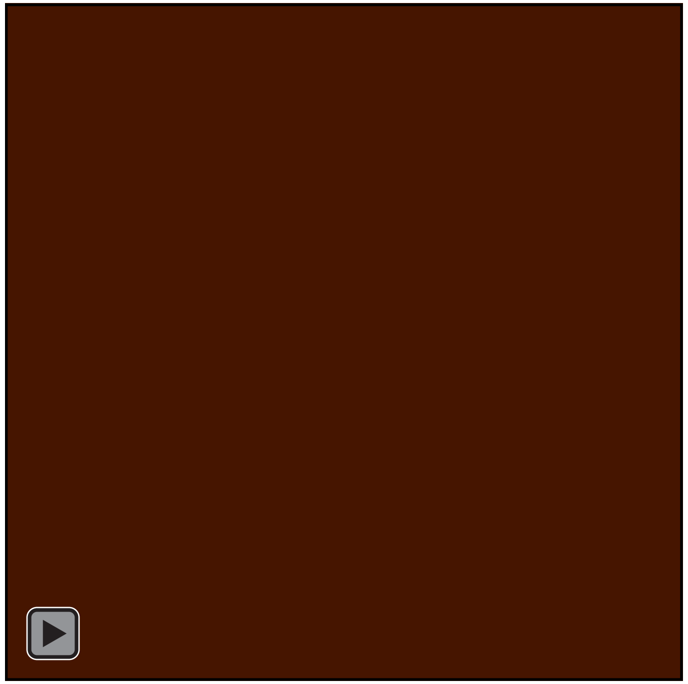

Predictive Coding in Terms of Spikes

Ballard, D, et al. "A model of predictive coding based on spike timing." 1999.

Spiking Neural Coding

- Spike-based predictive coding
  - Spike emission functions = leaky integrate-and-fire (LIF), Izhikevich, Hodgkin-Huxley models
  - **Key**: activation ***traces*** (correspond to concentrations of neurotransmitter bound to synaptic receptors)

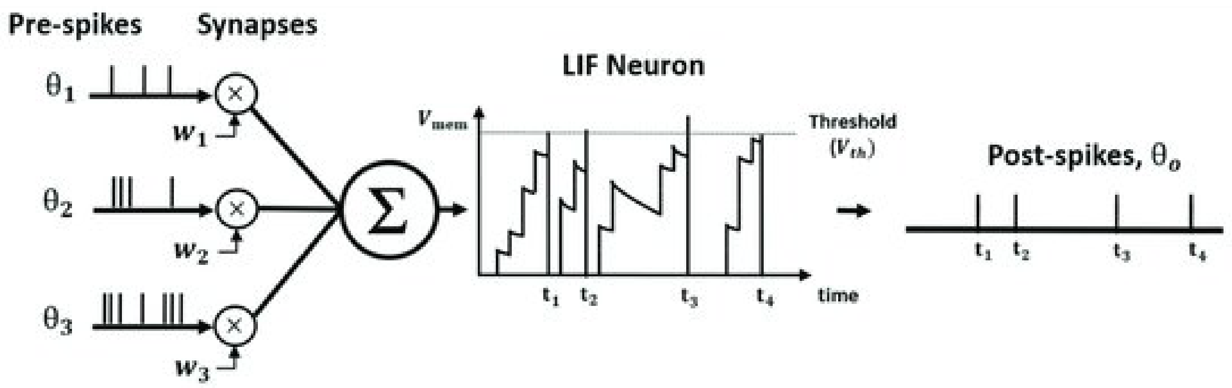

$$\tau_m \frac{\partial \mathbf{v}^\ell}{\partial t} = -\gamma_m \mathbf{v}^\ell(t) + R_m \mathbf{J}^\ell(t) \qquad \mathbf{s}^\ell(t) = \mathbf{v}^\ell(t) \geq \mathbf{v}_{thr}$$

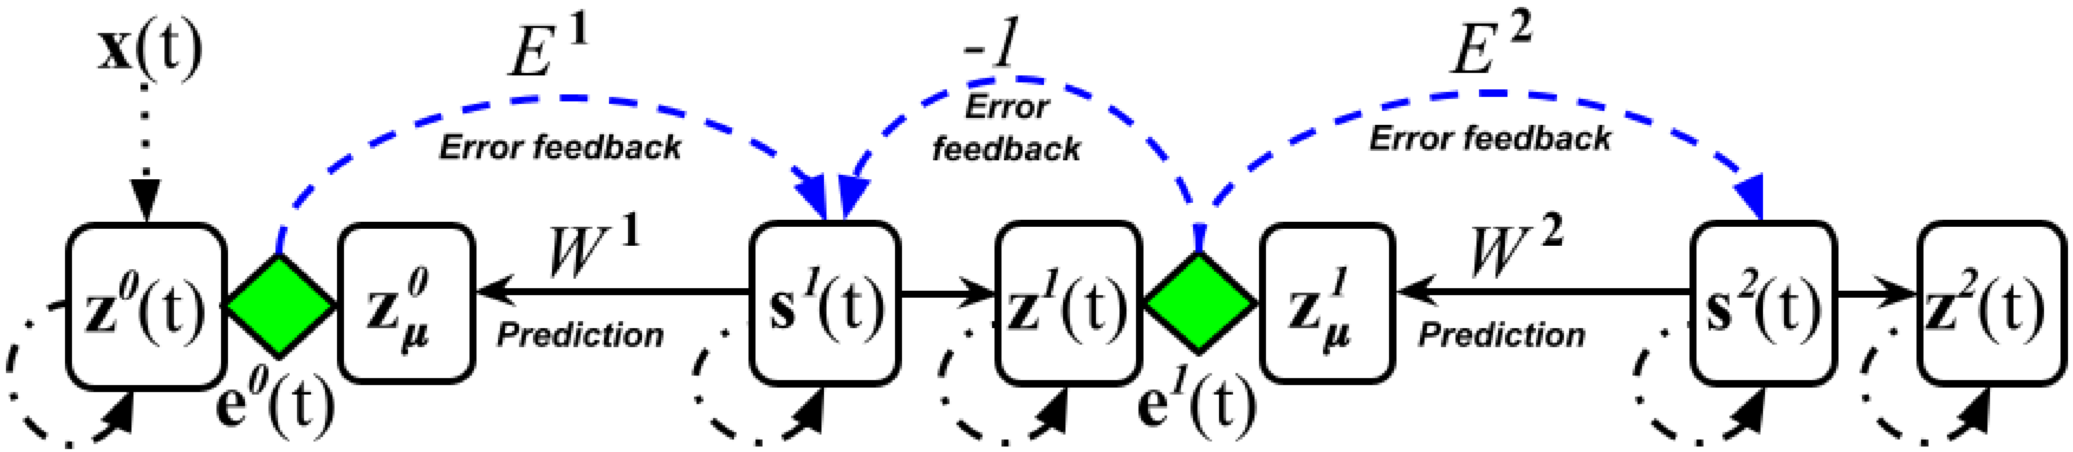

Ororbia, AG. "Spiking neural predictive coding for continually learning from data streams." 2023.

Step 1

Generate local hypotheses; deposit hypothesis signals into error neurons

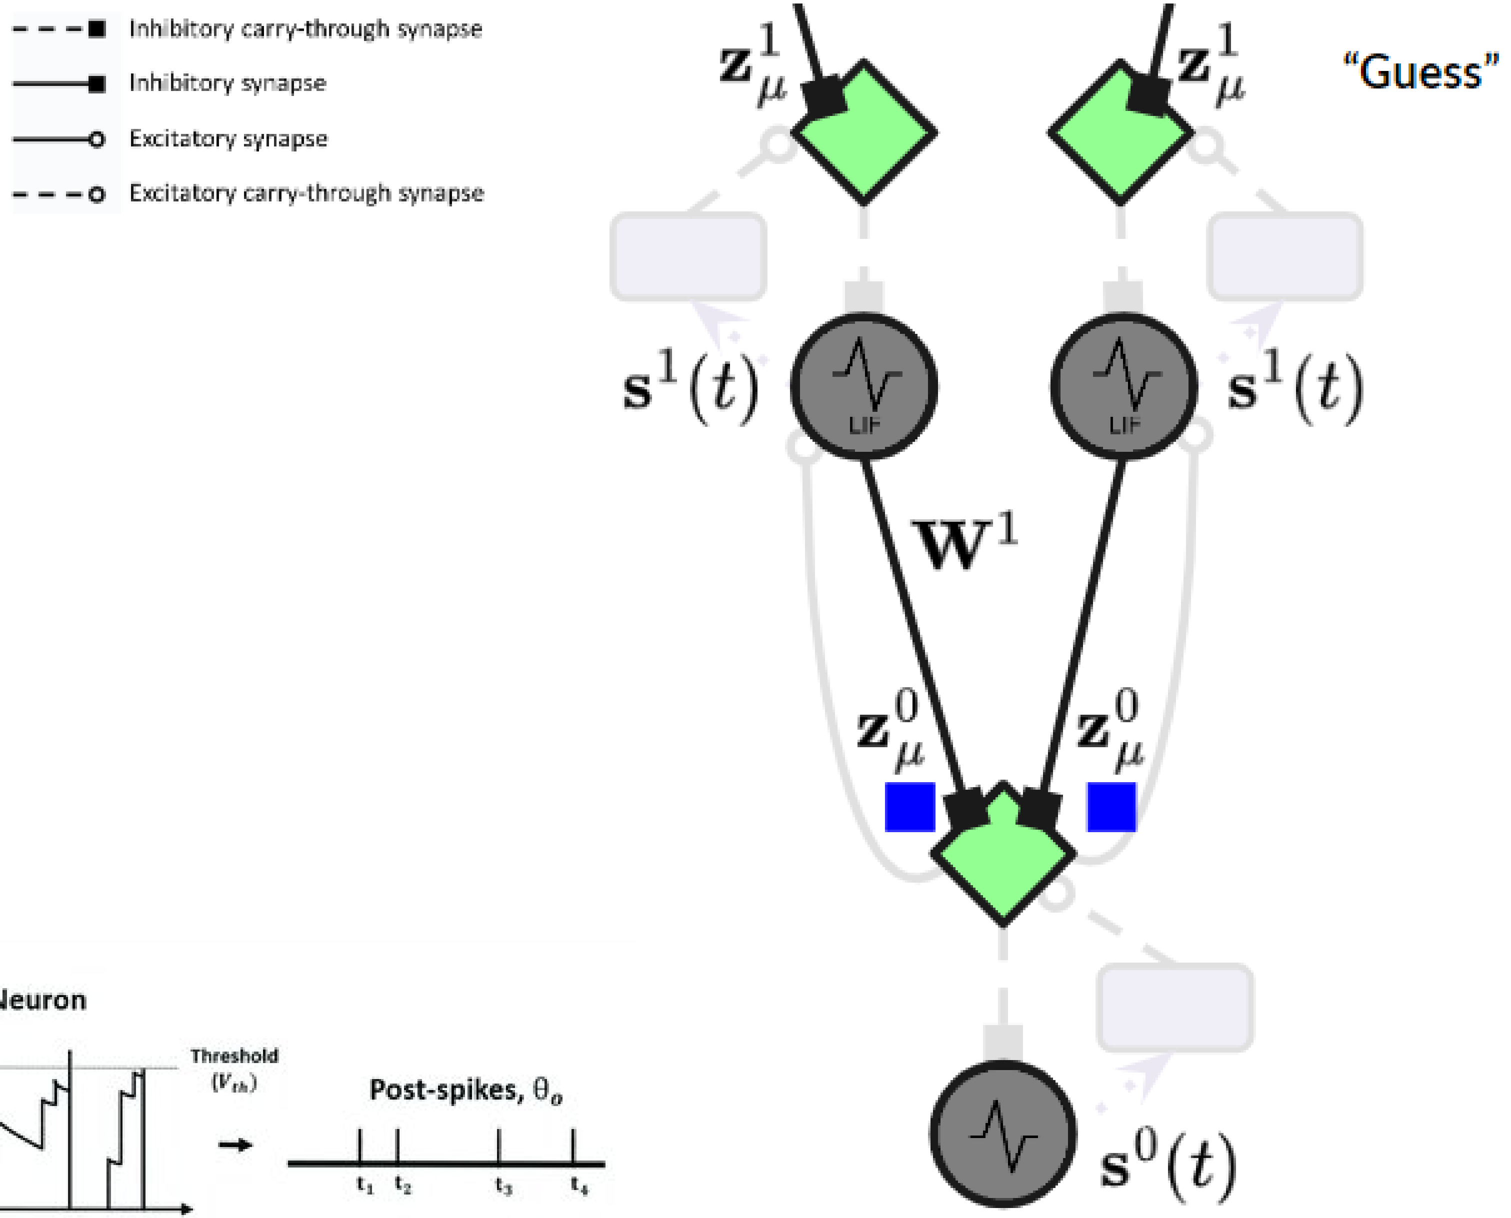

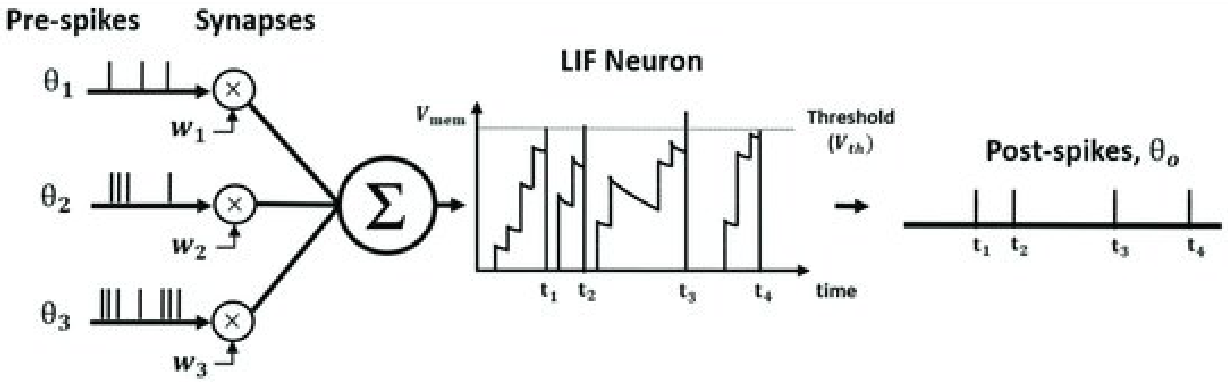

Step 2

Update neurotransmitter concentration traces

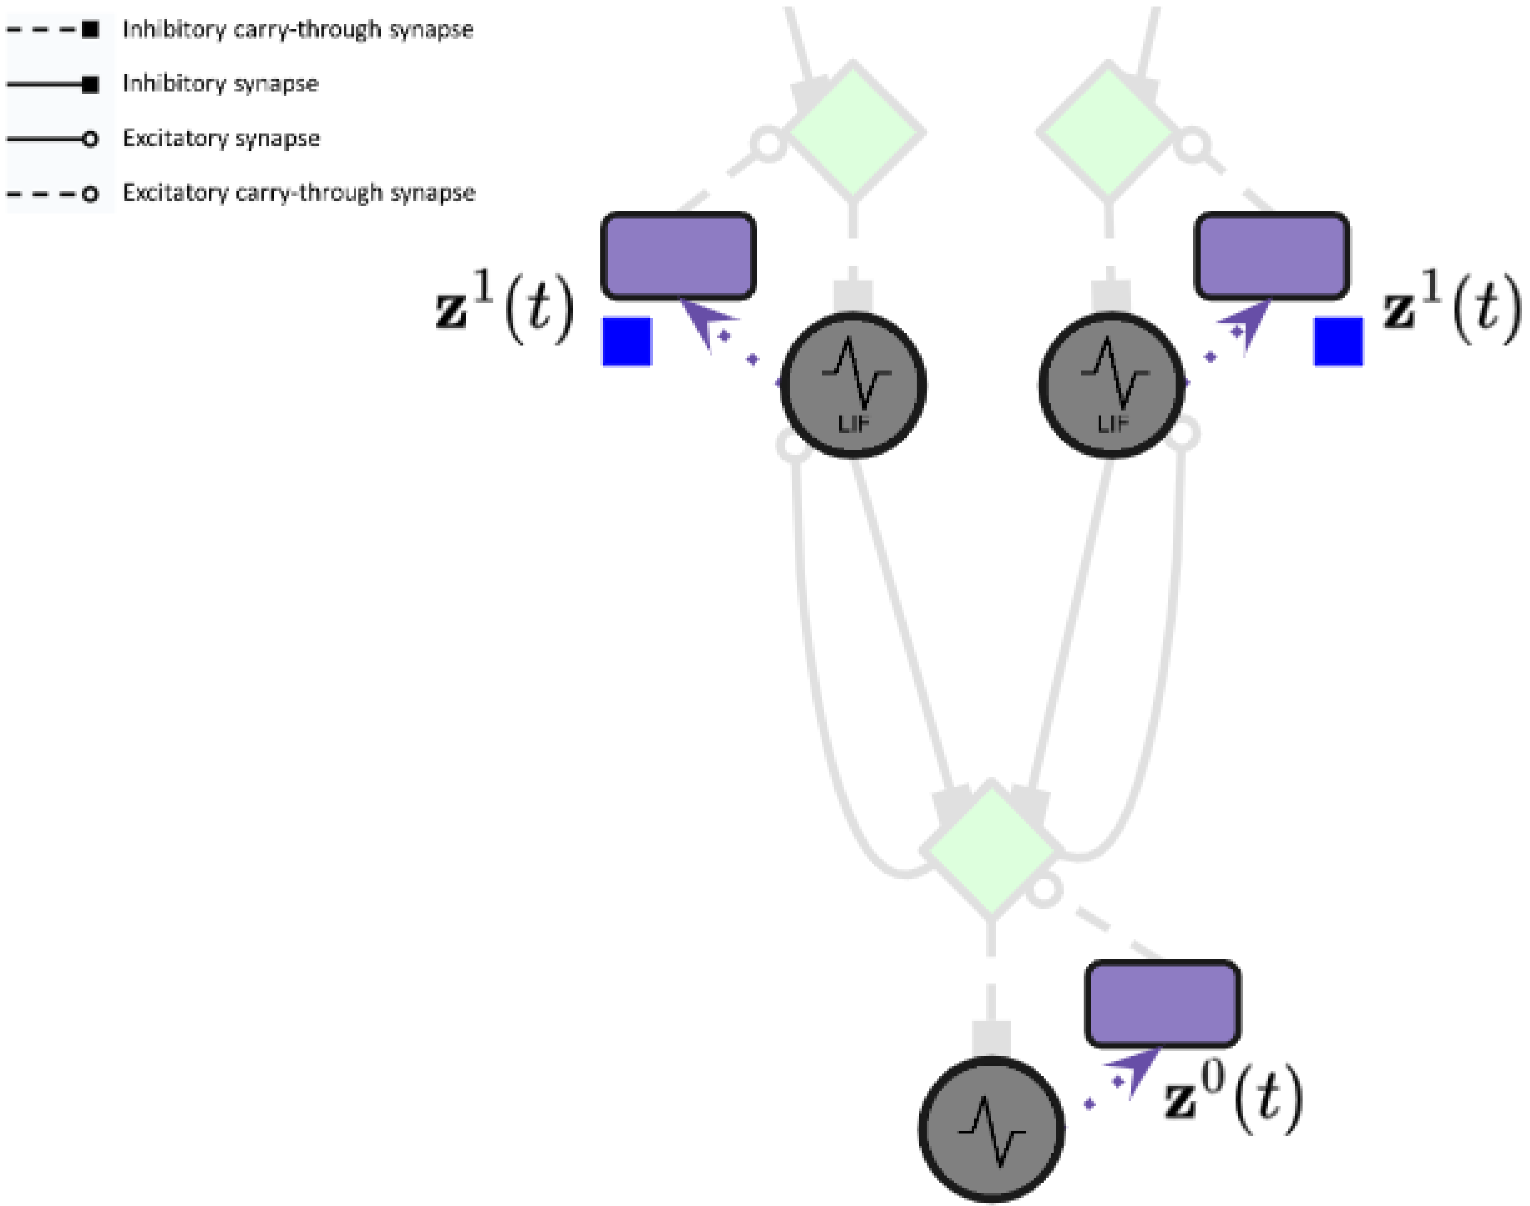

Step 3

Deposit top-down & bottom-up signals into SNM membrane (electrical current input)

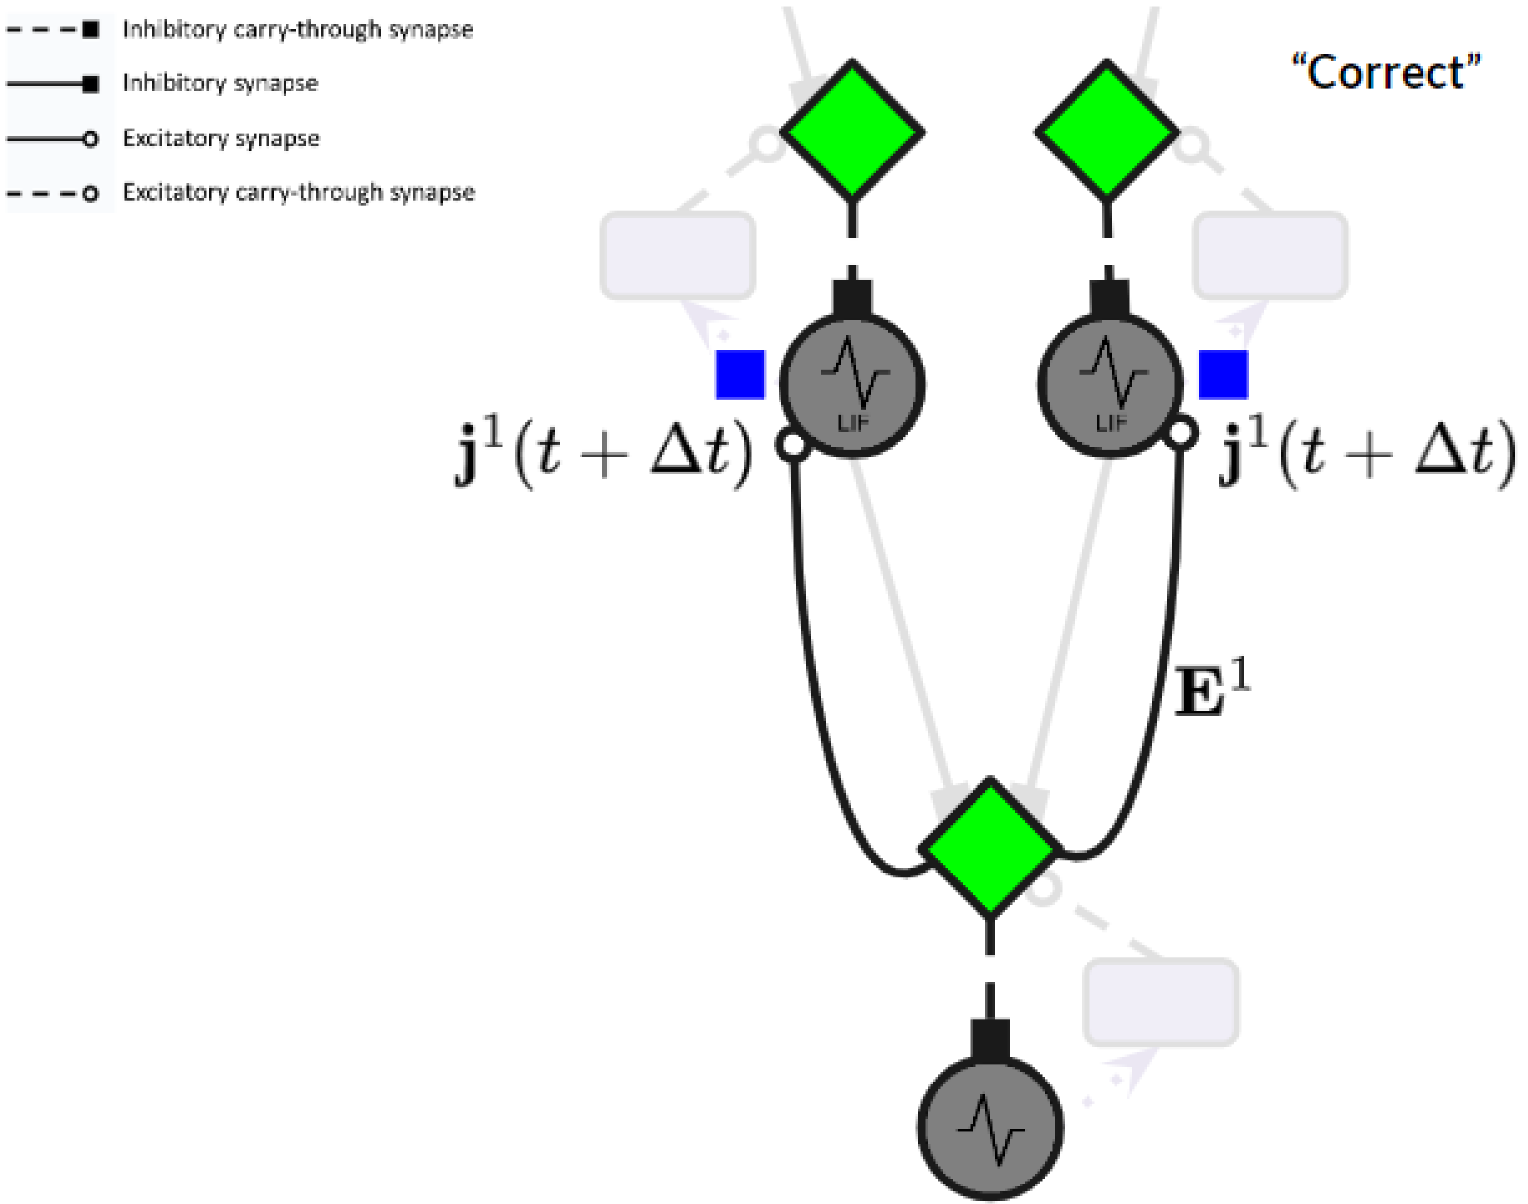

Step 4

Update membrane potential

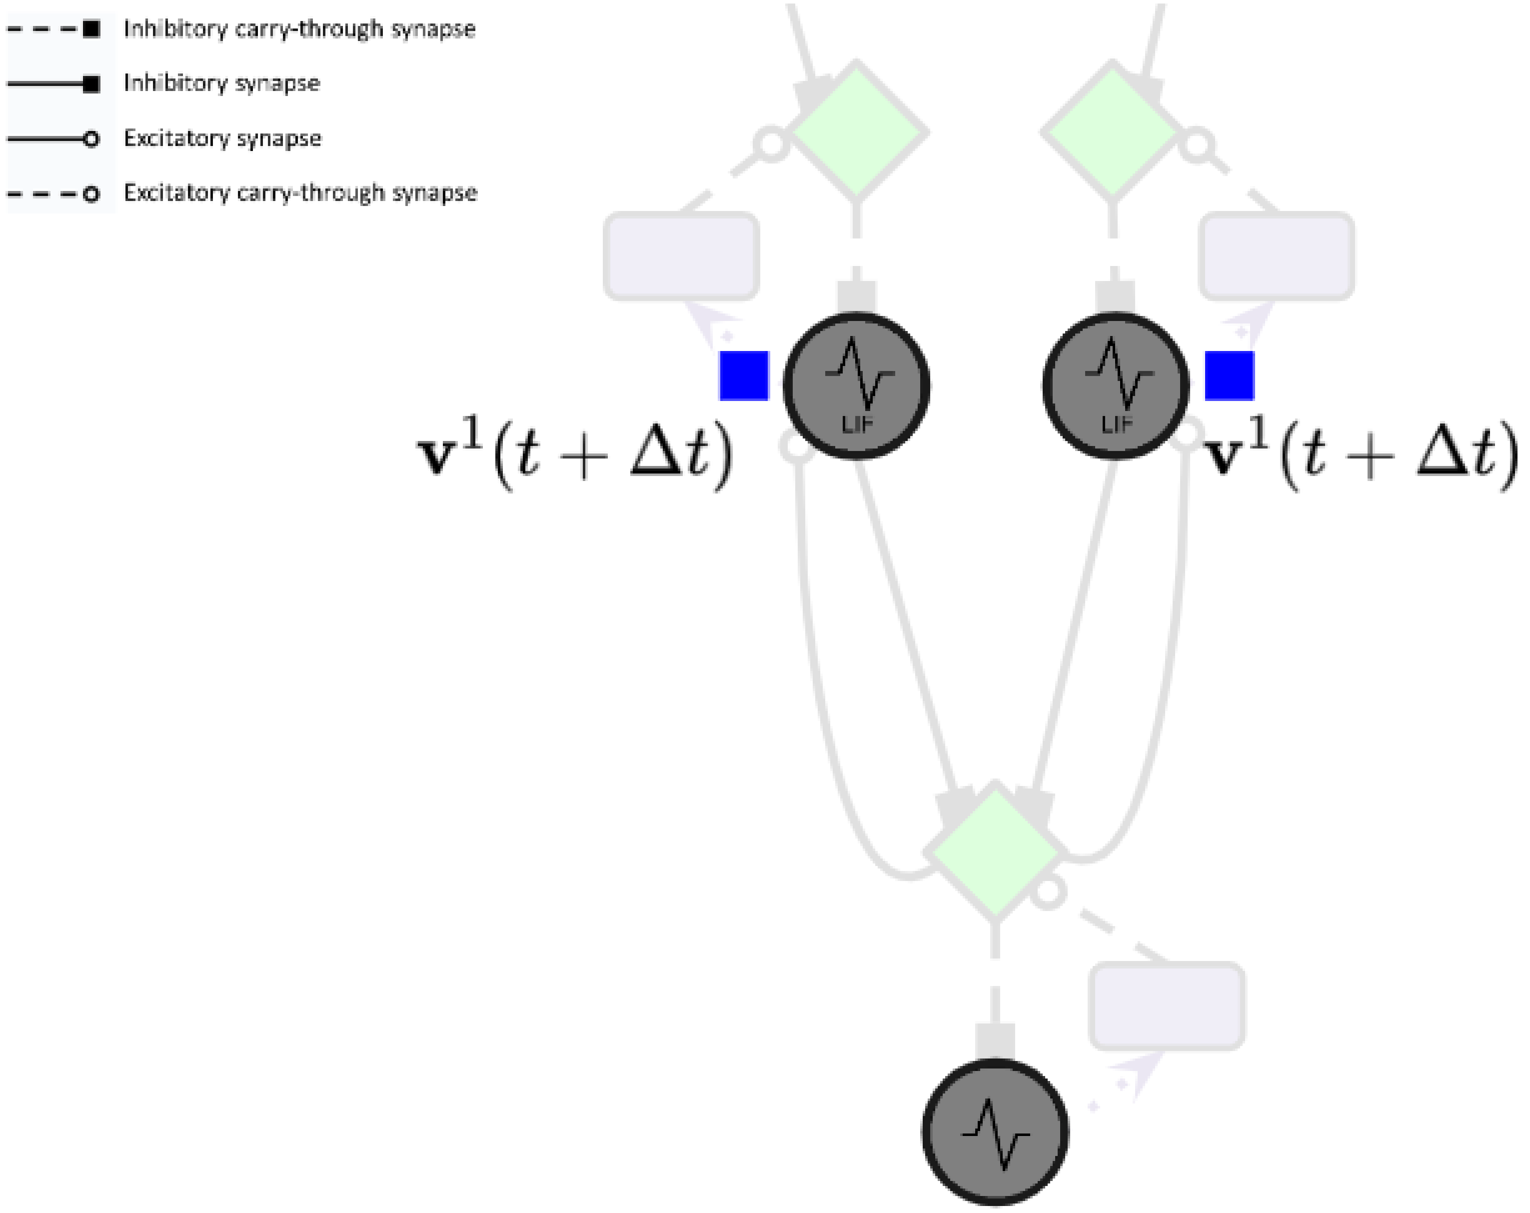

Step 5

Produce action potentials and **hyperpolarize**

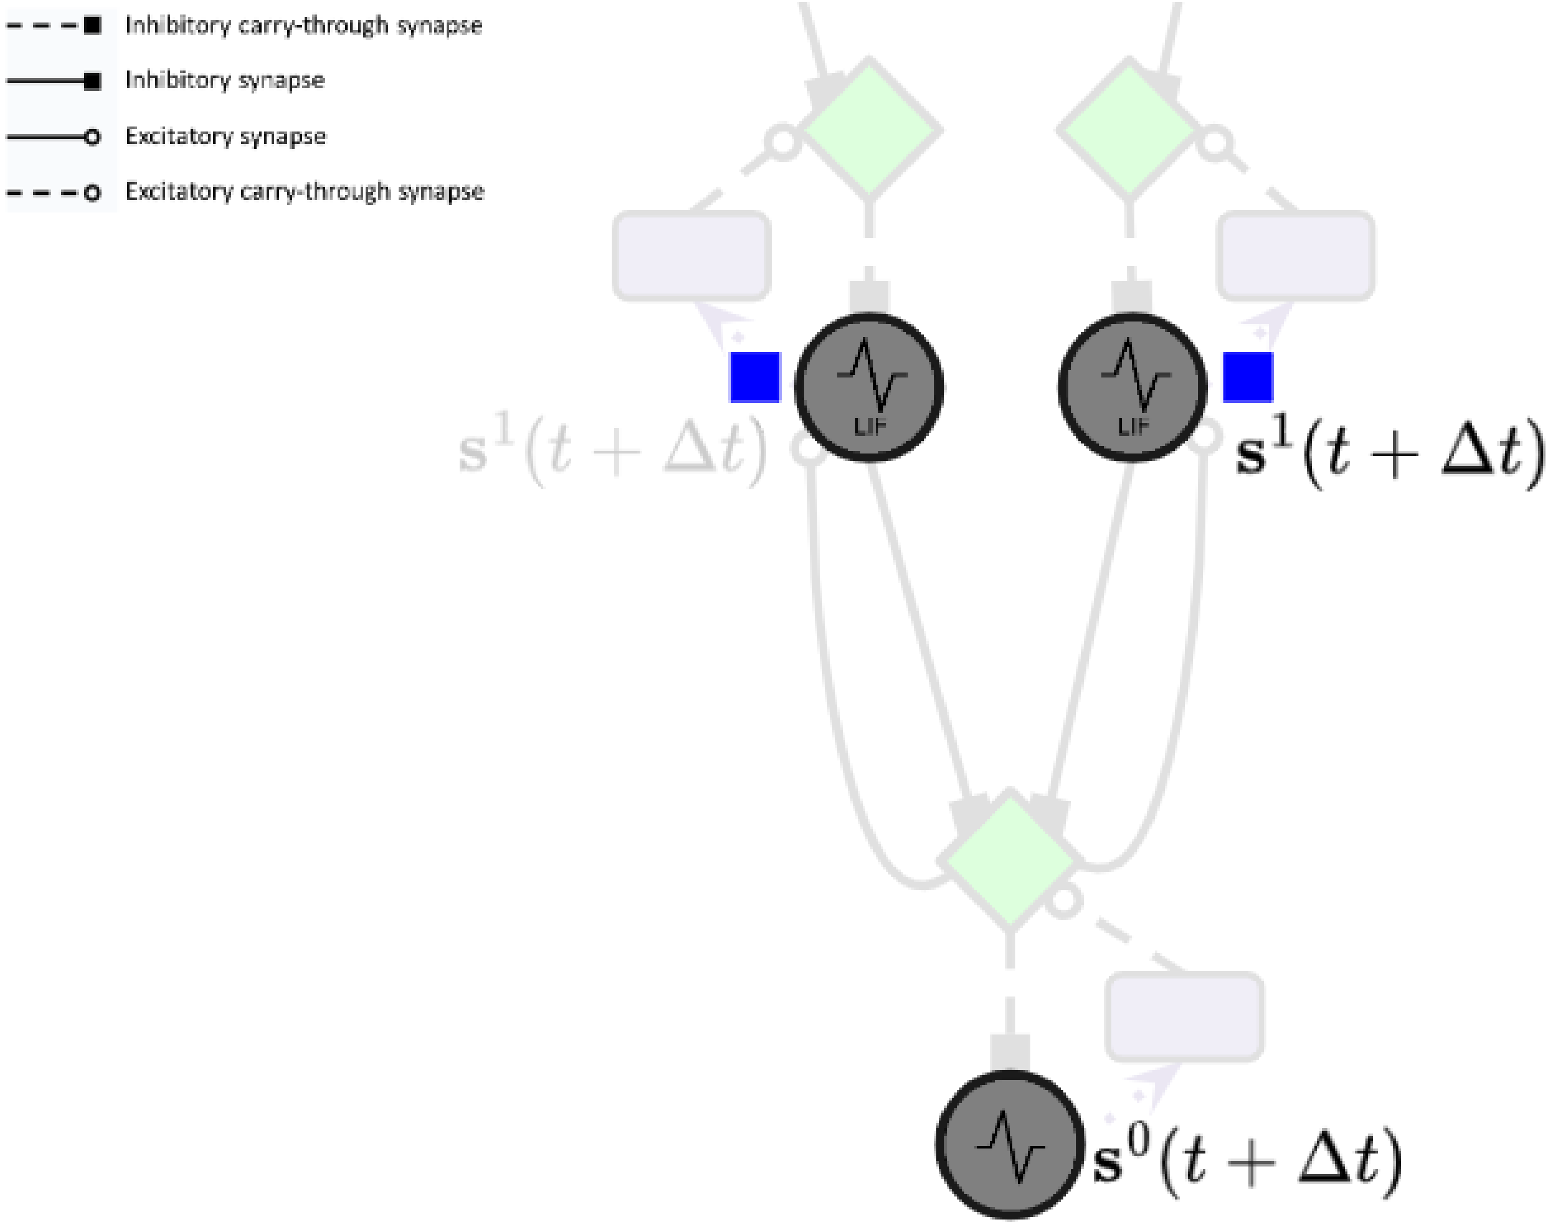

Go Back to:

(Step 1)

Generate local hypotheses…we have returned to Step 1 (but a step forward in time)

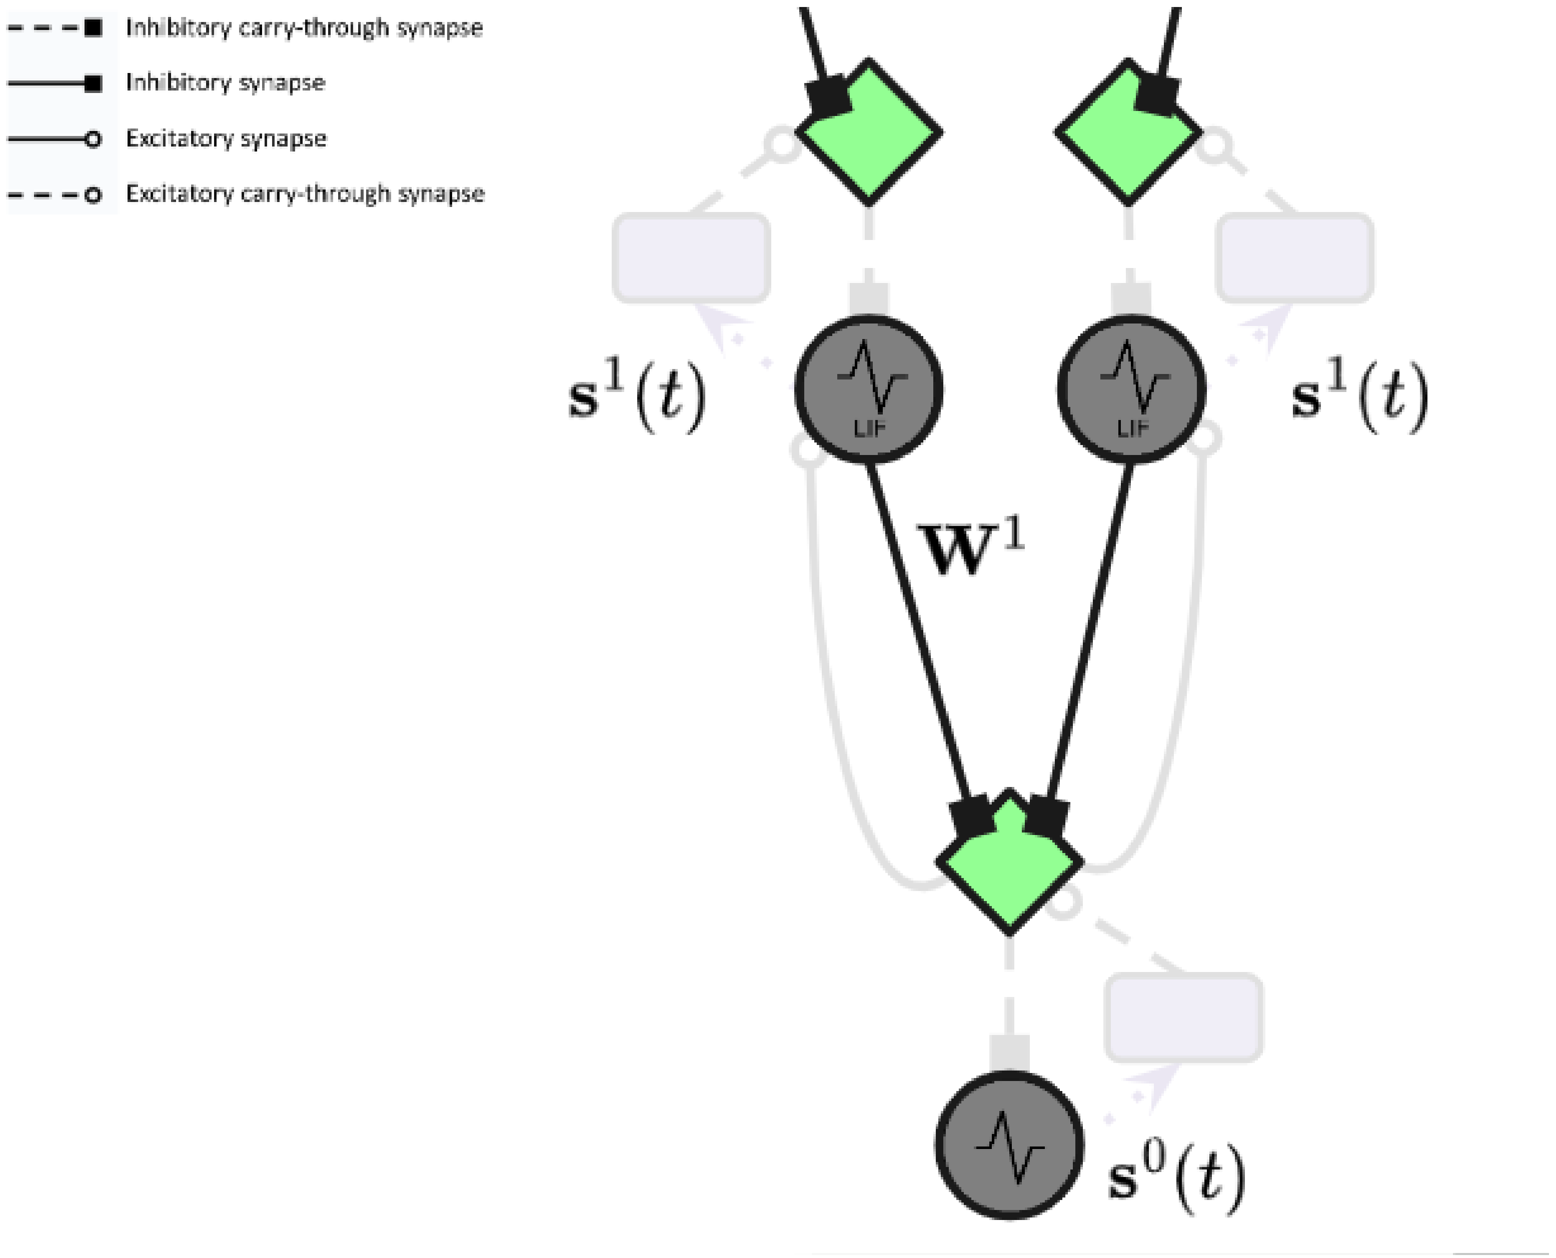

Synaptic Plasticity Dynamics

- (For circuit in last several slides, equations apply to any layer ℓ)
- Simple (error) Hebbian updates; local in space and time

$$\tau_w \frac{\partial \mathbf{W}^1(t)}{\partial t} = -\gamma_w \mathbf{W}^1(t) + \left(\mathbf{e}^0(t) \cdot \left(\mathbf{s}^1(t)\right)^T\right)$$

$$\tau_e \frac{\partial \mathbf{E}^1(t)}{\partial t} = -\gamma_e \mathbf{E}^1(t) + \left(\mathbf{s}^1(t) \cdot \left(\mathbf{e}^0(t)\right)^T\right)$$

Cross-Task Learning with the SpNCN

	MNIST		NotMNIST		FMNIST	
Model	**ACC**	**BWT**	**ACC**	**BWT**	**ACC**	**BWT**
EWC [51]	0.190 ± 0.030	-0.357	0.186 ± 0.020	-0.361	0.199 ± 0.06	-0.350
SI [107]	0.197 ± 0.110	-0.367	0.161 ± 0.030	-0.370	0.198 ± 0.100	-0.370
Lwf [58]	0.846 ± 0.340	-0.120	0.626 ± 0.091	-0.130	0.875 ± 0.300	-0.130
IMM [56]	0.951 ± 0.018	-0.007	0.925 ± 0.011	-0.006	0.950 ± 0.013	-0.005
GDumb [88]	0.978 ± 0.09	-0.005	0.940 ± 0.080	-0.004	0.973 ± 0.09	-0.006
SpNCN	0.735 ± 0.154	-0.302	0.776 ± 0.228	-0.228	0.8324 ± 0.097	-0.198
SpNCN-Buf	0.943 ± 0.451	-0.020	0.927 ± 0.331	-0.008	0.951 ± 0.329	-0.028
SpNCN-Lat	0.972 ± 0.297	-0.001	0.948 ± 0.311	-0.003	0.985 ± 0.216	-0.001

Model	MNIST Samples
SpNCN	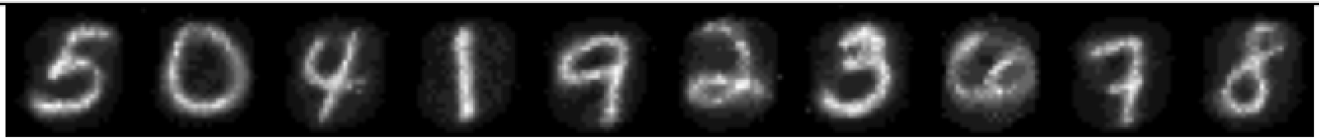

Yes, spiking neural coding is <u>not</u> limited to leaky integrator-and-fire cells…(any cell dynamics model can be used, e.g., FitzHugh–Nagumo model)

A Possible Pathway: Naturalist Machine Intelligence

A Free Energy Pathway

- System w/ non-equilibrium steady-state will behave *s.t.*:
    - Its internal density dynamics are conditionally independent of niche; system state is distinct from niche
    - It continues to self-evidence by returning to non-equilibrium

- *Corollary*: **Active inference**
    - Entity changes relationship with its niche via **action**

- *Corollary*: **Mortal computation**
    - Entity's "software" cannot be divorced from physical substrate
    - Imperative = remain in a non-equilibrium steady state (identity)

- Machine intelligence should be: ***elementary, embodied, enactive, embedded, and extended*** (5E Theory)

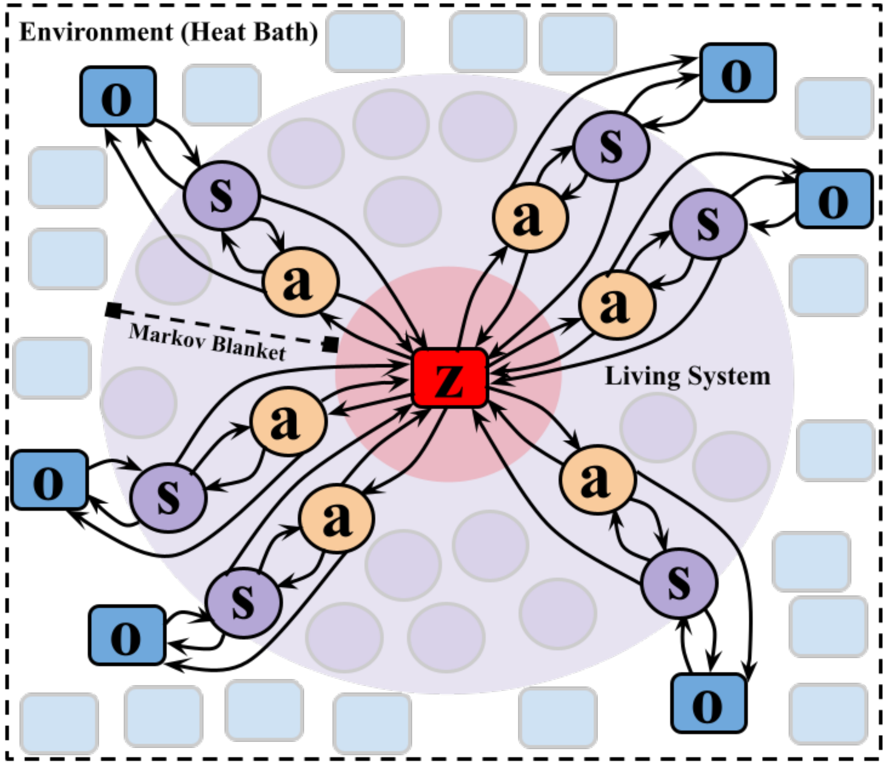

System enclosed by Markov blanket will evolve its internal generative model to minimize its variational free energy

Ororbia, AG, Friston, K. "Mortal computation: A foundation for biomimetic intelligence." 2023.

Thermodynamic Motivation for In-Memory Processing

- Inference on today's computers suffers from "memory wall"
  Realizing thermodynamic efficiency of computations requires belief updating **in memory**
  - Pathway for *Green AI* versus *Red AI* [Schwartz et al., 2020]

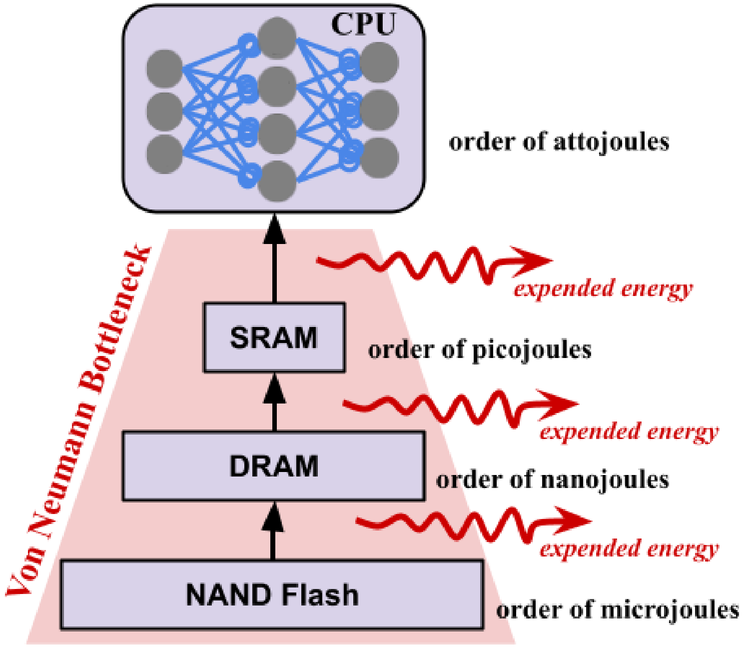

Inference/learning on a von Neumann Computer

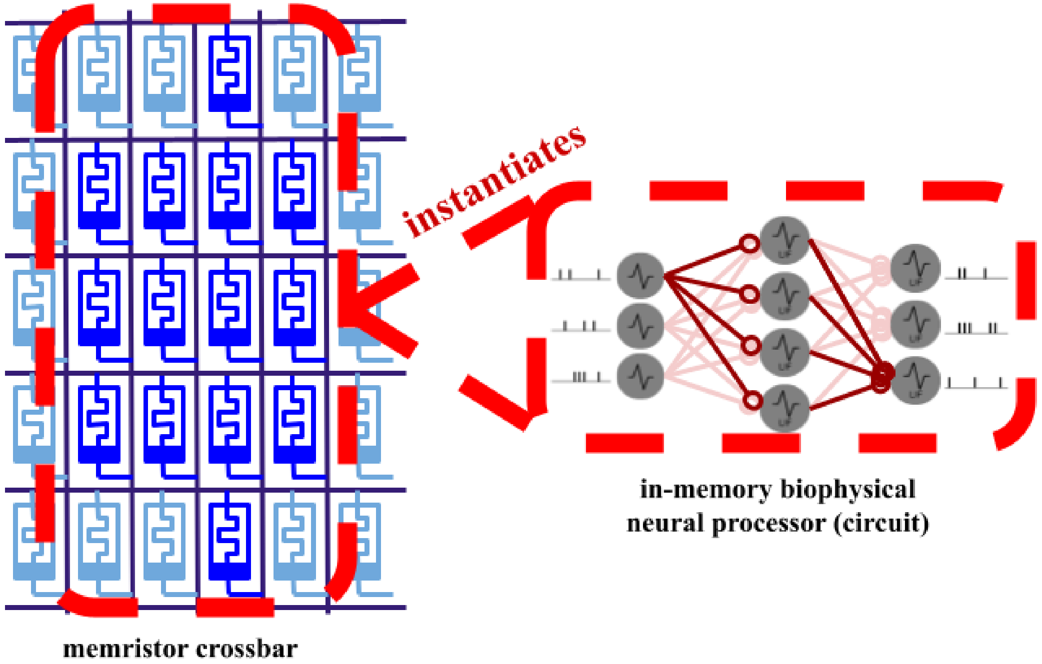

In-Memory (Neuromorphic) Computing

Challenges and Questions?

- ***Benchmarking*** –> we need to focus on foraging/exploration-centric tasks
- Degree of entanglement between architecture, credit assignment, and inference
  - How much architectural agnosticism are we willing to give up?
- Role of "***cognitive architecture prior***" and role of evolutionary processes
- How much neurobiological realism do we need to sufficiently generalize? How much is too much?
  - Enough to realize thermodynamic efficiency & rich temporal encoding properties and statefulness of neuronal systems (neurorobotics)
  - Neuromorphics:
    - Low energy required for computation; low energy for communication
    - Inputs arrive asynchronously, extreme sparsity

Challenges and Questions?

- ***Benchmarking*** –> we need to focus on foraging/exploration-centric tasks
- Degree of entanglement between architecture, credit assignment, and inference
  - How much architectural agnosticism are we willing to gi
- Role of "***cognitive architectu*** nary processes
- How much neu tly generalize? How much is to
  - Enough to reali temporal encoding properties and statefulness s (neurorobotics)
  - Neuromorphics:
    - Low energy required for computation; low energy for communication
    - Inputs arrive asynchronously, extreme sparsity

Guiding Principle: "Make everything as simple as possible, but not simpler."

Challenges and Unresolved Questions

- What is the ground-breaking “app” for PC and biological modeling?
  - There are barriers to wider-spread adoption of deep learning alternatives
- Model selection and sparsification:
  - Can we infer the best model from data using structural adaptation, model selection?
- Neuromorphic Hardware implementations:
  - Analog/memristor circuits?
  - Is there anything more exotic to consider, such as organoids?

Sidestepping Limitations of PC with Other Forms of Learning?

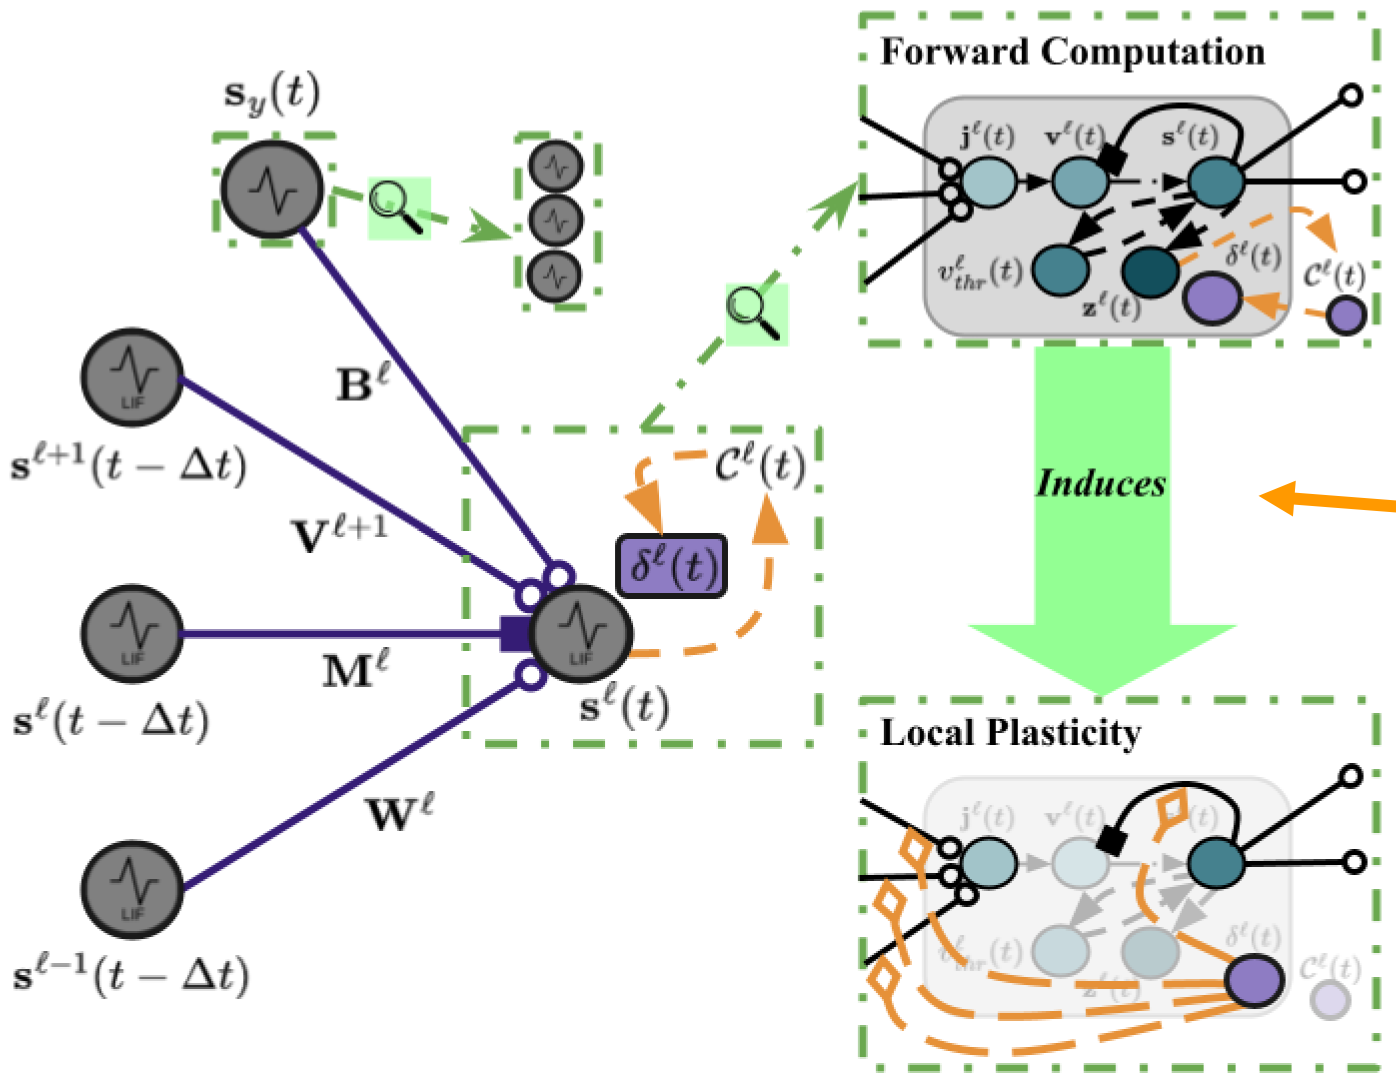

Hinton, G. "The forward-forward algorithm: Some preliminary investigations" 2022.

Ororbia, AG, Mali, A. "The predictive forward-forward algorithm." 2022

Ororbia, AG. "Contrastive-signal-dependent plasticity: Forward-forward learning of spiking neural systems." 2023.

Simulation Software: Building Your Own Biomimetic Models

- ***ngc-learn*** – Python simulation and design library for computational cognitive neuroscience, in ***JAX***
  - https://github.com/NACLab/ngc-learn
  - Supports arbitrary PC circuit design, biological credit assignment development, and spiking neuronal cell modeling
  - Model museum features historical models

NGC-Museum:
https://github.com/NACLab/ngc-museum

NGC-Lava: (for translating to Intel's Lava-NC/Loihi 2)
https://github.com/NACLab/ngc-lava

The NAC Lab

Mission: Create learning algorithms and computational architectures for biomimetic systems, motivated by models of cognition and biological circuitry

- Neurobiological credit assignment
- Predictive coding and processing
- Active inference, biophysical reinforcement learning, neurorobotics
- Spiking neural networks
- Continual machine, learning
- Neural-based cognitive architectures

- **Related Collaborators:**
Ben Goertzel (*SingularityNet*), Karl Friston (*UCL/VERSES*), Chris Buckley (*Sussex/VERSES*), Rajesh P. N. Rao (*UW*), Ankur Mali (*USF*), Daniel Kifer (*PSU*), C. Lee Giles (*PSU*), Hugo Latapie (*Cisco*), Mary Kelly (*Carleton*), Brett Fajen (*RPI*), Tommaso Salvatori (*VERSES*), Travis Desell (*RIT*), Daniel Krutz (*RIT*), Gabriel Diaz (*RIT*), Beren Millidge (*Oxford*), Adam Kohan (*UMass*)

NAC
Neural Adaptive Computing Laboratory

Doctoral Researchers:
William Gebhardt , Zhizhuo Yang, Mobina Ghorbaninejad, Faeze Habibi, Việt Dũng

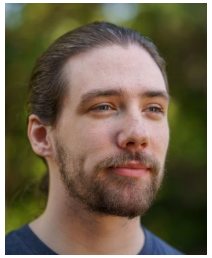



NGC Learn

IQrotech Meta

CISCO NSF

Questions?

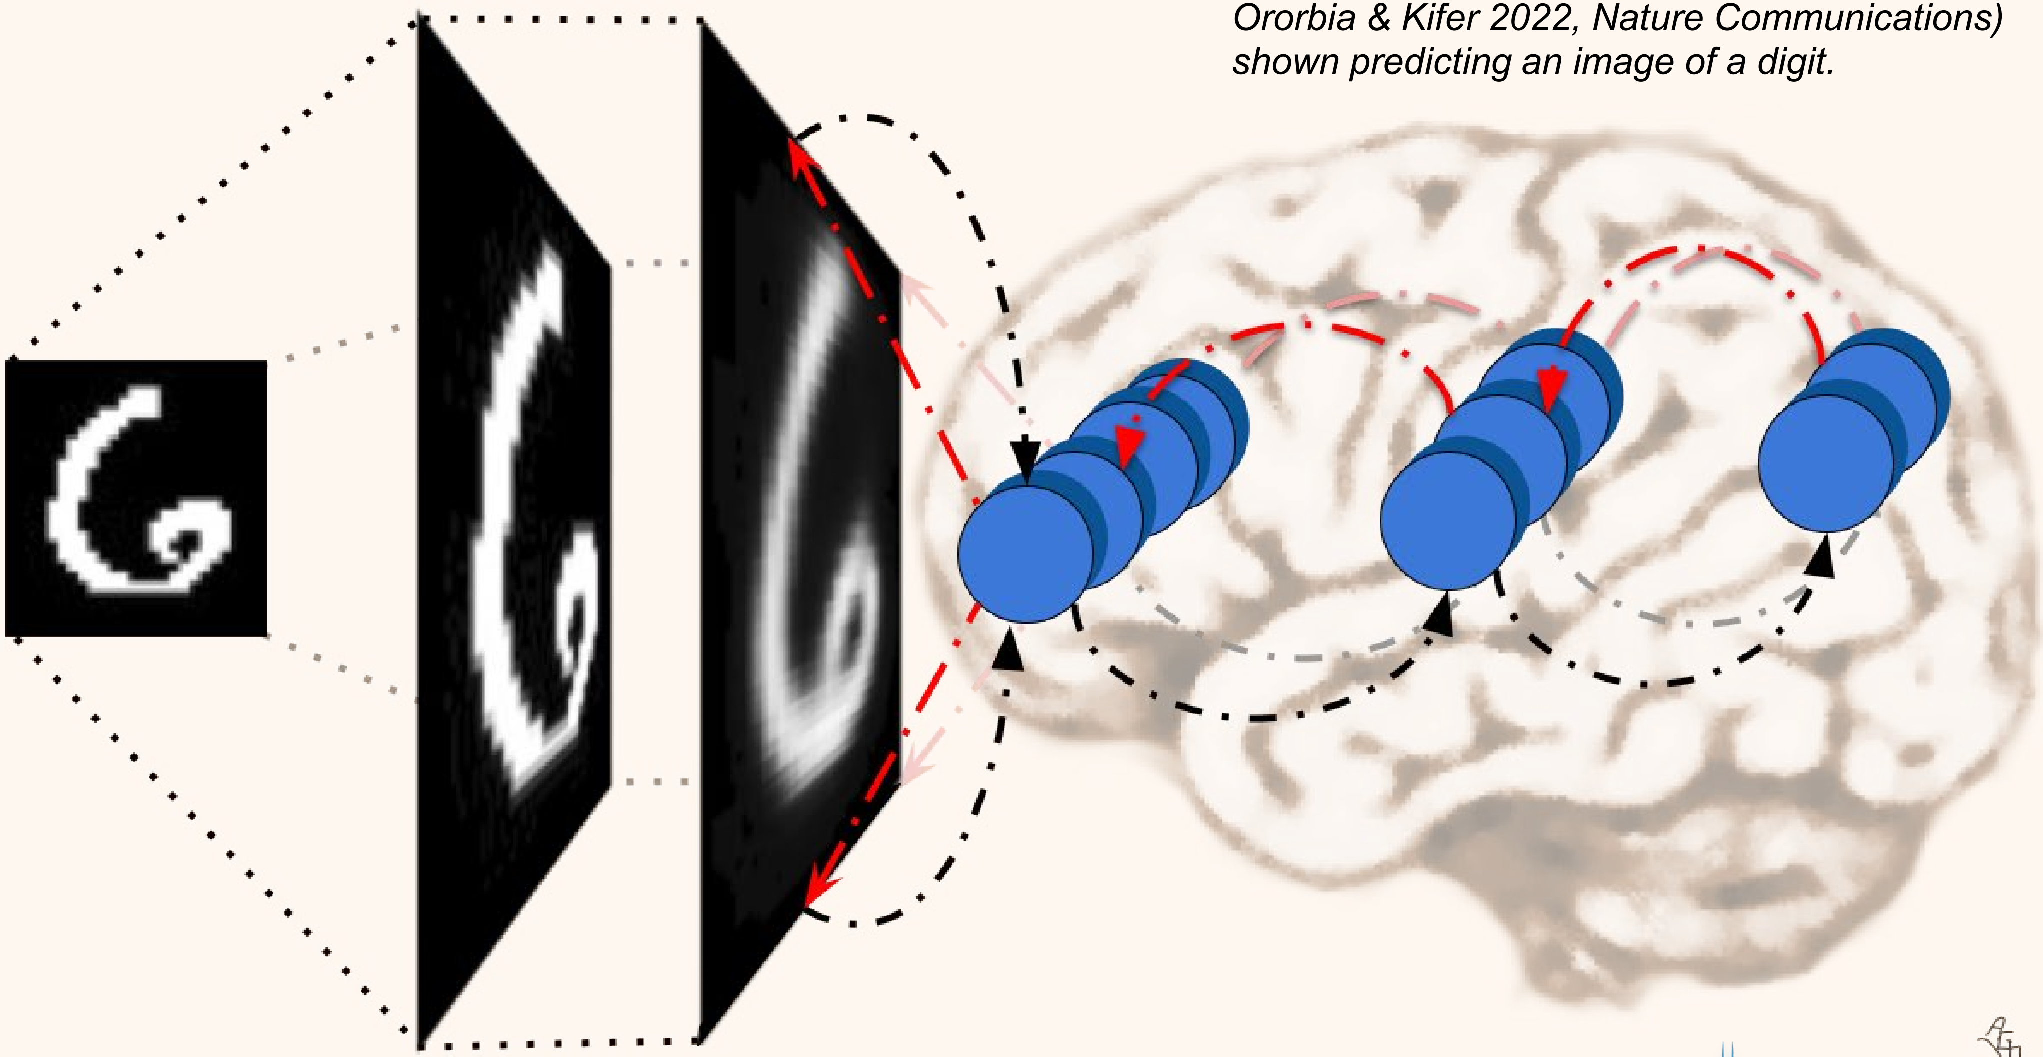

A neural generative coding circuit (NGC; Ororbia & Kifer 2022, Nature Communications) shown predicting an image of a digit.