Learned
Feedback & Feedforward
Perception & Control

Joseph Marino

August 9th, 2021

Caltech

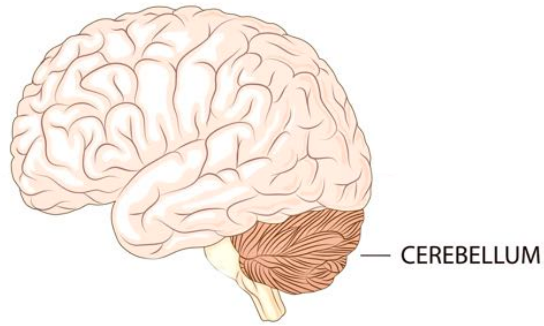

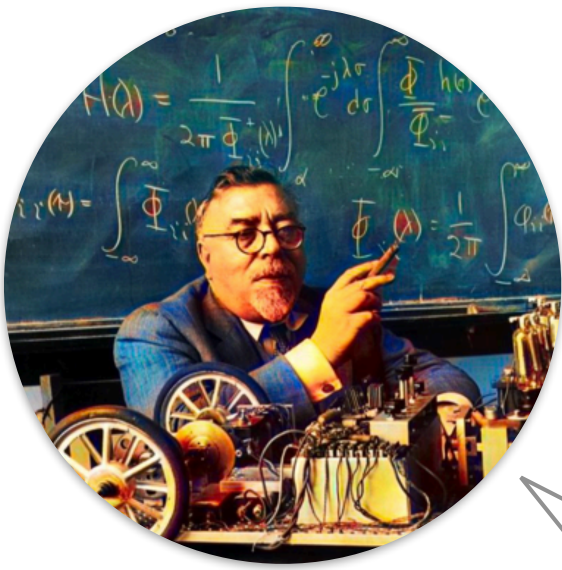

Norbert Wiener

1894 - 1964

*"…undamped **feedback** is strikingly similar to…a **cerebellar** patient. If he is asked to carry a glass of water from a table to his mouth, the hand carrying the glass will execute a series of **oscillatory** motions of increasing amplitude."*

Rosenblueth, Wiener, and Bigelow, 1943

youtube.com/ZkM-1MLn0_A

Norbert Wiener

1894 - 1964

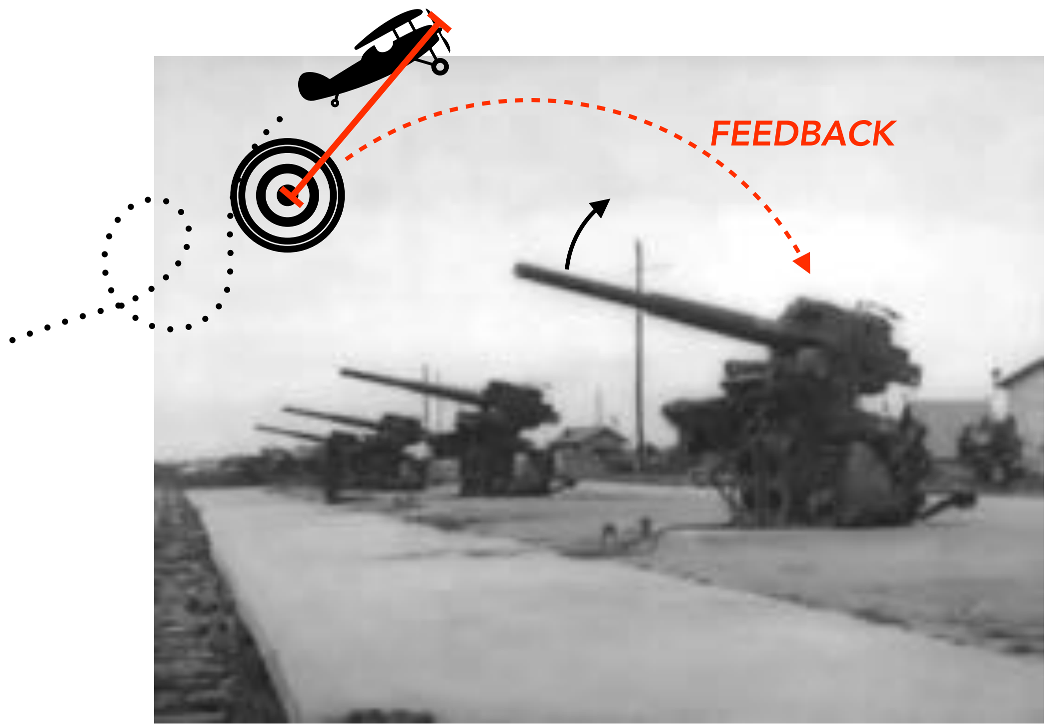

anti-aircraft guns

control as…

probability *maximization*

error *minimization*

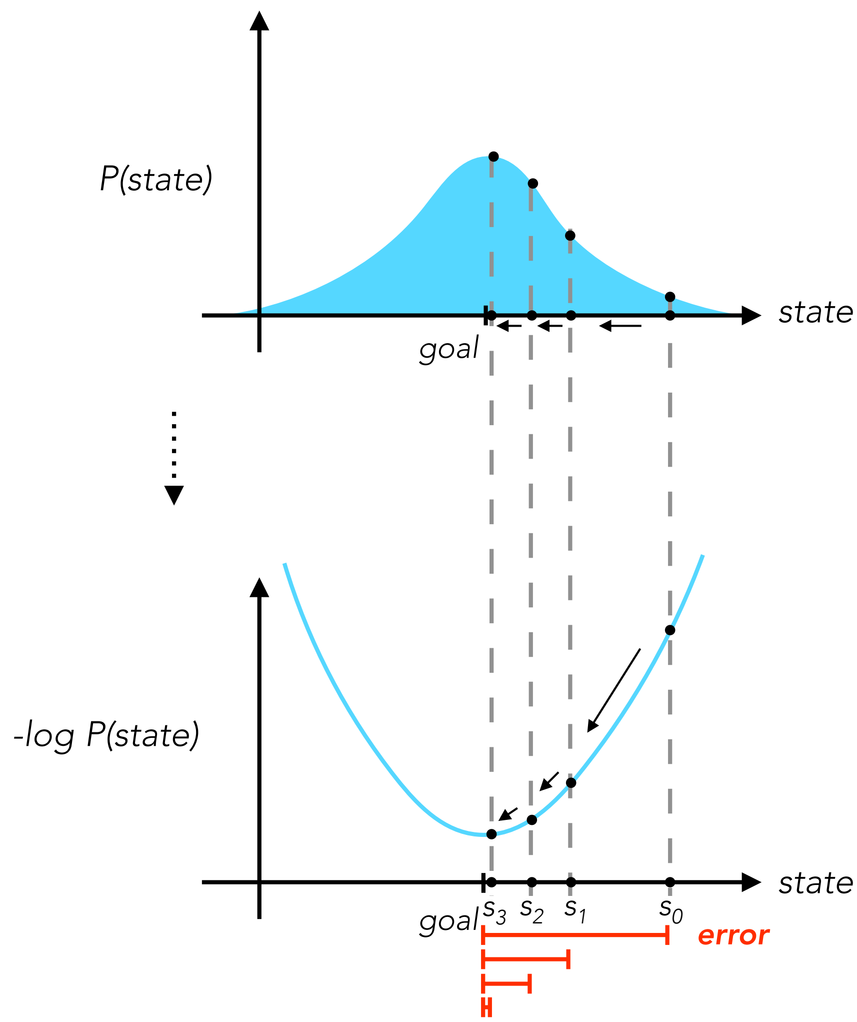

control as…

***probability** maximization*

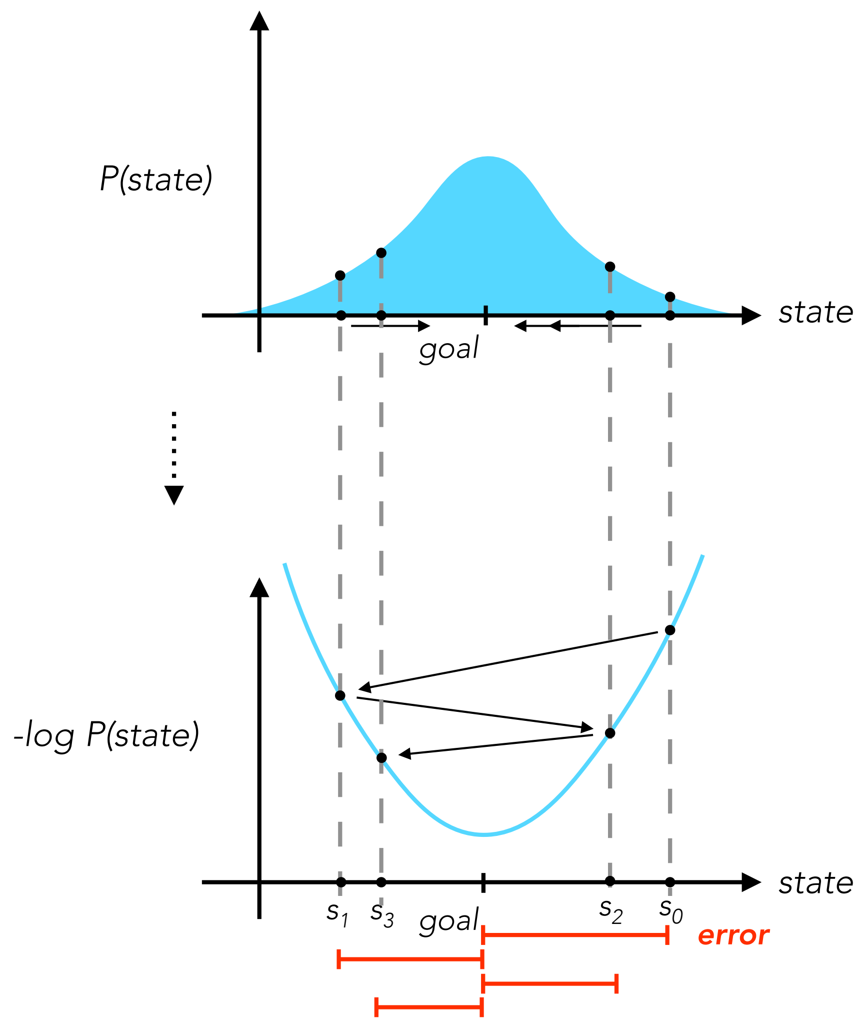

***error** minimization*

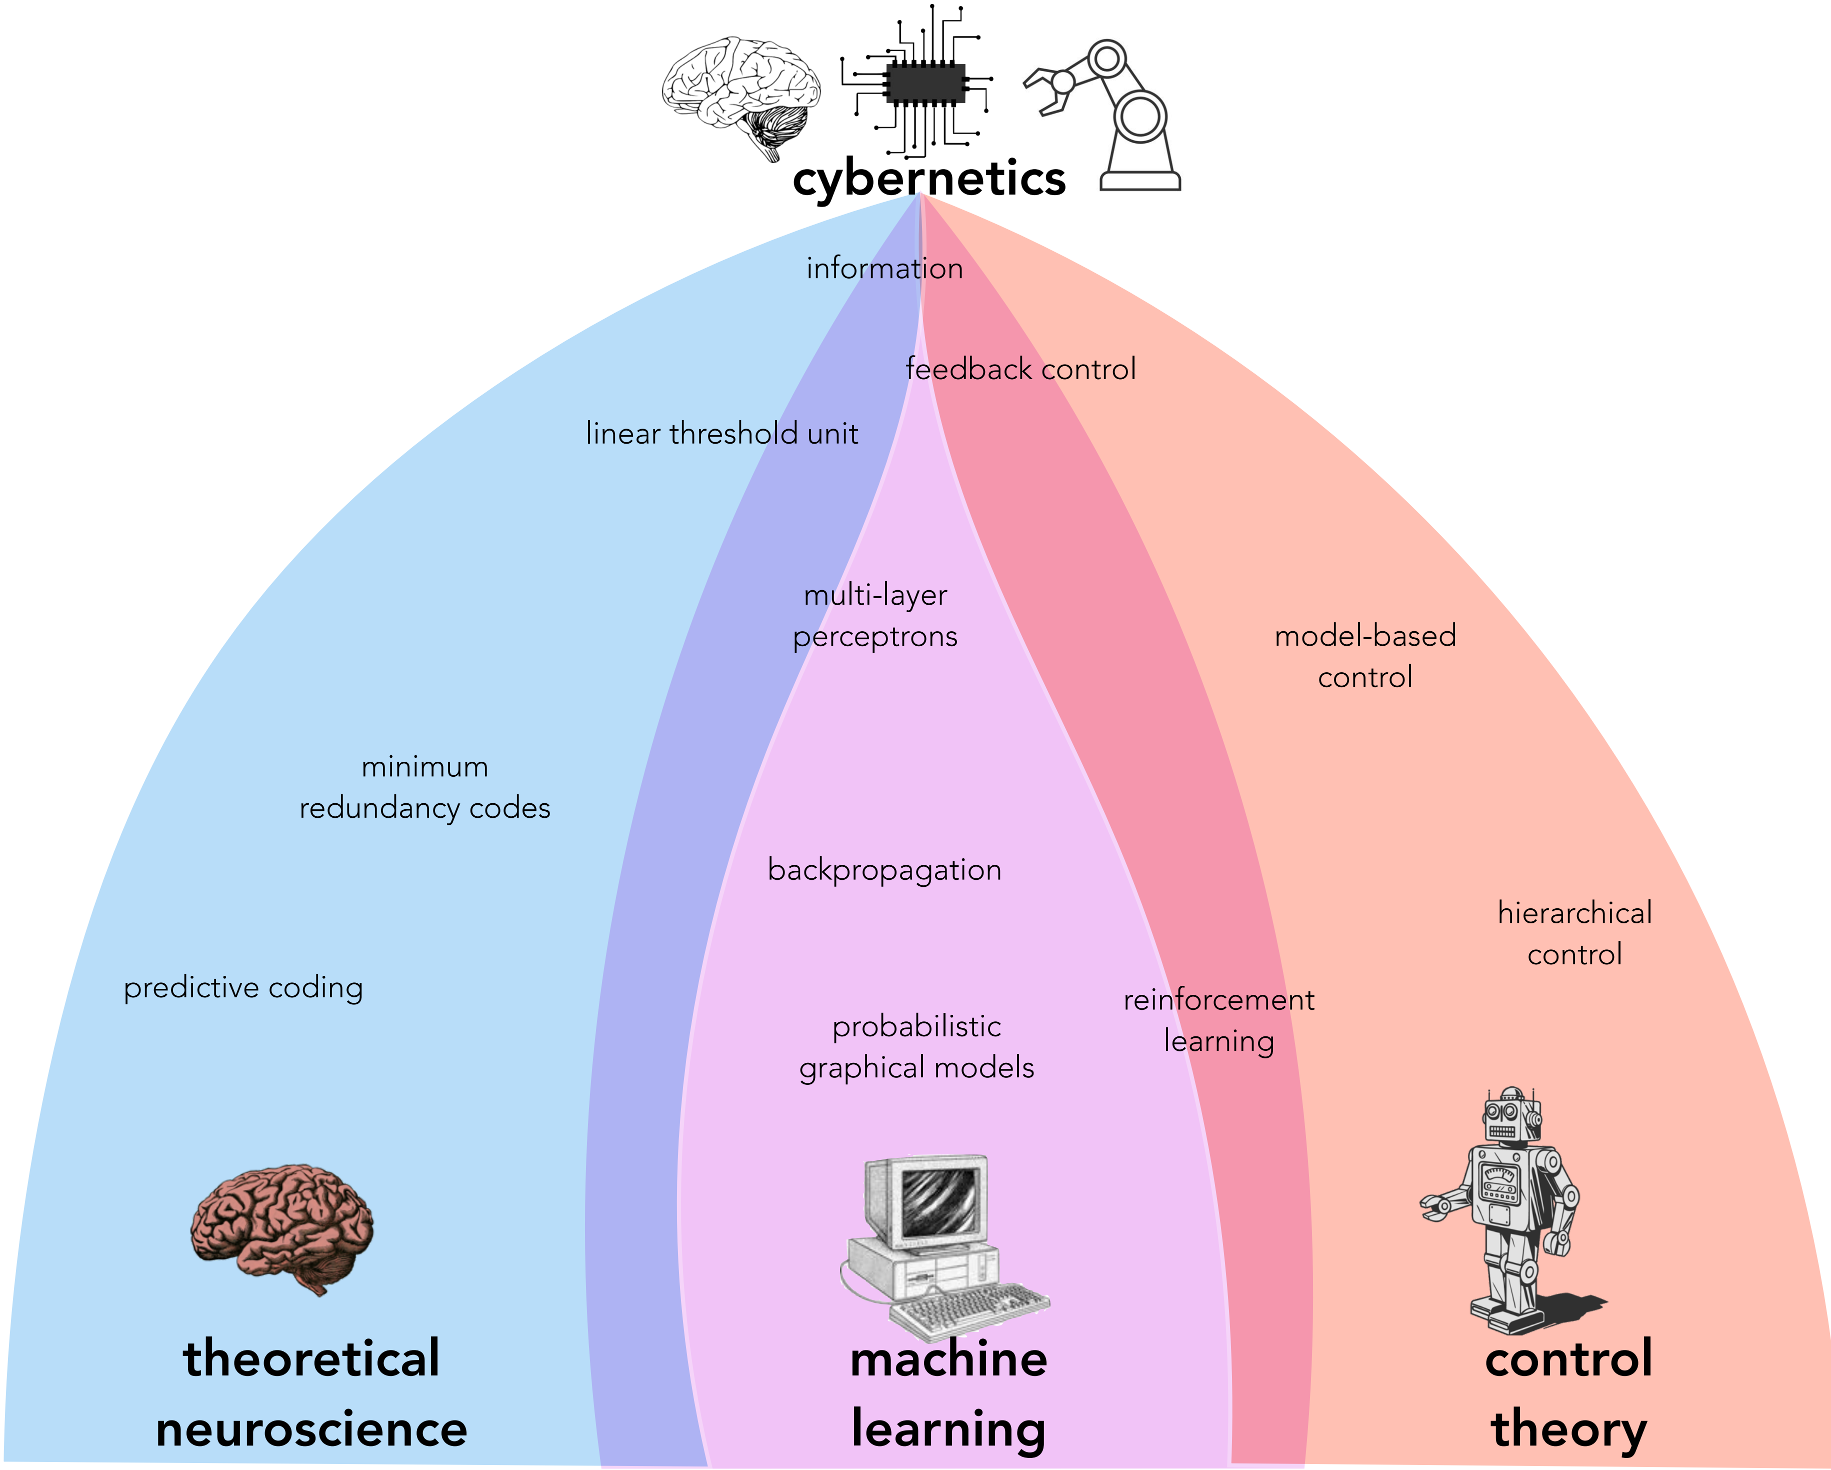
cybernetics
information
feedback control
linear threshold unit
multi-layer
perceptrons
model-based
control
minimum
redundancy codes
backpropagation
hierarchical
control
predictive coding
reinforcement
learning
probabilistic
graphical models
theoretical
neuroscience
machine
learning
control
theory

feedback

reducing an error in the moment

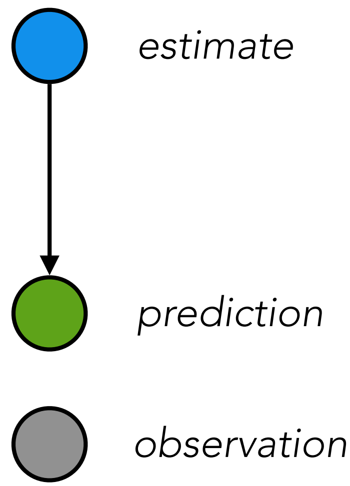

feedback

reducing an error in the moment

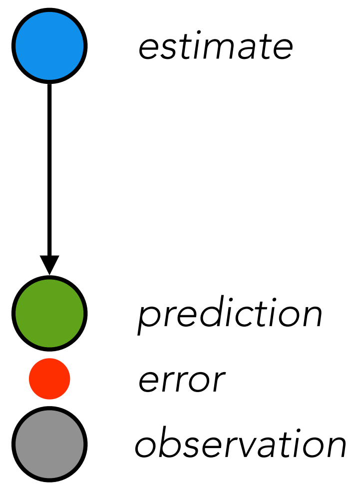

feedback

reducing an error in the moment

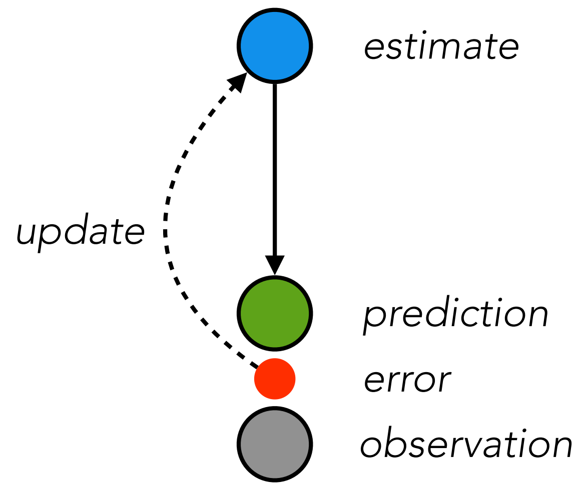

feedback

reducing an error in the moment

feedforward

preemptively reducing an error

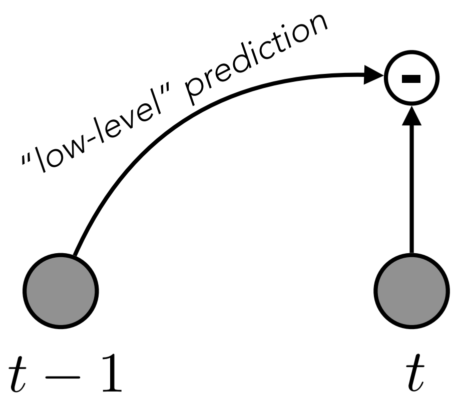

feedforward

preemptively reducing an error

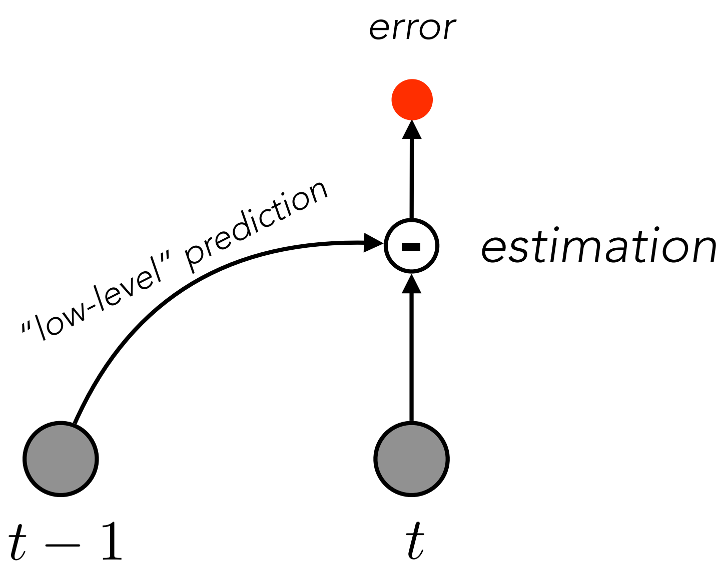

feedforward

preemptively reducing an error

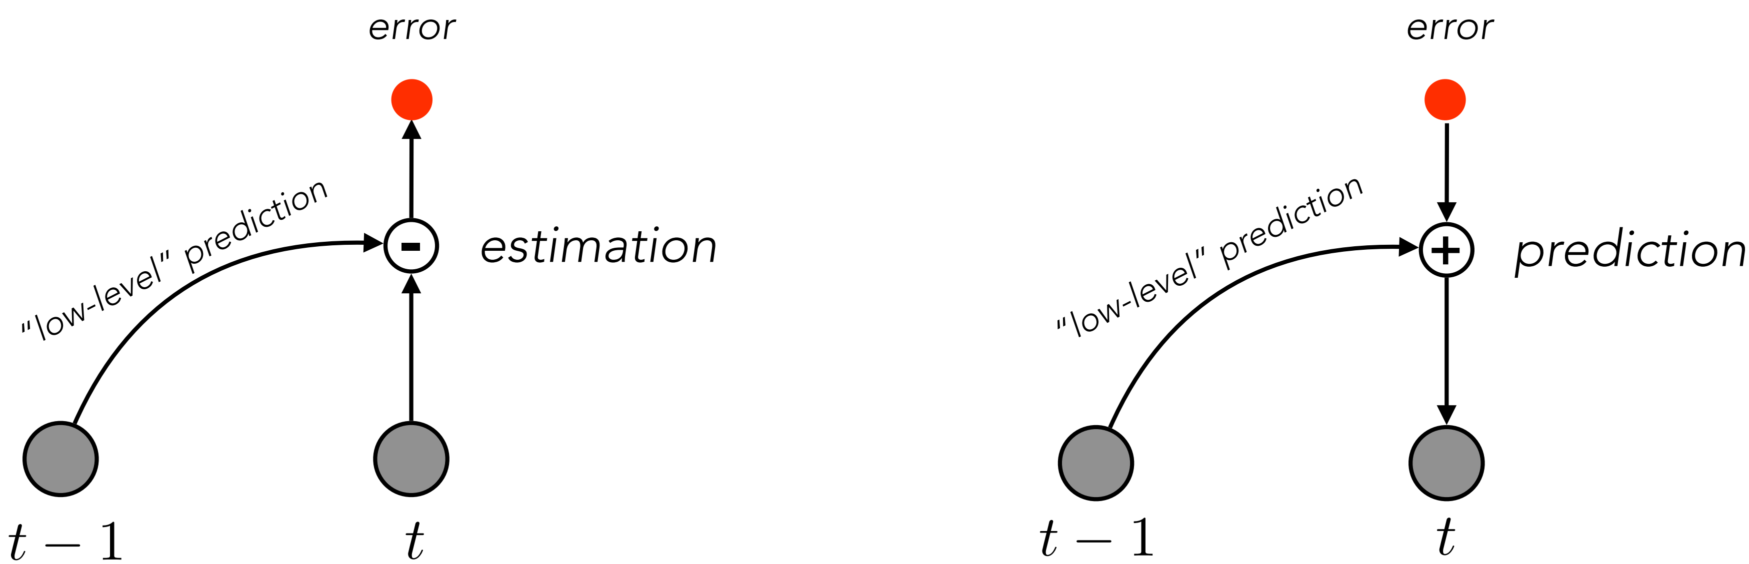

feedforward

preemptively reducing an error

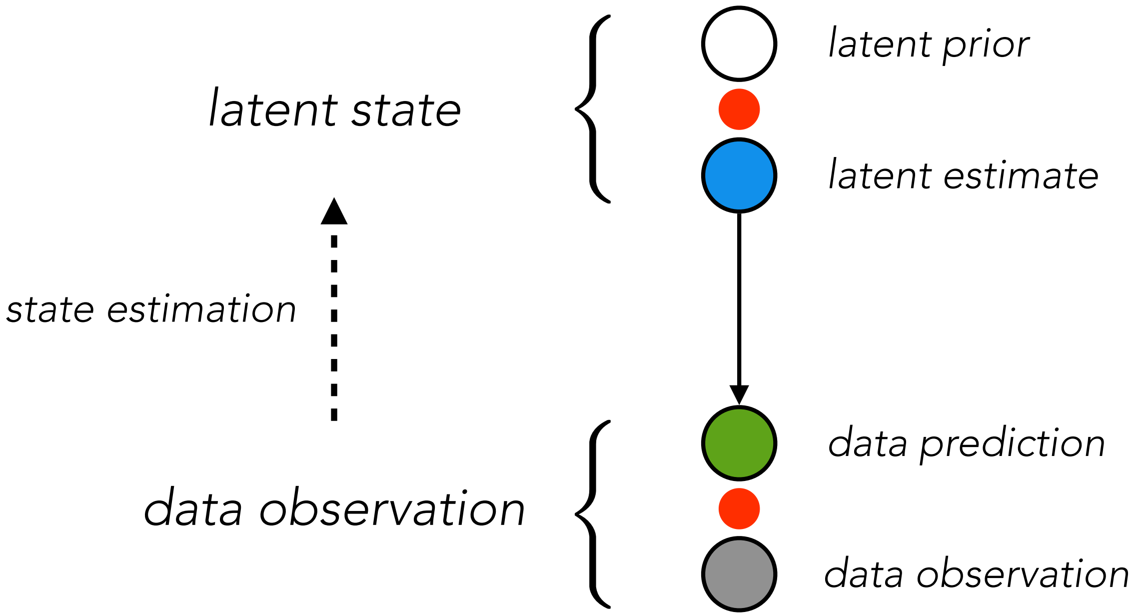

perception

state estimation

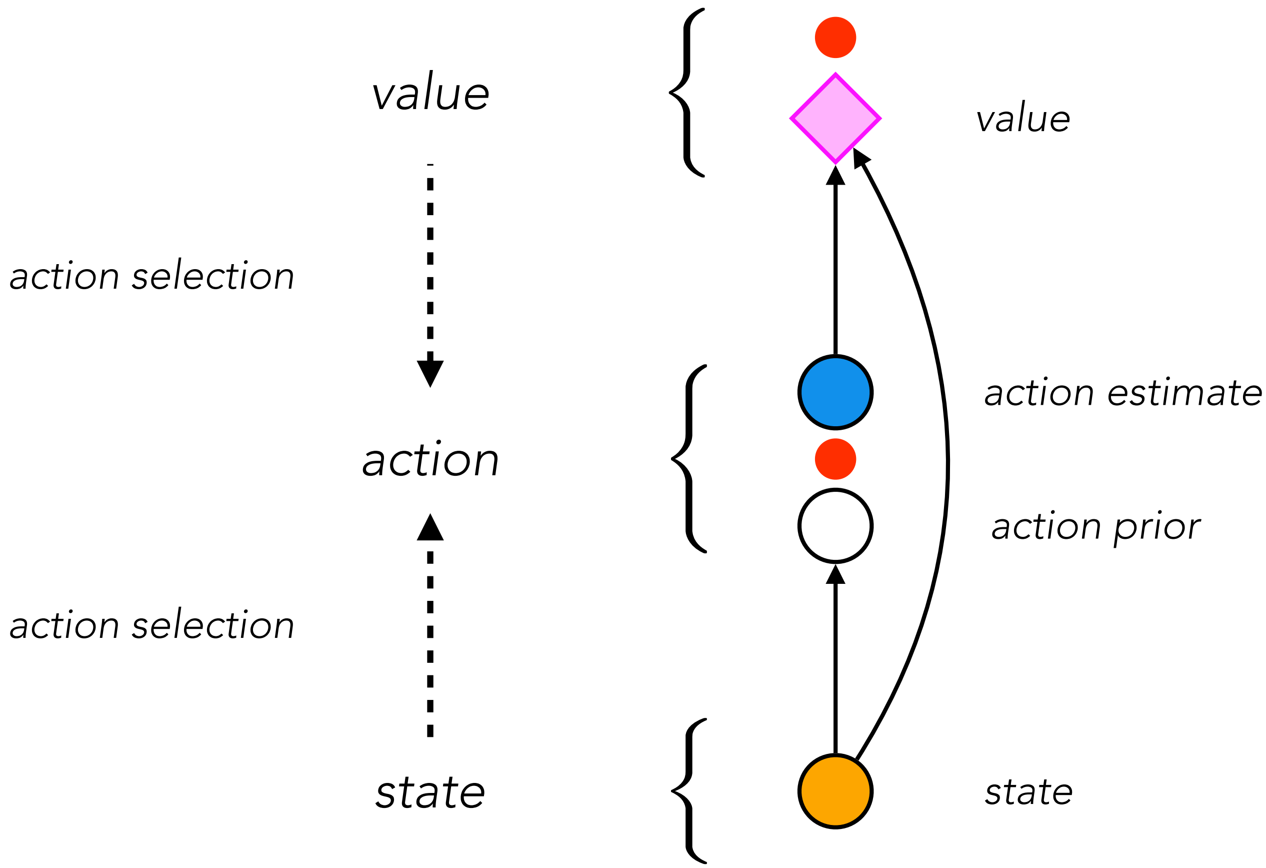

control

action selection

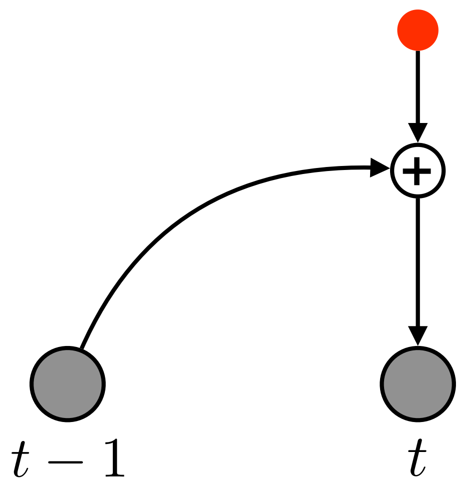

feedforward

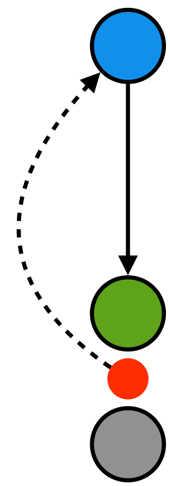

feedback

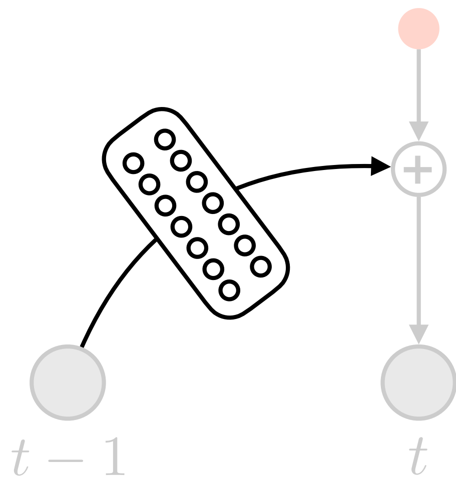

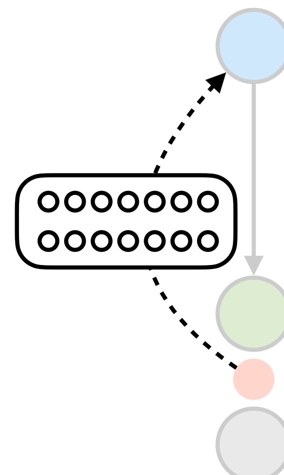

feedforward ***feedback***

my research:

we're going to ***learn*** feedforward & feedback functions

THESIS OVERVIEW

feedforward | ***feedback***

perception

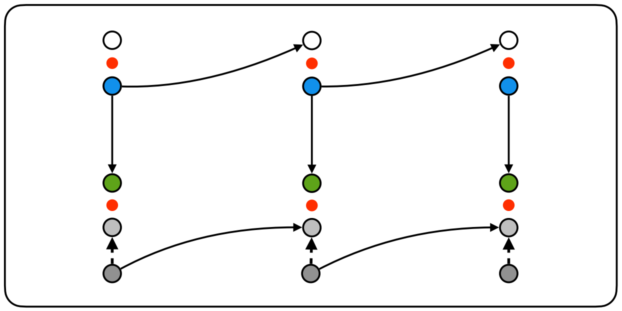

Marino, Chen, He, Mandt 2020
Yang, Yang, Marino, Mandt 2021

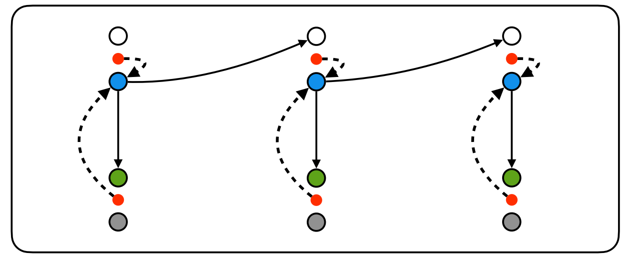

Marino, Yue, Mandt 2018
Marino, Cvitkovic, Yue, 2018

control

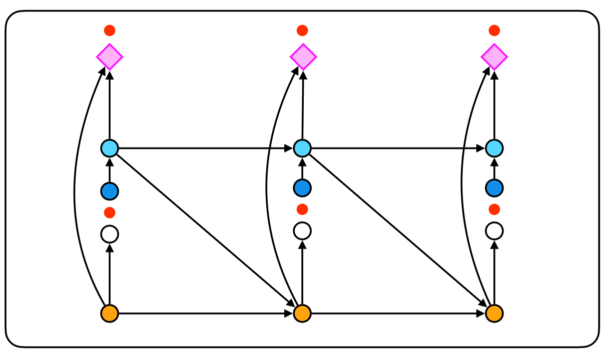

Guerra & Marino, 2020

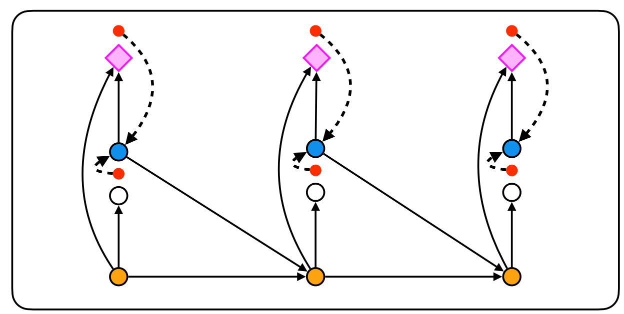

Marino, Piché, Ialongo, Yue 2020

THESIS OVERVIEW

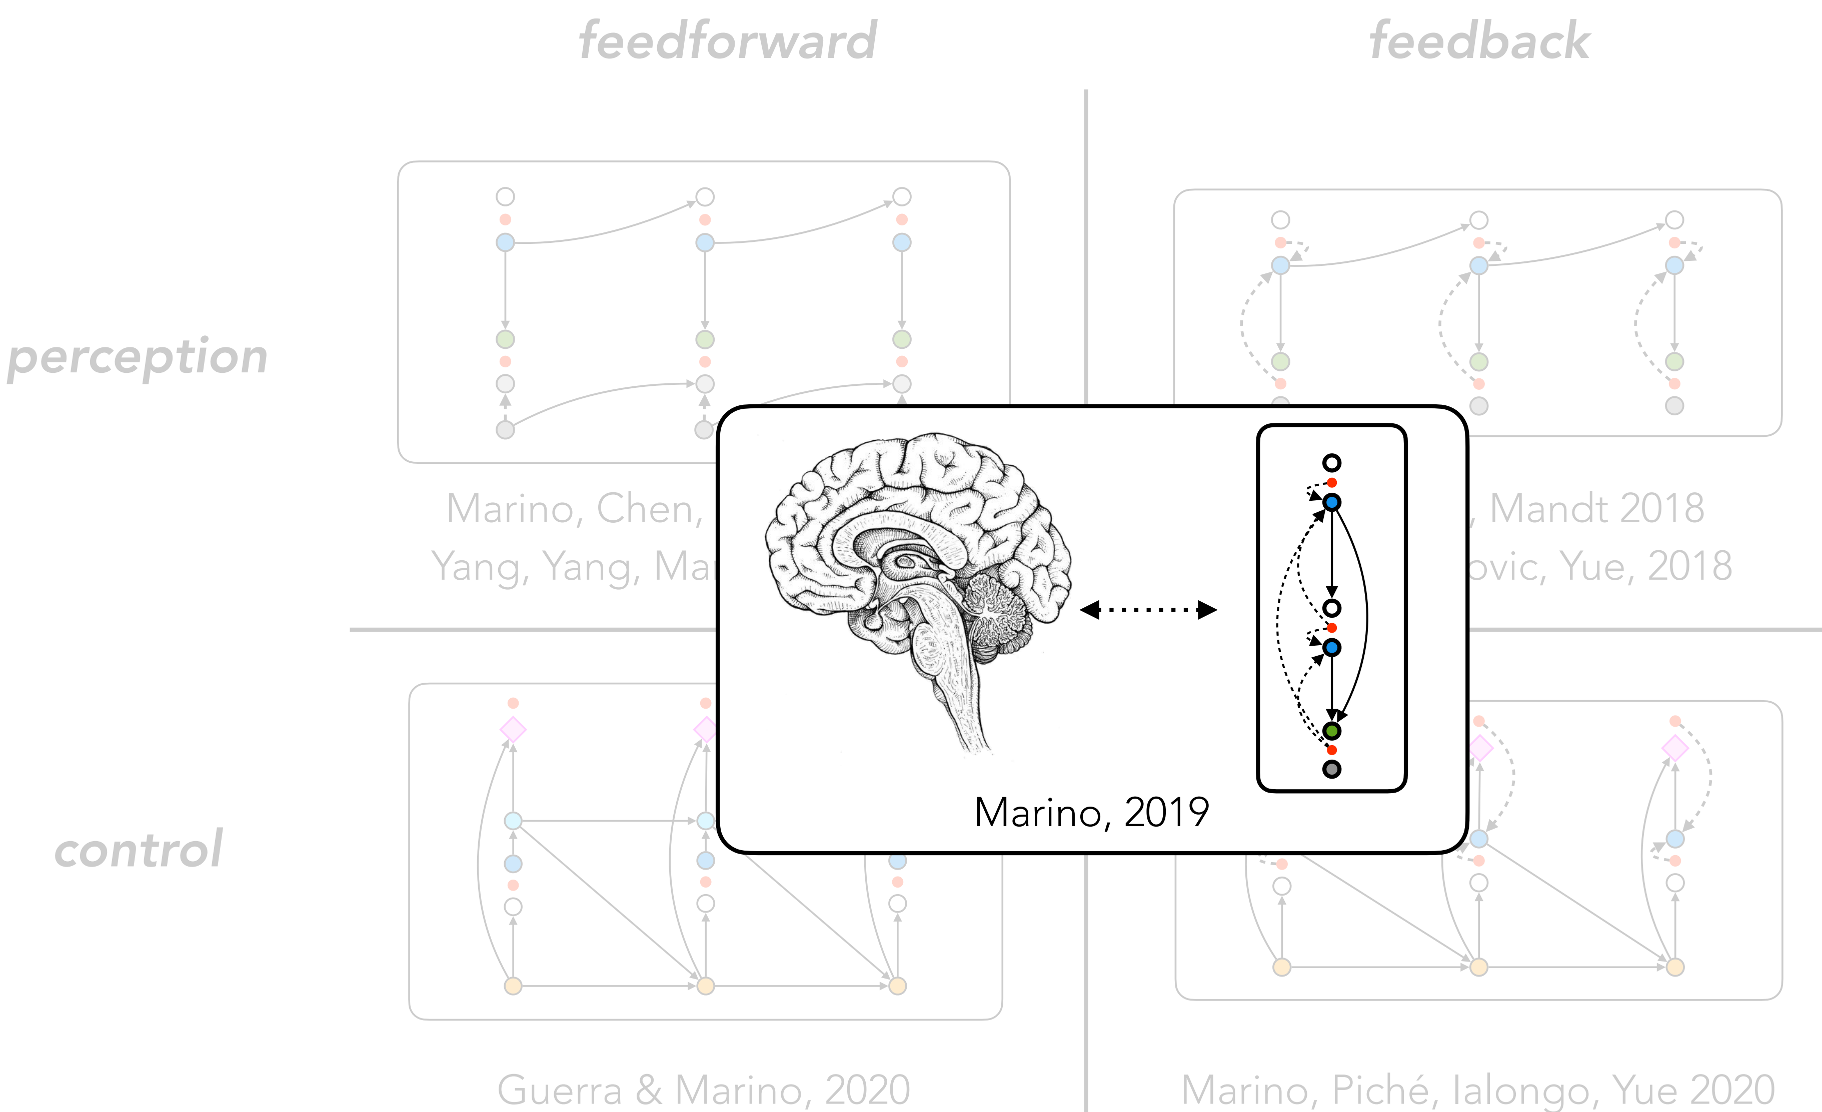

feedforward

PERCEPTION

SPATIOTEMPORAL PREDICTIVE CODING

spatiotemporal predictions ***normalize*** *sensory inputs*

Predictive coding: a fresh view of inhibition in the retina

By M. V. Srinivasan[1,2]†, S. B. Laughlin[1] and A. Dubs[1] 1982

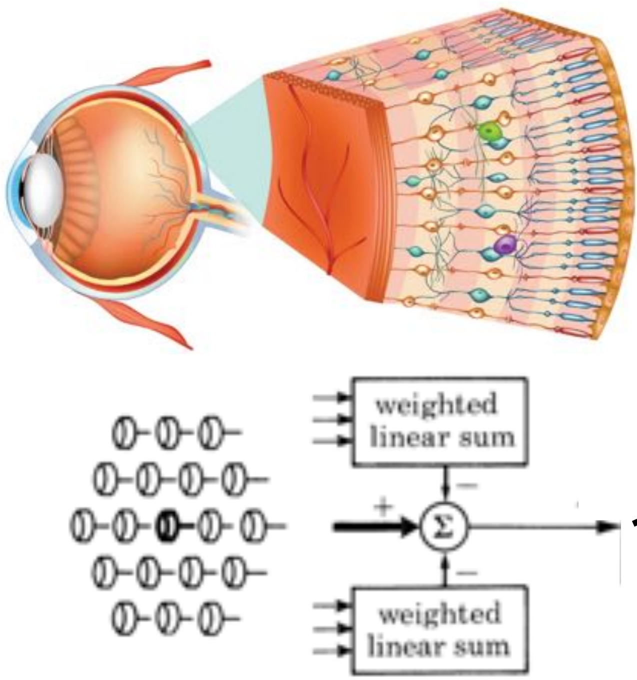

Temporal Decorrelation: A Theory of Lagged and Nonlagged Responses in the Lateral Geniculate Nucleus

Dawei W. Dong and Joseph J. Atick 1995

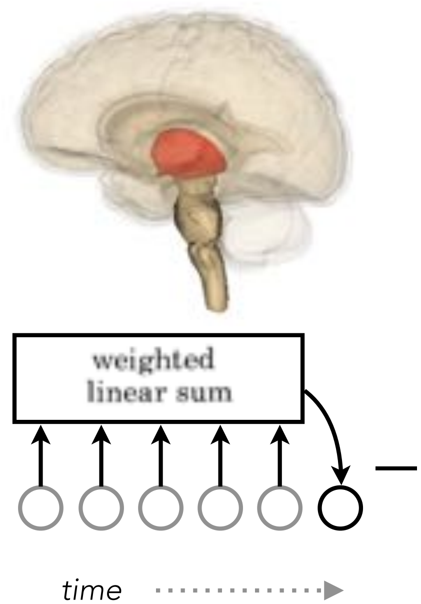

SPATIAL NORMALIZATION

TEMPORAL NORMALIZATION

FEEDFORWARD PERCEPTION

sequential latent variable model (SLVM)

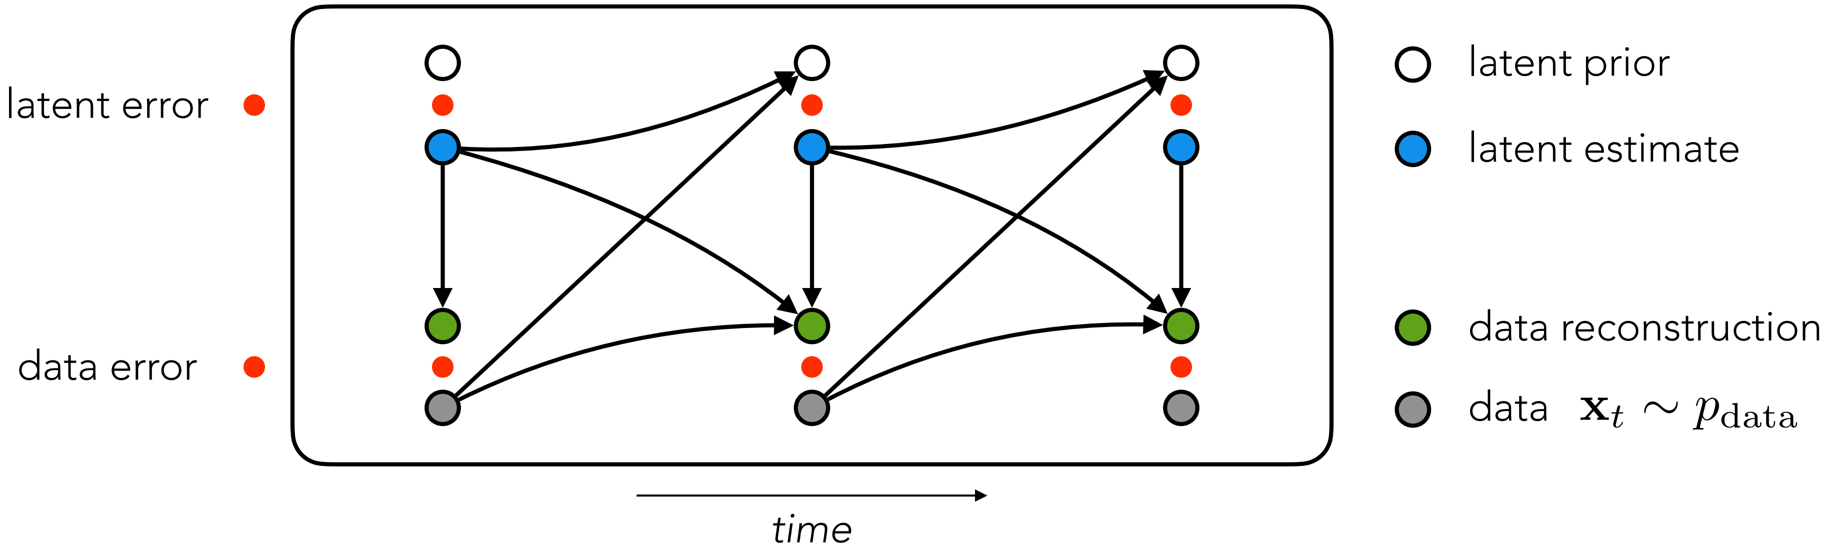

to improve this setup, we can…

FEEDFORWARD PERCEPTION

sequential latent variable model (SLVM)

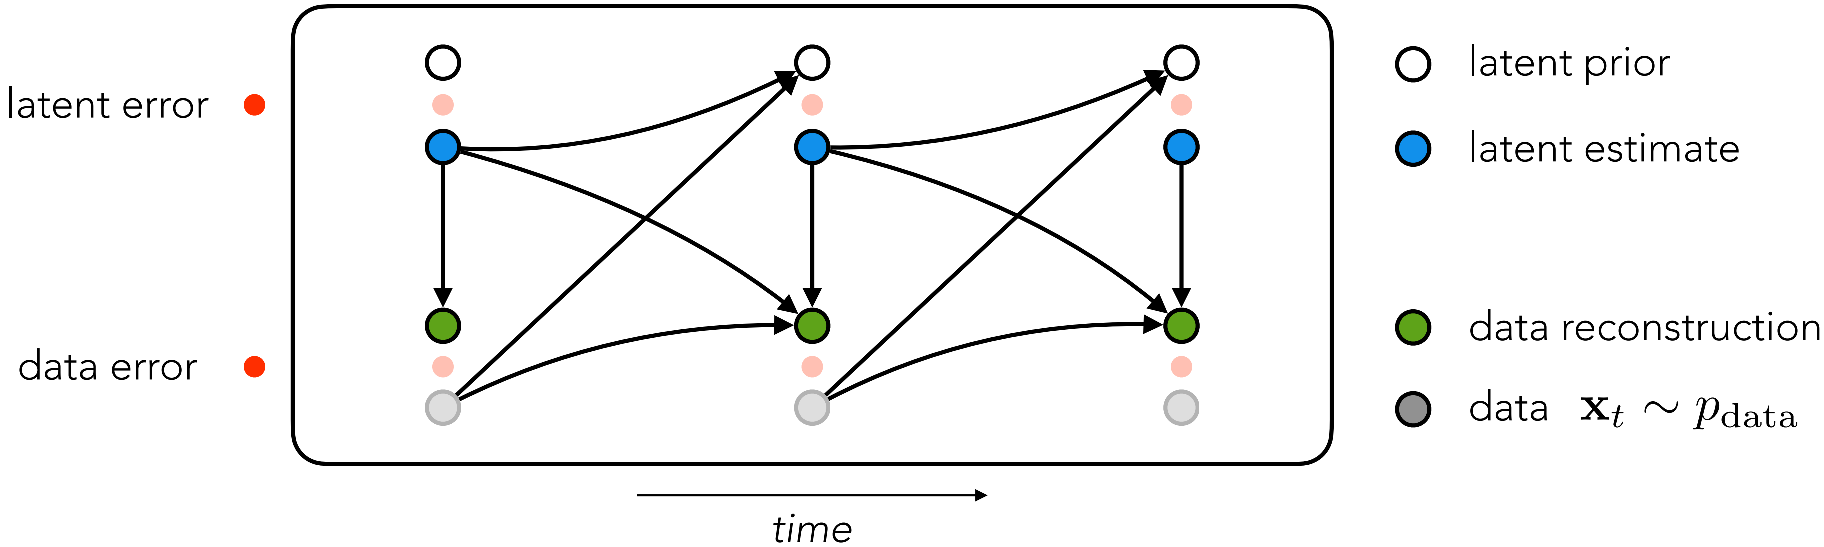

to improve this setup, we can…

improve the model

→ more dependencies, bigger functions

FEEDFORWARD PERCEPTION

sequential latent variable model (SLVM)

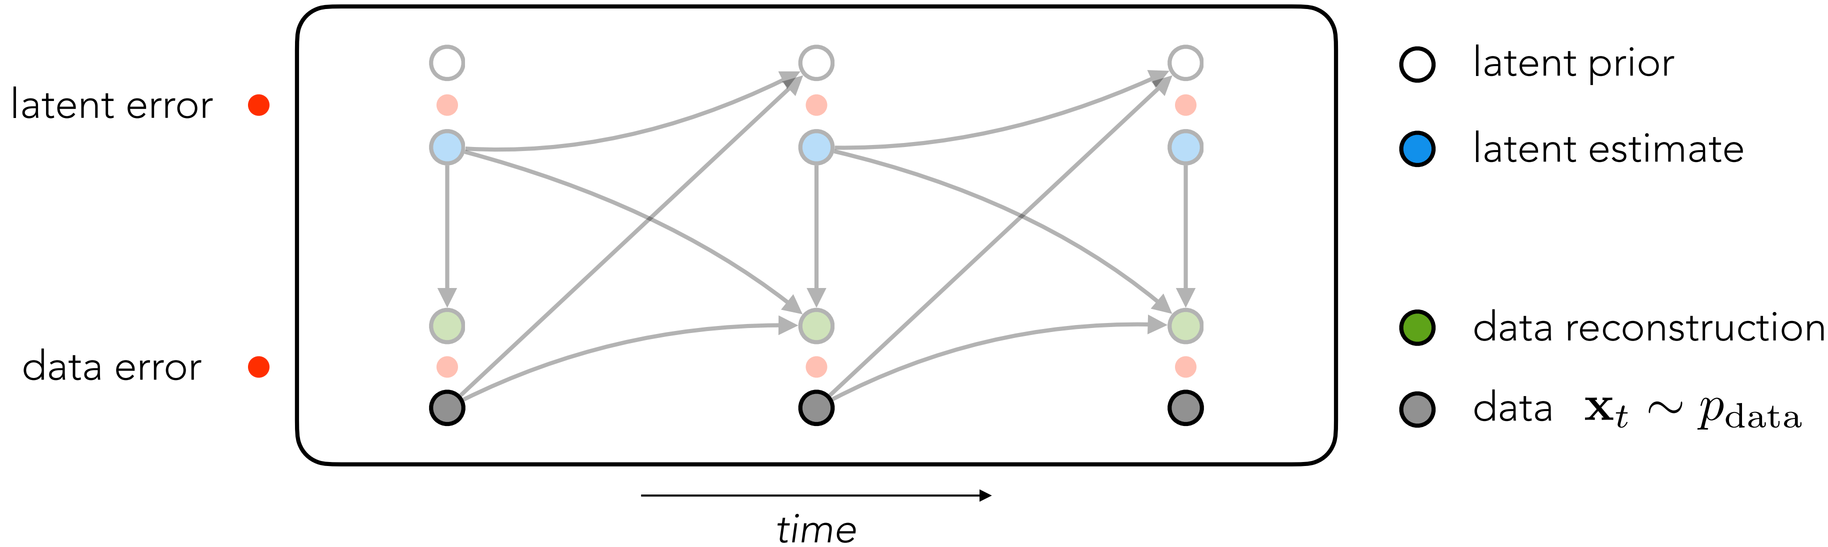

to improve this setup, we can…

improve the model

→ more dependencies, bigger functions

simplify the data (change the distribution)

→ e.g., model $\Delta \mathbf{x}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$

FEEDFORWARD PERCEPTION

temporal normalization

use an autoregressive model to remove a "low-level" prediction

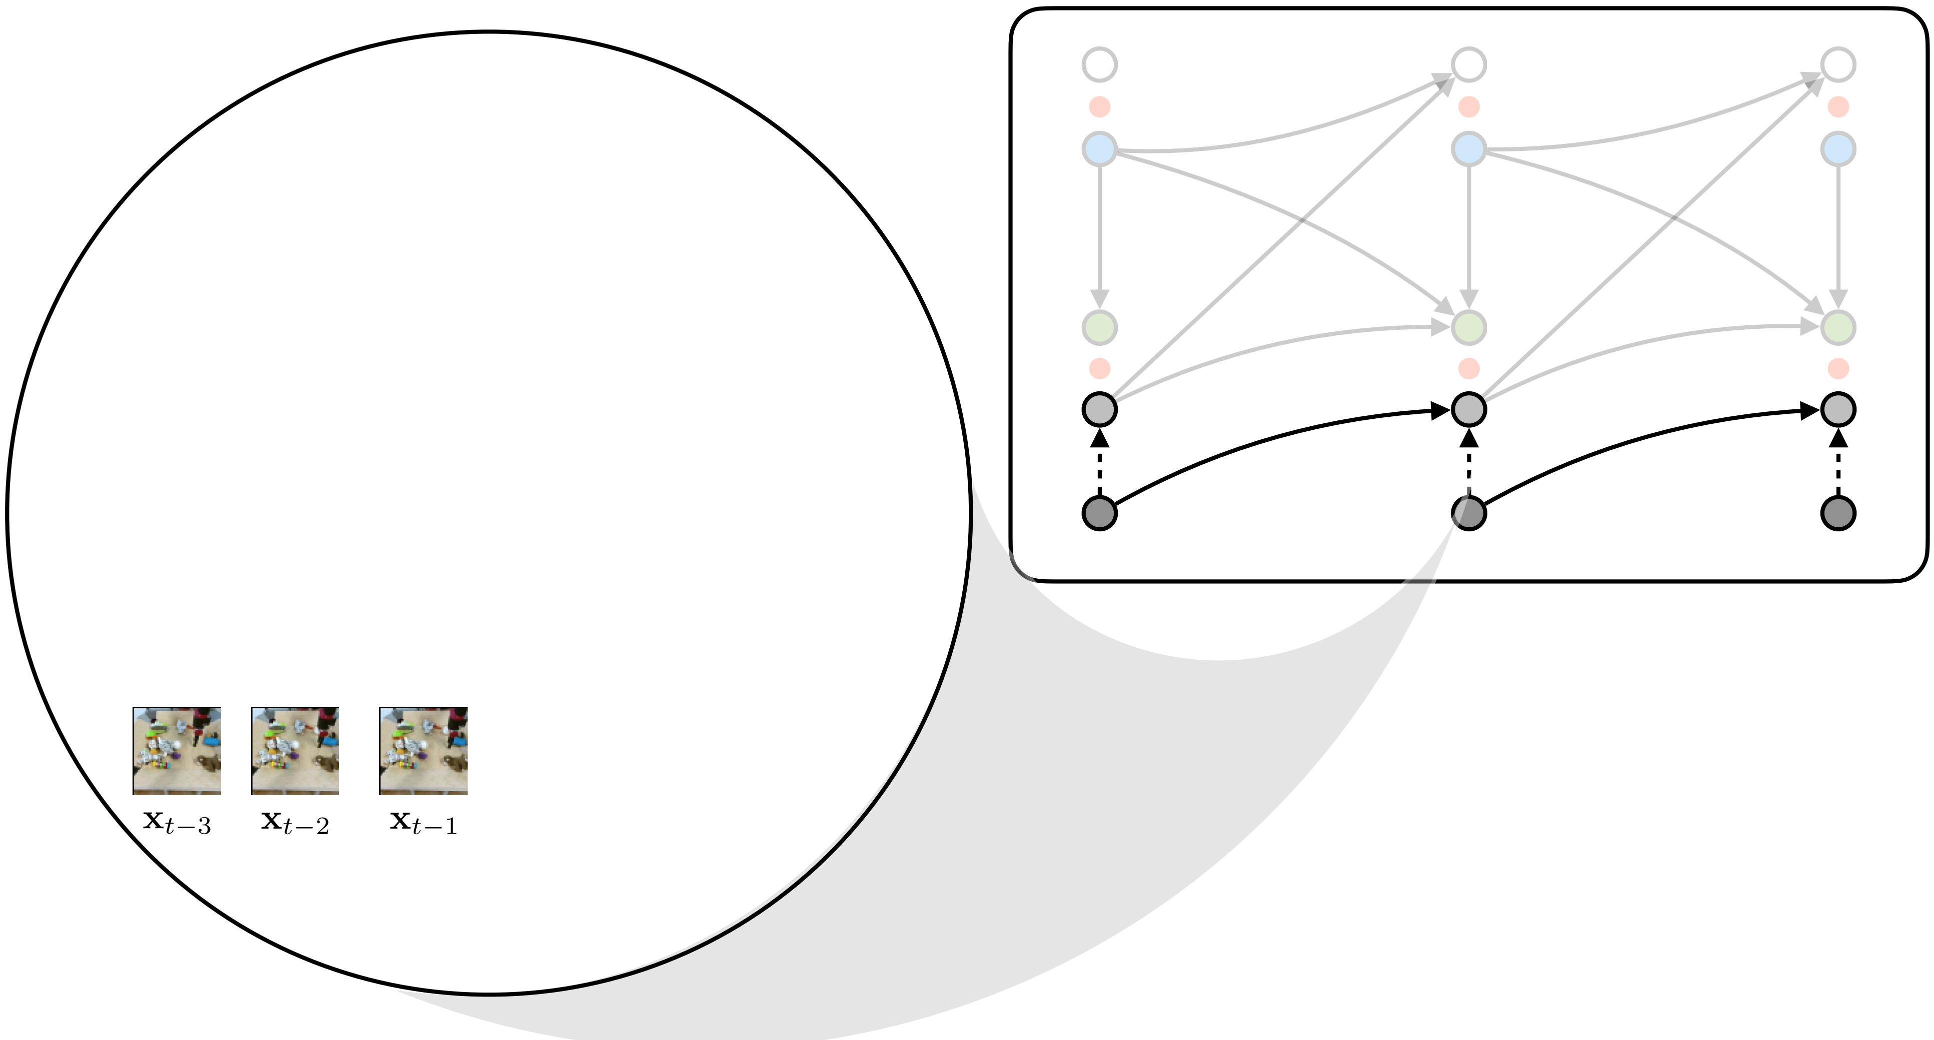

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

temporal normalization

use an autoregressive model to remove a “low-level” prediction

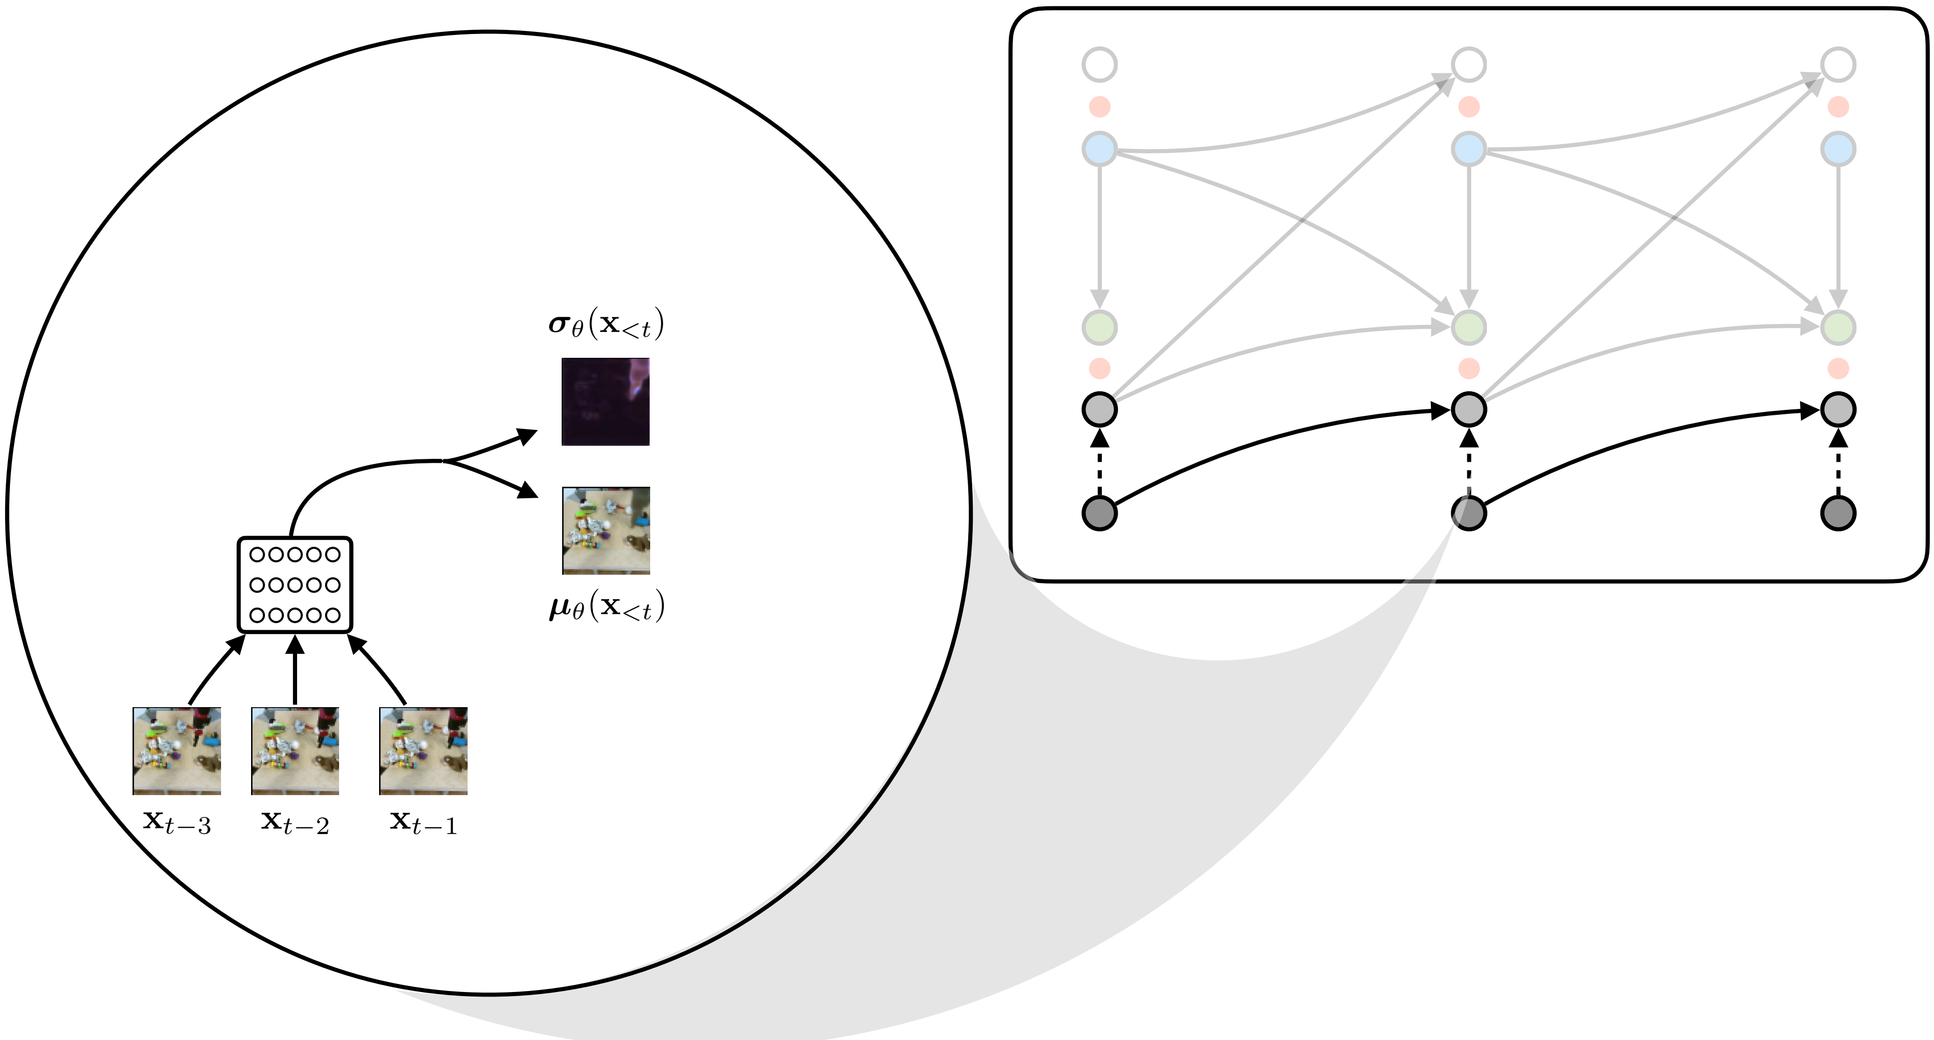

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

temporal normalization

use an autoregressive model to remove a "low-level" prediction

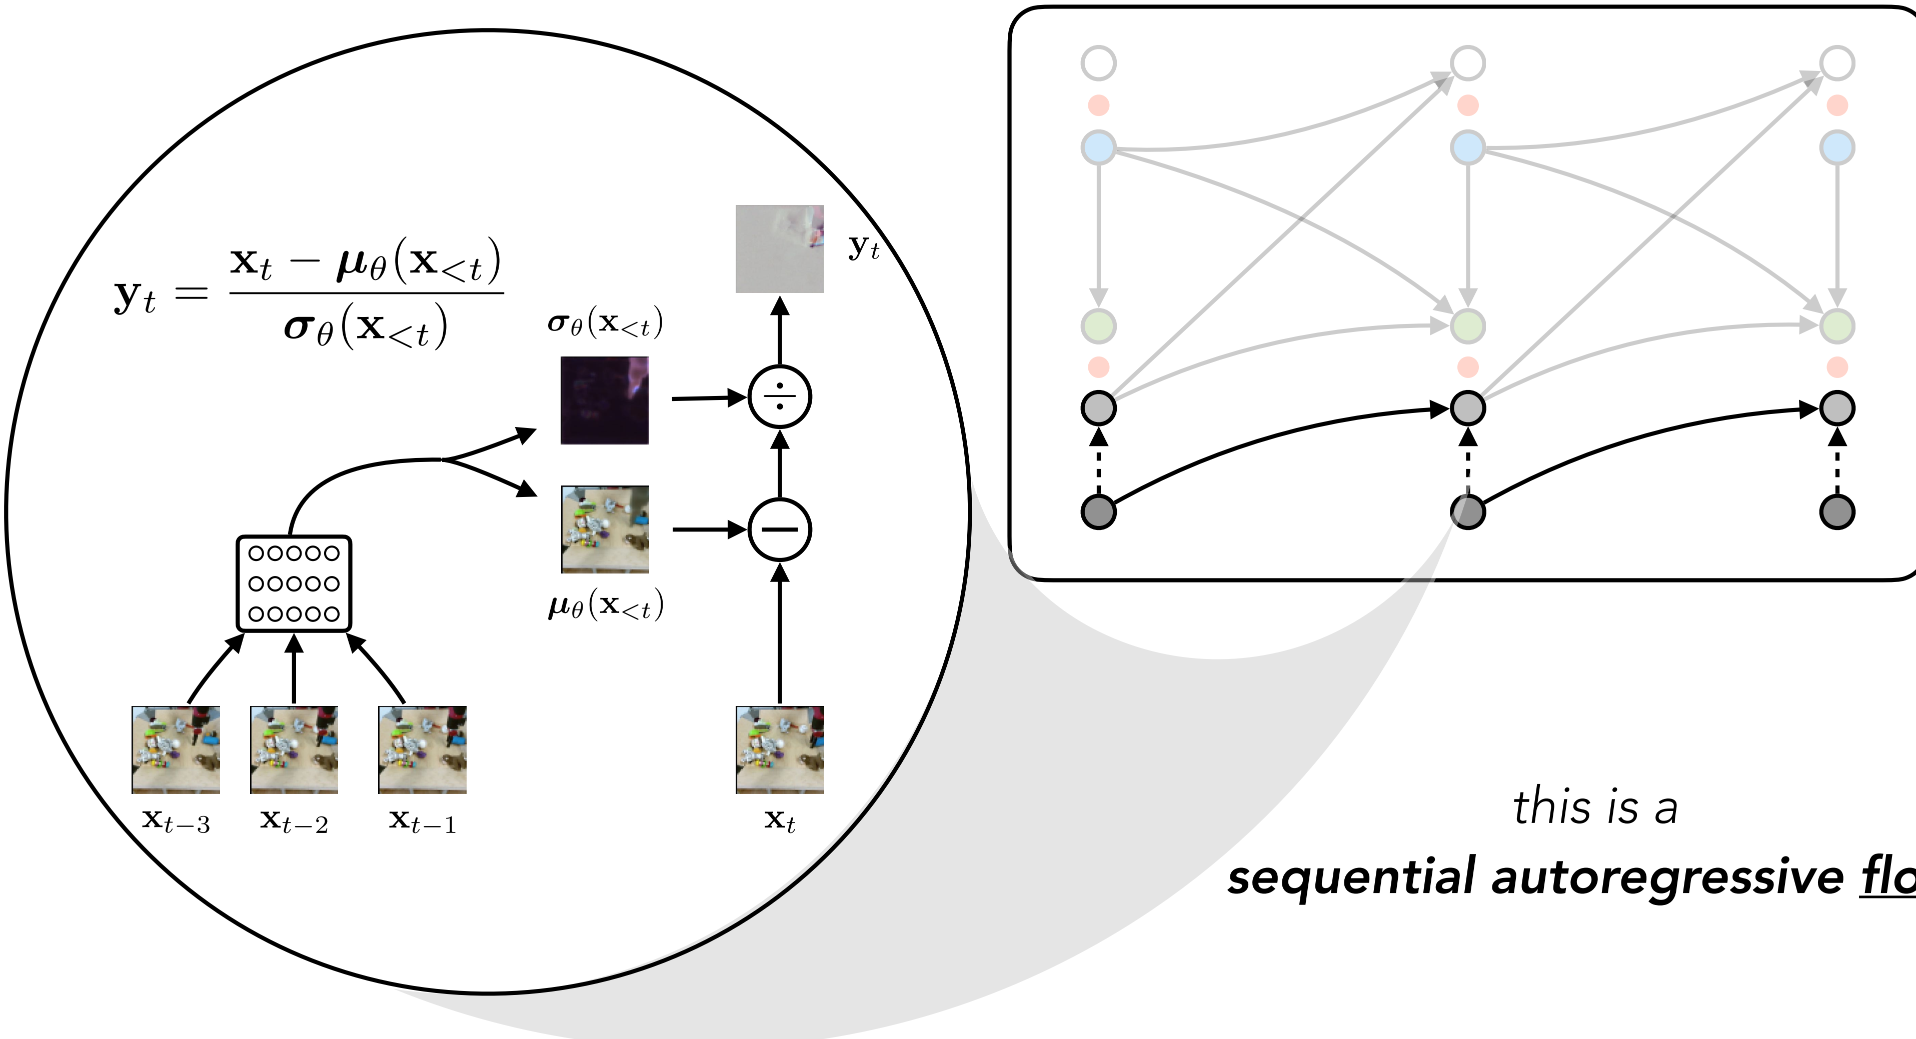

this is a

sequential autoregressive <u>flow</u>

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

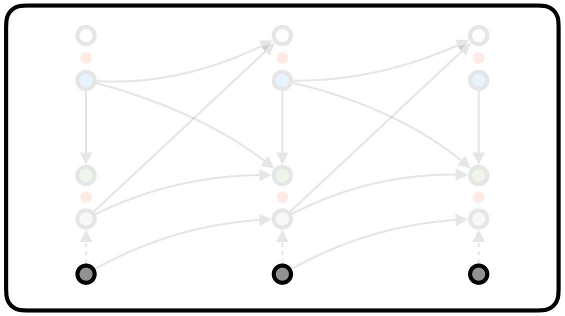

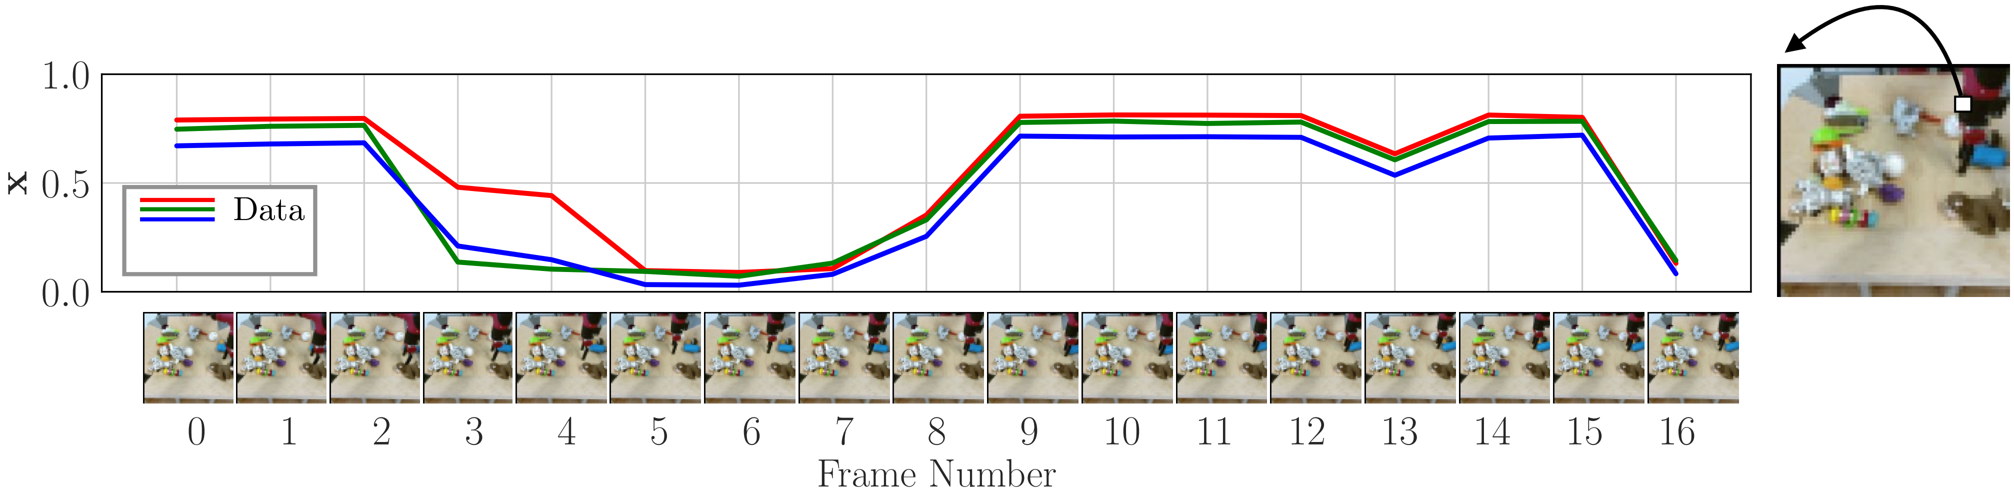

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

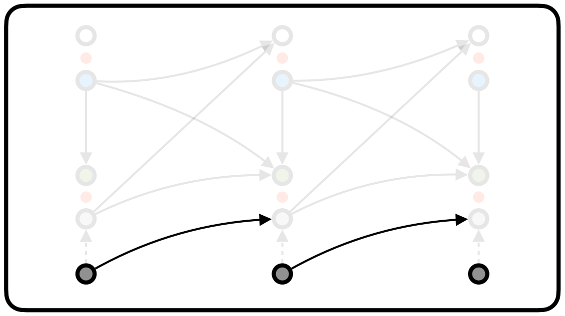

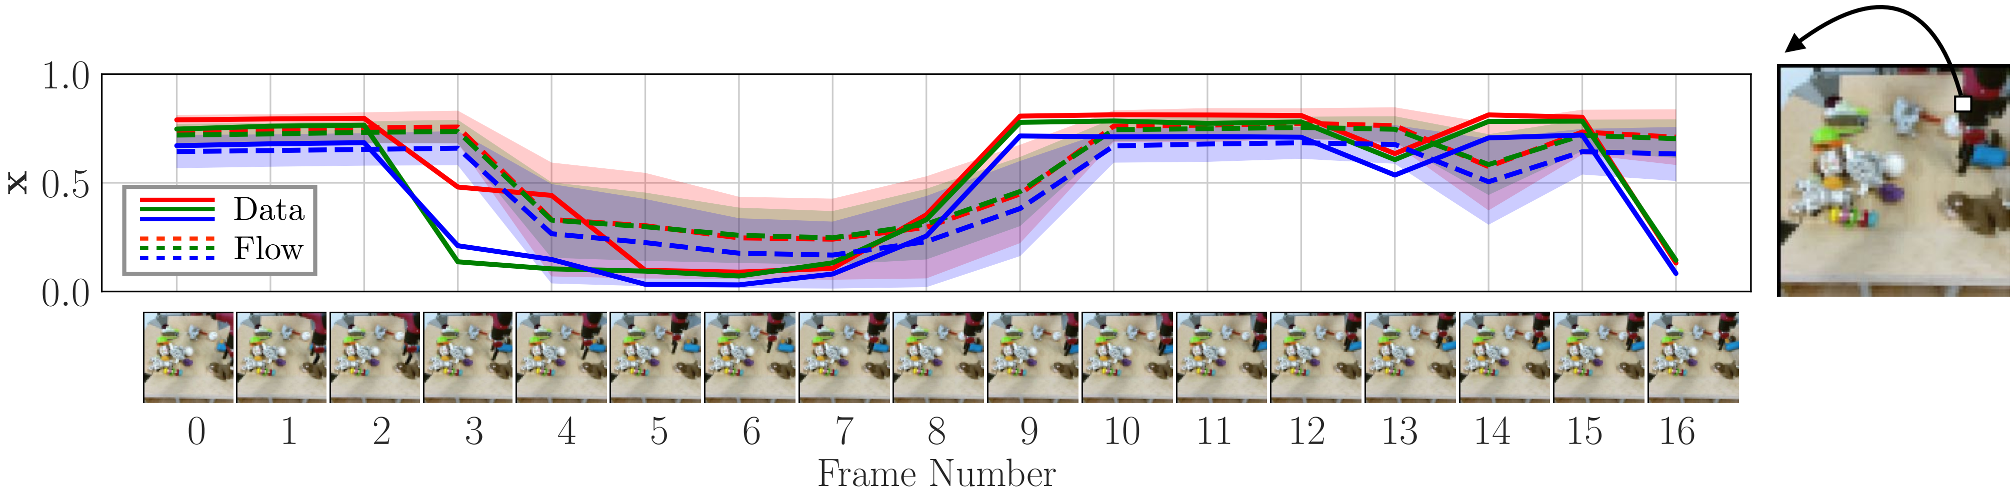

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

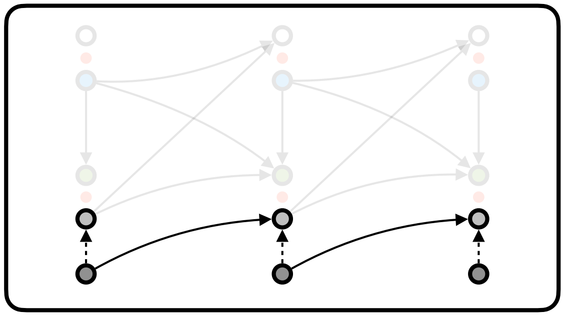

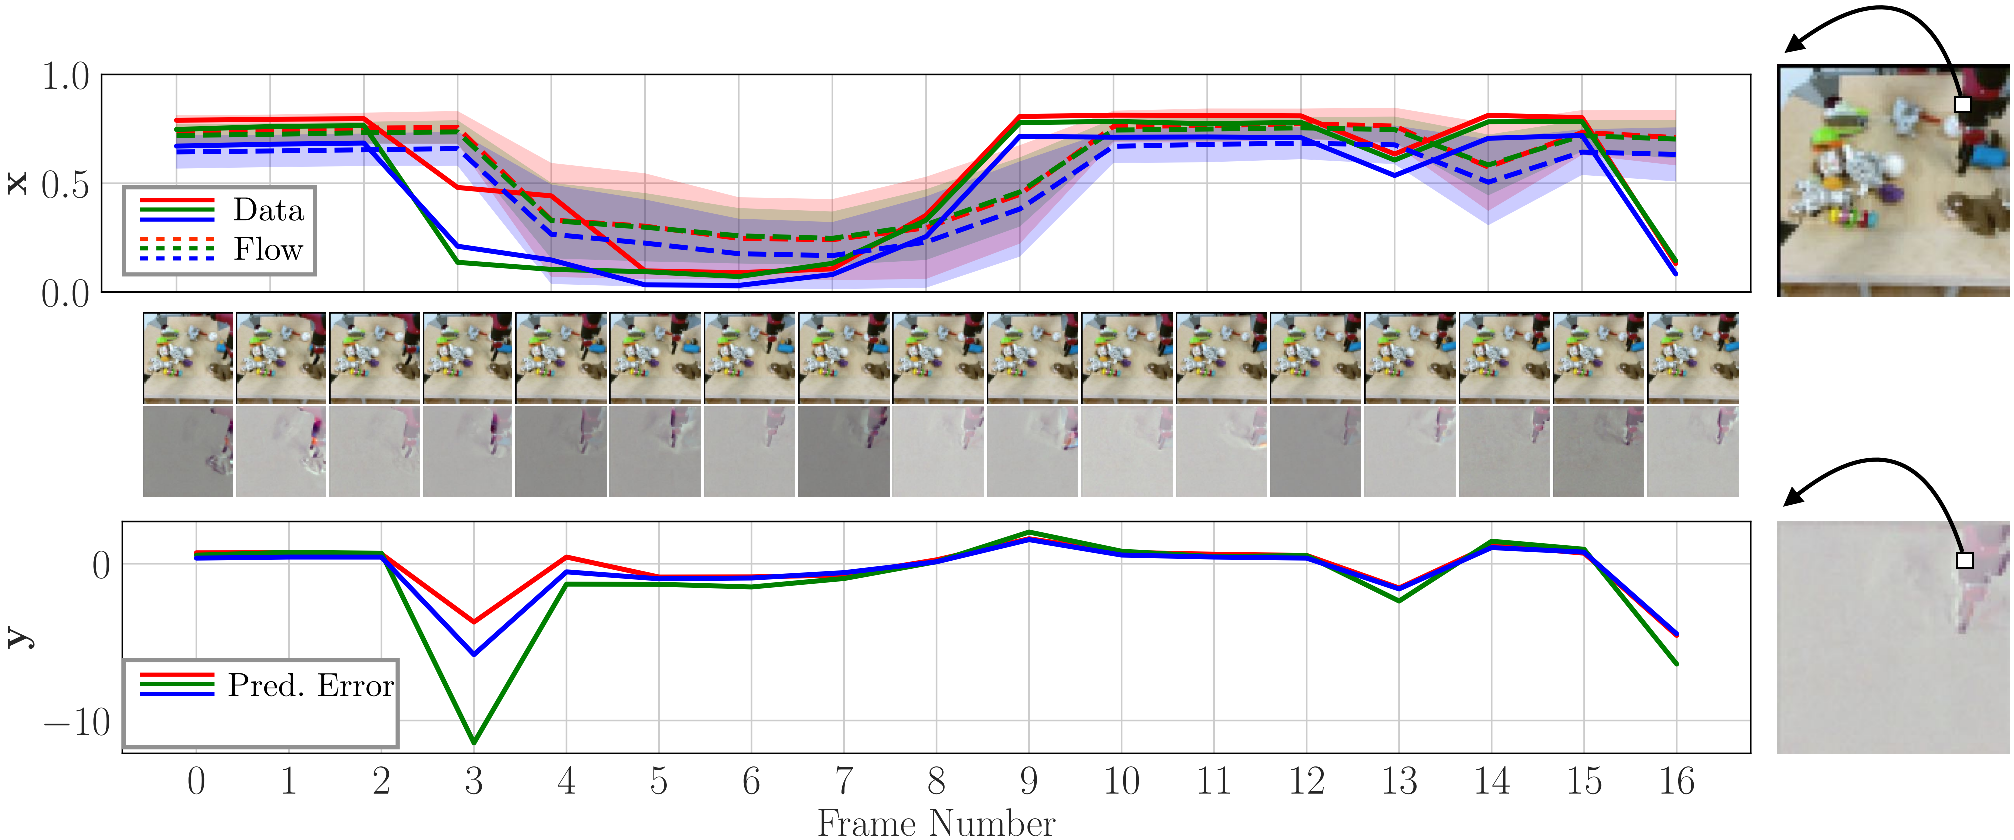

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

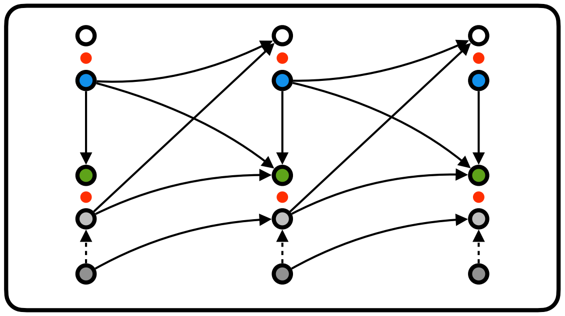

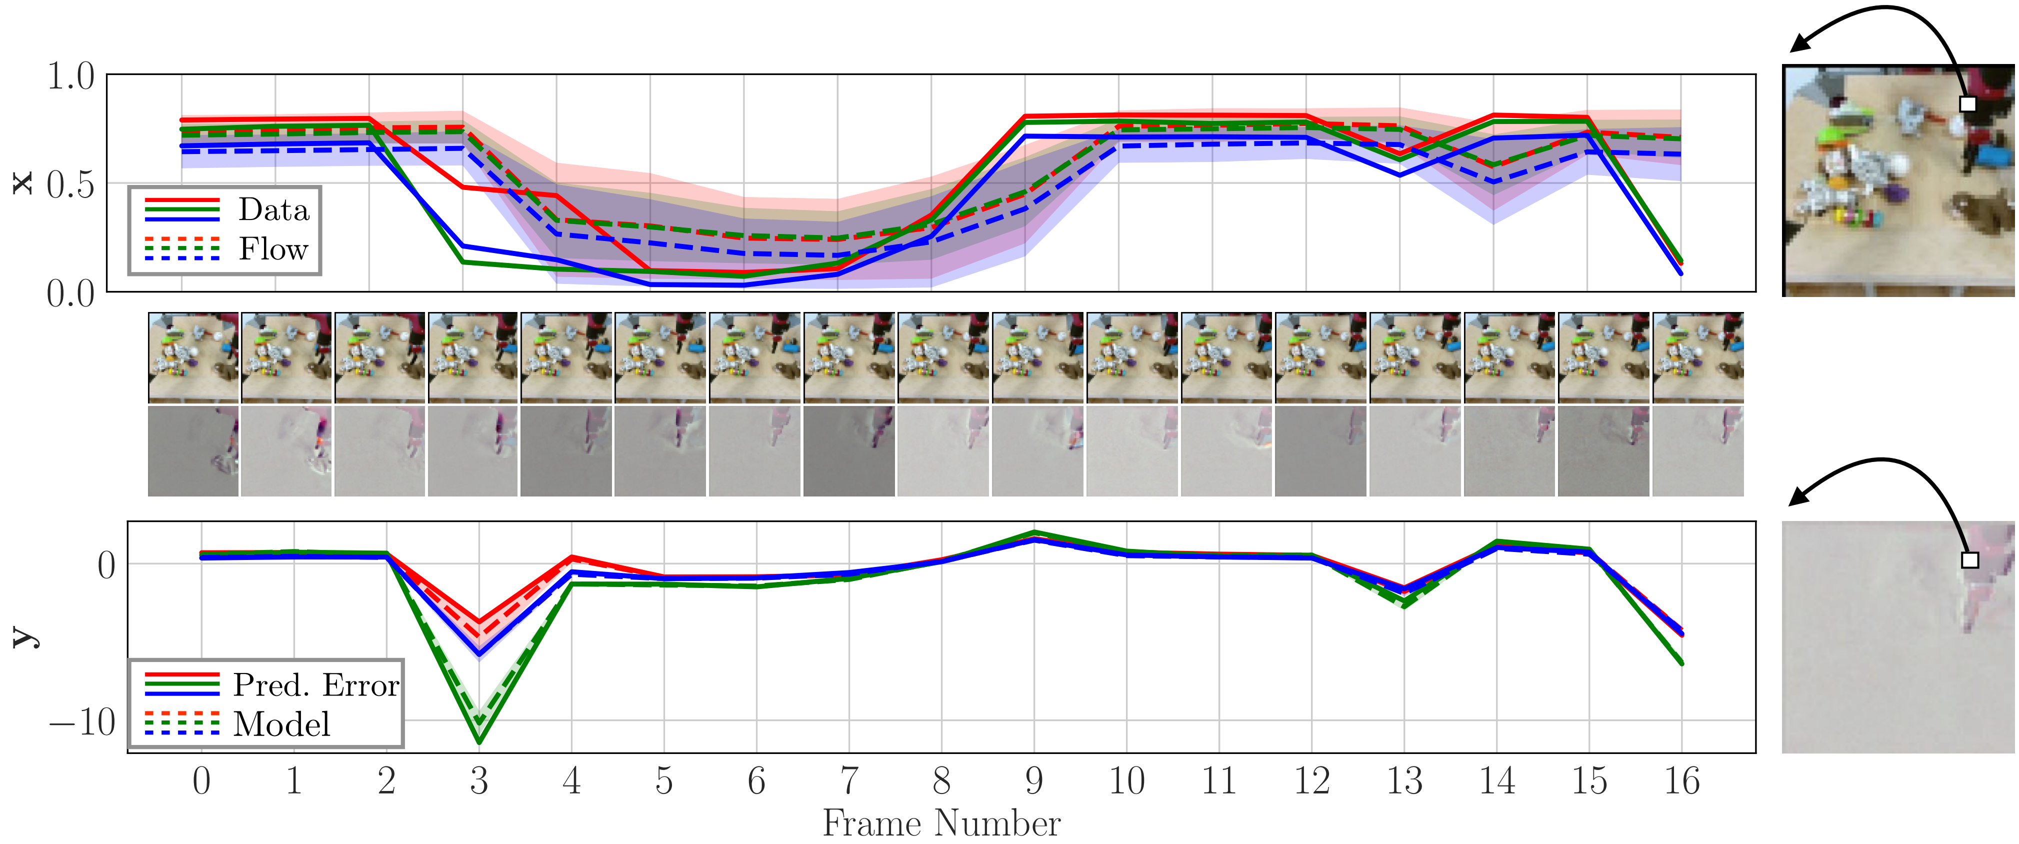

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

improves both ***performance*** & ***generalization***

KTH Actions (Schuldt et at., 2005)

different *people* across train & test data

histogram over data examples

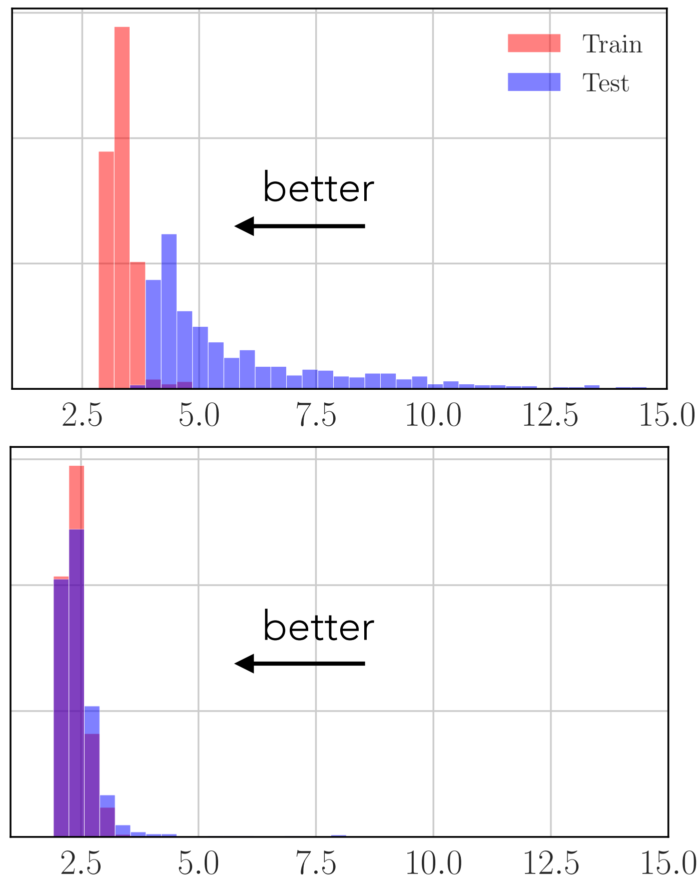

performance

SLVM

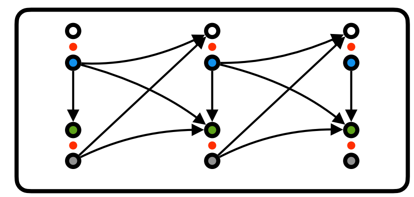

SLVM + AF

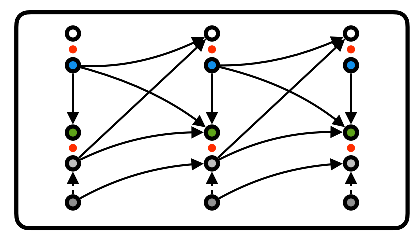

Marino, Chen, He, Mandt 2020

FEEDFORWARD PERCEPTION

state-of-the-art high-resolution video compression

UVG dataset (Mercat et al., 2020)

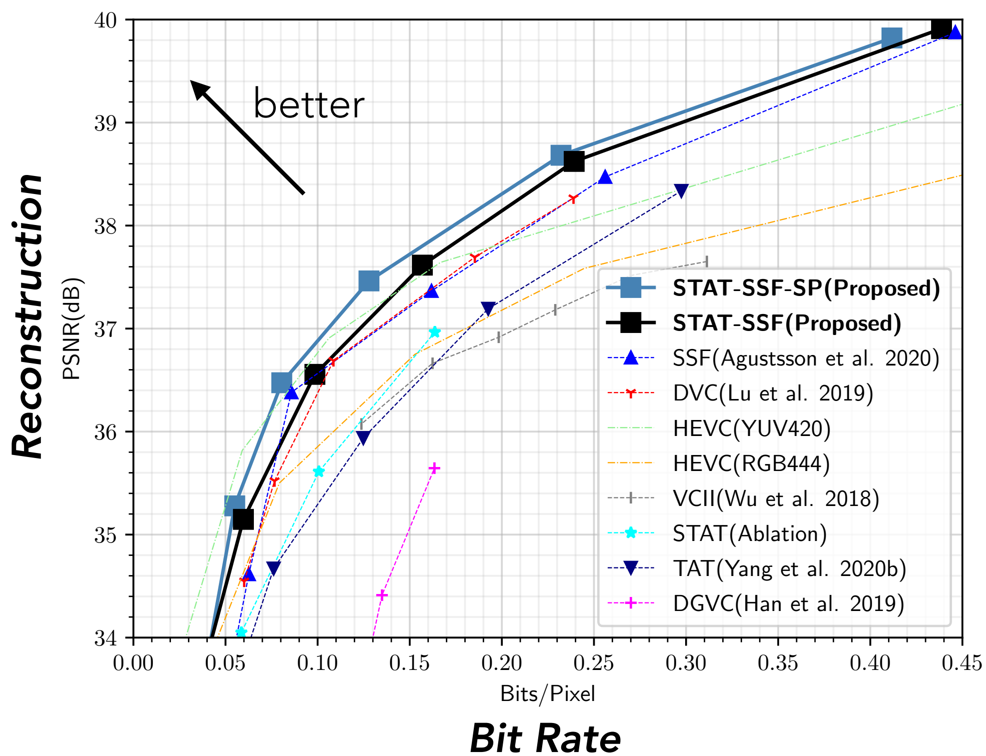

Yang, Yang, Marino, Mandt 2021

feedforward

CONTROL

FEEDFORWARD CONTROL

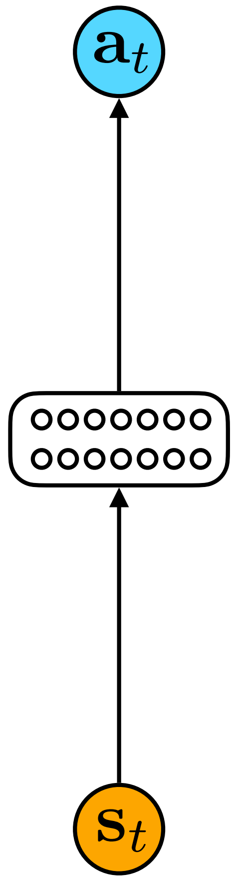

Direct Policy
$\pi(\mathbf{a}_t|\mathbf{s}_t)$

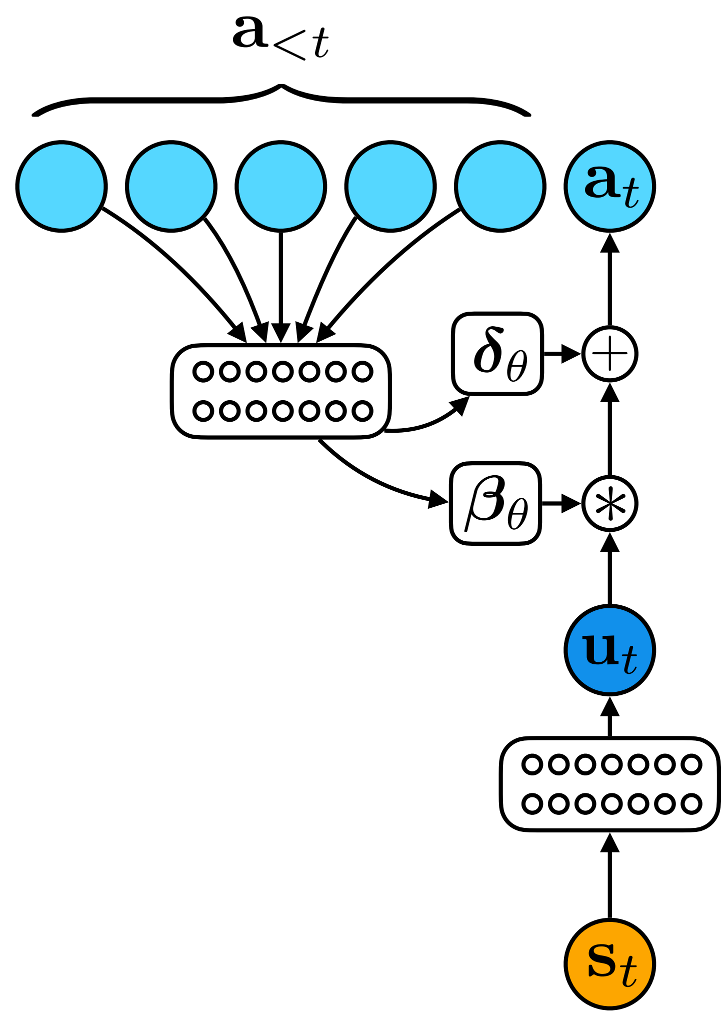

Autoregressive Policy
$\pi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{a}_{<t})$

Guerra & Marino, 2020

FEEDFORWARD CONTROL

simulated robotics environments from DeepMind control suite

`cheetah`

`quadruped`

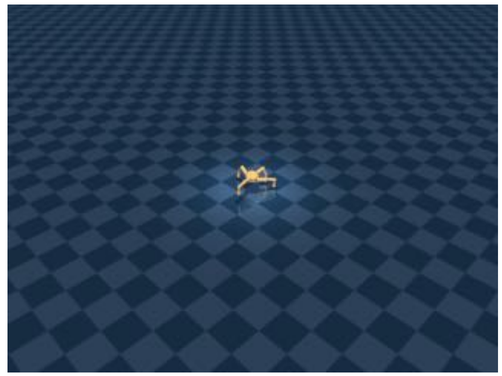

`hopper`

`walker`

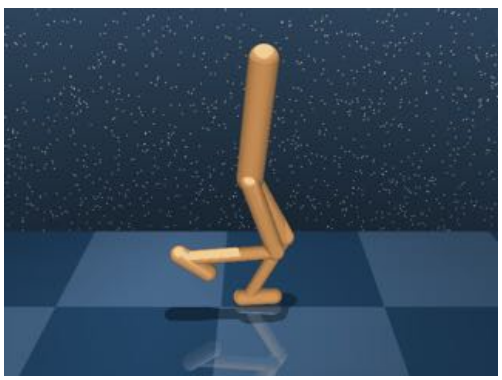

`humanoid`

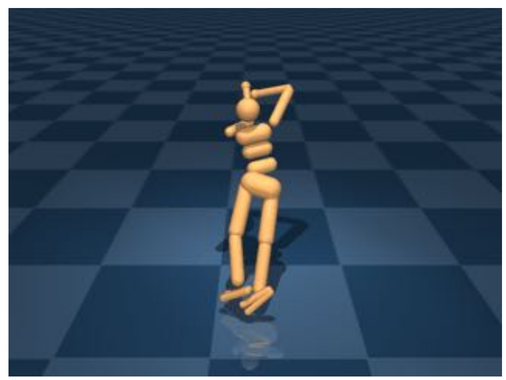

Guerra & Marino, 2020

FEEDFORWARD CONTROL

*improves **performance** across a range of environments*

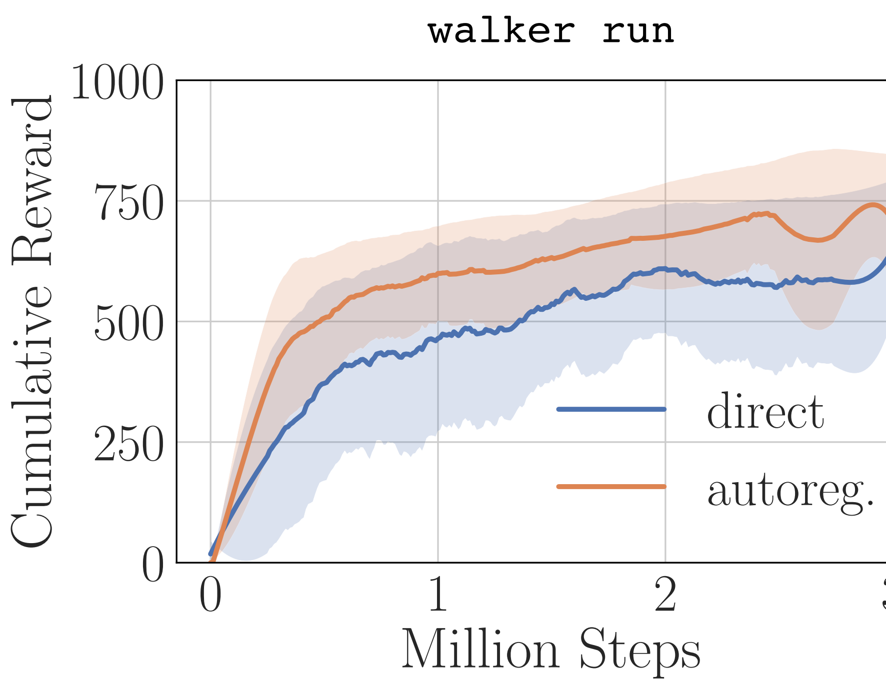

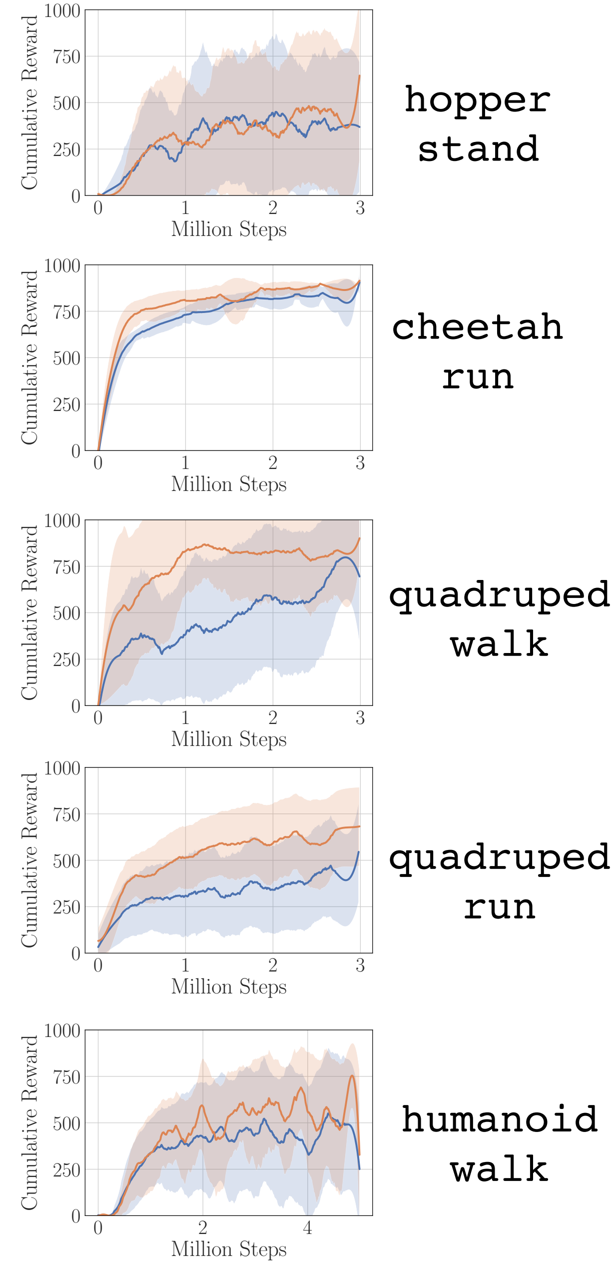

Guerra & Marino, 2020

FEEDFORWARD CONTROL

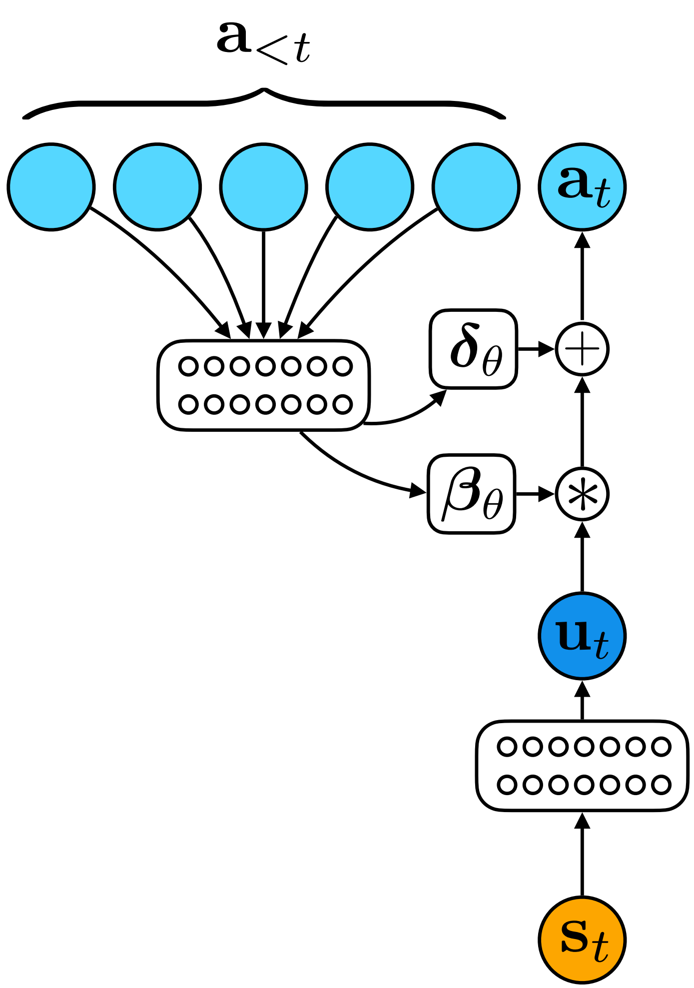

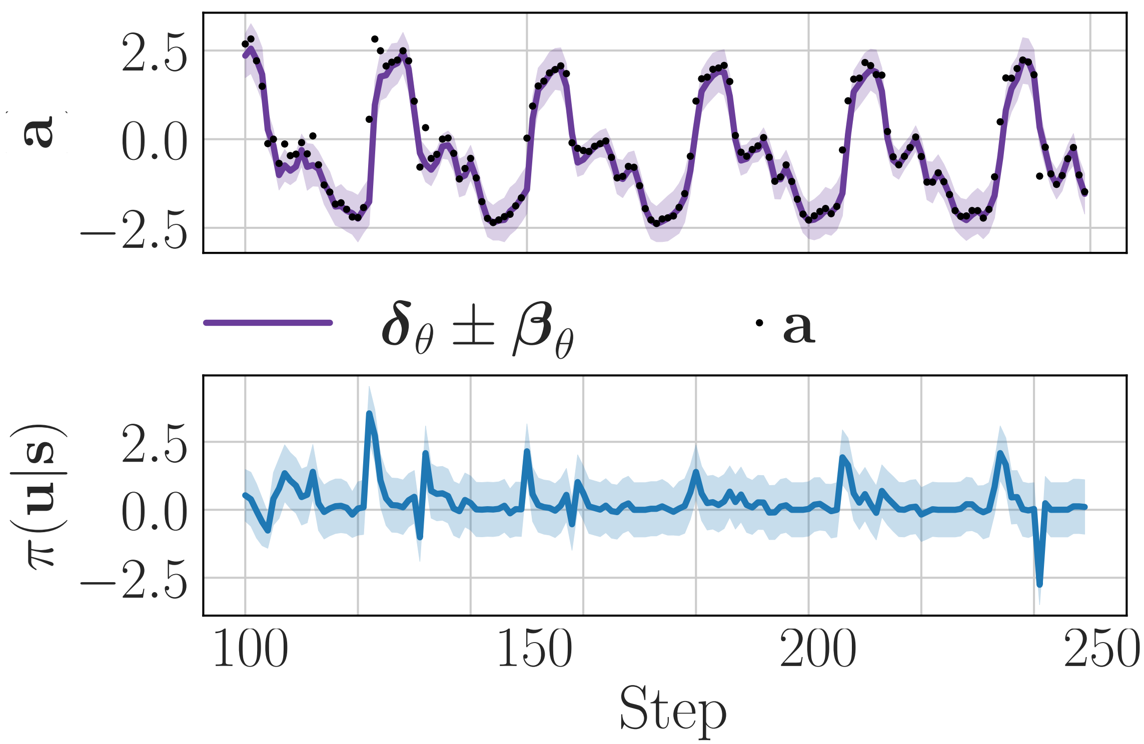

Guerra & Marino, 2020

feedback

PERCEPTION

HIERARCHICAL PREDICTIVE CODING

*neocortex forms hierarchical predictions of sensory inputs, using **prediction errors** for inference and learning*

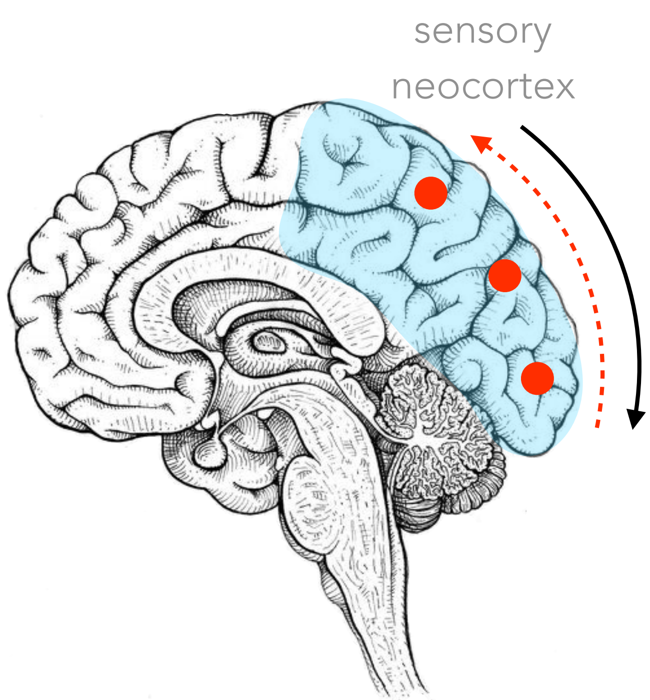

Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects

Rajesh P. N. Rao[1] and Dana H. Ballard[2]

1999

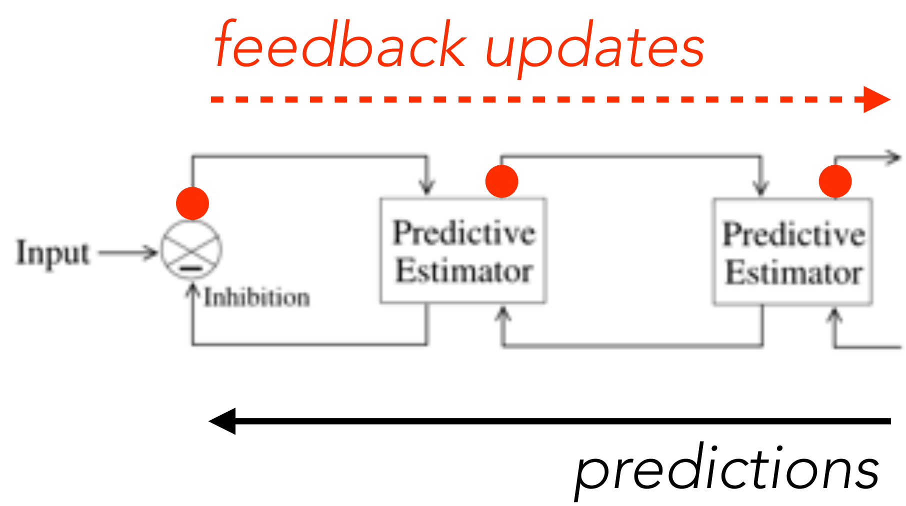

FEEDBACK PERCEPTION

state estimation as an *optimization* problem

$$\boldsymbol{\lambda} \leftarrow \arg\max_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}; q_{\boldsymbol{\lambda}})$$

restricting to parametric distributions with parameters $\boldsymbol{\lambda}$

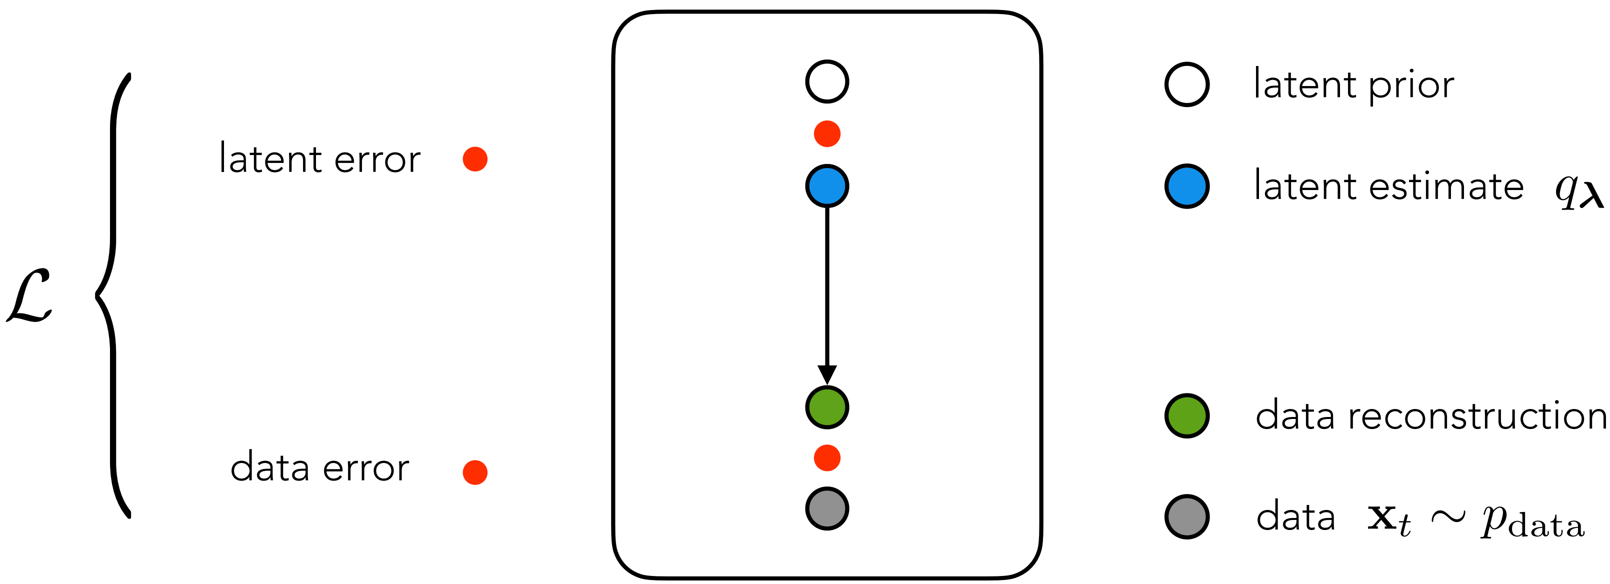

FEEDBACK PERCEPTION

state estimation as an *optimization* problem

$$\boldsymbol{\lambda} \leftarrow \arg\max_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}; q_{\boldsymbol{\lambda}})$$

restricting to parametric distributions with parameters $\boldsymbol{\lambda}$

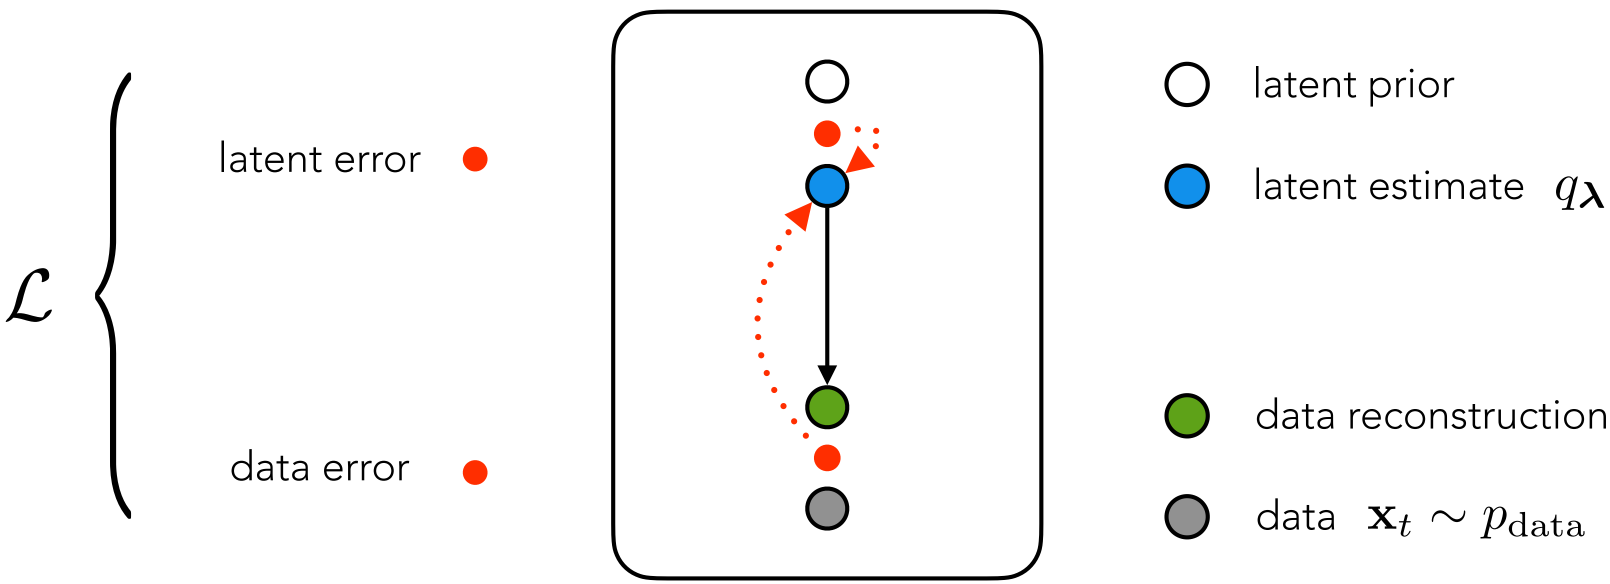

FEEDBACK PERCEPTION

but in practice…

variational autoencoder (VAE):

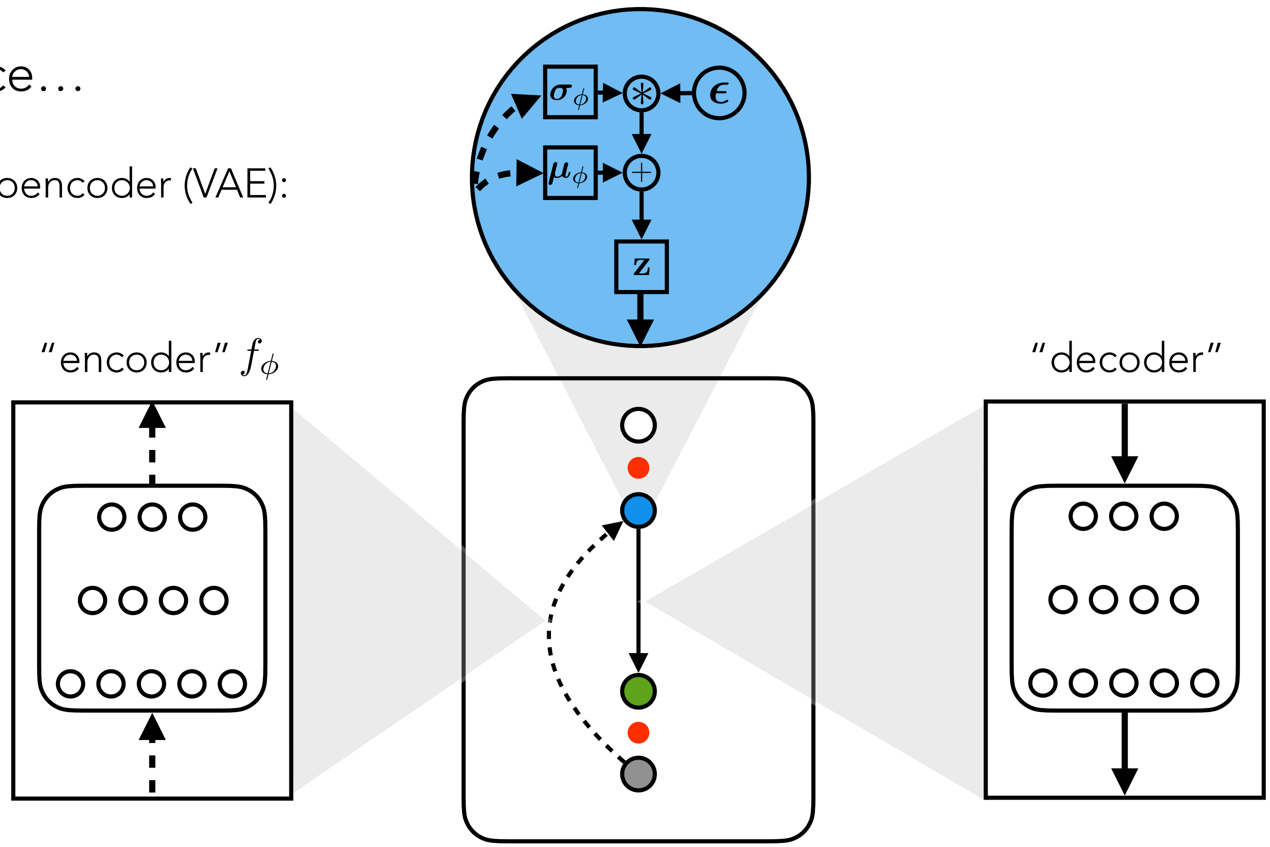

(direct) ***amortization:***

$$\boldsymbol{\lambda} \leftarrow f_\phi(\mathbf{x})$$

typically $\boldsymbol{\lambda} = [\boldsymbol{\mu}, \boldsymbol{\sigma}]$

optimize the "encoder" network f_ϕ

Kingma & Welling, 2014
Rezende et al., 2014

FEEDBACK PERCEPTION

direct inference models provide <u>suboptimal</u> estimates

"amortization gap"

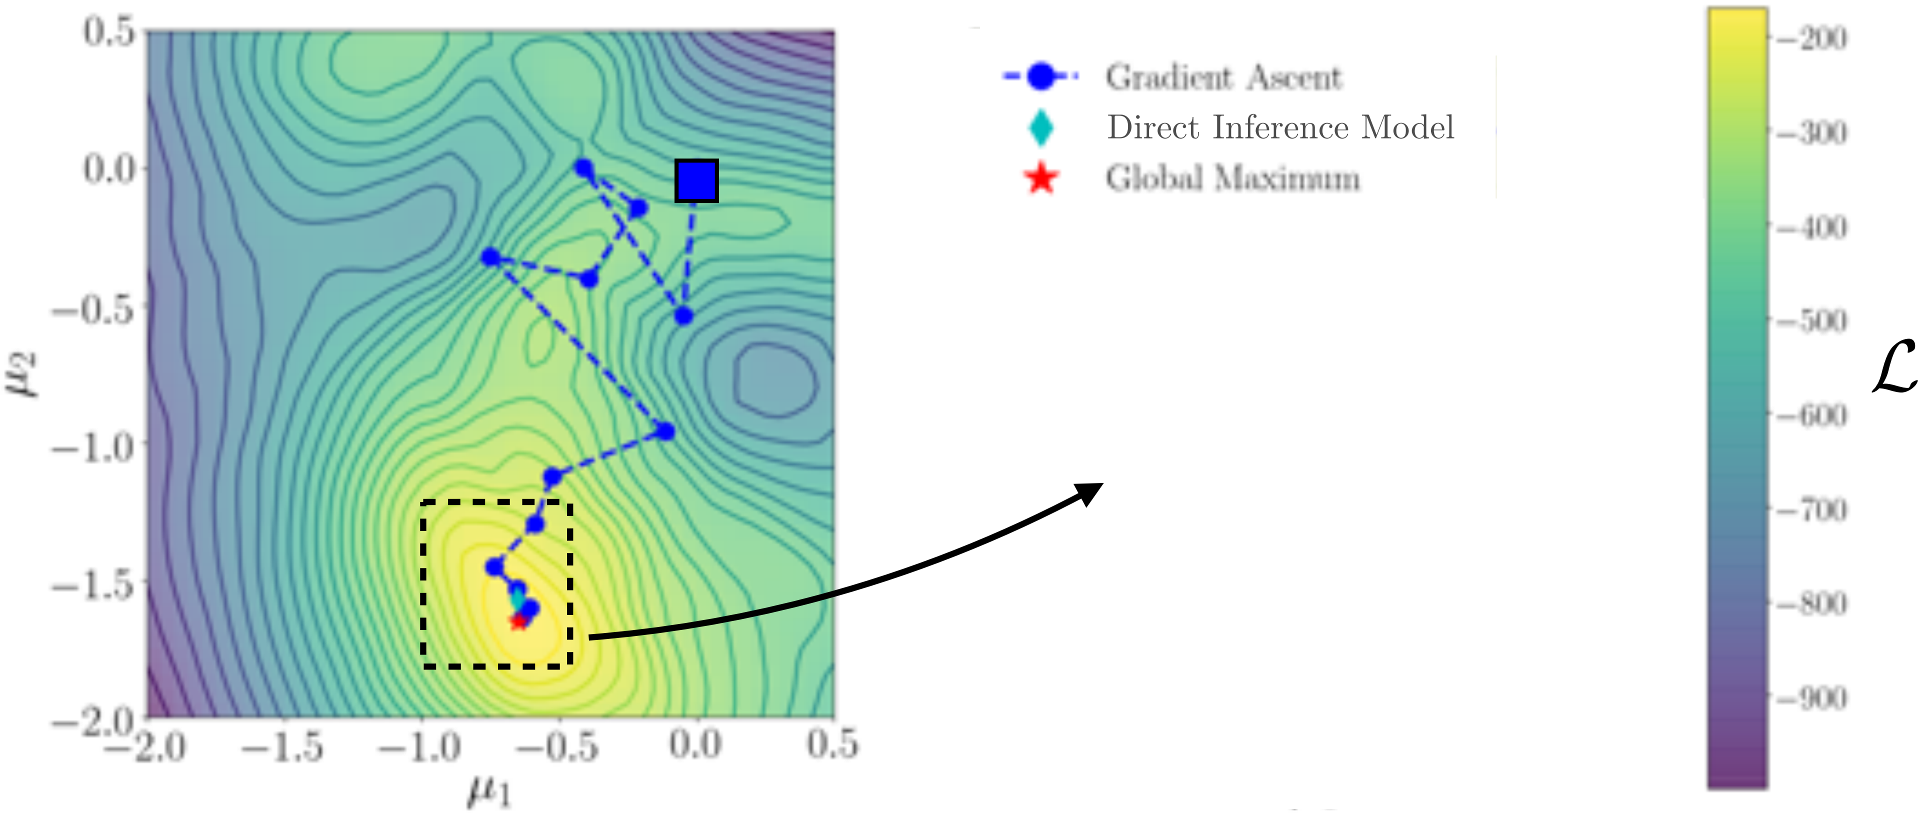

see also Cremer et al., 2018

Marino, Yue, Mandt, 2018

FEEDBACK PERCEPTION

formulate inference as an ***iterative*** amortized process

gradient-based form

$$\boldsymbol{\lambda} \leftarrow f_\phi(\boldsymbol{\lambda}, \nabla_{\boldsymbol{\lambda}} \mathcal{L})$$

error-based form (Gaussian)

$$\boldsymbol{\lambda} \leftarrow f_\phi(\boldsymbol{\lambda}, \underbrace{\boldsymbol{\xi}_{\mathbf{x}}, \boldsymbol{\xi}_{\mathbf{z}}}_{\text{weighted errors}})$$

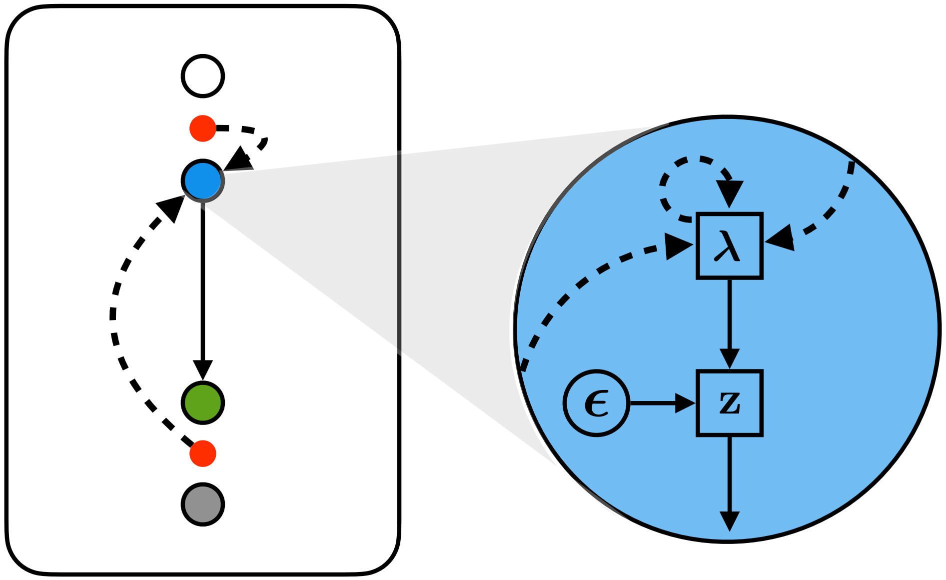

can *learn* to update

Marino, Yue, Mandt, 2018

FEEDBACK PERCEPTION

<u>accurate</u>

outperforms
direct amortization

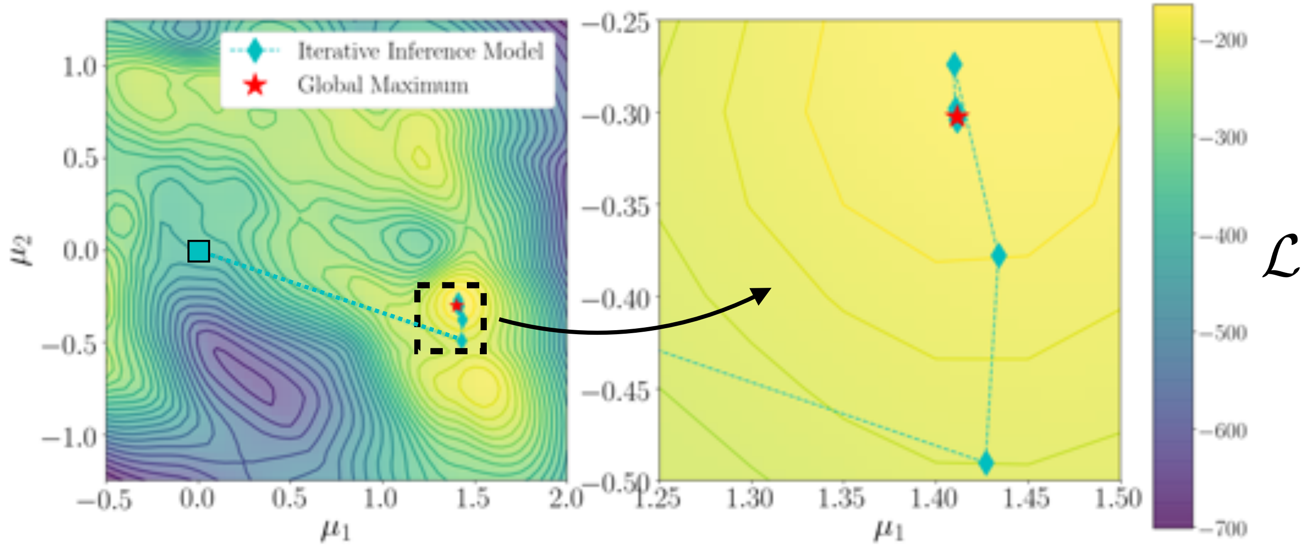

<u>*efficient*</u>

faster *than gradient-based optimization*

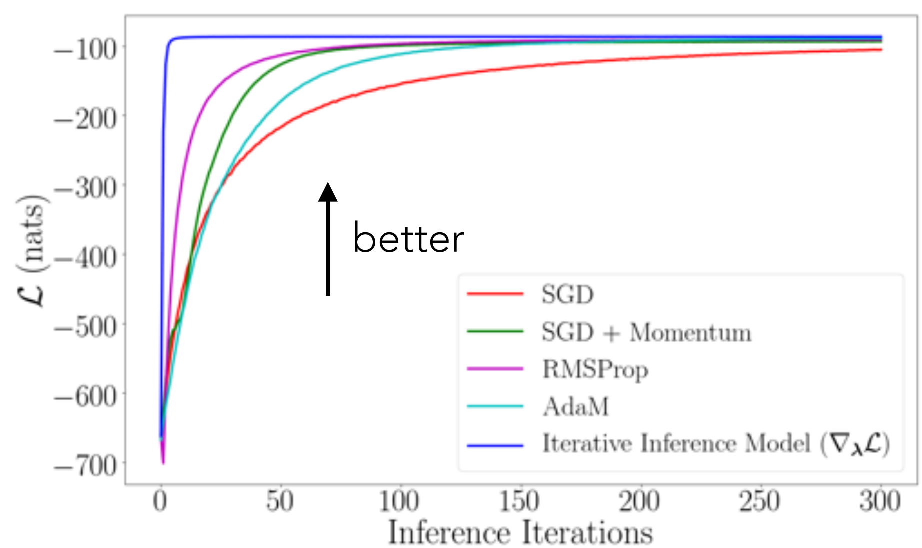

Marino, Yue, Mandt, 2018

FEEDBACK PERCEPTION

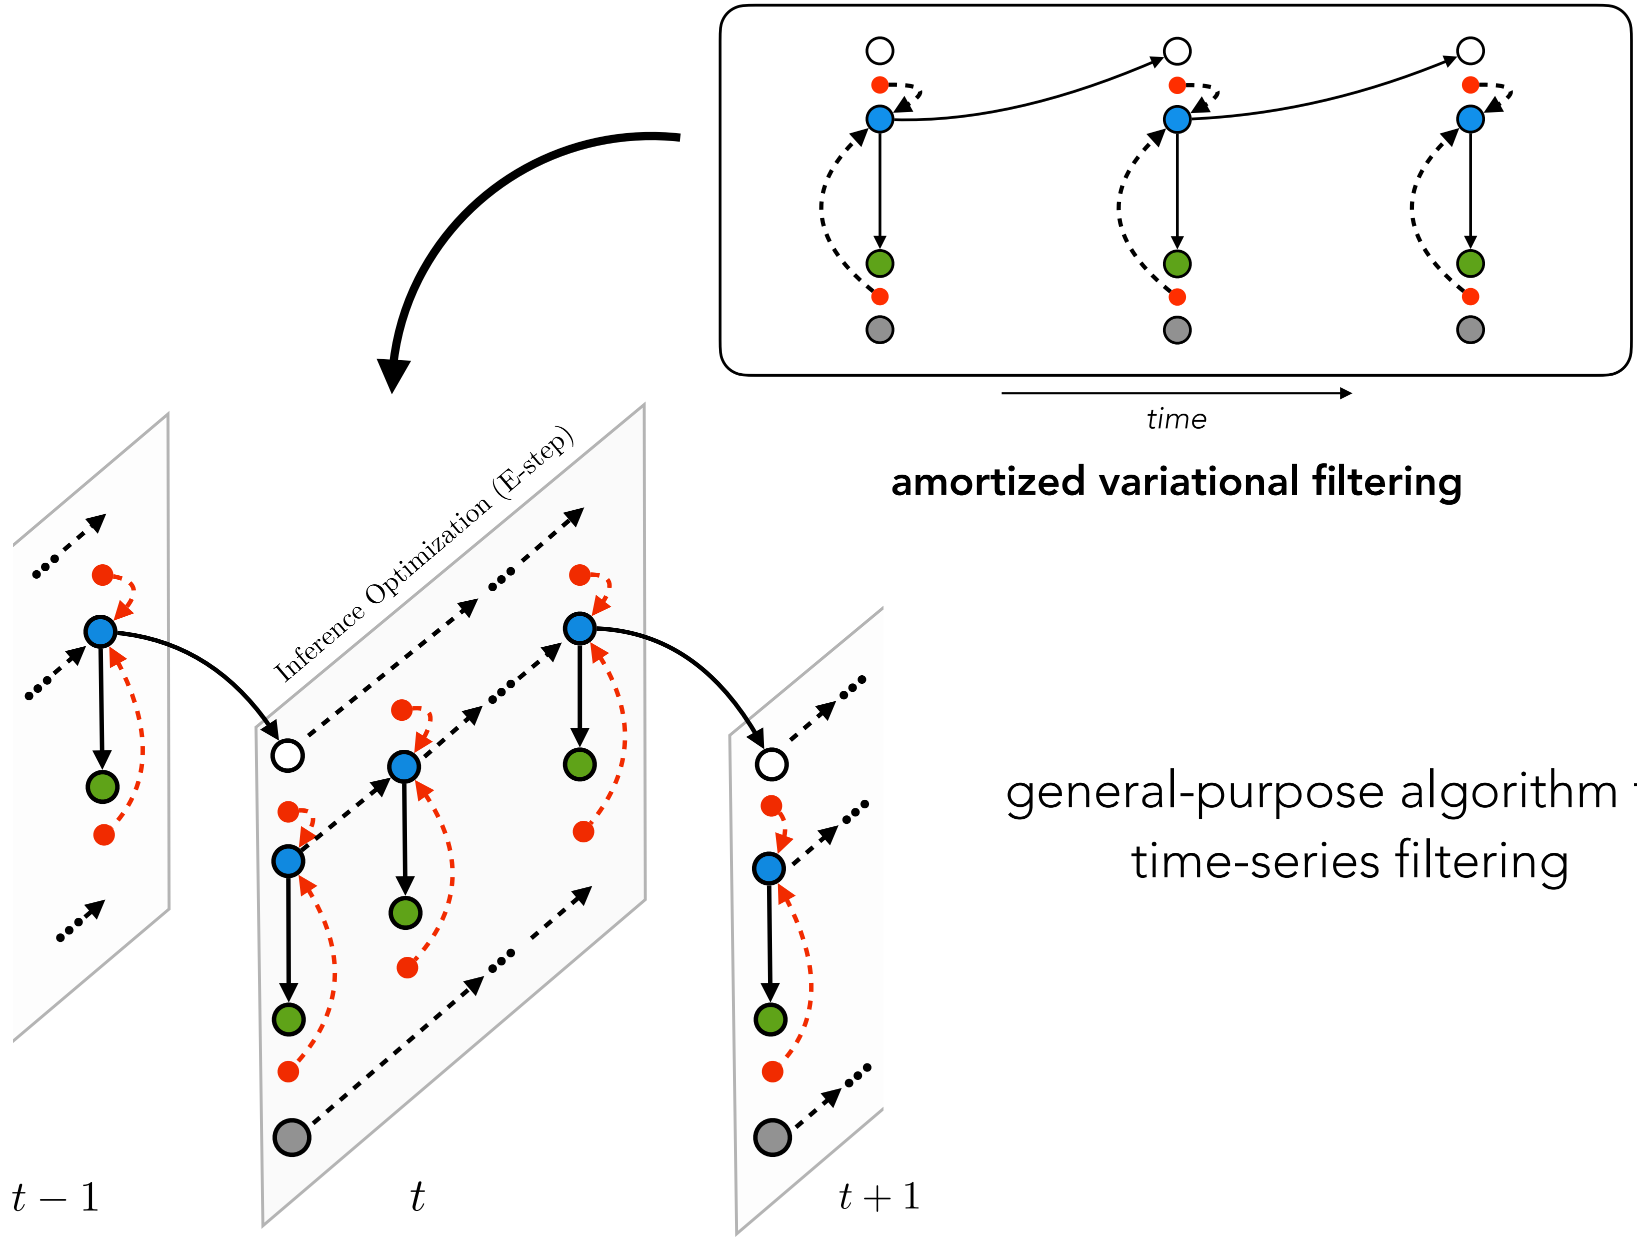

amortized variational filtering

general-purpose algorithm for time-series filtering

Marino, Cvitkovic, Yue, 2018

FEEDBACK PERCEPTION

custom (ad-hoc) amortized inference schemes

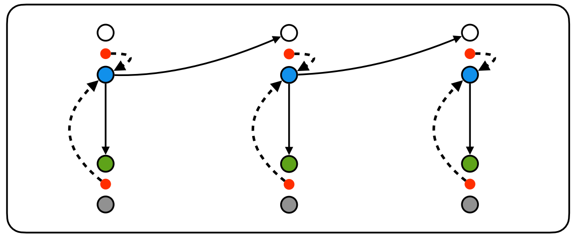

matches or outperforms each inference model
with a ***single* setup/architecture**

speech

video

MIDI

amortized variational filtering can be applied to <u>*any*</u> sequential LVM

Marino, Cvitkovic, Yue, 2018

FEEDBACK PERCEPTION

iterative improvement

more computation $\longrightarrow$ better performance

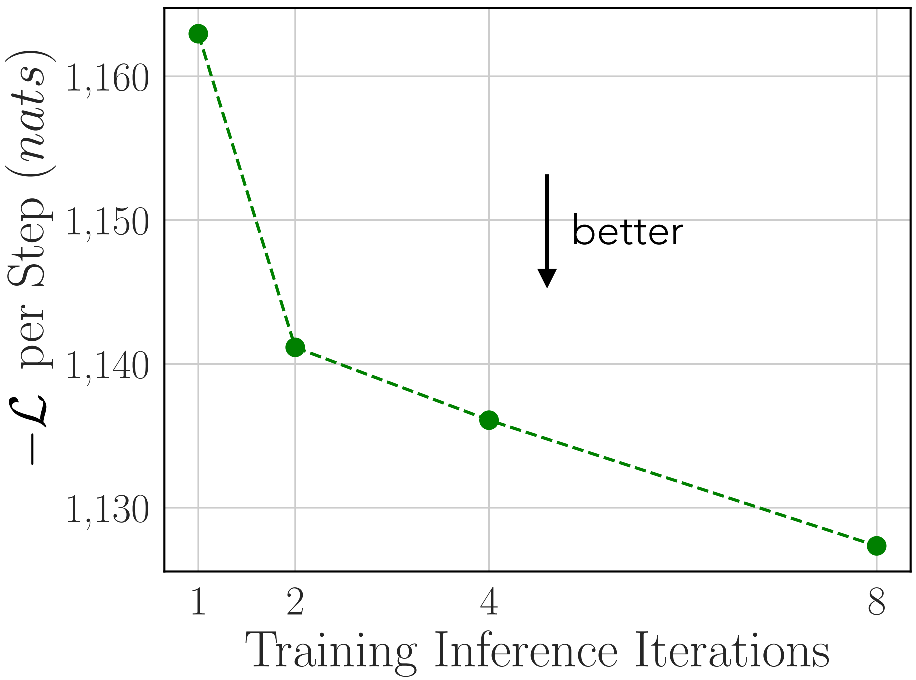

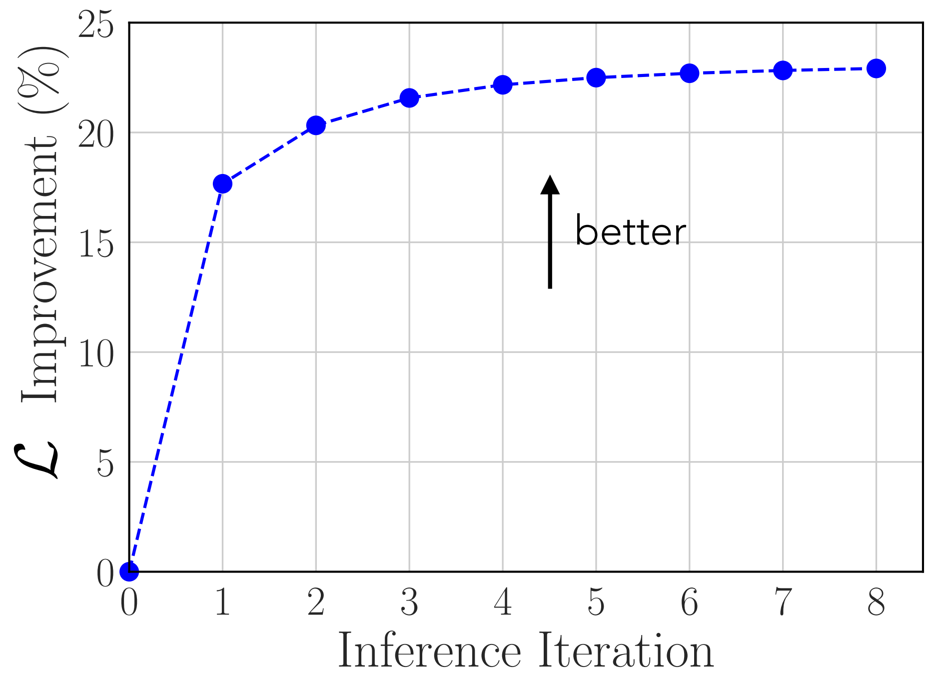

Marino, Cvitkovic, Yue, 2018

feedback

CONTROL

FEEDBACK CONTROL

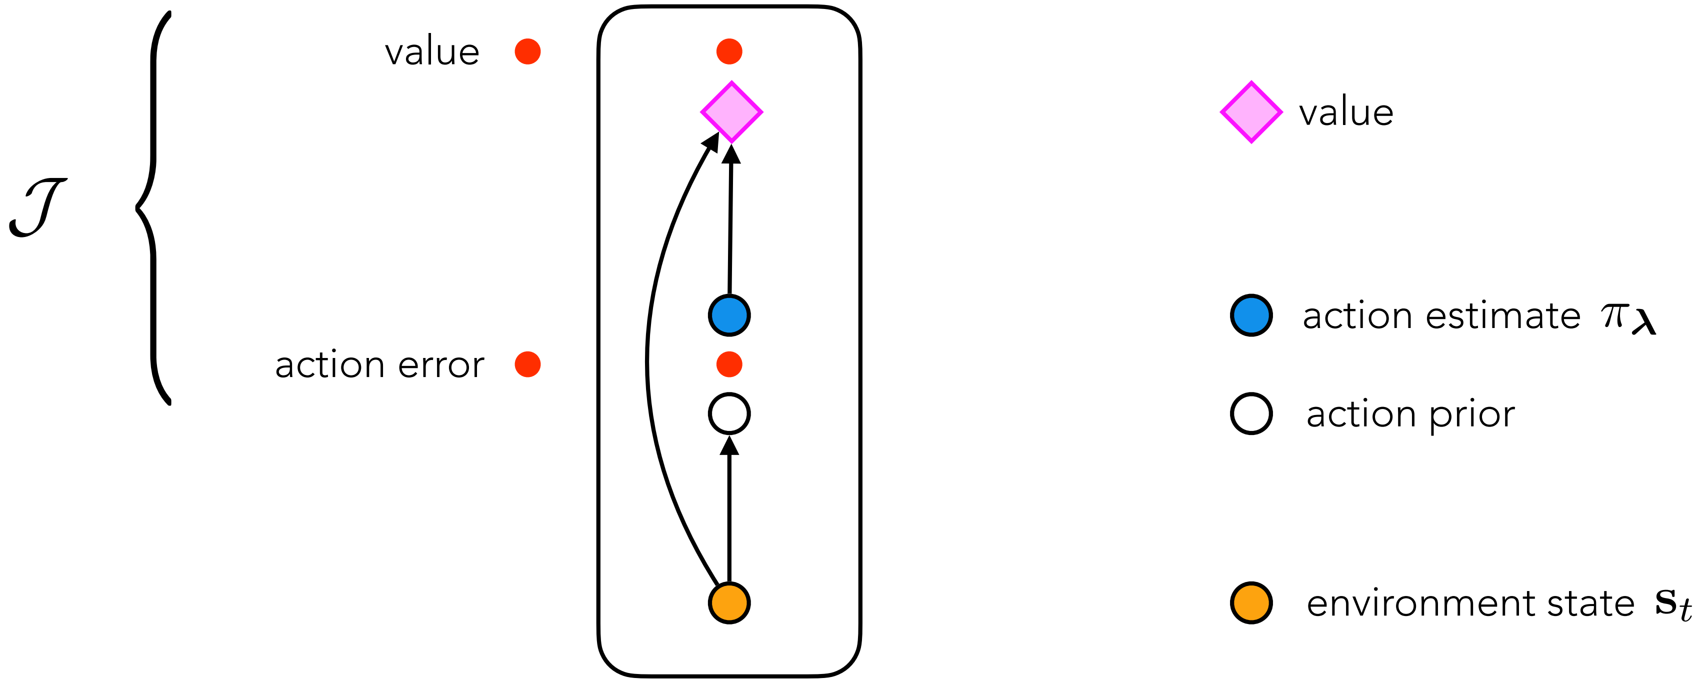

$$\mathrm{argmax}_{\boldsymbol{\lambda}} \mathcal{J}(\pi_{\boldsymbol{\lambda}}) \qquad \textit{e.g.,}\ \boldsymbol{\lambda} = [\boldsymbol{\mu}, \boldsymbol{\sigma}]$$

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

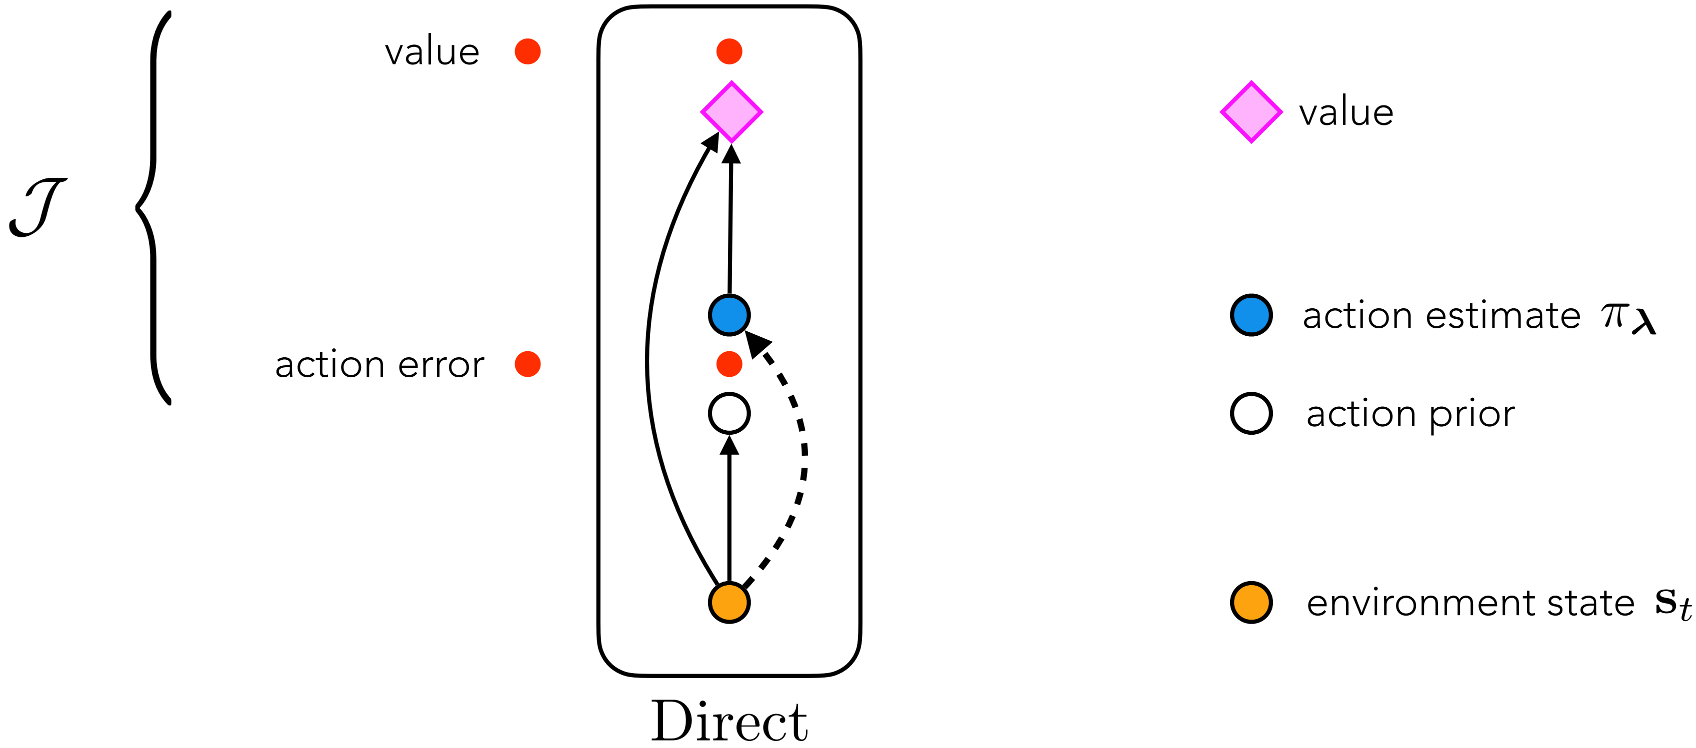

$$\mathrm{argmax}_{\boldsymbol{\lambda}} \mathcal{J}(\pi_{\boldsymbol{\lambda}}) \qquad \textit{e.g.}, \boldsymbol{\lambda} = [\boldsymbol{\mu}, \boldsymbol{\sigma}]$$

$$\boldsymbol{\lambda} \leftarrow f_{\phi}(\mathbf{s}_t)$$

direct

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

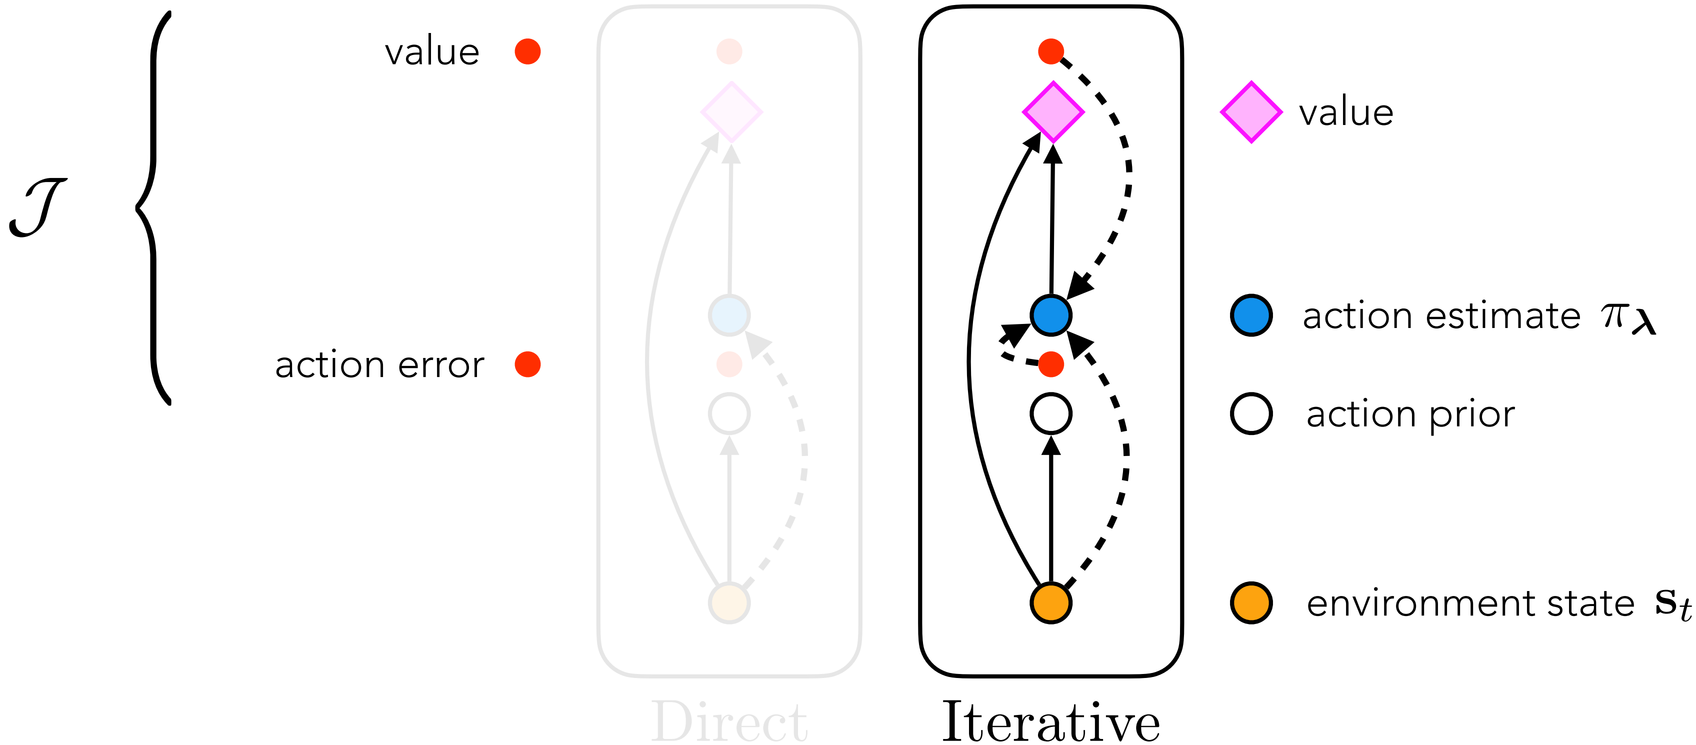

$$\mathrm{argmax}_{\boldsymbol{\lambda}} \mathcal{J}(\pi_{\boldsymbol{\lambda}}) \qquad e.g., \boldsymbol{\lambda} = [\boldsymbol{\mu}, \boldsymbol{\sigma}]$$

$$\boldsymbol{\lambda} \leftarrow f_\phi(\mathbf{s}_t)$$

direct

$$\boldsymbol{\lambda} \leftarrow f_\phi(\boldsymbol{\lambda}, \nabla_{\boldsymbol{\lambda}} \mathcal{J}, \mathbf{s}_t)$$

iterative

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

simulated robotics environments from OpenAI gym

`Hopper`

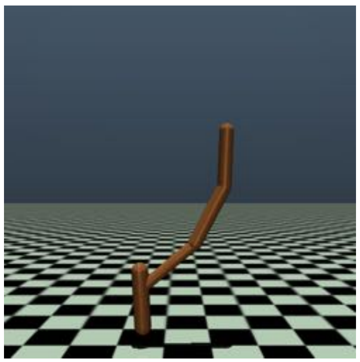

`HalfCheetah`

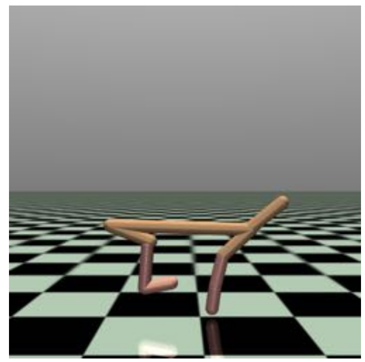

`Walker2d`

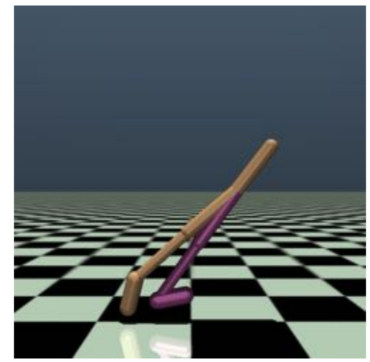

`Swimmer`

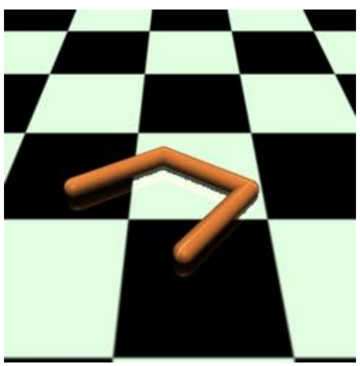

`Ant`

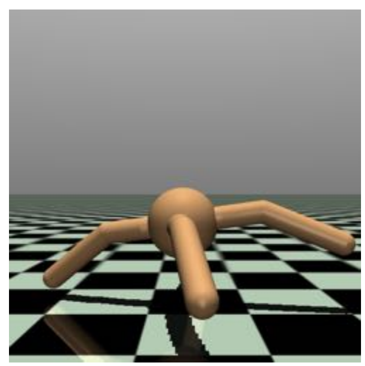

`Humanoid`

FEEDBACK CONTROL

direct policy networks yield suboptimal estimates

two dimensions of the policy

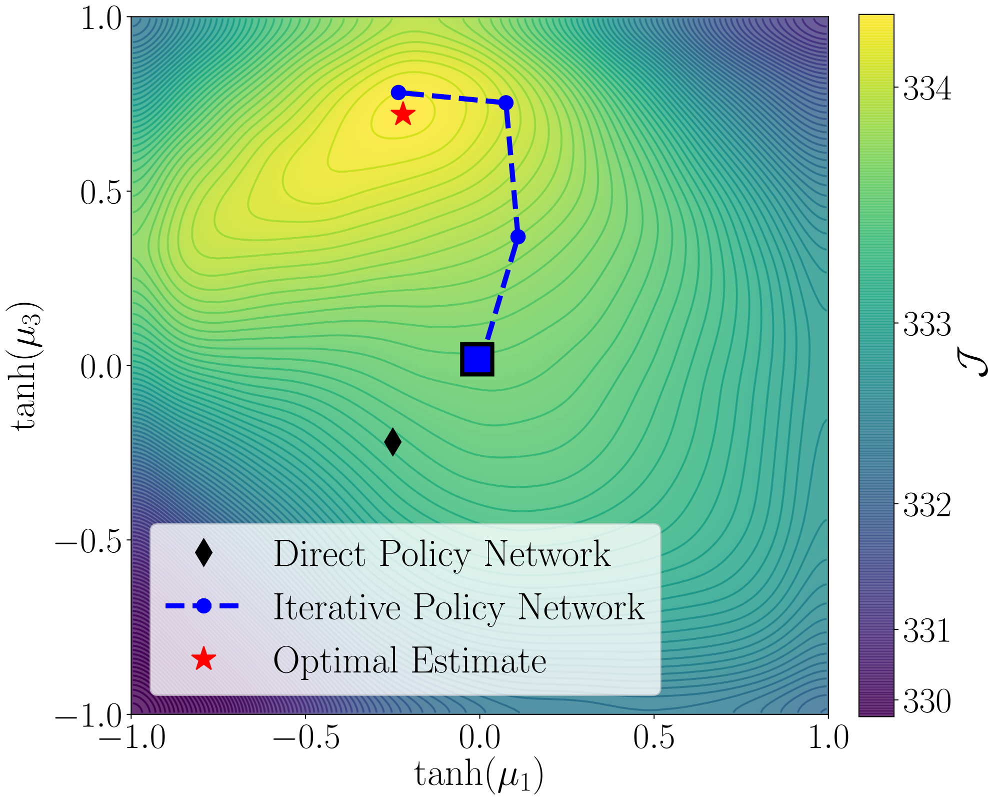

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

iterative optimizers yield <u>multiple</u> locally optimal estimates

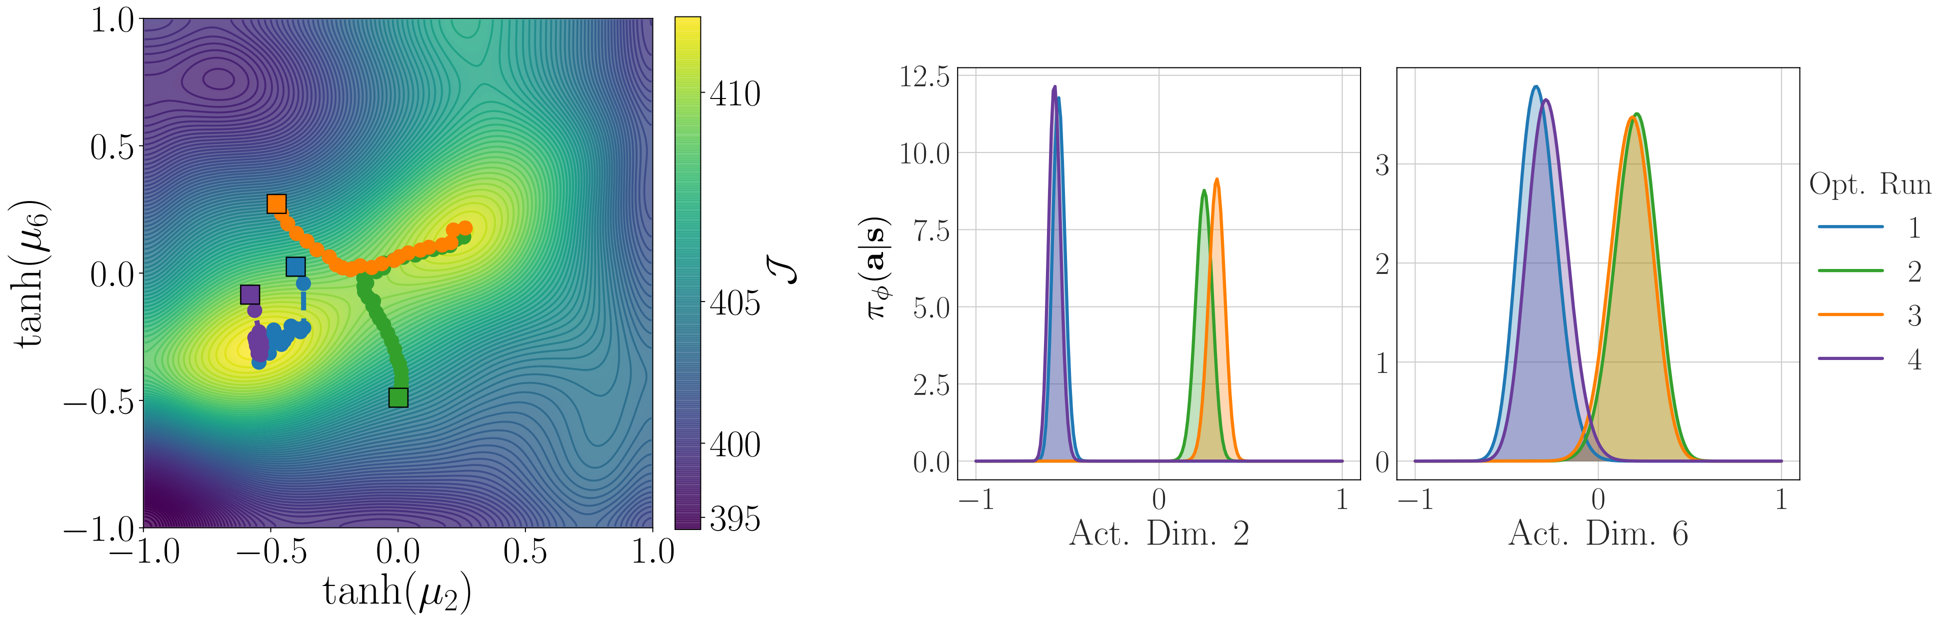

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

*improves **efficiency** over standard optimizers*

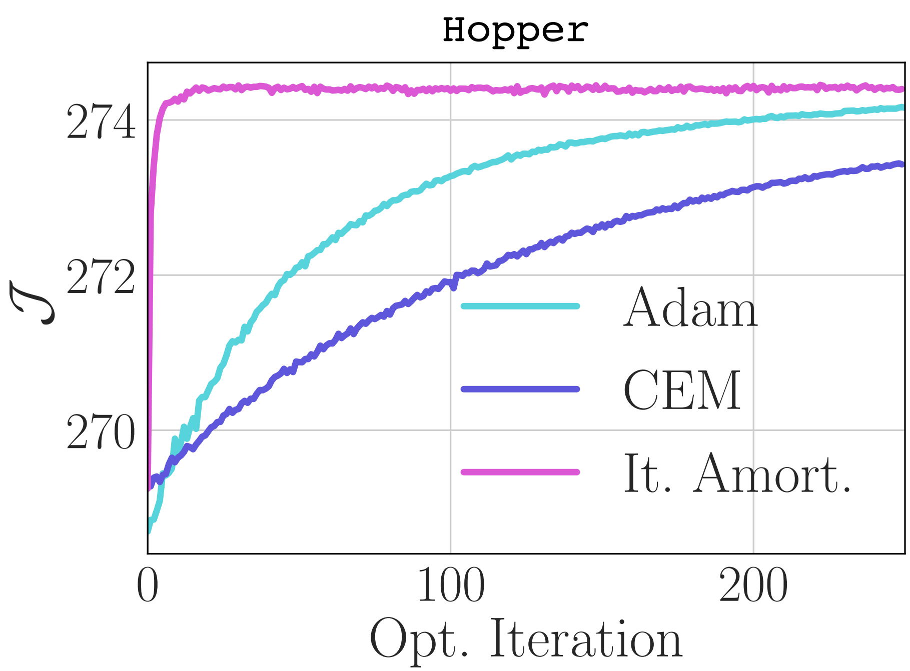

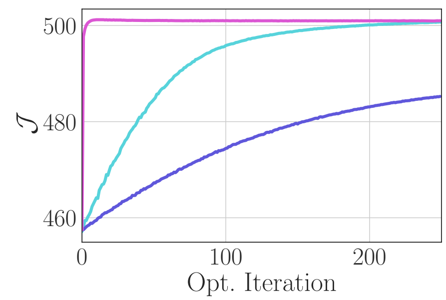

HalfCheetah

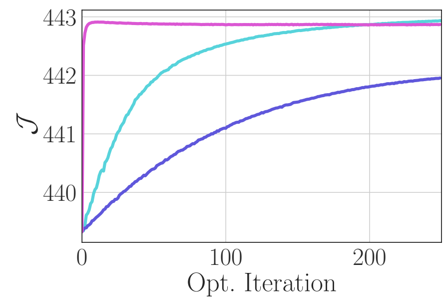

Walker2d

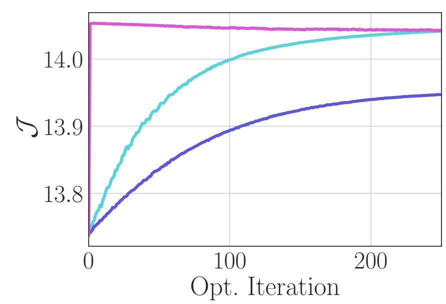

Swimmer

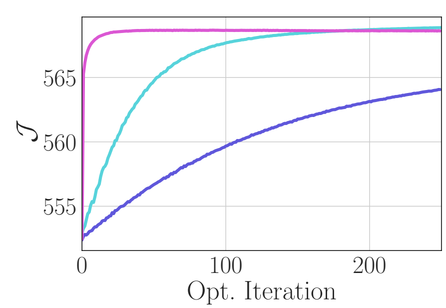

Ant

Marino, Piché, Ialongo, Yue 2020

FEEDBACK CONTROL

comparable or improved
***performance** over direct amortization*

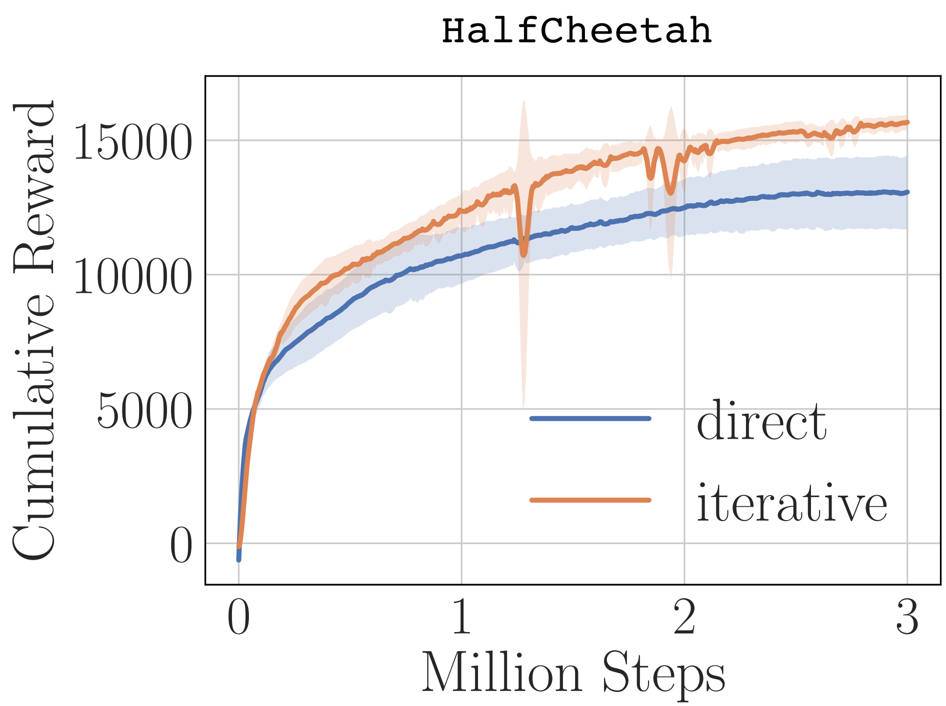

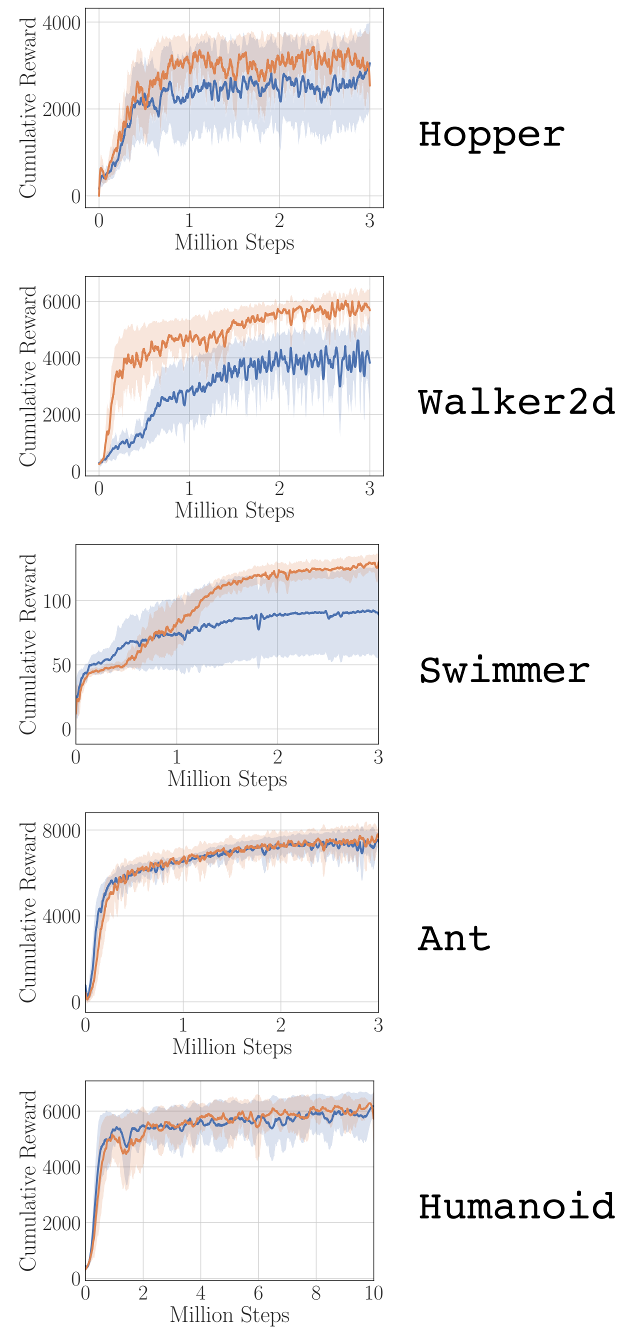

Marino, Piché, Ialongo, Yue 2020

connections to

NEUROSCIENCE

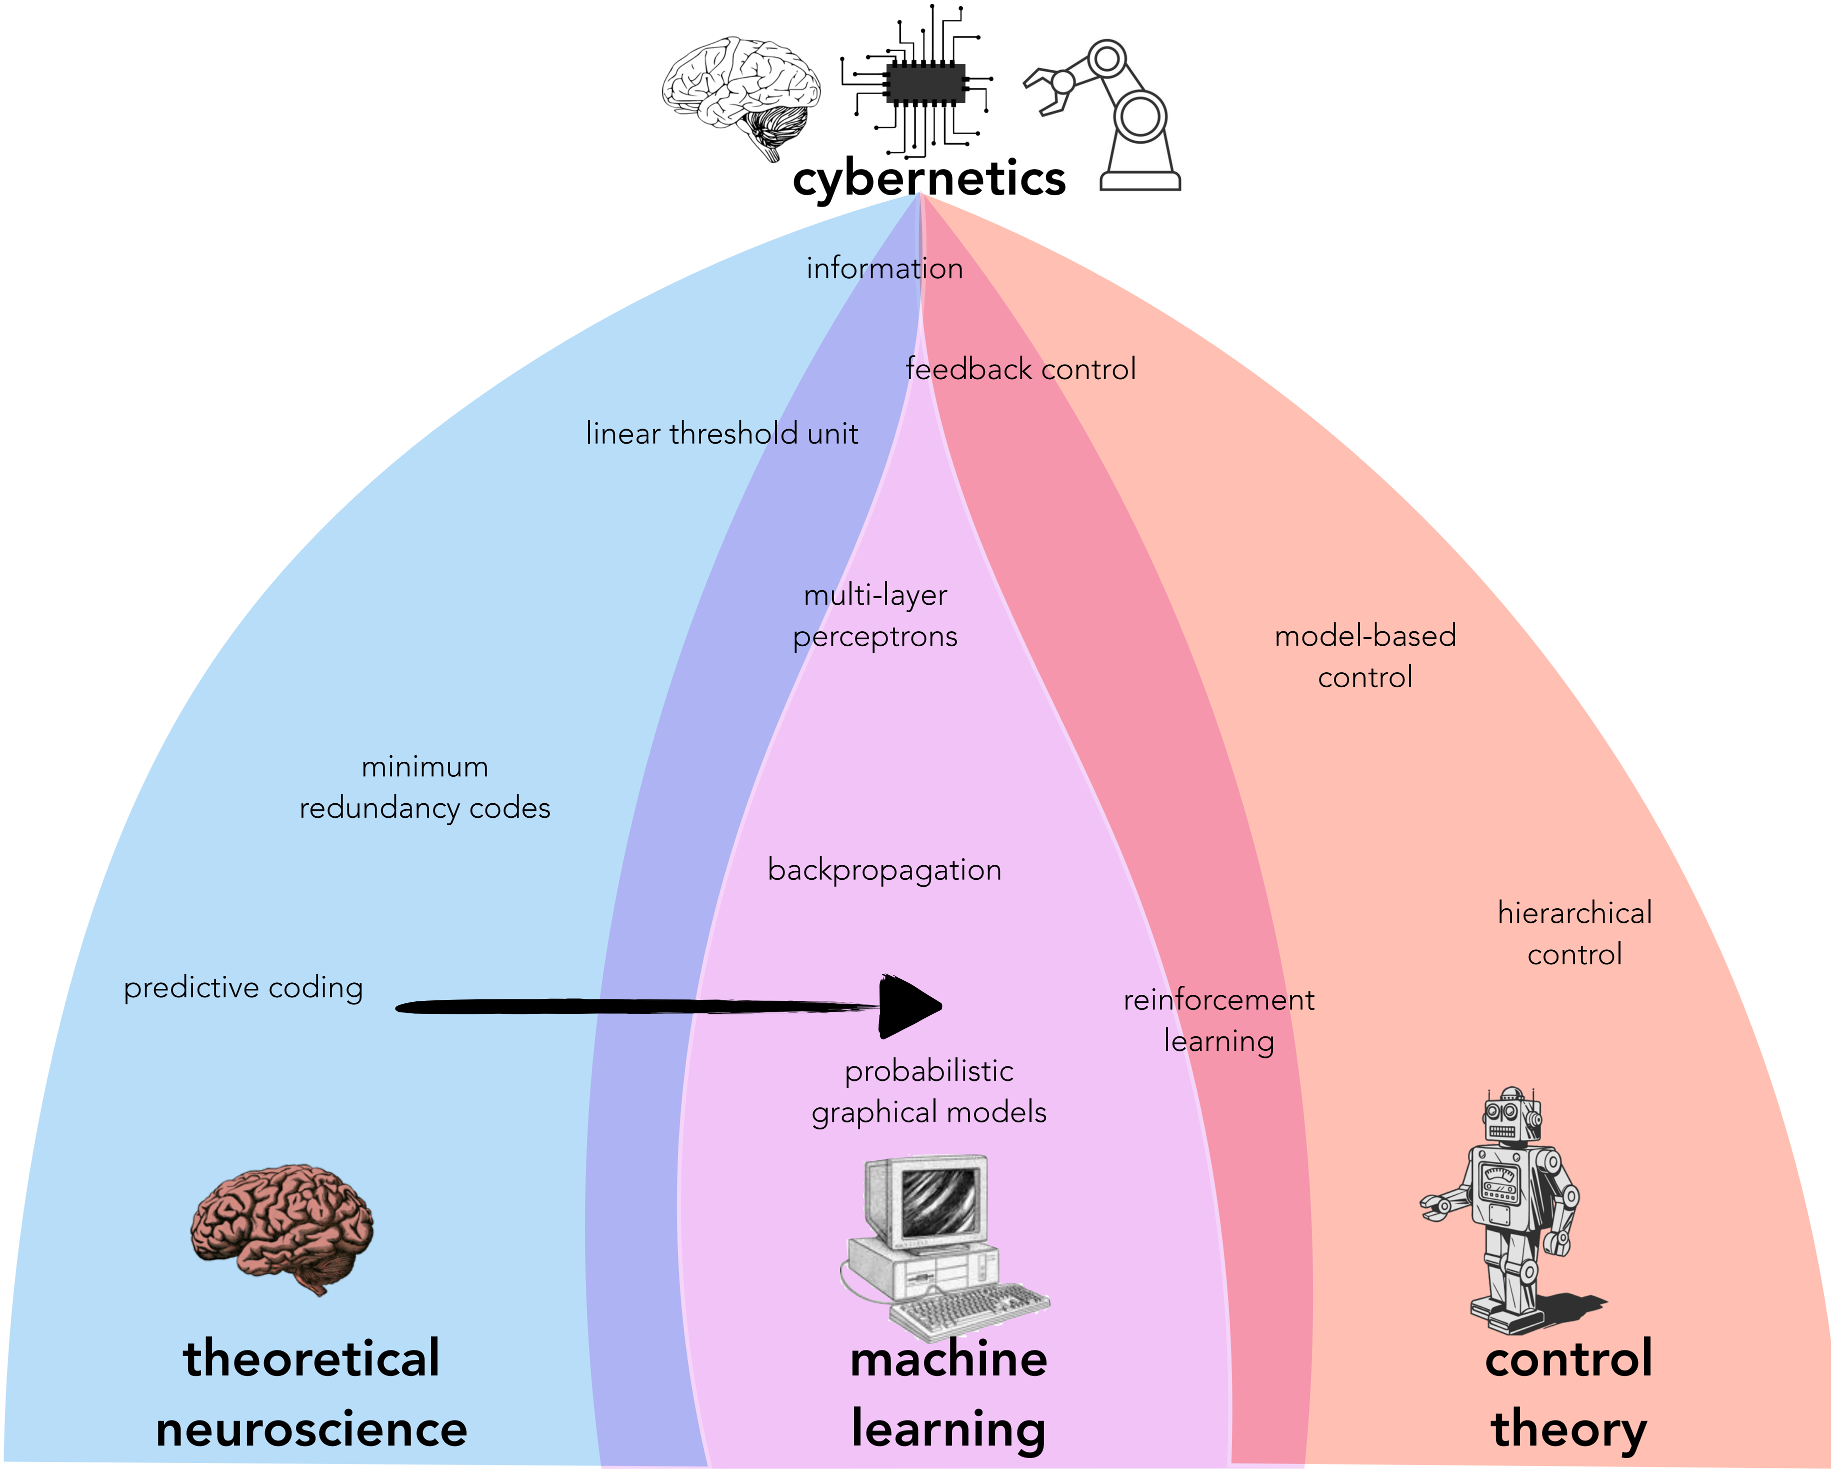
cybernetics
information
feedback control
linear threshold unit
multi-layer perceptrons
model-based control
minimum redundancy codes
backpropagation
hierarchical control
predictive coding
reinforcement learning
probabilistic graphical models
theoretical neuroscience
machine learning
control theory

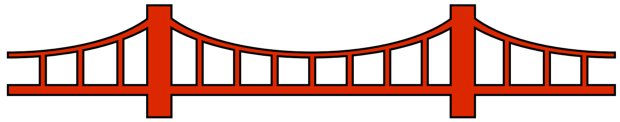

Neuroscience

predictive coding

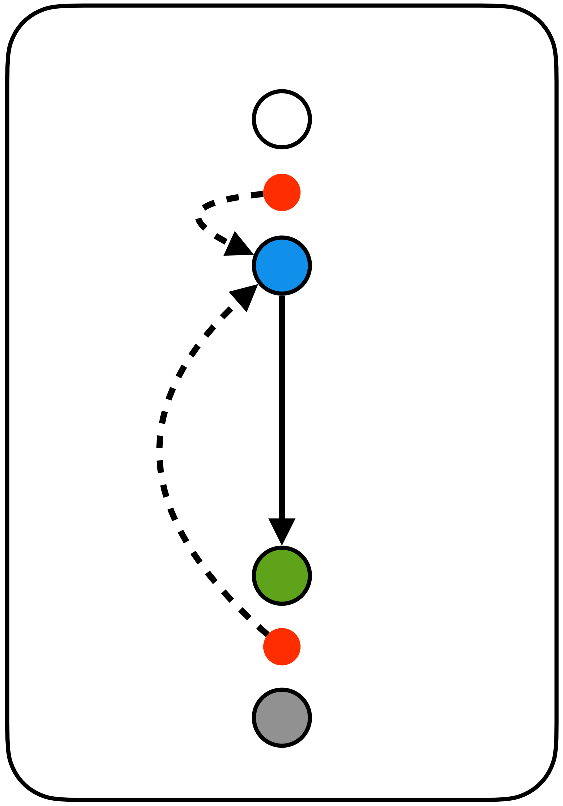

variational autoencoders

Machine Learning

providing new correspondences…

PYRAMIDAL NEURONS AS DEEP NETWORKS

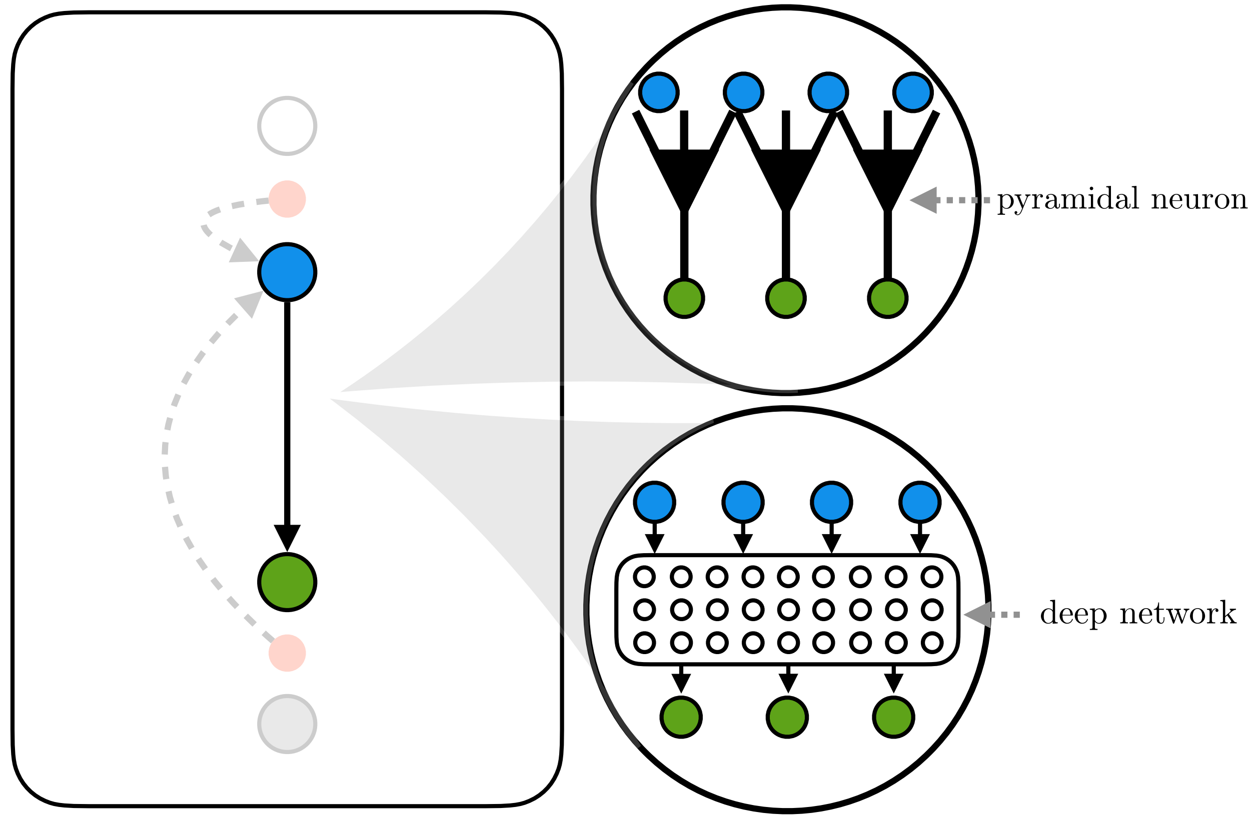

pyramidal neurons as multi-layer artificial networks

(Zador et al., 1992; Mel, 1992; Poirazi et al., 2003; Polsky et al., 2004; Gidon et al., 2020…)

PYRAMIDAL NEURONS AS DEEP NETWORKS

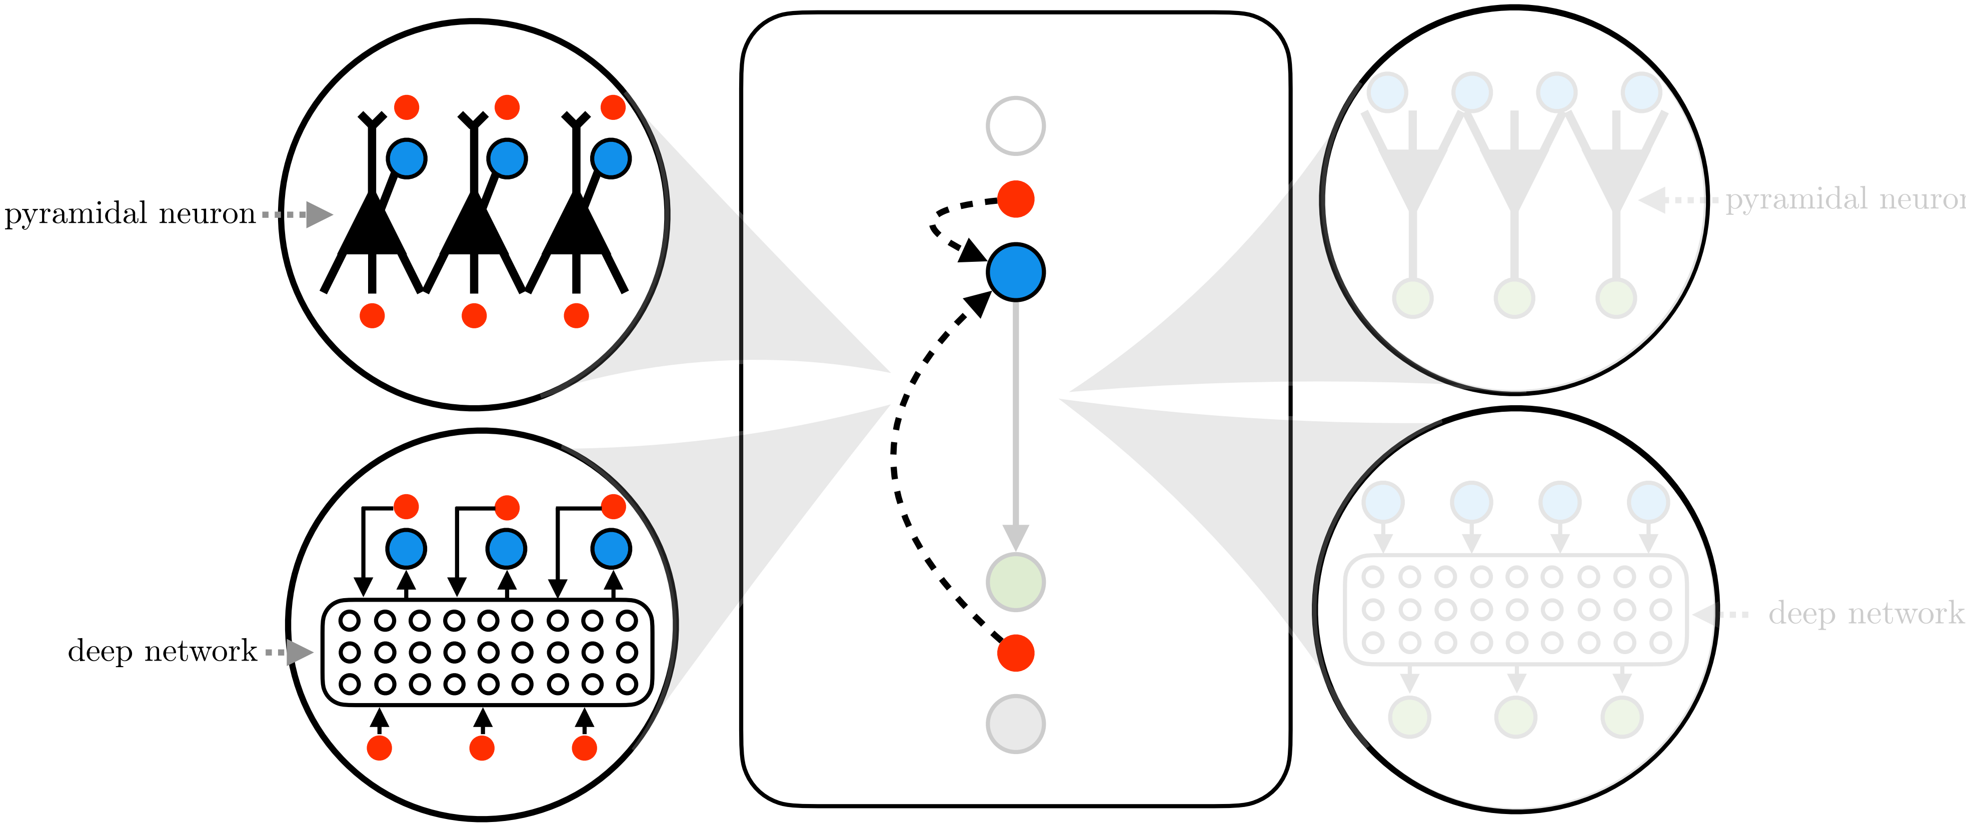

pyramidal neurons as multi-layer artificial networks

(Zador et al., 1992; Mel, 1992; Poirazi et al., 2003; Polsky et al., 2004; Gidon et al., 2020…)

Marino, 2019

PYRAMIDAL NEURONS AS DEEP NETWORKS

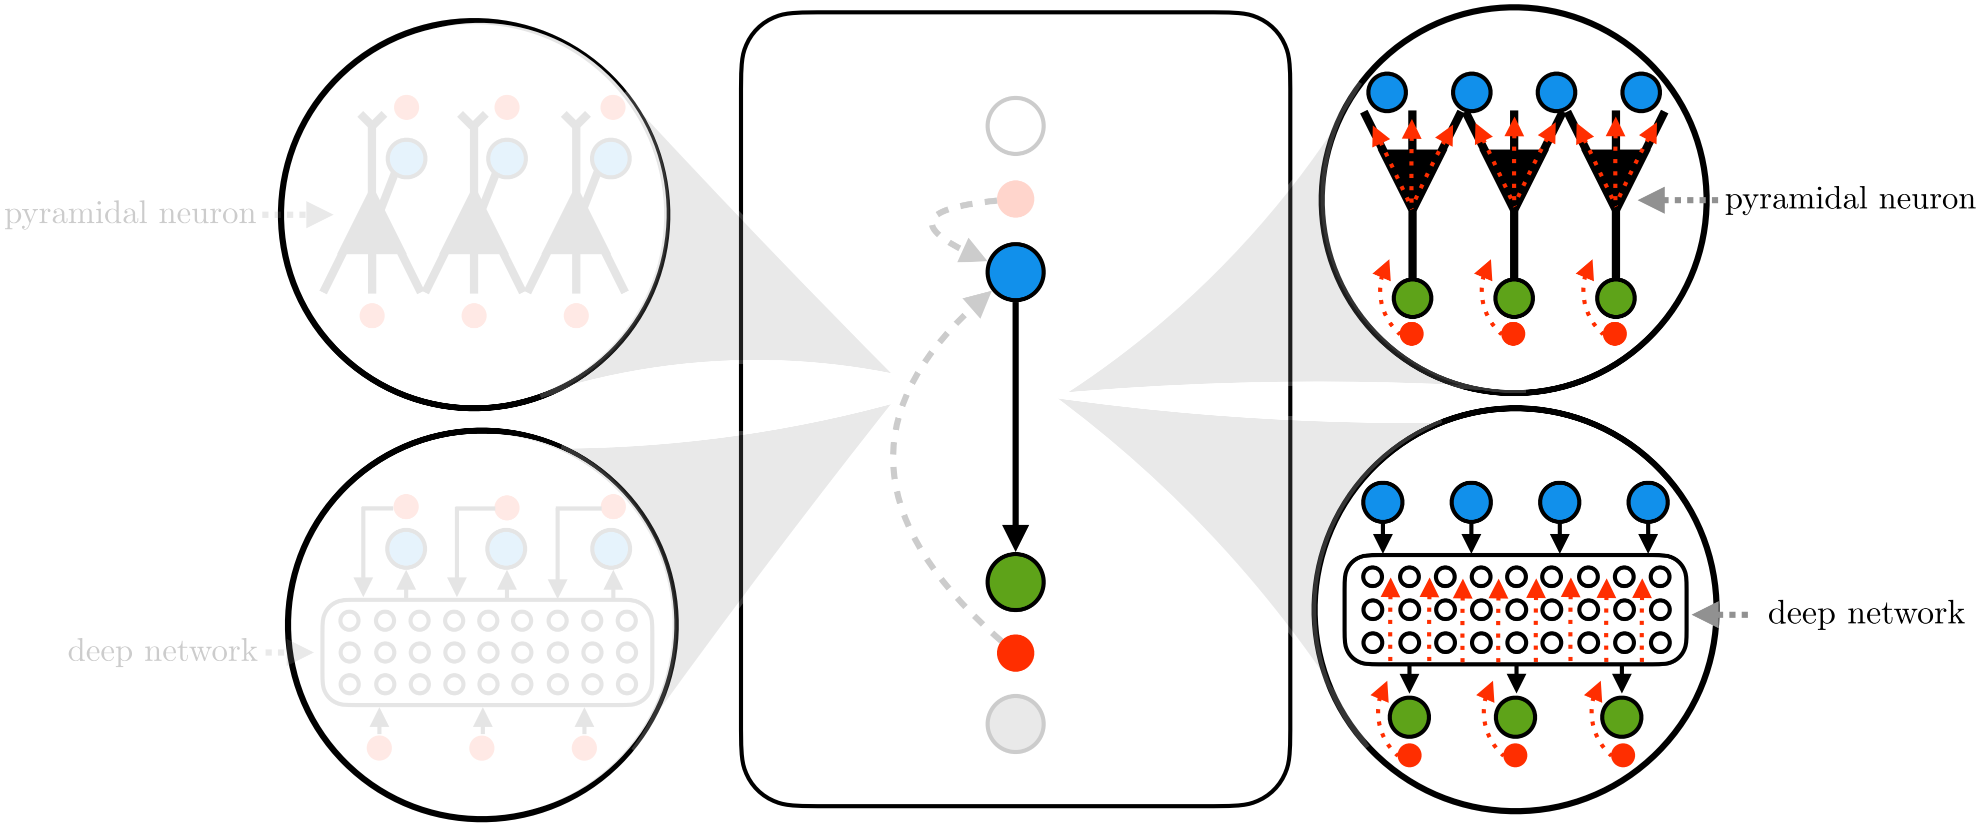

errors provide a ***local*** training signal

e.g., target propagation (Bengio, 2014)

Marino, 2019

LATERAL INHIBITION & NORMALIZING FLOWS

lateral **interneurons** add or remove spatiotemporal correlations between neurons

generalized by **normalizing flows** (from machine learning)

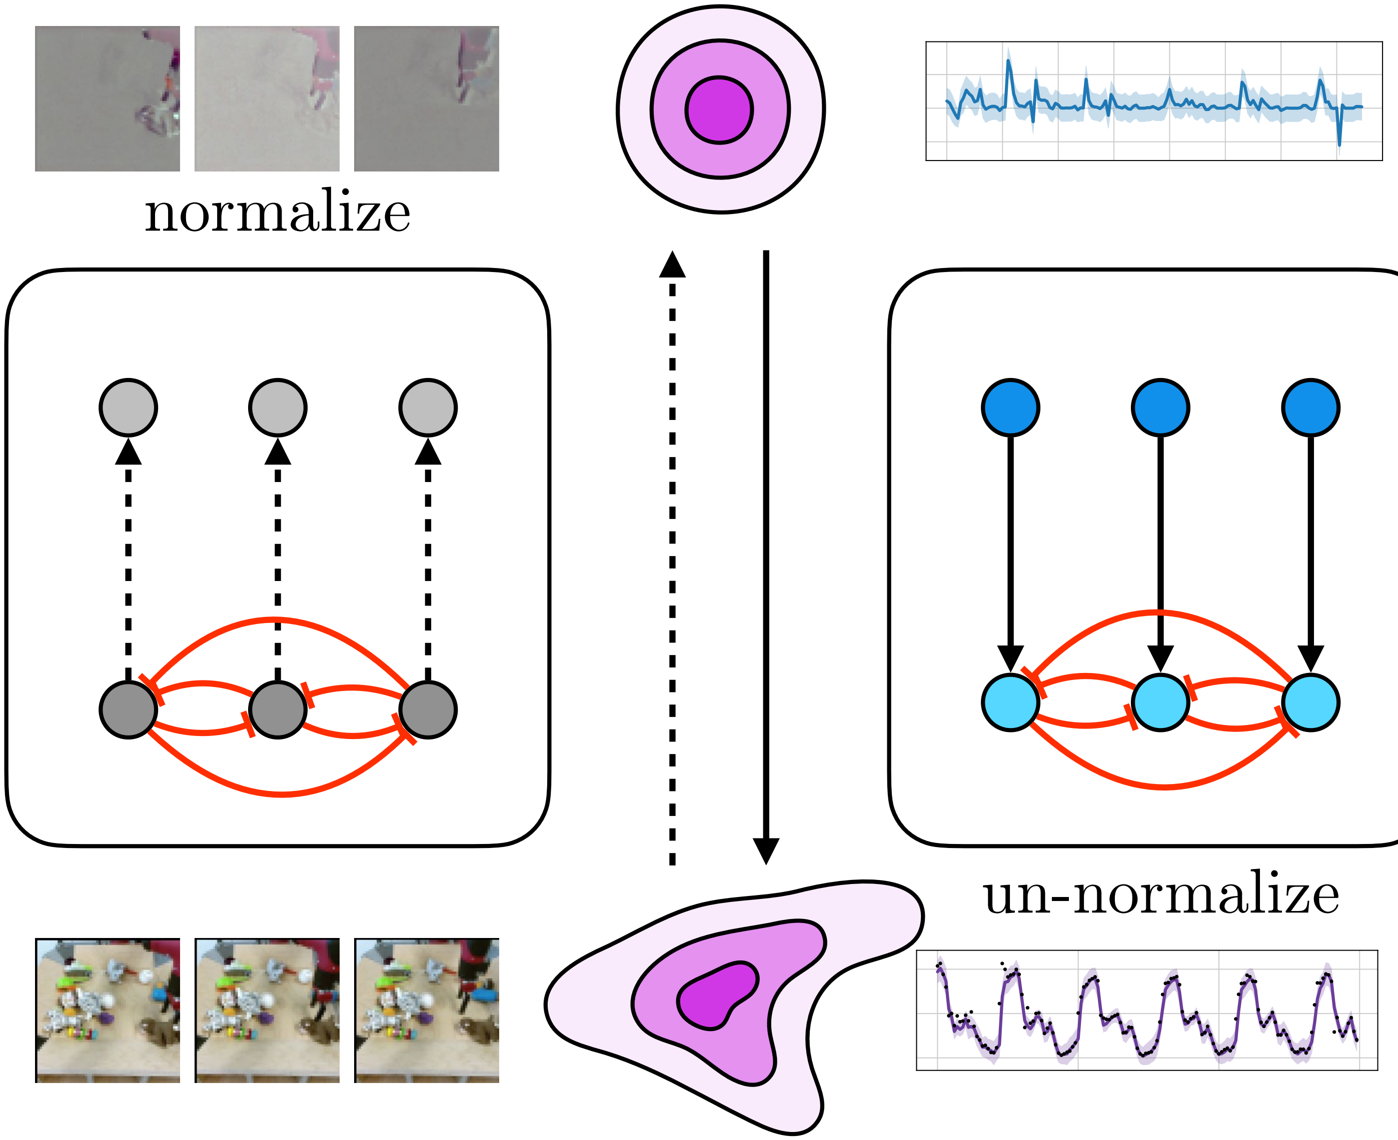

Marino, 2019

VISUAL PATHWAY

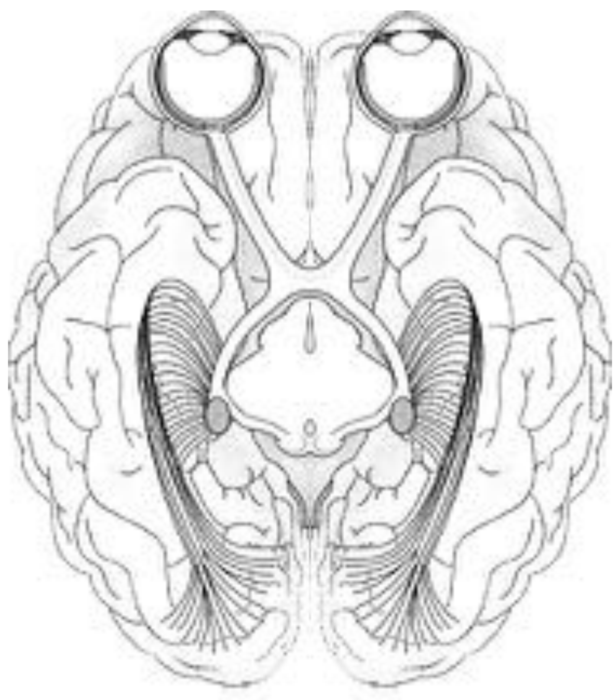

VISUAL PATHWAY

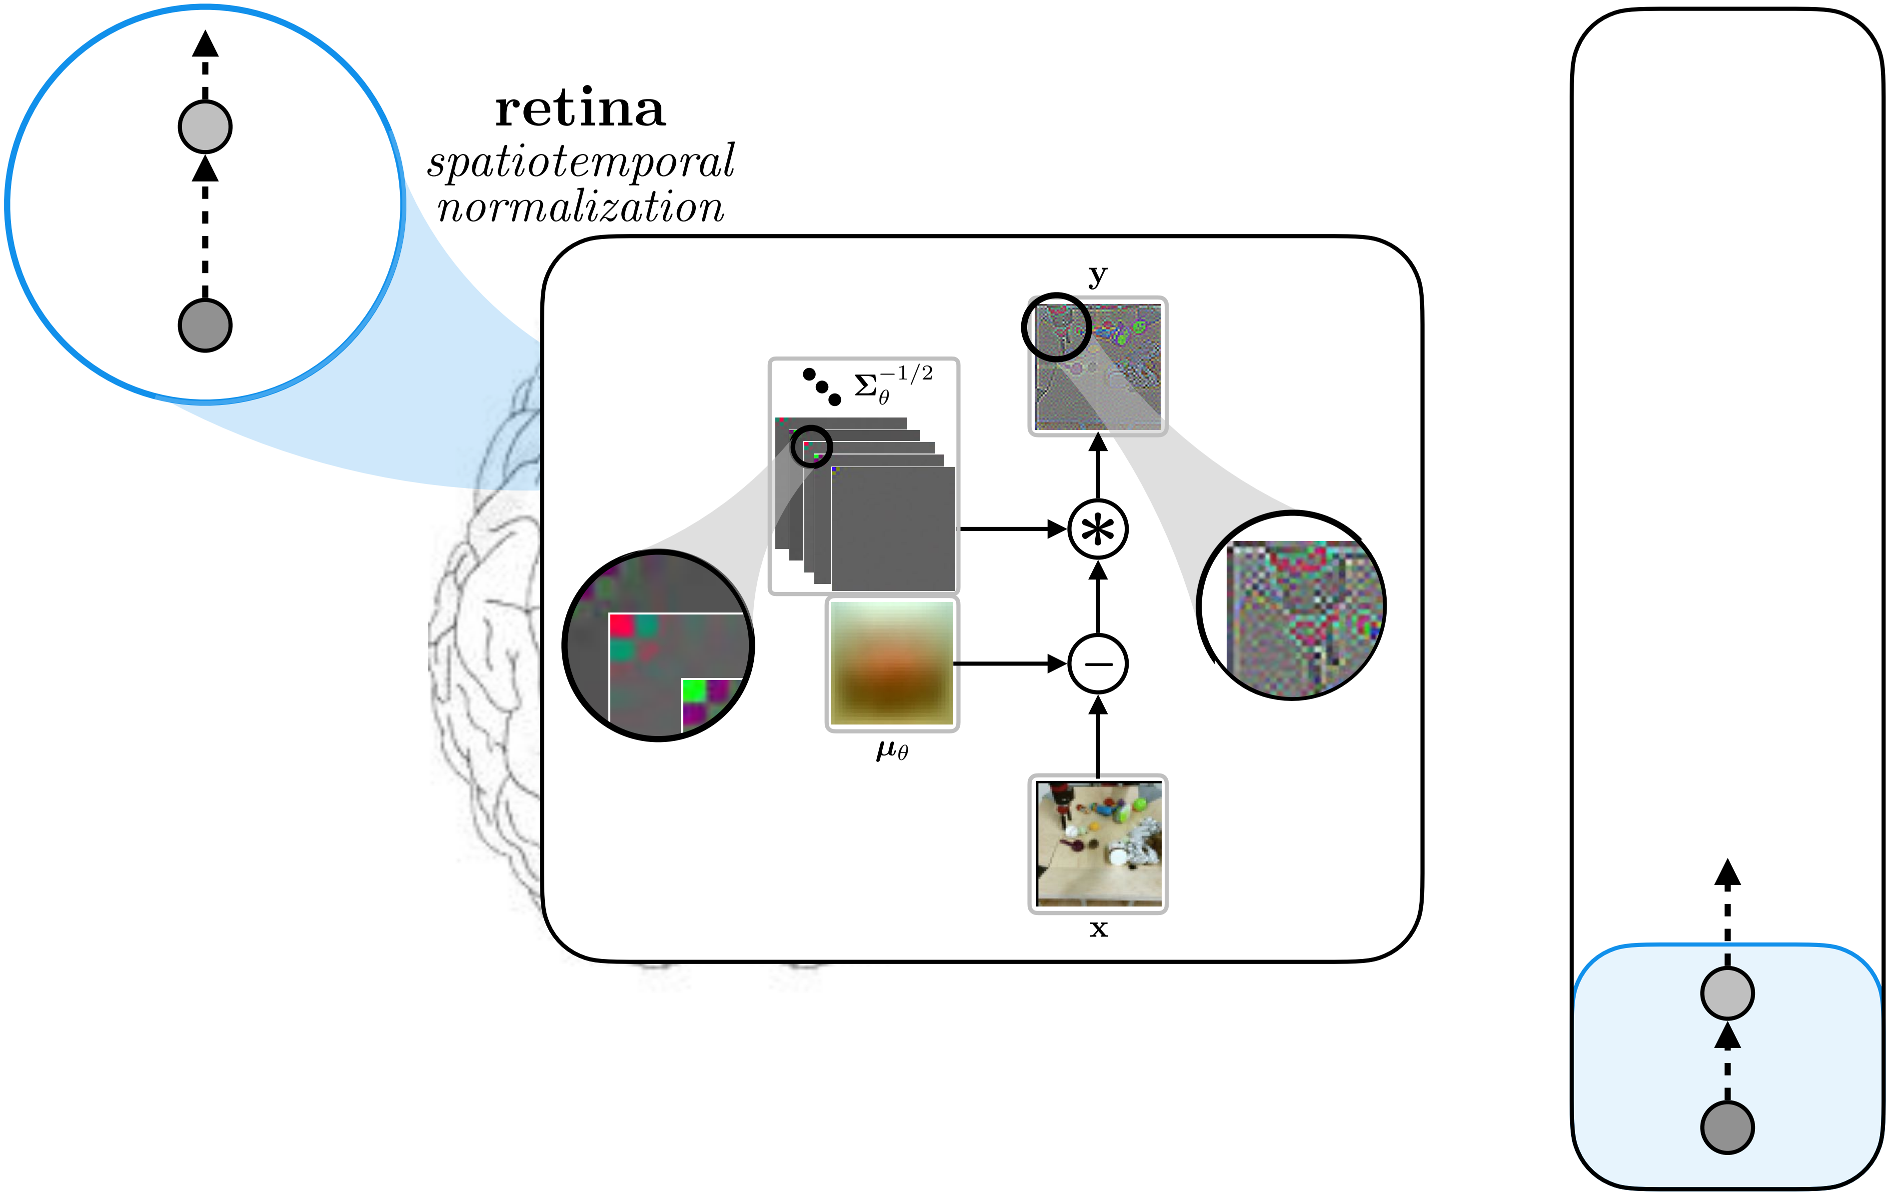

Marino, 2019

VISUAL PATHWAY

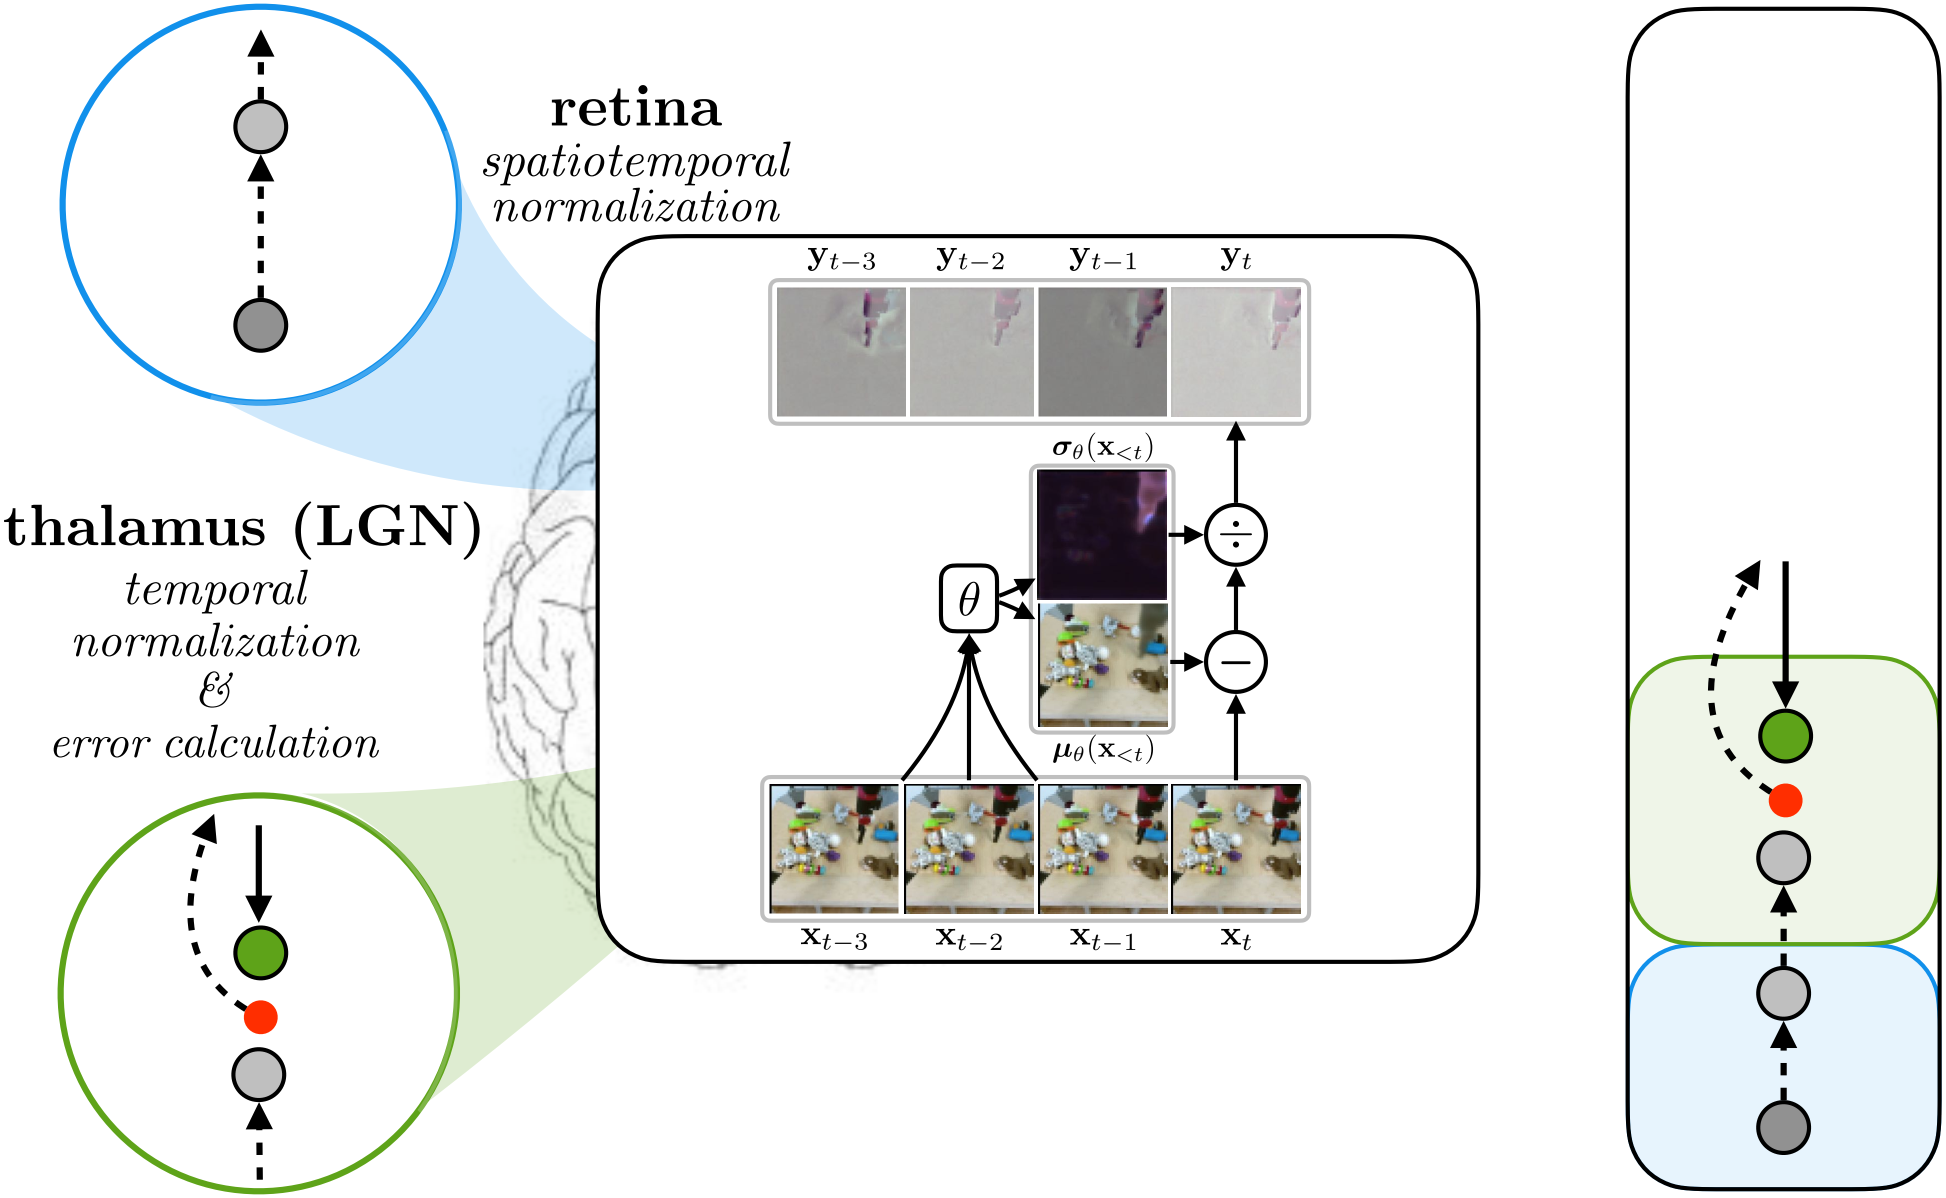

Marino, 2019

VISUAL PATHWAY

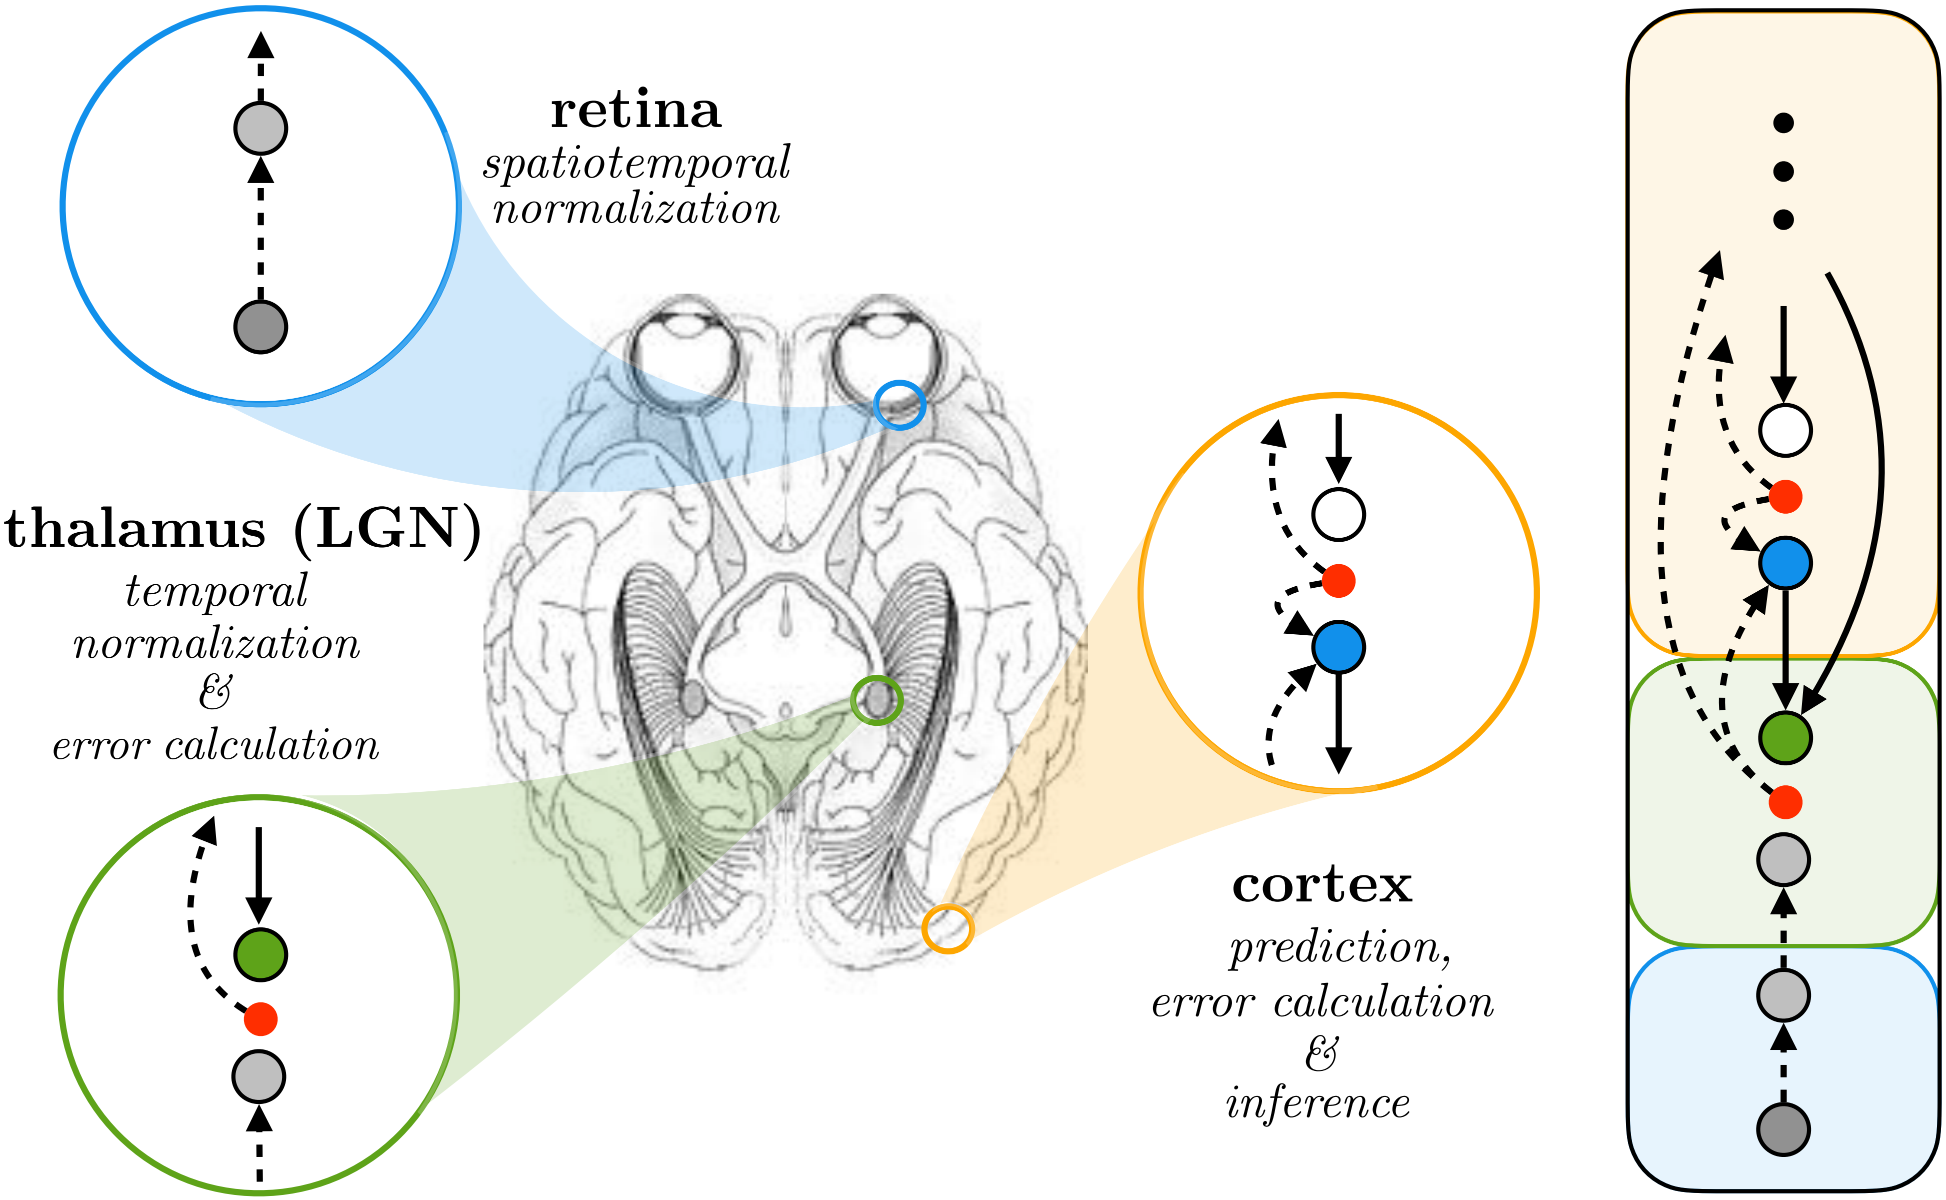

Marino, 2019

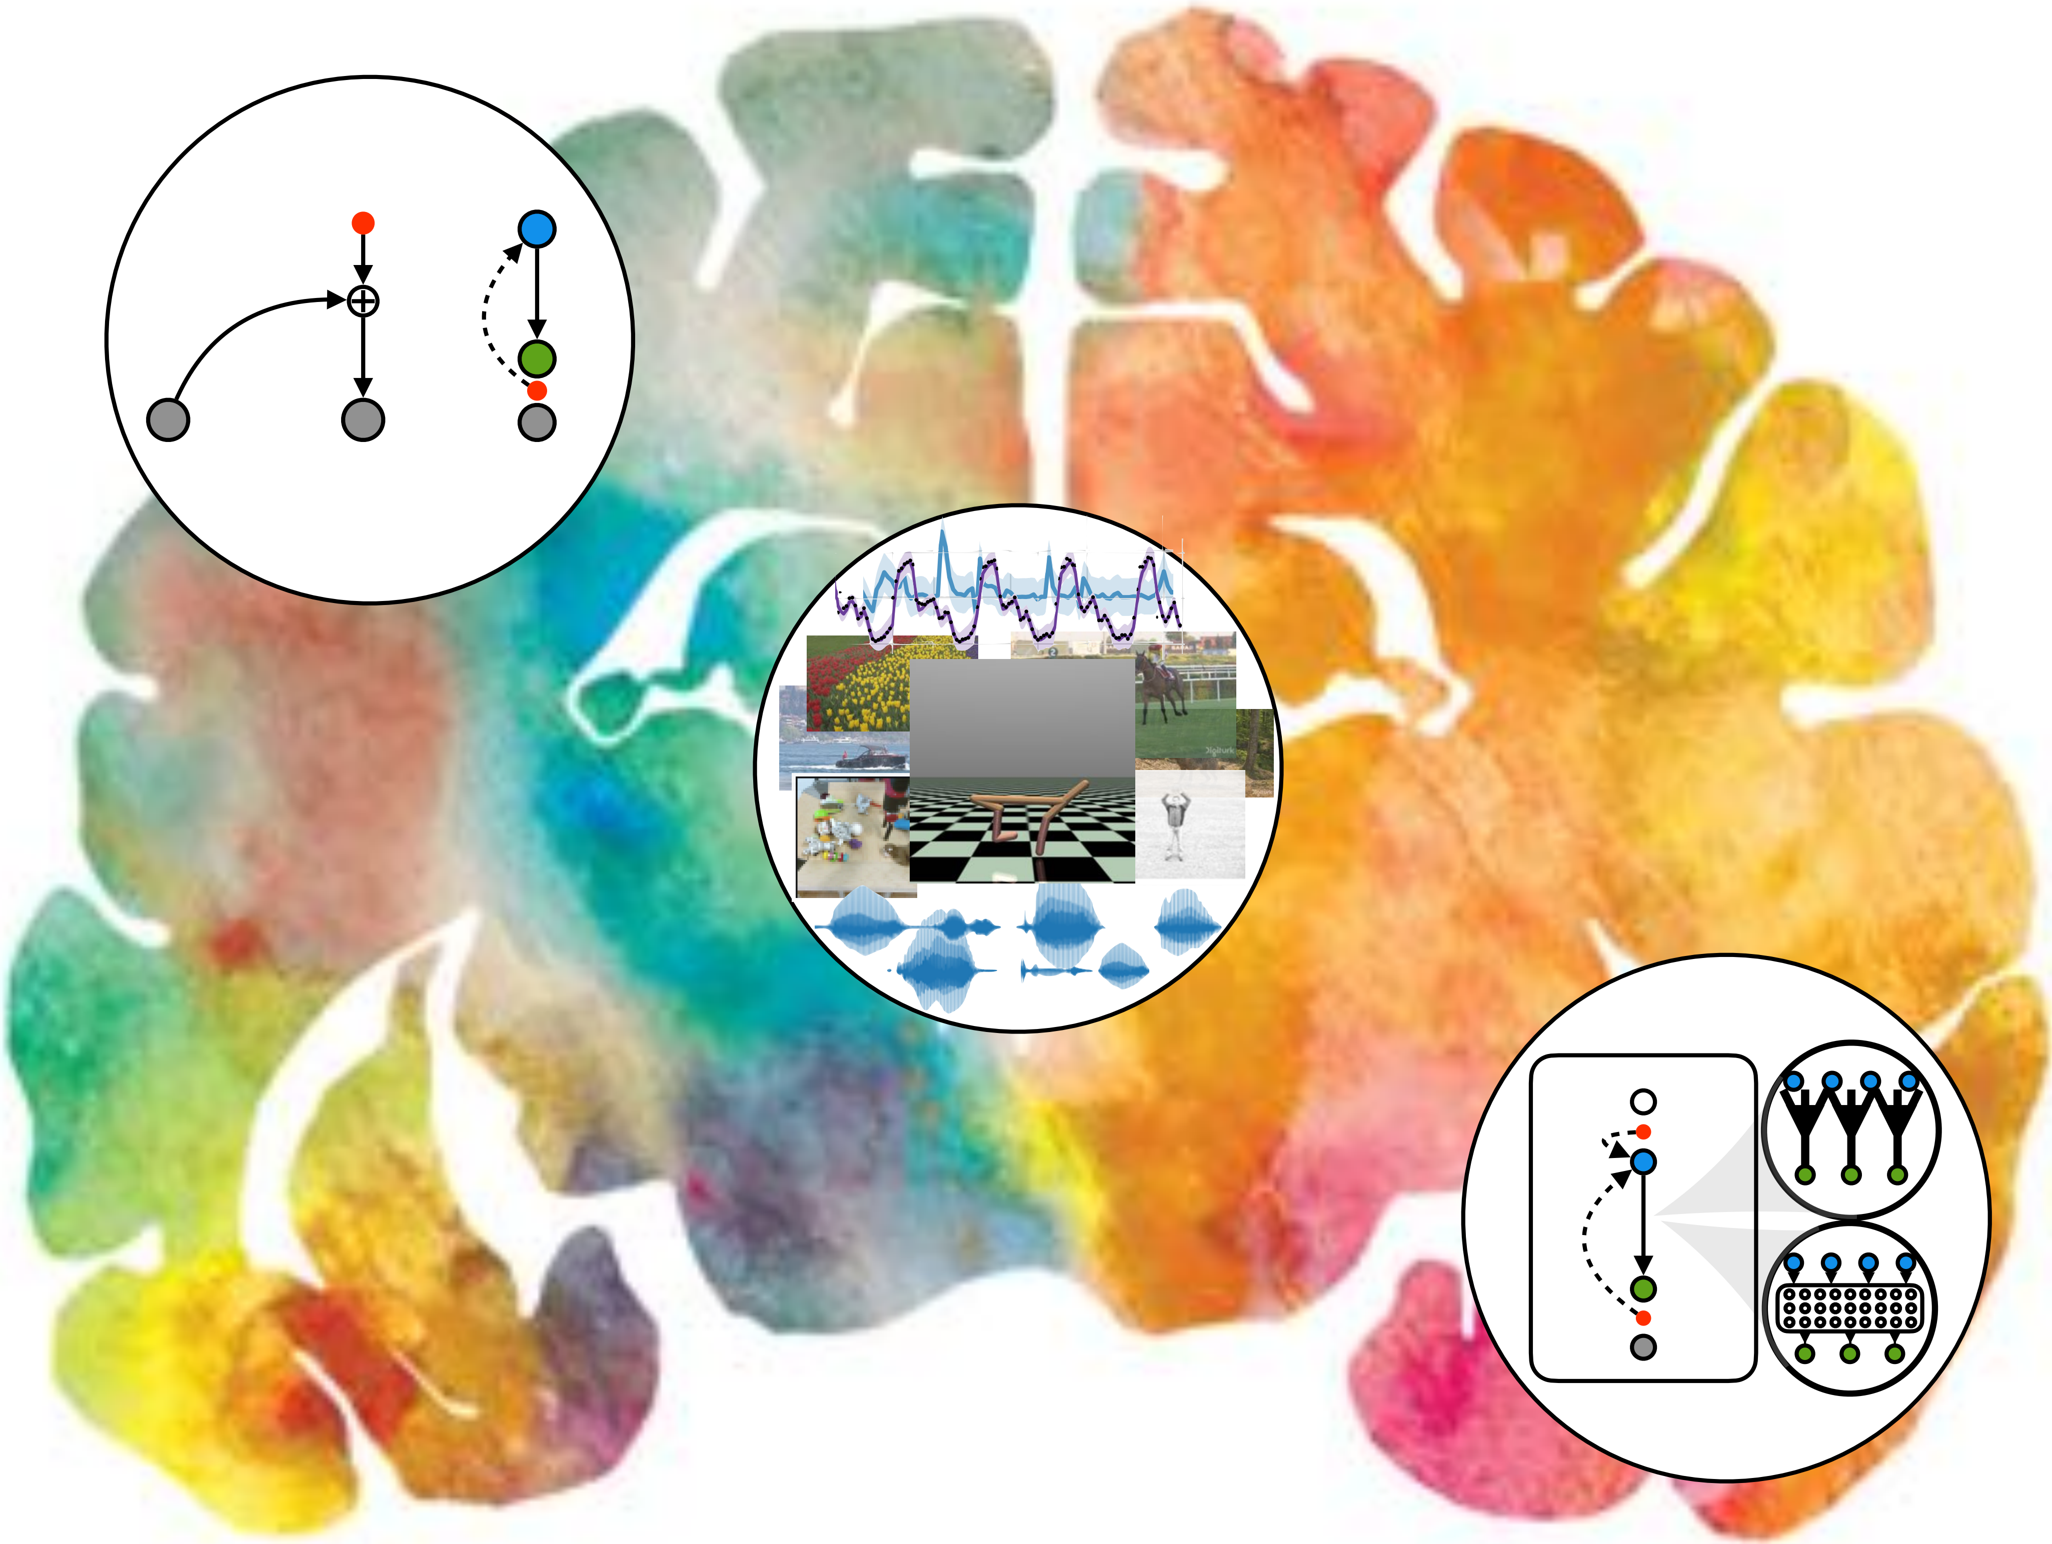

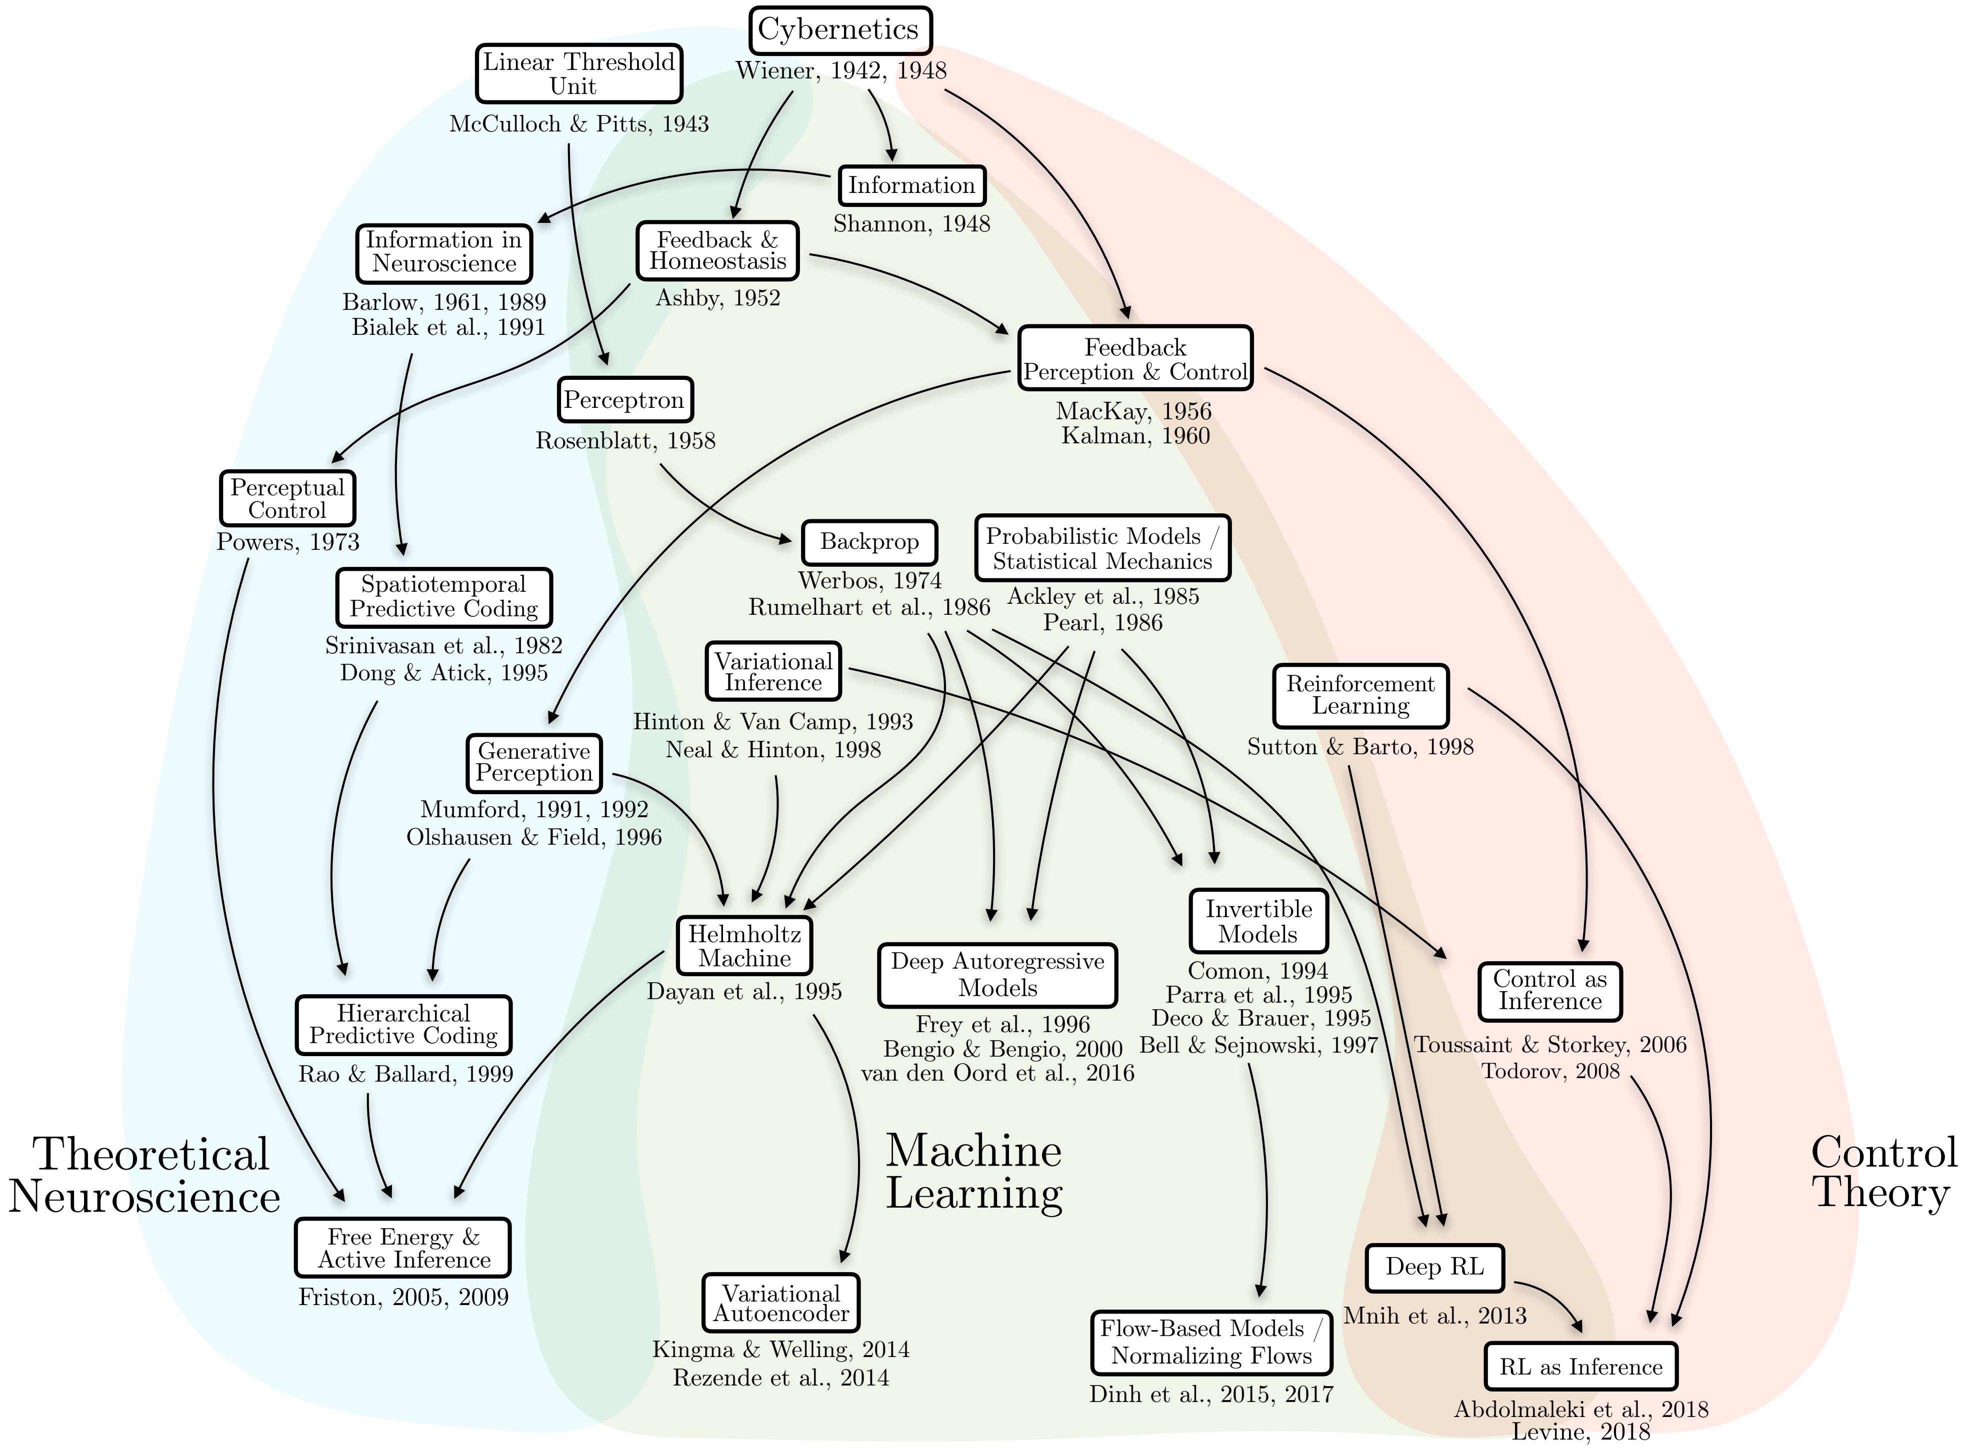

Cybernetics
Wiener, 1942, 1948
Linear Threshold Unit
McCulloch & Pitts, 1943
Information
Shannon, 1948
Information in Neuroscience
Barlow, 1961, 1989
Bialek et al., 1991
Feedback & Homeostasis
Ashby, 1952
Feedback Perception & Control
MacKay, 1956
Kalman, 1960
Perceptron
Rosenblatt, 1958
Perceptual Control
Powers, 1973
Backprop
Werbos, 1974
Rumelhart et al., 1986
Probabilistic Models / Statistical Mechanics
Ackley et al., 1985
Pearl, 1986
Spatiotemporal Predictive Coding
Srinivasan et al., 1982
Dong & Atick, 1995
Variational Inference
Hinton & Van Camp, 1993
Neal & Hinton, 1998
Reinforcement Learning
Sutton & Barto, 1998
Generative Perception
Mumford, 1991, 1992
Olshausen & Field, 1996
Helmholtz Machine
Dayan et al., 1995
Deep Autoregressive Models
Frey et al., 1996
Bengio & Bengio, 2000
van den Oord et al., 2016
Invertible Models
Comon, 1994
Parra et al., 1995
Deco & Brauer, 1995
Bell & Sejnowski, 1997
Control as Inference
Toussaint & Storkey, 2006
Todorov, 2008
Hierarchical Predictive Coding
Rao & Ballard, 1999
Theoretical Neuroscience
Machine Learning
Control Theory
Free Energy & Active Inference
Friston, 2005, 2009
Variational Autoencoder
Kingma & Welling, 2014
Rezende et al., 2014
Flow-Based Models / Normalizing Flows
Dinh et al., 2015, 2017
Deep RL
Mnih et al., 2013
RL as Inference
Abdolmaleki et al., 2018
Levine, 2018

Collaborators


Yisong Yue


Stephan Mandt


Jiawei He


Alexandre Piché


Alessandro Ialongo


Alex Guerra


Yibo Yang


Milan Cvitkovic


Yang Yang


Lei Chen

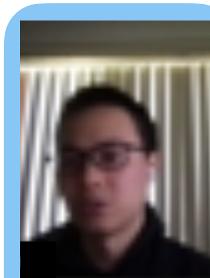
Ruihan Yang

- Caltech
- UC Irvine
- Simon Fraser Univ.
- MILA
- Cambridge/MPII
- Qualcomm

thank you