The least-control principle for learning at equilibrium

João Sacramento

with Alexander Meulemans*, Nicolas Zucchet*, Seijin Kobayashi*, Johannes von Oswald

IST Seminar in Mathematics, Physics and Machine Learning

ETH zürich

Neuroscience

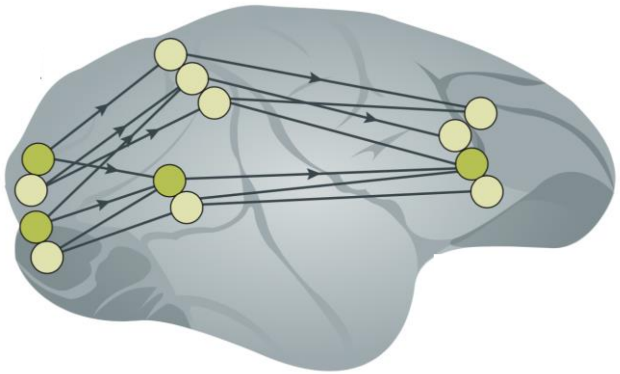

models, theories,
testable hypotheses

understand how the brain learns

novel ideas, insight

Machine learning

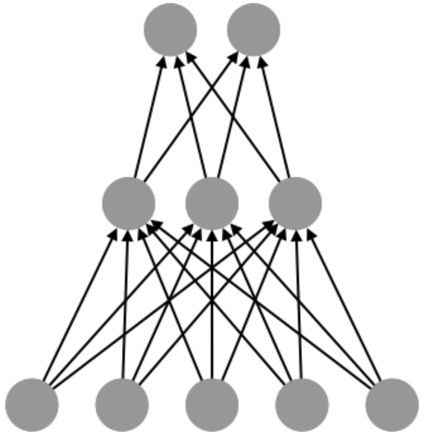

Exciting methodological advances
record thousands of neurons simultaneously
→ bridge the cellular and network level

Neural networks that perceive the world and act back upon it

Zador et al., arXiv 2022
Saxe et al., *Nat. Rev. Neurosci.* 2021
Richards et al., *Nat. Neurosci.* 2019
Hassabis et al., *Neuron* 2017

Outline

- Credit assignment in artificial and biological neural networks
- Assigning credit by backpropagation-of-error
- Assigning credit by controlling a neural dynamics
- Theory & experimental results
- Connections to probabilistic modeling and free-energy minimization

Neural networks

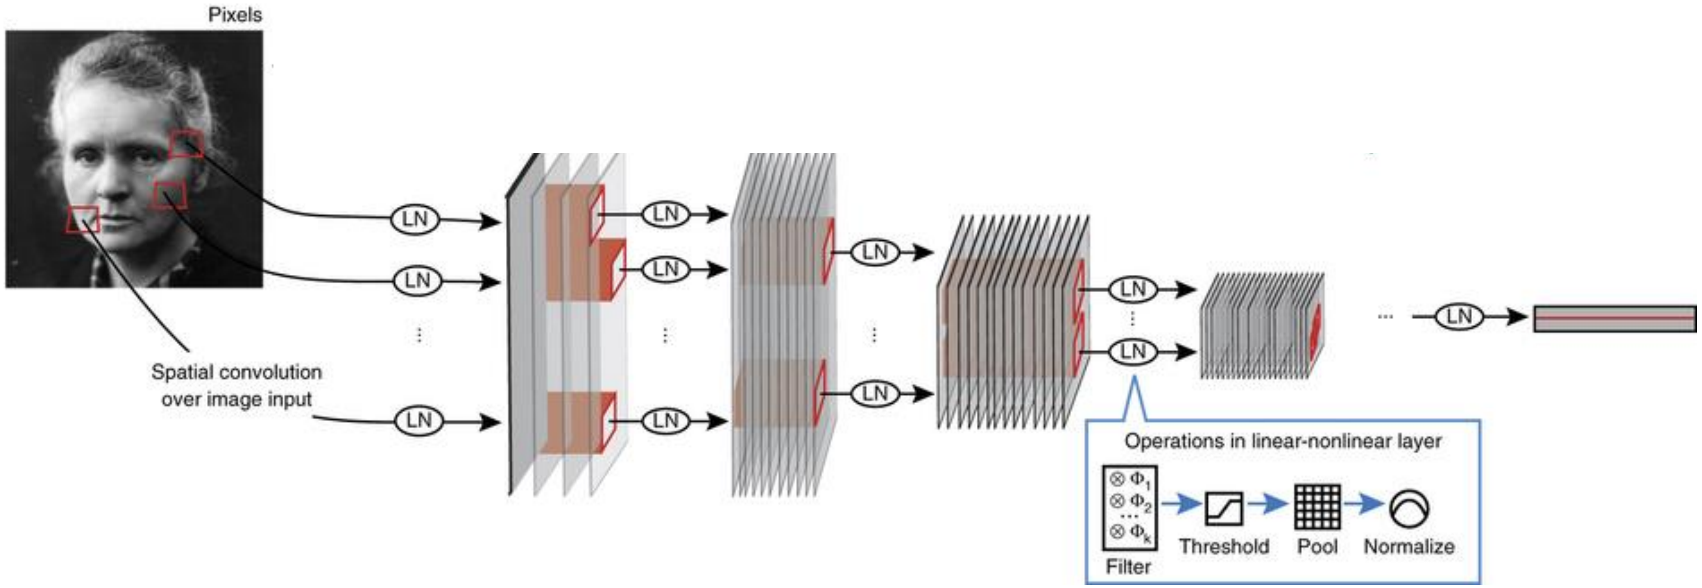

- Information processing done by in parallel by a network of basic computational units
- Output of each unit determined by its parameters
- **Parameters determined from data (the *learning* process; *plasticity*)**
- System is not fully hand-engineered, complex behavior emerges from interaction with data

Fig. adapted from Yamins & DiCarlo, *Nat. Neurosci.* 2016

The credit assignment problem

very large number of parameters

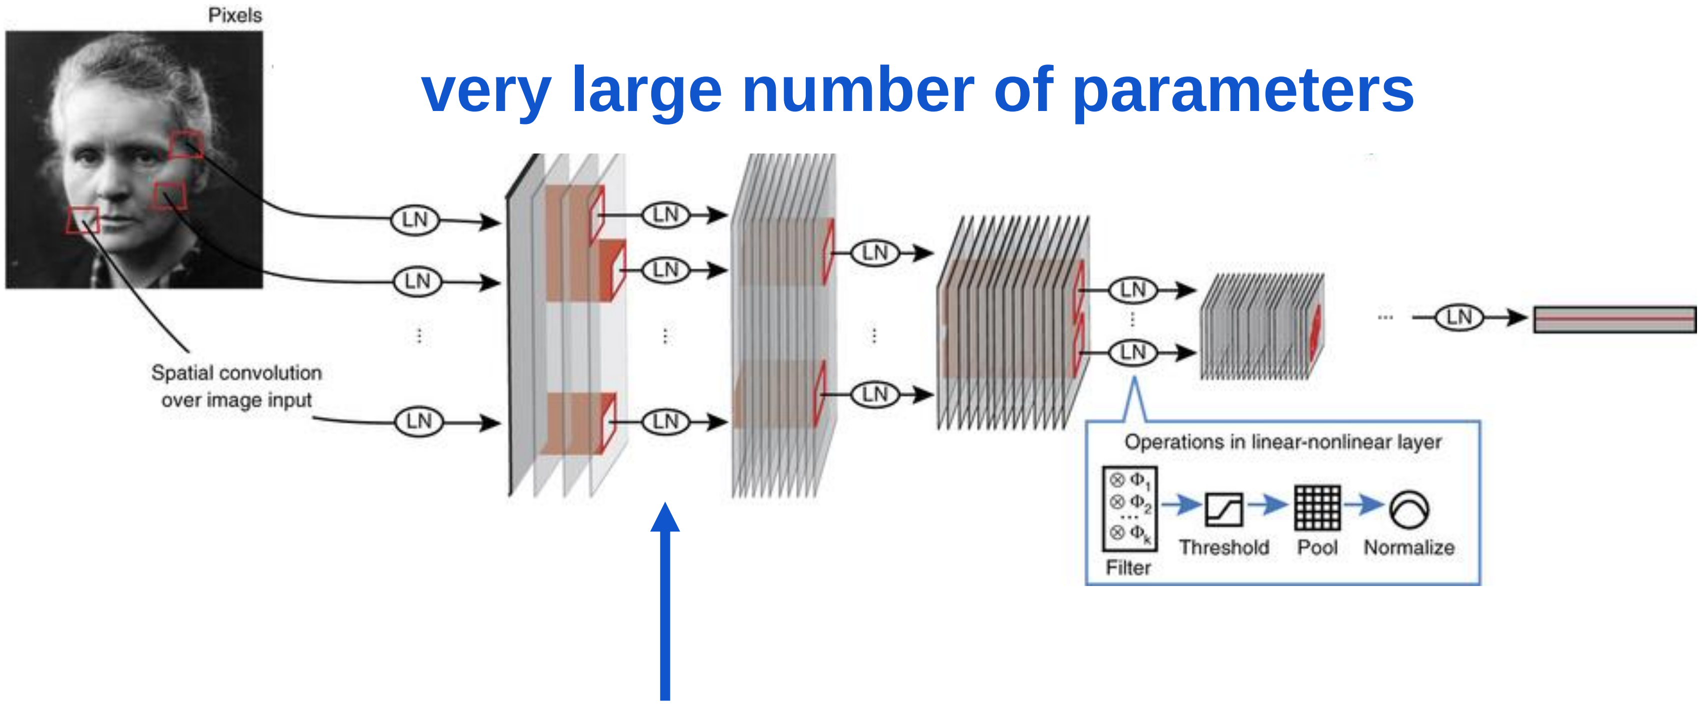

A change in early layers of processing affects the output of the network in a complicated way

Fig. adapted from Yamins & DiCarlo, *Nat. Neurosci.* 2016

Backpropagation

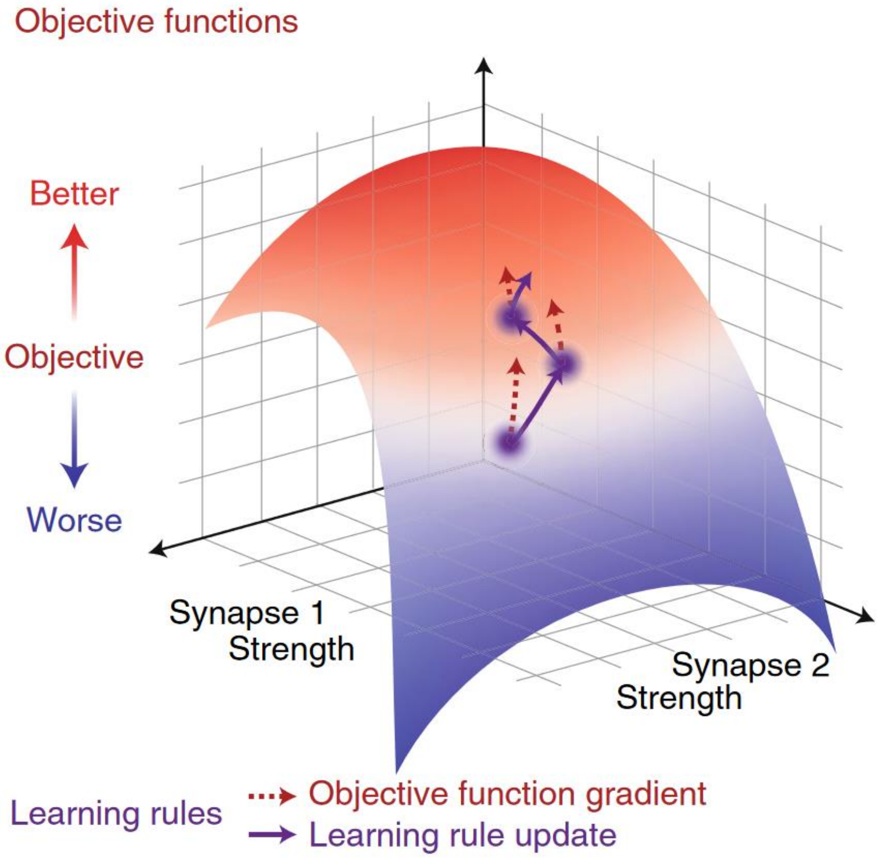

Learning by gradient ascent

Adjust synapses θ to maximize objective function *F* by taking steps along its gradient:

$$\theta_{t+1} = \theta_t + \eta \, \underline{\nabla F(\theta_t)}$$

almost always: computed by backpropagation-of-error
Bryson, *Harvard Symposium* 1961
Rumelhart et al., *Nature* 1986

- Scalable approach
- Unreasonably effective?
- Underlies great advances in supervised, unsupervised and reinforcement learning

Richards et al., *Nat. Neurosci.* 2019

Credit assignment in the brain

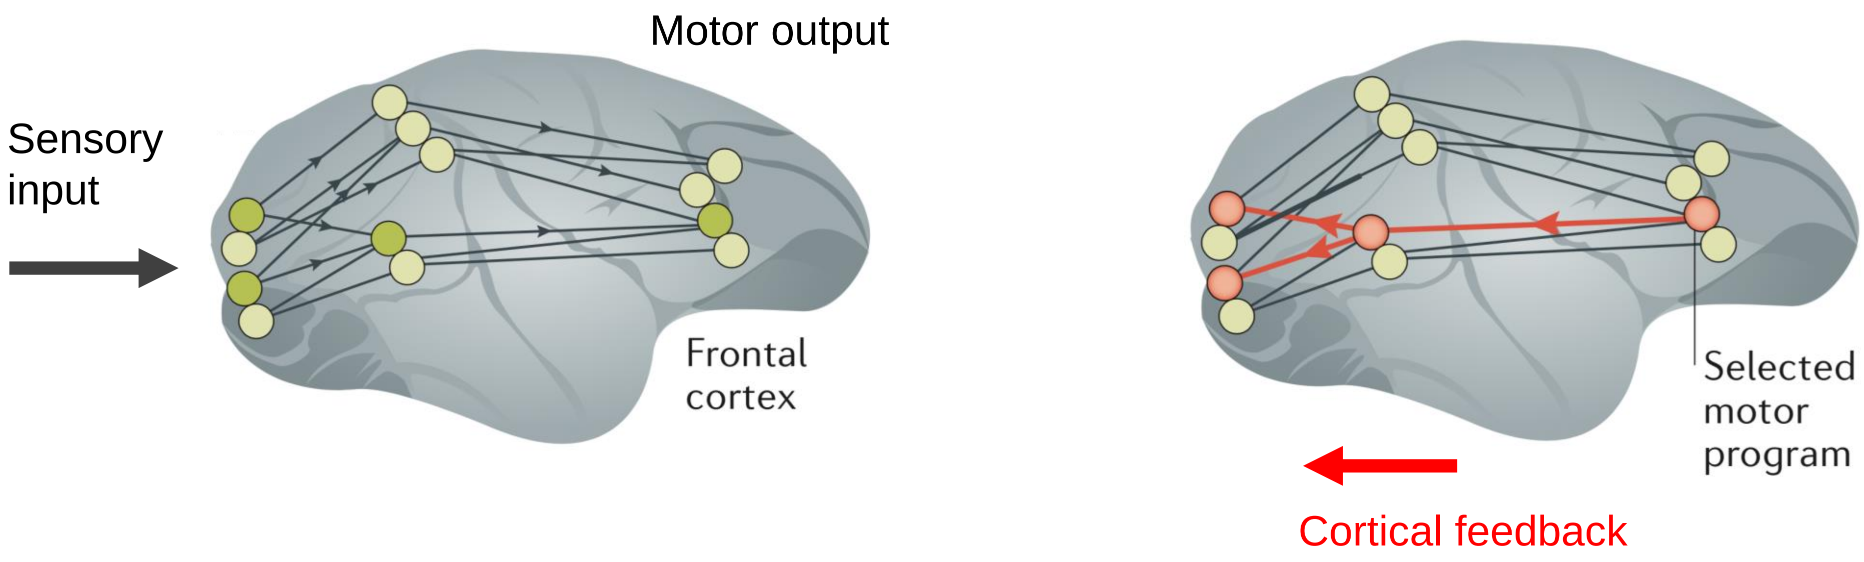

- The human brain has ~10^{15} synaptic connections, majority is plastic
- How does synaptic change eventually lead to improved behavior?
- Determining credit/blame with a parallel, distributed processing algorithm
- Does cortical feedback play a role in this process?

Fig. adapted from Roelfsema & Holtmaat, *Nat. Rev. Neurosci.* 2018

Backpropagation in the brain?

- **Weight transport problem (non-local learning rule)**: the change of a given synapse depends on the precise values of all others downstream of it

- Clocked **two-phase** algorithm: prediction followed by error computation

- Backpropagated errors instruct synaptic change, but **do not change activity**

- Restricted to acyclic computational graphs (but see Almeida-Pineda algorithm)

- Assumption of **differentiability**
  - But most neurons communicate with “spikes”, essentially binary events
  - Highly nonlinear dynamics can lead to vanishing/exploding derivatives
  - Model average (“firing rate”) activity? Use surrogate functions?

Grossberg, *Cogn. Sci.* 1987
Crick, *Nature* 1989
Lillicrap et al., *Nat. Rev. Neurosci.* 2020

Backpropagation in the brain?

- **Weight transport problem (non-local learning rule)**: the change of a given synapse depends on the precise values of all others downstream of it

- Clocked **two-phase** algorithm: prediction followed by error computation

- Backpropagated errors instruct synaptic change, but **do not change activity**

- Restricted to acyclic computational graphs (but see Almeida-Pineda algorithm)

- Assumption of **differentiability**
  - But most neurons communicate with “spikes”, essentially binary events
  - Highly nonlinear dynamics can lead to vanishing/exploding derivatives
  - Model average (“firing rate”) activity? Use surrogate functions?

Grossberg, *Cogn. Sci.* 1987
Crick, *Nature* 1989
Lillicrap et al., *Nat. Rev. Neurosci.* 2020

Outline

- Credit assignment in artificial and biological neural networks
- Assigning credit by backpropagation-of-error
- Assigning credit by controlling a neural dynamics
- Theory & experimental results
- Connections to probabilistic modeling and free-energy minimization

Desiderata for our theory

1. Gradient-based learning
2. Use feedback to embed credit assignment information in the neural activity
3. Local, activity-dependent learning rules
4. Single-phase learning
5. Ability to learn a recurrent dynamics (go beyond feedforward networks)

Going beyond feedforward networks

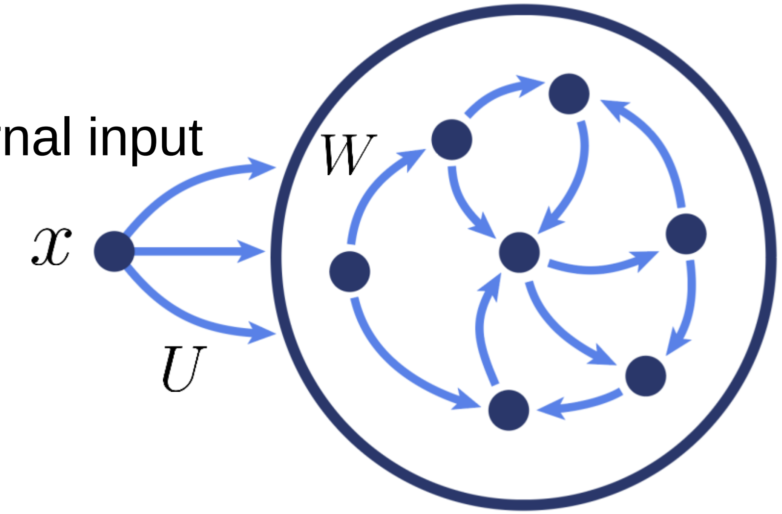

equilibrium state ϕ^*

$$\dot{\phi}(t) = -\phi(t) + W\sigma(\phi(t)) + Ux$$
$$\phi^* = W\sigma(\phi^*) + Ux$$
$$\theta = \{W, U\}$$

Key simplifying assumption:
dynamics reaches a fixed point

$$\dot{\phi}(t) = f_\theta(\phi(t), x)$$
$$0 = f_\theta(\phi^*, x)$$

In this talk, we will focus on
neural networks of firing rate neurons

Learning at equilibrium

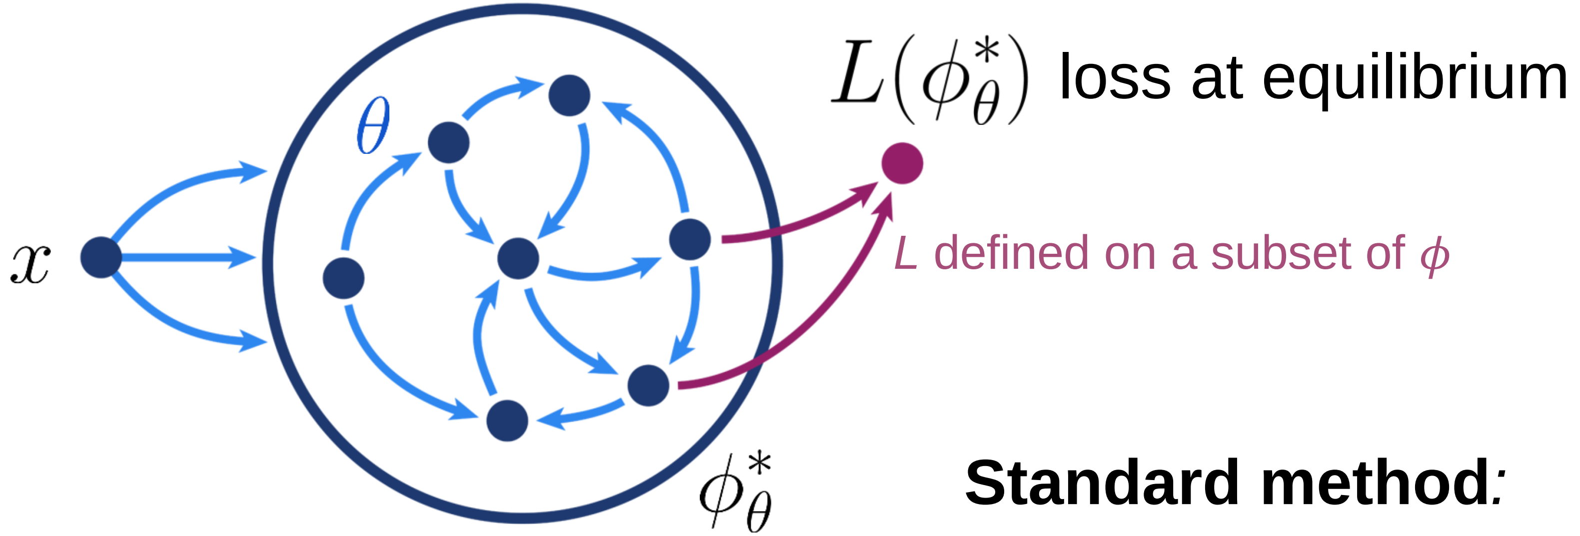

$$\phi^* = W\sigma(\phi^*) + Ux$$

$$L(\theta) = \|y^{\text{target}} - \underbrace{S_y \phi^*_\theta}_{\text{selects output neurons}}\|^2$$

Standard method:
Compute $\nabla_\theta L(\phi^*_\theta)$ using the Almeida-Pineda (recurrent backpropagation) algorithm

Critique of conventional backpropagation applies

Almeida, *Neural Computers* 1989
Pineda, *Neural Computation* 1989
Bai et al., *NeurIPS* 2019

Learning by minimizing control

Controlling neural activity

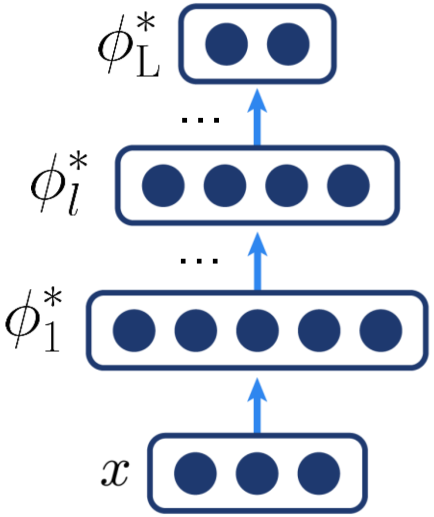

Feedforward neural network

$$\phi^*_l = W_l \sigma(\phi^*_{l-1})$$

Continuous-time feedforward neural network

$$\dot{\phi}_l(t) = -\phi_l(t) + W_l \sigma(\phi_{l-1}(t))$$

Sussillo & Abbott, *Neuron* 2009
Gilra & Gerstner, *eLife* 2017
Alemi et al., *AAAI* 2018
Podlaski & Machens, *NeurIPS* 2020
Meulemans et al., *NeurIPS* 2021
Meulemans et al., *ICML* 2022

Controlling neural activity

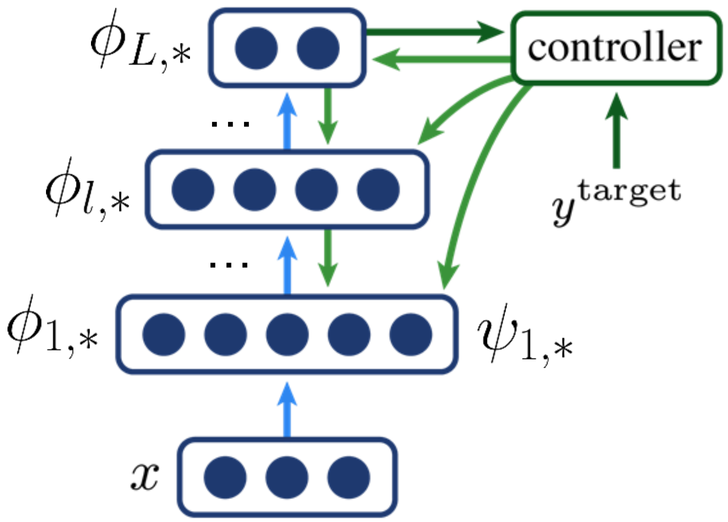

Controlled feedforward neural network

$$\dot{\phi}_l(t) = -\phi_l(t) + W_l \sigma(\phi_{l-1}(t)) + \psi_l(t)$$

network becomes recurrent due to feedback control

Controlled equilibrium

$$\phi_{l,*} = W_l \sigma(\phi_{l-1,*}) + \psi_{l,*}$$

Degeneracy: for our neural networks there are typically infinitely many configurations of the activity for which the output is at the desired target

Sussillo & Abbott, *Neuron* 2009
Gilra & Gerstner, *eLife* 2017
Alemi et al., *AAAI* 2018
Podlaski & Machens, *NeurIPS* 2020
Meulemans et al., *NeurIPS* 2021
Meulemans et al., *ICML* 2022

Least-control states

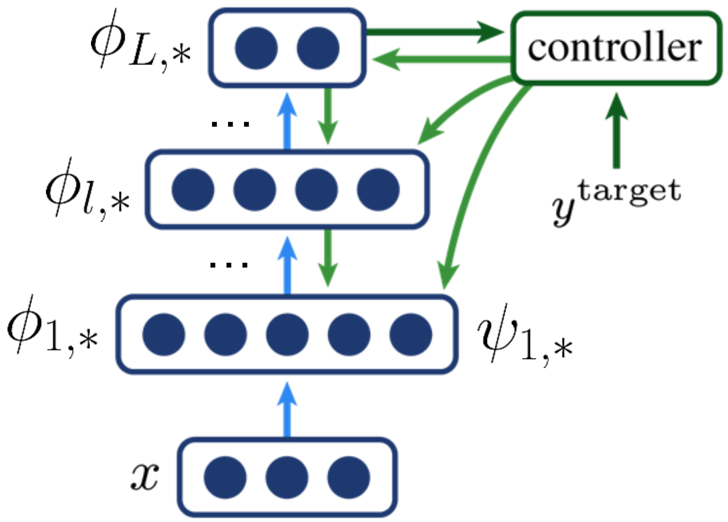

Controlled feedforward neural network

$$\dot{\phi}_l(t) = -\phi_l(t) + W_l\sigma(\phi_{l-1}(t)) + \psi_l(t)$$

network becomes recurrent due to feedback control

Controlled equilibrium

$$\phi_{l,*} = W_l\sigma(\phi_{l-1,*}) + \psi_{l,*}$$

Least-control states: reach y^{target} with the minimal amount of control

Learning: decrease the amount of control needed to reach target

Learning by minimizing control

Least-control states: reach y^{target} with the minimal amount of control

Learning: decrease the amount of control needed to reach target

Formalizing the least-control principle

We want to **minimize the amount of control** while $\min \|\psi\|^2$

- **the controlled dynamics are at equilibrium** $0 = f_\theta(\phi, x) + \psi$
- **the loss is minimized at this equilibrium** $0 = \nabla_\phi L(\phi)$

$$\boxed{\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0, \quad \nabla_\phi L(\phi) = 0}$$

the least-control principle

Solving least-control problems

$$\boxed{\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0, \quad \nabla_\phi L(\phi) = 0}$$

Solving least-control problems

$$\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0, \quad \nabla_\phi L(\phi) = 0$$

Step 1: find an optimal control

The minimal amount of control possible for the current parameter setting

Solving least-control problems

$$\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0, \quad \nabla_\phi L(\phi) = 0$$

Step 1: find an optimal control

The minimal amount of control possible for the current parameter setting

Step 2: parameter update

Take a gradient step w.r.t. θ

Solving least-control problems

$$\boxed{\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad -\phi + W\sigma(\phi) + Ux + \psi = 0, \quad \nabla_\phi L(\phi) = 0}$$

neural network model

Step 1: find an optimal control

The minimal amount of control possible for the current parameter setting

Step 2: parameter update

Take a gradient step w.r.t. θ

$$\Delta W = \psi_* \sigma(\phi_*)^\top$$

local (in space and time) "Hebbian" learning rule

Learning rule for the general case

least-control state, optimal control	$(\phi_\star, \psi_\star) = \arg\min_{\phi,\psi} \left\{\frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0,\ \nabla_\phi L(\phi) = 0\right\}$
learning objective	$\mathcal{H}(\theta) = \frac{1}{2}\|\psi_\star\|^2$
objective gradient	$\nabla_\theta \mathcal{H}(\theta) = -\frac{\partial f}{\partial \theta}(\phi_\star, x, \theta)^\top \psi_\star$

(needs some technical assumptions
to ensure existence of implicit functions)

- **Single-phase, activity-dependent, gradient-based learning**
- How does optimizing for least-control relate to the original learning problem?
- How do we find the optimal control?

Linking learning objectives

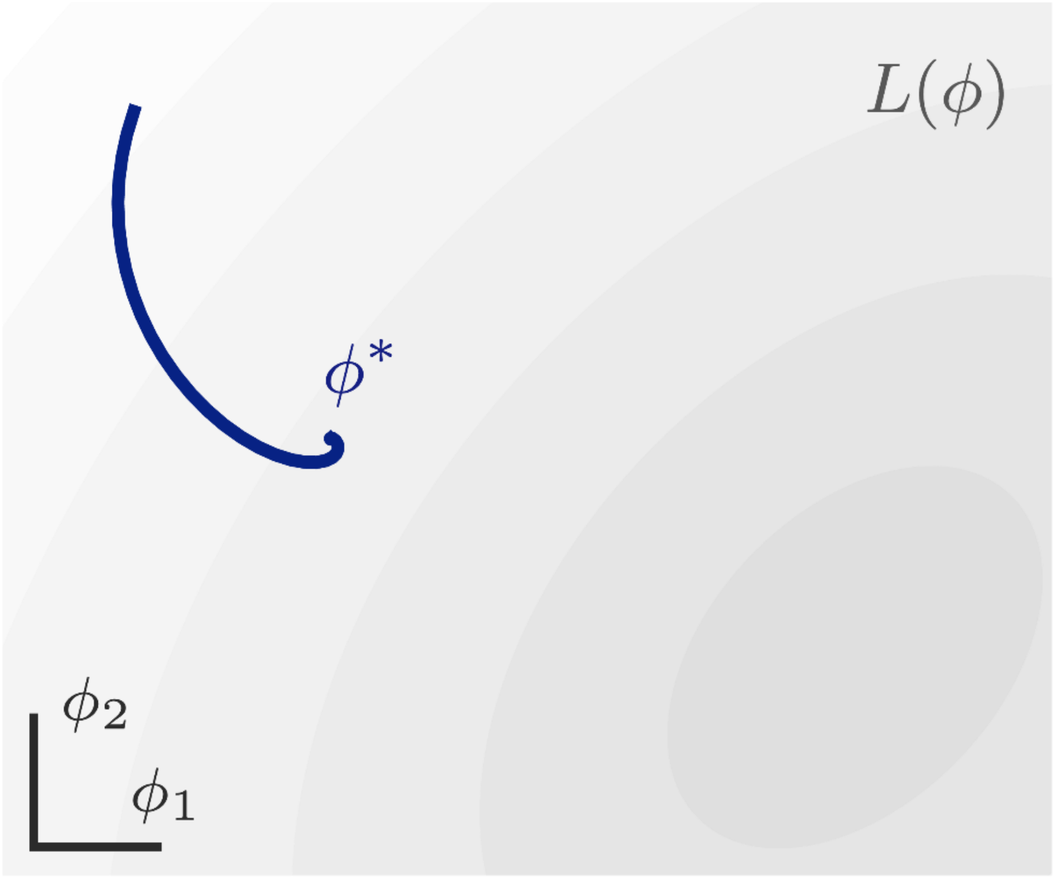

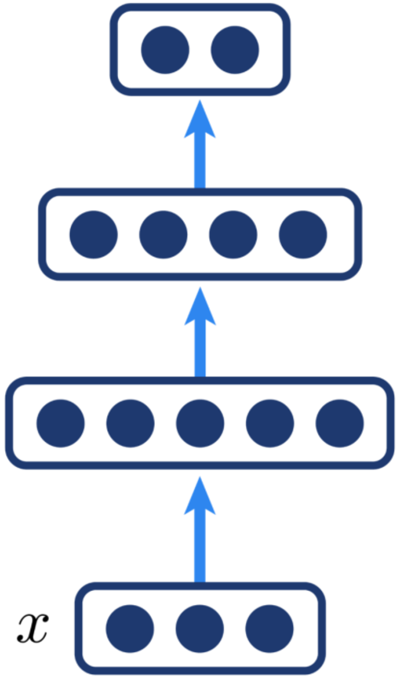

Linking learning objectives

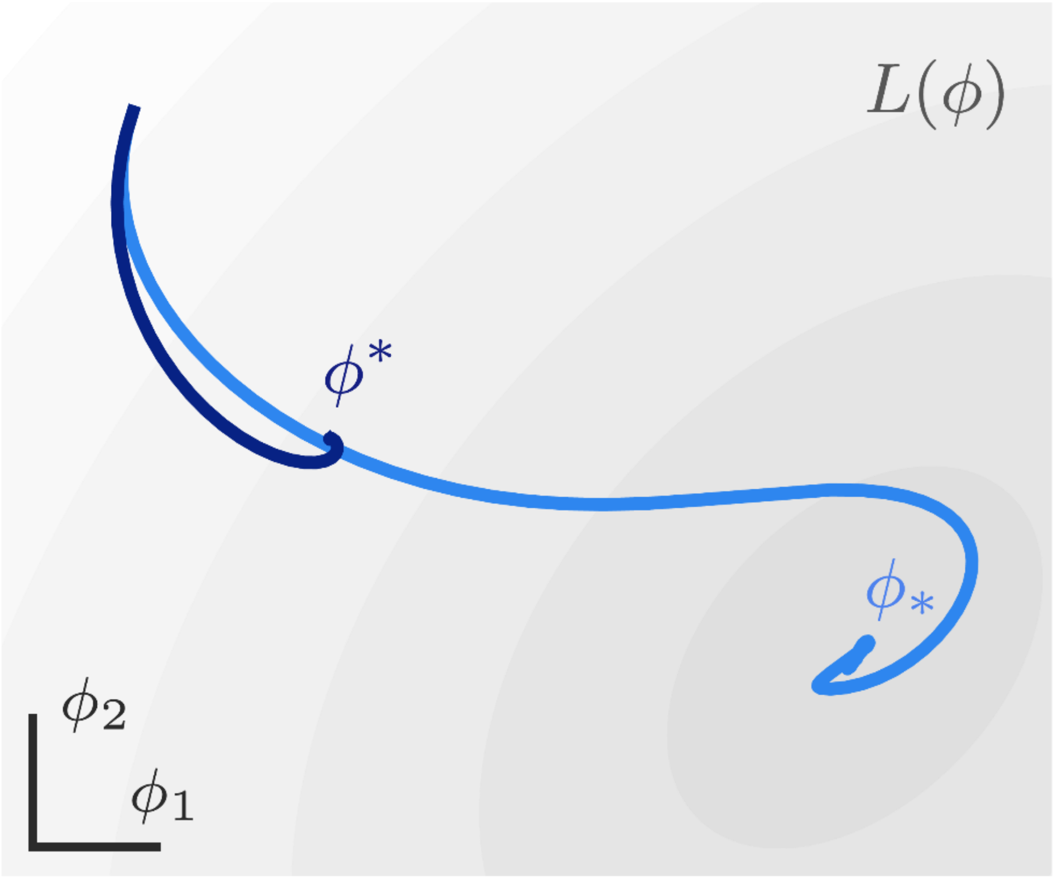

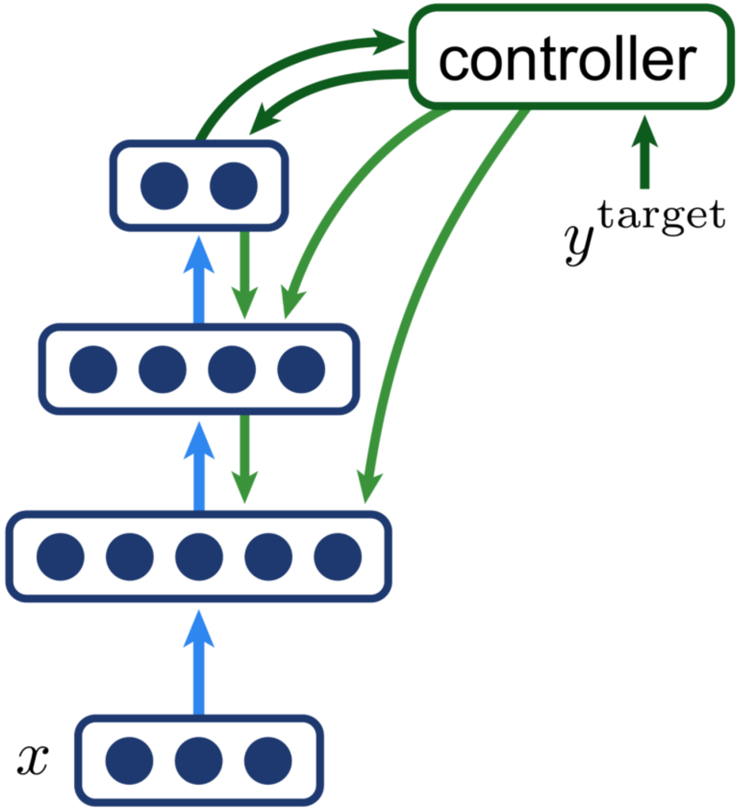

Linking learning objectives

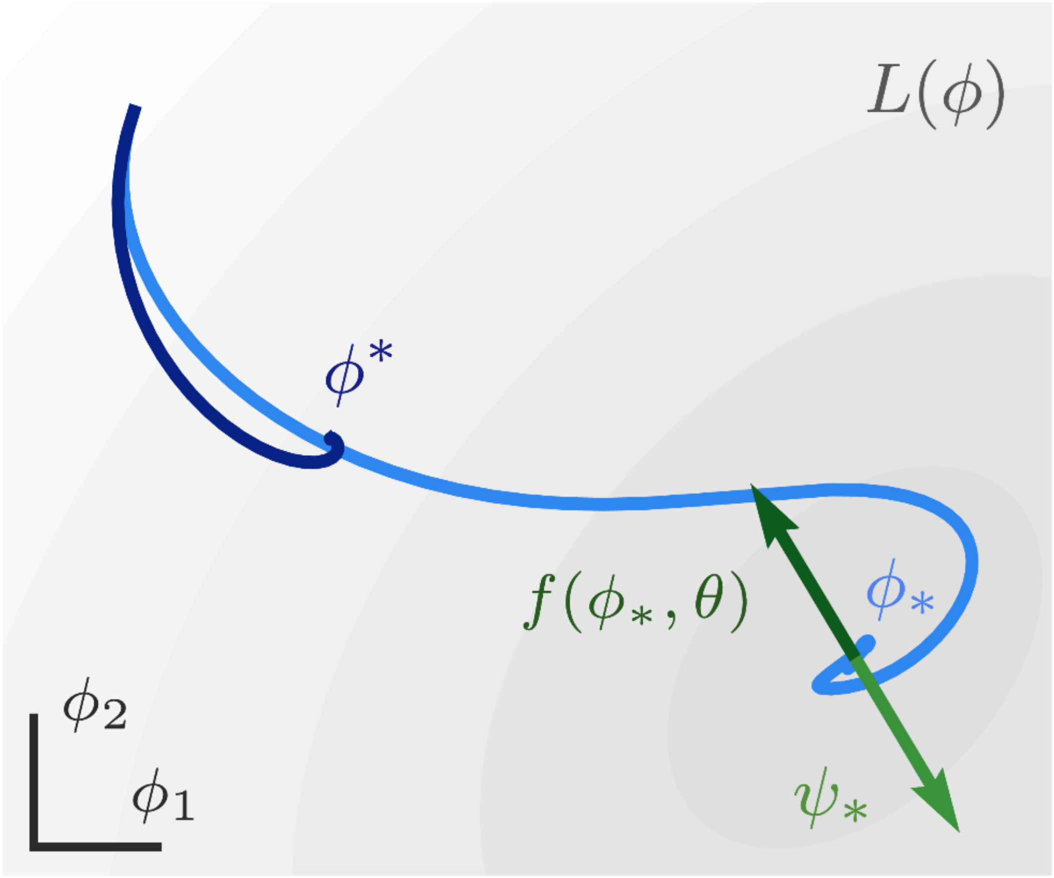

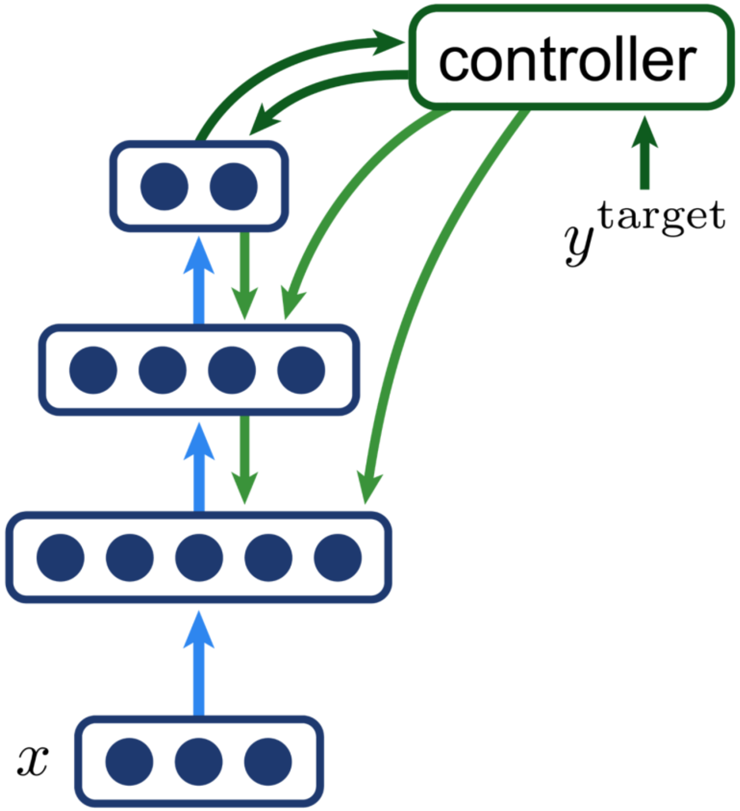

Linking learning objectives

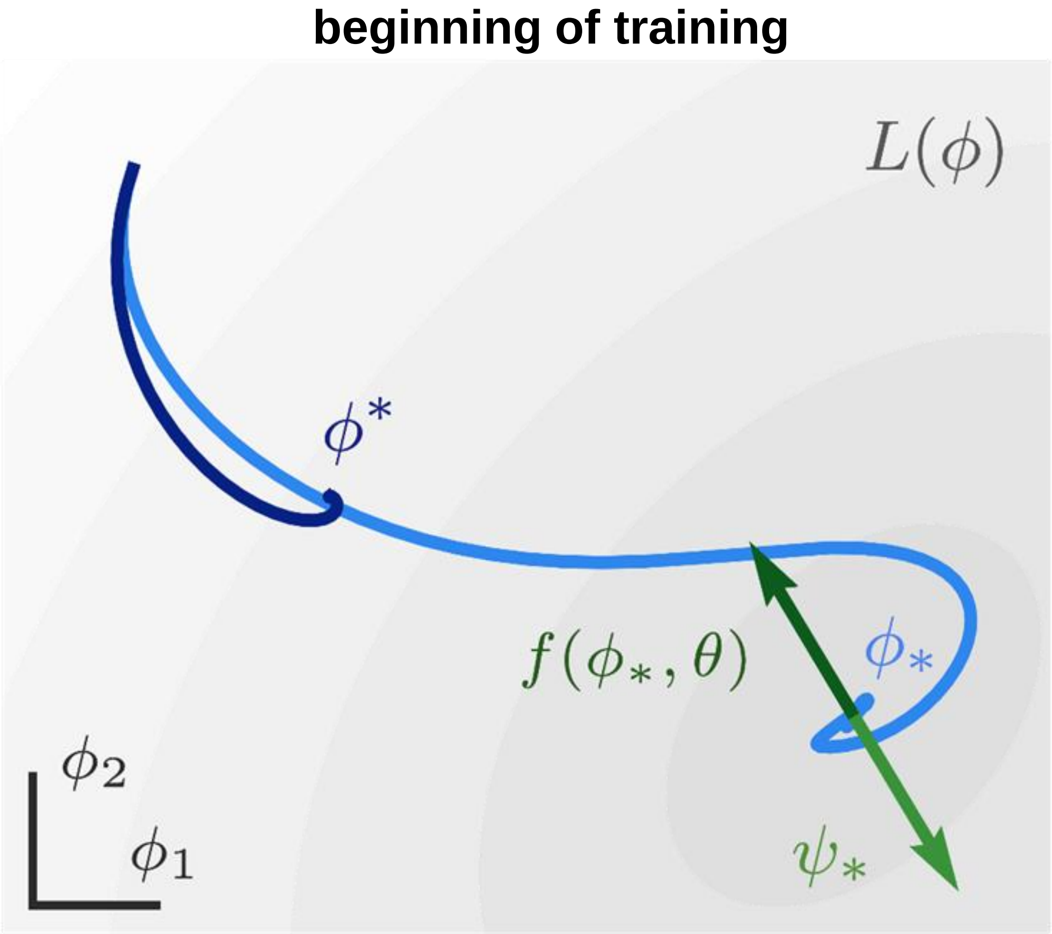

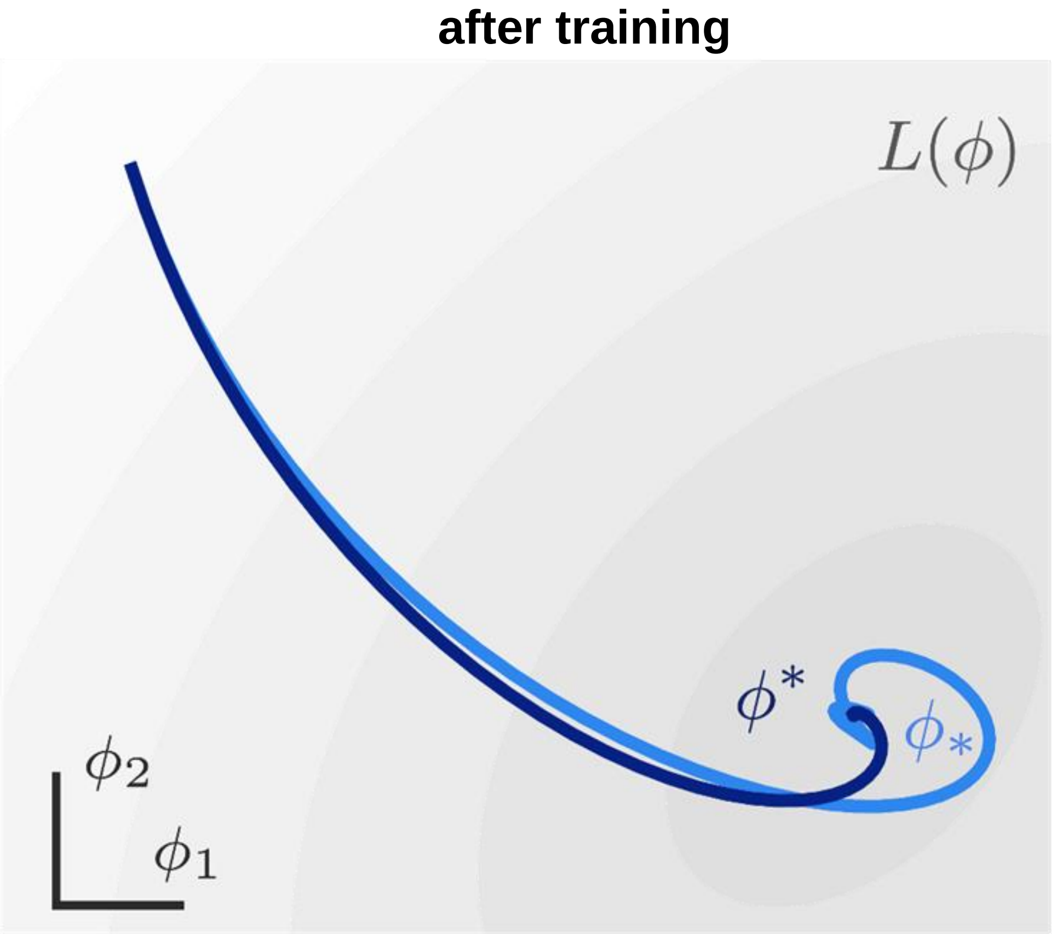

If we minimize the amount of control $\|\psi_*\|^2$ to zero, the original objective is solved

Formal propositions

original problem $$\min_\theta L(\phi^*) \quad \text{s.t.} \;\; f(\phi^*, \theta) = 0, \tag{1}$$

least-control objective $$\mathcal{H}(\theta) = \frac{1}{2}\|\psi_\star\|^2$$

Proposition 3. *Assuming L is convex on the system output, we have that the optimal control $\psi_\star$ is equal to 0 iff. the free equilibrium ϕ^* minimizes L.*

Proposition 4. *Assuming L is convex on the system output, a local minimizer θ of the least control objective $\mathcal{H}$ is a global minimizer of the original learning problem* (1)*, under the sufficient condition that $\partial_\theta f(\phi_\star, \theta)$ has full row rank.*

Condition met in the *overparameterized* regime: $|\theta| \geq |\phi|$

Computing the optimal control:
I. Output control

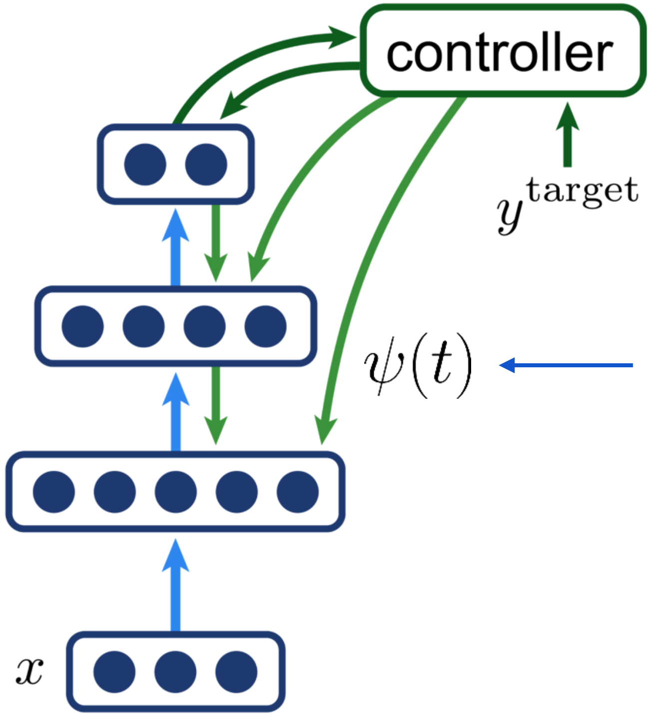

integral output controller

$$\dot{u}(t) = \underbrace{-\nabla_{\phi_L} L(\phi_L(t))}_{\text{output error } y^{\text{target}} - \phi_L} - \alpha u(t)$$

enforces constraint
on the network output

?
should enforce minimum norm constraint on the activity

Computing the optimal control:
II. Conditions for hidden control at equilibrium

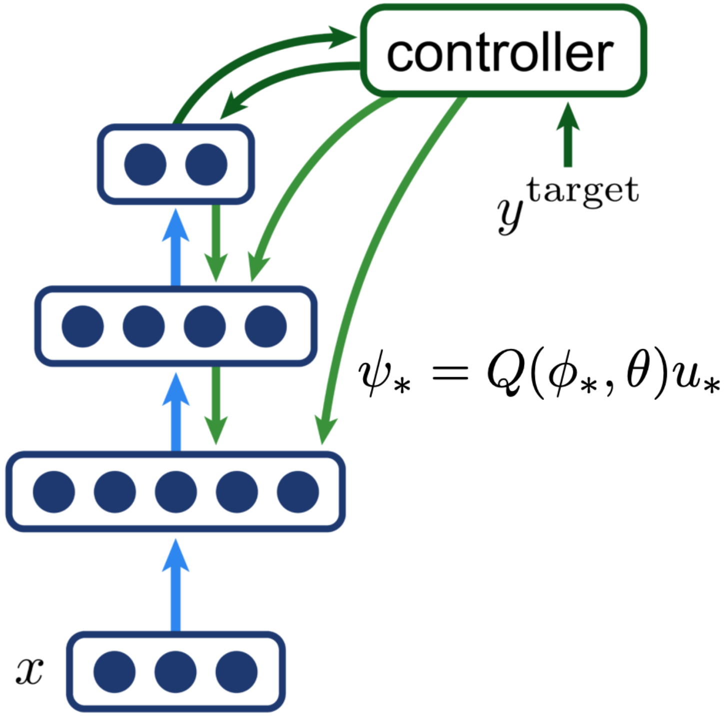

Generalized weight alignment conditions
(recall weight transport problem)

Feedforward network:

$$\mathrm{Col}\big[Q(\phi_*, \theta)\big] = \mathrm{Col}\big[J_*^\top\big]$$

$$J_* = \left[\frac{\partial \phi_L}{\partial \phi_1} \quad \cdots \quad \frac{\partial \phi_L}{\partial \phi_L}\right]\Bigg|_{\phi=\phi_*}$$

⚠ Abuse of partials

Recurrent network:

$$\mathrm{Col}\left[Q(\phi_*, \theta)\right] = \mathrm{Col}\left[(\mathrm{Id} - \sigma'(\phi_*)W^\top)^{-1}D^\top\right]$$

Computing the optimal control: III. Neural control circuits

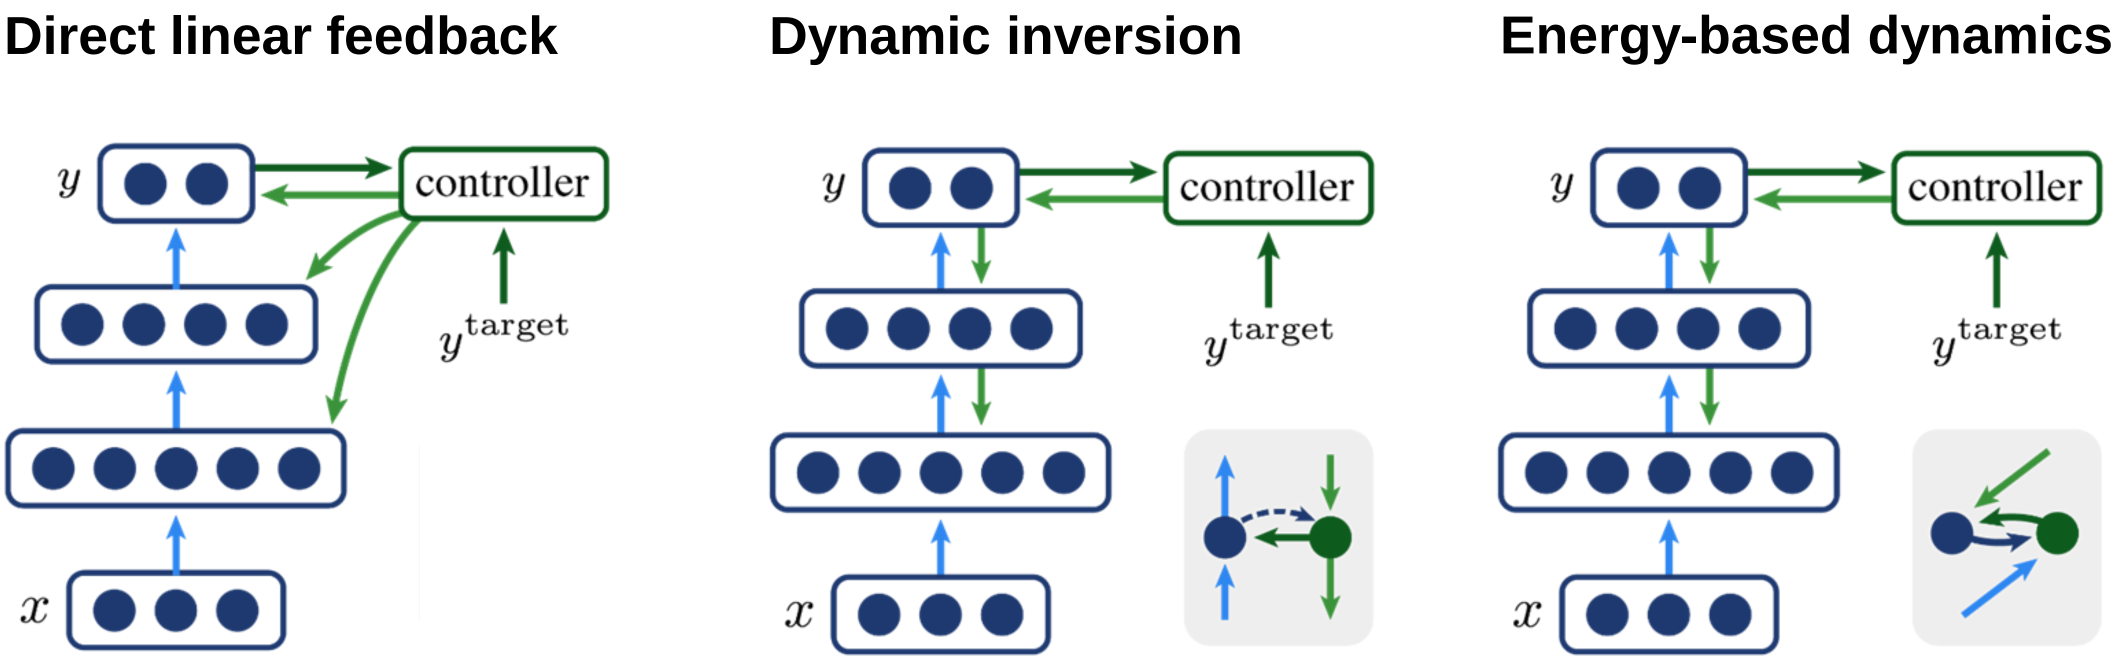

Central idea: combine network dynamics and controller dynamics to settle down to an optimally controlled state

Approach 1: direct linear feedback

Feedforward network

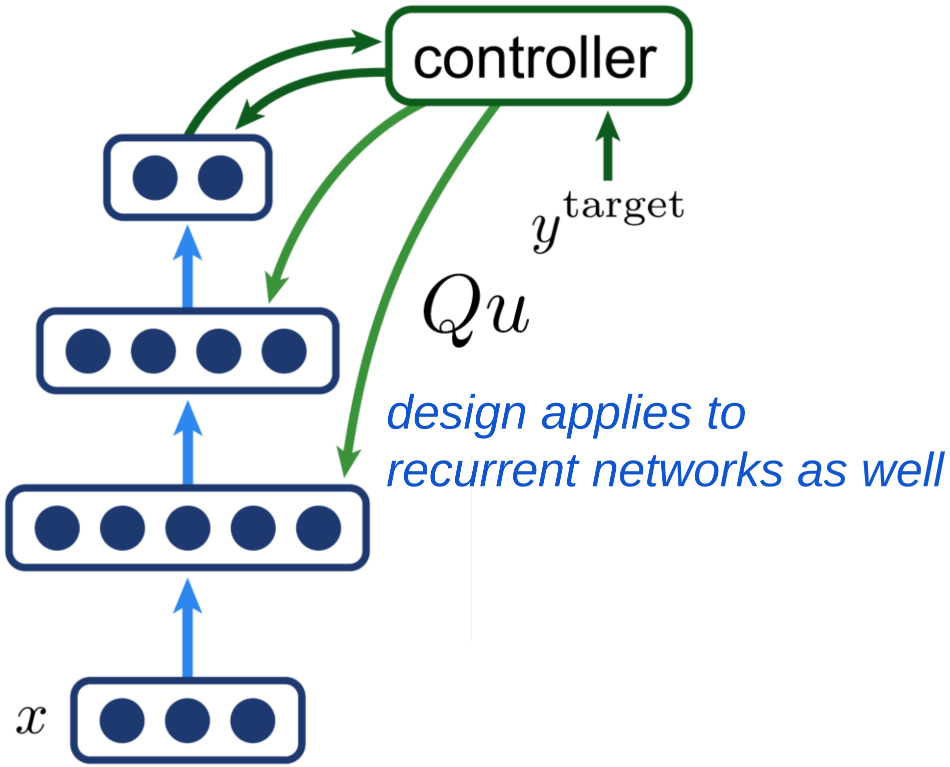

Central idea: Take the simplest feedback controller architecture, and train *Q* to approximately satisfy optimal control conditions

$$\dot{\phi}_l(t) = -\phi_l(t) + W_l\sigma\big(\phi_{l-1}(t)\big) + Q_l u(t)$$

– Learn feedback weights *Q* with the help of noise

– Approximate control for nonlinear networks

Lillicrap et al., *Nat. Comm.* 2016; Nokland *NeurIPS* 2016;
Meulemans et al. *NeurIPS* 2021, ‘22; Akrout et al. *NeurIPS* 2019

Approach 2: dynamic inversion

Feedforward neural network

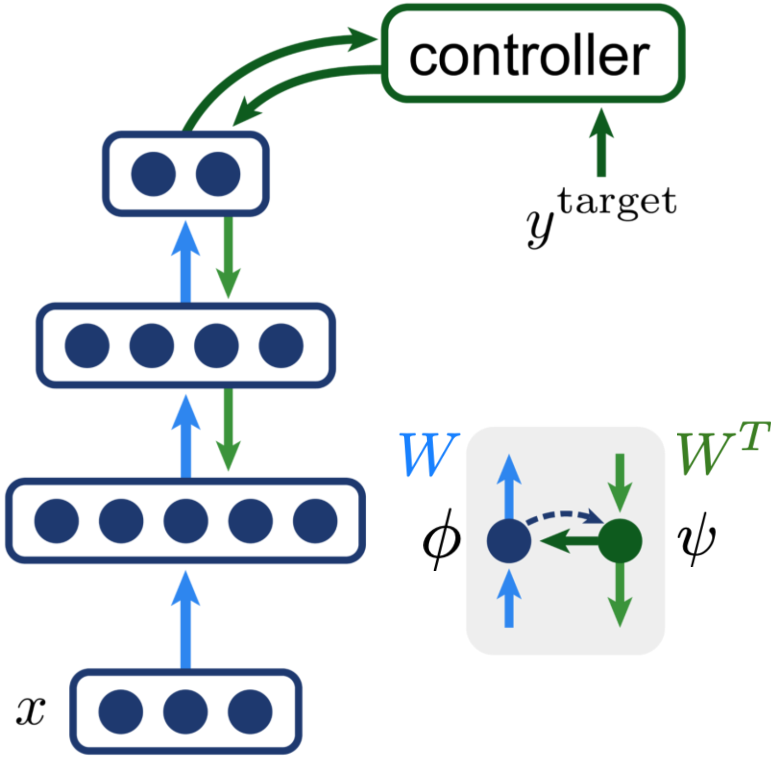

Central idea: combine network dynamics $\dot{\phi}$ and controller dynamics $\dot{\psi}$ to jointly converge to an optimal control

$$\dot{\phi}_l(t) = -\phi_l(t) + W_l\sigma\big(\phi_l(t)\big) + \psi_l(t)$$

$$\dot{\psi}_l(t) = -\psi_l(t) + \sigma' W_{l+1}^\top \psi_{l+1}(t)$$

- Compare to backpropagation
- Feedback weights can be learned too
- Exact optimal control

Approach 2: dynamic inversion

Recurrent neural network

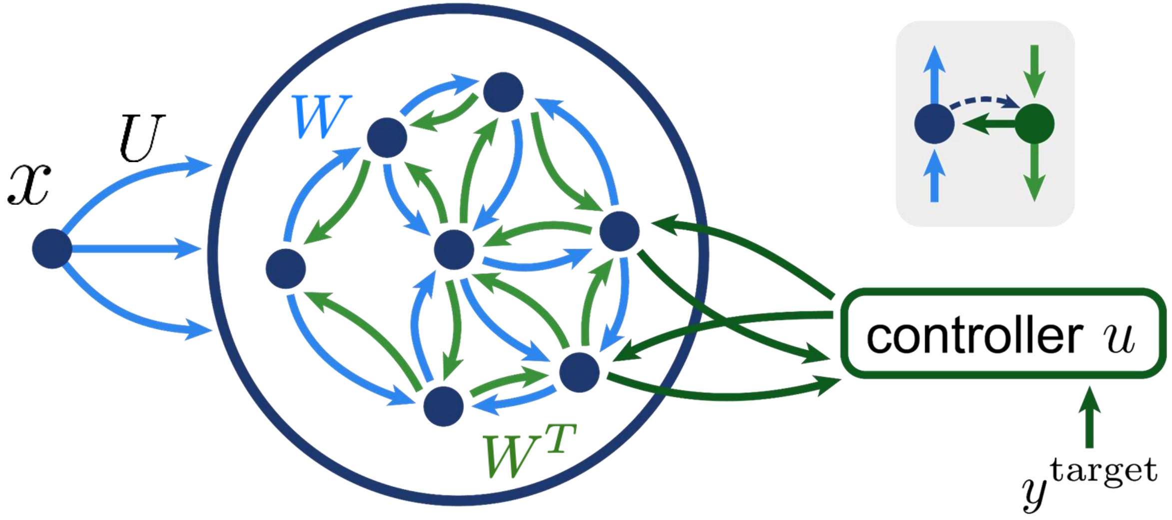

Central idea: combine network dynamics $\dot{\phi}$ and controller dynamics $\dot{\psi}$ to jointly converge to an optimal control

$$\dot{\phi} = -\phi + W\sigma(\phi) + Ux + \psi$$
$$\dot{\psi} = -\psi + \sigma'(\phi)W^\top \psi + S_y^\top u$$

- Compare to recurrent backpropagation
- Feedback weights can be learned too
- Exact optimal control

Canonical microcircuits of the cortex

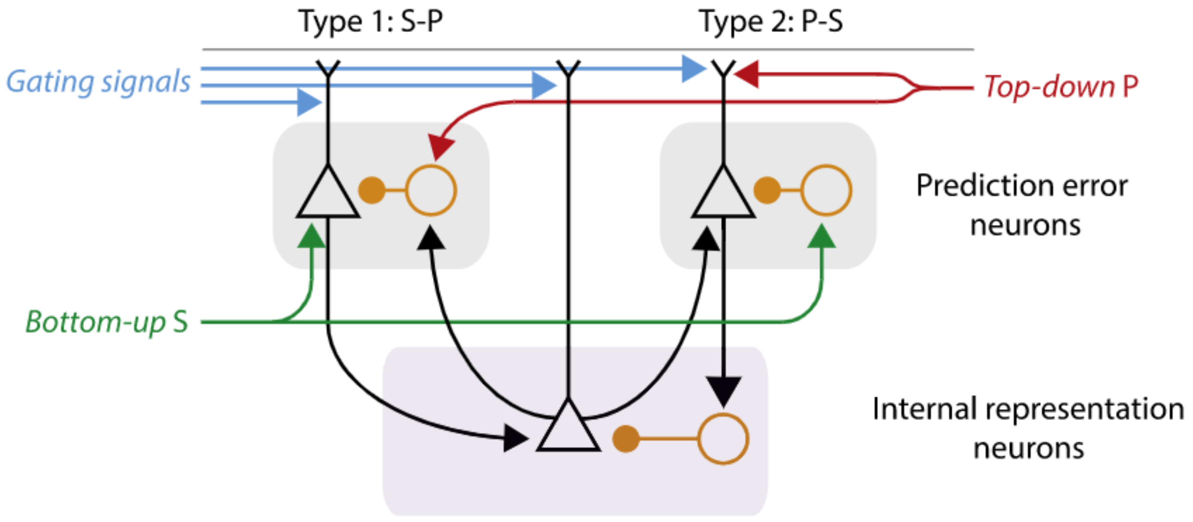

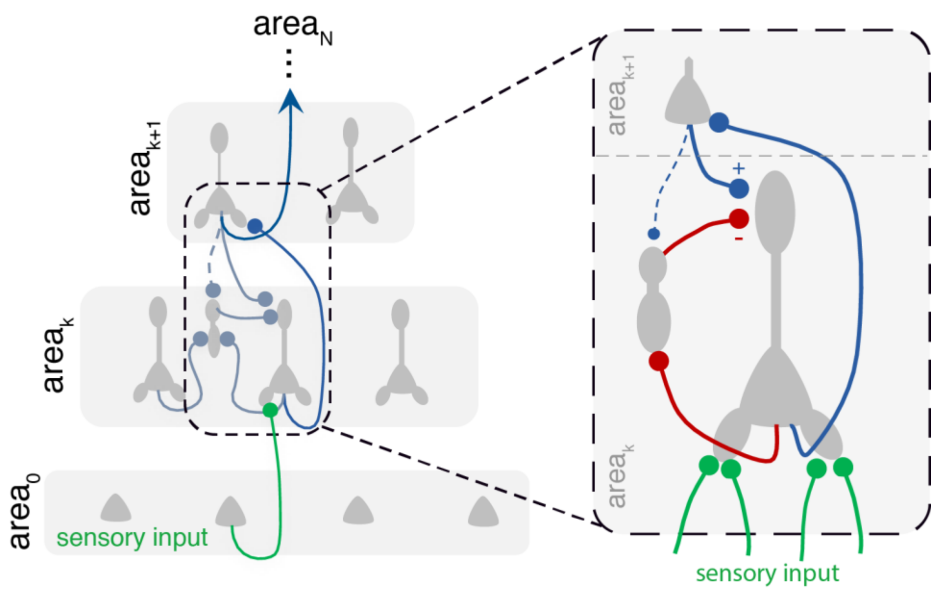

Douglas et al., *Neural Comput.* 1989
Bastos et al., *Neuron* 2012
Keller & Mrsic-Flogel, *Neuron* 2018
Sacramento et al., *NeurIPS* 2018
Whittington & Bogacz, *Trends Cogn. Sci.* 2019
Payeur et al., *Nat. Neurosci.* 2021

- Flexible control conditions suggest new feedback circuits
- Compatible with previous microcircuit “motifs”
- Feedback can now strongly influence neural activity

Interim summary

We have:

- Embedded credit information within a dynamics as an optimal control
  - Control brings network towards a minimal-loss equilibrium state

- Reformulated learning as minimizing the amount of control at equilibrium

- Shown that (minimum-norm) least-control states enjoy desirable properties:
  - Yield local gradient-based learning rule for adapting parameters
  - Allow for flexible design of neural control circuits
  - Admit optimal parameters that also solve the original learning problem

Competitive performance in practice

	Feedforward (test accuracy, %)		Recurrent (test accuracy, %)	
	MNIST	CIFAR-10	MNIST	CIFAR-10
LCP (linear feedback)	97.73 ± 0.07	/	97.70 ± 0.11	/
LCP (dynamic inversion)	98.11 ± 0.07	77.28 ± 0.10	97.58 ± 0.16	80.26 ± 0.17
LCP (dynamic inversion + learned feedback)	98.14 ± 0.09	77.16 ± 0.10	97.75 ± 0.11	/
(R)BP	98.29 ± 0.14	77.58 ± 0.14	97.87 ± 0.19	80.14 ± 0.20
	2-hidden-layer feedforward network	convolutional network	fully-connected equilibrium RNN	convolutional equilibrium RNN

Bai et al., *NeurIPS* 2020

Outline

- Credit assignment in artificial and biological neural networks
- Assigning credit by backpropagation-of-error
- Assigning credit by controlling a neural dynamics
- Theory & experimental results
- Connections to probabilistic modeling and free-energy minimization

Connection to energy-based learning

$$\boxed{\min_{\phi,\psi,\theta} \frac{1}{2}\|\psi\|^2 \quad \text{s.t.} \quad f_\theta(\phi, x) + \psi = 0, \quad \nabla_\phi L(\phi) = 0}$$

the least-control principle

Substituting the first constraint into the objective, we get, for feedforward neural networks:

$$\min_{\phi,\theta} \frac{1}{2} \sum_l \|\phi_l - W_l \sigma(\phi_{l-1})\|^2 + \beta L(\phi)$$

augmented energy

$$= \min_{\phi,\theta} \{- \ln p(\phi, y \mid x; \theta)\}$$

variational expectation maximization
under a conditional Gaussian probabilistic model

Letting $\beta \to \infty$ we retrieve the least-control objective

- Solving a least-control problem is equivalent to free-energy minimization
- Gradient flow on the free-energy yields a least-control state
- New results for classical hierarchical predictive coding models

Rao & Ballard, *Nat. Neurosci.* 1999
Friston, *Trends Cogn. Sci.* 2009
Whittington & Bogacz, *Trends Cogn. Sci.* 2019

Unifying the spectrum

$$\boxed{\frac{1}{2}\|f(\phi,\theta)\|^2 + \beta L(\phi) = 0}$$

augmented energy

Weak nudging		*Perfect control*
$\beta \to 0$	β	$\beta \to \infty$

Equilibrium propagation	**Augmented energy**	**Least-control principle**
$\min_\theta L(\phi^*)$ s.t. $\phi^* \in \arg\min_\phi \|f(\phi,\theta)\|^2$	$\min_{\phi,\theta} \frac{1}{2}\|f(\phi,\theta)\|^2 + \beta L(\phi)$	$\min_{\phi,\theta} \|f(\phi,\theta)\|^2$ s.t. $\nabla_\phi L(\phi) = 0$

Scellier & Bengio
Front. Comput. Neurosci. 2017

Conclusions

- Local learning rules by embedding credit assignment information within a neural dynamics
- Credit assignment results from controlling dynamics; no weak perturbations
- Flexible conditions for designing control circuits
- Strong performance on standard supervised learning benchmarks
- Dynamics at equilibrium generalize feedforward networks.
  Learning away from equilibrium?

Thank you

Alexander Meulemans

Nicolas Zucchet

Seijin Kobayashi

Johannes von Oswald

Angelika Steger

Funding:

ETH zürich

Schweizerischer Nationalfonds

More results in:

The least-control principle for local learning at equilibrium

A. Meulemans*, N. Zucchet*, S. Kobayashi*, J. von Oswald, J. Sacramento

NeurIPS 2022