Contraction in Neural Networks and Biologically Plausible Optimization

Giovanni Russo

University of Salerno

giovarusso@unisa.it

December 11, 2025

Conference on Decision and Control
CDC2025
Rio de Janeiro, Brazil. Dec. 09 - 12

CDC Tutorial *"Contraction Theory in Control, Optimization and Learning"*

Acknowledgments

Francesco Bullo

UCSB

Veronica Centorrino

SSM -> ETH

Alexander Davydov

UCSB -> Rice Univ.

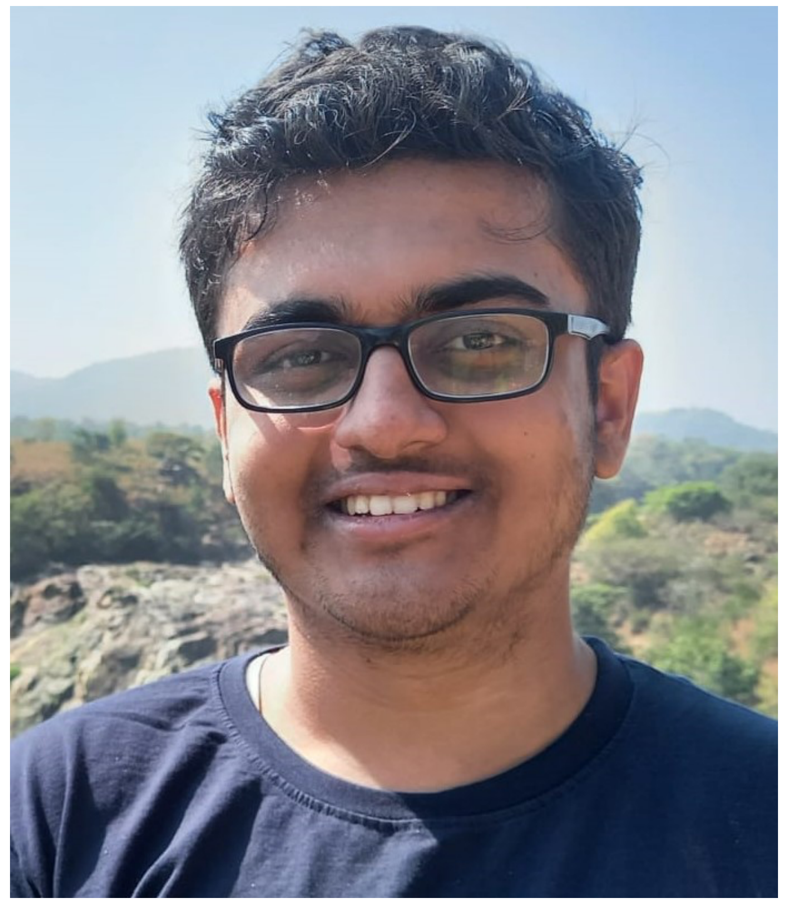

Anand Gokhale

UCSB

Francesca Rossi

SSM

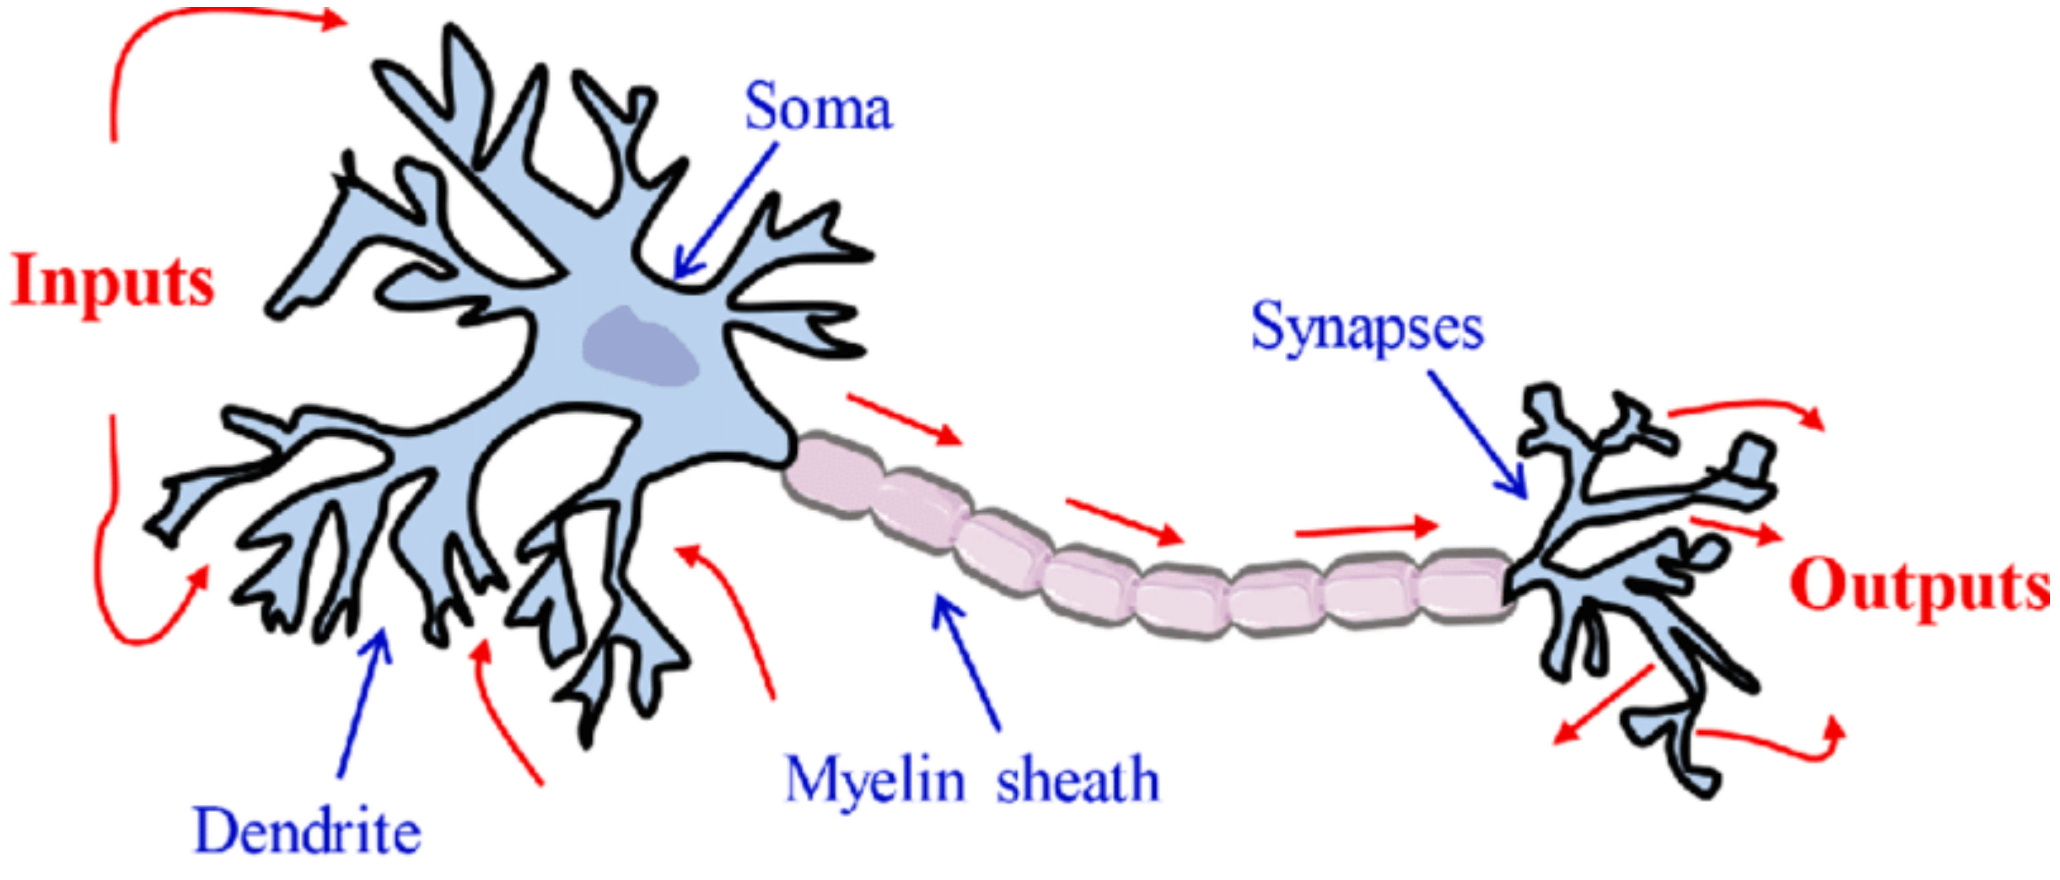

(a) Biological neuron

Introduction and Motivation

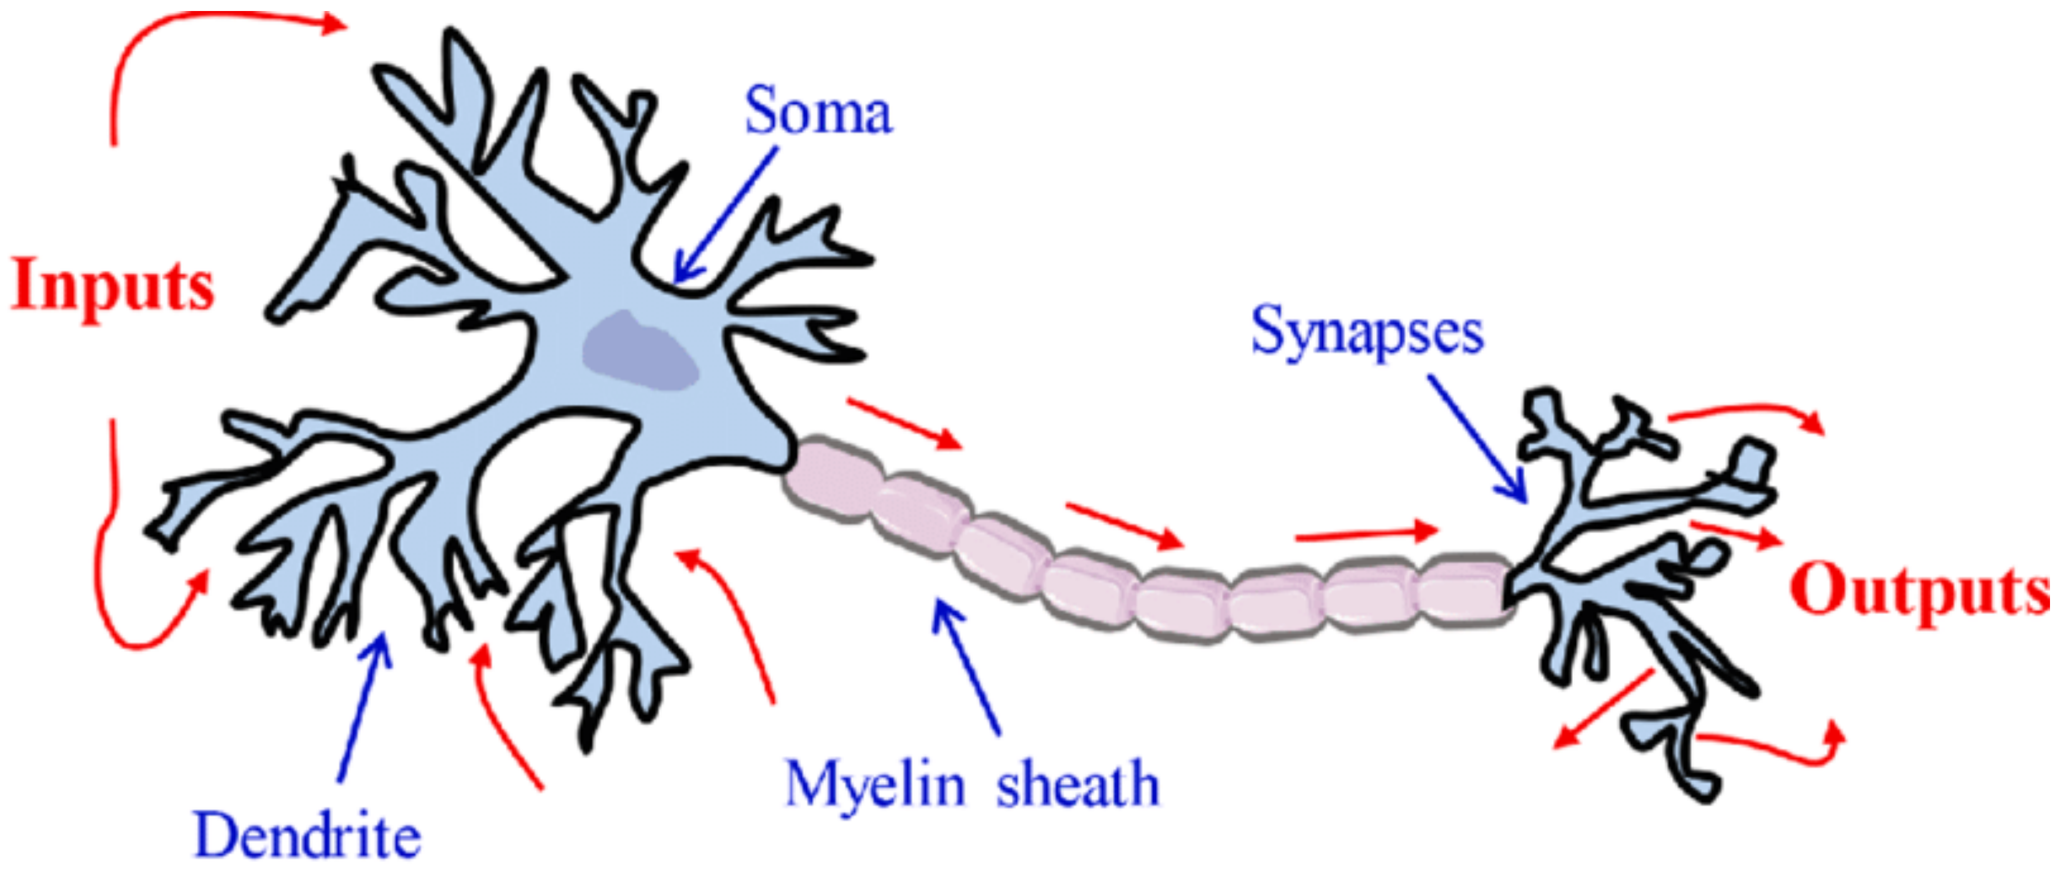

(a) Biological neuron

ANNs have diverged significantly from biological realism, driven by performance criteria

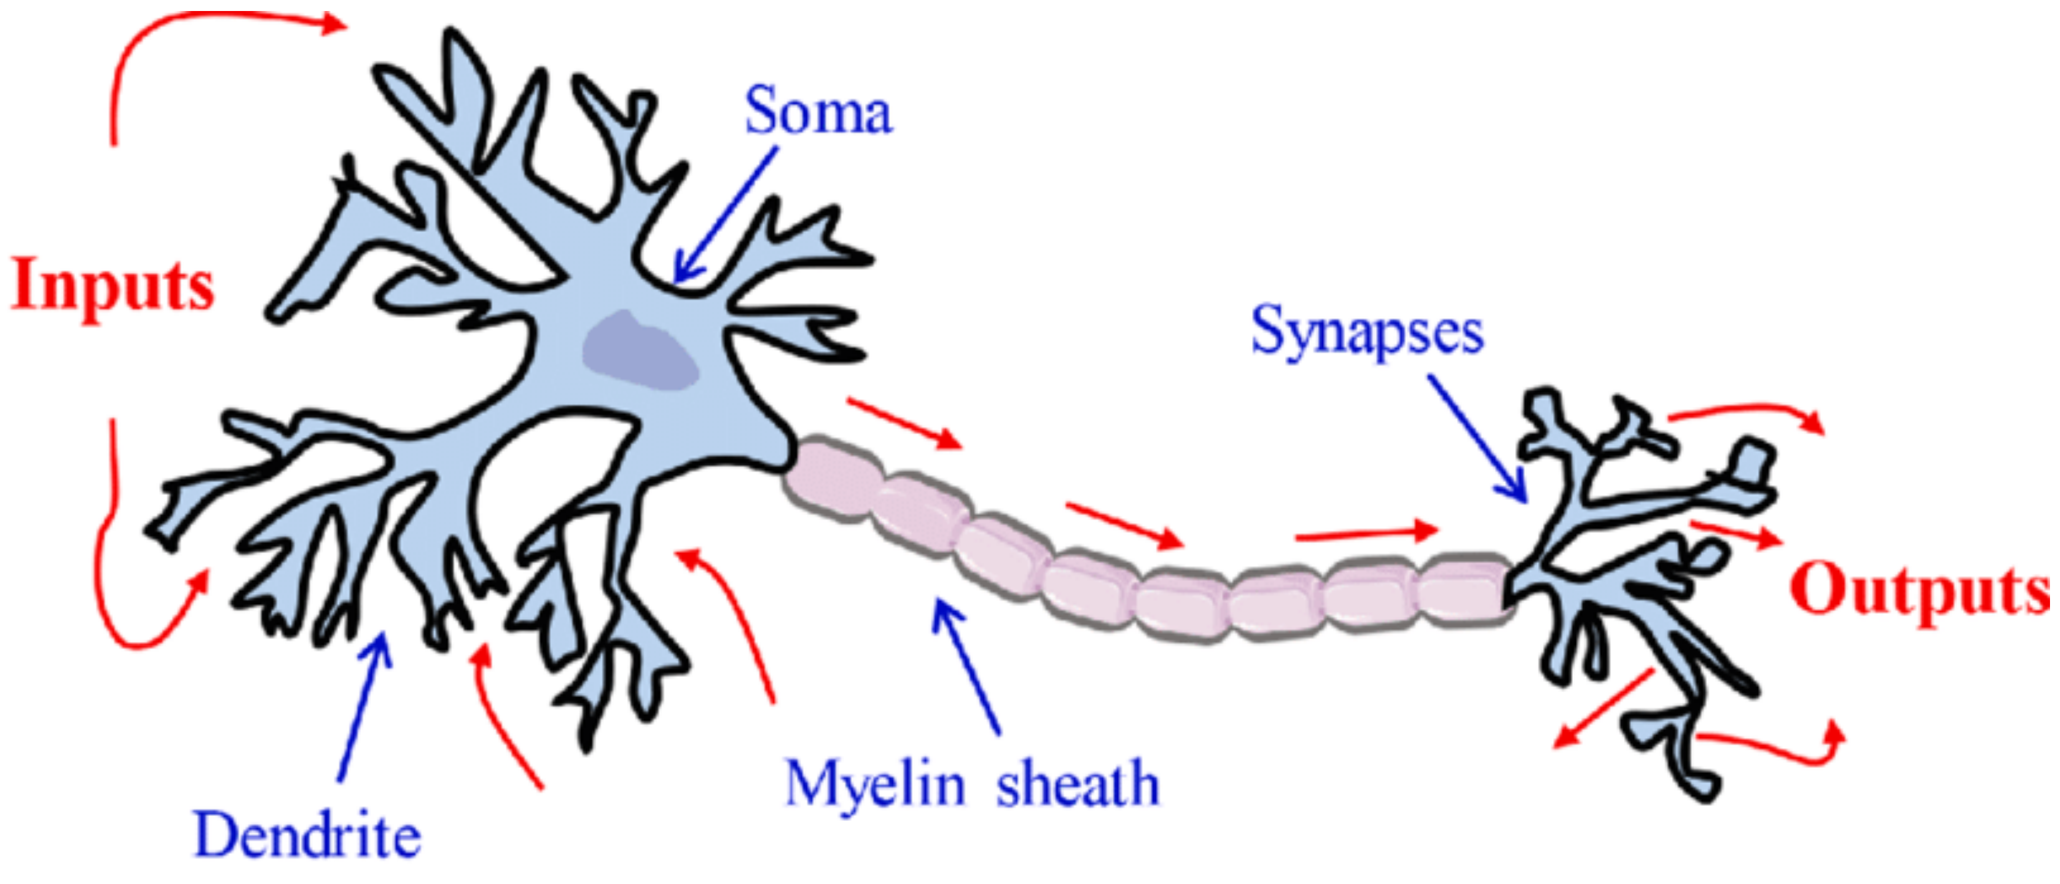

(a) Biological neuron

ANNs have diverged significantly from biological realism, driven by performance criteria

Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation

Benjamin Scellier and Yoshua Bengio †*

Département d'Informatique et de Recherche Opérationnelle, Montreal Institute for Learning Algorithms, Université de Montréal, Montreal, QC, Canada

Frontiers in Computational Neuroscience, 11:24, 2017. doi:10.3389/fncom. 2017.00024.

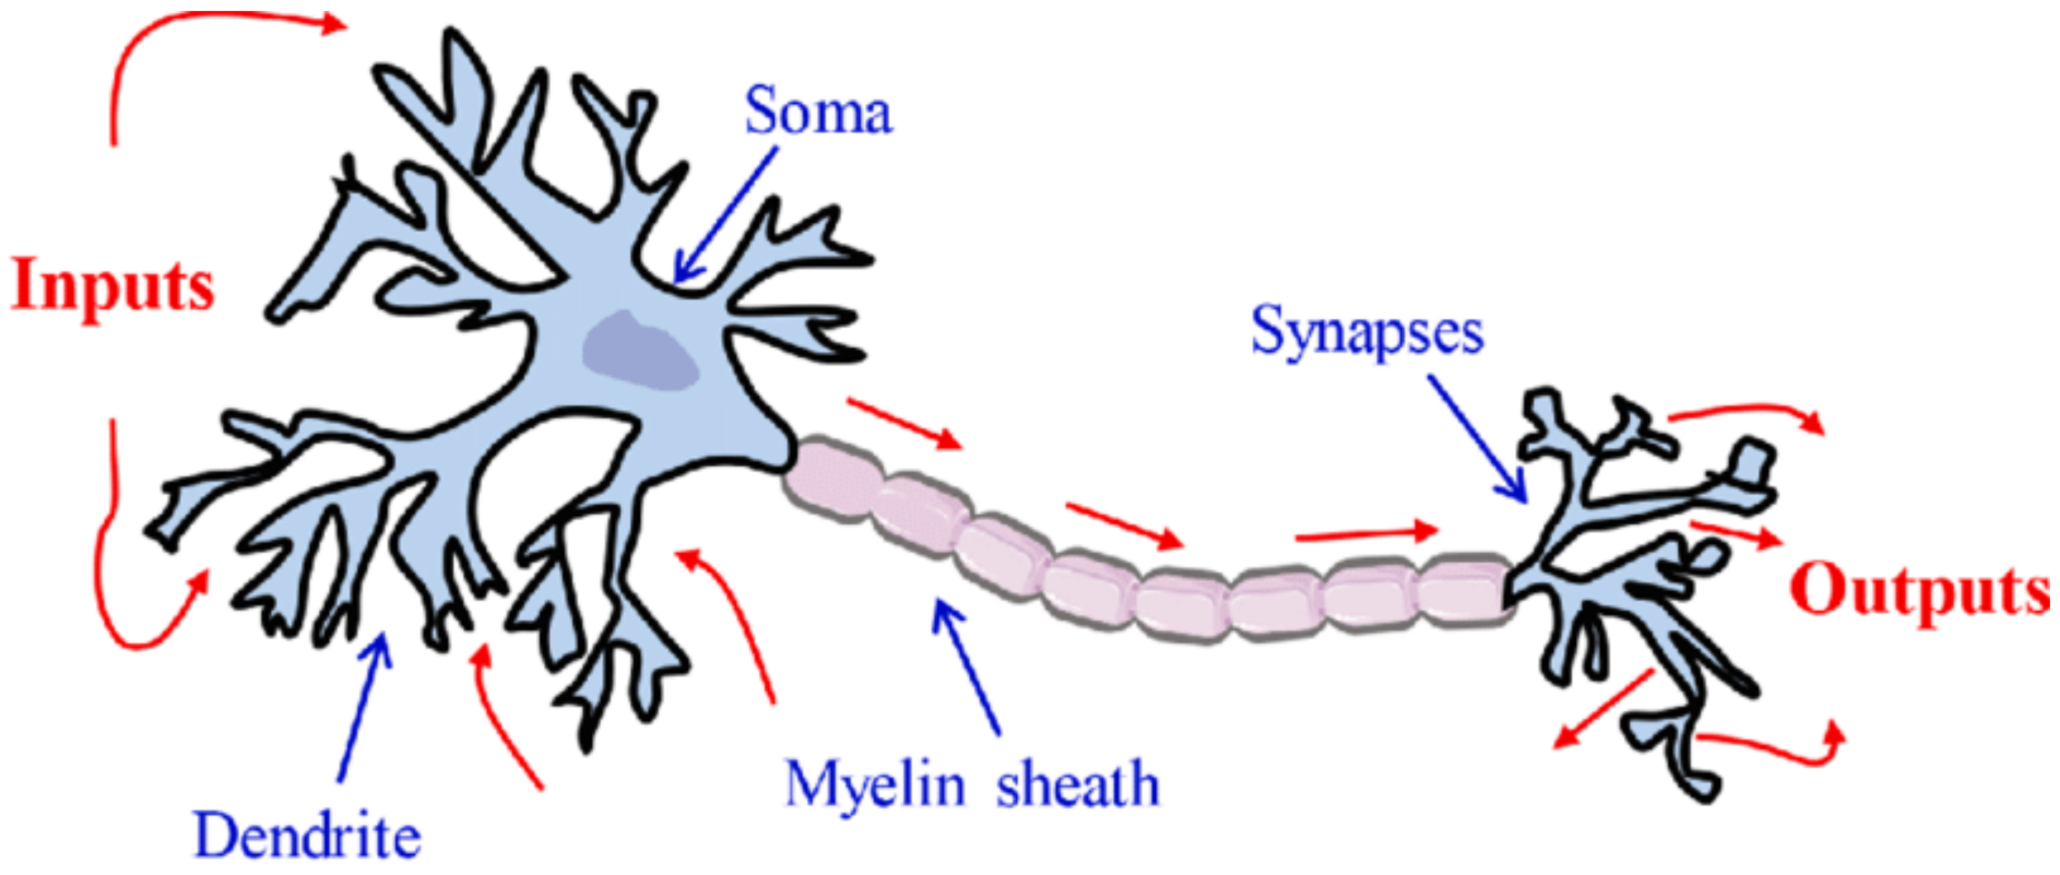

(a) Biological neuron

ANNs have diverged significantly from biological realism, driven by performance criteria

- **Interpretable**
- **Local learning rules**
- **Convergence**

Outline

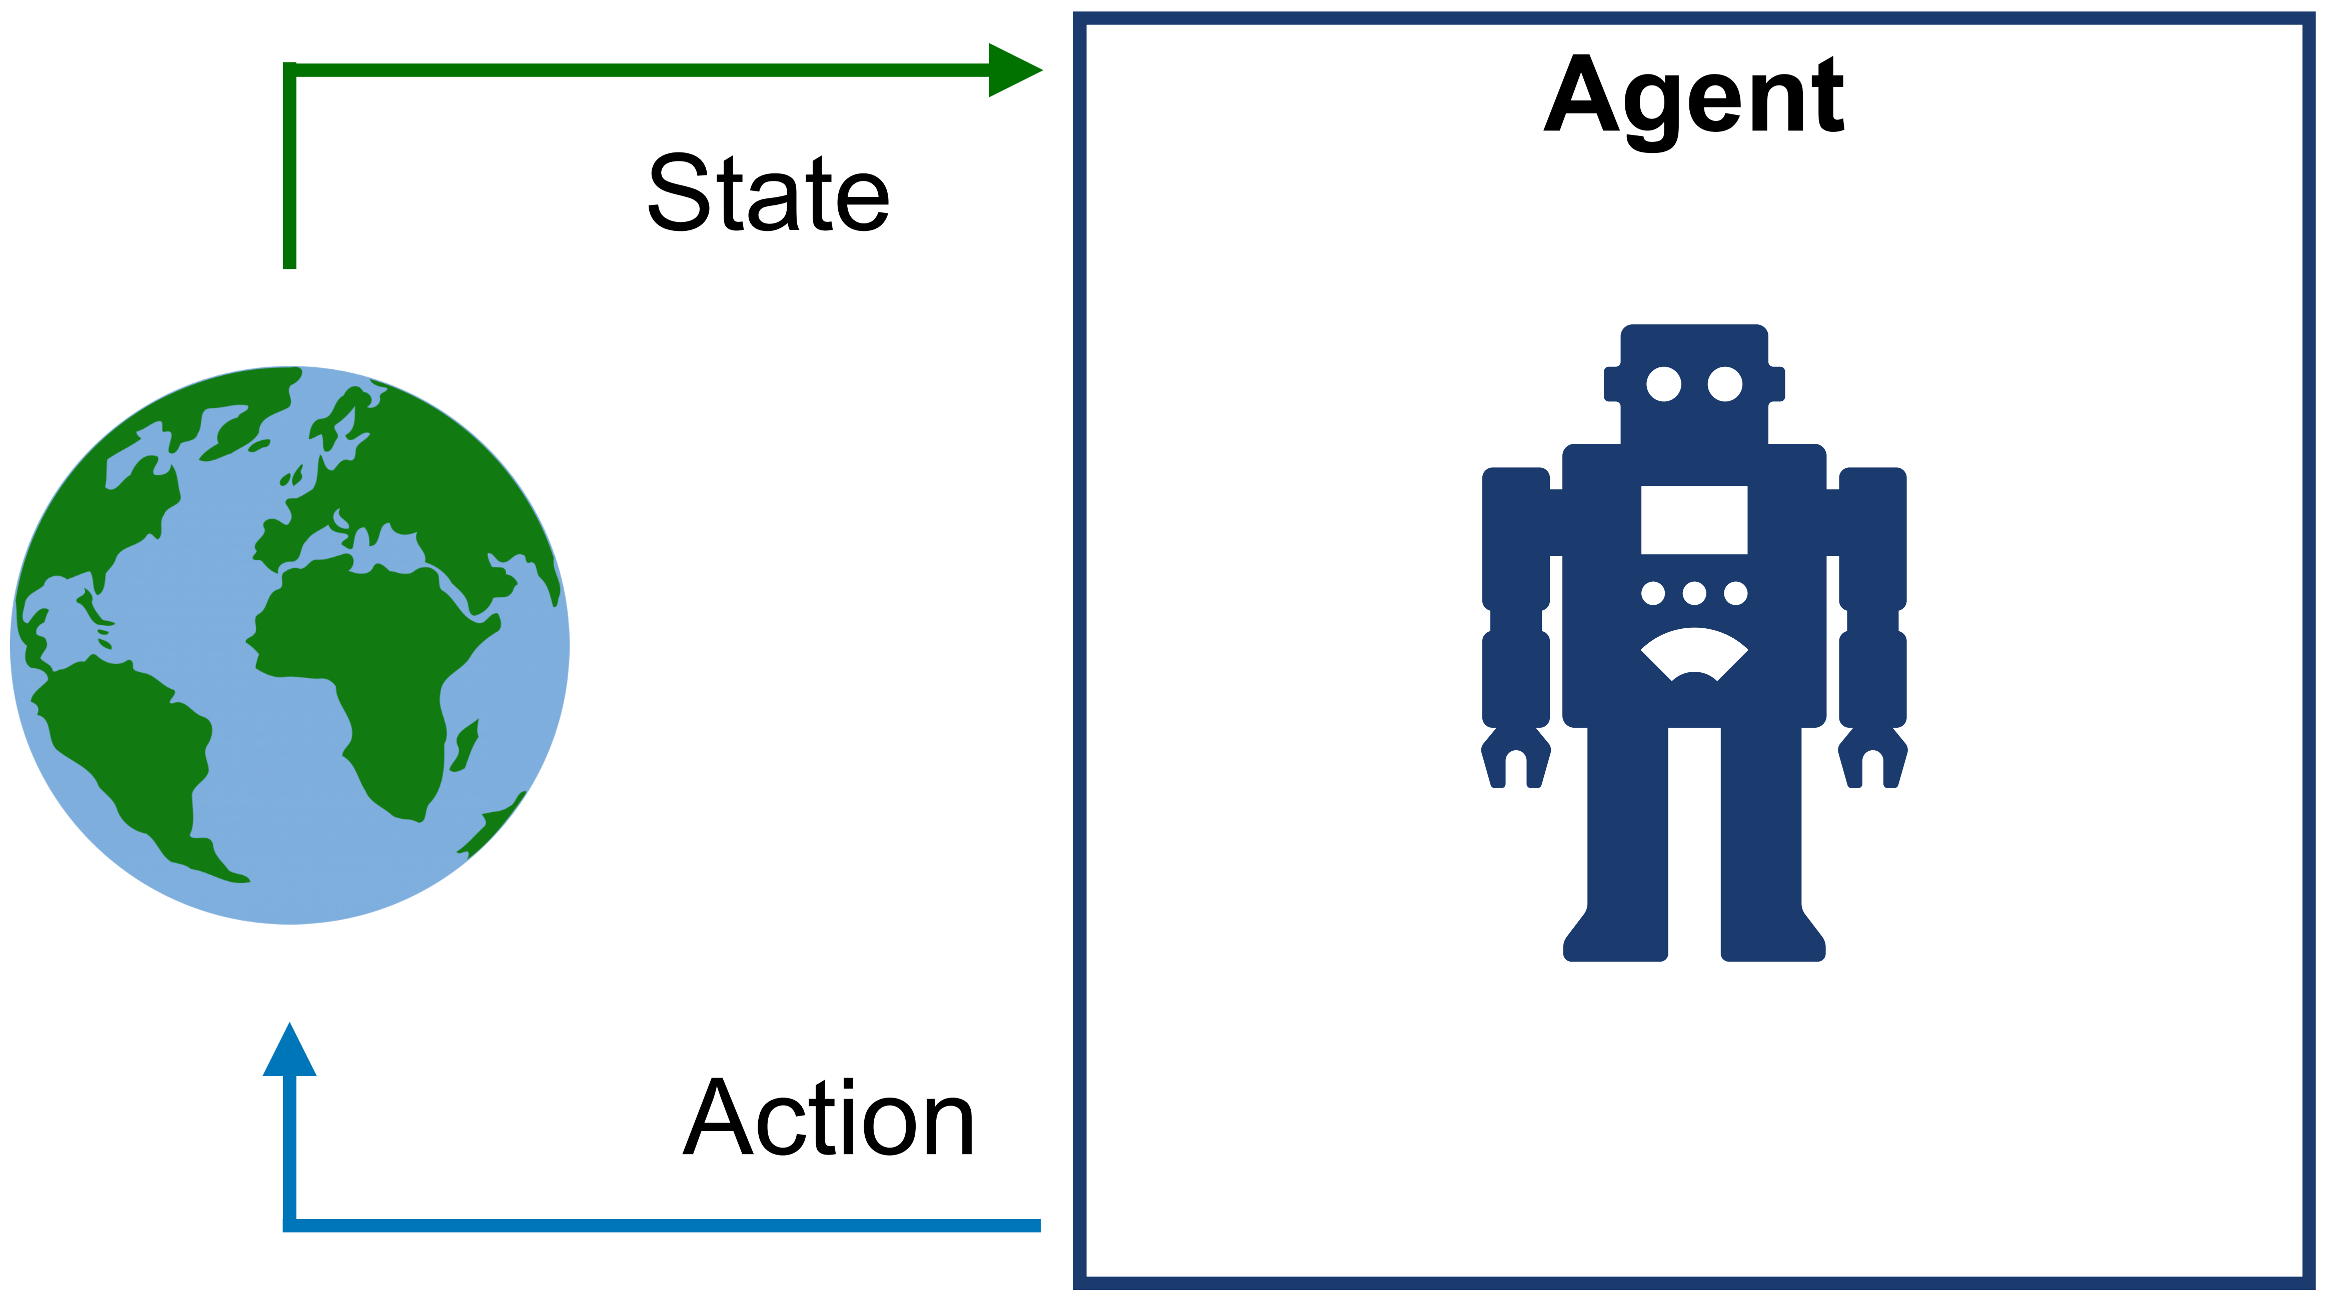

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Outline

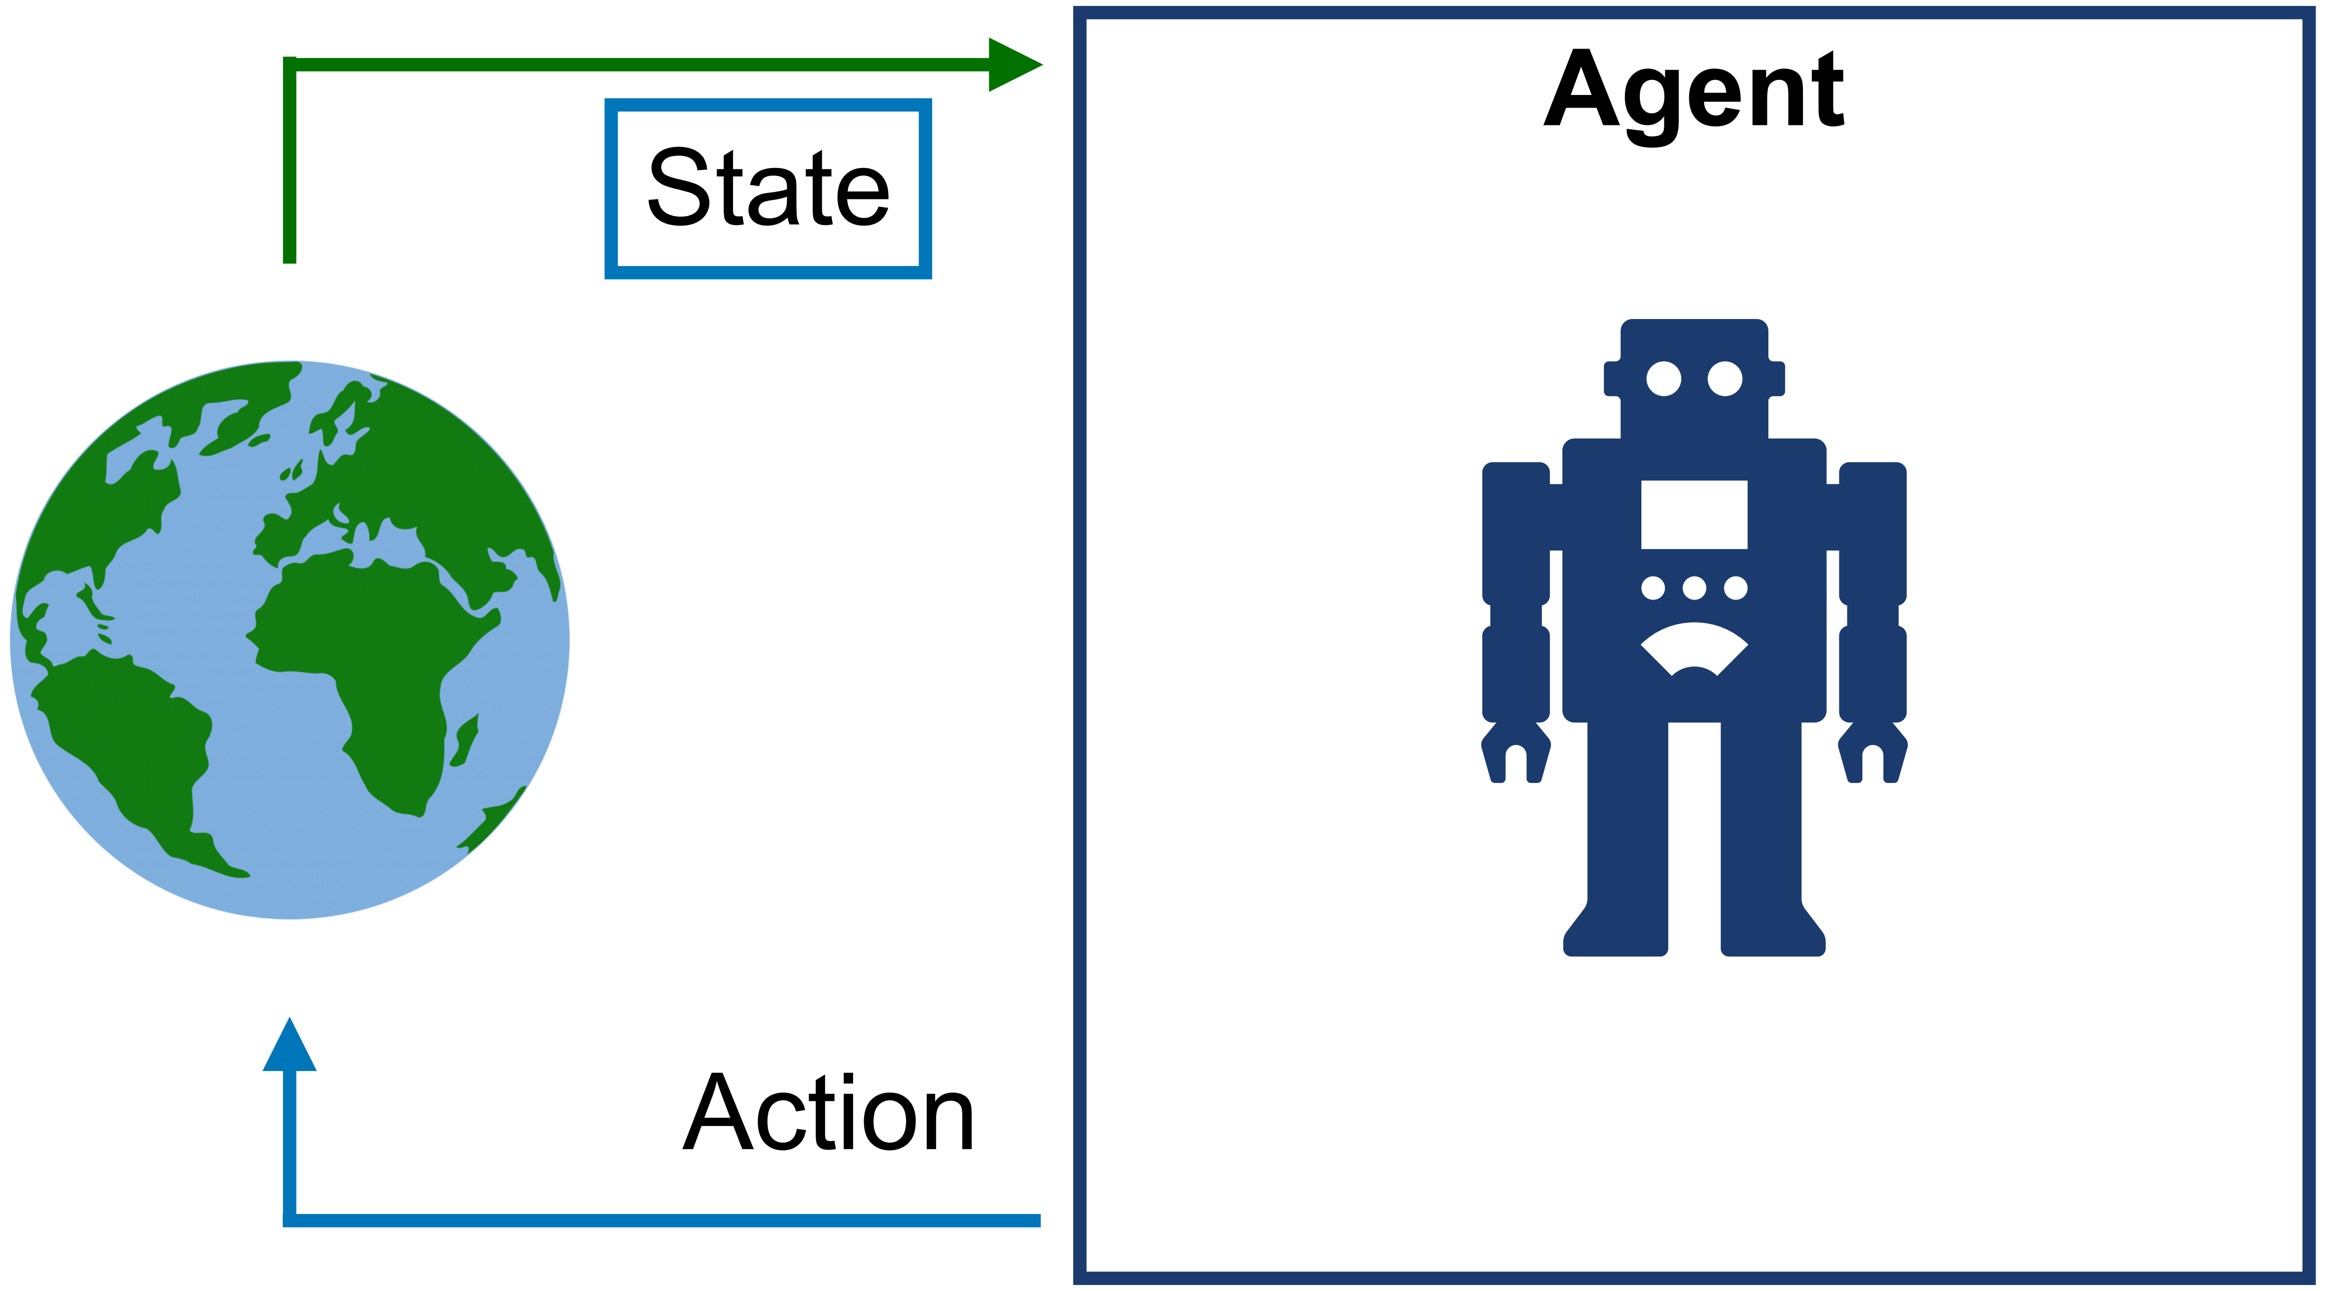

1) Sparse reconstruction

- [x] Optimization-based explanation
- [x] Competitive positive firing rate dynamics

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Outline

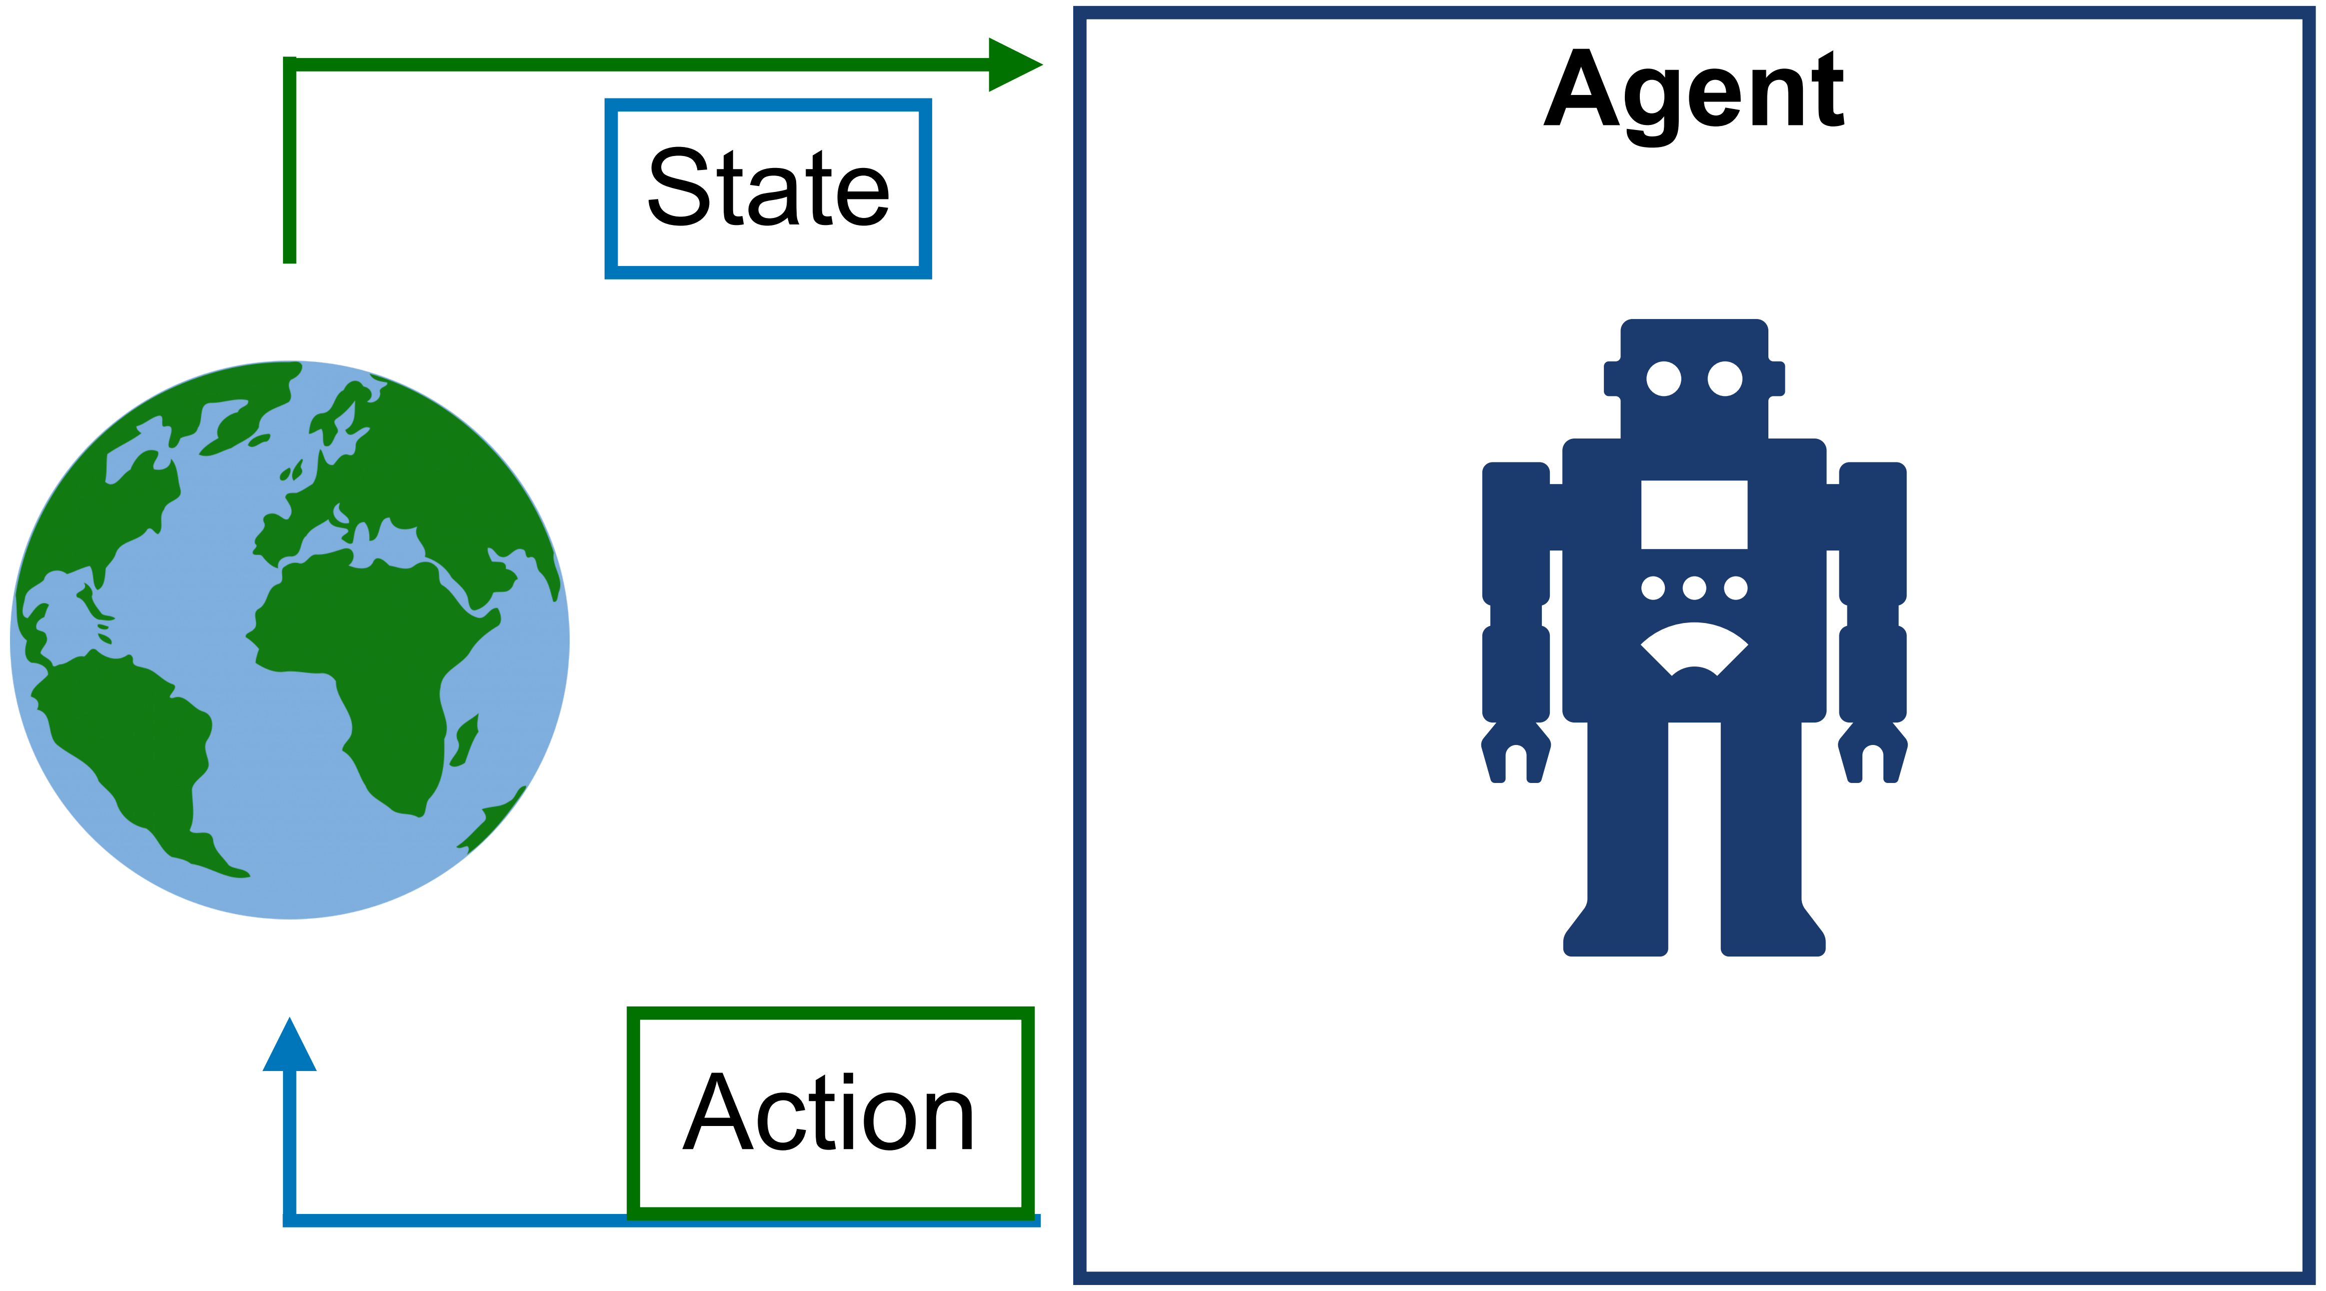

1) Sparse reconstruction

- ☑ Optimization-based explanation
- ☑ Competitive positive firing rate dynamics

2) Policy Composition

- ☑ Free energy minimization
- ☑ Softmax flows from entropy max.

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Outline

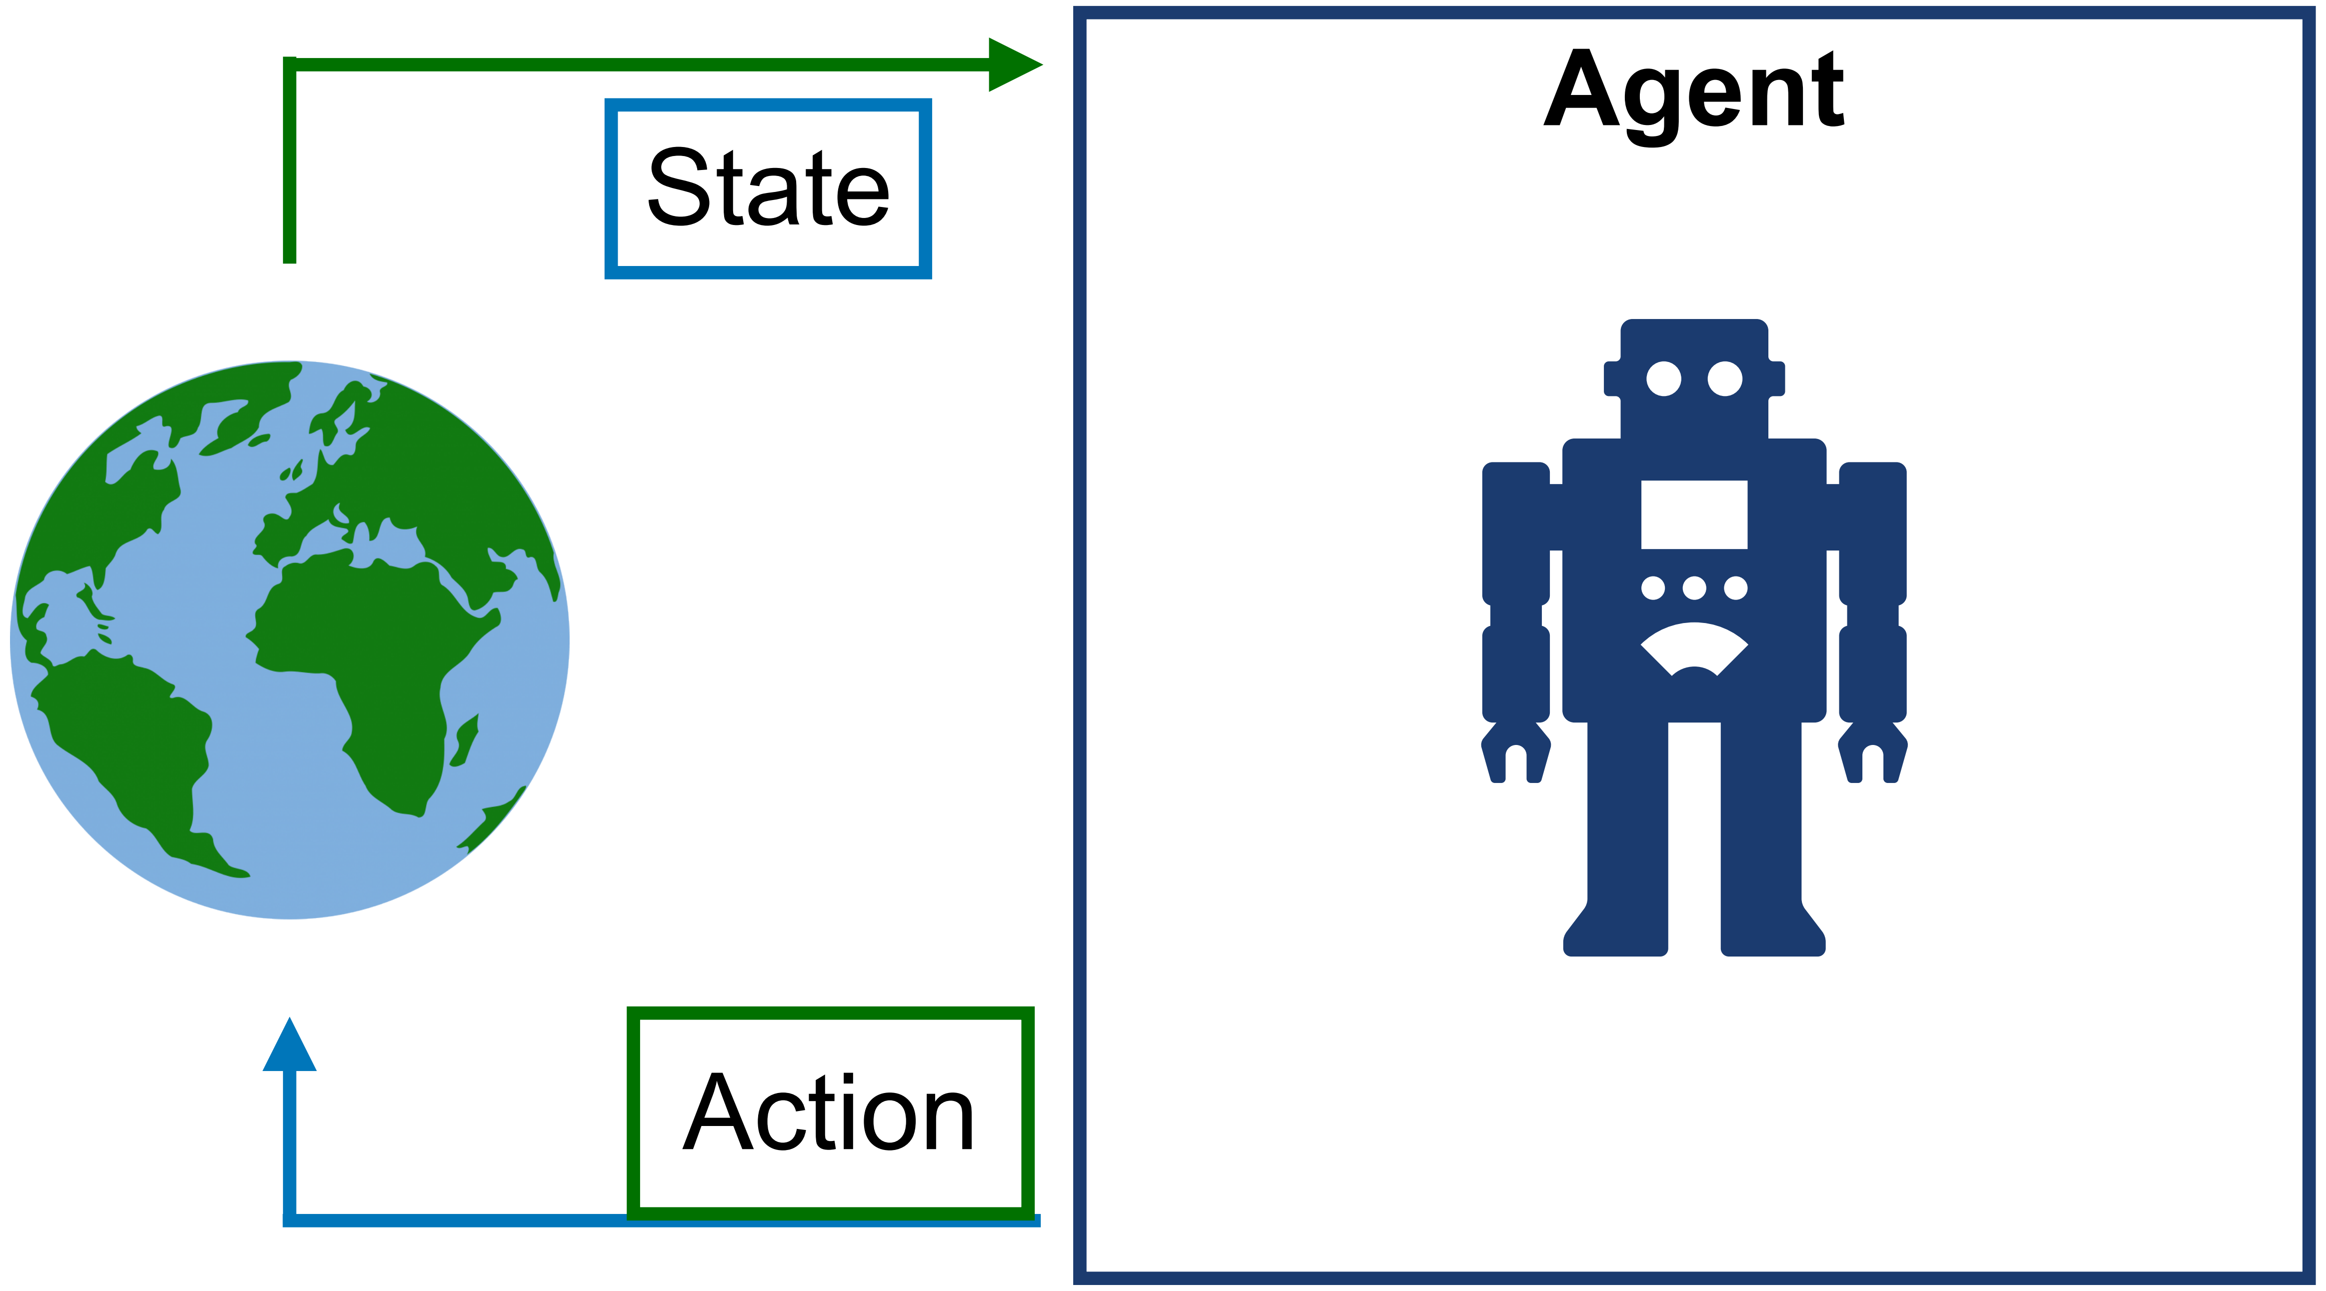

0) Proximal Flows and Neural Networks

1) Sparse reconstruction

- ☑ Optimization-based explanation
- ☑ Competitive positive firing rate dynamics

2) Policy Composition

- ☑ Free energy minimization
- ☑ Softmax flows from entropy max.

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Consider a composite minimization problem

$$\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$$

$f(x, u)$ is convex and differentiable in x

$g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex, closed, and proper (ccp)

Consider a composite minimization problem

$$\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$$

$f(x, u)$ is convex and differentiable in x

$g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex, closed, and proper (ccp)

Equivalence property:

1. $x^\star$ is minimizer for : $\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$
2. $x^\star$ is fixed point for : $x = \mathrm{prox}_{\gamma g}(x - \gamma \nabla f(x, u))$ for all γ

Consider a composite minimization problem

$$\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$$

$f(x, u)$ is convex and differentiable in x

$g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex, closed, and proper (ccp)

Equivalence property:

1. $x^\star$ is minimizer for : $\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$
2. $x^\star$ is fixed point for : $x = \mathrm{prox}_{\gamma g}(x - \gamma \nabla f(x, u))$ for all γ

Proximal Operator:

$$\mathrm{prox}_{\gamma g}(z) := \arg\min_{x \in \mathbb{R}^n} g(x) + \frac{1}{2\gamma}\|x - z\|_2^2$$

The Proximal Gradient Flow

Consider a composite minimization problem

$$\min_{x \in \mathbb{R}^n} \; f(x, u) \; + \; g(x)$$

$f(x, u)$ is convex and differentiable in x

$g : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex, closed, and proper (ccp)

Equivalence property:

1. $x^\star$ is minimizer for : $\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$
2. $x^\star$ is fixed point for : $x = \mathrm{prox}_{\gamma g}(x - \gamma \nabla f(x, u))$ for all γ

Proximal gradient dynamics: $\dot{x} = -x + \mathrm{prox}_{\gamma g}(x - \gamma \nabla f(x, u))$

S. Hassan-Moghaddam, M. R. Jovanovic, *"Proximal gradient flow and Douglas-Rachford splitting dynamics: Global exponential stability via integral quadratic constraints"*, *Automatica*, 2021.

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

$$x^\star = \arg\min_{x \in \mathbb{R}^n} f(x, u) + g(x)$$

$$\Longleftrightarrow$$

$x^\star$ is an equilibrium point of:

$$\dot{x} = -x + \mathrm{prox}_g(x - \nabla f(x, u))$$

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

$$x^\star = \arg\min_{x \in \mathbb{R}^n} f(x,u) + g(x) \iff$$

$x^\star$ is an equilibrium point of:

$$\dot{x} = -x + \mathrm{prox}_g(x - \nabla f(x,u))$$

$$\downarrow$$

$$\dot{x} = -x + \mathrm{prox}_g(Wx + Bu)$$

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Proximal Flows as Neural Networks

$$x^\star = \arg\min_{x \in \mathbb{R}^n} f(x,u) + g(x) \quad \Longleftrightarrow$$

$x^\star$ is an equilibrium point of:

$$\dot{x} = -x + \mathrm{prox}_g(x - \nabla f(x,u))$$

$$\dot{x} = -x + \mathrm{prox}_g(Wx + Bu)$$

Firing-Rate Neural Network!

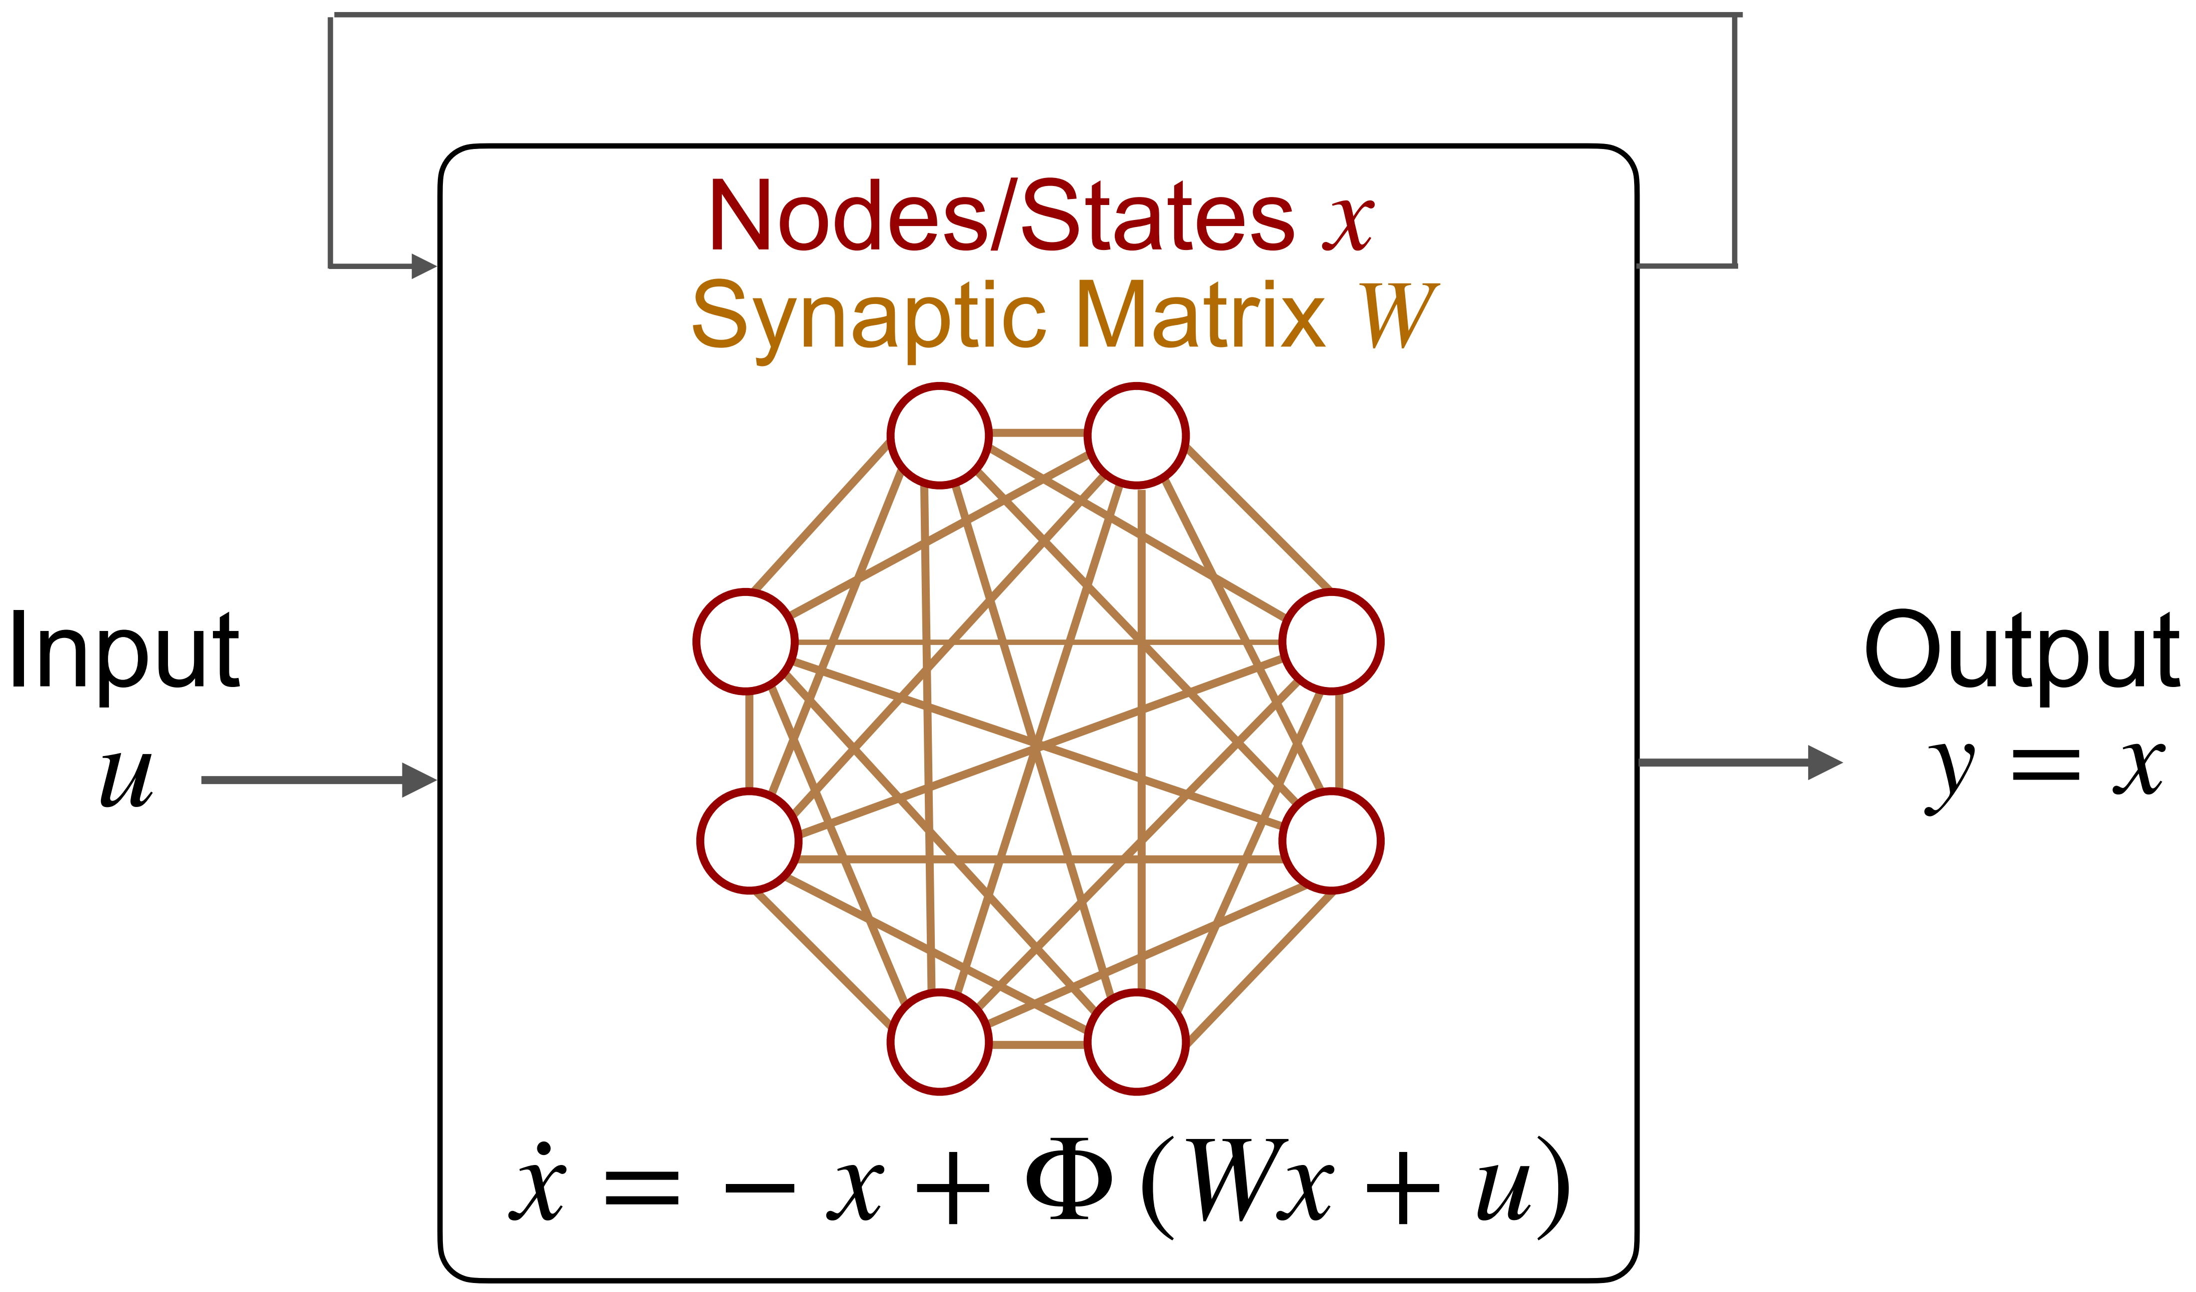

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Firing-Rate Neural Network

$$\dot{x} = -x + \Phi(Wx + u) := f_{\mathrm{F}}(x, u)$$

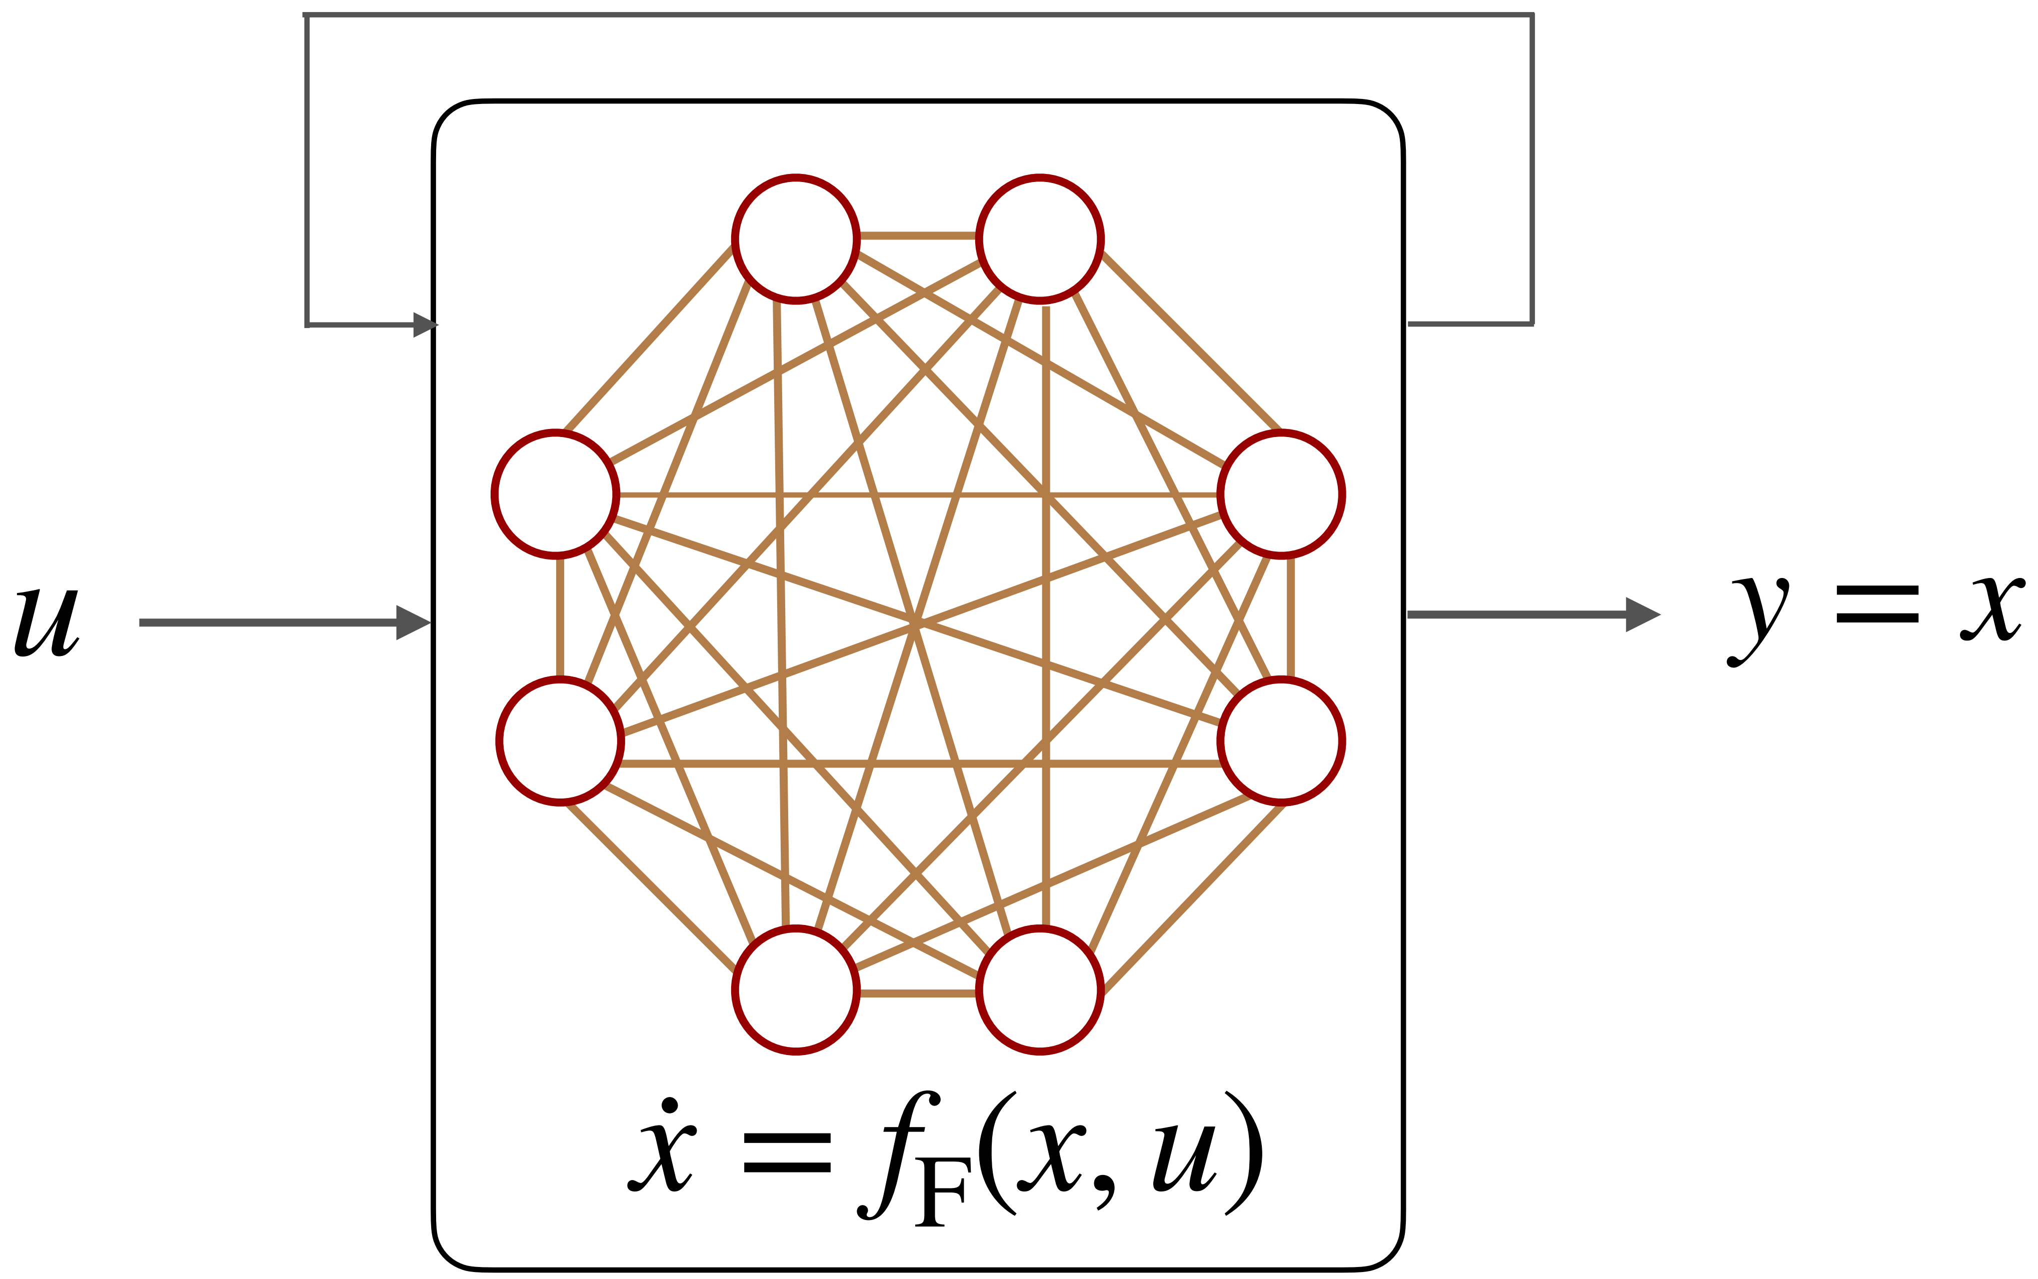

Firing-Rate Neural Network

$$\dot{x} = -\boxed{x} + \Phi(W x + u) := f_{\mathrm{F}}(x, u)$$

state

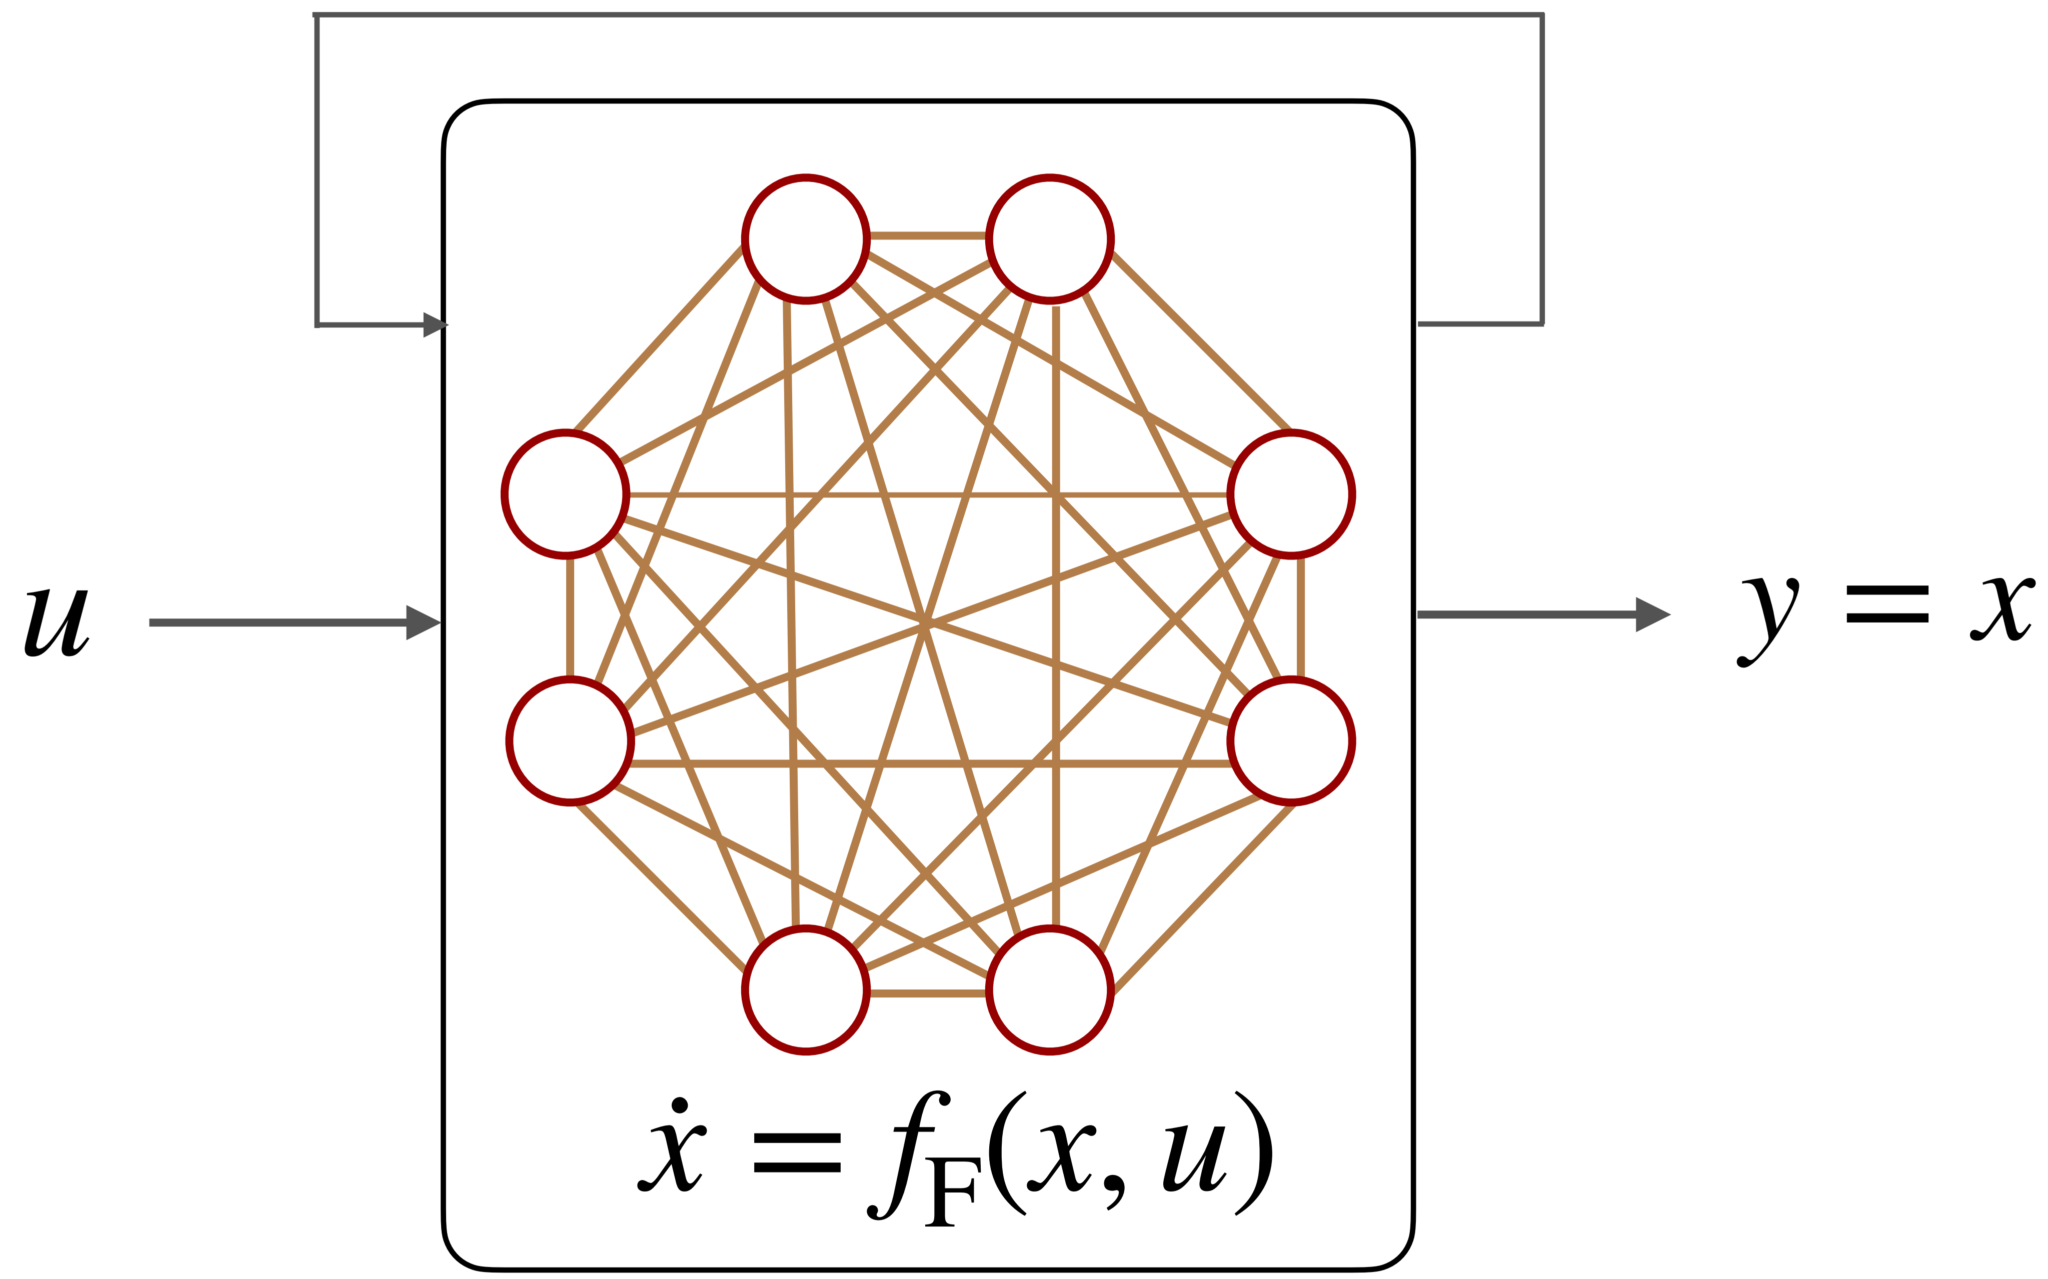

Firing-Rate Neural Network

$$\dot{x} = -\boxed{x} + \boxed{\Phi}(\ W\ x + \ u) := f_{\mathrm{F}}(x, u)$$

state activation function

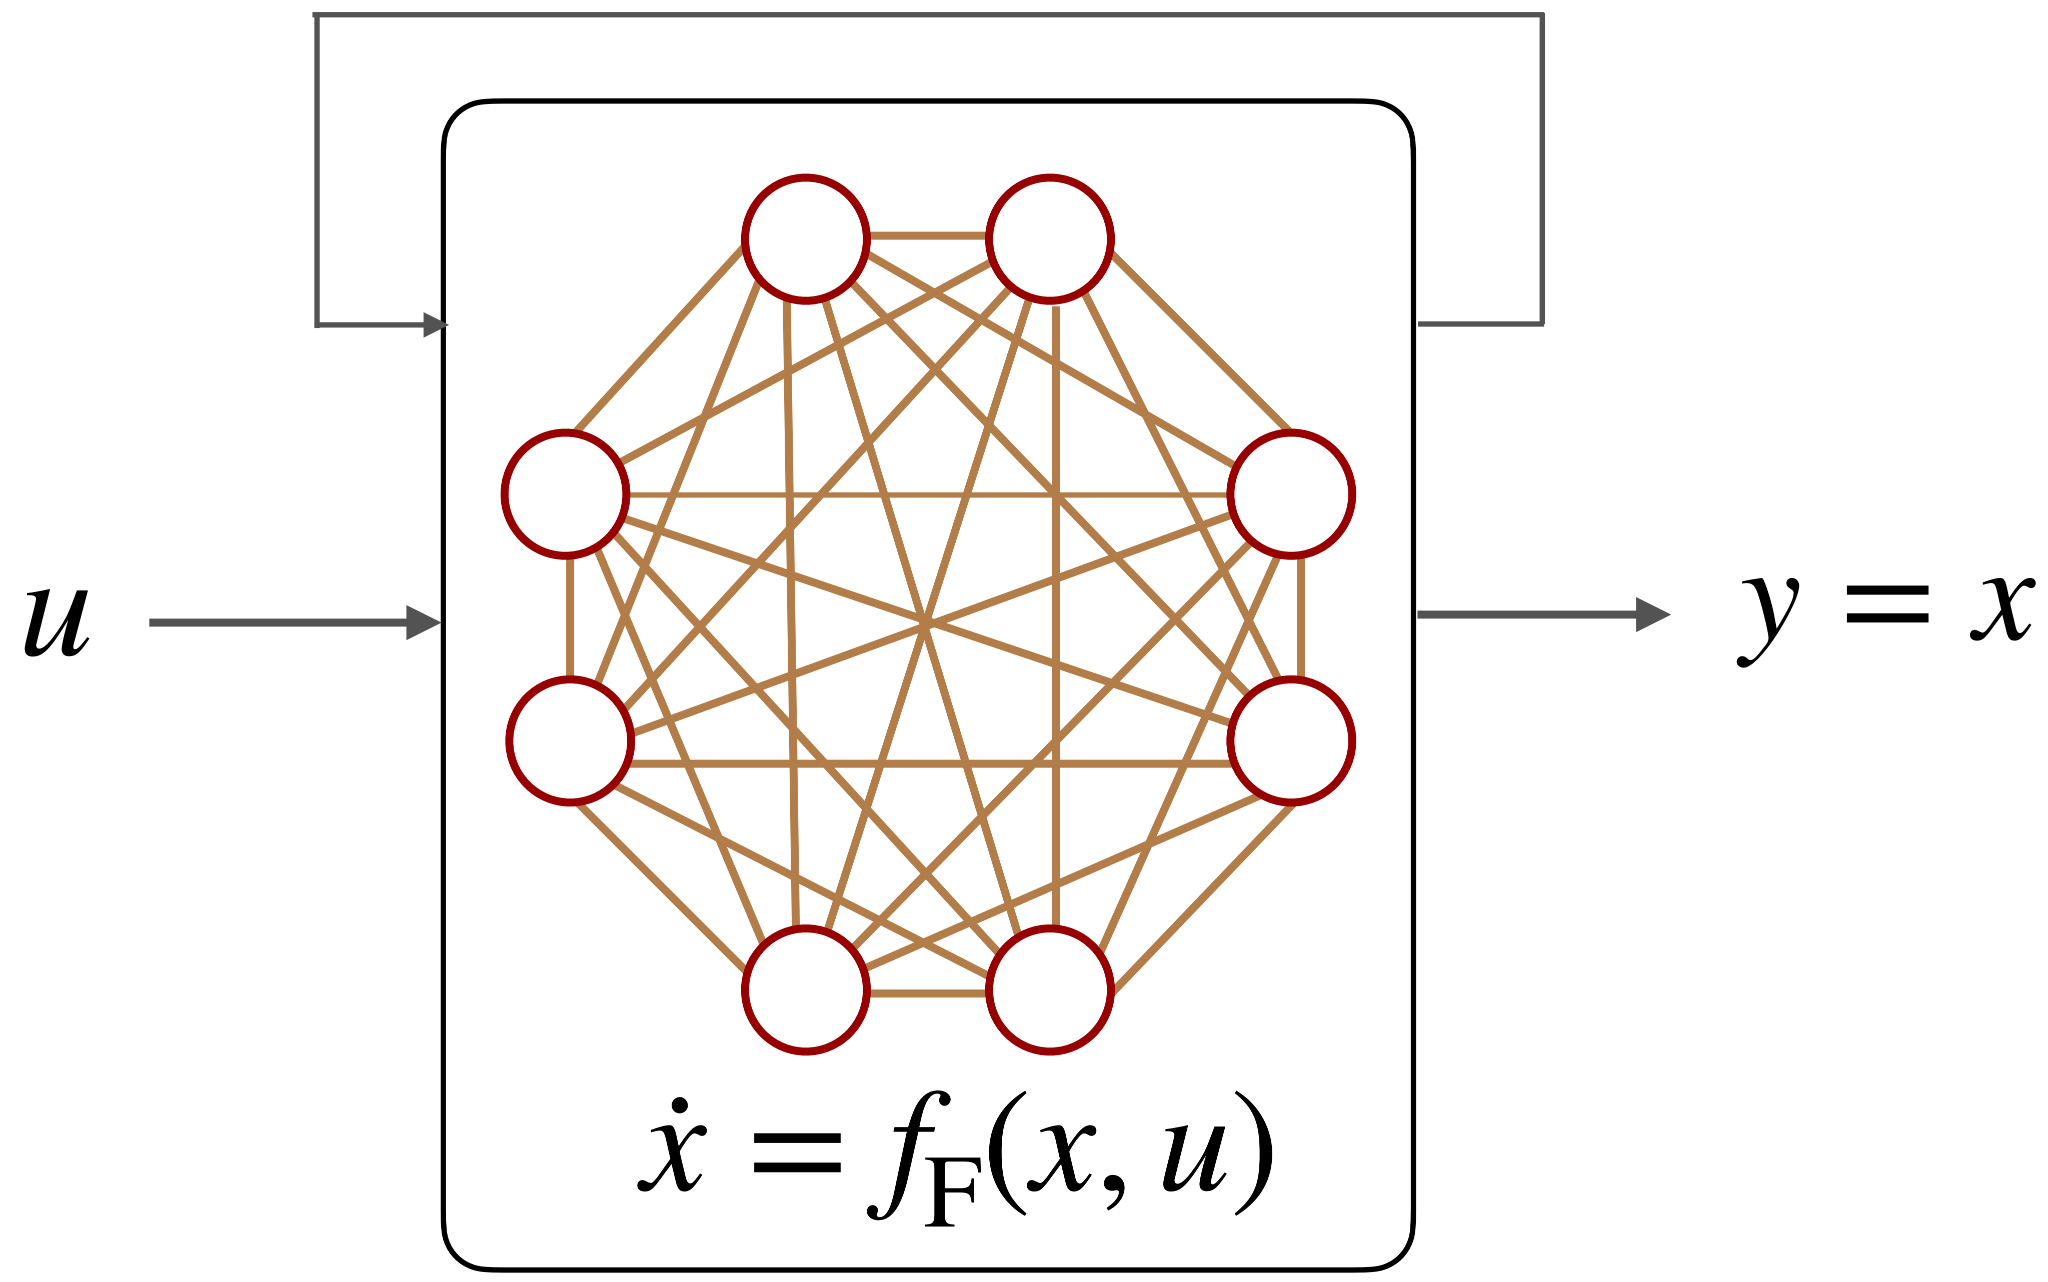

Firing-Rate Neural Network

$$\dot{x} = -\boxed{x} + \boxed{\Phi}(\boxed{W}x + u) := f_{\mathrm{F}}(x, u)$$

state activation function synaptic matrix

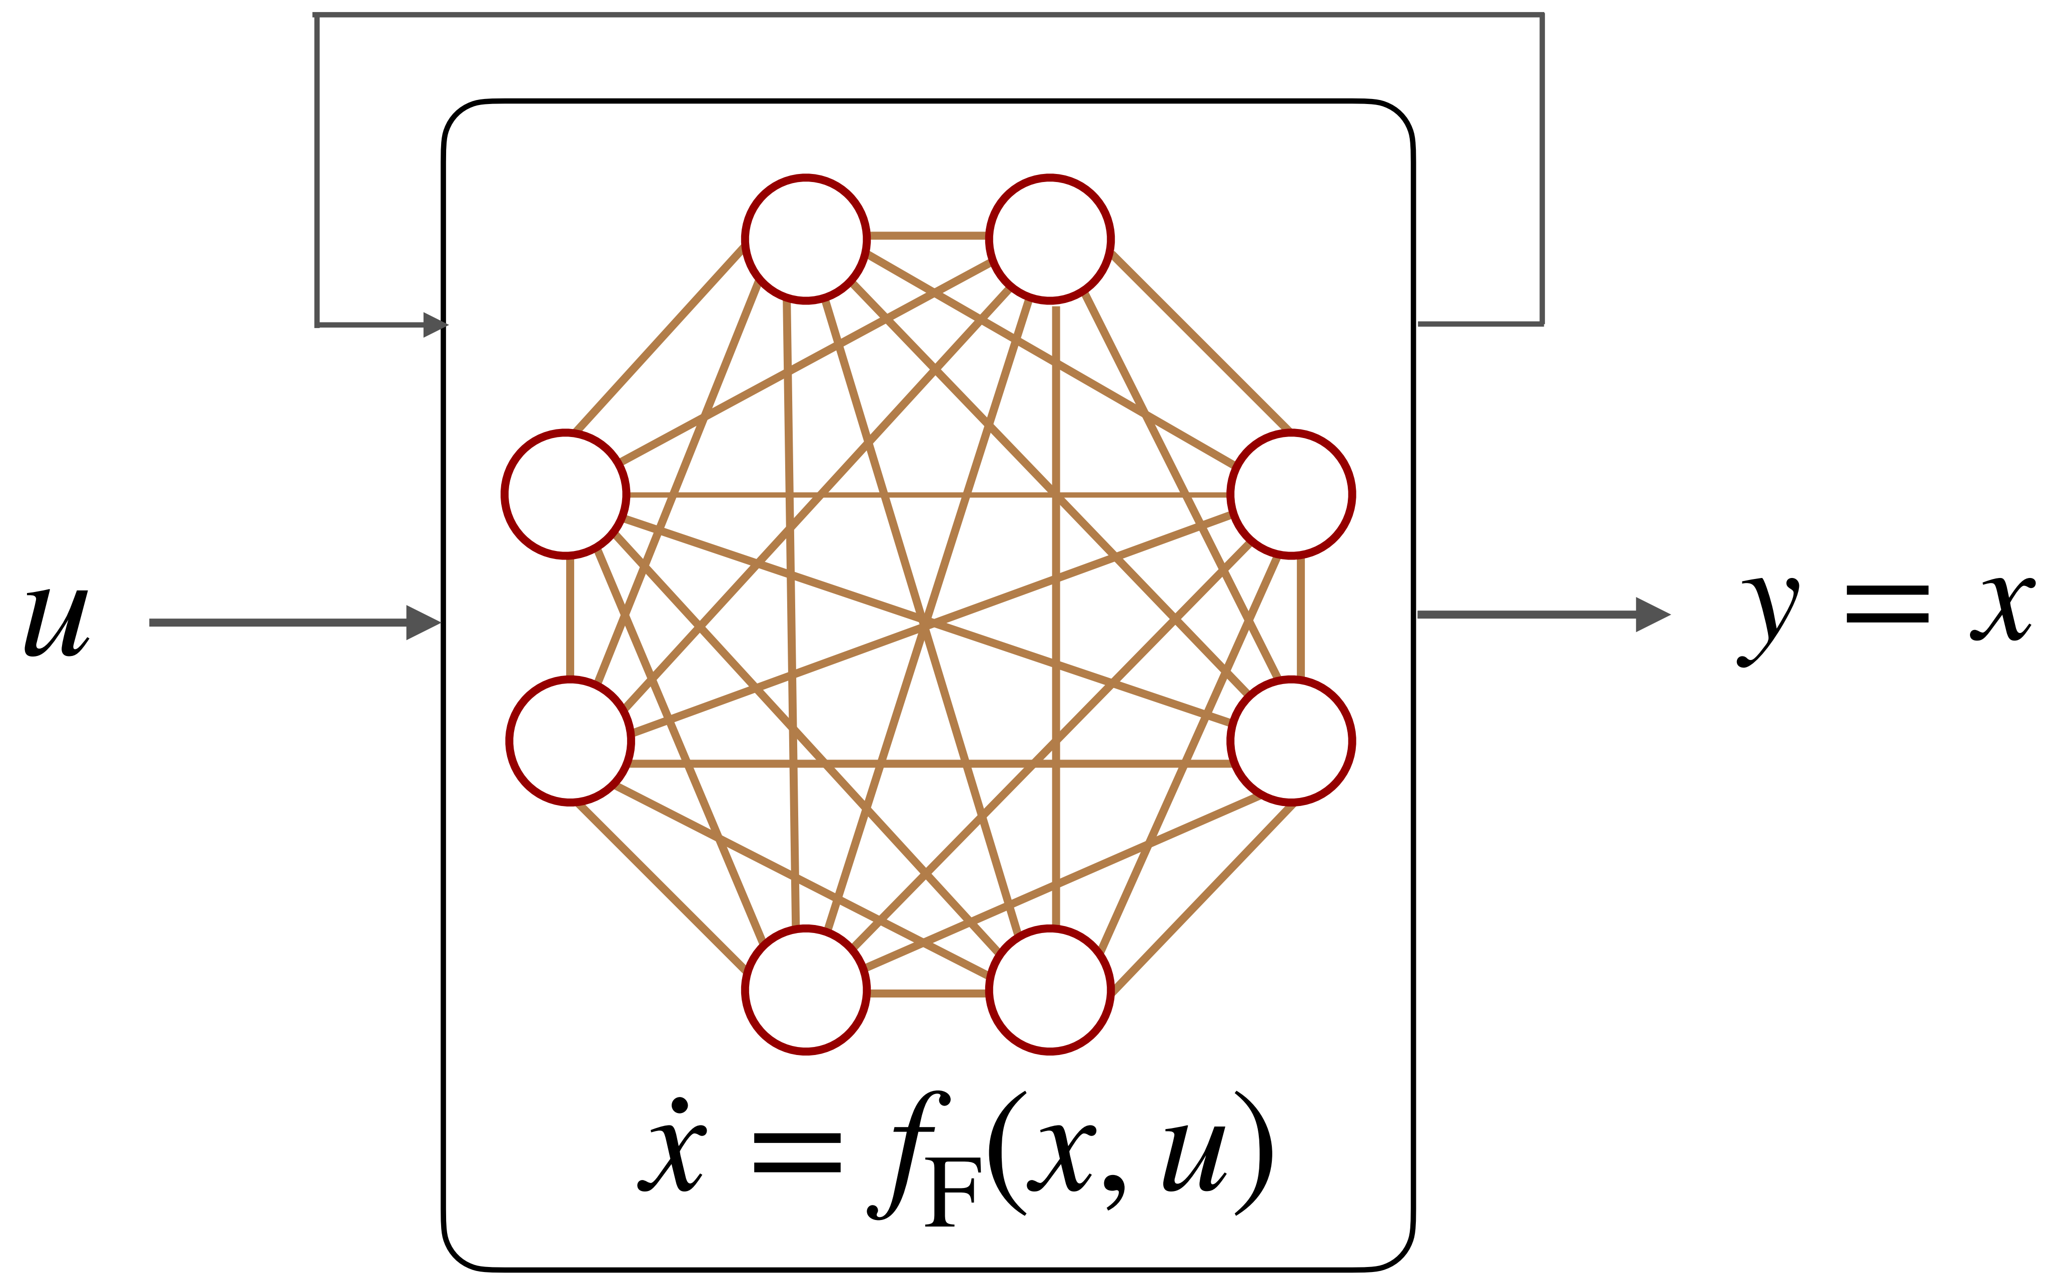

Firing-Rate Neural Network

$$\dot{x} = -\underbrace{\boxed{x}}_{\text{state}} + \underbrace{\boxed{\Phi}}_{\text{activation function}}(\underbrace{\boxed{W}}_{\text{synaptic matrix}}x + \underbrace{\boxed{u}}_{\text{stimulus}}) := f_{\mathrm{F}}(x, u)$$

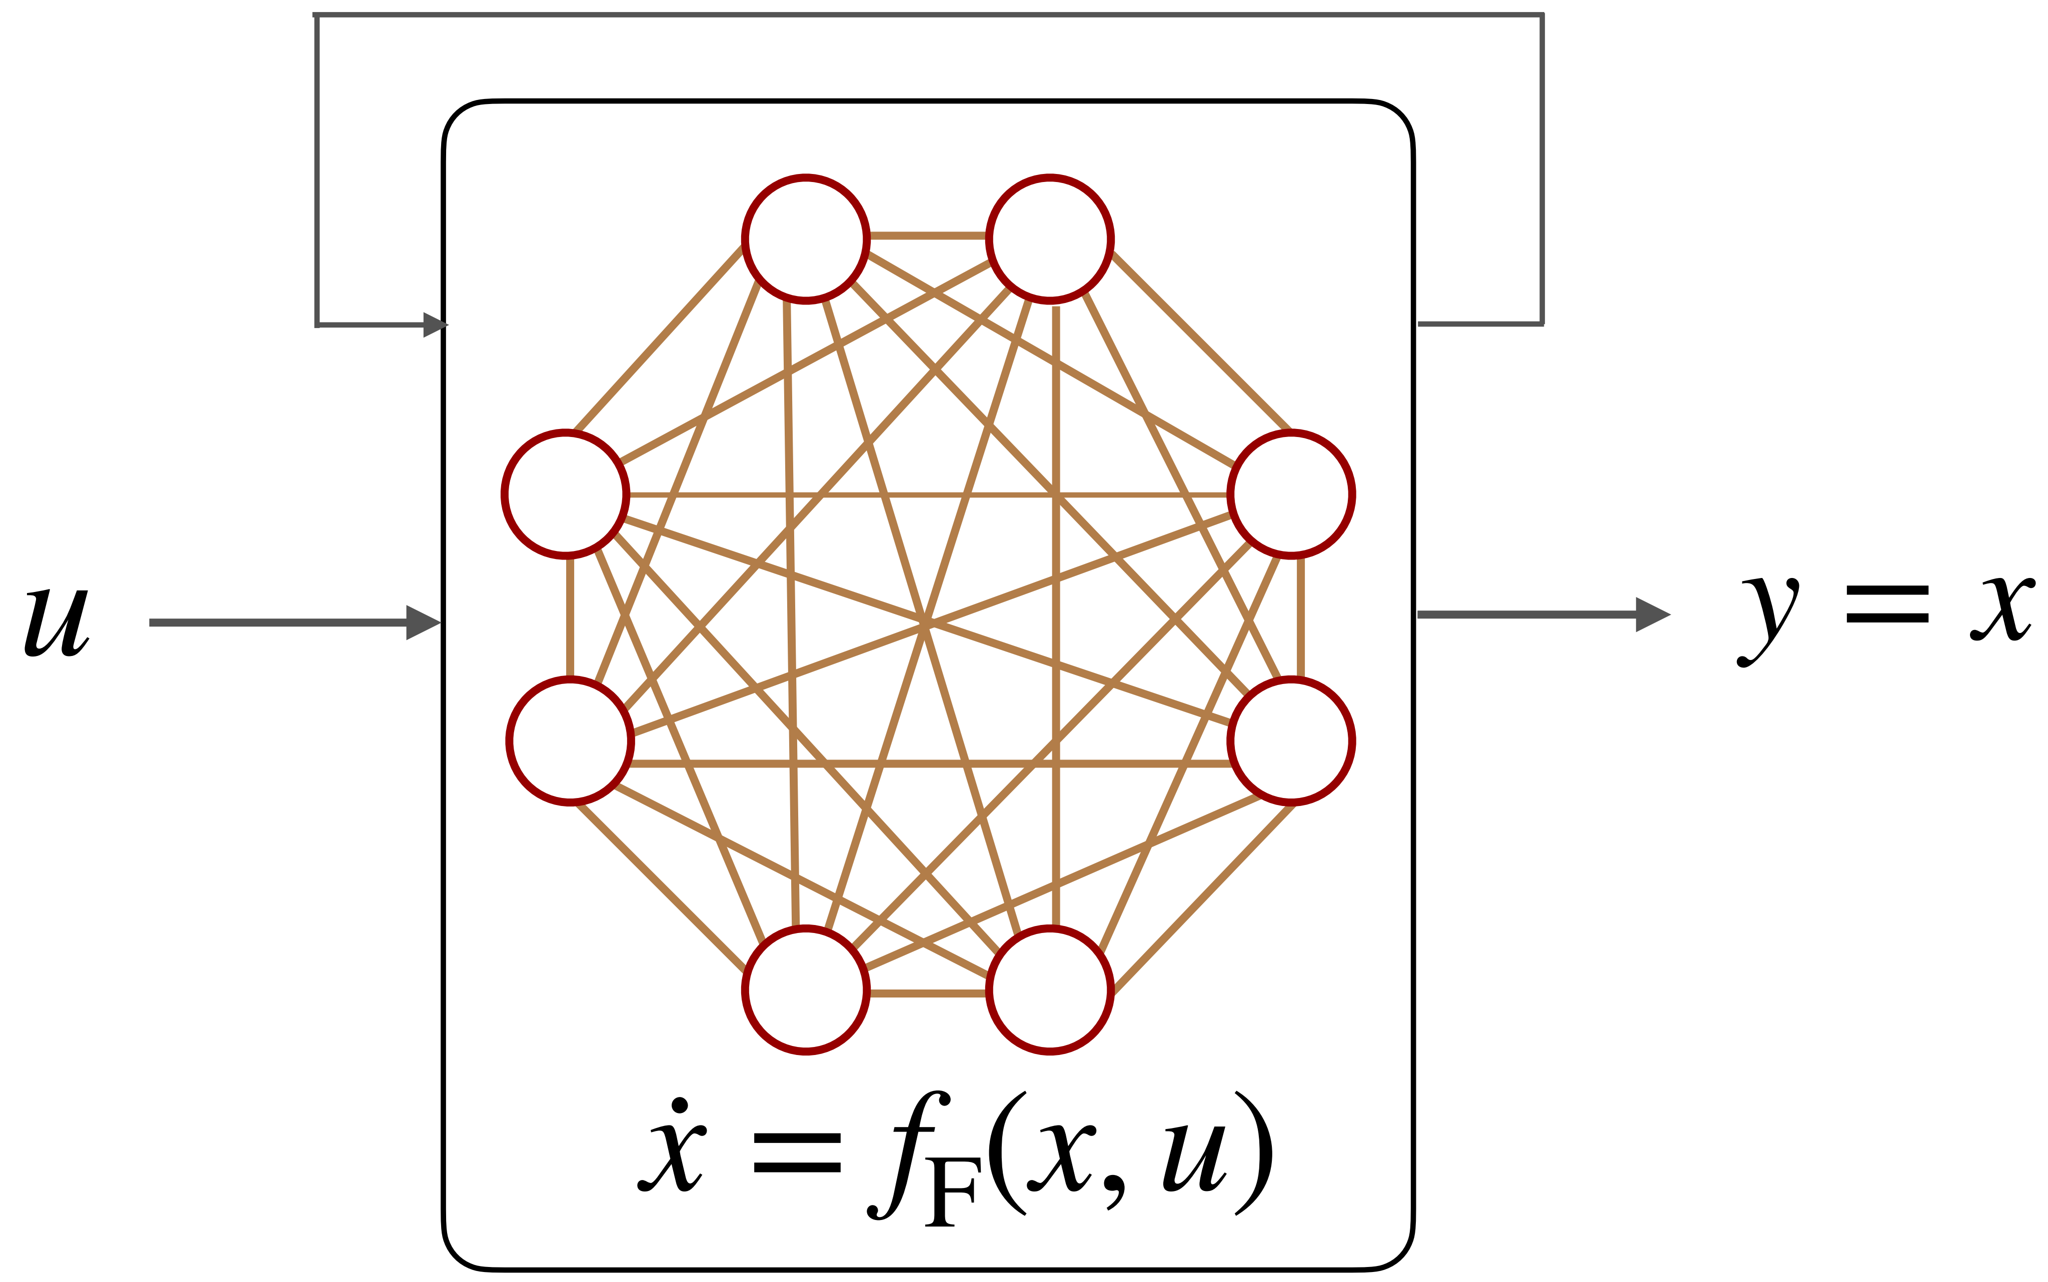

Firing-Rate Neural Network

$$\dot{x} = -\boxed{x} + \boxed{\Phi}(\boxed{W}x + \boxed{u}) := f_{\mathrm{F}}(x, u)$$

state   activation function   synaptic matrix   stimulus

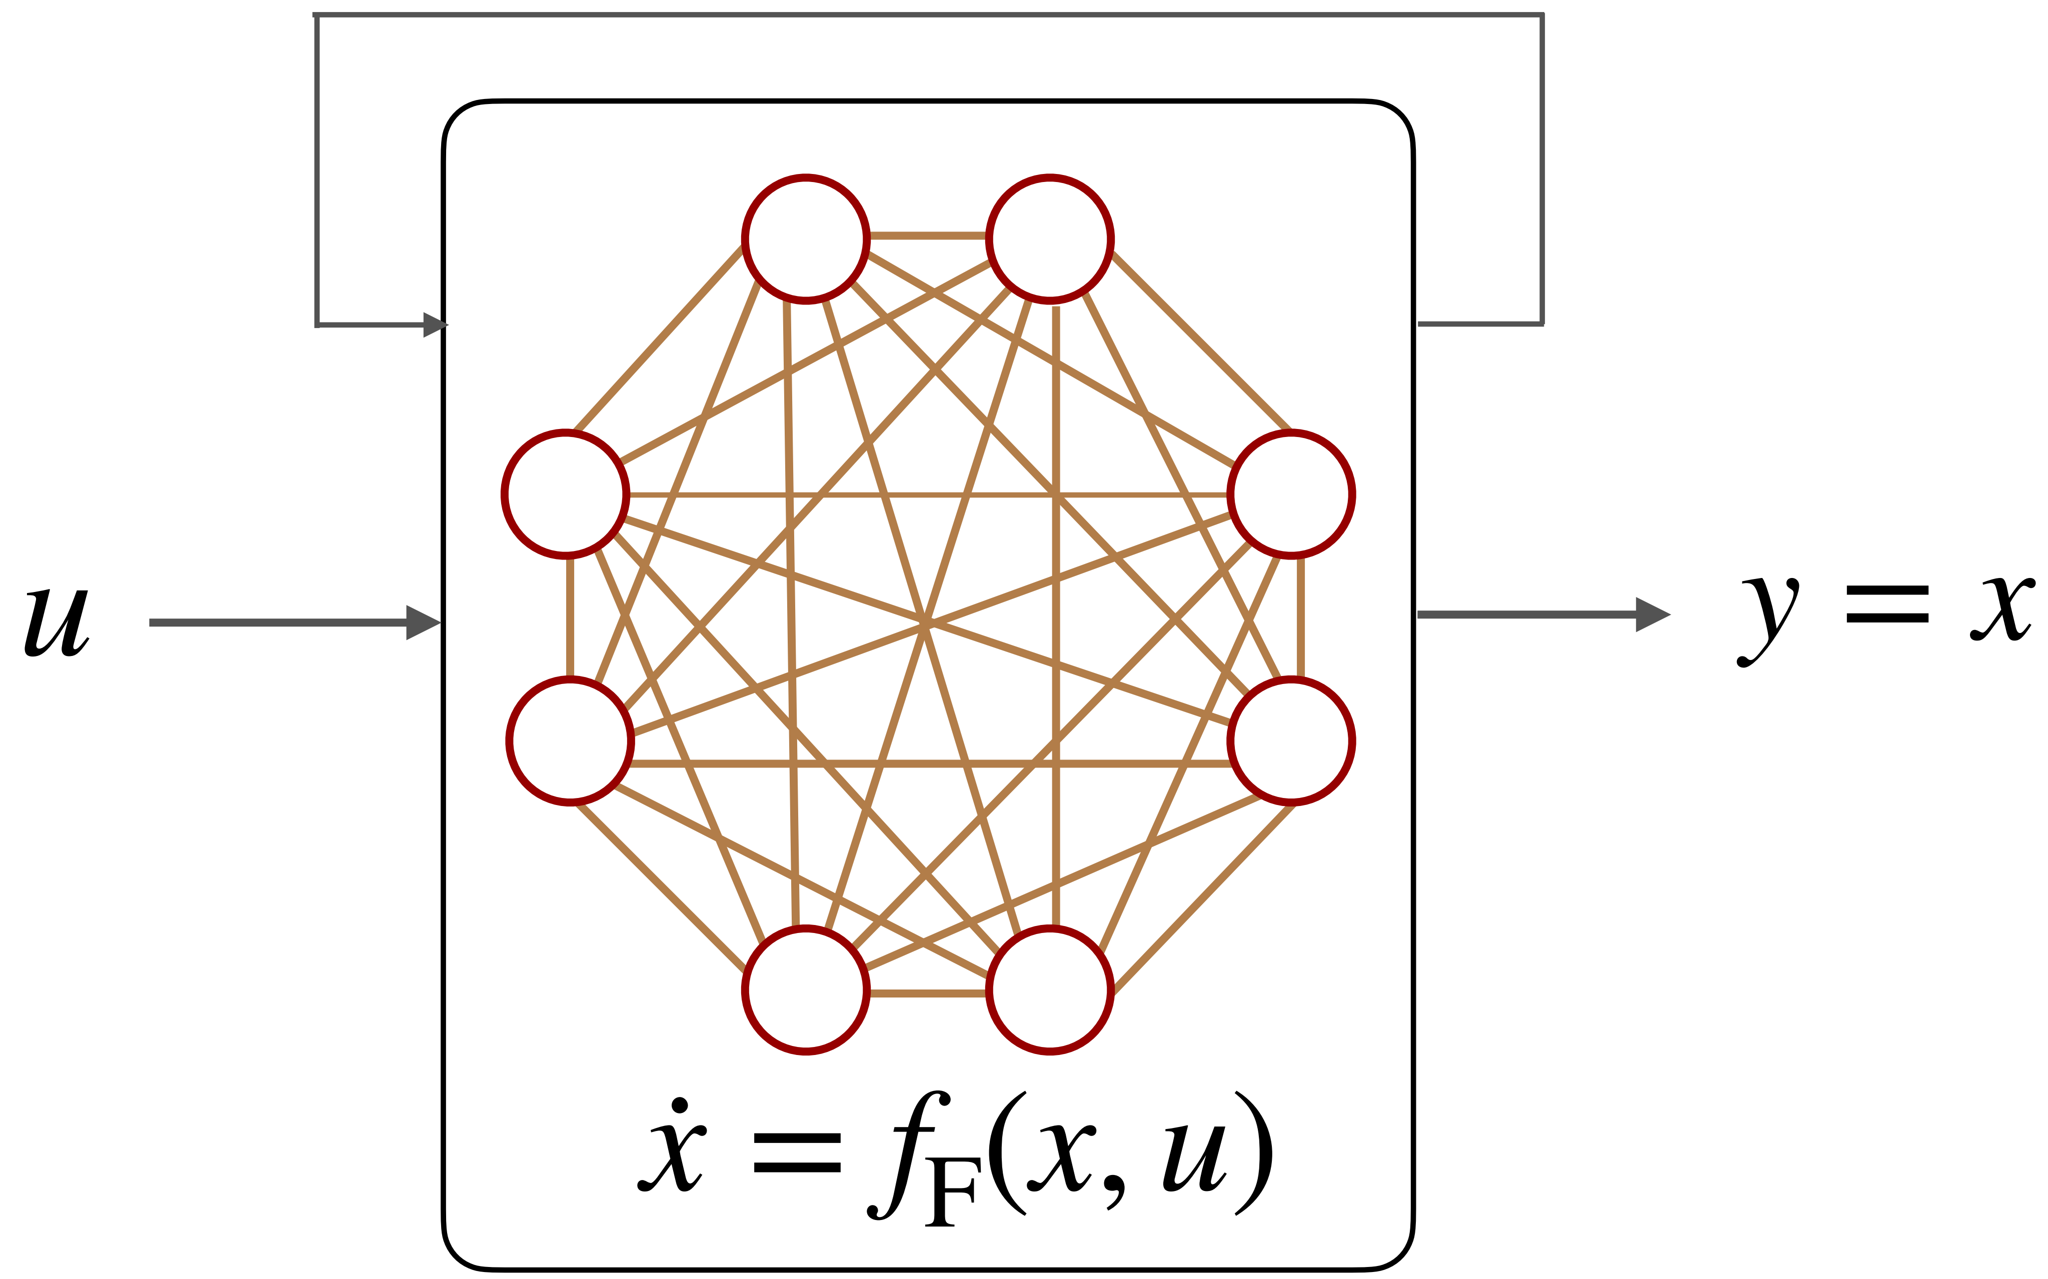

- FNN can be positive (when Φ is positive)
- Naturally models inhibition and excitation

Outline

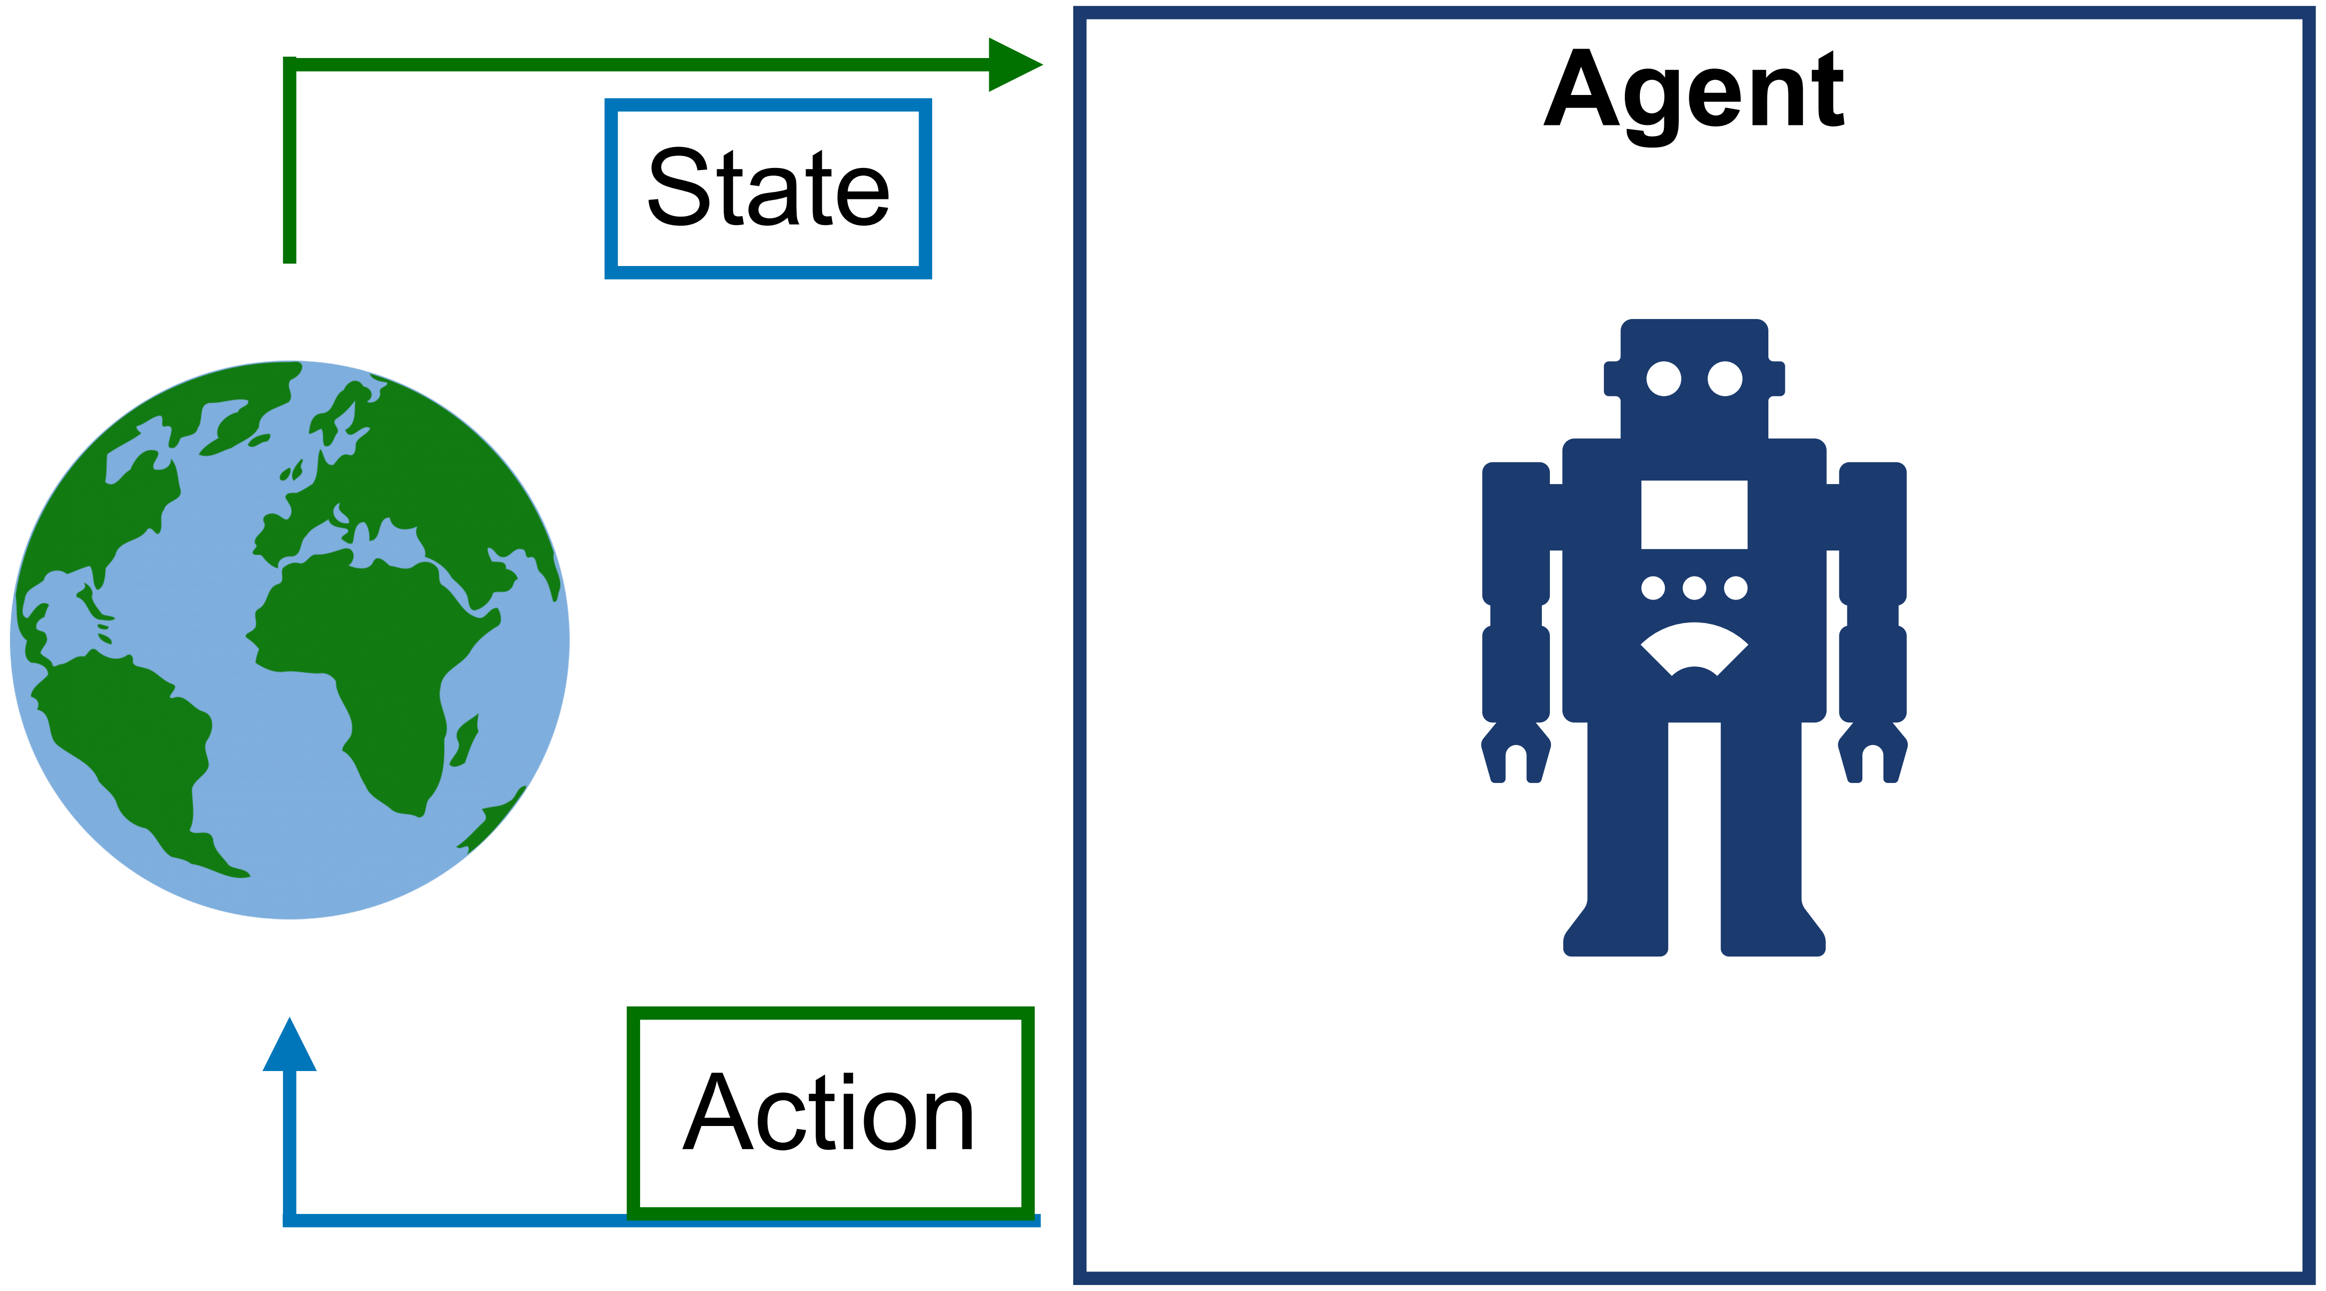

0) Proximal Flows and Neural Networks

1) Sparse reconstruction

- ☑ Optimization-based explanation
- ☑ Competitive positive firing rate dynamics

2) Policy Composition

- ☑ Free energy minimization
- ☑ Softmax flows from entropy max.

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Outline

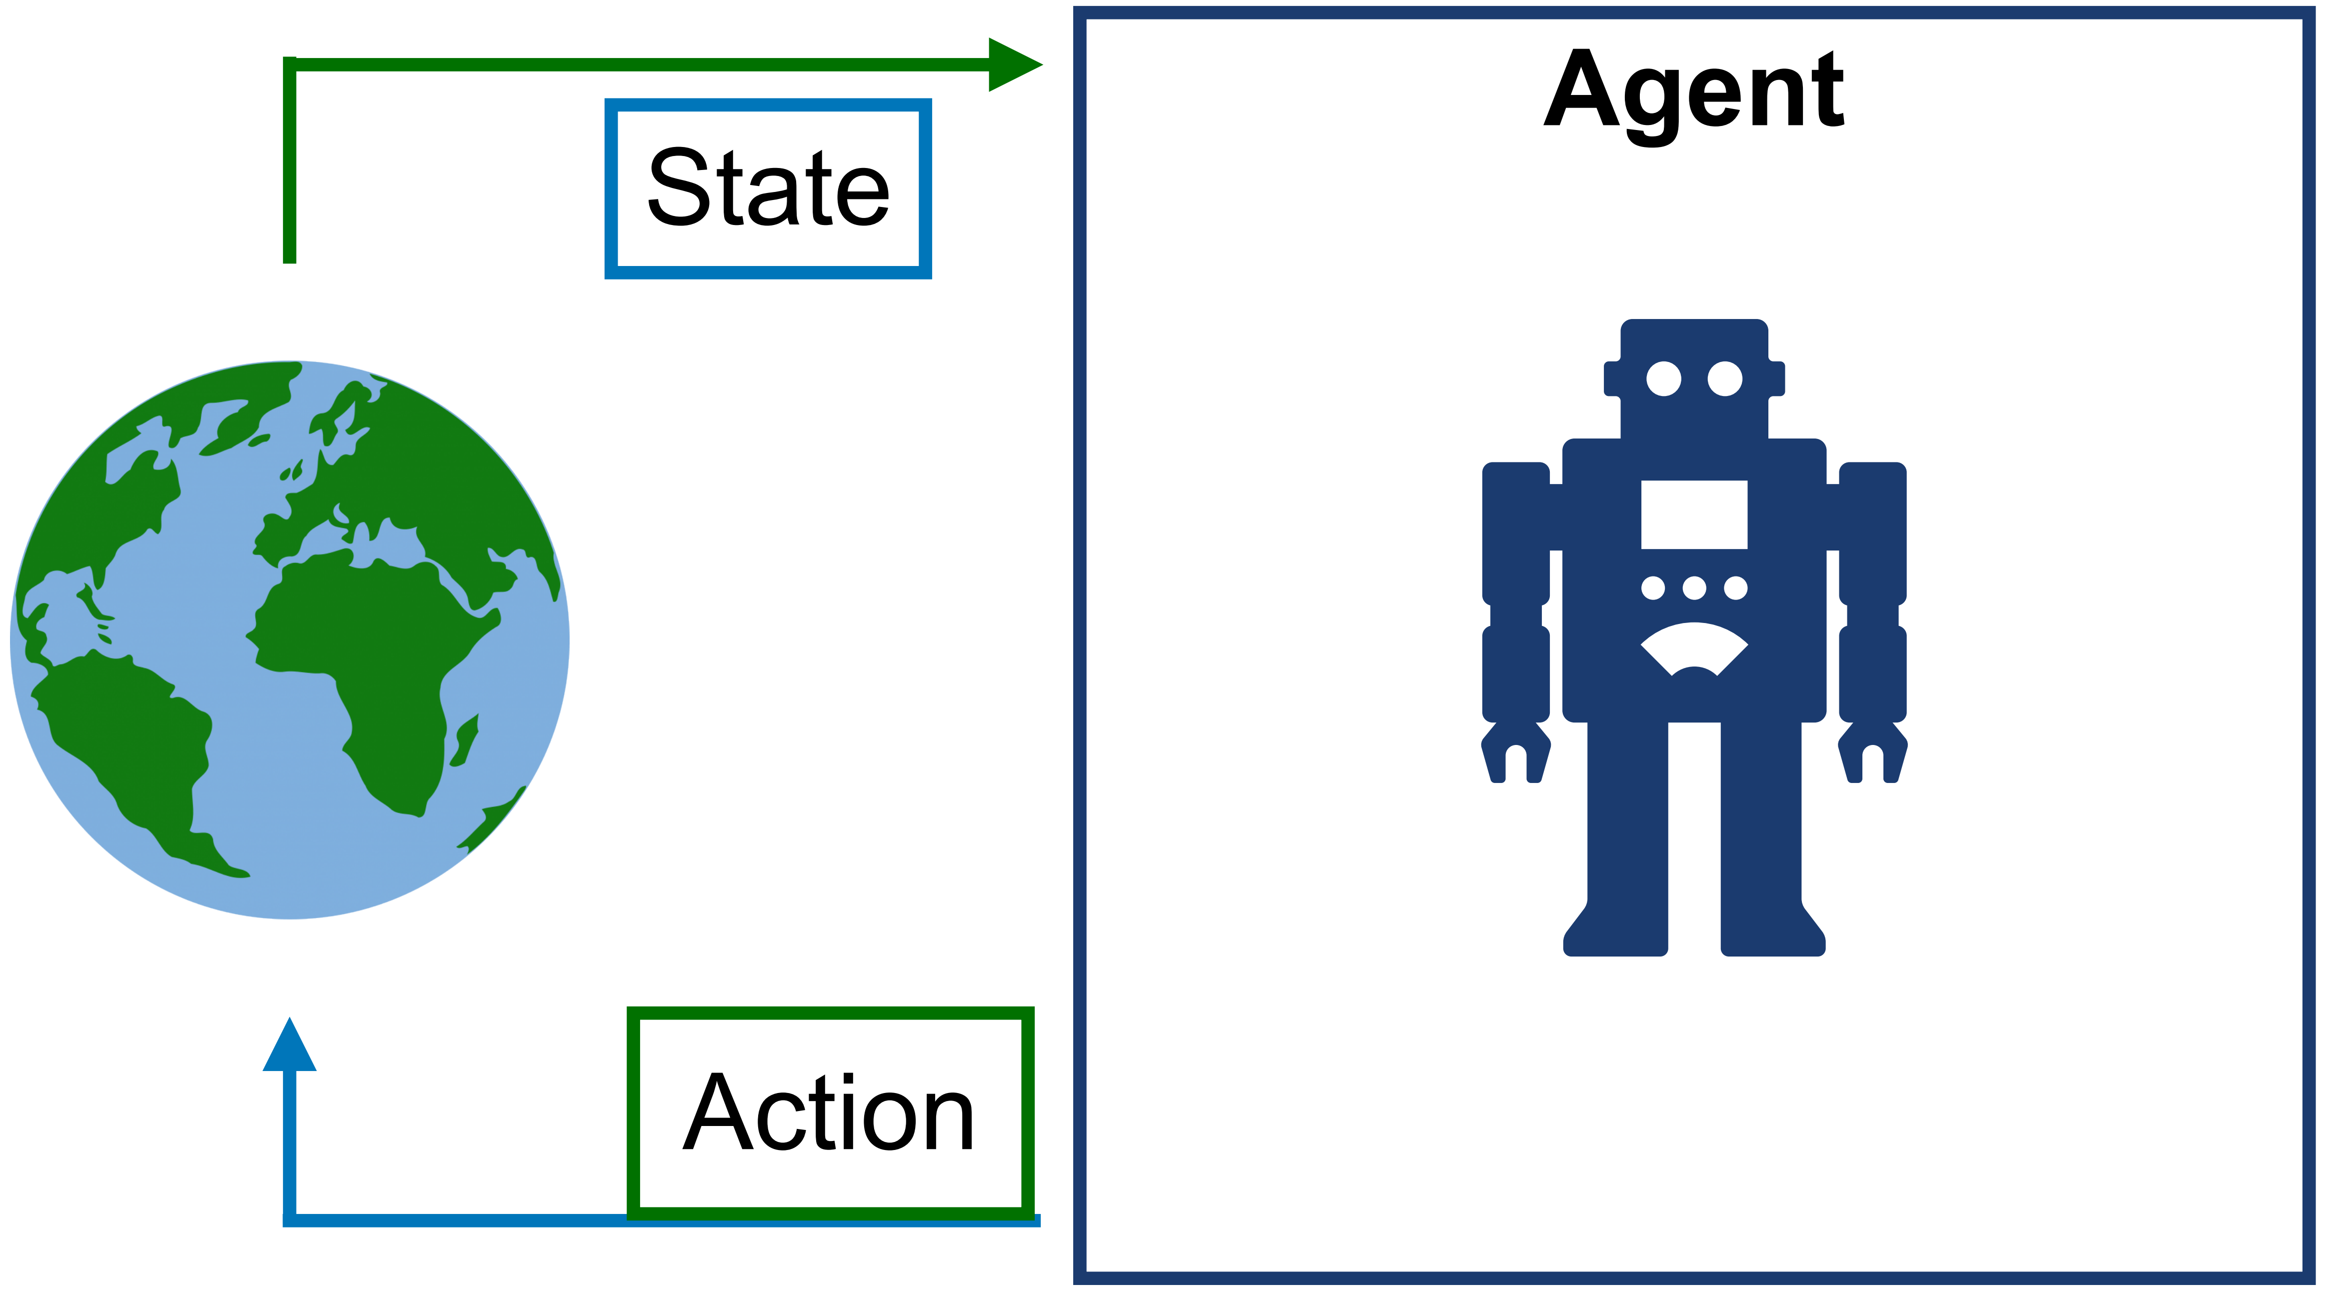

1) Sparse reconstruction

- ☑ Optimization-based explanation
- ☑ Competitive positive firing rate dynamics

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Input

u_1	u_2	$\cdots$	
		$\cdots$	u_m

Input image

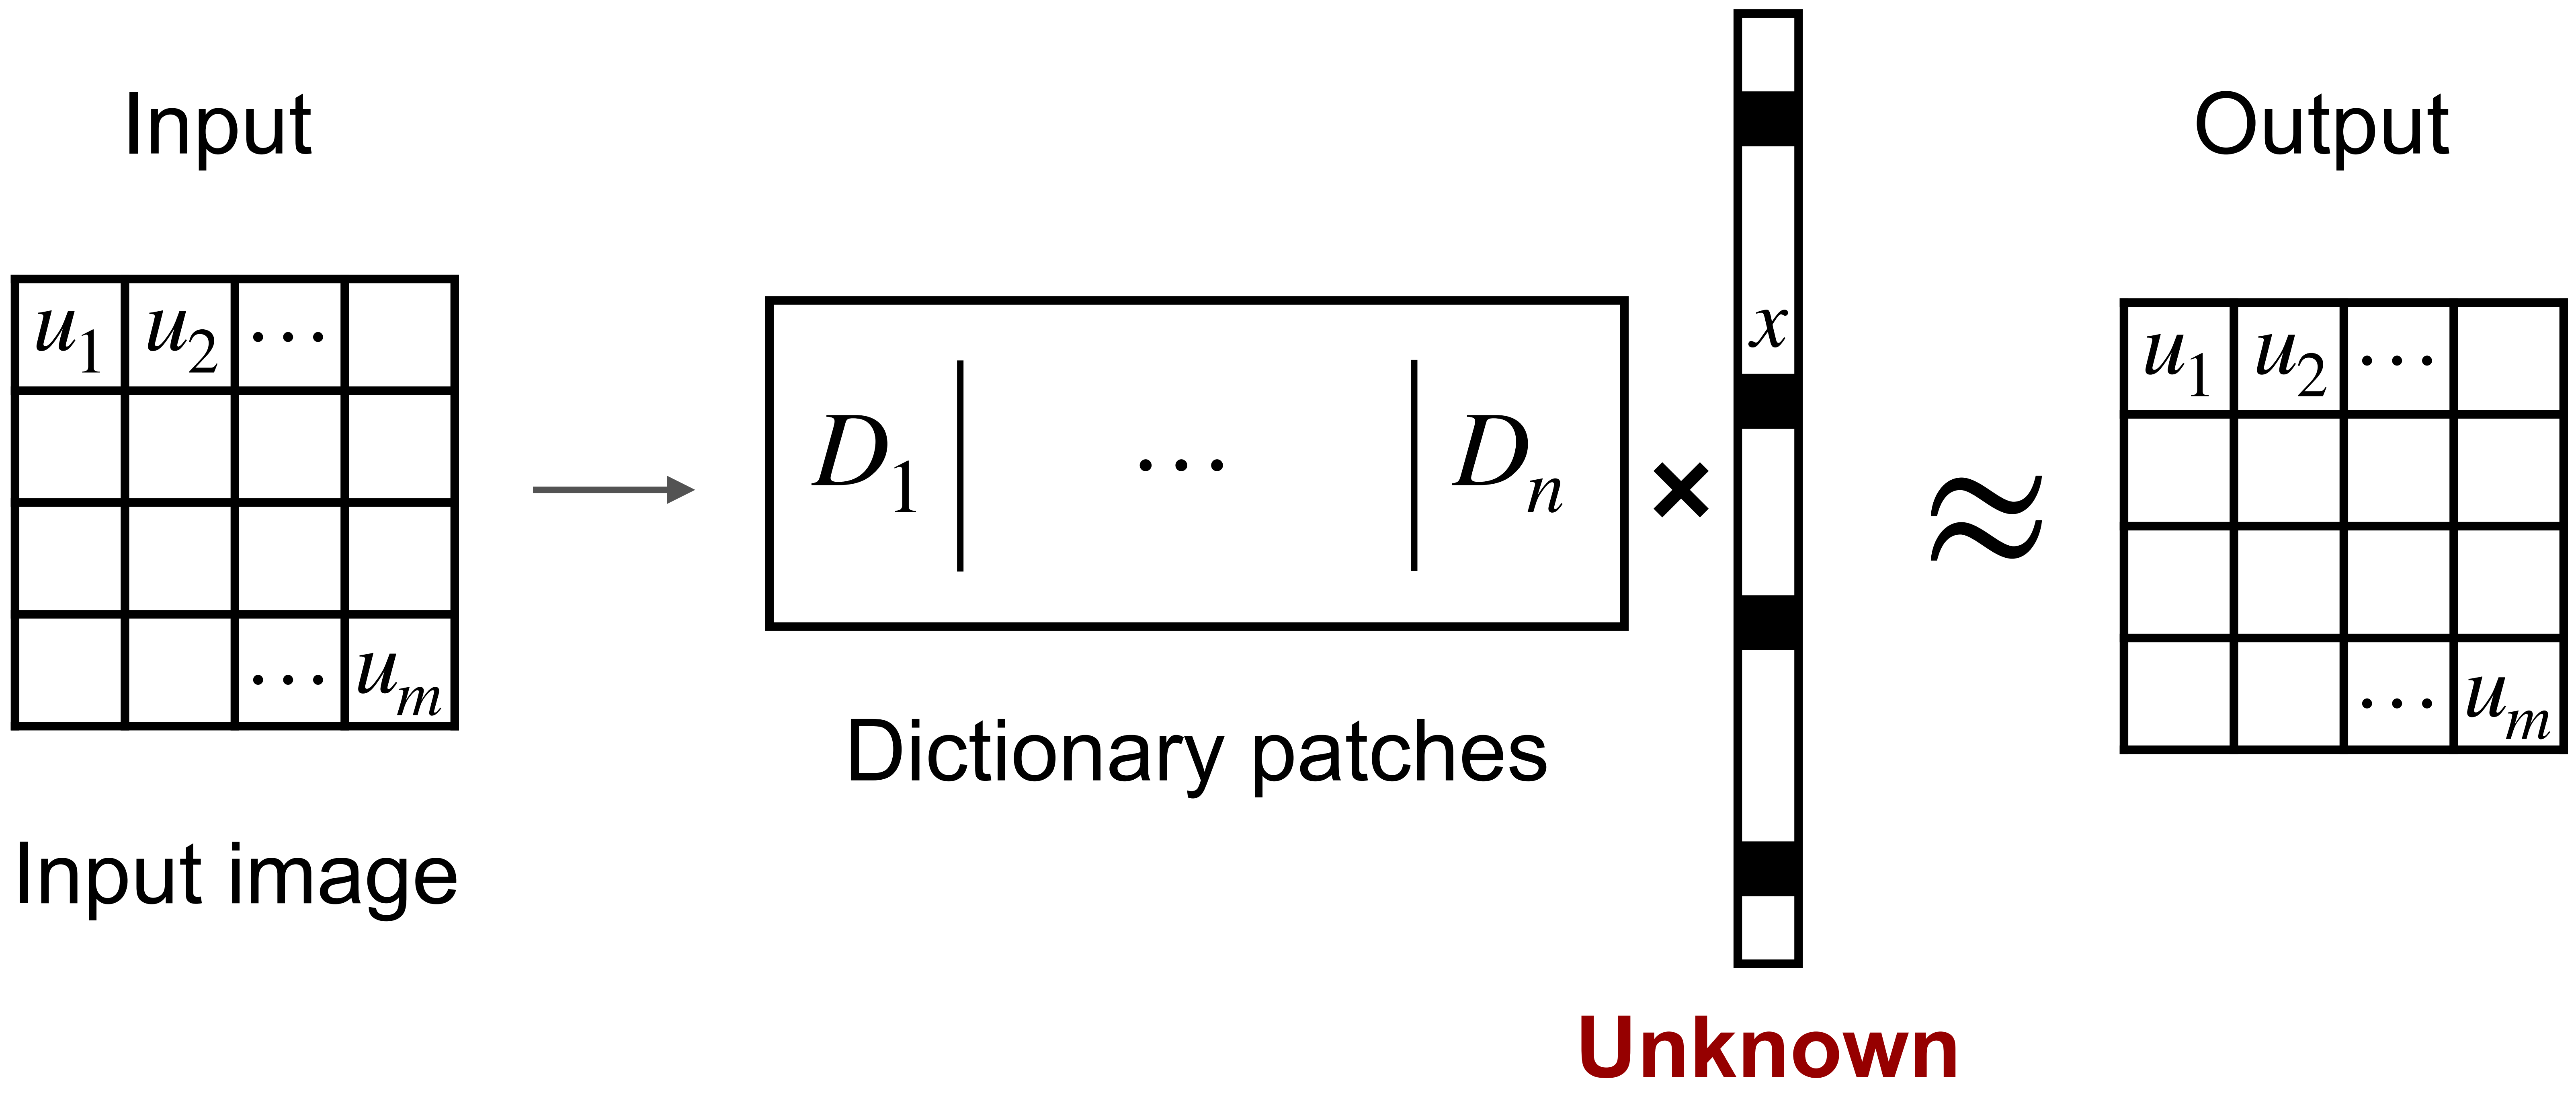

J. Wright and Y. Ma, *"High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications"*, Cambridge University Press, 2022.

Sparse Reconstruction

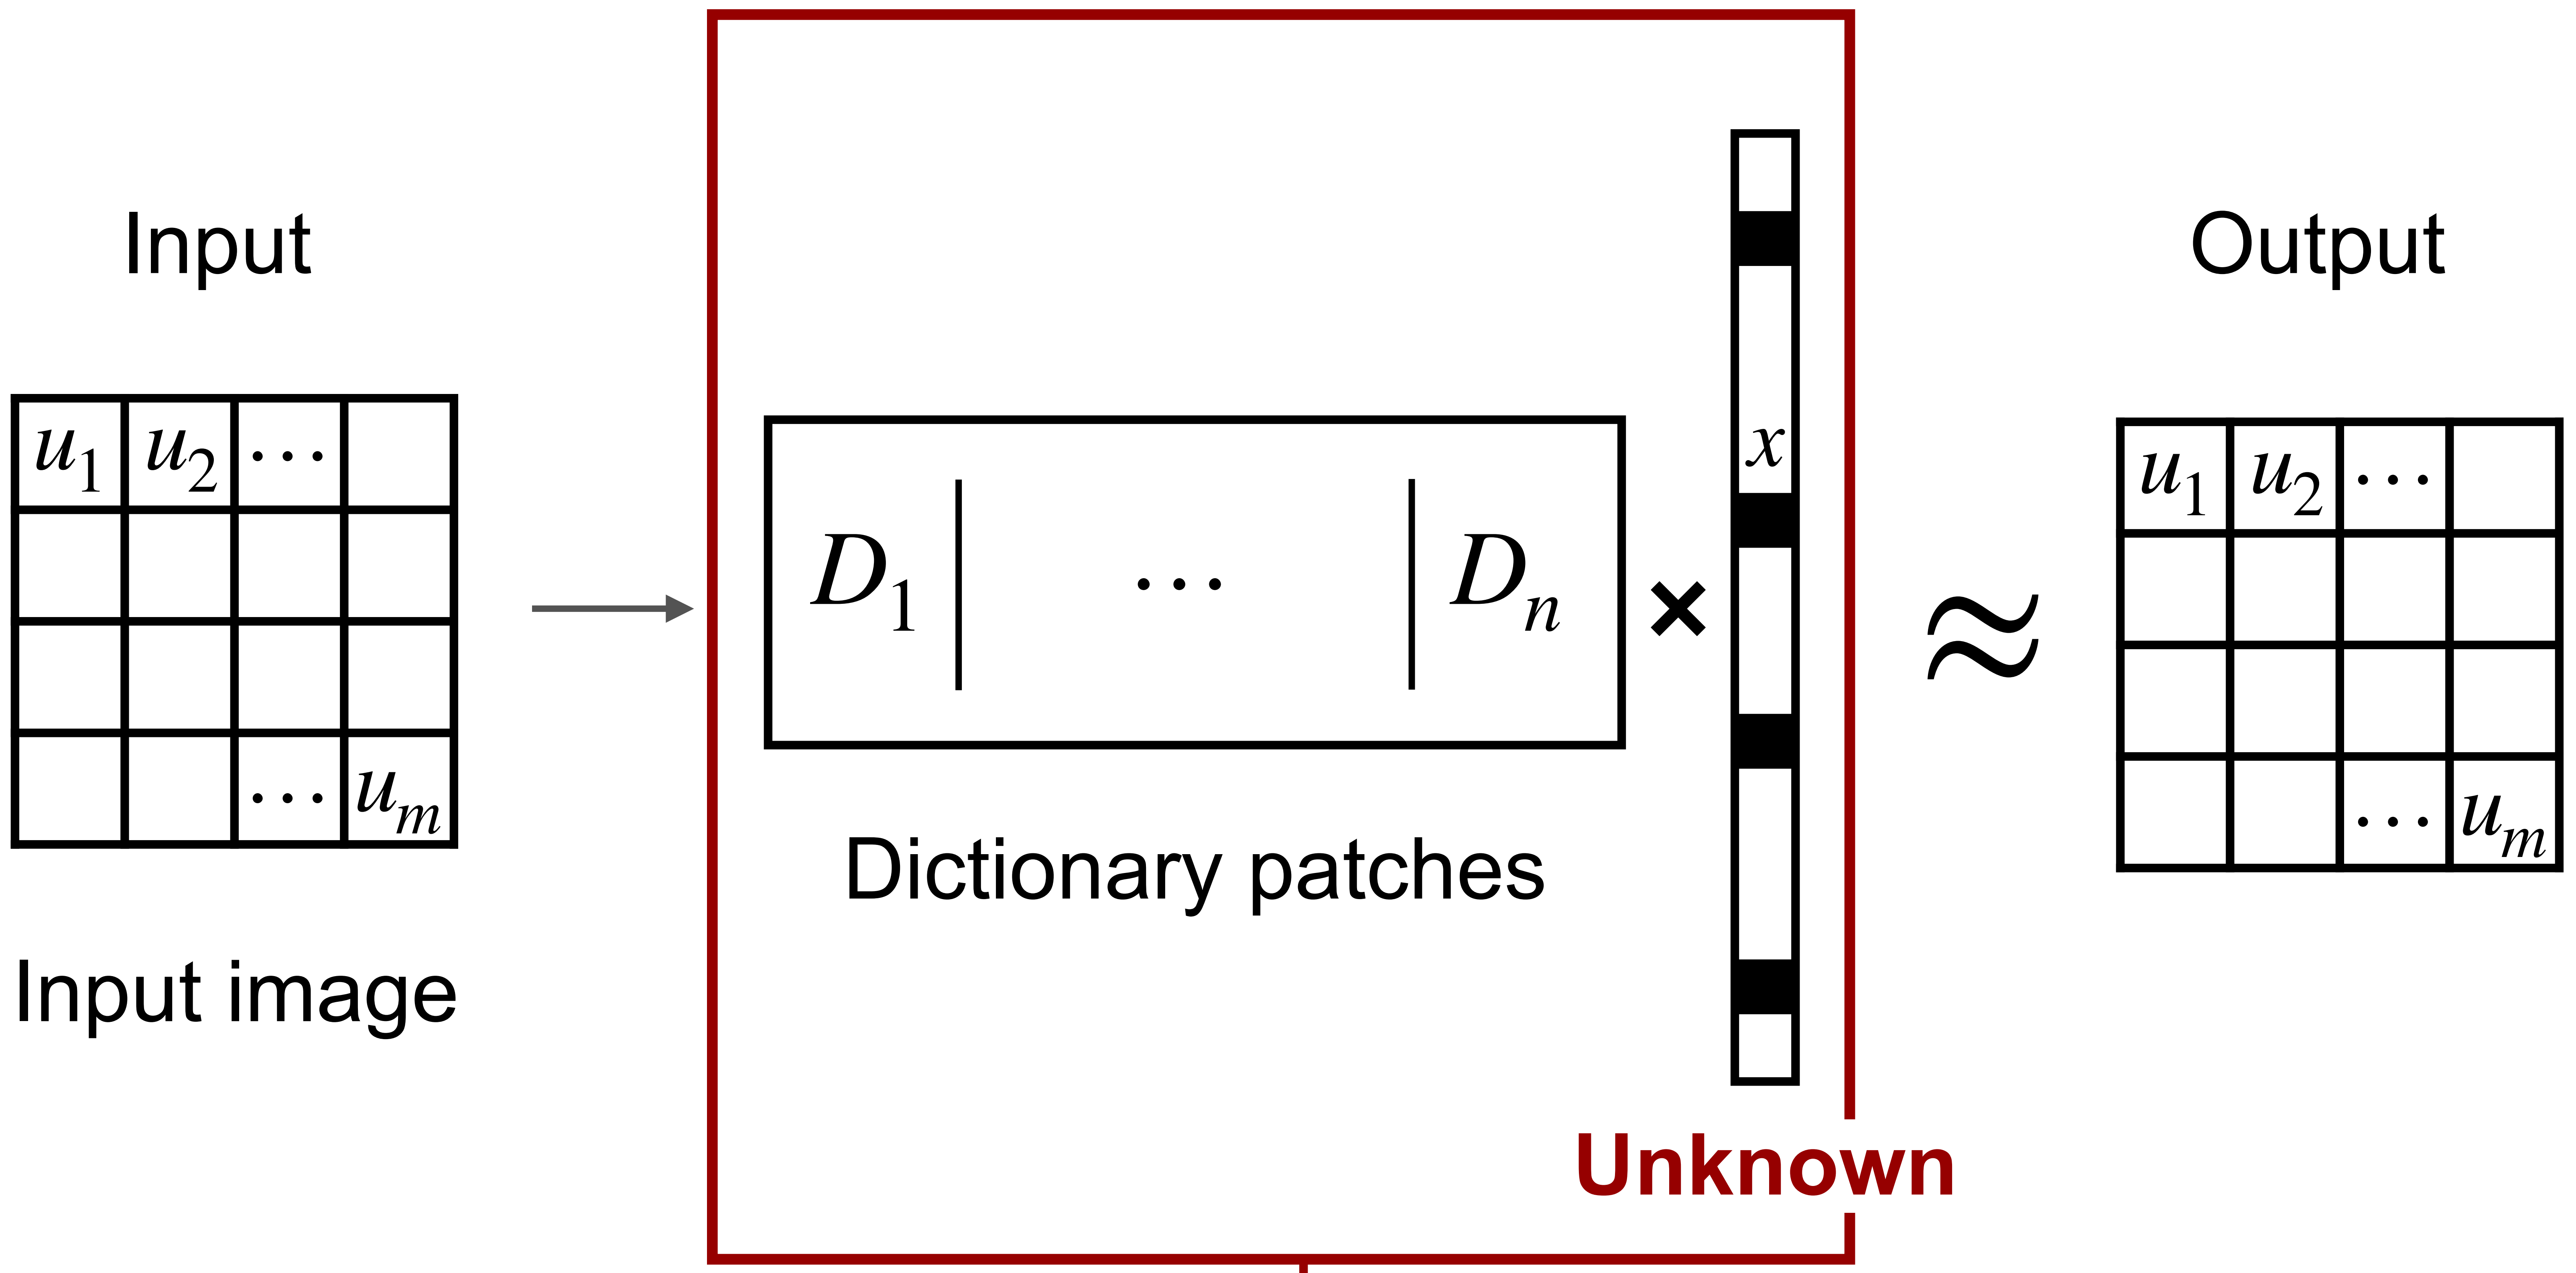

$$\min\left(\textbf{Reconstruction Error} + \textbf{Sparsity constraint}\right)$$

J. Wright and Y. Ma, *"High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications"*, Cambridge University Press, 2022.

Sparse Reconstruction

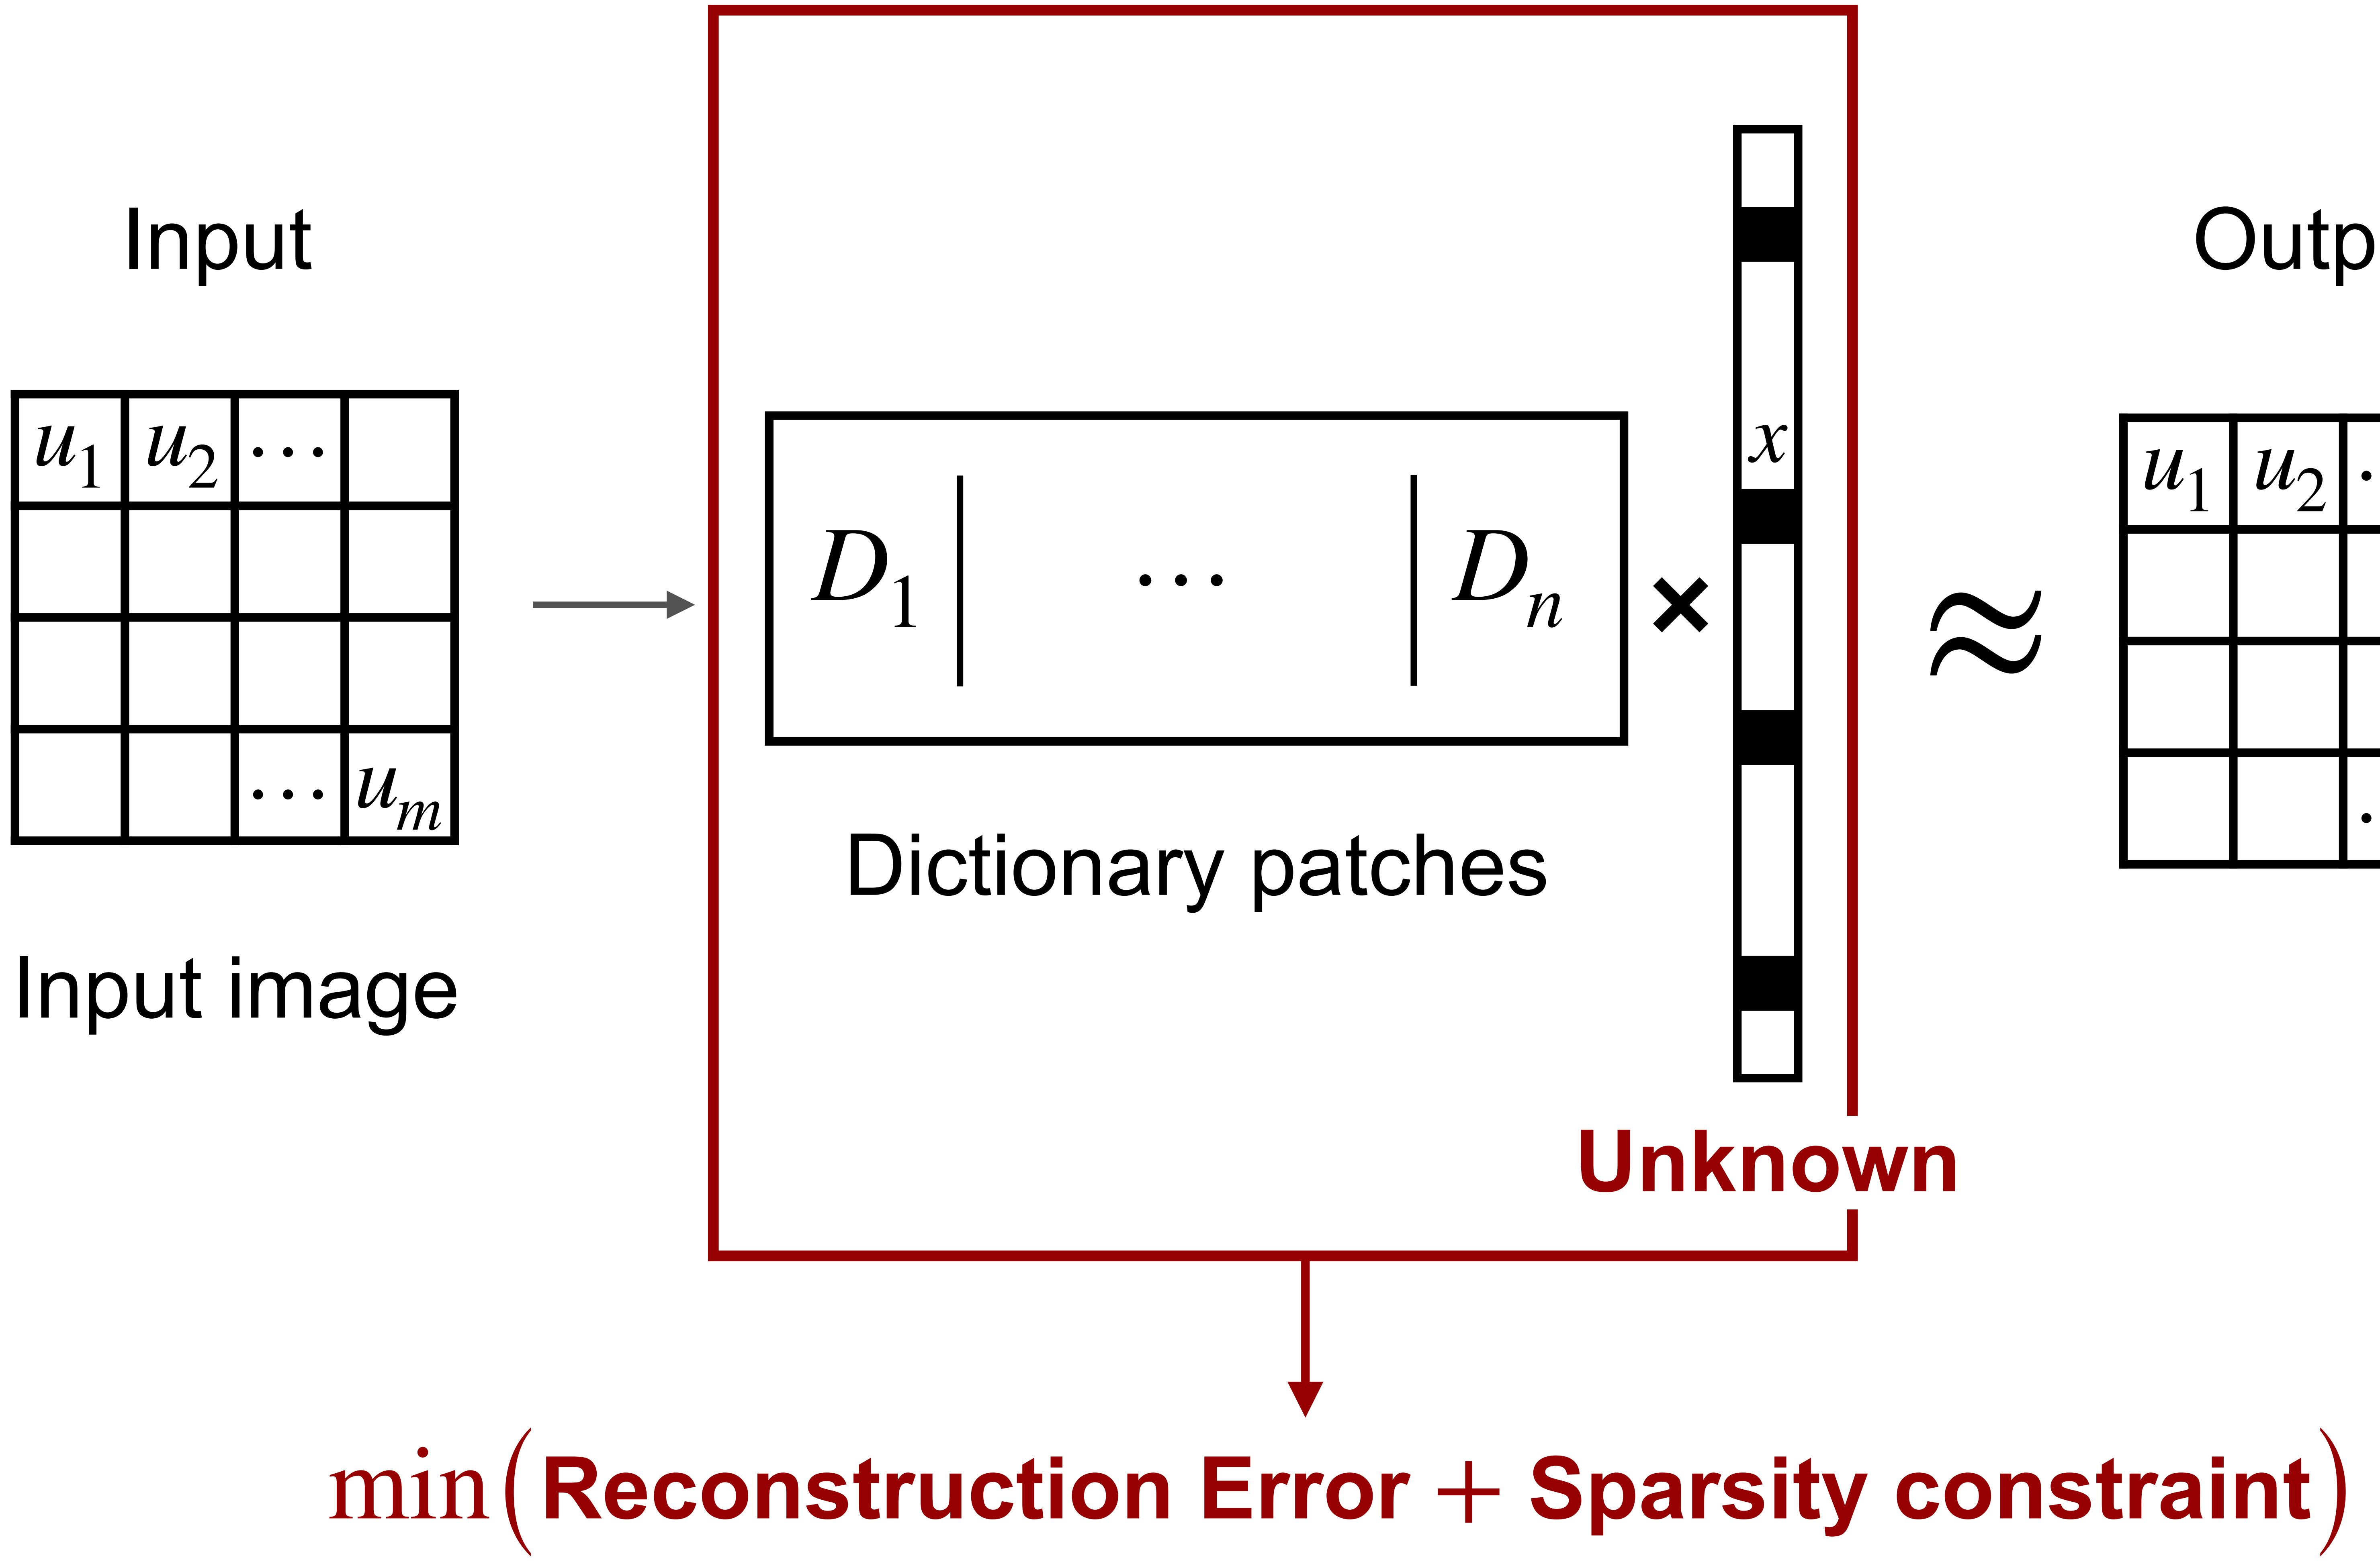

Applications in:

- machine learning
- signal processing
- compressed sensing
- neuroscience

J. Wright and Y. Ma, *"High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications"*, Cambridge University Press, 2022.

Mathematical Formulation

GOAL : Reconstruct $u \in \mathbb{R}^m$ with a linear combination of $D \in \mathbb{R}^{m \times n}$ and a sparse vector $x \in \mathbb{R}^n$.

GOAL : **Reconstruct** $u \in \mathbb{R}^m$ with a linear combination of $D \in \mathbb{R}^{m \times n}$ and a sparse vector $x \in \mathbb{R}^n$.

$$\text{(SR)} \quad \min_{x \in \mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x) := E(x)$$

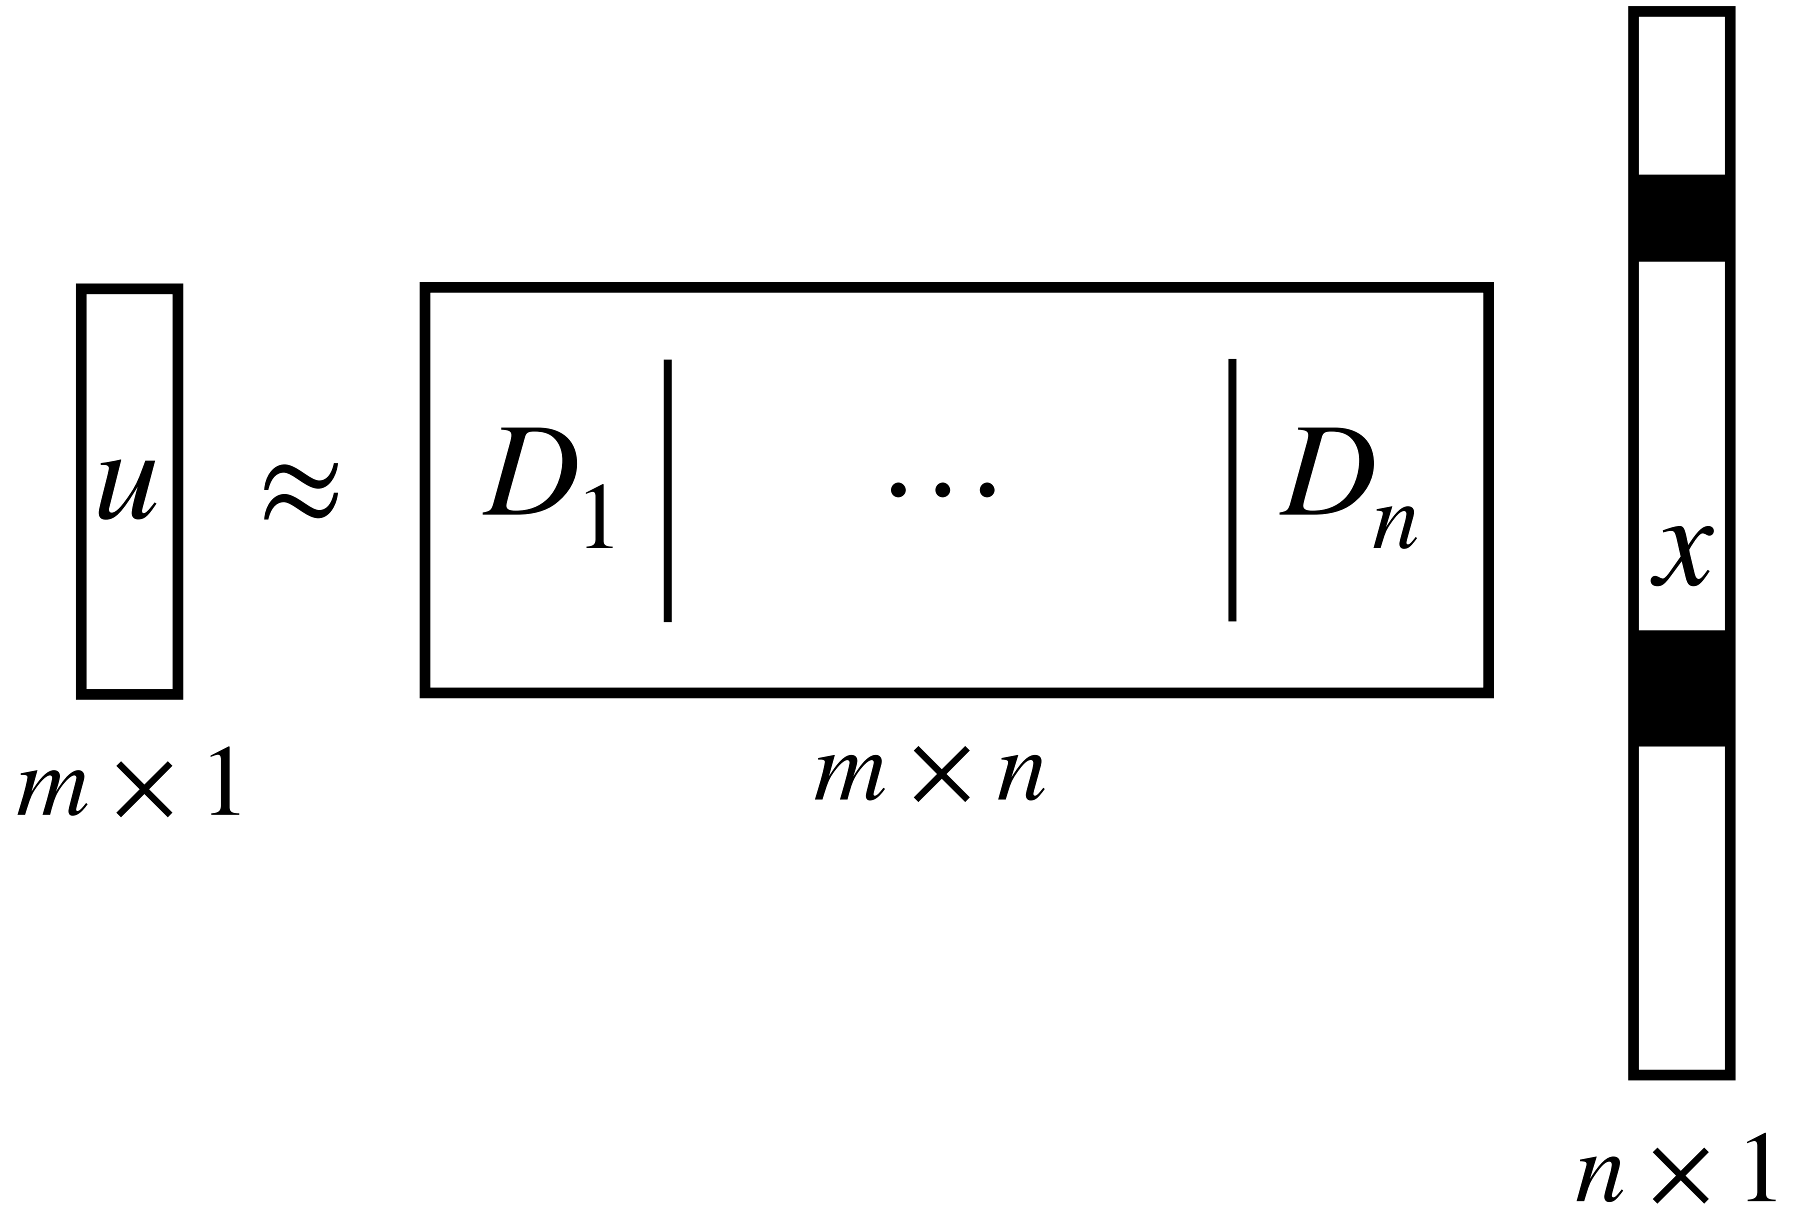

GOAL : Reconstruct $u \in \mathbb{R}^m$ with a linear combination of $D \in \mathbb{R}^{m\times n}$ and a sparse vector $x \in \mathbb{R}^n$.

$$\text{(SR)} \quad \min_{x\in\mathbb{R}^n} \boxed{\frac{1}{2}\|u - Dx\|_2^2} + \lambda S(x) := E(x)$$

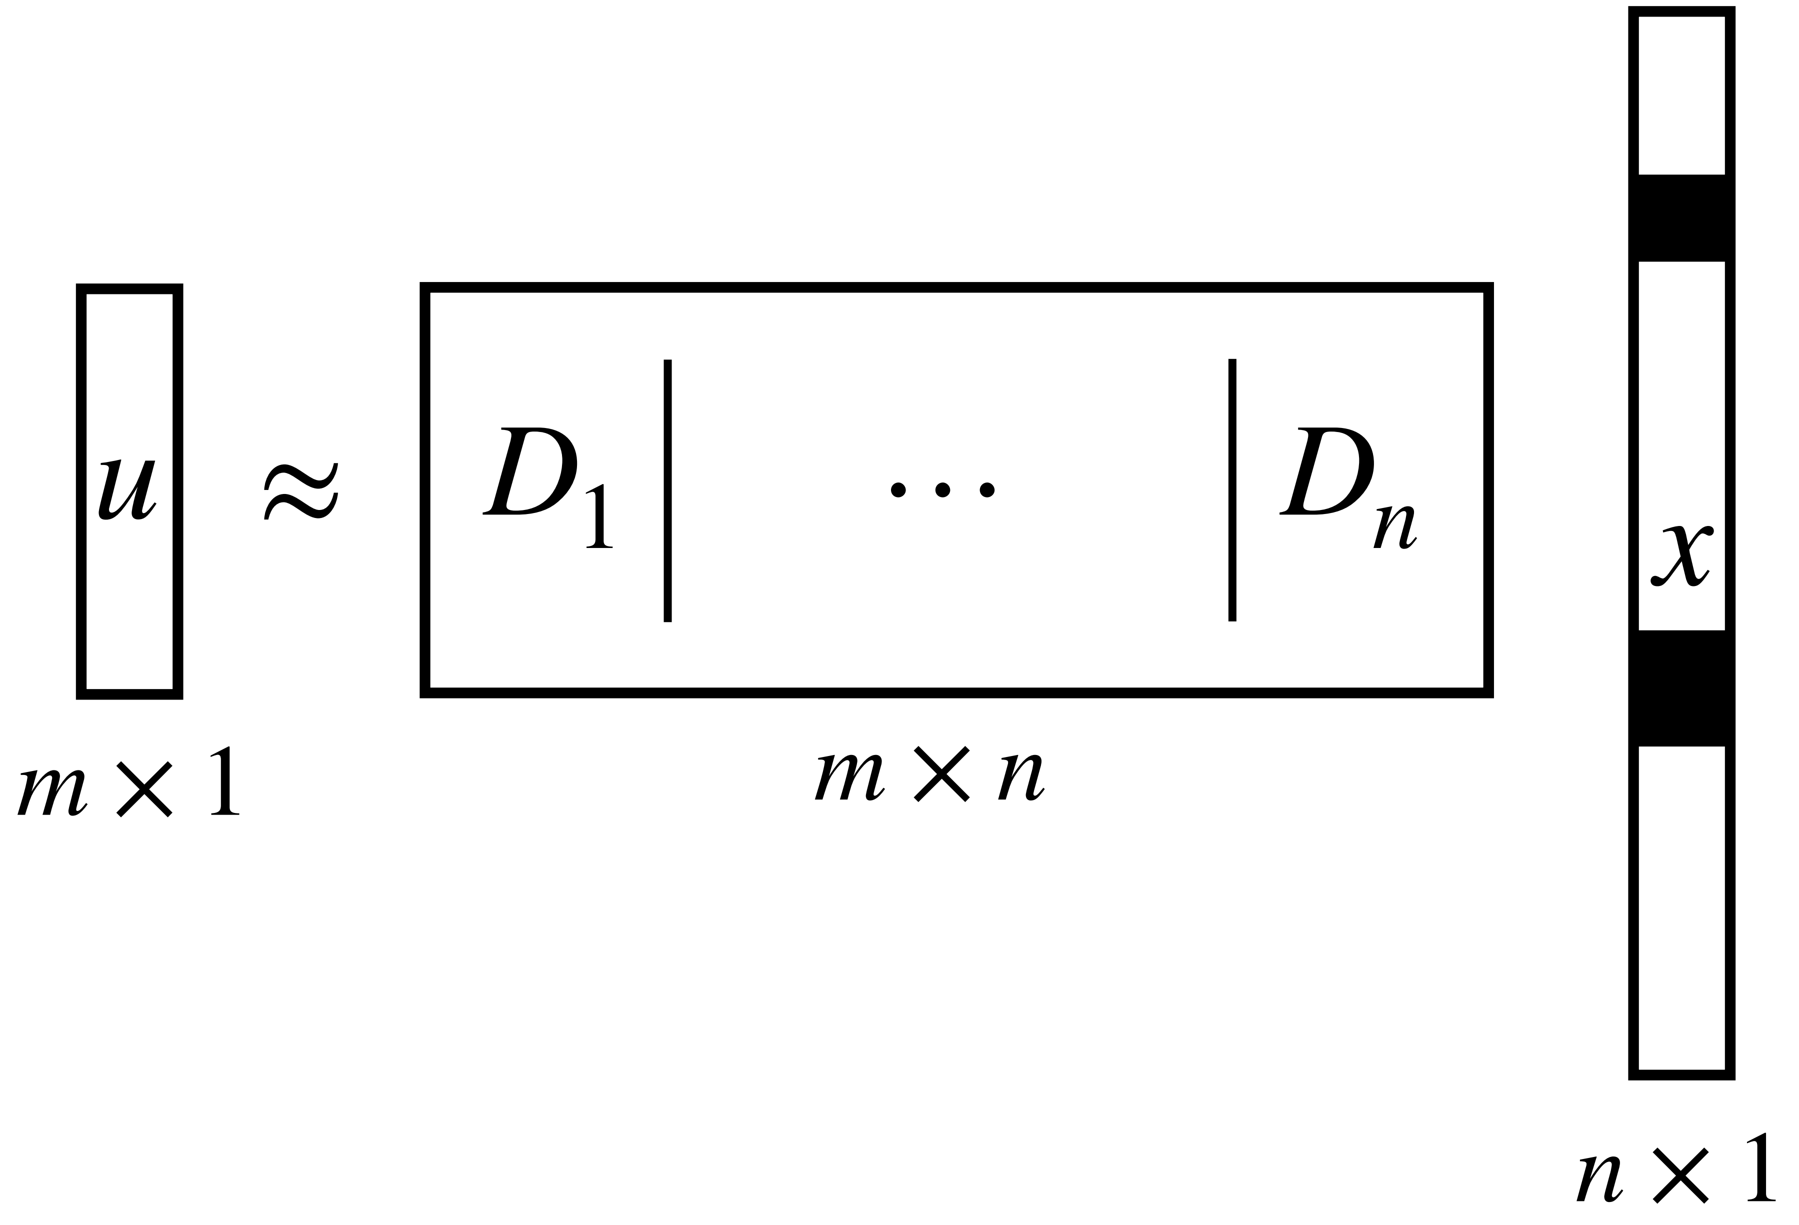

GOAL : Reconstruct $u \in \mathbb{R}^m$ with a linear combination of $D \in \mathbb{R}^{m \times n}$ and a sparse vector $x \in \mathbb{R}^n$.

$$\text{(SR)} \quad \min_{x \in \mathbb{R}^n} \boxed{\frac{1}{2}\|u - Dx\|_2^2} + \boxed{\lambda S(x)} := E(x)$$

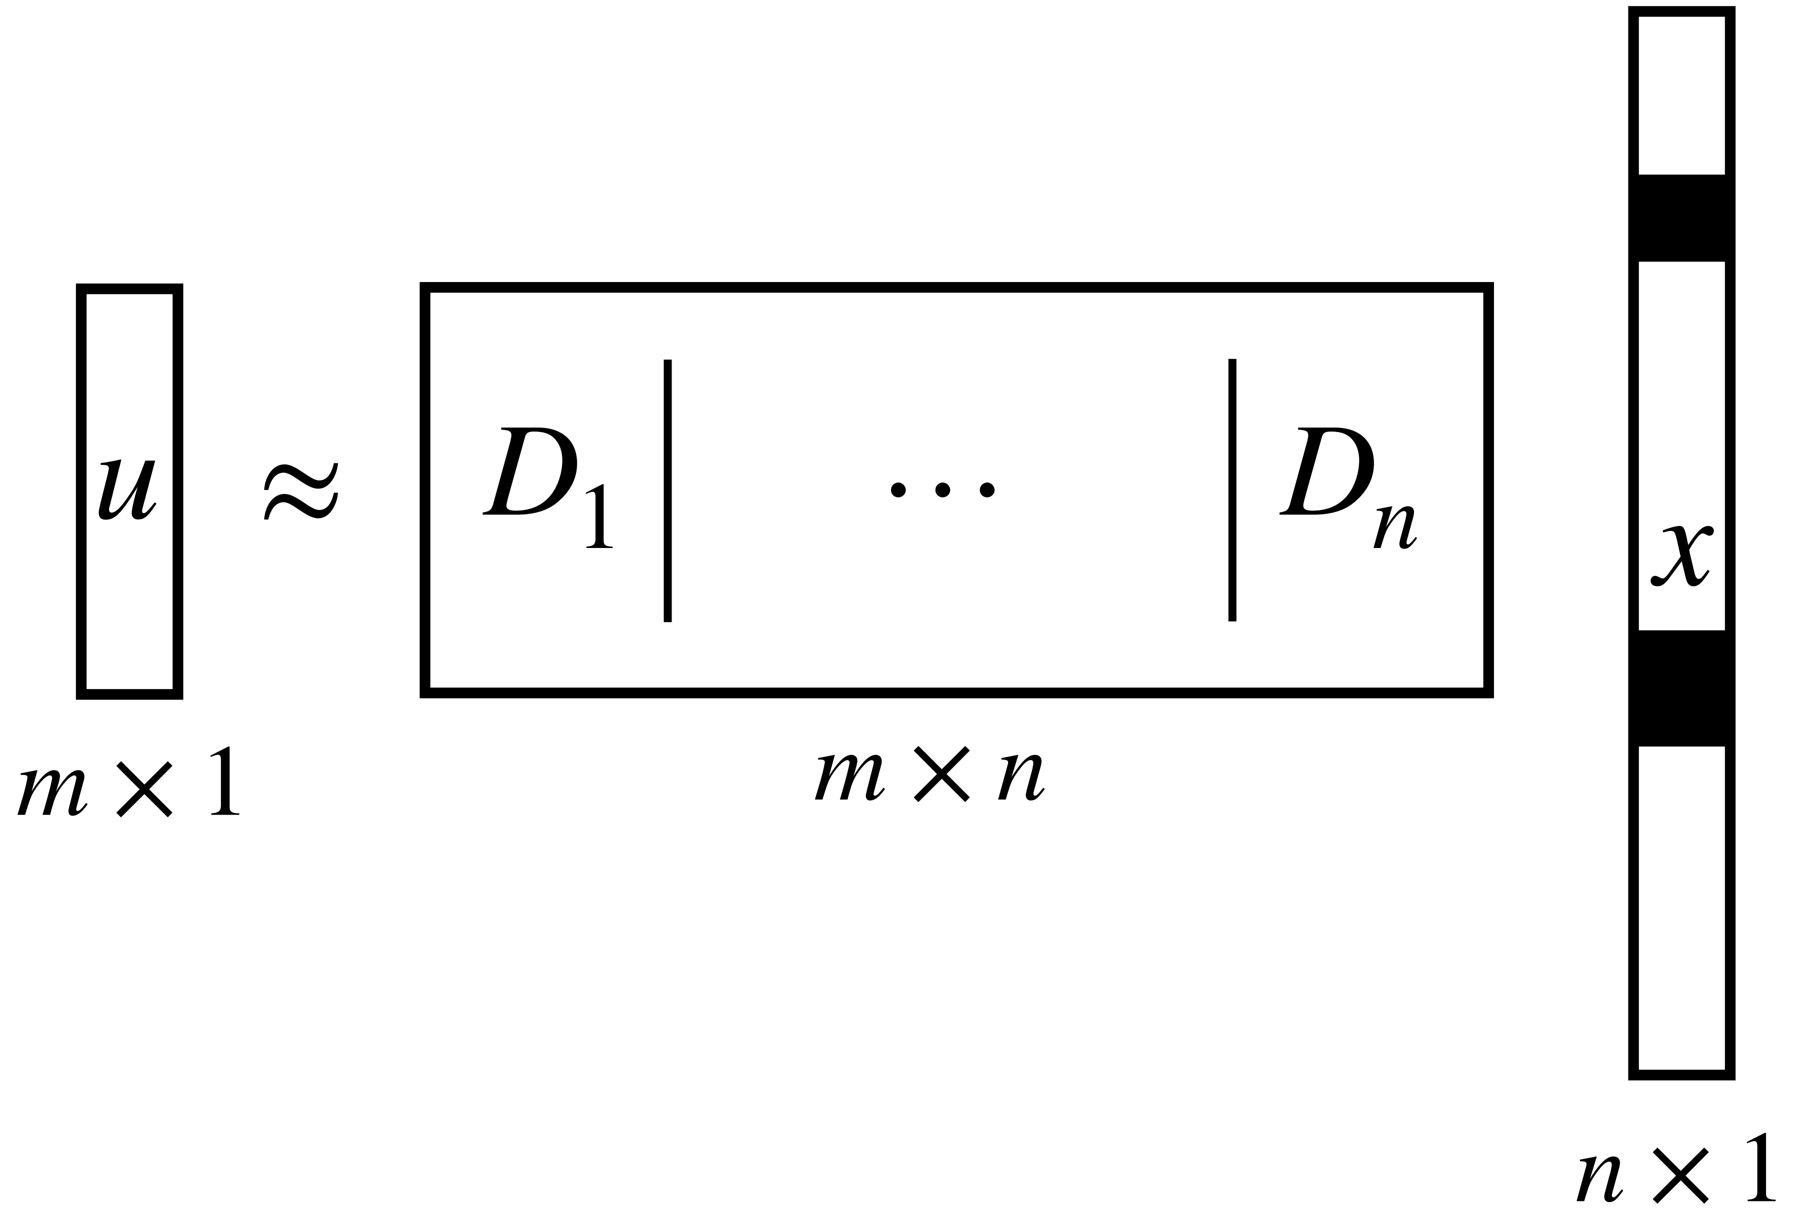

Sparse Reconstruction Problem

$$\min_{x \in \mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x)$$

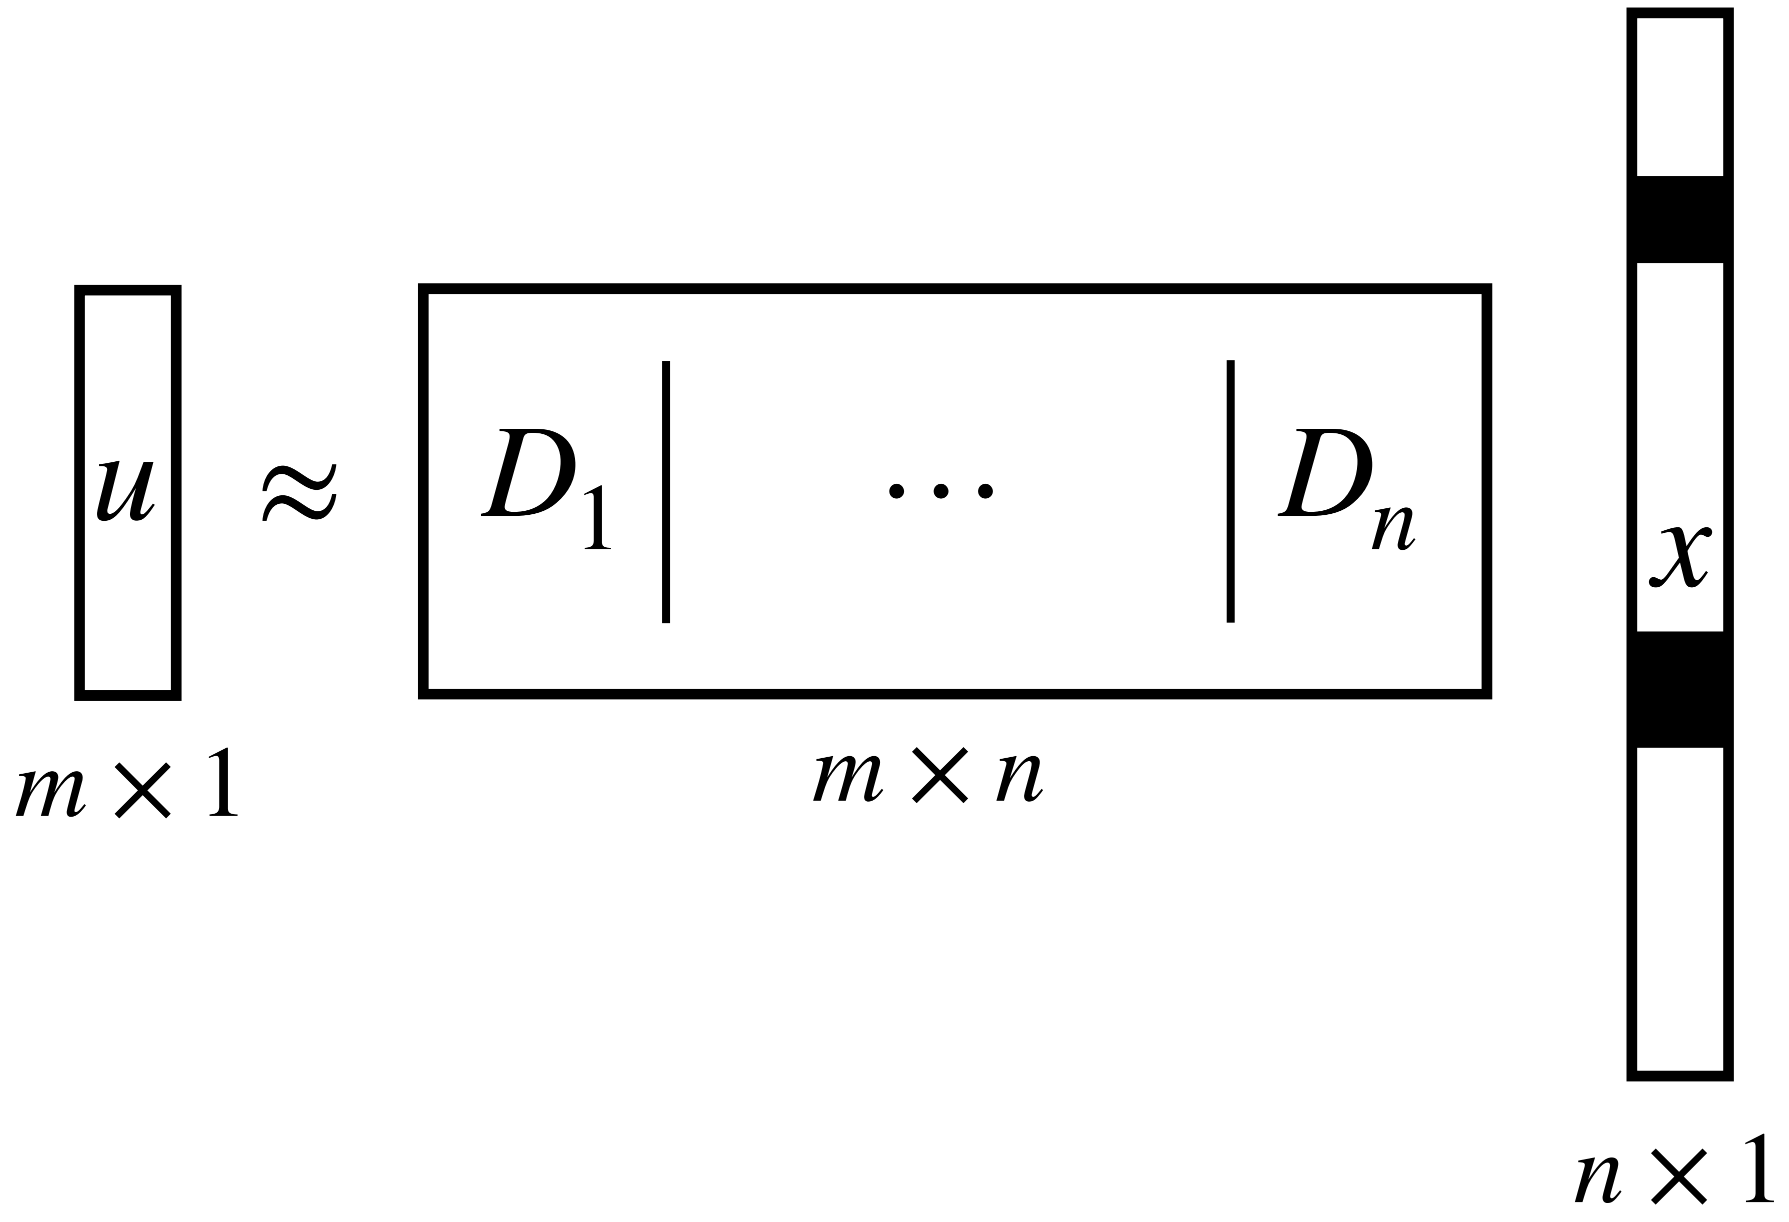

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Sparse Reconstruction Problem

$$\min_{x\in\mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x)$$

$\Longleftrightarrow$

Firing-Rate Competitive Networks

$$\dot{x} = -x + \mathrm{prox}_{\lambda S}\big((I_n - D^\top D)x + D^\top u\big)$$

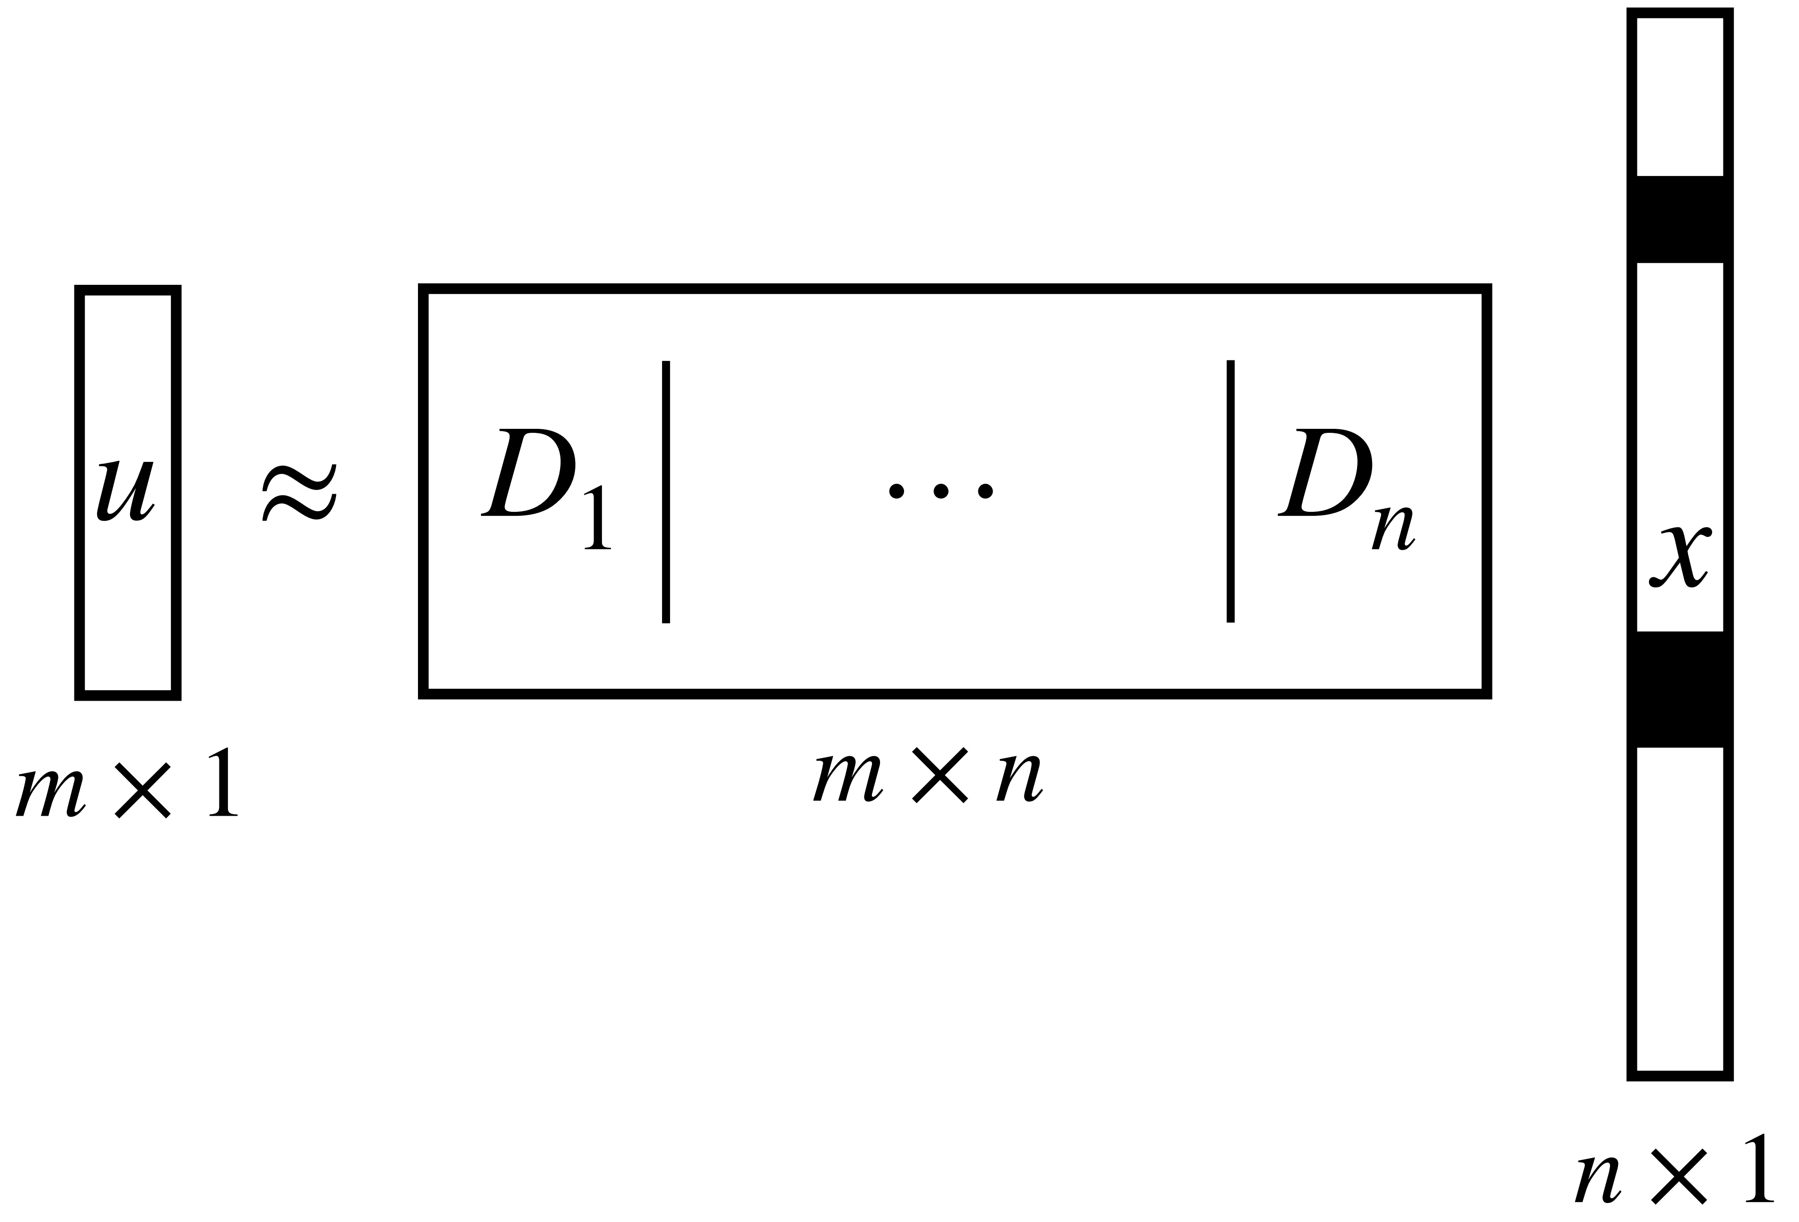

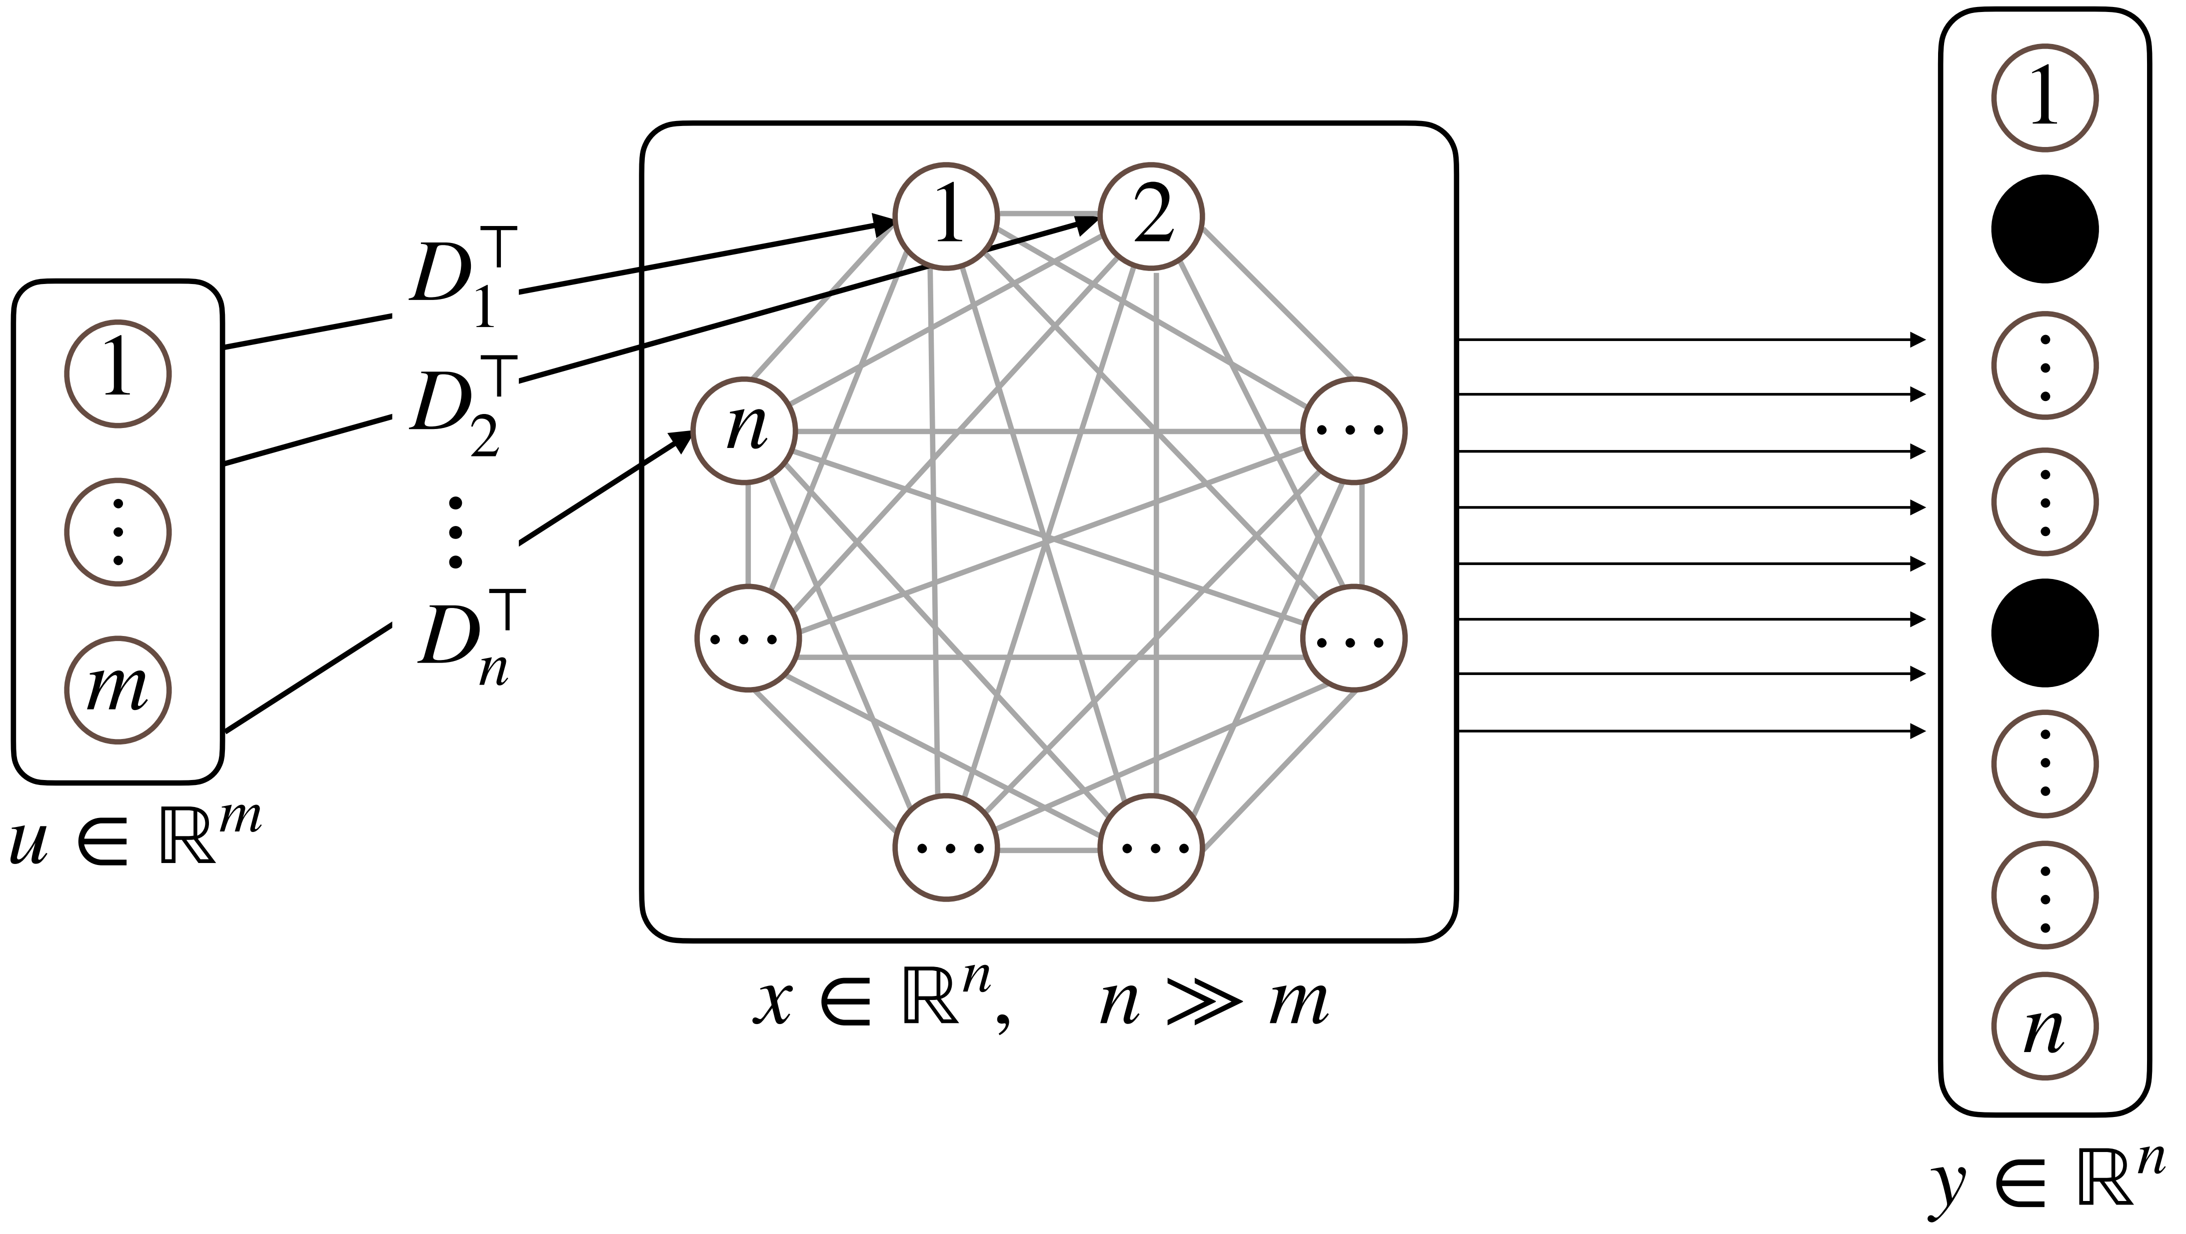

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Sparse Reconstruction Problem

$$\min_{x\in\mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x)$$

$\Longleftrightarrow$

Firing-Rate Competitive Networks

$$\dot{x} = -x + \mathrm{prox}_{\lambda S}\big((I_n - D^\top D)x + D^\top u\big)$$

Lasso: $S_L(x) = \|x\|_1$

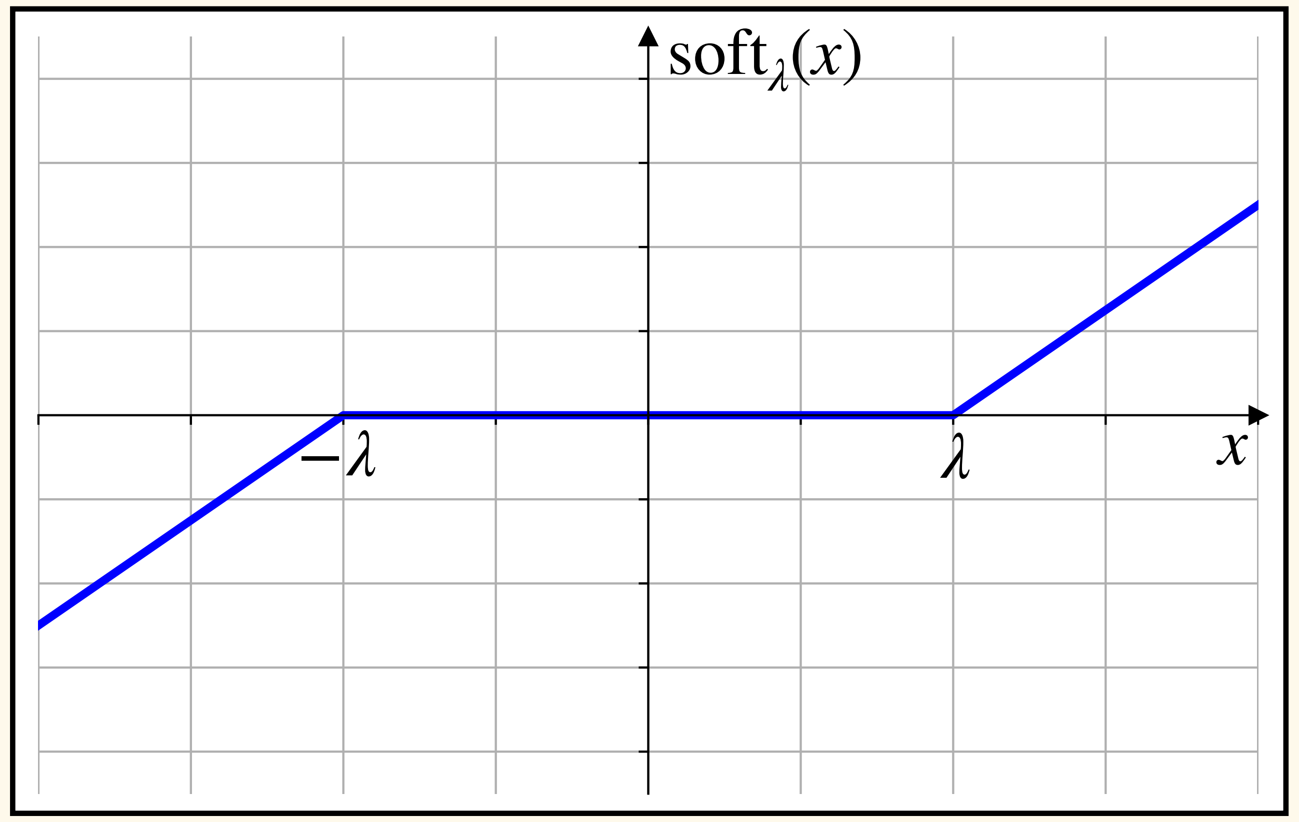

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Sparse Reconstruction Problem

$$\min_{x\in\mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x)$$

$\Longleftrightarrow$

Firing-Rate Competitive Networks

$$\dot{x} = -x + \mathrm{prox}_{\lambda S}\big((I_n - D^\top D)x + D^\top u\big)$$

Lasso: $S_L(x) = \|x\|_1$

$$\dot{x} = -x + \mathrm{soft}_\lambda\big((I_n - D^\top D)x + D^\top u\big)$$

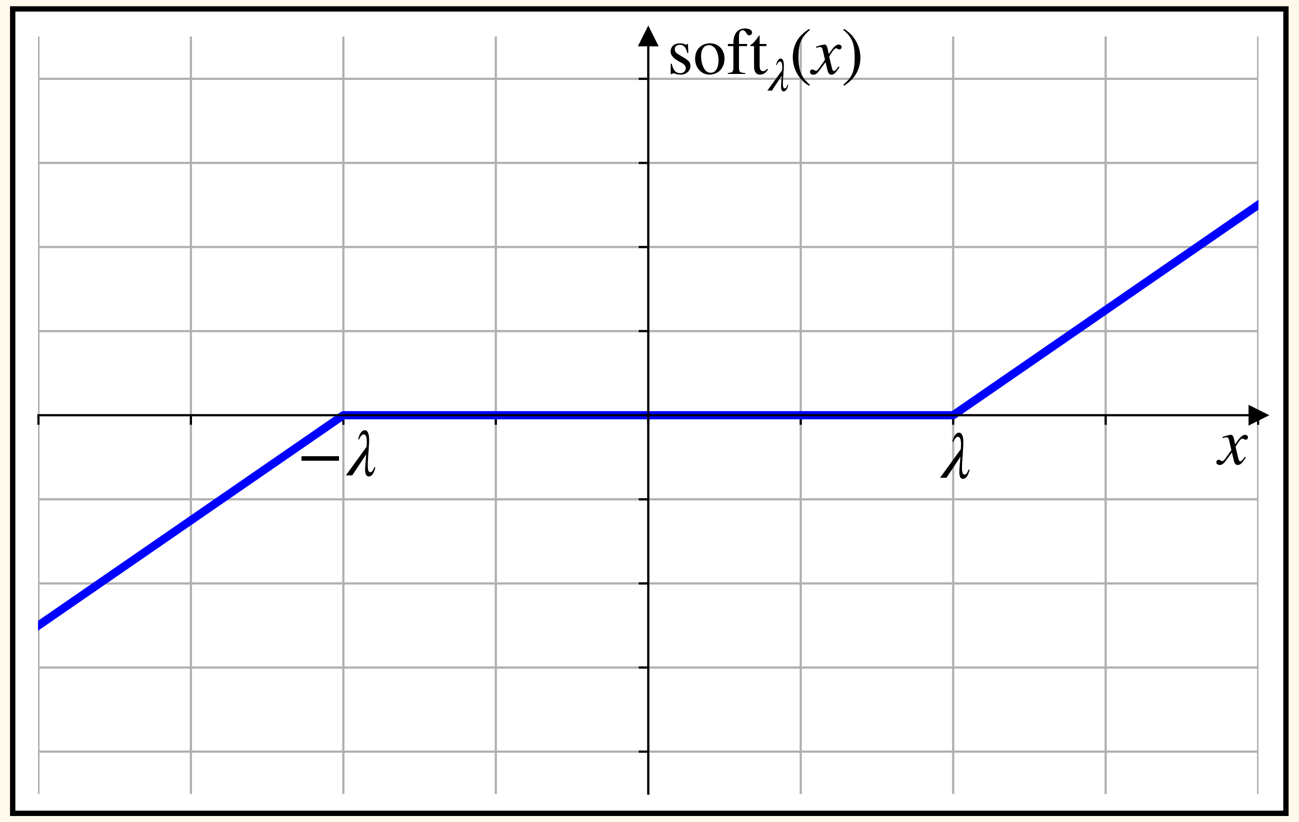

Positive Firing-Rate Competitive Network (PFCN)

Sparse Reconstruction Problem

Firing-Rate Competitive Networks

$$\min_{x\in\mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x) \iff \dot{x} = -x + \mathrm{prox}_{\lambda S}\big((I_n - D^\top D)x + D^\top u\big)$$

Lasso: $S_L(x) = \|x\|_1$

$$\dot{x} = -x + \mathrm{soft}_\lambda\big((I_n - D^\top D)x + D^\top u\big)$$

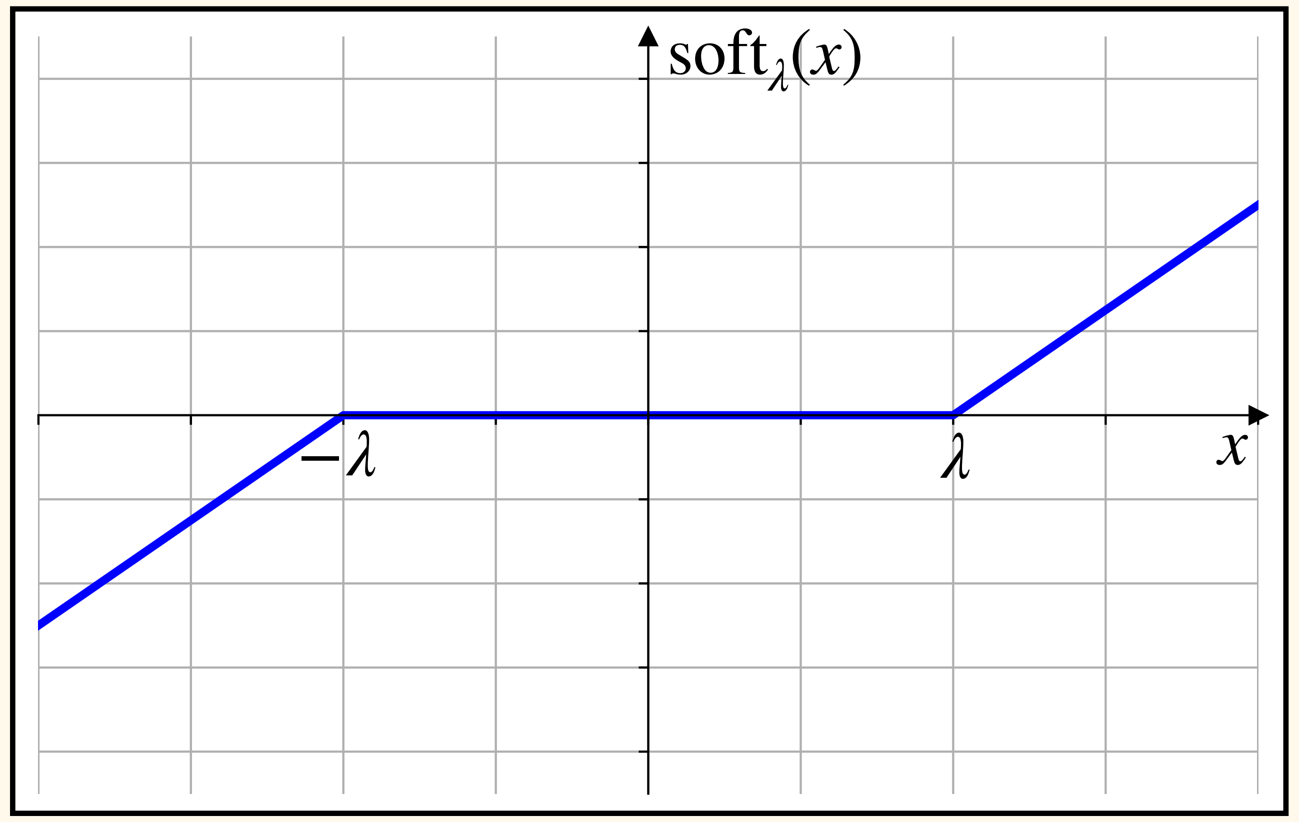

Positive SR: $S_1(x) = \lambda\|x\|_1 + \iota_{\mathbb{R}^n_{\geq 0}}(x)$

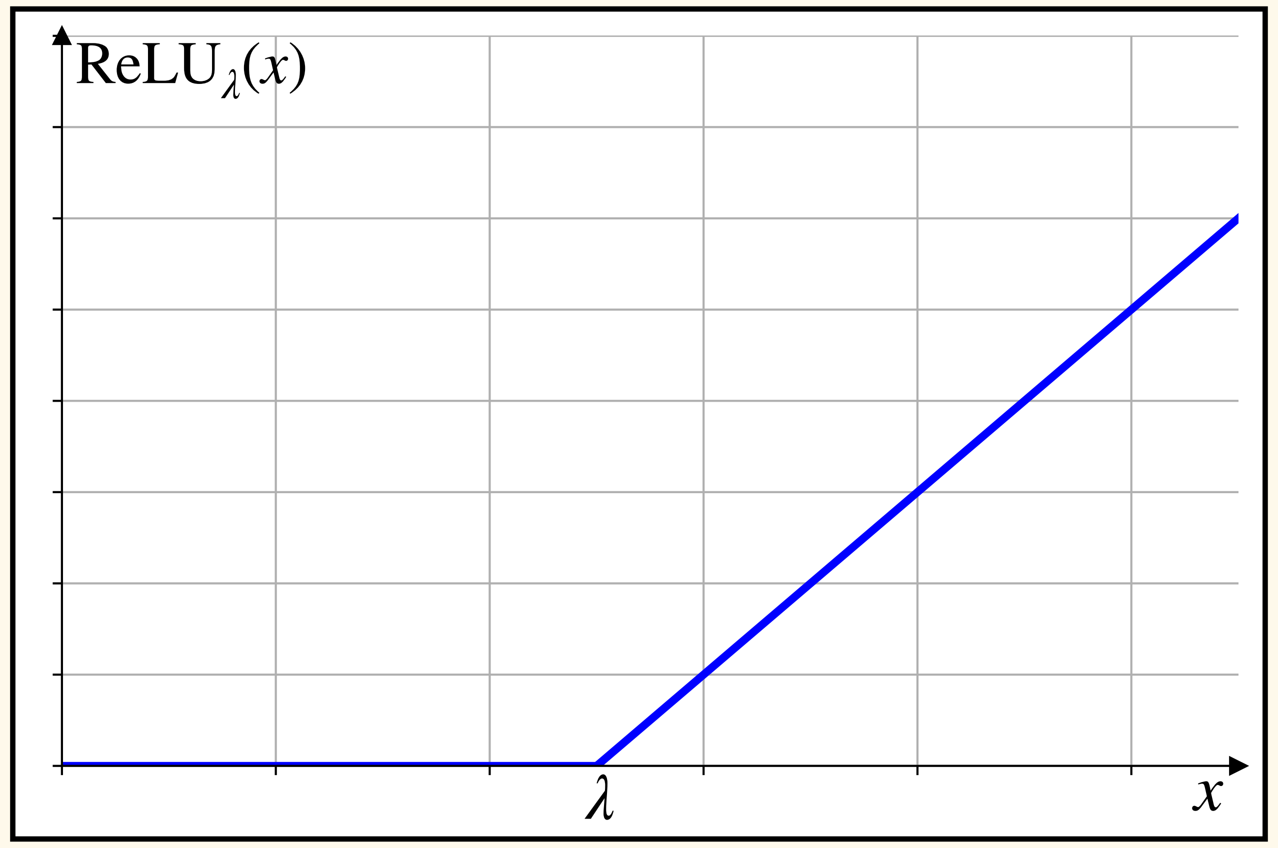

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Positive Firing-Rate Competitive Network (PFCN)

Sparse Reconstruction Problem

Firing-Rate Competitive Networks

$$\min_{x \in \mathbb{R}^n} \frac{1}{2}\|u - Dx\|_2^2 + \lambda S(x) \iff \dot{x} = -x + \mathrm{prox}_{\lambda S}\big((I_n - D^\top D)x + D^\top u\big)$$

Lasso: $S_L(x) = \|x\|_1$

$$\dot{x} = -x + \mathrm{soft}_\lambda\big((I_n - D^\top D)x + D^\top u\big)$$

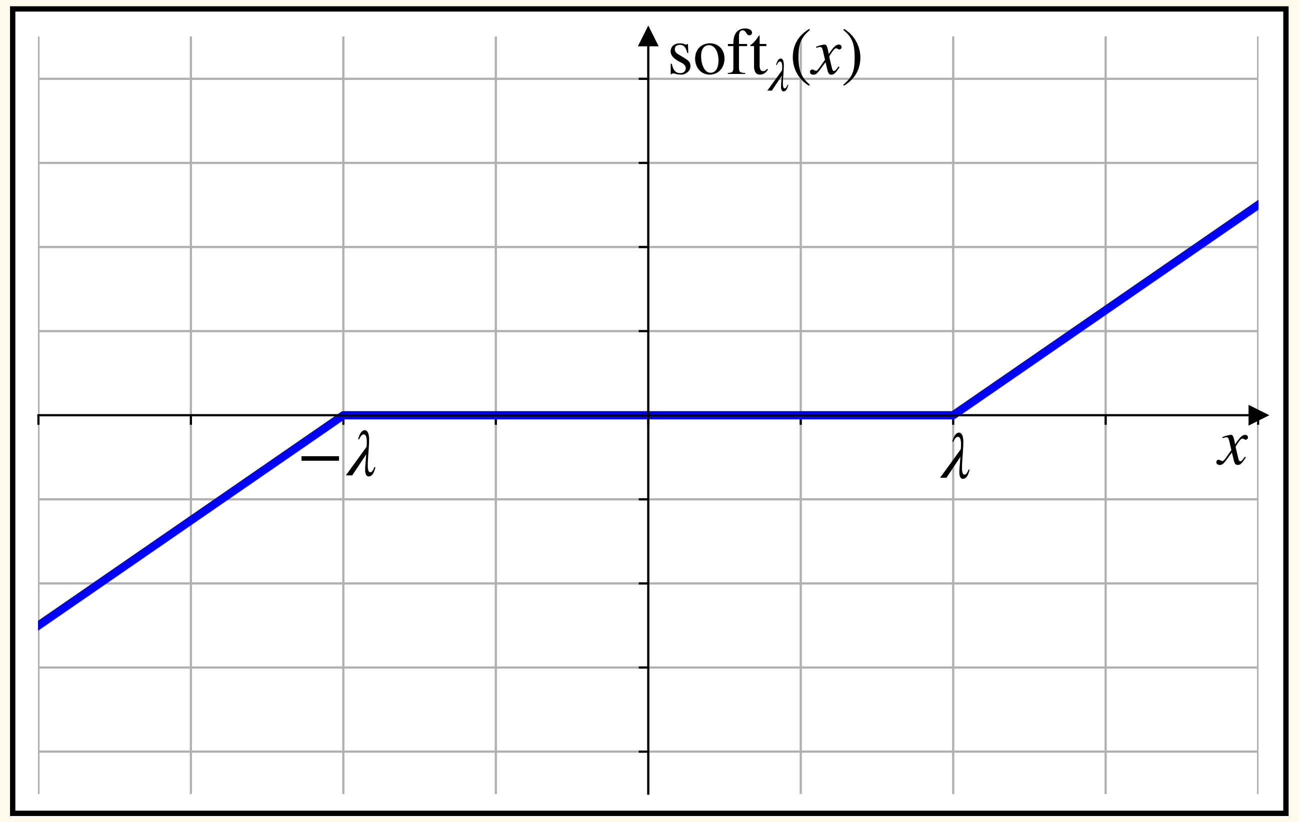

Positive SR: $S_1(x) = \lambda\|x\|_1 + \iota_{\mathbb{R}^n_{\geq 0}}(x)$

$$\dot{x} = -x + \mathrm{ReLU}_\lambda\big((I_n - D^\top D)x + D^\top u\big)$$

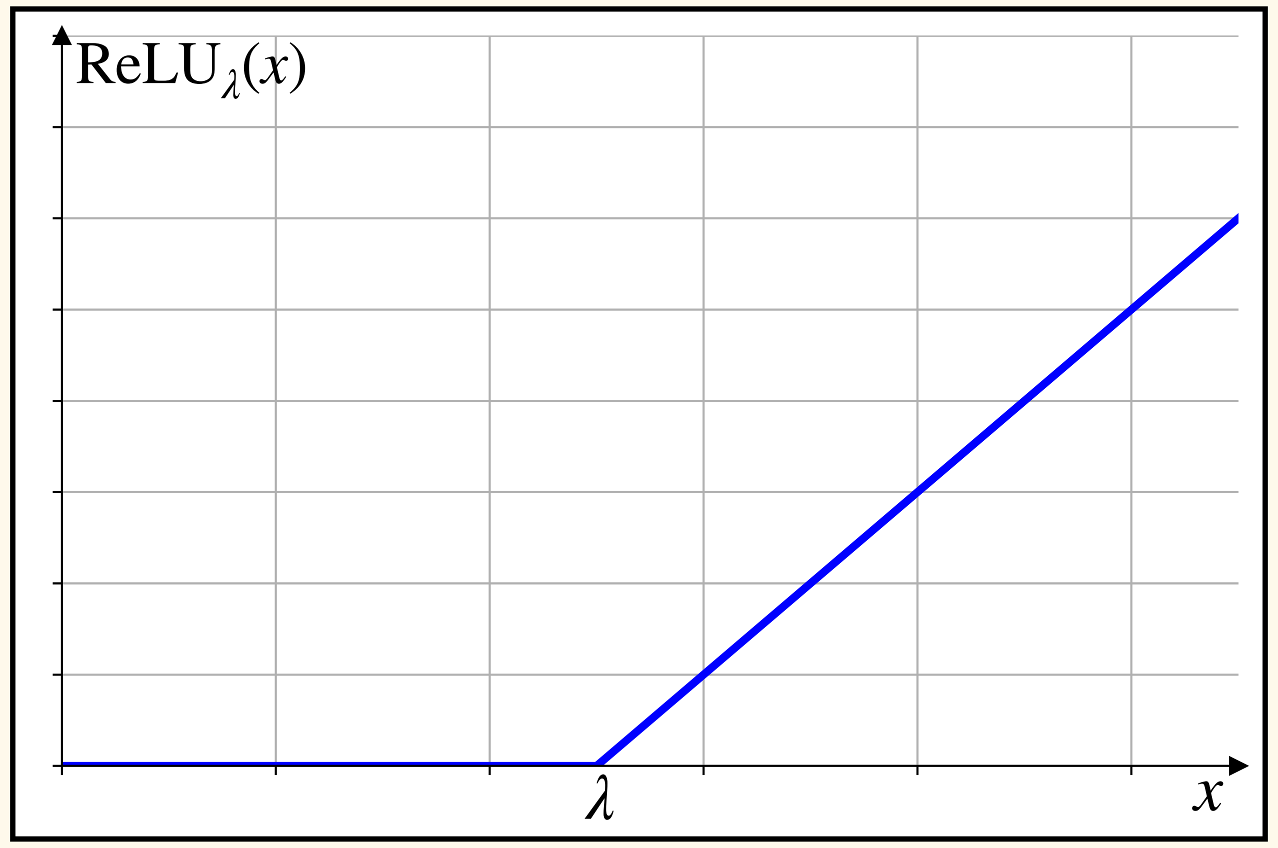

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

$$\dot{x} = -x + \mathrm{ReLU}\big((I_n - D^\top D)x + D^\top u - \lambda\big) := f_{PFCN}(x) \qquad \text{(PFCN)}$$

Property *(On the positiveness of the PFCN)*

The PFCN is a **positive system**, that is the state variables are never negative, given a non-negative initial state

Positive Firing-Rate Competitive Network

$$\dot{x} = -x + \mathrm{ReLU}\big((I_n - D^\top D)x + D^\top u - \lambda\big) := f_{PFCN}(x) \qquad \text{(PFCN)}$$

Property *(On the positiveness of the PFCN)*

The PFCN is a **positive system**, that is the state variables are never negative, given a non-negative initial state

- ☑ Recurrent Connection
- ☑ Exhibit non-negative states
- ☑ Effective model inhibition and excitation

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

Contractivity of Competitive Firing-Rate Neural Networks

$$\dot{x} = -x + \mathrm{ReLU}\Big((I_n - D^\top D)x + D^\top u - \lambda\Big) \qquad \text{(PFCN)}$$

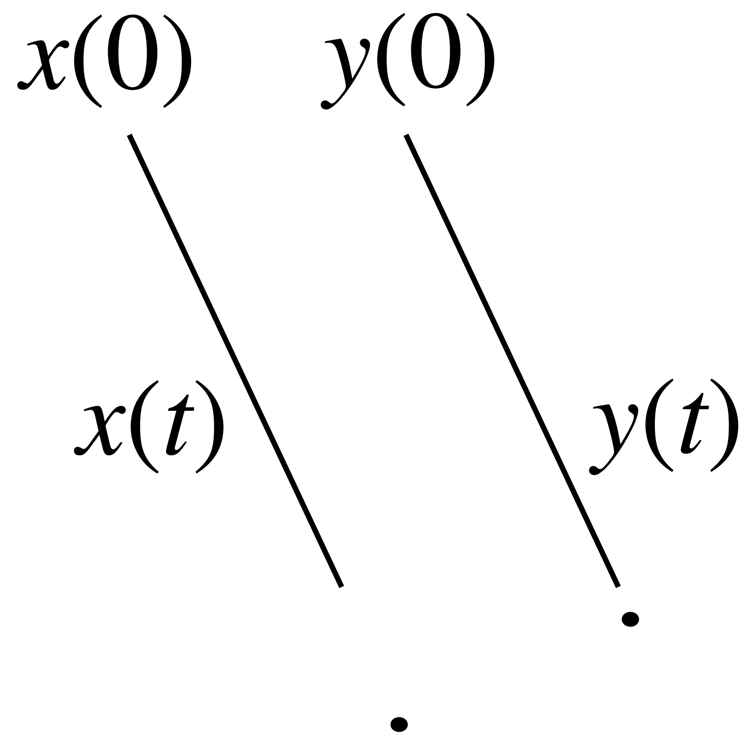

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Euclidean Contractivity of Neural Networks with Symmetric Weights", IEEE Control Systems Letters, vol. 7, pp. 1724-1729, 2023

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

$$\dot{x} = -x + \mathrm{ReLU}\Big((I_n - D^\top D)x + D^\top u - \lambda\Big) \qquad \text{(PFCN)}$$

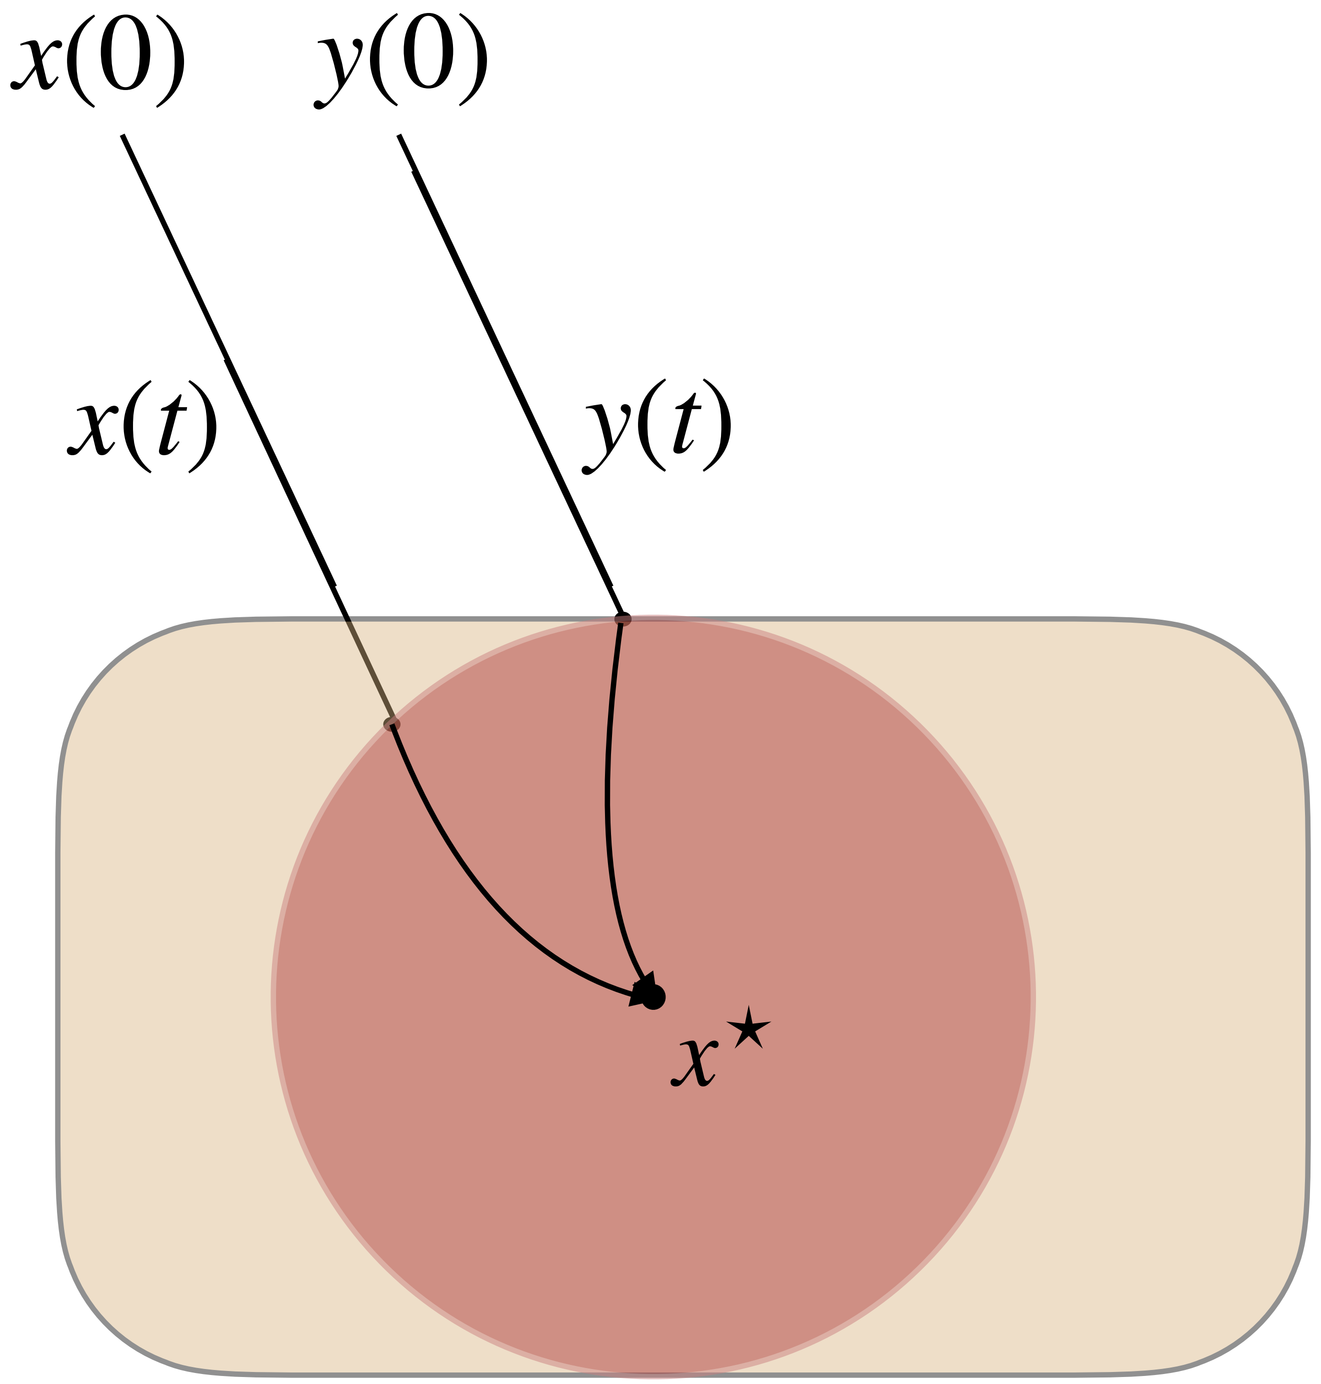

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Euclidean Contractivity of Neural Networks with Symmetric Weights", IEEE Control Systems Letters, vol. 7, pp. 1724-1729, 2023

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

$$\dot{x} = -x + \mathrm{ReLU}\Big((I_n - D^\top D)x + D^\top u - \lambda\Big) \qquad \text{(PFCN)}$$

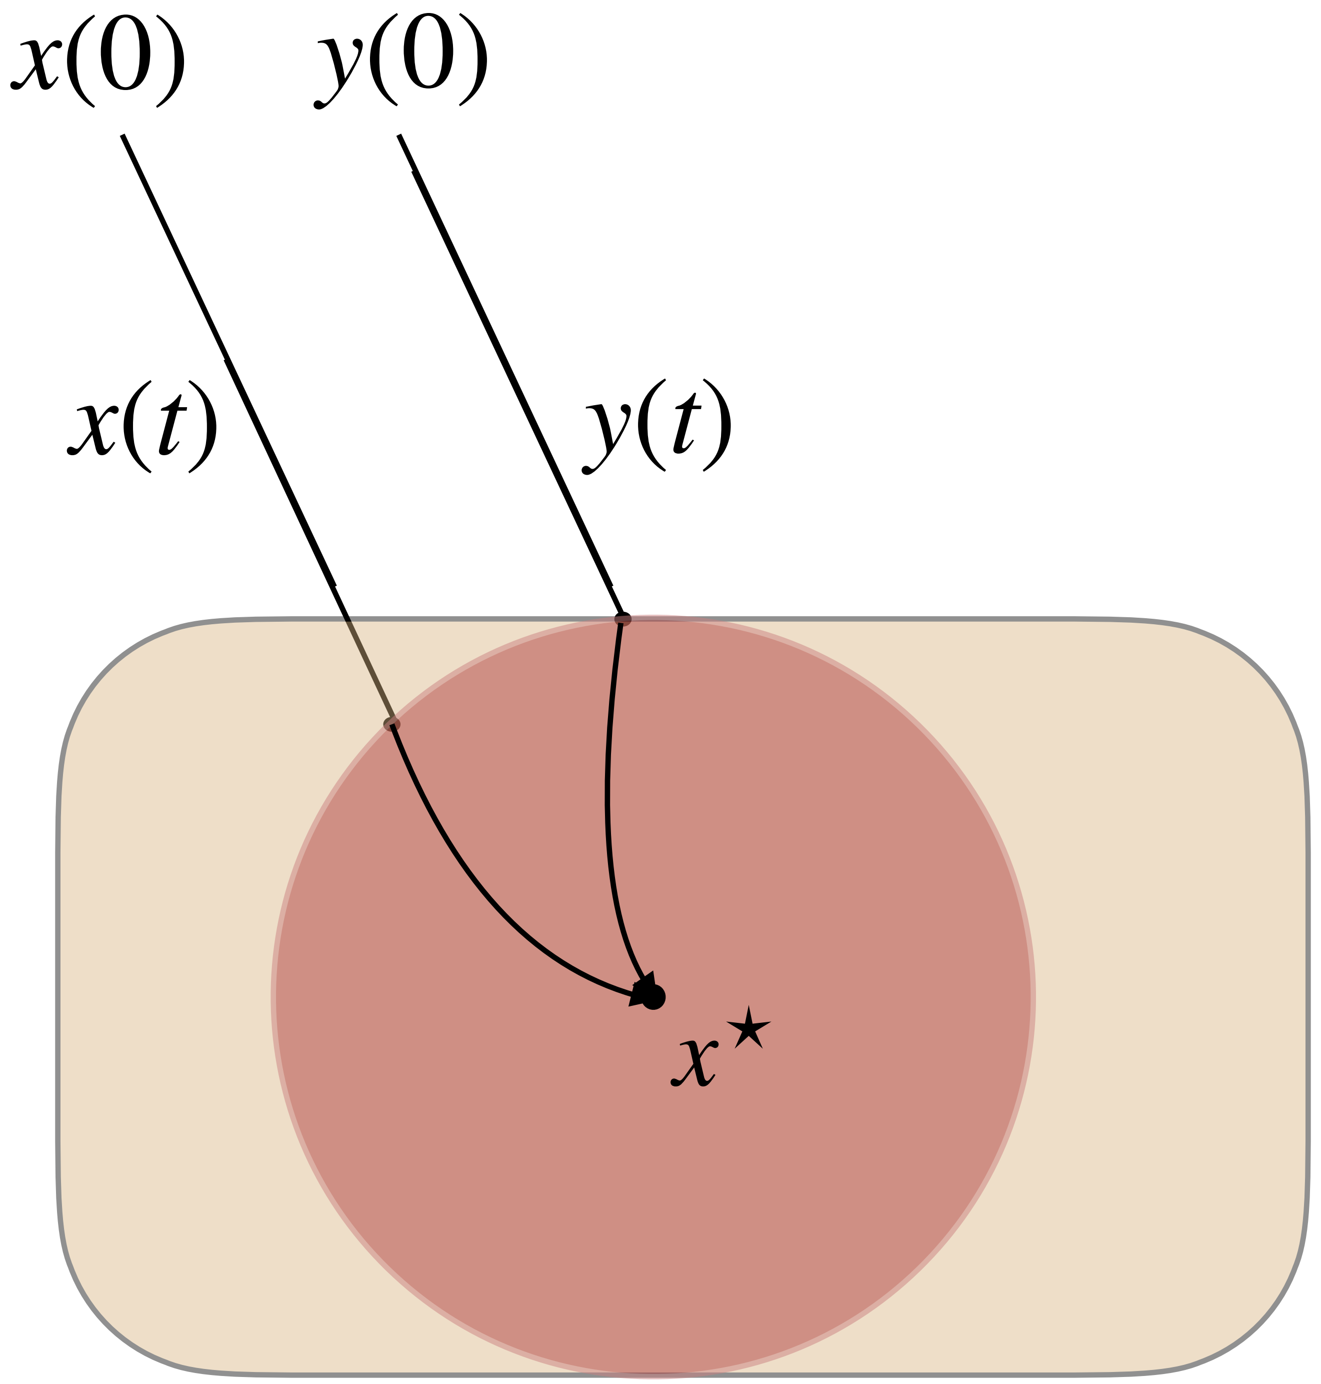

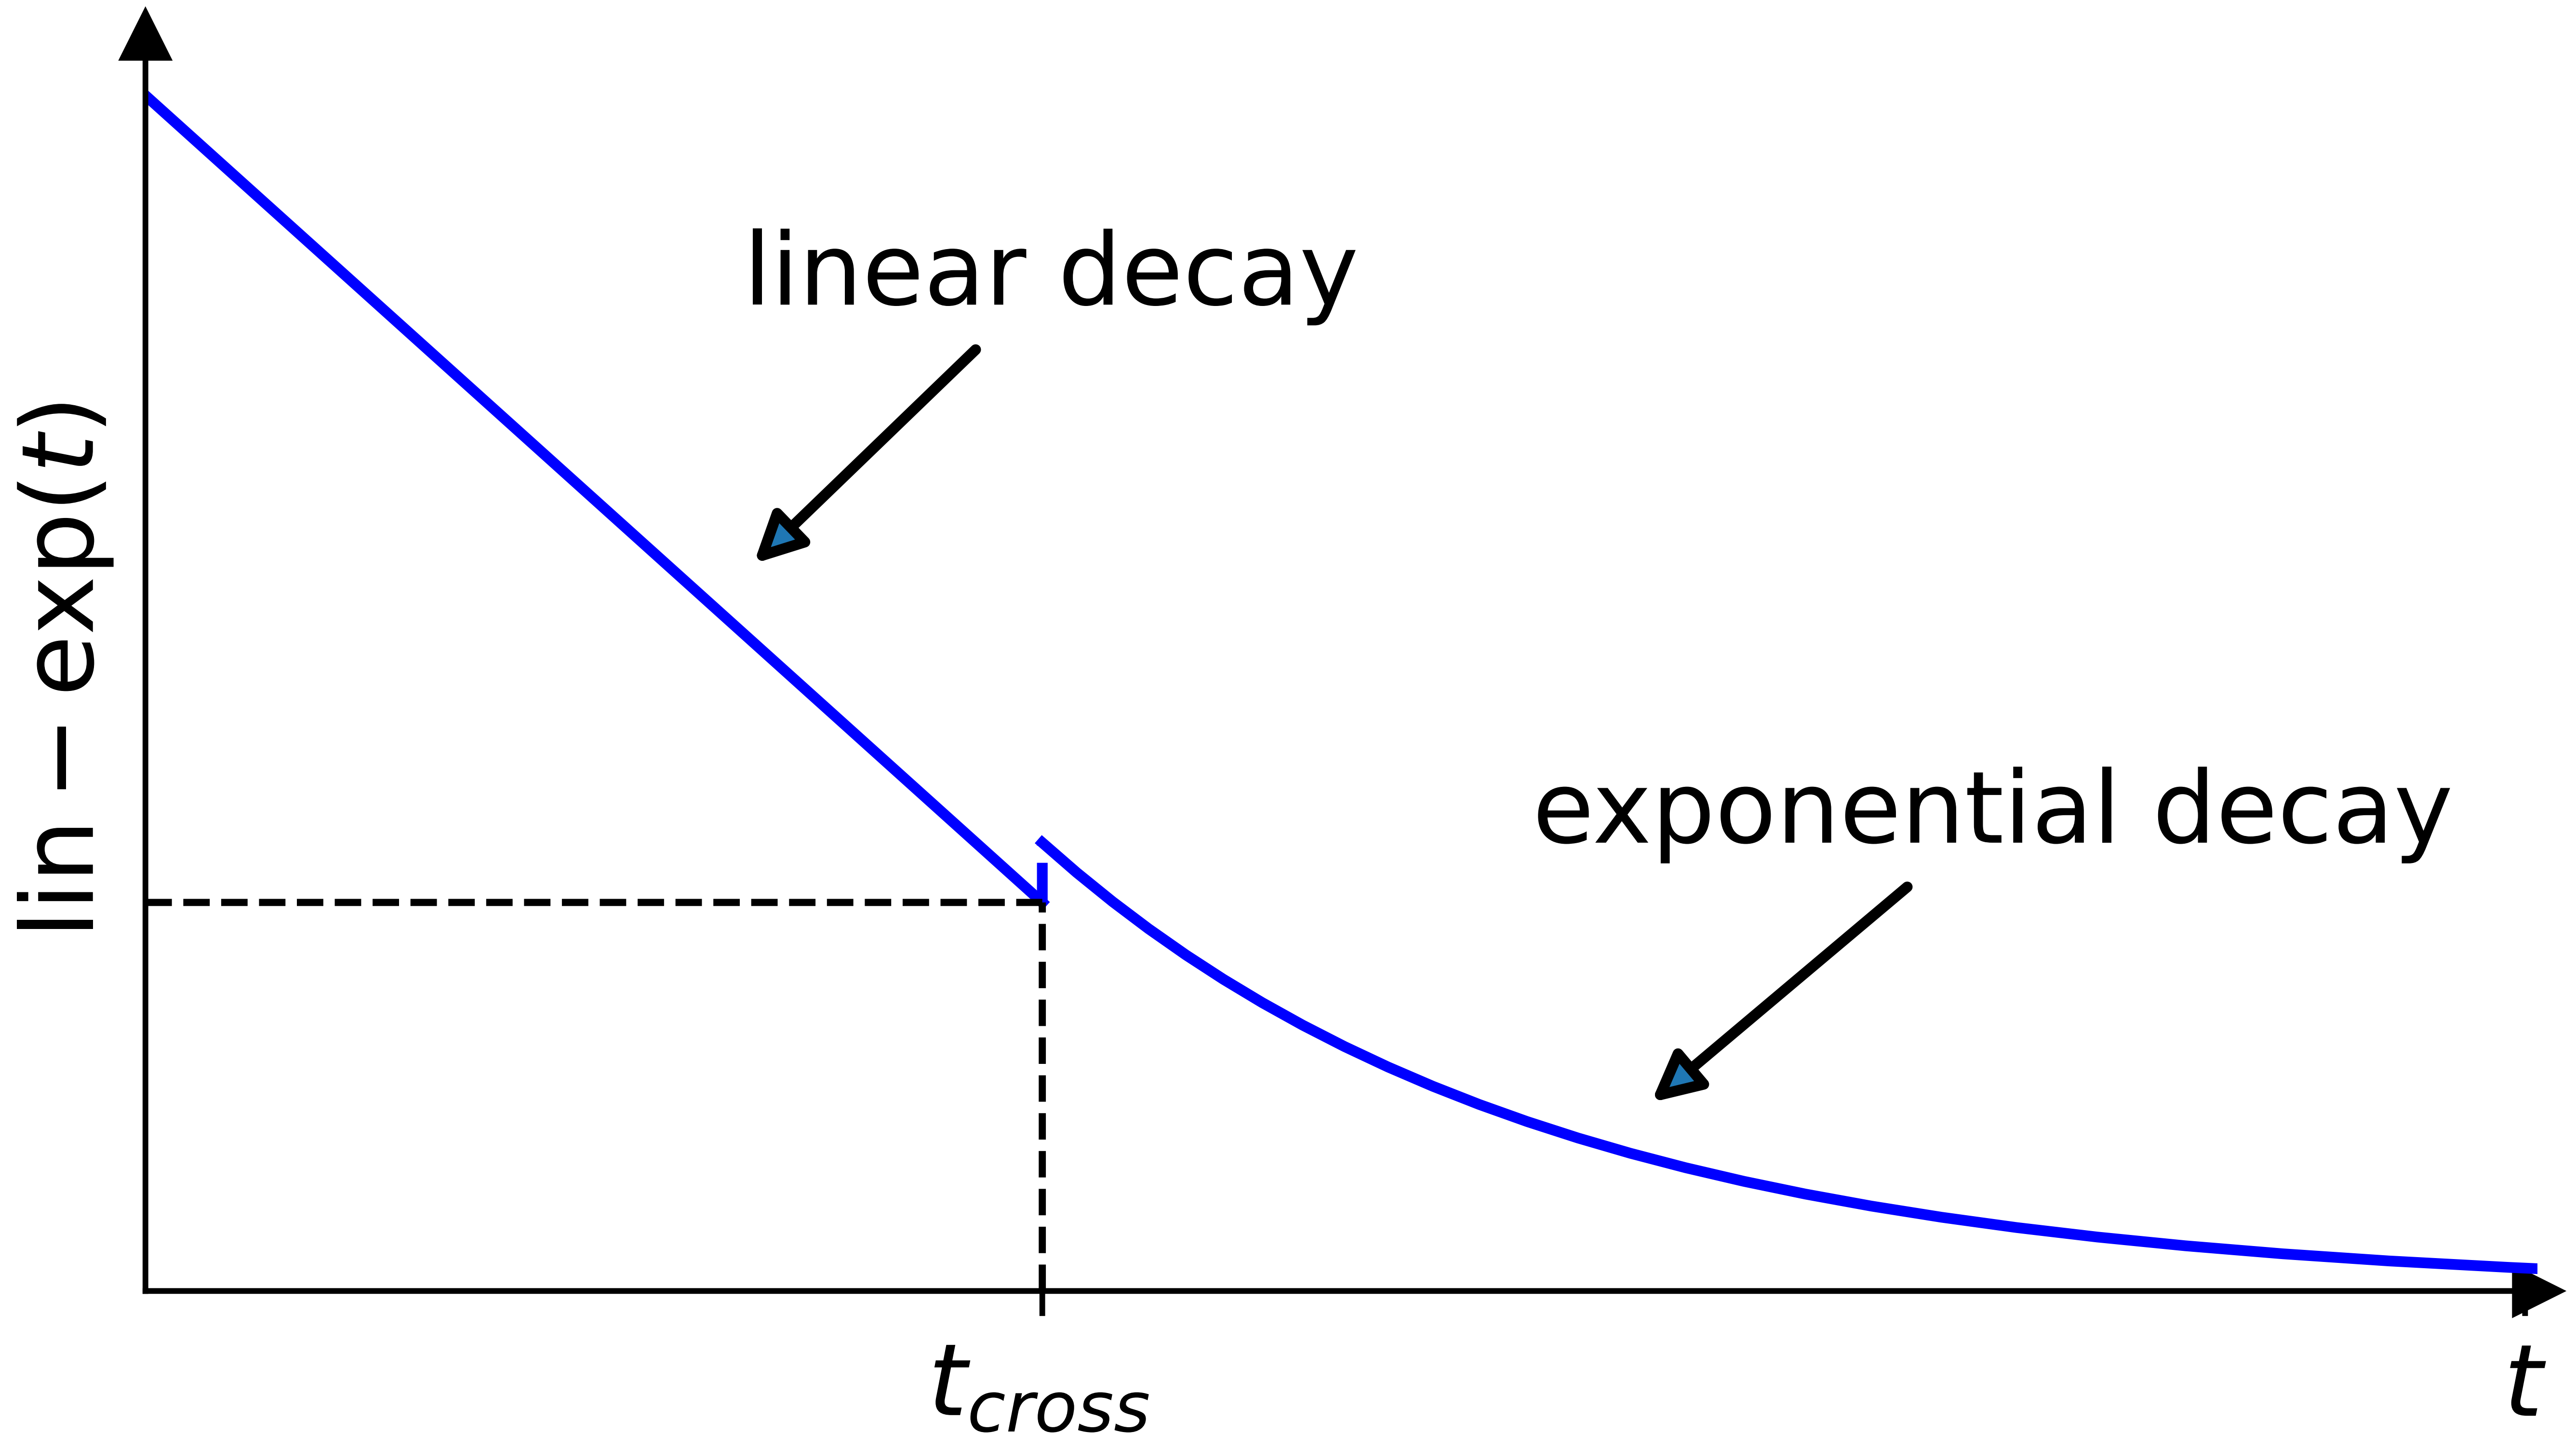

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Euclidean Contractivity of Neural Networks with Symmetric Weights", IEEE Control Systems Letters, vol. 7, pp. 1724-1729, 2023

V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

V. Centorrino, F. Bullo, GR, "Modelling and Contractivity of Neural-Synaptic Networks with Hebbian Learning", Automatica, vol. 164, 111636, 2024

Sparse Coding

$$\min_{x,D} \frac{1}{2}\|u - Dx\|^2 + \lambda S(x)$$

V. Centorrino, F. Bullo, GR, "Modelling and Contractivity of Neural-Synaptic Networks with Hebbian Learning", Automatica, vol. 164, 111636, 2024

Sparse Coding

$$\min_{x,D} \frac{1}{2}\|u - Dx\|^2 + \lambda S(x) \quad \Longrightarrow \quad \begin{cases} \dot{x}_i = -x_i + \Phi(t, x, W) \\ \dot{W}_{ij} = h(t, W_{ij}, x_i, x_j) \end{cases}$$

Embed Learning

Sparse Coding

$$\min_{x,D} \frac{1}{2}\|u - Dx\|^2 + \lambda S(x)$$

$$\Longrightarrow$$

$$\begin{cases} \dot{x}_i = -x_i + \Phi(t, x, W) \\ \dot{W}_{ij} = h(t, W_{ij}, x_i, x_j) \end{cases}$$

1. There is a synaptic change between two neurons when both are active.
2. The learning rule for the synapse W_{ij} should depend only on the activity of neuron j and neuron i.

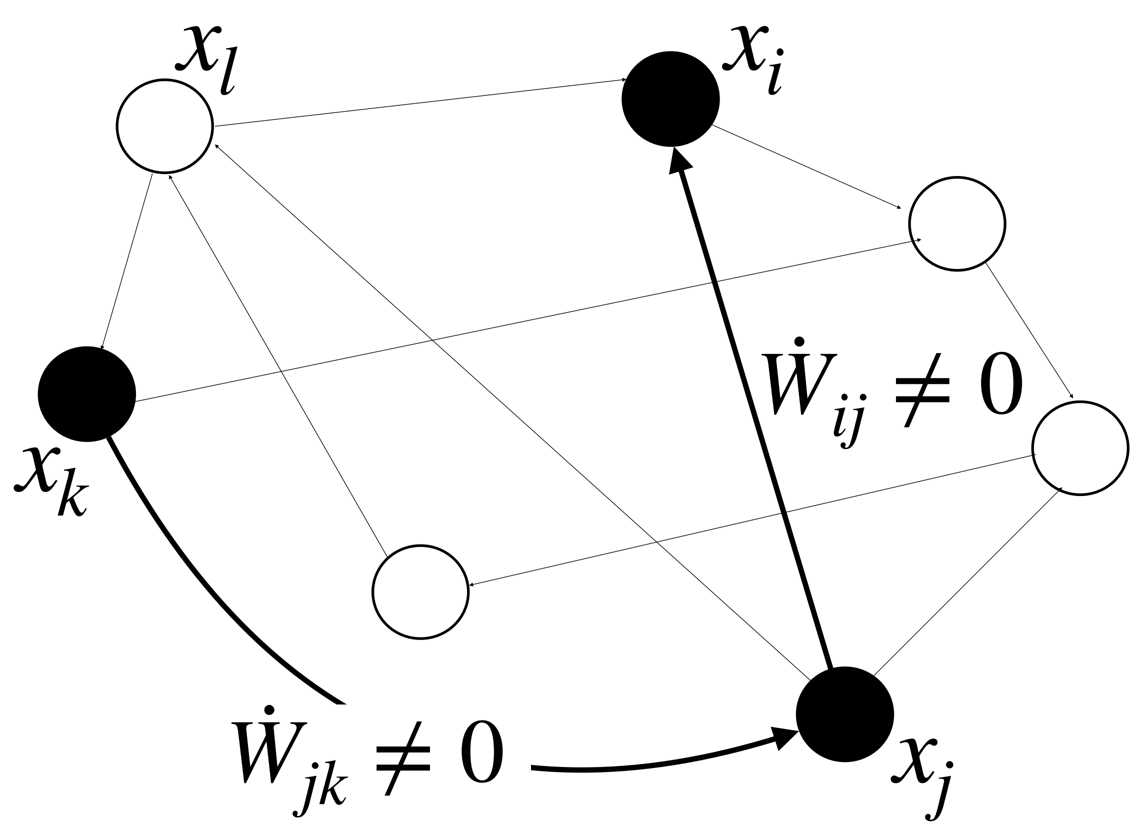

V. Centorrino, F. Bullo, GR, "Modelling and Contractivity of Neural-Synaptic Networks with Hebbian Learning", Automatica, vol. 164, 111636, 2024

Outline

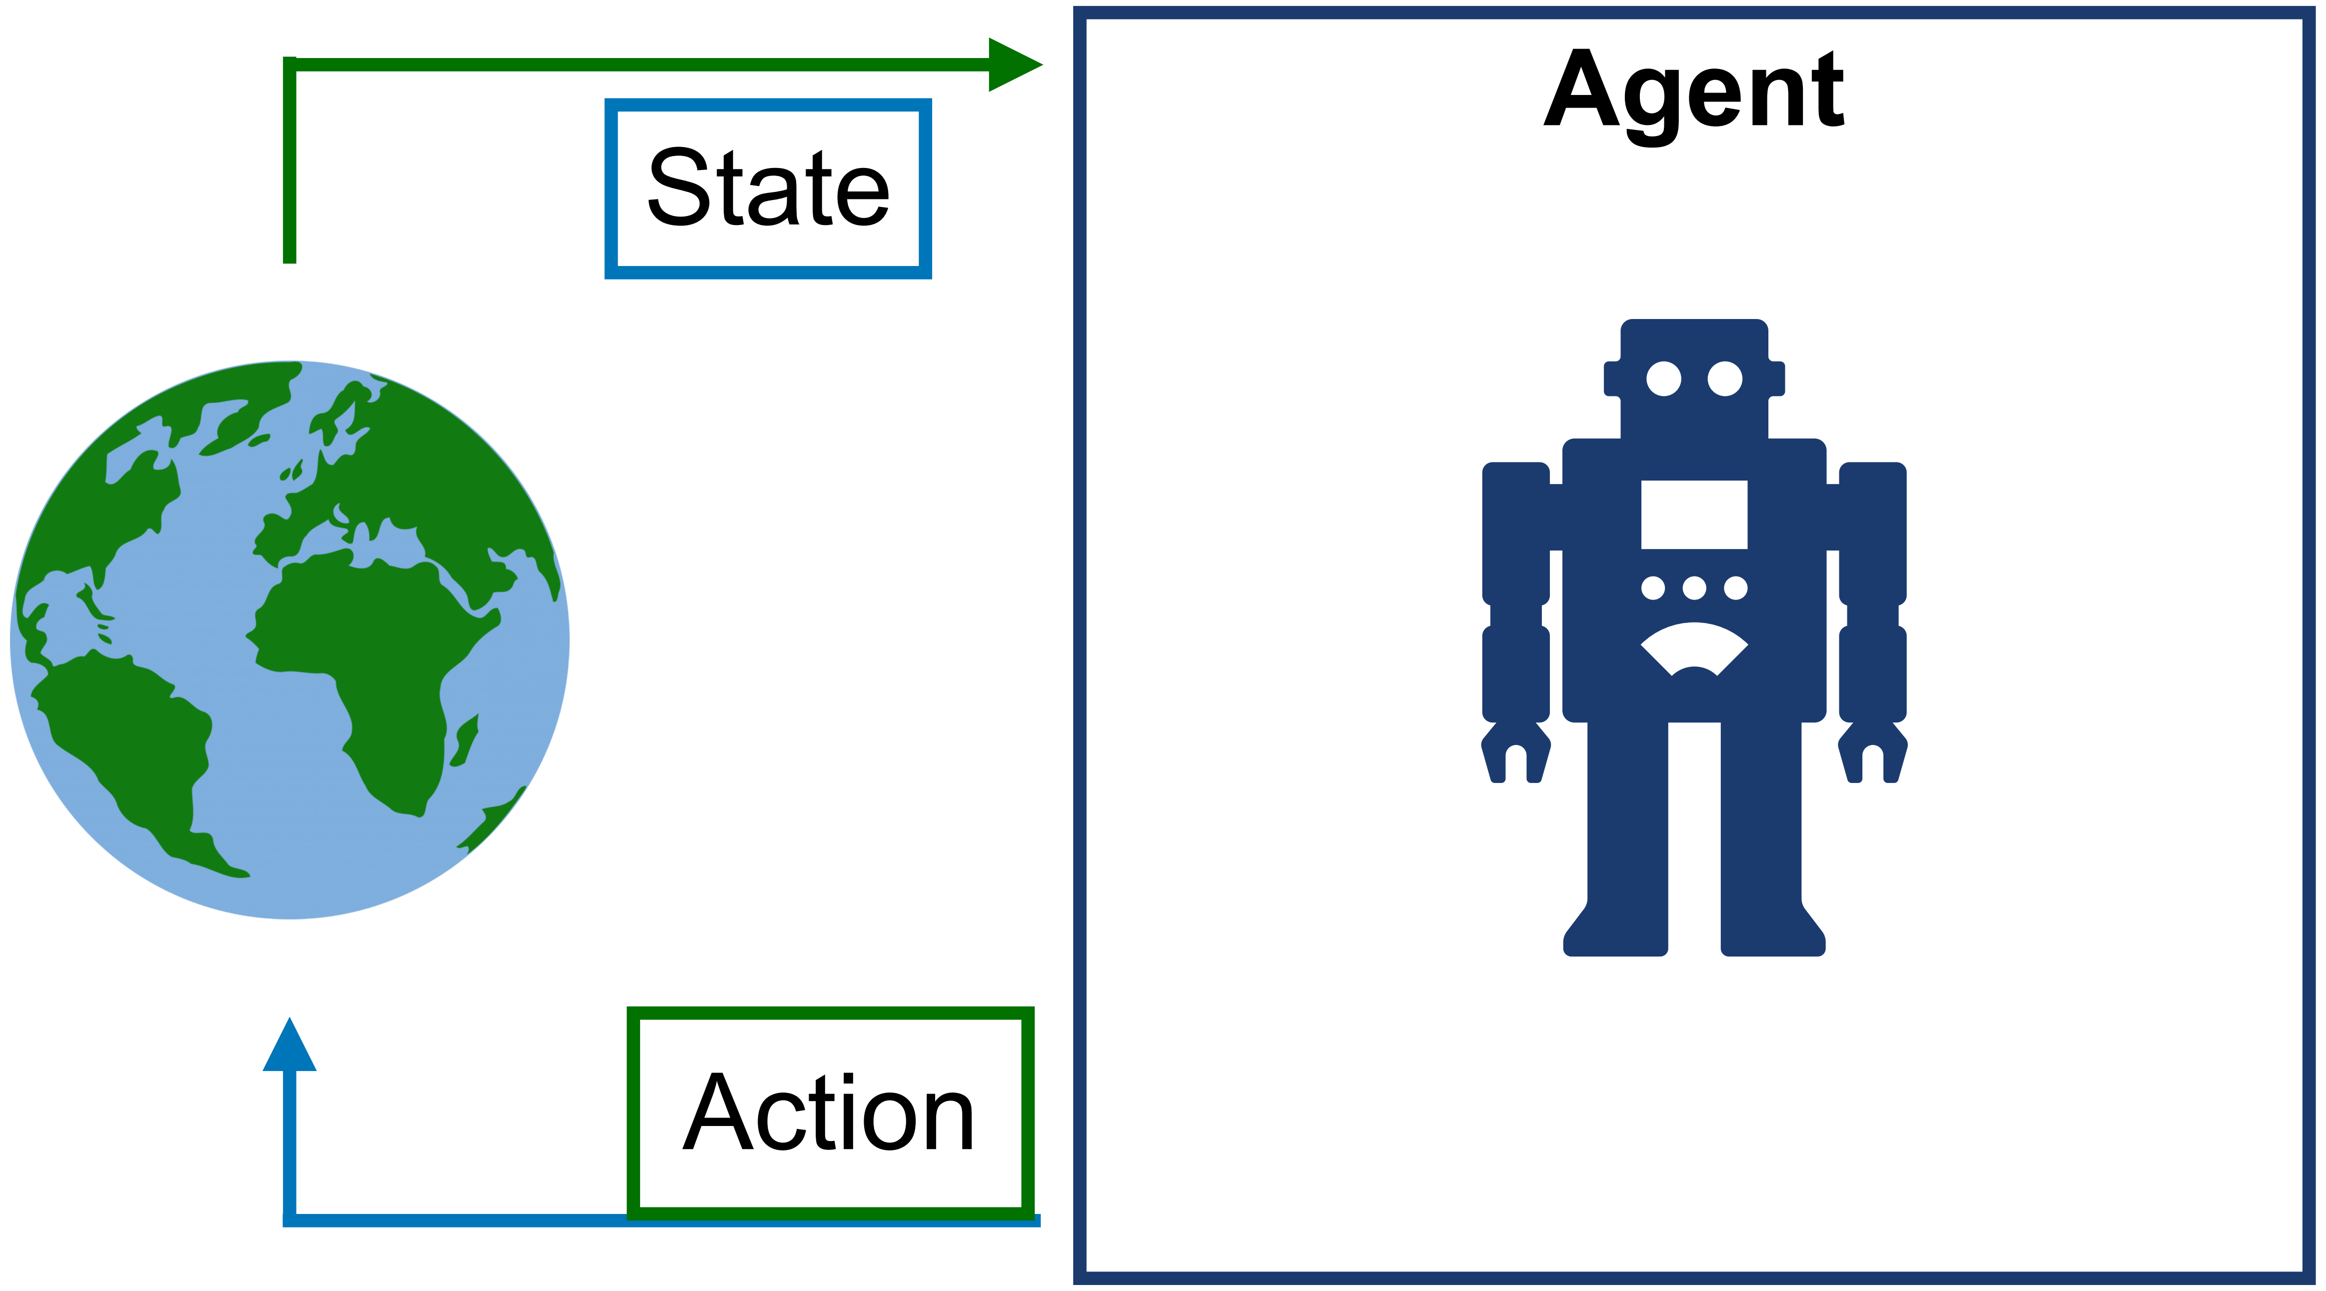

0) Proximal Flows and Neural Networks

1) Sparse reconstruction

- ☑ Optimization-based explanation
- ☑ Competitive positive firing rate dynamics

2) Policy Composition

- ☑ Free energy minimization
- ☑ Softmax flows from entropy max.

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Outline

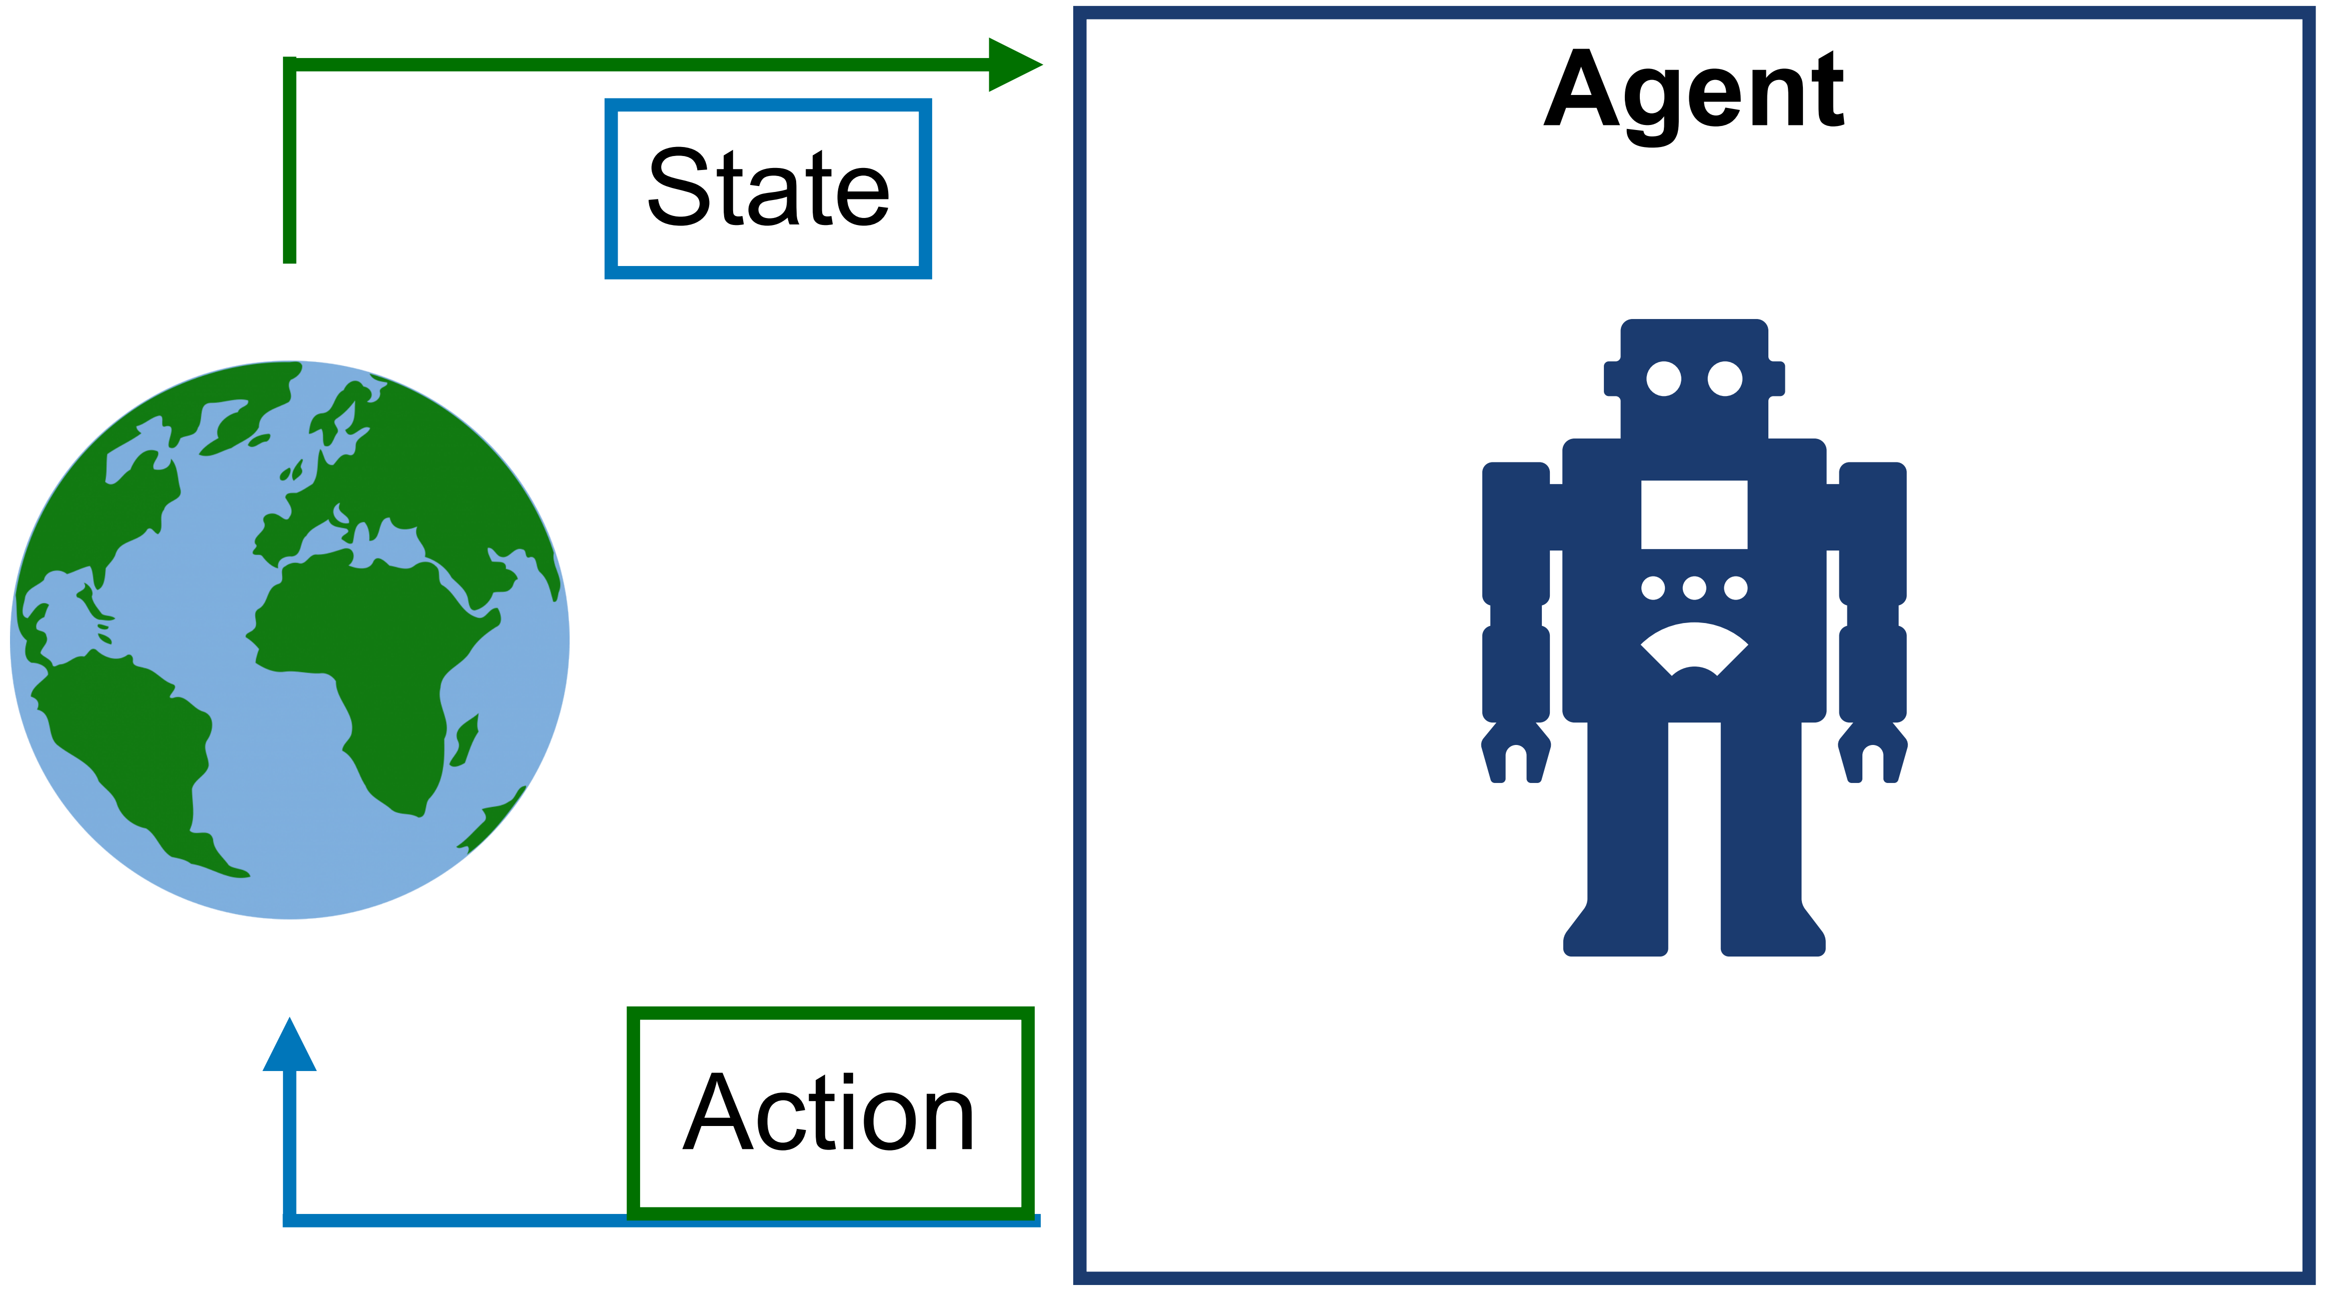

2) Policy Composition

- ☑ Free energy minimization
- ☑ Softmax flows from entropy max.

[1] V. Centorrino, A. Gokhale, A. Davydov, GR, F. Bullo, "On Weakly Contracting Dynamics for Convex Optimization", IEEE Control Systems Letters, vol. 8, pp. 1745-1750, 2024
[2] V. Centorrino, F. Bullo, GR, "Similarity Matching Networks: Hebbian Learning and Convergence Over Multiple Time Scales", 2025
[3] A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, 2025
[4] F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Free Energy Minimization

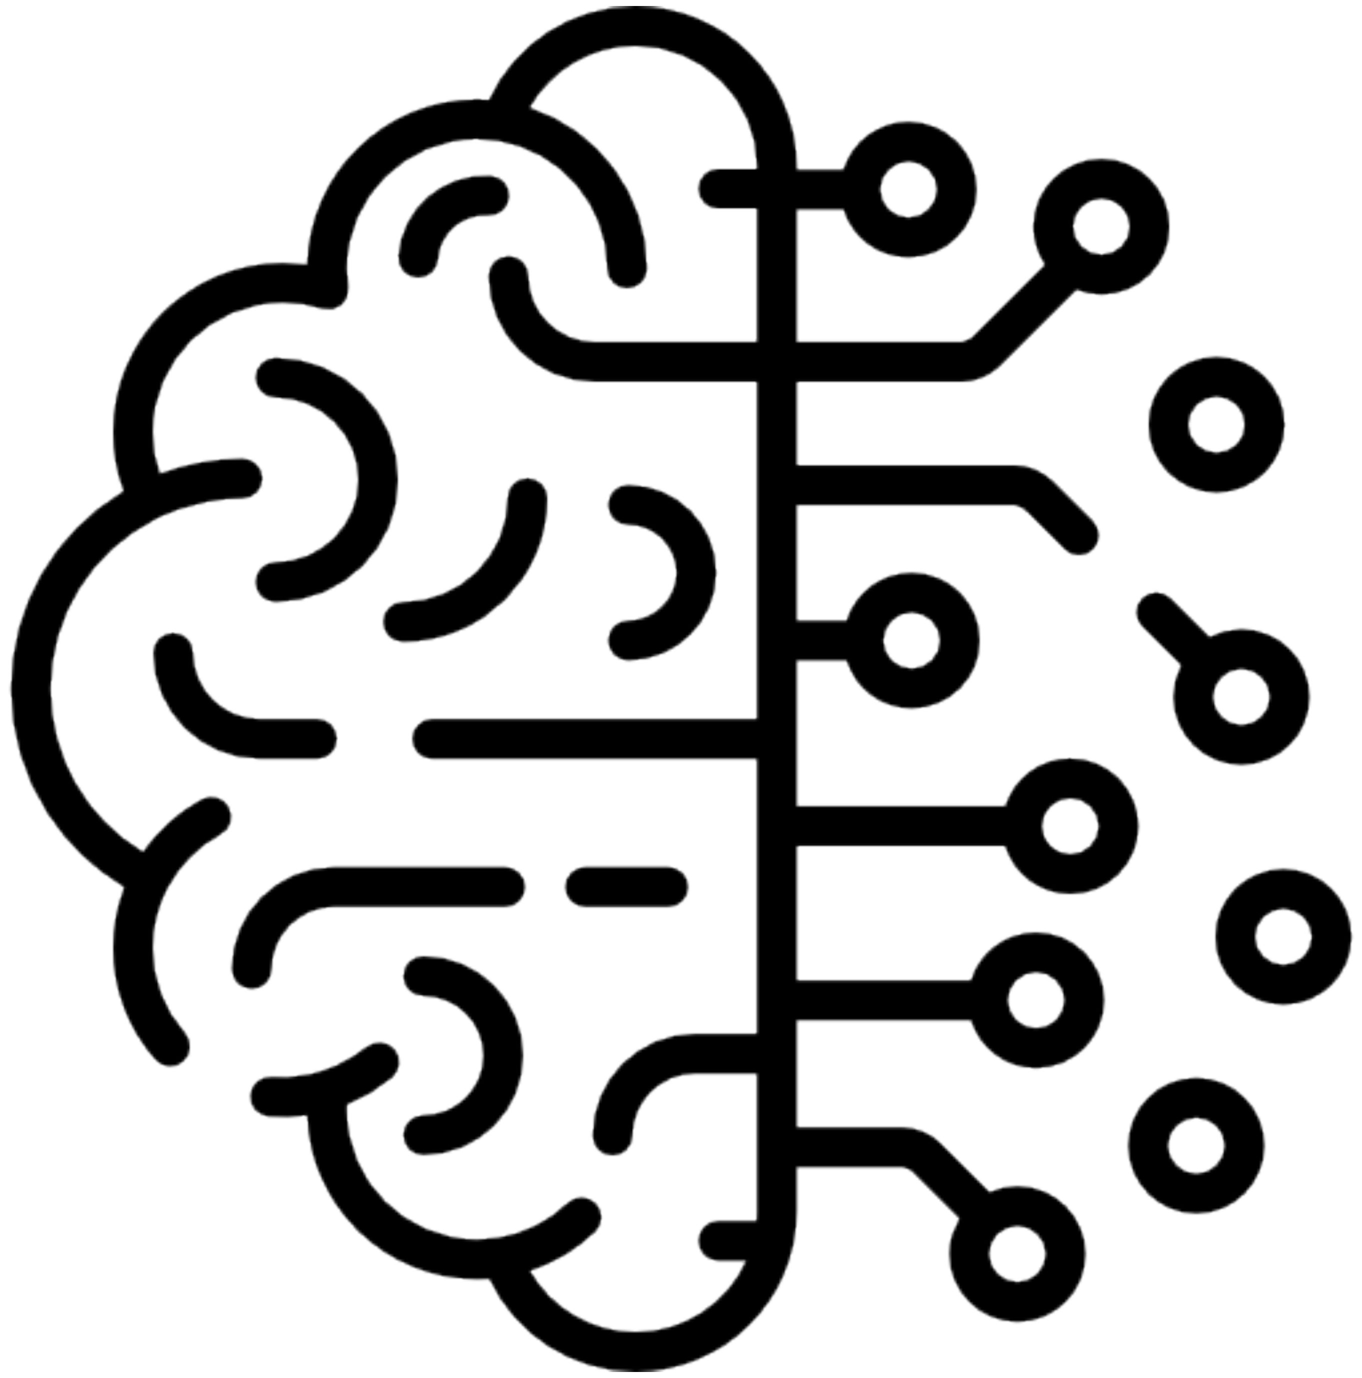

A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, in press

A unifying theory for:

- ✓self-organization
- ✓brain function
- ✓sentient behaviors
- ✓…

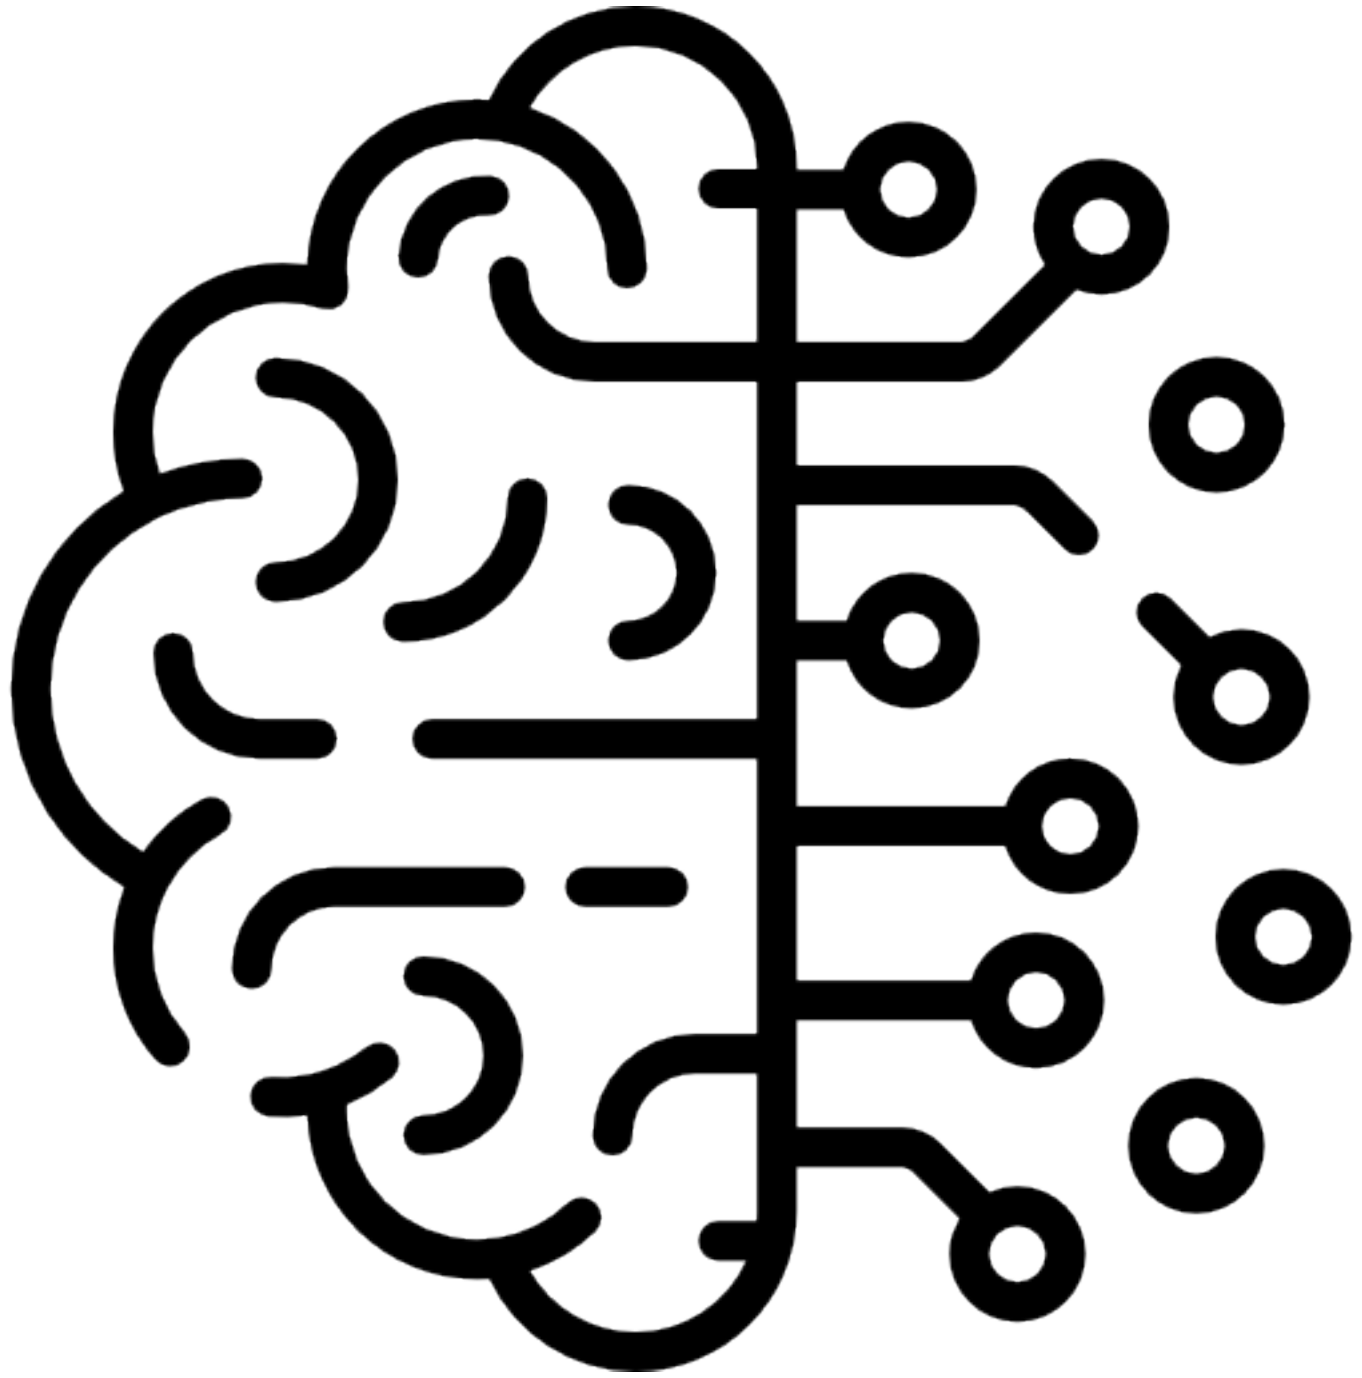

A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, in press

Free Energy Minimization

A unifying theory for:

- ✓self-organization
- ✓brain function
- ✓sentient behaviors
- ✓…

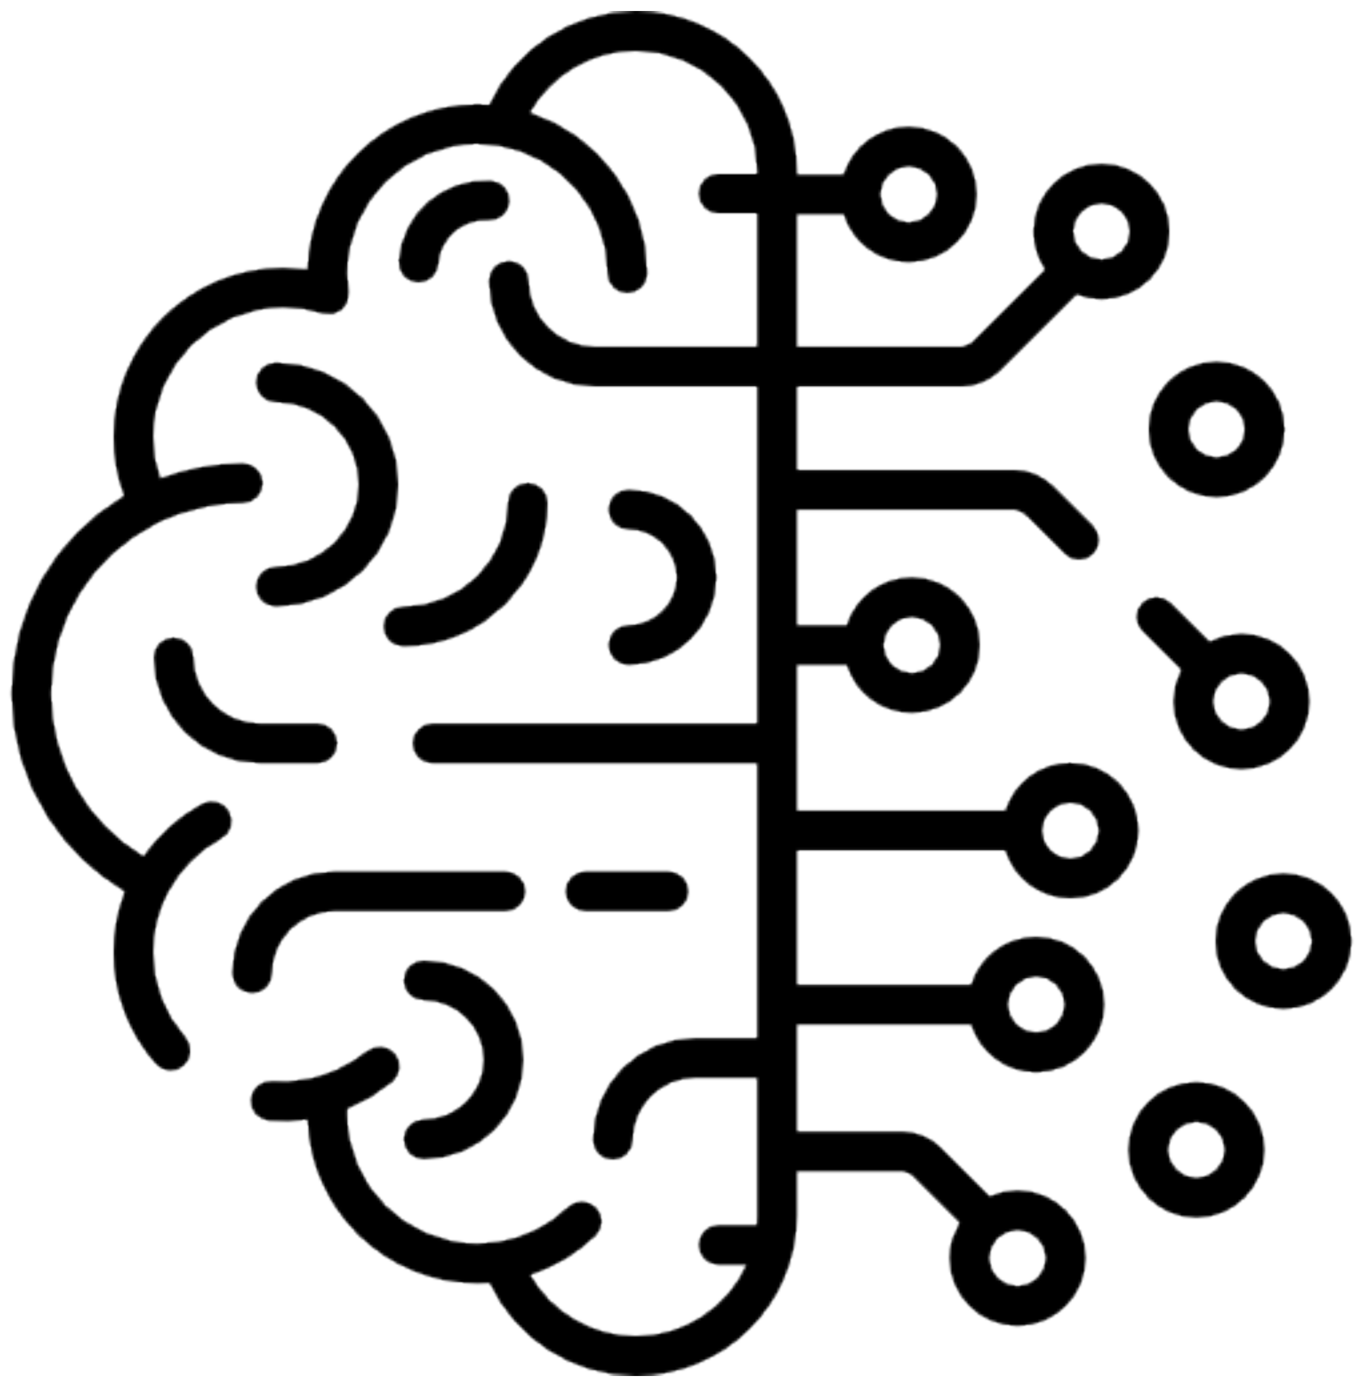

A unifying framework for:

- ✓mirror descent
- ✓maximum entropy
- ✓diffusion processes
- ✓…

A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, in press

Free Energy Minimization

A unifying theory for:

- ✓self-organization
- ✓brain function
- ✓sentient behaviors
- ✓…

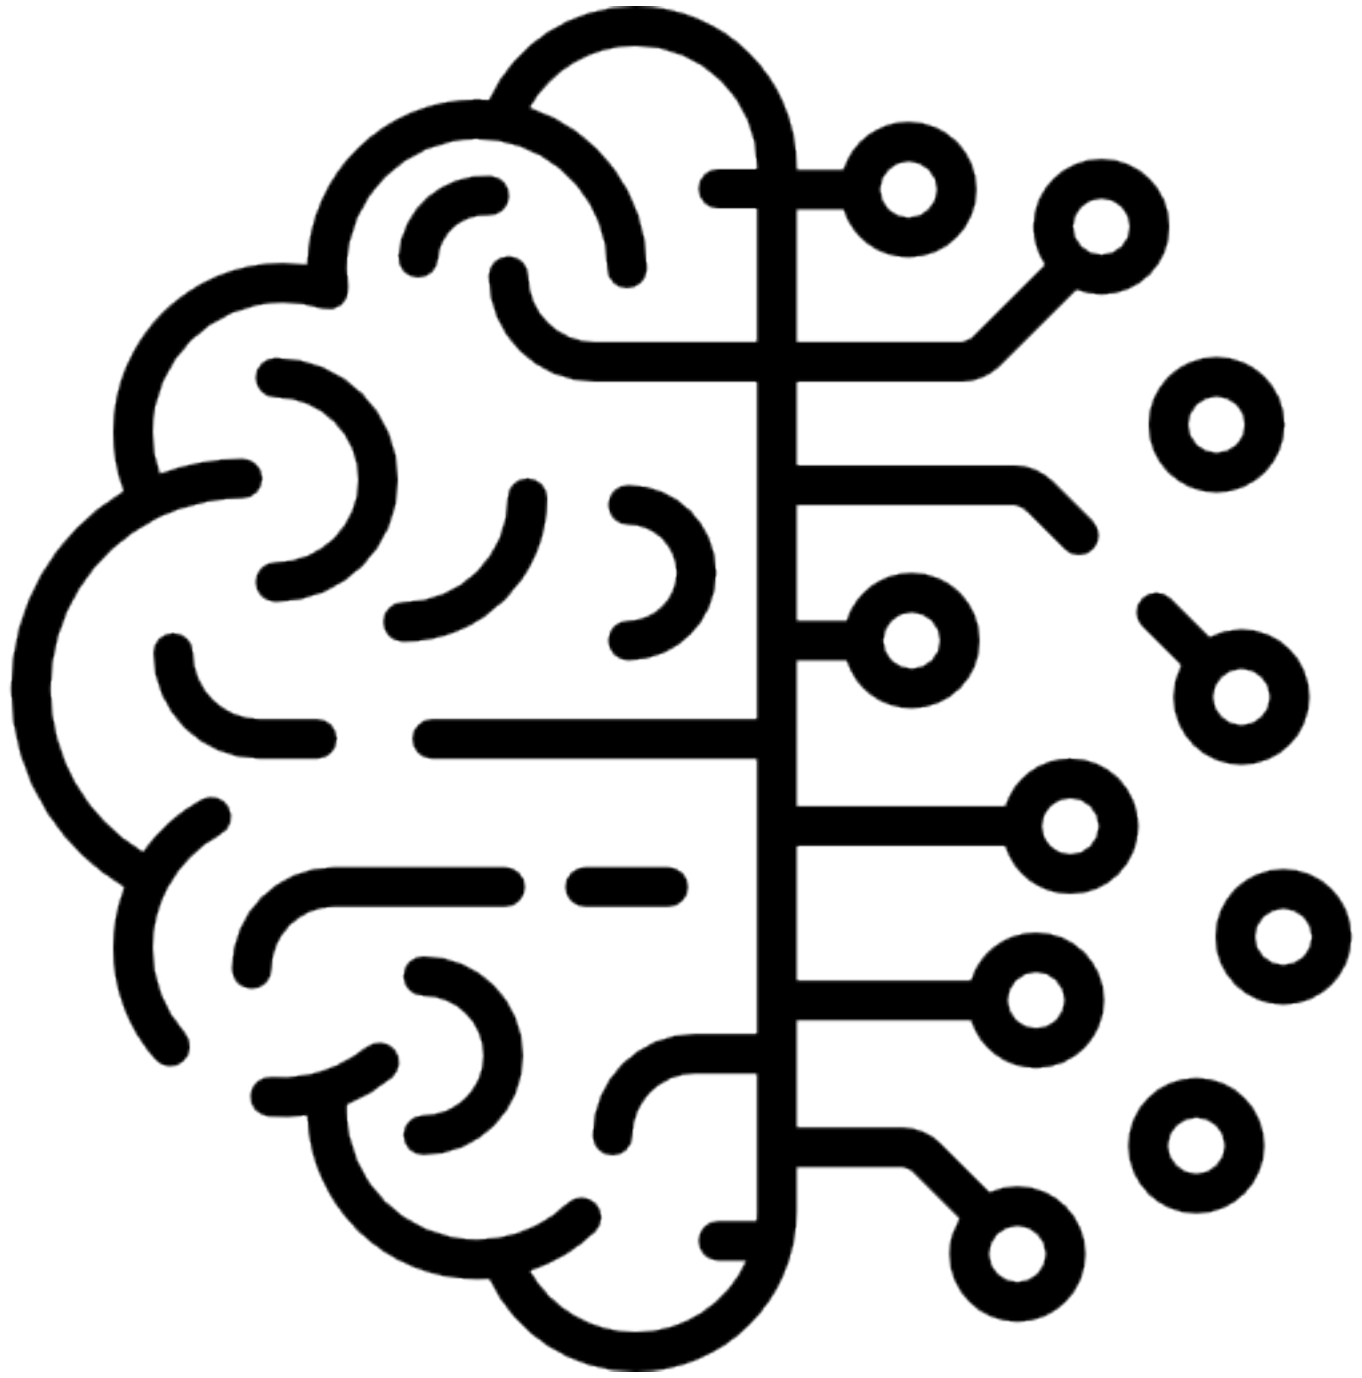

A unifying framework for:

- ✓mirror descent
- ✓maximum entropy
- ✓diffusion processes
- ✓…

Autonomous Decision-Making!

A. Shafiei, H. Jesawada, K. Friston, GR, "Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, in press

Normative Framework for Policy Composition

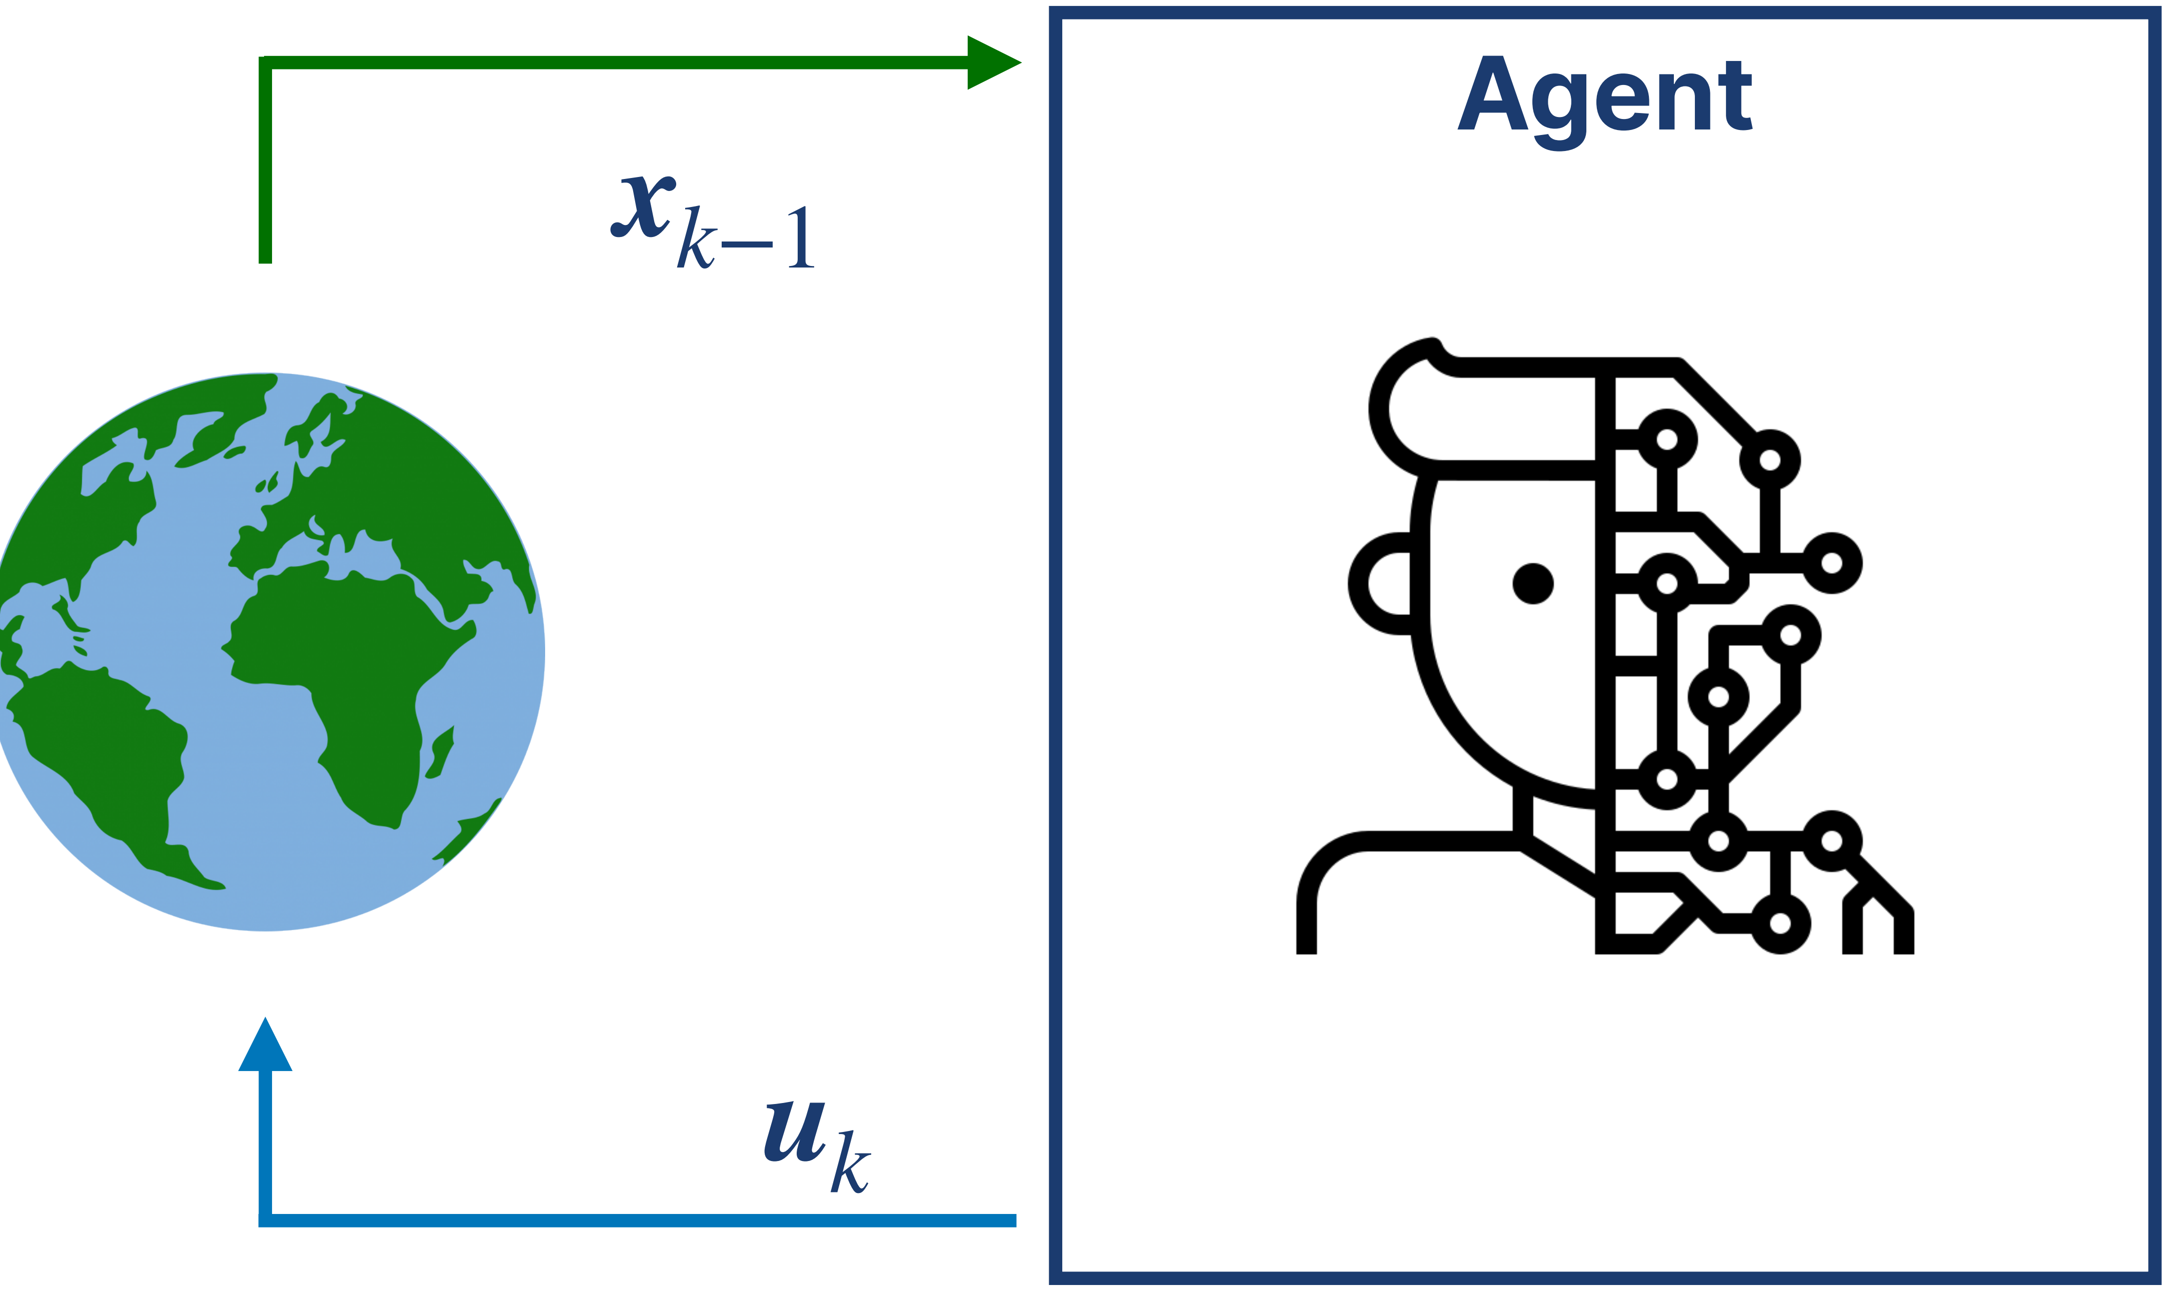

Agent/Environment dynamics:

$$p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k\right) p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

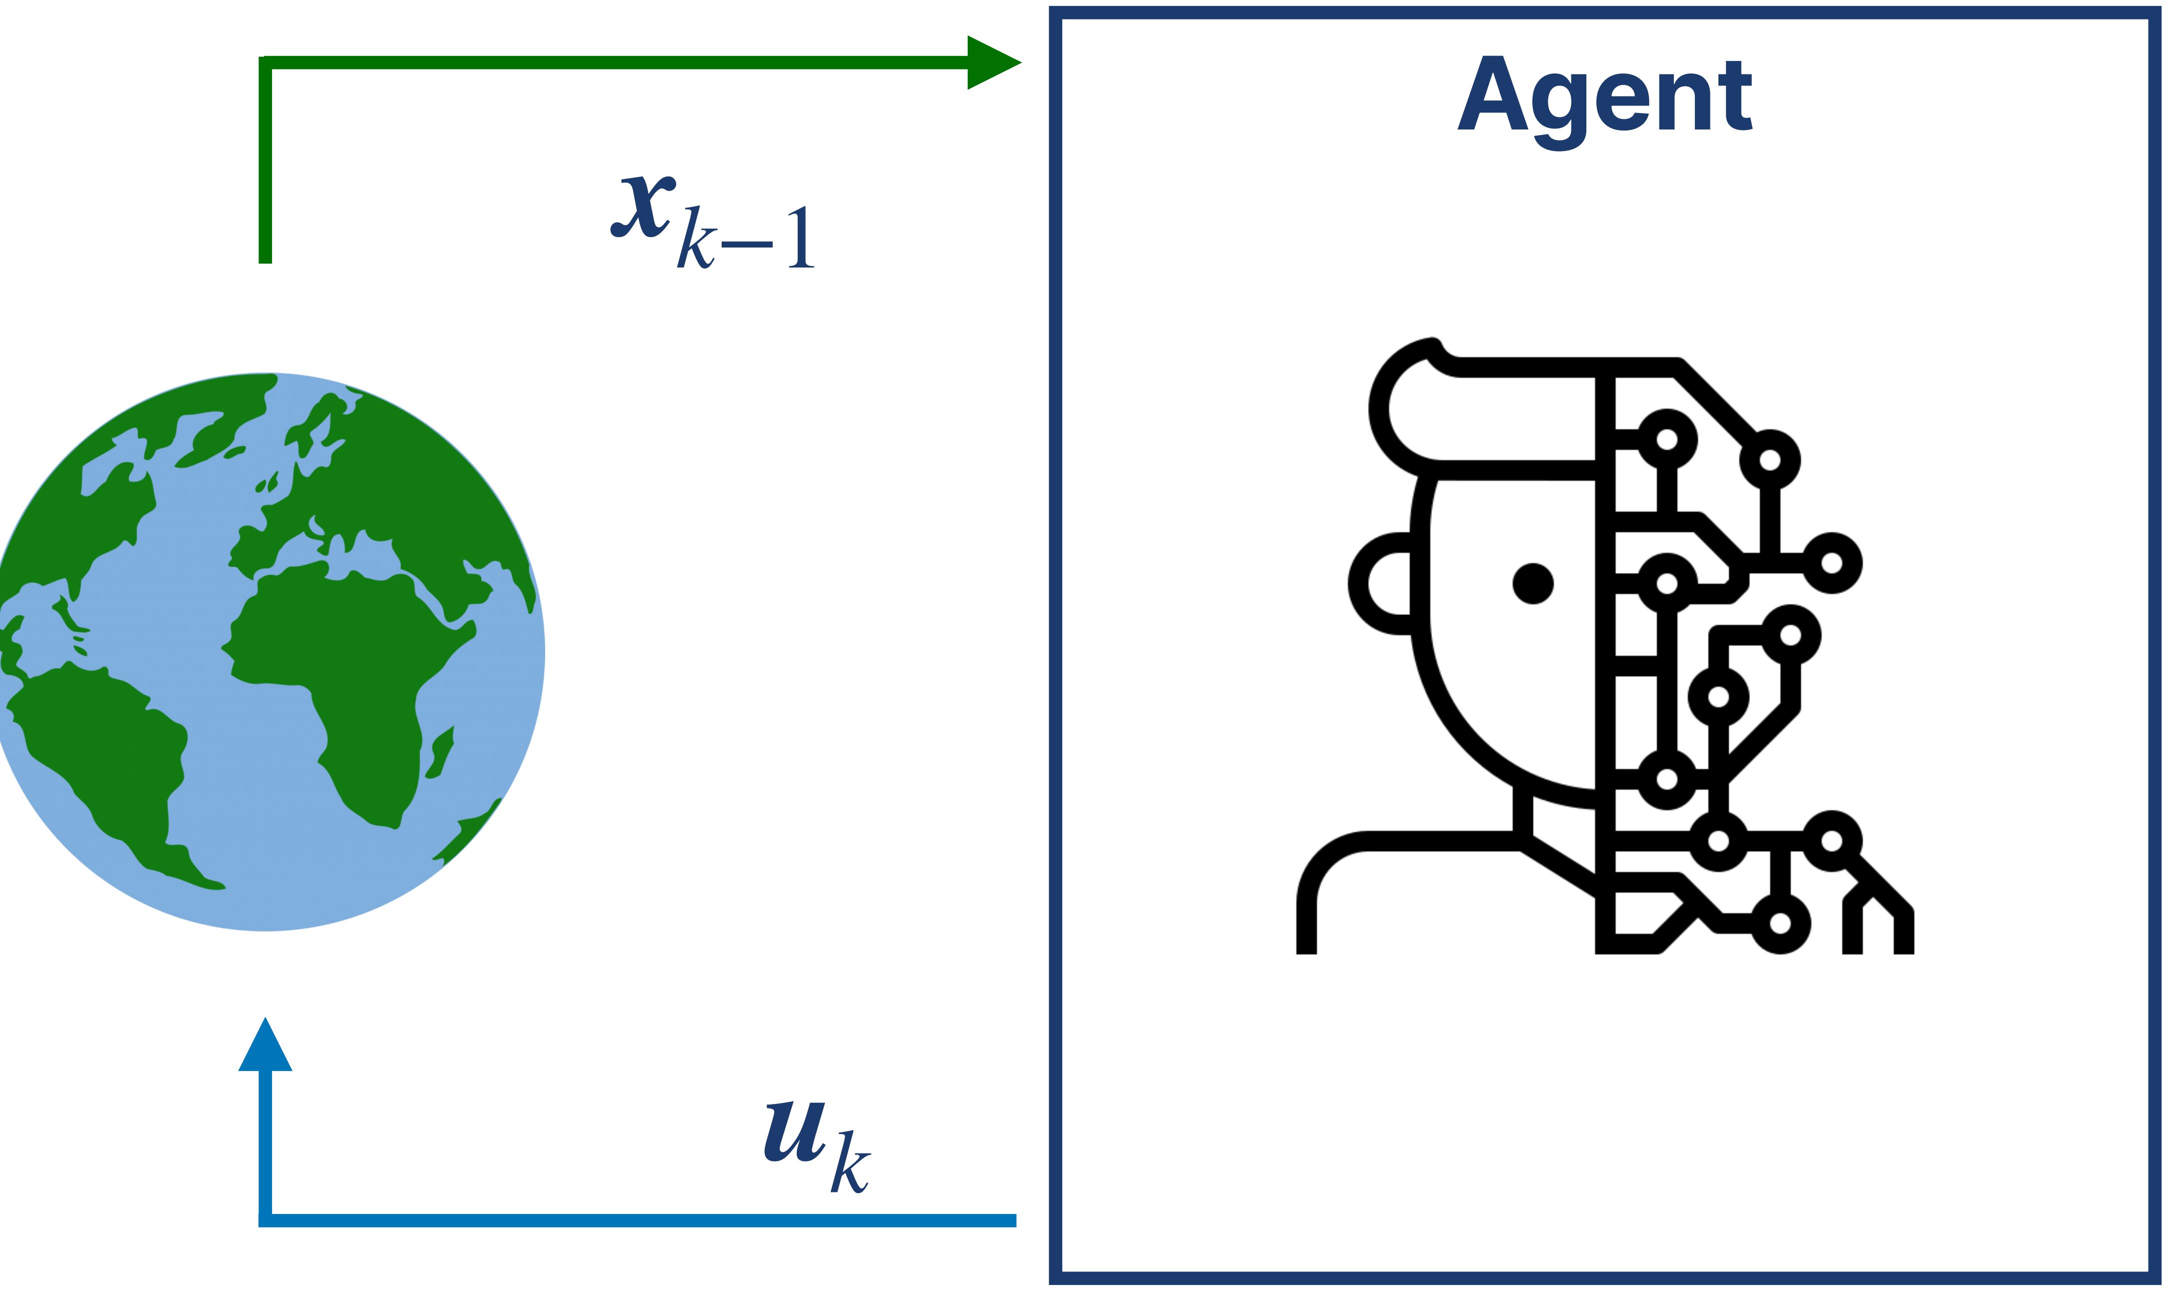

Minimize $\mathbf{w}_k$: **Error + Entropy**

Subject to: **policy = mixture($\mathbf{w}_k$) of primitives**

Agent/Environment dynamics:

$$p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k\right) p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

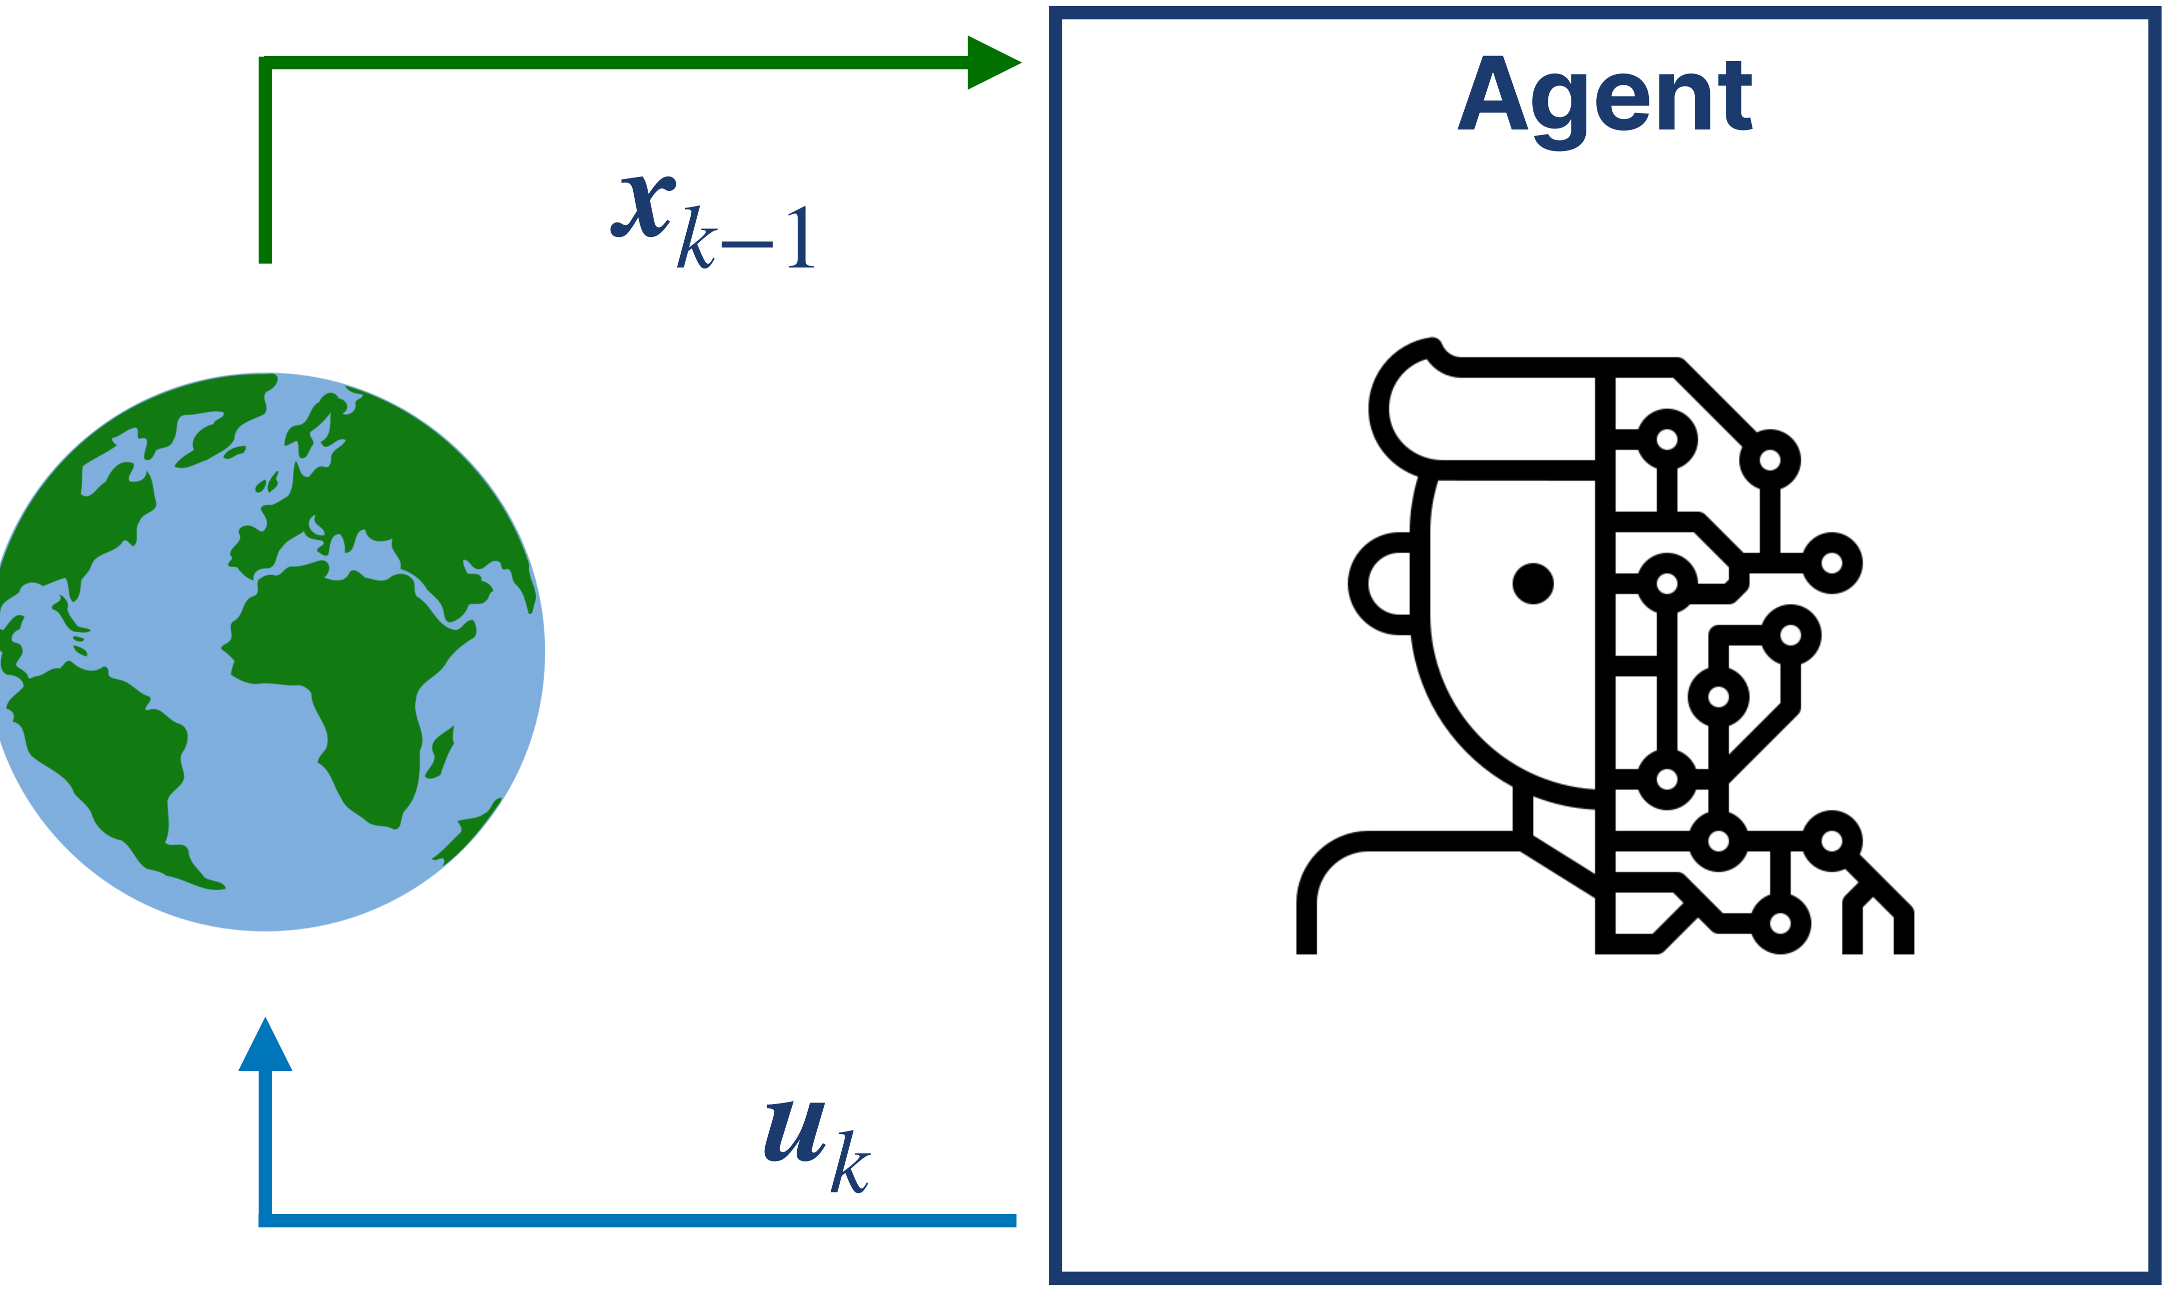

Minimize $\mathbf{w}_k$: **Error + Entropy**

Subject to: **policy = mixture($\mathbf{w}_k$) of primitives**

- Error: $D_{\mathrm{KL}}\left(p(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}) \mid\mid q(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1})\right)$

Agent/Environment dynamics:

$$p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k\right) p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Normative Framework for Policy Composition

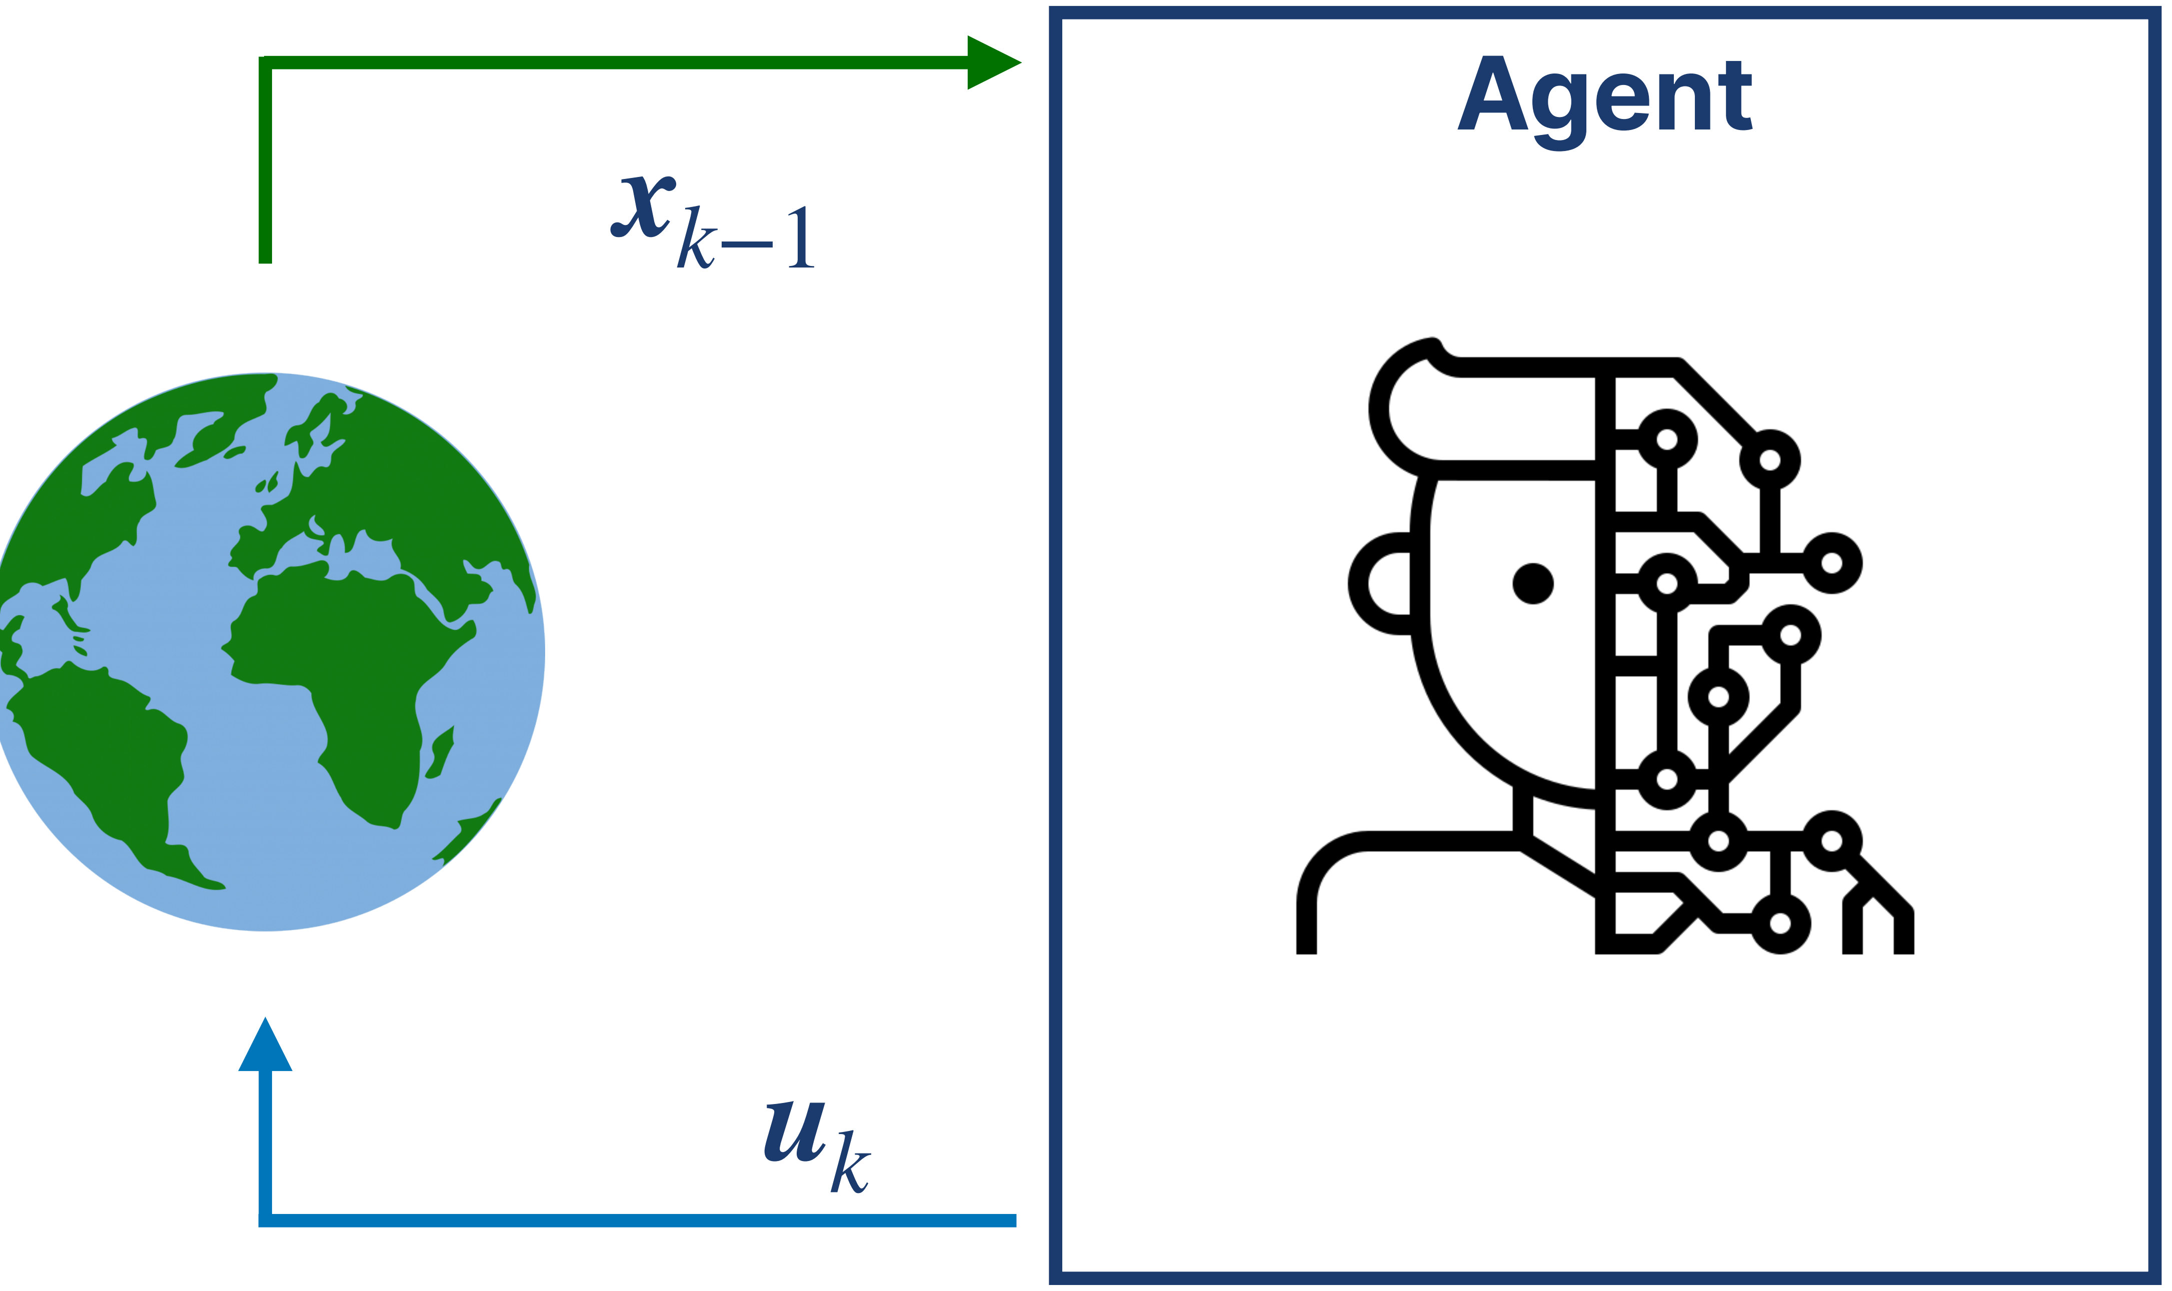

Minimize $\mathbf{w}_k$: **Error + Entropy**

Subject to: **policy = mixture($\mathbf{w}_k$) of primitives**

- Error: $D_{\mathrm{KL}}\left(p(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}) \mid\mid q(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1})\right)$

Agent/Environment dynamics:

$$p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k\right) p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)$$

$$\boxed{D_{\mathrm{KL}}\left(p \mid\mid q\right) := \int p(\mathbf{z}) \ln \frac{p(\mathbf{z})}{q(\mathbf{z})} d\mathbf{z}}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Optimization Problem

$$\min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\mathrm{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \mid\mid q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \overbrace{\mathrm{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\mathsf{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \mid\mid q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \overbrace{\mathrm{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\text{s.t. } p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}_{\text{Mixture-of-primitives}}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

$$\mathbf{w}_k^\star \in \arg \min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\mathrm{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \,||\, q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \overbrace{\mathrm{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\text{s.t. } p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha \left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}_{\text{Mixture-of-primitives}}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Optimization Problem

$$\boxed{\mathbf{w}_k^\star \ \in \ \text{arg}} \ \min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\text{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \,||\, q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \ \overbrace{\text{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\text{s.t.} \ \ p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha \left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}_{\text{Mixture-of-primitives}}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Optimization Problem

$$\boxed{\mathbf{w}_k^\star \in \arg} \min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\mathrm{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \,||\, q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \overbrace{\mathrm{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\text{s.t.} \quad p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}_{\text{Mixture-of-primitives}}$$

$$\boxed{p_{\mathbf{u}}^\star\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \sum_{\alpha=1}^{n_\pi} \mathbf{w}_k^{\alpha,\star} \pi^\alpha\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}$$

$$\boxed{\mathbf{w}_k^\star \in \text{arg}} \min_{\mathbf{w}_k \in \Delta_{n_\pi}} \overbrace{D_{\mathrm{KL}}\left(p\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right) \,||\, q\left(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}\right)\right)}^{\text{Error}} - \varepsilon \overbrace{\mathrm{H}(\mathbf{w}_k)}^{\text{Entropy}}$$

$$\text{s.t. } p\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}_{\text{Mixture-of-primitives}}$$

Agent Action

$$\mathbf{u}_k^\star \sim \boxed{p_{\mathbf{u}}^\star\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right) = \sum_{\alpha=1}^{n_\pi} \mathbf{w}_k^{\alpha,\star} \pi^\alpha\left(\mathbf{u}_k \mid \mathbf{x}_{k-1}\right)}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

The Softmax Gradient Flow

Goal: dynamical system that provably converges to optimal solution

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Goal: dynamical system that provably converges to optimal solution

$$\tau\dot{\mathbf{w}}_k = -\,\mathbf{w}_k + \mathrm{softmax}\left(-\varepsilon^{-1}\,\nabla\,\mathsf{Error}(\mathbf{w}_k)\right)$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

The Softmax Gradient Flow

Goal: dynamical system that provably converges to optimal solution

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \mathrm{softmax}\left(-\varepsilon^{-1} \nabla \mathsf{Error}(\mathbf{w}_k)\right)$$

$$\mathrm{softmax}(\mathbf{x})_i = \frac{\exp(x_i)}{\sum_i \exp(x_i)}$$

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

The Softmax Gradient Flow

Goal: dynamical system that provably converges to optimal solution

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \mathrm{softmax}\left(-\varepsilon^{-1} \nabla \mathsf{Error}(\mathbf{w}_k)\right)$$

- Equilibrium is the **optimal solution**
- Simplex is **forward invariant**
- Naturally interpreted as **firing-rate**

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

The Softmax Gradient Flow

Goal: dynamical system that provably converges to optimal solution

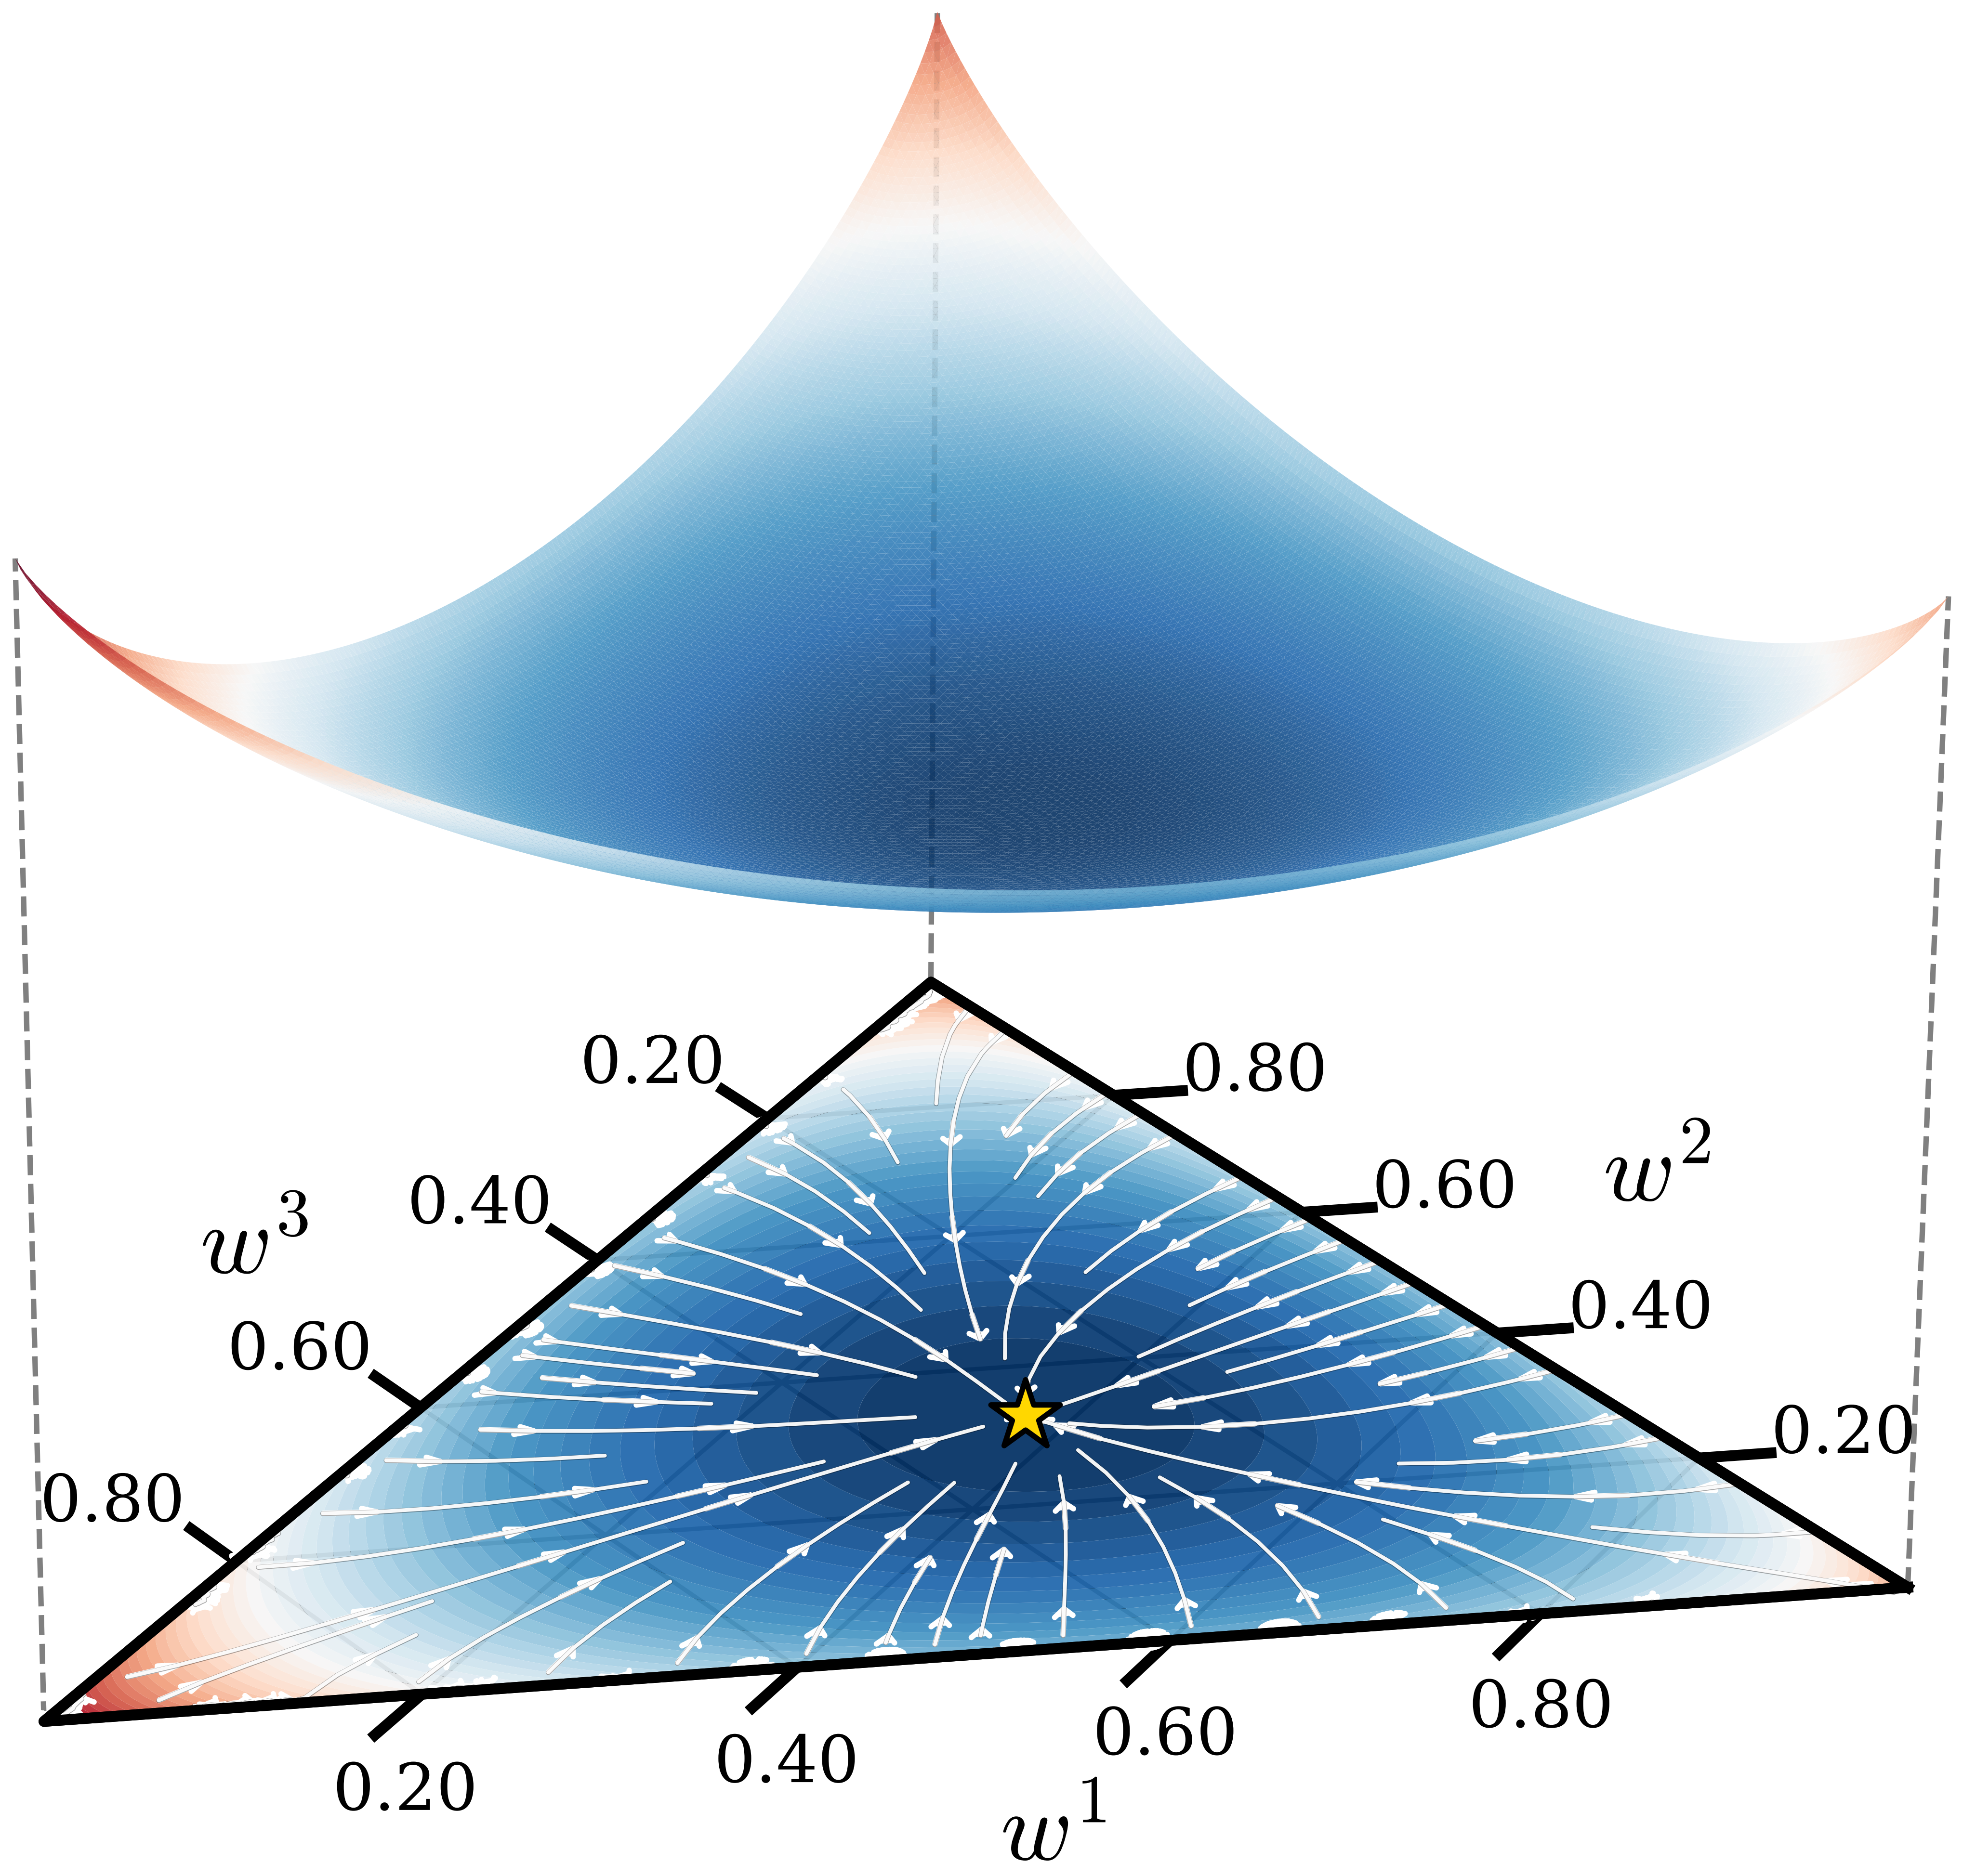

$$\tau \dot{\mathbf{w}}_k = -\,\mathbf{w}_k + \text{softmax}\left(-\varepsilon^{-1} \nabla \text{Error}(\mathbf{w}_k)\right)$$

- Equilibrium is the **optimal solution**
- Simplex is **forward invariant**
- Naturally interpreted as **firing-rate**

The Softmax Gradient Flow

Goal: dynamical system that provably converges to optimal solution

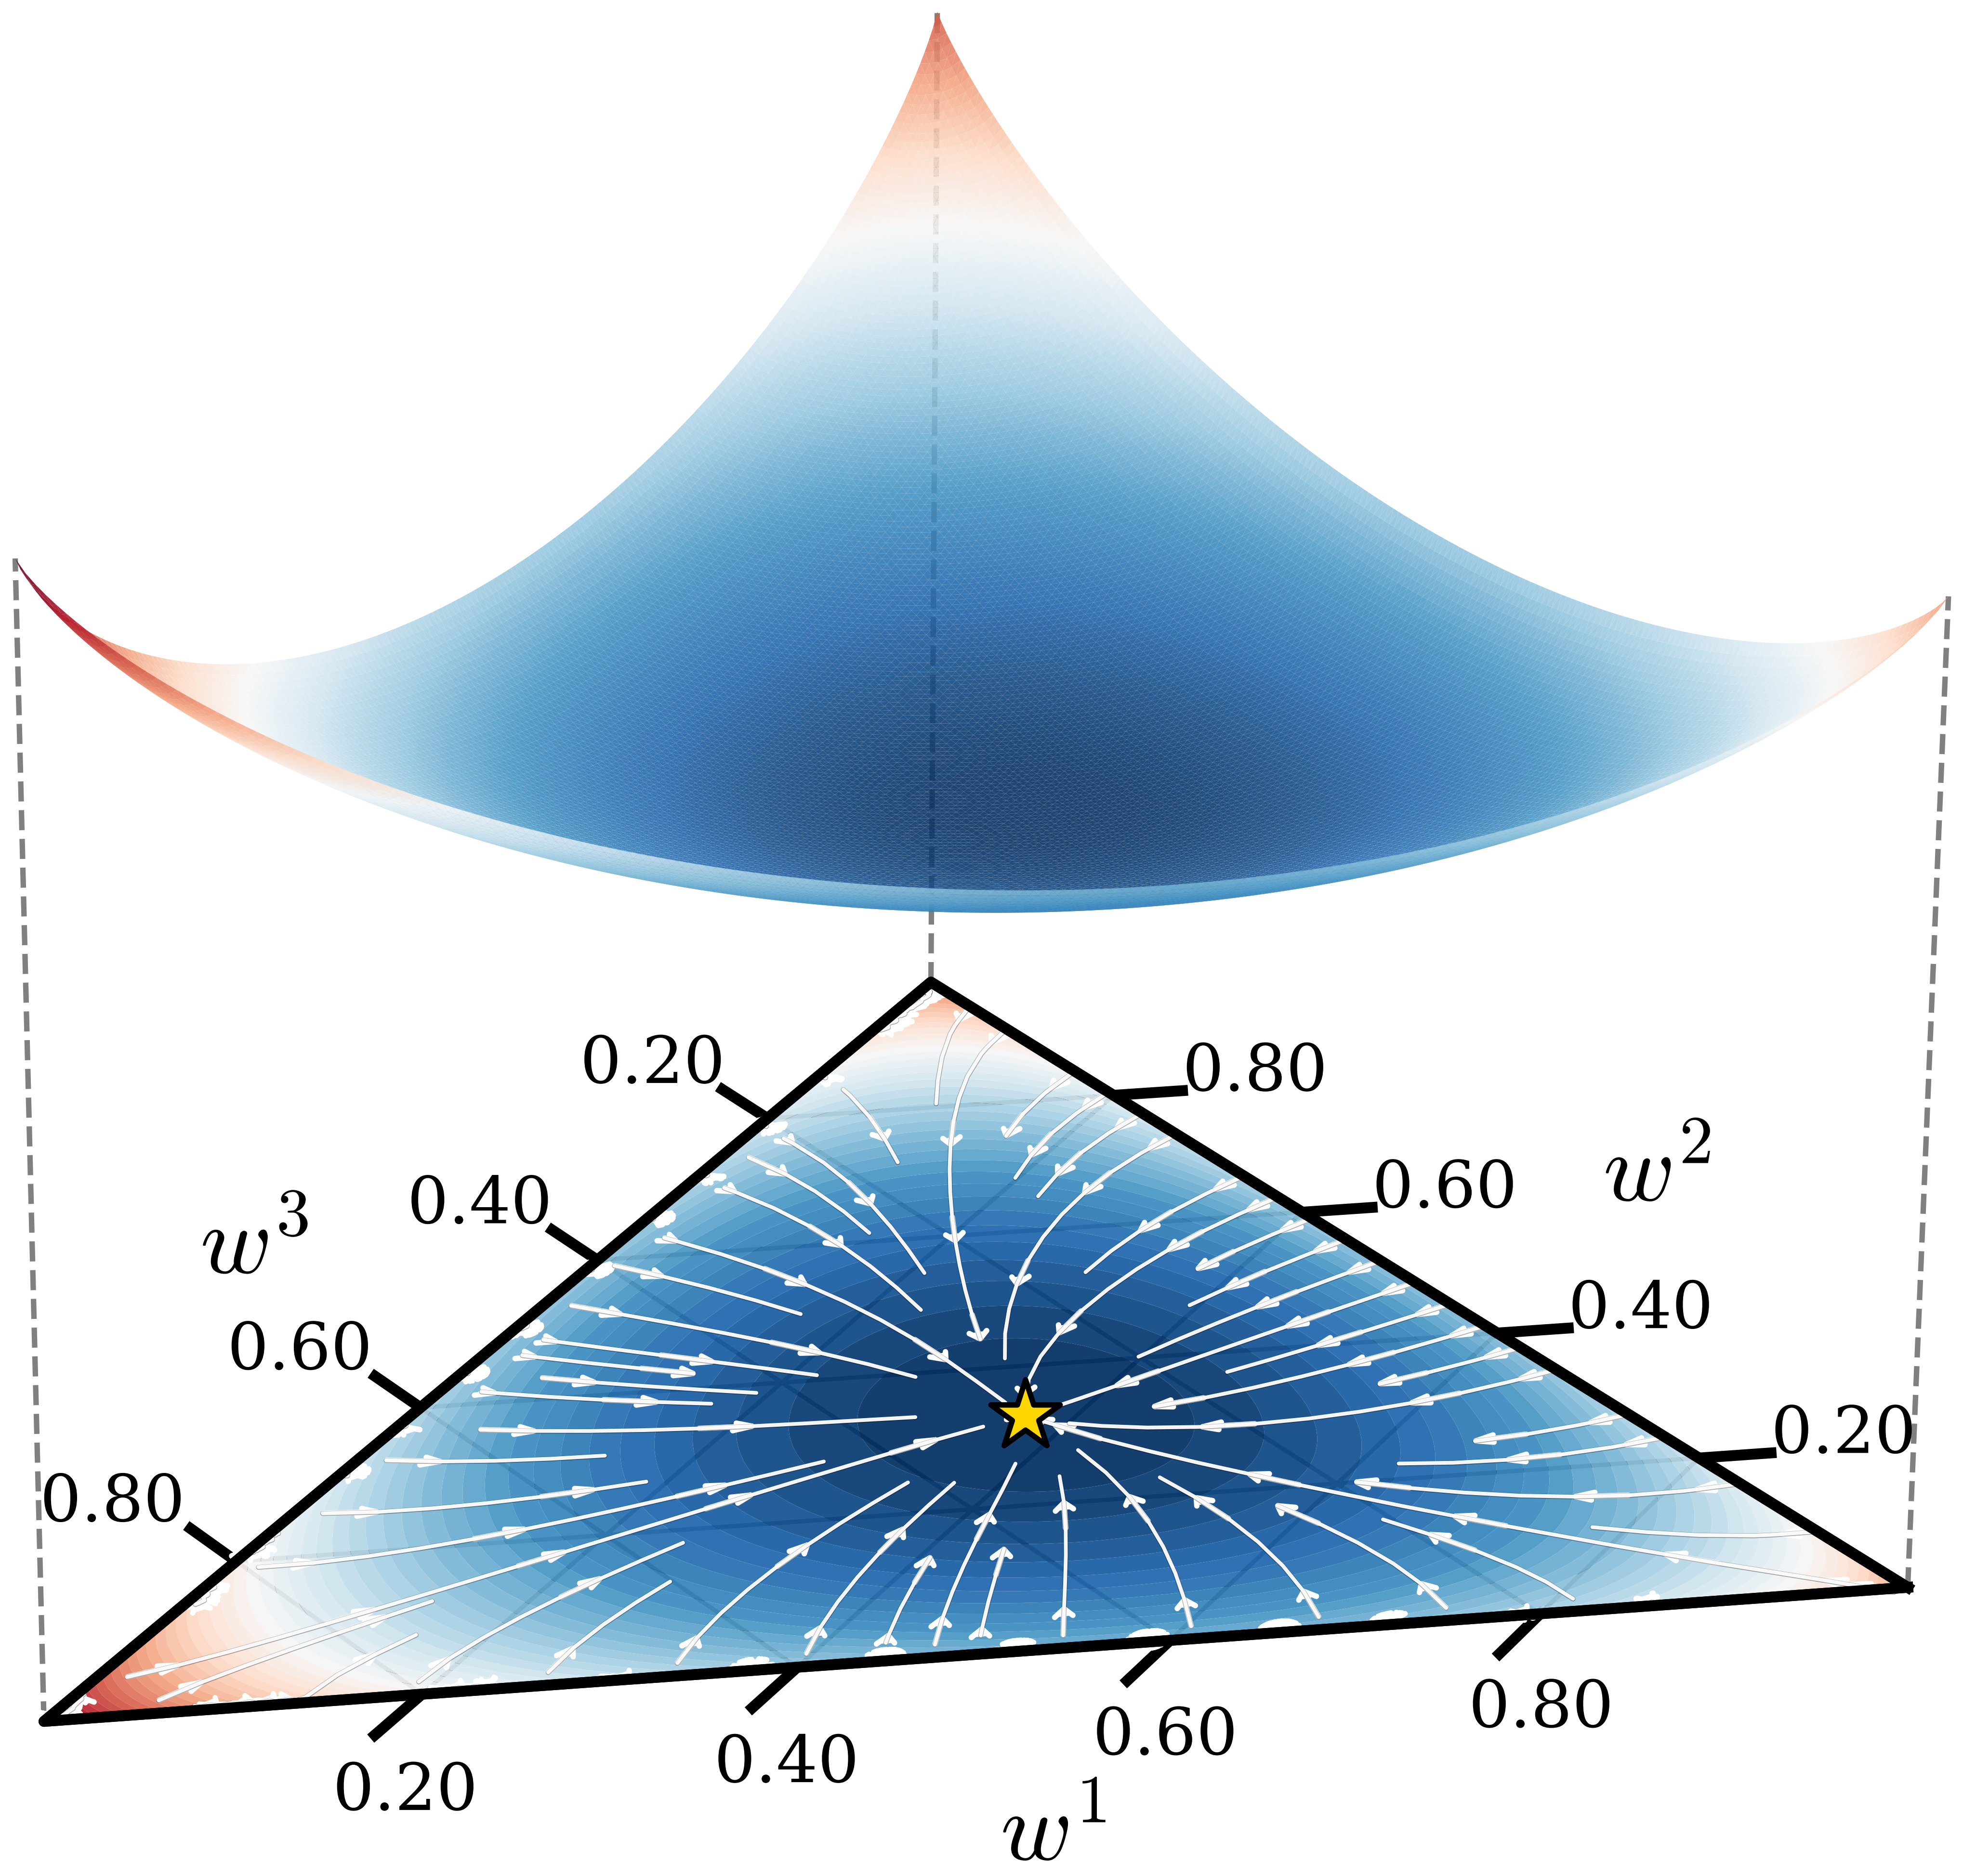

$$\tau \dot{\mathbf{w}}_k = -\,\mathbf{w}_k + \mathrm{softmax}\left(-\varepsilon^{-1}\nabla \mathsf{Error}(\mathbf{w}_k)\right)$$

- Equilibrium is the **optimal solution**
- Simplex is **forward invariant**
- Naturally interpreted as **firing-rate**

- **Contracting energy model!**

F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025

Example: Group Coordination

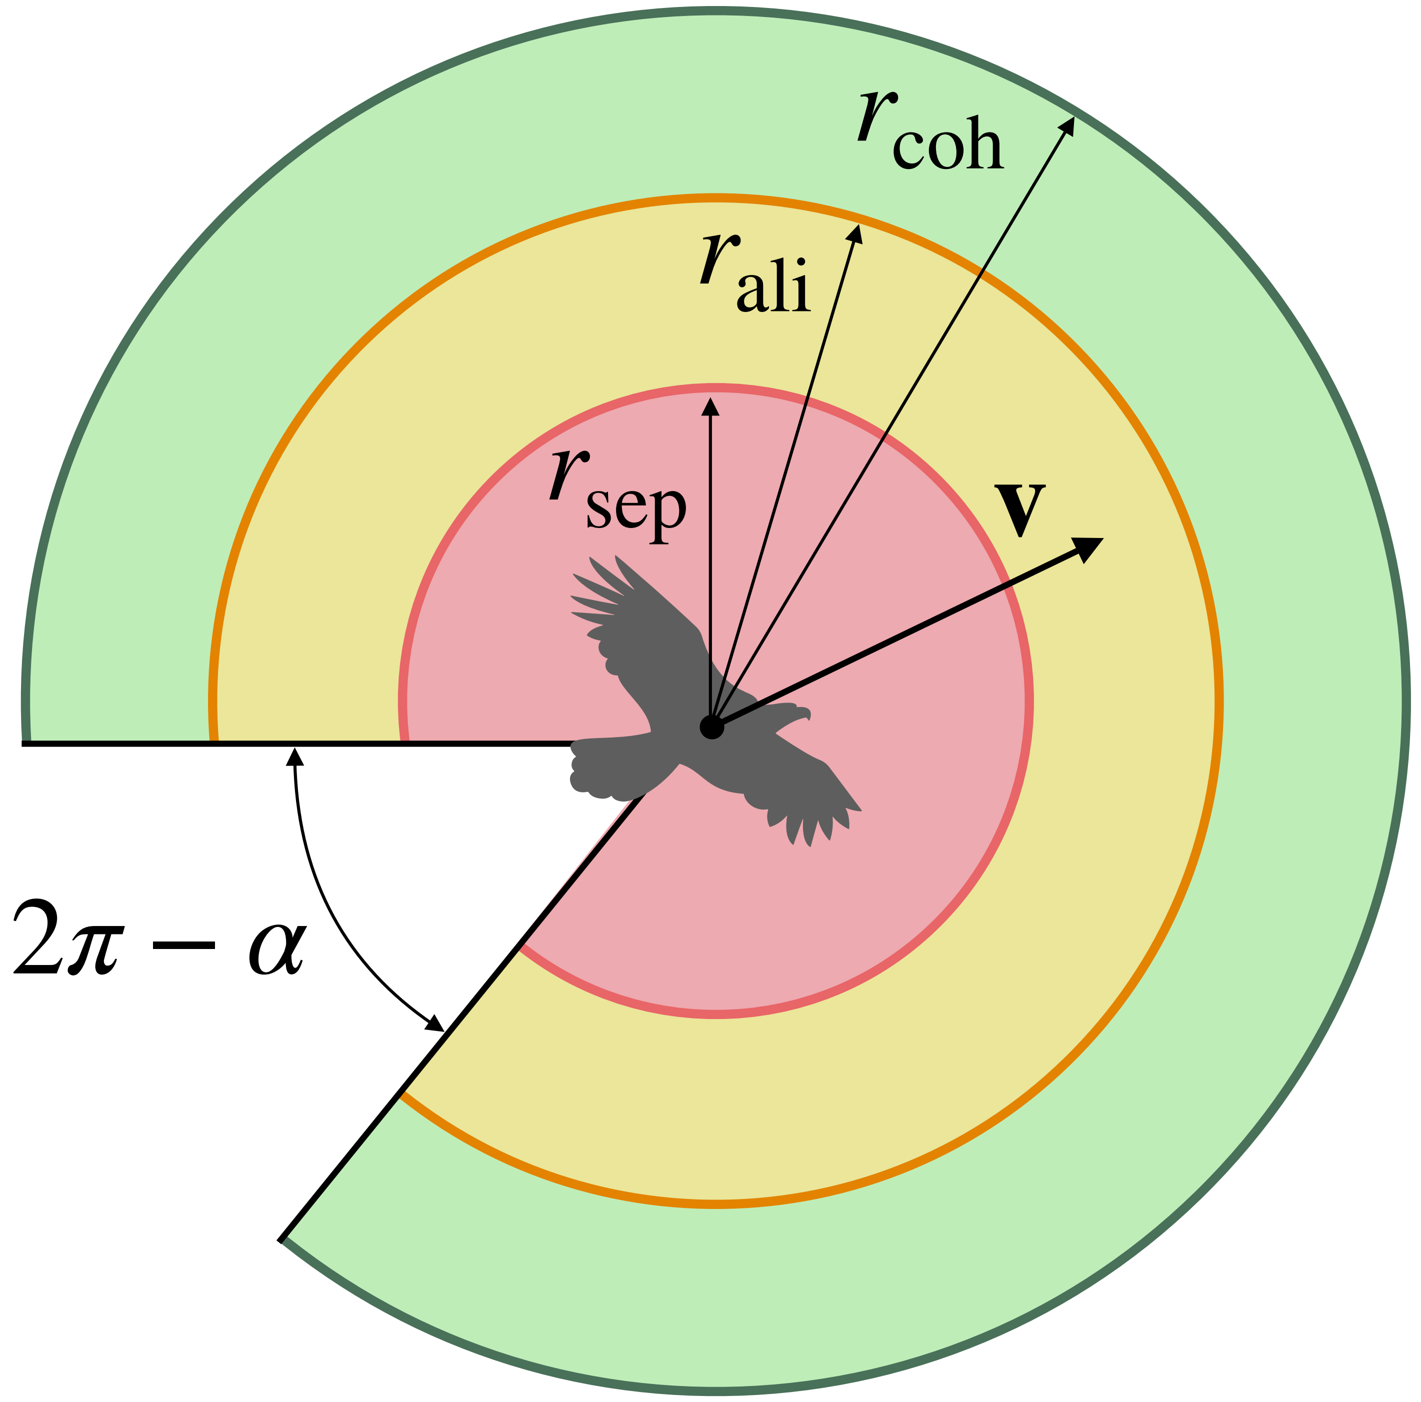

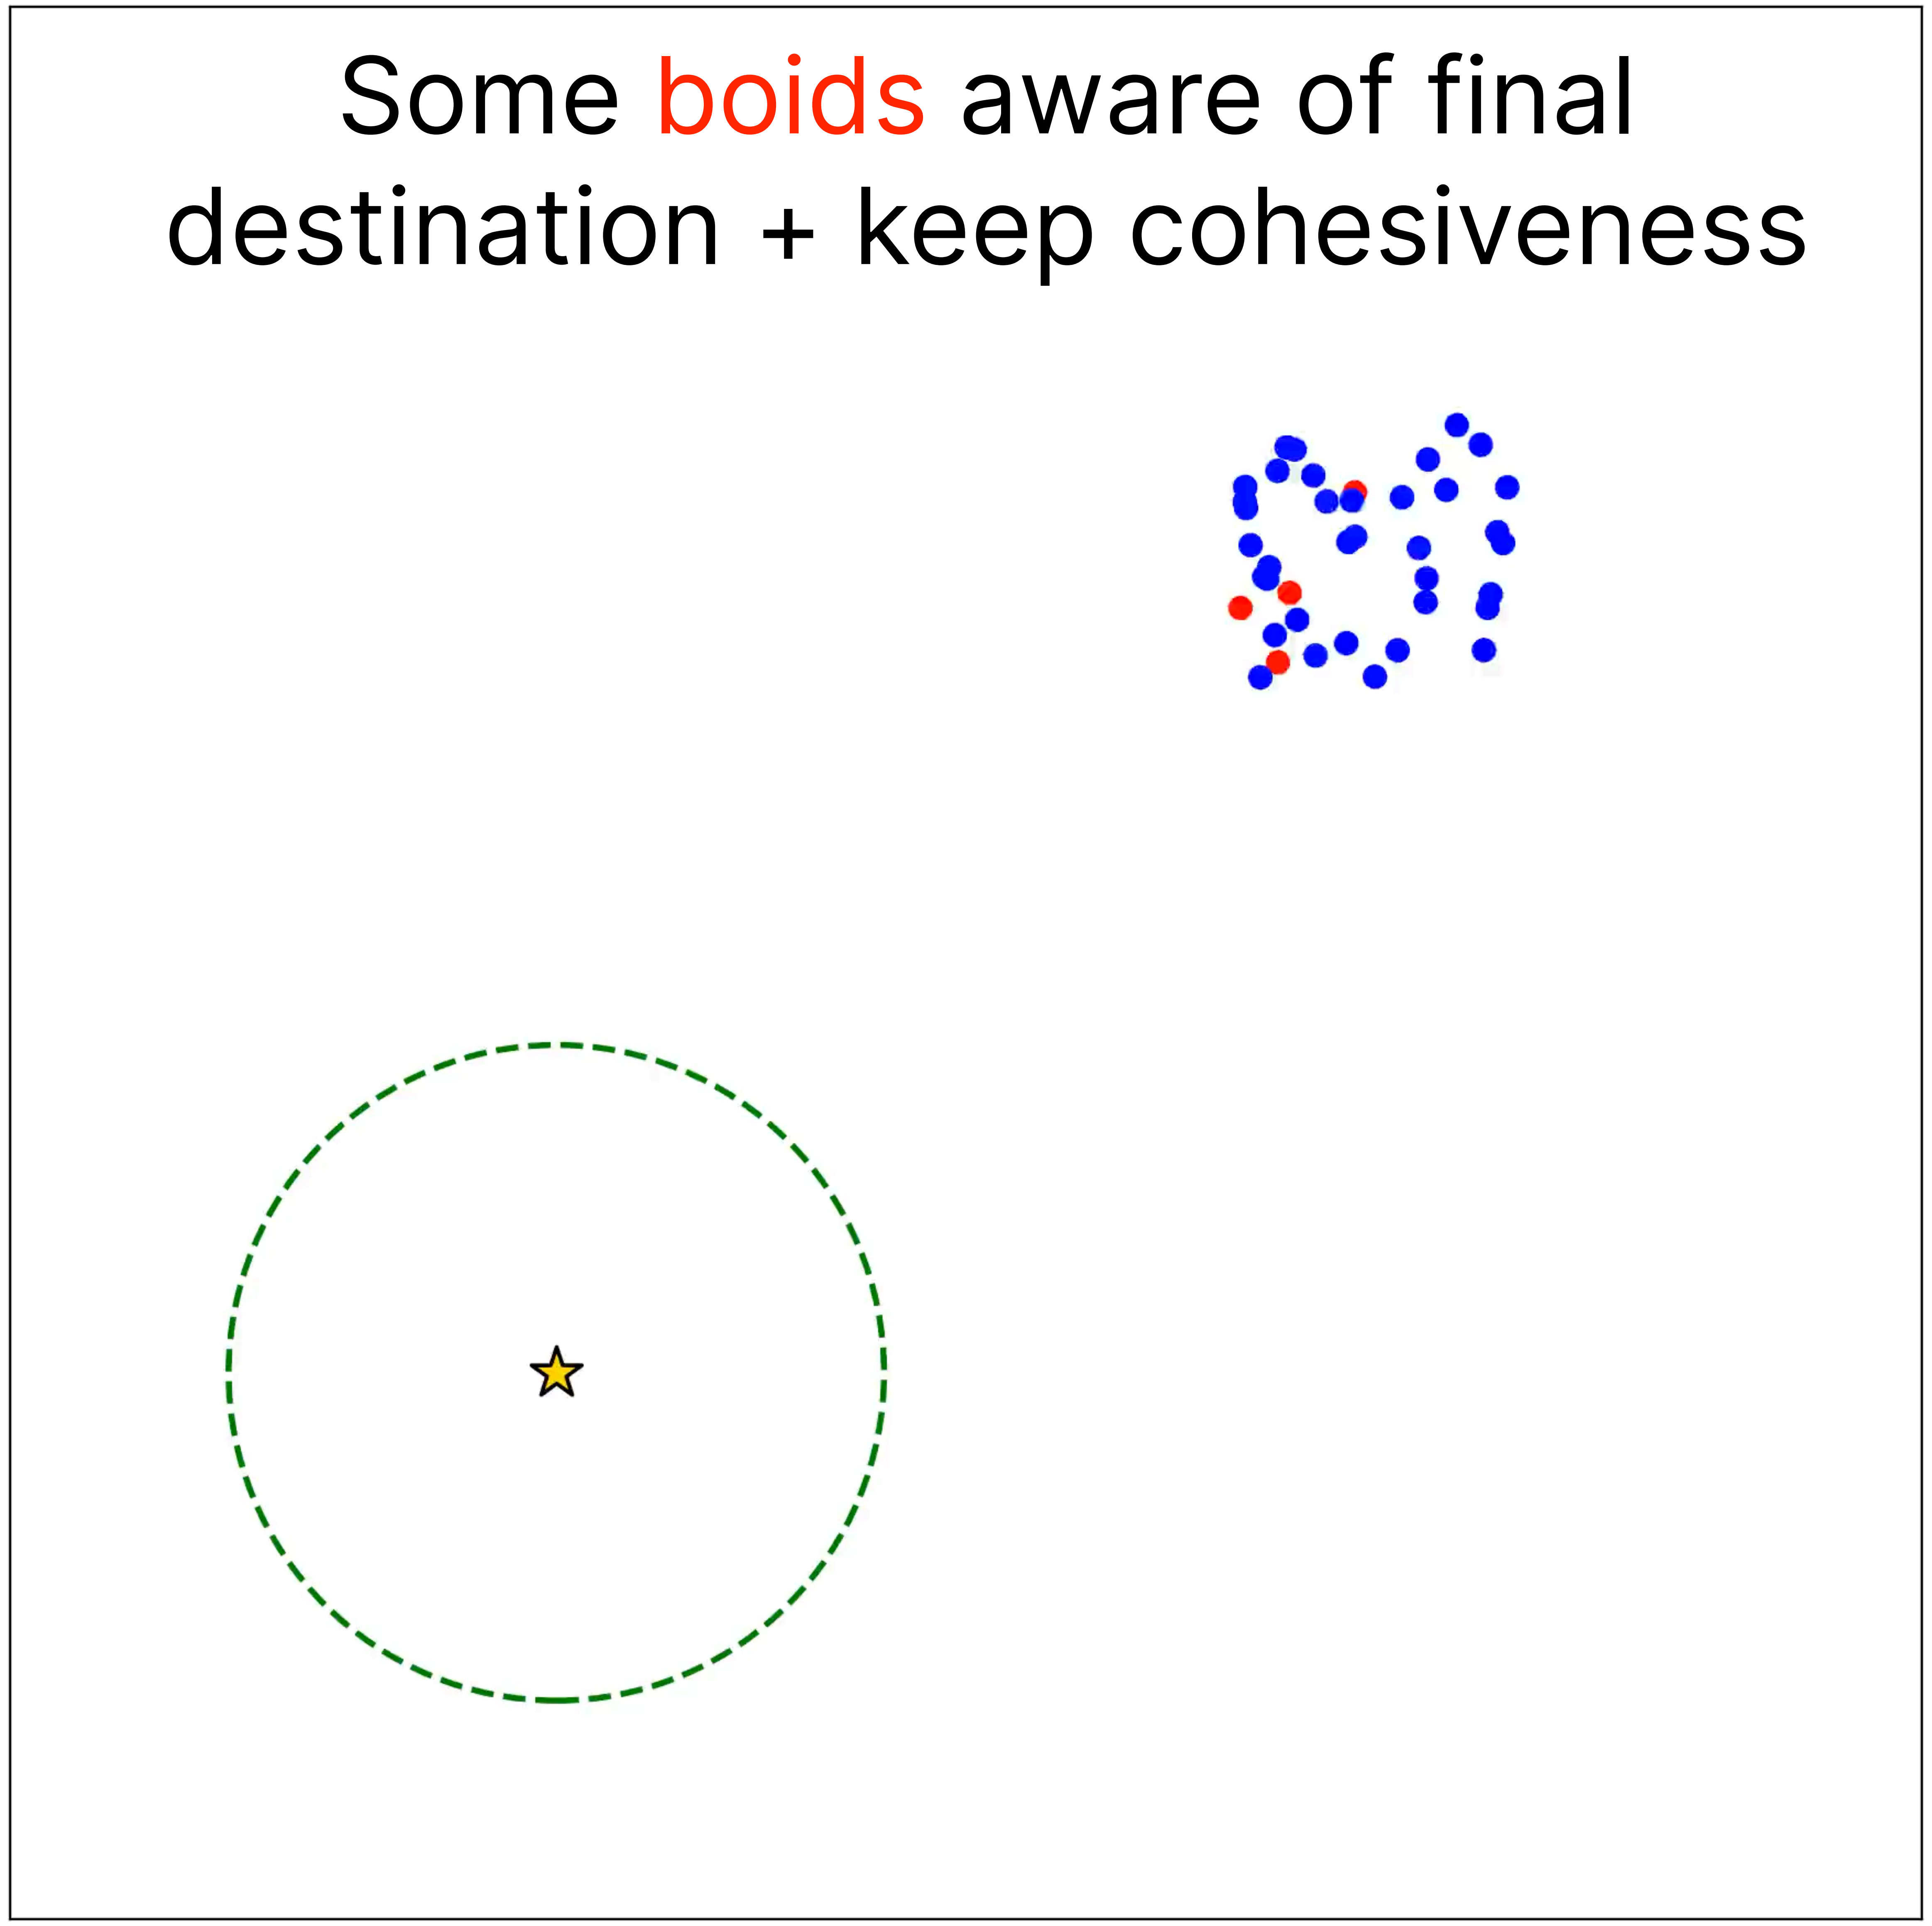

See Movie

C Heins, et al., Collective behavior from surprise minimization. PNAS, vol. 121, e2320239121, 2025

$\boldsymbol{x}_k$: position, velocity

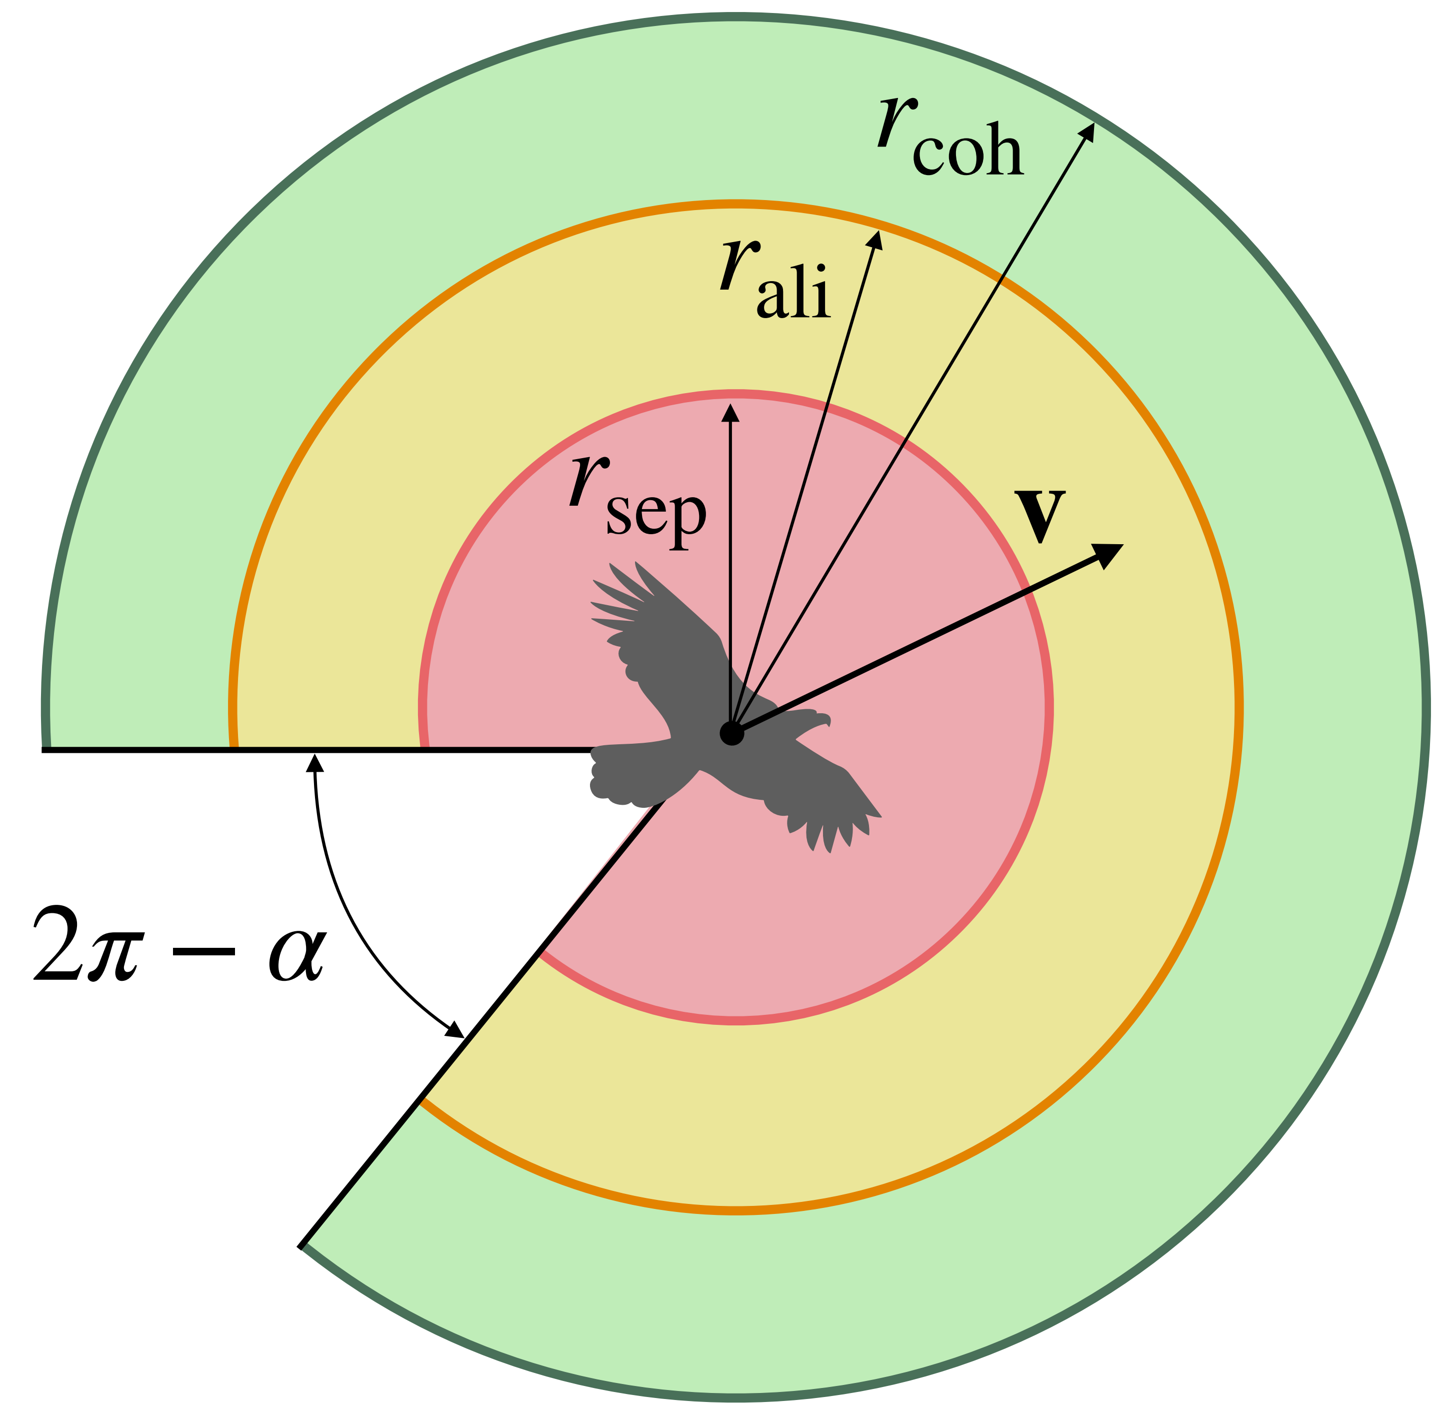

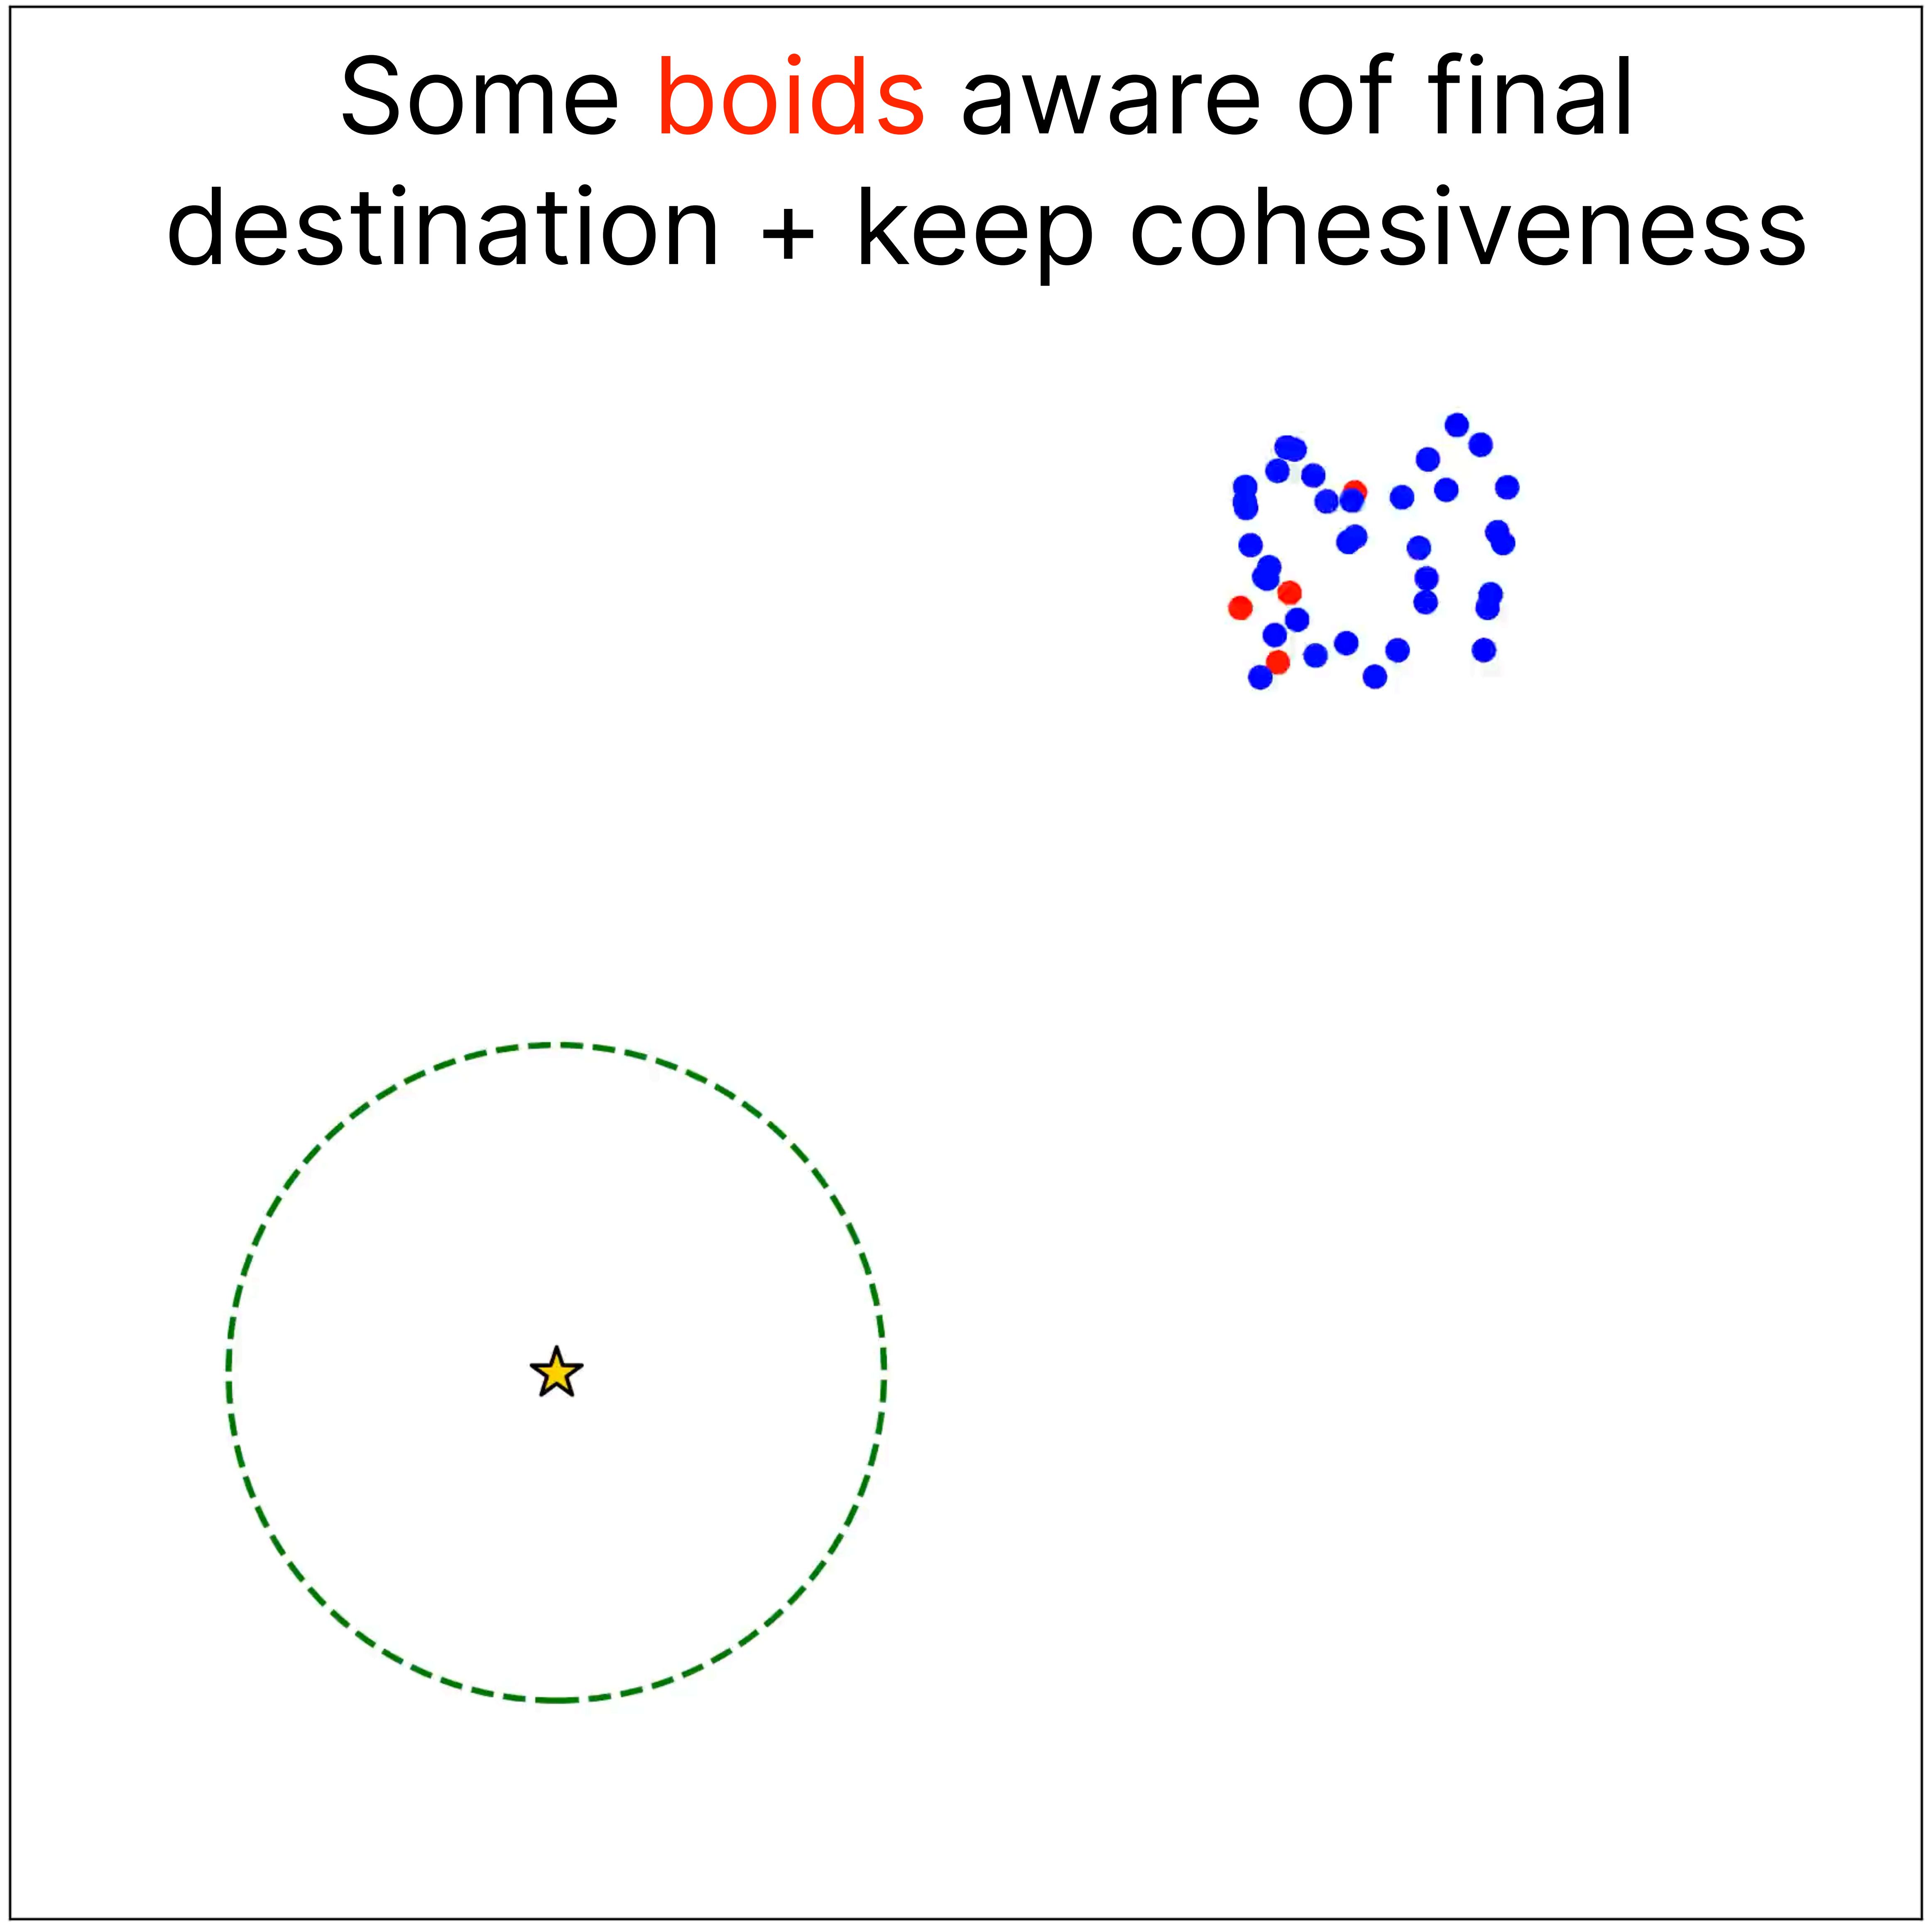

See Movie

C Heins, et al., Collective behavior from surprise minimization. PNAS, vol. 121, e2320239121, 2025

Example: Group Coordination

$\boldsymbol{x}_k$: position, velocity

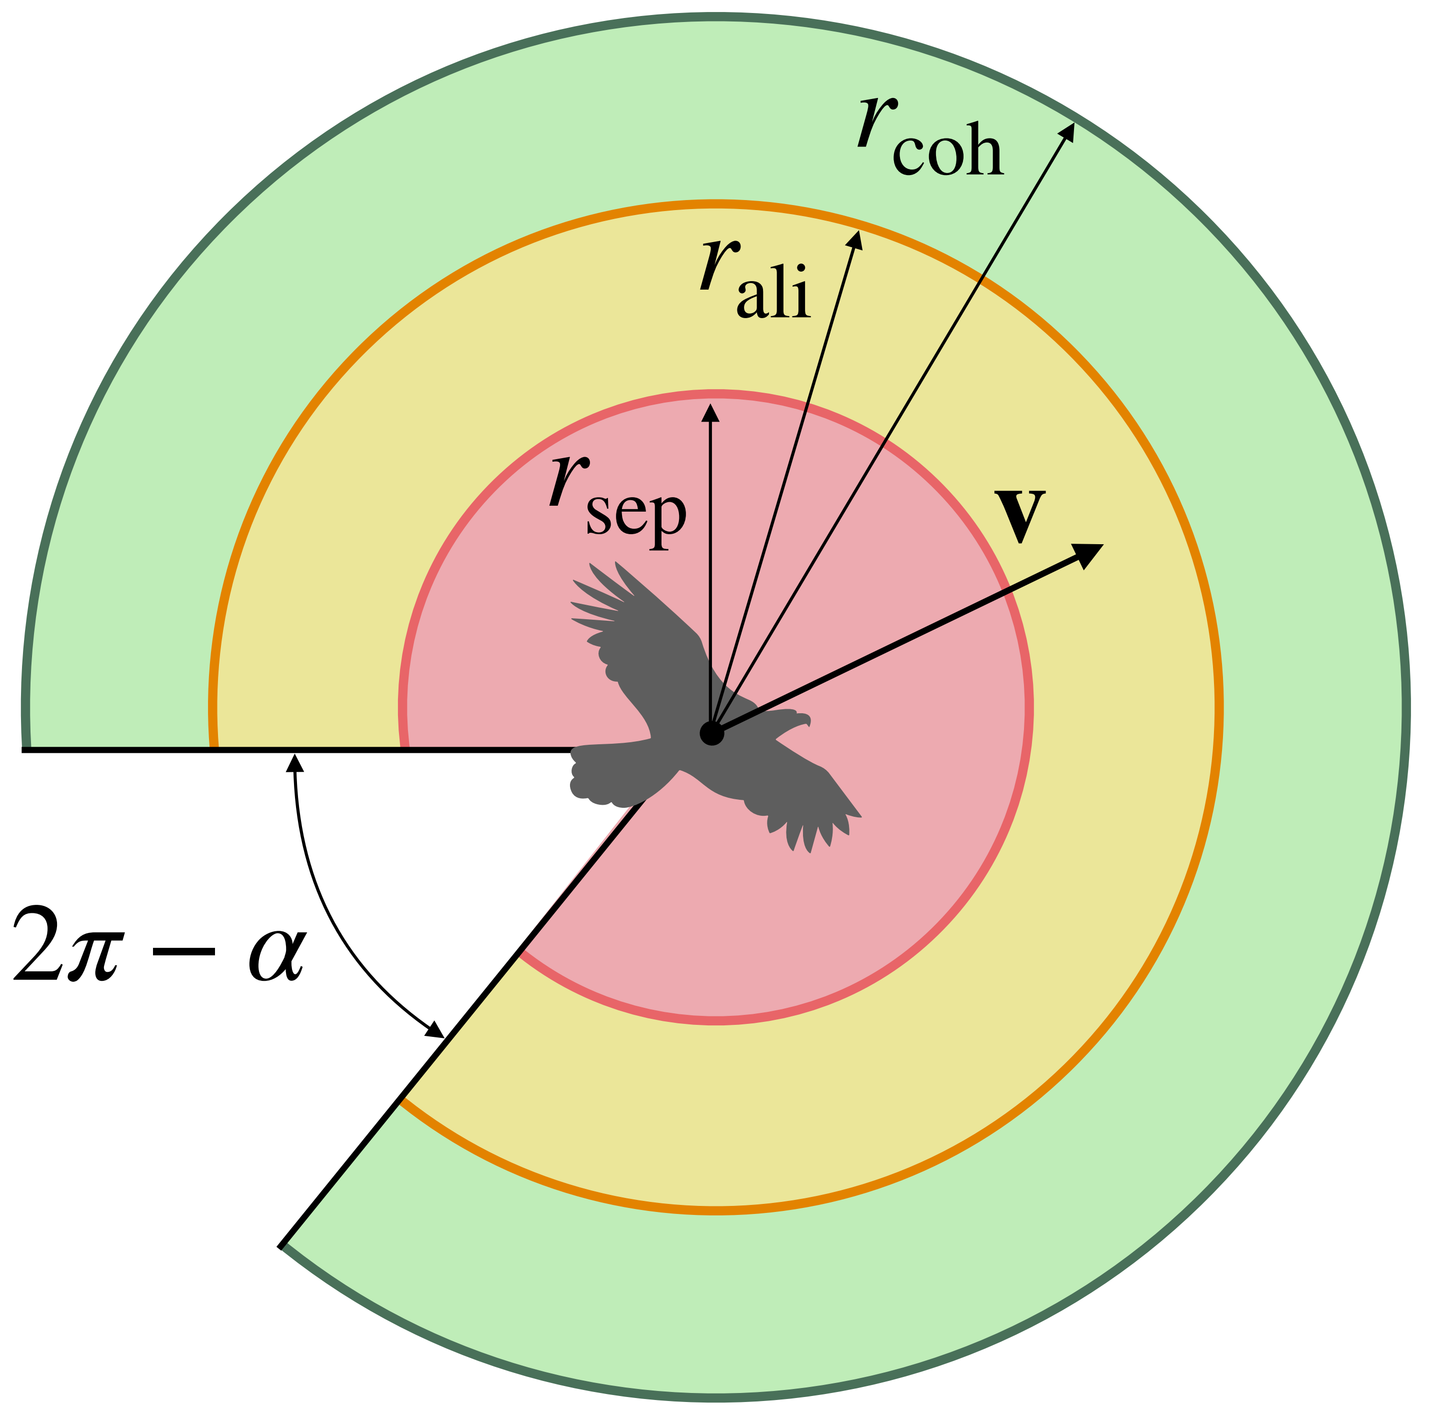

$\boldsymbol{u}_k$: acceleration

Some boids aware of final destination + keep cohesiveness

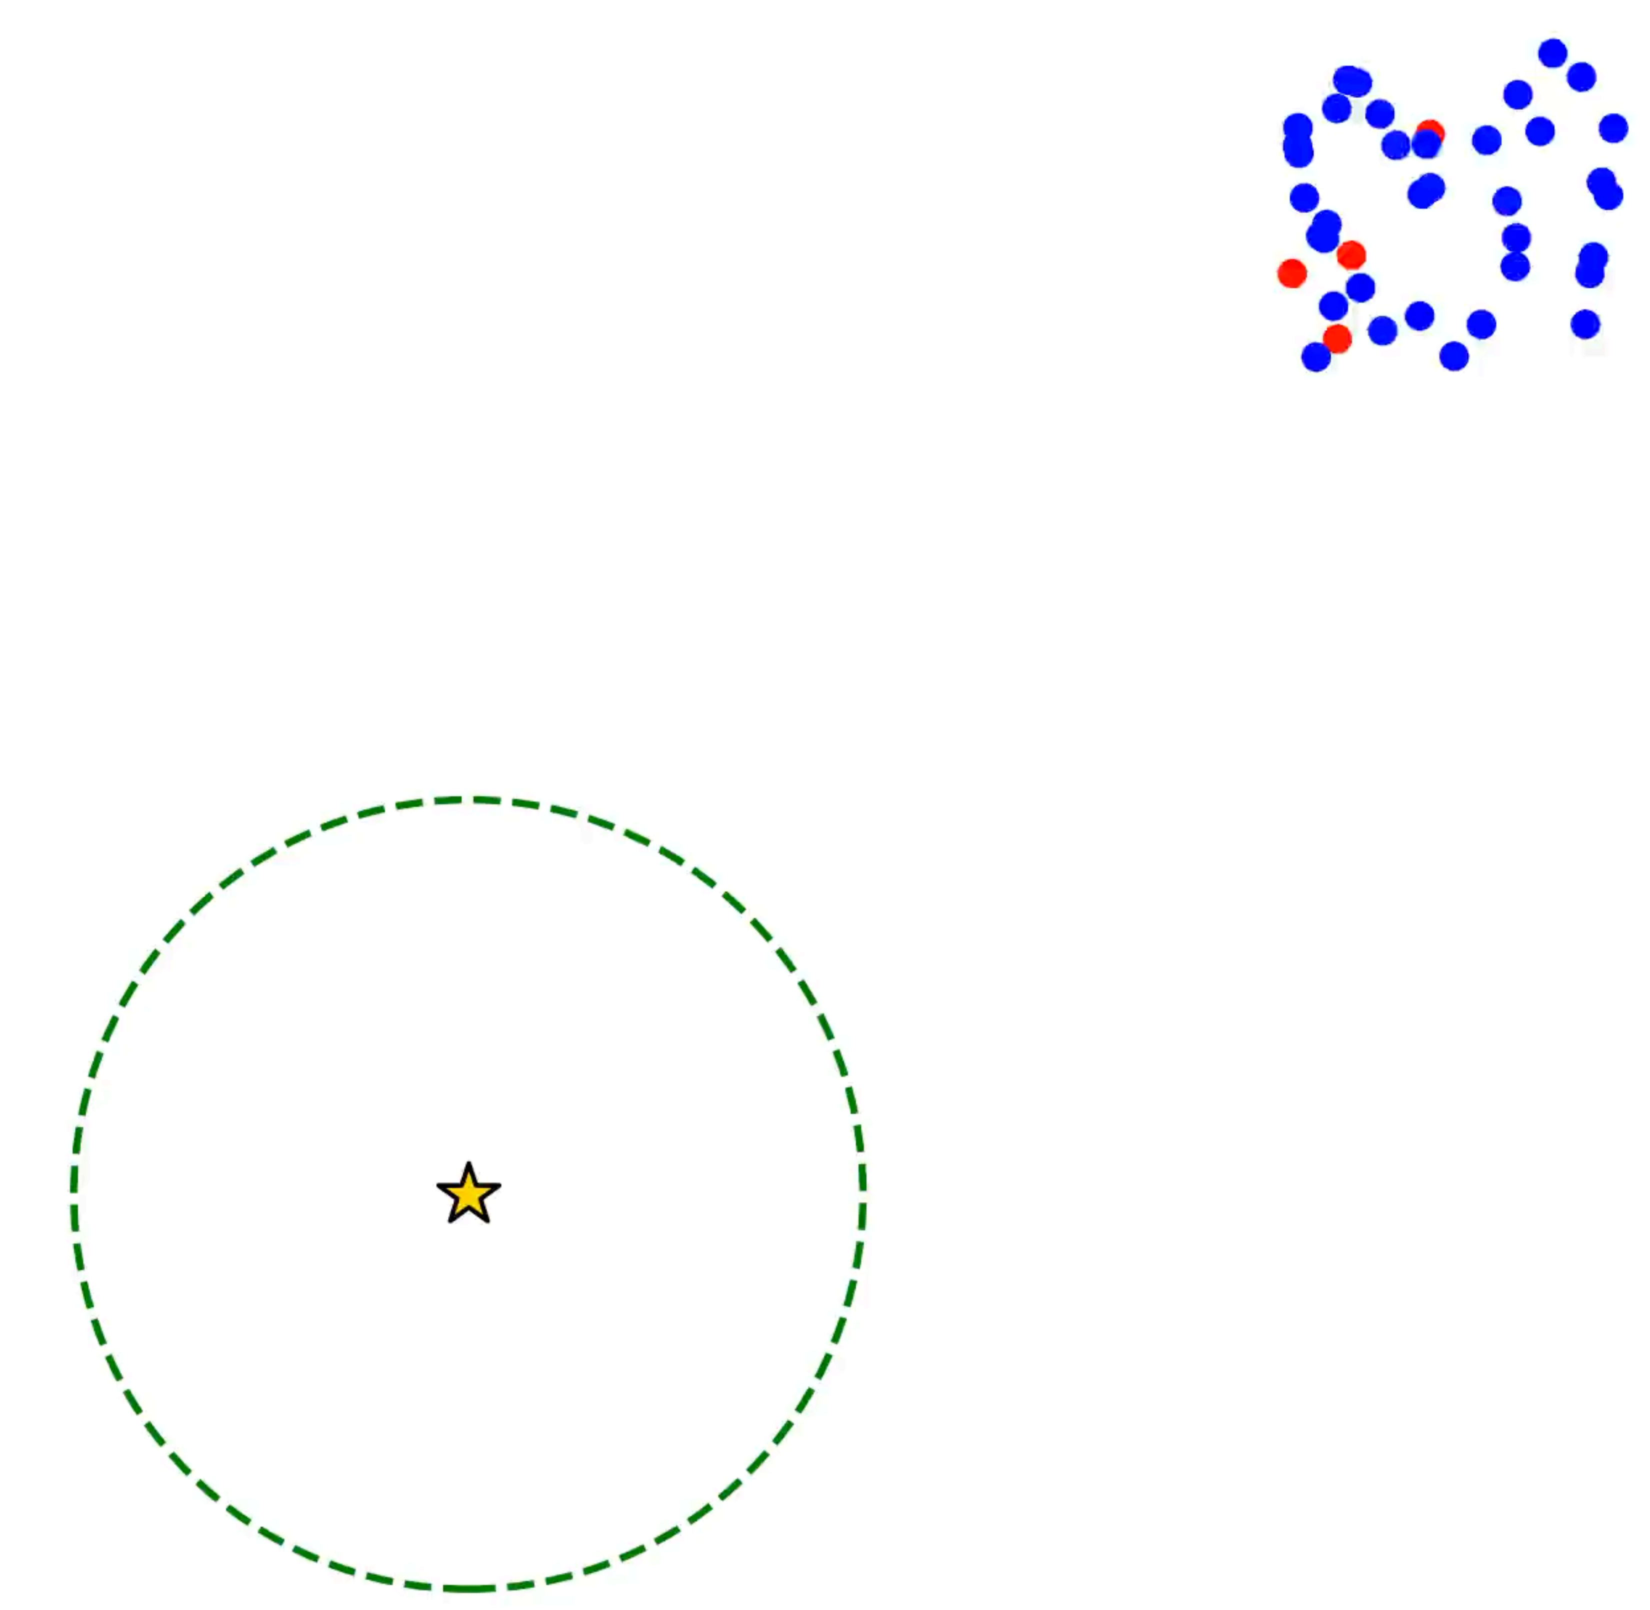

See Movie

C Heins, et al., Collective behavior from surprise minimization. PNAS, vol. 121, e2320239121, 2025

Example: Group Coordination

$\boldsymbol{x}_k$: position, velocity

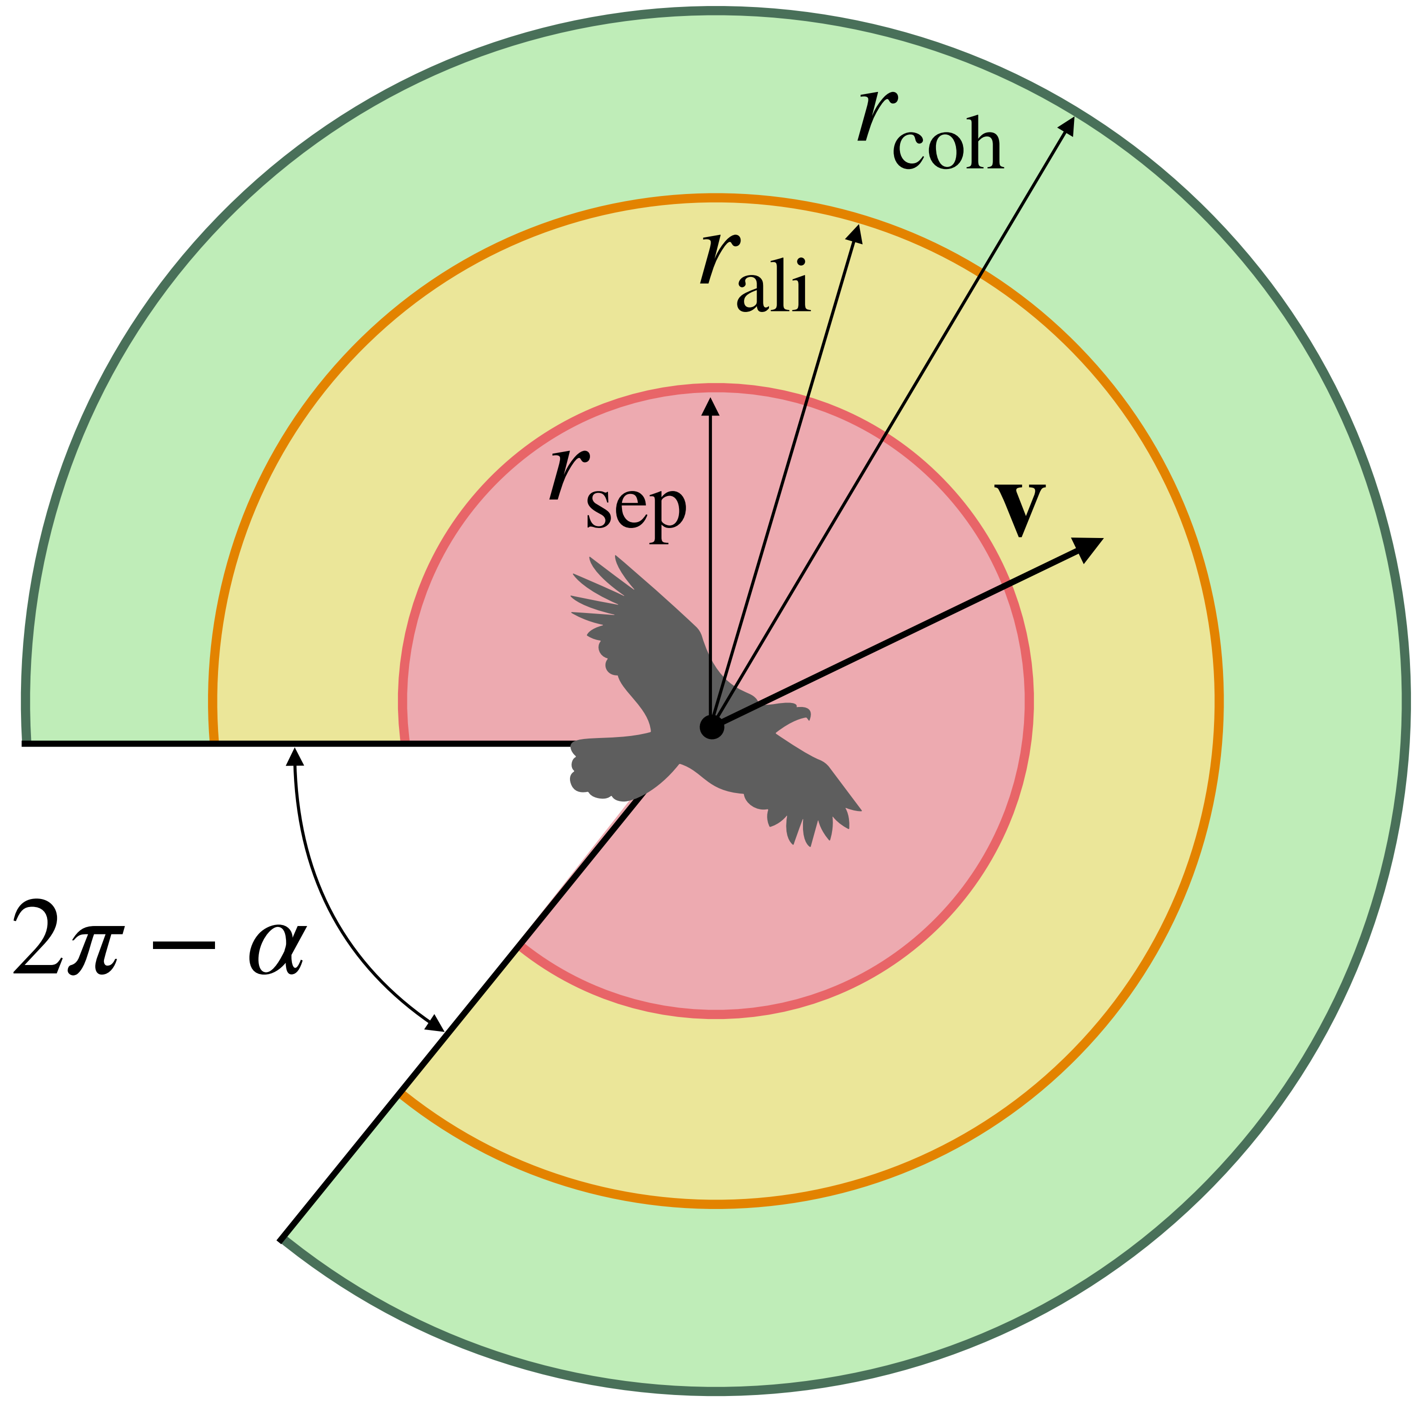

$\boldsymbol{u}_k$: acceleration

Some boids aware of final destination + keep cohesiveness

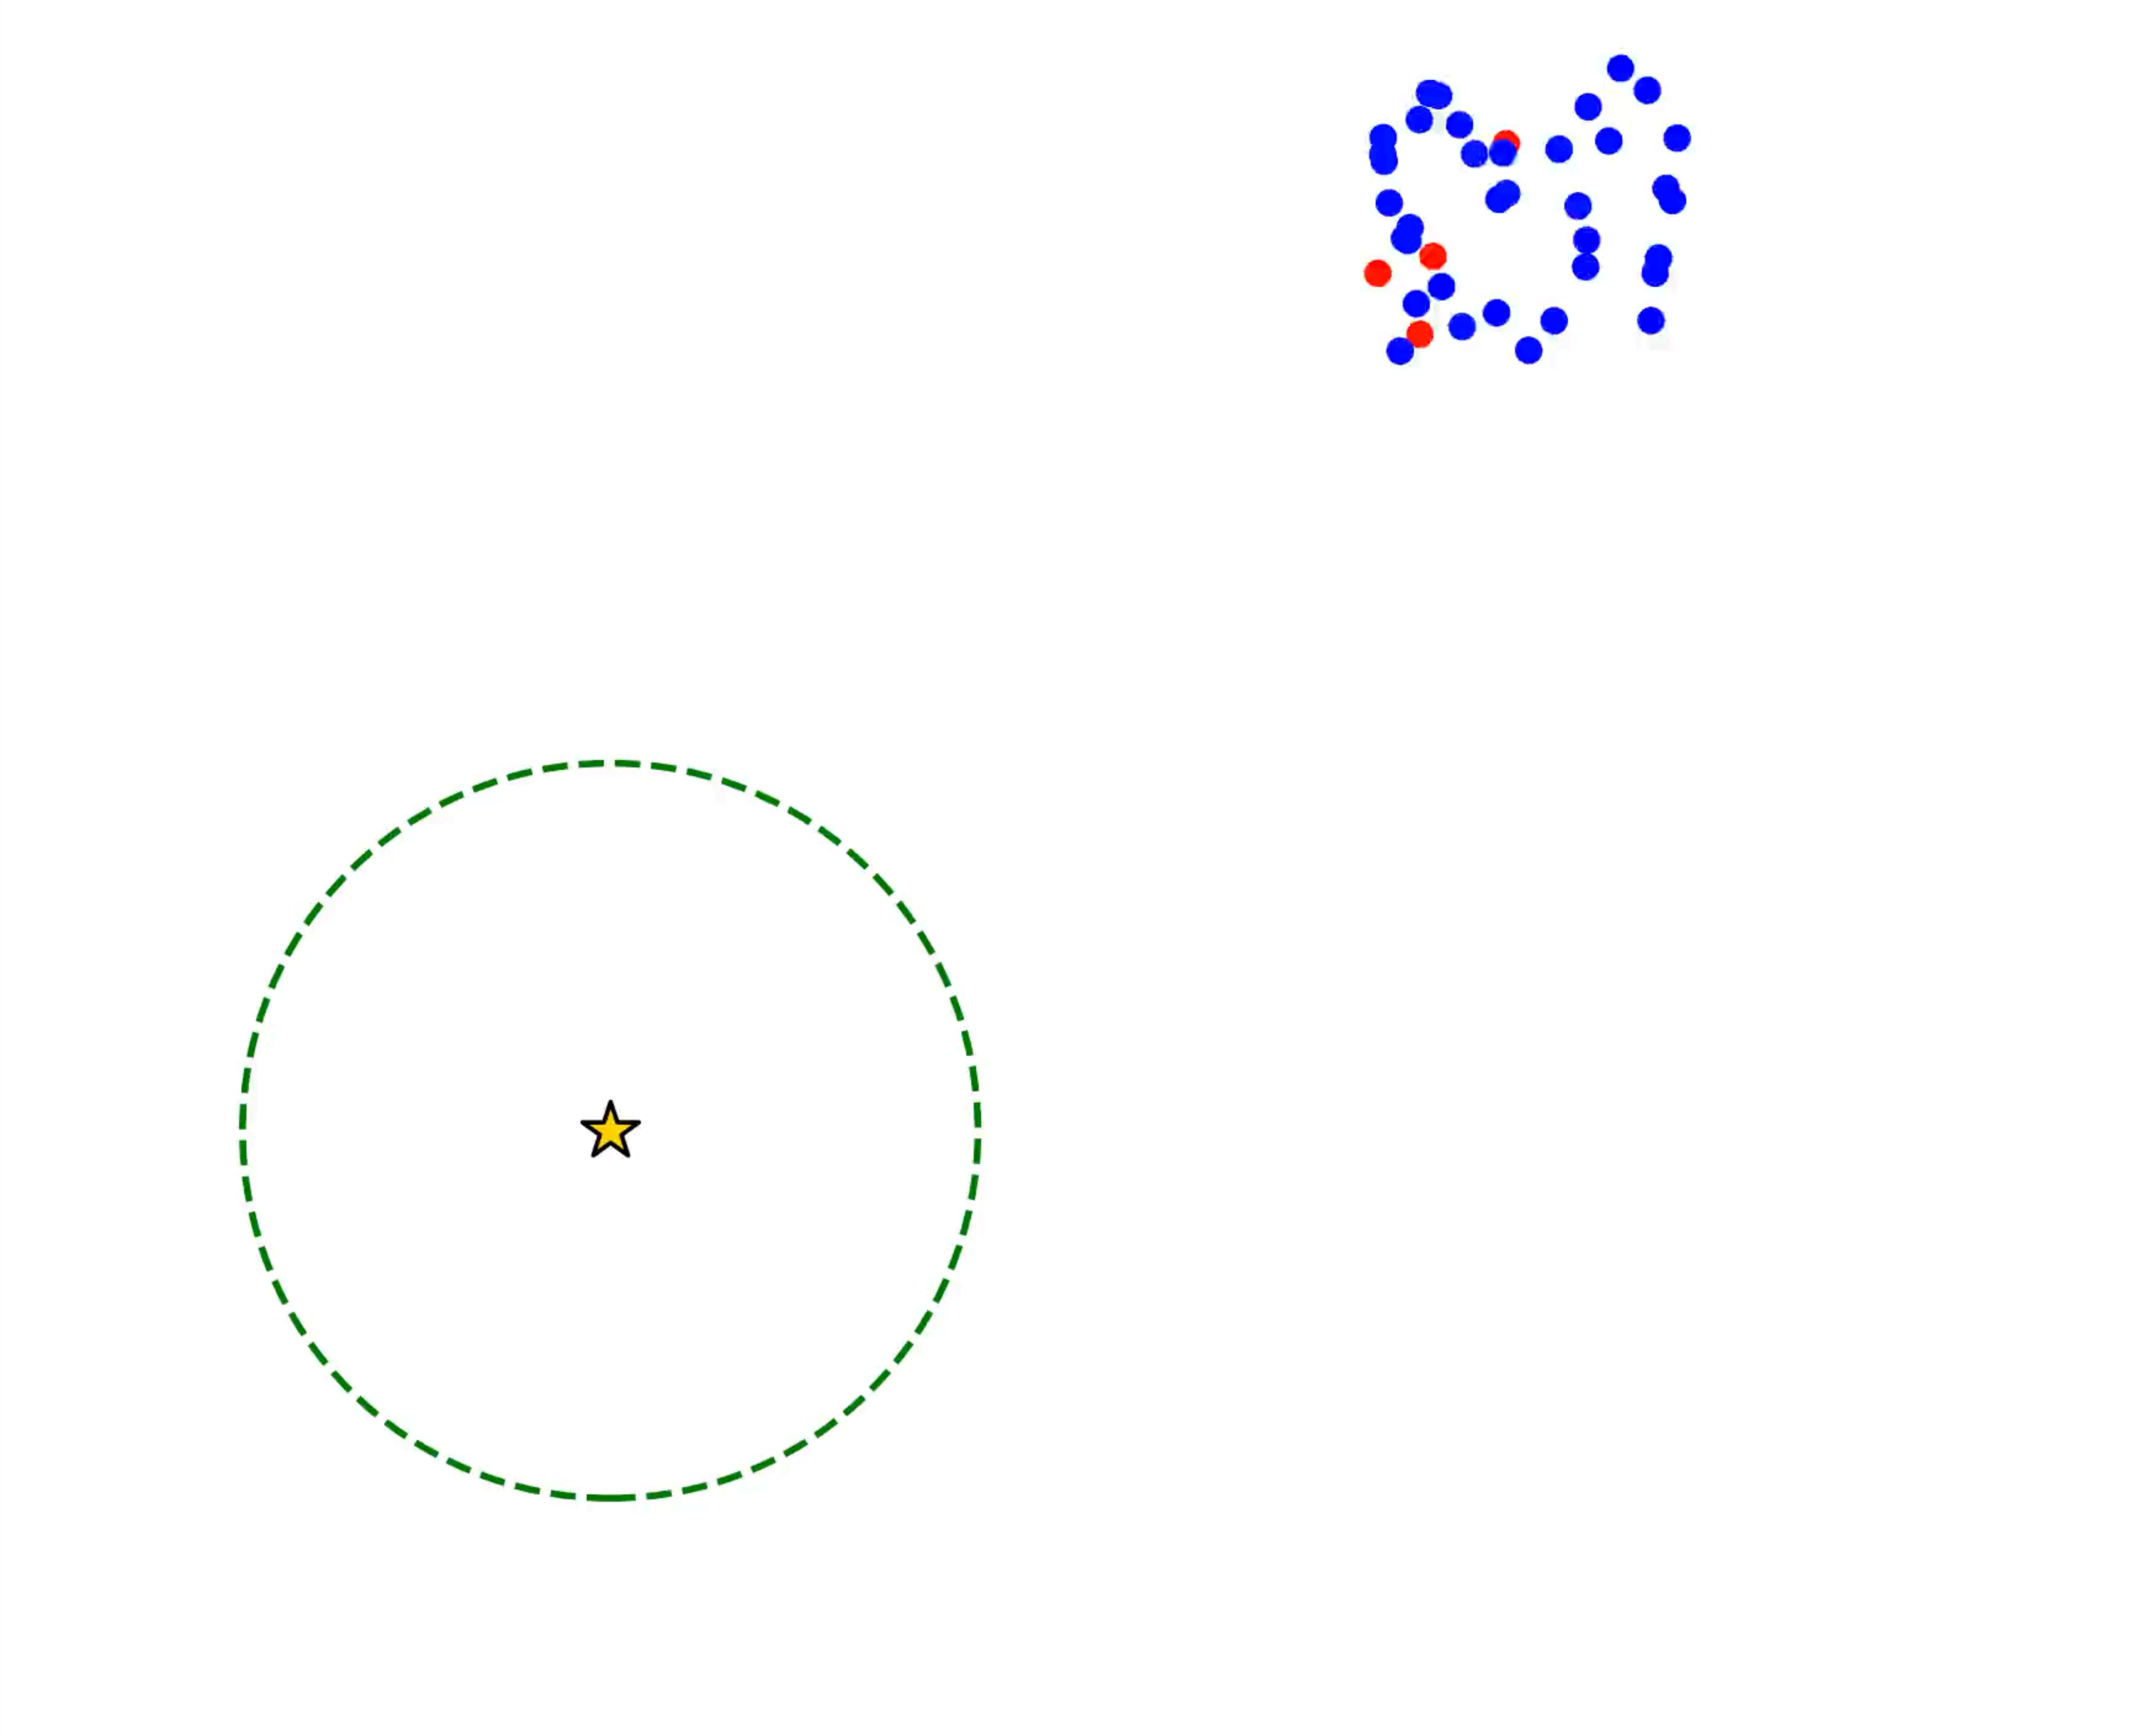

See Movie

C Heins, et al., Collective behavior from surprise minimization. PNAS, vol. 121, e2320239121, 2025

Thank you for your attention!

giovarusso@unisa.it