

Closing the loop between nonlinear distributed control and learning-based optimization

Dr. Luca Furieri



- SNSF Ambizione Fellow at EPF Lausanne (Jan. 2023 to present)
 - Principal investigator of «*Reliable Machine Learning for Distributed Control*»
- Postdoc at EPF Lausanne (Nov. 2020 to Dec. 2022)
 - Working with Prof. Giancarlo Ferrari Trecate
- PhD at ETH Zurich (Nov. 2016 to Sep 2020)
 - Supervised by Prof. Maryam Kamgarpour

Optimal control for large-scale dynamical systems

Target challenges




Optimal control for large-scale dynamical systems



Target challenges

1) Optimality in distributed tasks




■ *Myopic*, local improvement \rightarrow global effect

Optimal control for large-scale dynamical systems



Target challenges

1) Optimality in distributed tasks



- *Myopic*, local improvement \rightarrow global effect

2) Uncertain system models


```
11100001011
10011000011
10111001116
10000100000
10000011111
10001101001
11011100117
00110110001
00101011001
10111111001
```

- *Noisy data* \rightarrow optimal control policy
- Tension: performance VS guarantees

Optimal control for large-scale dynamical systems

Target challenges

1) Optimality in distributed tasks



- *Myopic*, local improvement \rightarrow global effect

2) Uncertain system models

```
11100001011
10011000011
10111001116
10000100000
10000011111
10001101001
11011100117
00110110001
00101011001
10111111000
```

- *Noisy* data \rightarrow optimal control policy
- Tension: performance VS guarantees

3) Guarantees of learning-based control



- Neural network policies, nonlinear objectives
- Closed-loop stability? Safety?



Selected Contributions: Part 1

Optimal Distributed Control (ODC) for Linear Systems

- [1] “*Sparsity invariance for convex design of distributed controllers*”, [Furieri](#), Zheng, Papachristodoulou, Kamgarpour, TCNS, 2020
- [2] “*Learning the globally optimal distributed LQ regulator*”, [Furieri](#), Zheng, Kamgarpour, L4DC, 2020

Beyond long-standing limitations of linear ODC^[1]

[1] "Sparsity Invariance for Convex Design of Distributed Controllers", [Furieri, Zheng, Papachristodoulou, Kamgarpour, \[TCNS20\]](#)*
[*IEEE Transactions on Control of Network Systems Best Paper Award, 2022](#)



$\min_{\mathbf{K} \text{ stabilizing}}$


subject to

$$\text{cost}(\mathbf{G}, \mathbf{K}) = \|(w, v) \xrightarrow{\mathbf{G}, \mathbf{K}} (u, y)\|_{\mathcal{H}_2/\mathcal{H}_\infty}$$

$$\mathbf{K} = \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$$

Beyond long-standing limitations of linear ODC^[1]

[1] “Sparsity Invariance for Convex Design of Distributed Controllers”, Furieri, Zheng, Papachristodoulou, Kamgarpour, [TCNS20]*
 *IEEE Transactions on Control of Network Systems Best Paper Award, 2022




Limitation: convex reformulation available **only if Quadratic Invariance (QI) holds** [Rotkowitz et al., 2006]


 “*design sparse closed-loop maps*”


 “*sparse controller*” \iff “*sparse closed-loop maps*”

Beyond long-standing limitations of linear ODC[1]

[1] "Sparsity Invariance for Convex Design of Distributed Controllers", Furieri, Zheng, Papachristodoulou, Kamgarpour, [TCNS20]*
 *IEEE Transactions on Control of Network Systems Best Paper Award, 2022



Limitation: convex reformulation available only if *Quadratic Invariance (QI)* holds [Rotkowitz et al., 2006]

↓
 “design sparse closed-loop maps”

↓
 “sparse controller” \iff “sparse closed-loop maps”


My Contribution: near-optimal convex restrictions

1. QI as a special case of **Sparsity Invariance (SI)** \leftarrow Plant-independent
2. **Globally optimal and sparse** K for some non-QI cases
3. Near-optimality guarantees for **arbitrary sparsity patterns**

\implies **QI no-longer a limitation for sparse controller design**

Beyond long-standing limitations of linear ODC[1]

[1] "Sparsity Invariance for Convex Design of Distributed Controllers", Furieri, Zheng, Papachristodoulou, Kamgarpour, [TCNS20]*
 *IEEE Transactions on Control of Network Systems Best Paper Award, 2022



Limitation: convex reformulation available only if *Quadratic Invariance (QI)* holds [Rotkowitz et al., 2006]



Next question: unknown system model?

My Contribution: near-optimal convex restrictions

1. QI as a special case of **Sparsity Invariance (SI)** \leftarrow Plant-independent
2. **Globally optimal and sparse** K for some non-QI cases
3. Near-optimality guarantees **for arbitrary sparsity patterns**

\Rightarrow **QI no-longer a limitation for sparse controller design**


Learning-based control philosophies

Model-based

Model-free
(RL-like)

Learning-based control philosophies

Model-based



Noisy trajectories

System Identification



Optimal Control

$$\begin{array}{ll} \min_{\mathbf{K} \text{ stabilizing}} & \text{cost}(\hat{\mathbf{G}}, \mathbf{K}) \\ \text{subject to} & \mathbf{K} = \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix} \end{array}$$

Model-free
(RL-like)

Learning-based control philosophies

Model-based

System Identification

Optimal Control

???

Noisy

Aspects I address


- Suboptimality (as compared to ground-truth)
- Sample-complexity (how many data?)
- Safety (for the real system)

$\min \text{cost}(\hat{G}, K)$

$$= \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$$

Apply policy

Model-free
(RL-like)



Sample the cost

$$\text{cost}(\pi_i) = \sum_{t=0}^T y_t^T Q y_t + u_t^T R u_t$$

Policy update

$$\pi_{i+1} = \pi_i - \eta \nabla \text{cost}(\pi_i)$$

Model-free learning of the globally optimal distributed LQ regulator^[1]

[1] [Furieri](#), Zheng, Kamgarpour, "Learning the globally optimal distributed LQ regulator", L4DC, 2020

Result: Favorable optimization landscape^[1]

The cost is **nonconvex** in distributed policies $\pi \in \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$. However;

Model-free learning of the globally optimal distributed LQ regulator^[1]

[1] [Furieri](#), Zheng, Kamgarpour, "Learning the globally optimal distributed LQ regulator", L4DC, 2020

Result: Favorable optimization landscape^[1]

The cost is **nonconvex** in distributed policies $\pi \in \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$. However;


Model-free learning of the globally optimal distributed LQ regulator^[1]

[1] Furieri, Zheng, Kamgarpour, "Learning the globally optimal distributed LQ regulator", L4DC, 2020

Result: Favorable optimization landscape^[1]

The cost is **nonconvex** in distributed policies $\pi \in \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$. However;

- Whenever QI holds, $\text{cost}(\pi)$ is **Polyak-Lojasiewicz (PL)** →
- Several non-QI cases are also PL!
- ... PL holds for arbitrary sparsities? **No.**




Model-free learning of the globally optimal distributed LQ regulator^[1]

[1] Furieri, Zheng, Kamgarpour, "Learning the globally optimal distributed LQ regulator", LADC, 2020

Result: Favorable optimization landscape^[1]

The cost is **nonconvex** in distributed policies $\pi \in \begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix}$. However;

- Whenever QI holds, $\text{cost}(\pi)$ is **Polyak-Lojasiewicz (PL)** →
- Several non-QI cases are also PL!
- ... PL holds for arbitrary sparsities? **No.**



Result: Sample-complexity for global optimality^[1]

Assume that PL holds. Run model-free policy gradient for T steps with

$$T \propto \left(\text{card} \left(\begin{bmatrix} \star & \star & 0 \\ 0 & \star & \star \end{bmatrix} \right) \right)^2 \epsilon^{-2} \delta^{-4}$$

Then:

$$\text{cost}(\pi_T) - \text{cost}(\pi^*) \leq \epsilon$$


with high probability $1 - \delta$.

Selected Contributions: Part 2

Learning to control with stability guarantees for nonlinear systems

- [1] [Furieri](#), Galimberti, F. Trecate, “*Neural system level synthesis: learning over all stabilizing policies for nonlinear systems*”, [CDC 2022]
- [2] [Furieri](#), Galimberti, F. Trecate, “*Learning to boost the performance of stable nonlinear systems*”, [ArXiv, 2024]

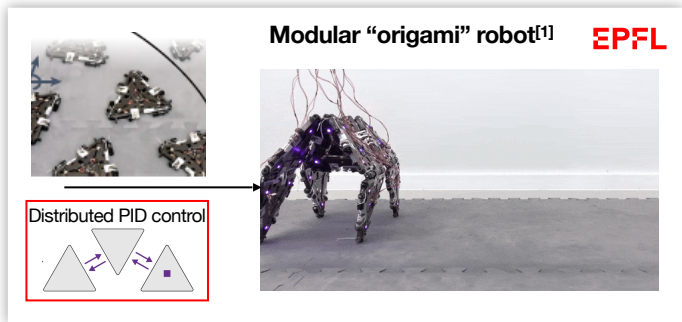
- Complex real-world systems are nonlinear
- Frequent **availability of stabilizing controllers** around equilibrium or a reference



[1] Wisth, Camurri, Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots.” IEEE Transactions on Robotics, 2022

[2] Belke, Holdcroft, Sigrist, Paik, “Morphological flexibility in robotic systems through physical polygon meshing.” Nature Machine Intelligence, 2023

- Complex real-world systems are nonlinear
- Frequent **availability of stabilizing controllers** around equilibrium or a reference

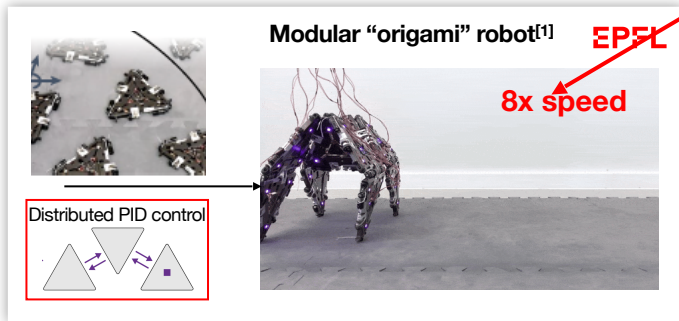


[1] Wisth, Camurri, Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots.” IEEE Transactions on Robotics, 2022

[2] Belke, Holdcroft, Sigrist, Paik, “Morphological flexibility in robotic systems through physical polygon meshing.” Nature Machine Intelligence, 2023

- Complex real-world systems are nonlinear
- Frequent **availability of stabilizing controllers** around equilibrium or a reference

... however, stability is not enough

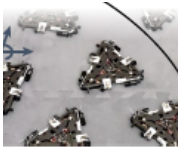


[1] Wisth, Camurri, Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots.” IEEE Transactions on Robotics, 2022

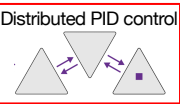
[2] Belke, Holdcroft, Sigrist, Paik, “Morphological flexibility in robotic systems through physical polygon meshing.” Nature Machine Intelligence, 2023

- Complex real-world systems are nonlinear
- Frequent **availability of stabilizing controllers** around equilibrium or a reference


... however, stability is not enough



Distributed PID control




Modular “origami” robot^[1]



8x speed

All-terrain legged-robots^[2]



>4x speed

CROSSING 25CM BARRIERS BY INVERTING ITS LEGS

ORI

Improve performance without compromising stability?

[1] Wisth, Camurri, Fallon, “VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots.” IEEE Transactions on Robotics, 2022

[2] Belke, Holdcroft, Sigrist, Paik, “Morphological flexibility in robotic systems through physical polygon meshing.” Nature Machine Intelligence, 2023

Parametrization of all nonlinear stabilizing controllers^[1]

[1] Eufieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]

[2] Eufieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{cases} x_t = f_t(x_{t-1}, u_{t-1}) + w_t \\ u_t = K_t(x_{t:0}) \end{cases} \quad \begin{matrix} \mathbf{x} = (x_0, x_1, \dots) \\ \mathbf{K}(\mathbf{x}) = (K_0(x_0), K_1(x_{1:0}), \dots) \end{matrix} \longrightarrow$$

Operator model

$$\begin{cases} \mathbf{x} = \mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{w} \\ \mathbf{u} = \mathbf{K}(\mathbf{x}) \end{cases}$$

Main result:^[1,2] Parametrization of robustly stabilizing controllers + completeness

Parametrization of all nonlinear stabilizing controllers^[1]

[1] Eufieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]


[2] Eufieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{cases} x_t = f_t(x_{t-1}, u_{t-1}) + w_t \\ u_t = K_t(x_{t:0}) \end{cases} \quad \begin{matrix} \mathbf{x} = (x_0, x_1, \dots) \\ \mathbf{K}(\mathbf{x}) = (K_0(x_0), K_1(x_{1:0}), \dots) \end{matrix} \longrightarrow$$

Operator model

$$\begin{cases} \mathbf{x} = \mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{w} \\ \mathbf{u} = \mathbf{K}(\mathbf{x}) \end{cases}$$

Main result:^[1,2] Parametrization of robustly stabilizing controllers + completeness



Assume the open-loop plant $\mathbf{x} = \mathcal{F}(\mathbf{u}, \mathbf{w})$ stable, or pre-stabilized. Then:

Parametrization of all nonlinear stabilizing controllers^[1]

[1] Eufier, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]


[2] Eufier, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{cases} x_t = f_t(x_{t-1}, u_{t-1}) + w_t \\ u_t = K_t(x_{t:0}) \end{cases} \quad \begin{matrix} \mathbf{x} = (x_0, x_1, \dots) \\ \mathbf{K}(\mathbf{x}) = (K_0(x_0), K_1(x_{1:0}), \dots) \end{matrix} \longrightarrow$$

Operator model

$$\begin{cases} \mathbf{x} = \mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{w} \\ \mathbf{u} = \mathbf{K}(\mathbf{x}) \end{cases}$$

Main result:^[1,2] Parametrization of robustly stabilizing controllers + completeness



Assume the open-loop plant $\mathbf{x} = \mathcal{F}(\mathbf{u}, \mathbf{w})$ stable, or pre-stabilized. Then:

- **Sufficiency:** the controller $\mathbf{K}(\mathbf{x})$ stabilizes the real system $\mathbf{F}(\mathbf{x}, \mathbf{u})$ $\text{gain}(\mathcal{M}) < \frac{1}{\text{gain}(\Delta)(\text{gain}(\mathcal{F}) + 1)}$

Parametrization of all nonlinear stabilizing controllers^[1]

[1] Eufieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]


[2] Eufieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{cases} x_t = f_t(x_{t-1}, u_{t-1}) + w_t \\ u_t = K_t(x_{t:0}) \end{cases} \quad \begin{matrix} \mathbf{x} = (x_0, x_1, \dots) \\ \mathbf{K}(\mathbf{x}) = (K_0(x_0), K_1(x_{1:0}), \dots) \end{matrix} \longrightarrow$$

Operator model

$$\begin{cases} \mathbf{x} = \mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{w} \\ \mathbf{u} = \mathbf{K}(\mathbf{x}) \end{cases}$$

Main result:^[1,2] Parametrization of robustly stabilizing controllers + completeness



Assume the open-loop plant $\mathbf{x} = \mathcal{F}(\mathbf{u}, \mathbf{w})$ is stable, or pre-stabilized. Then:

- **Sufficiency:** the controller $\mathbf{K}(\mathbf{x})$ stabilizes the real system $\mathbf{F}(\mathbf{x}, \mathbf{u})$ $\text{gain}(\mathcal{M}) < \frac{1}{\text{gain}(\Delta)(\text{gain}(\mathcal{F}) + 1)}$
- **Necessity:** if the system model is known ($\Delta = 0$), any stable closed-loop behavior is achieved by appropriately selecting a stable \mathcal{M}

Parametrization of all nonlinear stabilizing controllers^[1]

[1] Eufieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]


[2] Eufieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{cases} x_t = f_t(x_{t-1}, u_{t-1}) + w_t \\ u_t = K_t(x_{t:0}) \end{cases} \quad \begin{matrix} \mathbf{x} = (x_0, x_1, \dots) \\ \mathbf{K}(\mathbf{x}) = (K_0(x_0), K_1(x_{1:0}), \dots) \end{matrix} \longrightarrow$$

Operator model

$$\begin{cases} \mathbf{x} = \mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{w} \\ \mathbf{u} = \mathbf{K}(\mathbf{x}) \end{cases}$$

Main result:^[1,2] Parametrization of robustly stabilizing controllers + completeness



Assume the open-loop plant $\mathbf{x} = \mathcal{F}(\mathbf{u}, \mathbf{w})$ is stable, or pre-stabilized. Then:

- **Sufficiency:** the controller $\mathbf{K}(\mathbf{x})$ stabilizes the real system $\mathbf{F}(\mathbf{x}, \mathbf{u})$ $\text{gain}(\mathcal{M}) < \frac{1}{\text{gain}(\Delta)(\text{gain}(\mathcal{F}) + 1)}$
- **Necessity:** if the system model is known ($\Delta = 0$), any stable closed-loop behavior is achieved by appropriately selecting a stable \mathcal{M}

Importance of main result for nonlinear optimal control

[1] Furieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]

[2] Furieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

$$\begin{aligned} & \max_{\mathbf{K} \text{ is stabilizing}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ \text{subject to } & x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & u_t = K_t(x_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

Importance of main result for nonlinear optimal control

[1] Furieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]

[2] Furieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

Nonlinear generalization of "Youla", "System Level Synthesis",...

$$\begin{aligned} & \max_{\mathbf{K} \text{ is stabilizing}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ \text{subject to } & x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & u_t = K_t(x_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

\Leftrightarrow

$$\begin{aligned} & \max_{\mathcal{M} \text{ is stable}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ \text{subject to } & x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & \hat{w}_t = x_t - f_t(x_{t-1}, u_{t-1}) \\ & u_t = \mathcal{M}_t(\hat{w}_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

- **Theory:** directly optimize over stable closed-loop operator \mathcal{M}

Importance of main result for nonlinear optimal control

[1] Euzeni, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]

[2] Euzeni, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

Nonlinear generalization of "Youla", "System Level Synthesis",...

$$\begin{aligned} & \max_{\mathbf{K} \text{ is stabilizing}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ \text{subject to } & x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & u_t = K_t(x_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

\Leftrightarrow

$$\begin{aligned} & \max_{\mathcal{M} \text{ is stable}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ \text{subject to } & x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & \hat{w}_t = x_t - f_t(x_{t-1}, u_{t-1}) \\ & u_t = \mathcal{M}_t(\hat{w}_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

- **Theory:** directly optimize over **stable closed-loop operator** \mathcal{M}
- **Implementation:** compatible with **deep neural network** stable **operators**! [3]



[3] "Recurrent equilibrium networks: Flexible dynamic models with guaranteed stability and robustness", Revay, Wang, Manchester, [TAC23]

Importance of main result for nonlinear optimal control

[1] Furieri, Galimberti, F. Trecate, "Neural system level synthesis: learning over all and only the stabilizing controllers for nonlinear systems", [CDC 2022]

[2] Furieri, Galimberti, F. Trecate, "Learning to boost the performance of stable nonlinear systems", [ArXiv, 2024]

Nonlinear generalization of "Youla", "System Level Synthesis",...

$$\begin{aligned} & \max_{\mathbf{K} \text{ is stabilizing}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ & \text{subject to } x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & \quad u_t = K_t(x_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$



$$\begin{aligned} & \max_{\mathcal{M} \text{ is stable}} \mathbb{E}_{\mathbf{w}} [\text{performance}(\mathbf{x}, \mathbf{u})] \\ & \text{subject to } x_t = f_t(x_{t-1}, u_{t-1}) + w_t, \\ & \quad \hat{w}_t = x_t - f_t(x_{t-1}, u_{t-1}) \\ & \quad u_t = \mathcal{M}_t(\hat{w}_{t:0}), \quad t = 1, 2, \dots \end{aligned}$$

- **Theory:** directly optimize over **stable closed-loop operator** \mathcal{M}
- **Implementation:** compatible with **deep neural network** stable operators! [3]



[3] "Recurrent equilibrium networks: Flexible dynamic models with guaranteed stability and robustness", Revay, Wang, Manchester, [TAC23]


Remark: results compatible with distributed nonlinear control

[4] Massai, Saccani, Furieri, F. Trecate, "Unconstrained learning of **networked nonlinear systems** via free parametrization of stable interconnected operators", [ECC24]

[5] Saccani, Massai, Furieri, F. Trecate, "**Optimal distributed control** with stability guarantees by training a network of neural closed-loop maps", [ArXiv 2024]


Numerical experiments: mobile robots

...before performance-boosting



Numerical experiments: mobile robots



...before performance-boosting




Numerical experiments: mobile robots

$$\text{cost}(\mathbf{x}, \mathbf{u}) = \text{cost}_{\text{target}}(\mathbf{x}, \mathbf{u}) + \text{cost}_{\text{collisions}}(\mathbf{x})$$

...before performance-boosting





After performance-boosting



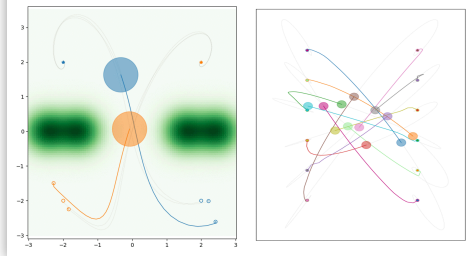
Numerical experiments: mobile robots

$$\text{cost}(\mathbf{x}, \mathbf{u}) = \text{cost}_{\text{target}}(\mathbf{x}, \mathbf{u}) + \text{cost}_{\text{collisions}}(\mathbf{x})$$

...before performance-boosting




After performance-boosting



CL stability guaranteed even with partial training

25% training



75% training

