

Nonparametric Analysis and Control of Dynamical Systems

Stability, Safety, and Policy Improvement

Enrique Mallada



CDS Seminar, Caltech

April 2nd, 2025

Acknowledgements



Yue Shen



Roy Siegelmann



Agustin Castellano



Sohrab Rezaei



Fernando Paganini



Maxim Bichuch



Hussein Sibai




Jared Markowitz



A Dream World of Success Stories

2017 Google DeepMind's DQN



LETTER

Human-level control through deep reinforcement learning

Vladimir Mnih¹, Koray Kavukcuoglu^{2*}, David Silver^{1*}, Andrei A. Ruus¹, Joel Veness¹, Marc G. Bellemare¹, Alex Graves¹, Martin Riedmiller¹, Andreas K. Földiars¹, Georg Ostrovski¹, Srik Petersen¹, Charles Beattie¹, Amir Sadik¹, Ioannis Antonoglou¹, Helen King¹, Dhruv Kumar¹, Quan Vuorio¹, Shaze Legg¹ & Demis Hassabis¹

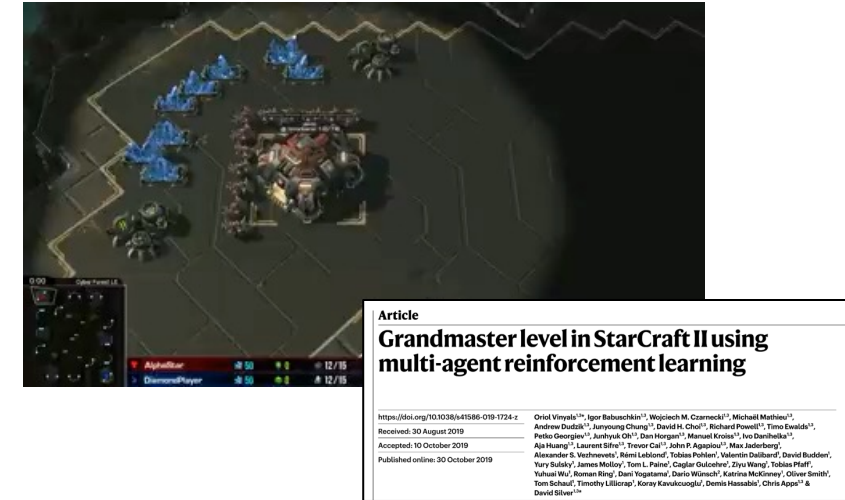
2017 AlphaZero – Chess, Shogi, Go




Boston Dynamics




2019 AlphaStar – Starcraft II



OpenAI – Rubik's Cube



Waymo



Reality Kicks In

Angry Residents, Abrupt Stops: Waymo Vehicles Are Still Causing Problems in Arizona

RAY STERN | MARCH 31, 2021 | 8:26AM

DeepMind's Losses and the Future of Artificial Intelligence

Alphabet's DeepMind unit, conqueror of Go and other games, is losing lots of money. Continued deficits could imperil investments in AI.

AARIAN MARSHALL | BUSINESS | 12.07.2020 04:06 PM

Uber Gives Up on the Self-Driving Dream

The ride-hail giant invested more than \$1 billion in autonomous vehicles. Now it's selling the unit to Aurora, which makes self-driving tech.

Tesla Recalls Nearly All Vehicles Due to Autopilot Failures

Tesla disagrees with feds' analysis of glitches

BY LINA FISHER, 2:54PM, WED. DEC. 13, 2023

CRUISE KNEW ITS SELF-DRIVING CARS HAD PROBLEMS RECOGNIZING CHILDREN — AND KEPT THEM ON THE STREETS


According to internal materials reviewed by The Intercept, Cruise cars were also in danger of driving into holes in the road.



OpenAI disbands its robotics research team

Kyle Wiggers | @Kyle_L_Wiggers | July 16, 2021 11:24 AM

f t in




Self-driving Uber car that hit and killed woman did not recognize that pedestrians jaywalk

The automated car lacked "the capability to classify an object as a pedestrian unless that object was near a crosswalk," an NTSB report said.



Fundamental challenge: The curse of dimensionality

- **Statistical: Sampling in d dimension with resolution ϵ**



Sample complexity:

$$O(\epsilon^{-d})$$

For $\epsilon = 0.1$ and $d = 100$, we would need 10^{100} points.
Atoms in the universe: 10^{78}

- **Computational: Verifying non-negativity of polynomials**

Copositive matrices:

$$\begin{bmatrix} x_1^2 & \dots & x_d^2 \end{bmatrix} A \begin{bmatrix} x_1^2 & \dots & x_d^2 \end{bmatrix}^T \geq 0$$


Murty&Kadabi [1987]: Testing co-positivity is NP-Hard

Sum of Squares (SoS):

$$z(x)^T Q z(x) \geq 0, \quad z_i(x) \in \mathbb{R}[x], \quad x \in \mathbb{R}^d, Q \succcurlyeq 0$$

Artin [1927] (Hilbert's 17th problem):

Non-negative polynomials are sum of square of *rational* functions



Motzkin [1967]:

$$p = x^4y^2 + x^2y^4 + 1 - 3x^2y^2$$


is nonnegative,

not a sum of squares,

but $(x^2 + y^2)^2 p$ is SoS

Methodological challenges

- Focused on a ***design-then-deploy*** philosophy
 - Most methods have a strict separation between control synthesis and deployment
- Synthesis usually aims for the ***best*** (optimal) controller
 - Lack of exploration of the benefits of designing sub-optimal controllers
- Policy ***parameters*** can ***drastically affect*** the system's ***behavior***
 - The params to behavior maps are highly sensitive to perturbations



RL:


$$\max_{\pi} J(\pi) = \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]$$

$$\text{s.t. } s_{t+1} \sim P(\cdot | s_t, a_t), \quad a_t \sim \pi(\cdot | s_t)$$

Optimal Control:

$$\min_{u(\cdot)} J = \int_0^T L(x(t), u(t), t) dt + \Phi(x(T))$$

$$\text{s.t. } \dot{x}(t) = f(x(t), u(t), t), \quad x(0) = x_0$$



Aspirational Goals



To design control policies as nature does...



refining post deployment self improving, with each trial avoiding poor decisions reinforcing good ones

Research Goals

- To develop analysis and design methods that *trade off complexity and performance*.
- To allow for *continual improvement*, without the need for redesign, retune, or retrain
- To design control policies with controlled sensitivity to parameter changes



Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations


Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t : $\phi(t, x_0)$.

Ω -Limit Set $\Omega(f)$:

$$x \in \Omega(f) \iff \exists x_0, \{t_n\}_{n \geq 0}, \text{ s.t. } \lim_{n \rightarrow \infty} t_n = \infty \text{ and } \lim_{n \rightarrow \infty} \phi(t_n, x_0) = x$$



Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t : $\phi(t, x_0)$.


Ω -Limit Set $\Omega(f)$:

$$x \in \Omega(f) \iff \exists x_0, \{t_n\}_{n \geq 0}, \text{ s.t. } \lim_{n \rightarrow \infty} t_n = \infty \text{ and } \lim_{n \rightarrow \infty} \phi(t_n, x_0) = x$$


Types of Ω -limit set




equilibrium



limit cycle



limit torus






chaotic attractor

Remark: invariance is a shared property, thus a natural tool for analysis

Invariant sets

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S}, \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times




Invariant sets: Merits

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S}, \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

- Invariant sets approximate regions of attraction**

Compact invariant set $\mathcal{S}$, containing **only** $\{x^*\} = \Omega(f) \cap \mathcal{S}$ must be in the region of attraction $\mathcal{A}(x^*)$ ($\mathcal{S} \subset \mathcal{A}(x^*)$)



Invariant sets: Merits

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S}, \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

- Invariant sets approximate regions of attraction**



Compact invariant set $\mathcal{S}$, containing **only** $\{x^*\} = \Omega(f) \cap \mathcal{S}$ must be in the region of attraction $\mathcal{A}(x^*)$ ($\mathcal{S} \subset \mathcal{A}(x^*)$)

- Invariant sets guarantee stability**

Lyapunov stability: solutions starting "close enough" to the equilibrium (within a distance δ) remain "close enough" forever (within a distance ε)

- Invariant sets further certify asymptotic stability via Lyapunov's direct method**

Asymptotic stability: solutions that start close enough, remain close enough, and eventually converge to equilibrium.




Invariant sets: Challenges

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S}, \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times



- $\mathcal{S}$ is topologically constrained
 - If $\mathcal{S} \cap \Omega(f) = \{x^*\}$, then $\mathcal{S}$ is connected
- $\mathcal{S}$ is geometrically constrained
 - f should not point outwards for $x \in \partial\mathcal{S}$
- $\mathcal{S}$ geometry can be wild
 - $\mathcal{A}(\Omega(f))$ can be fractal

$\mathcal{S}$: 
 $\mathcal{A}(x^*)$: 



A not invariant trajectory: 

Basin of $\Omega(f)$



Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations


Recurrent sets: Letting things go, and come back


A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Property of Recurrent Sets

- $\mathcal{R}$ need **not** be **connected**
- $\mathcal{R}$ does **not** require f to **point inwards** on all $\partial\mathcal{R}$

Recurrent sets, while not invariant,
guarantee that solutions that start in this set,
will come back **infinitely often, forever!**





Recurrent set $\mathcal{R}$: 

A recurrent trajectory: 

Recurrent sets: Letting things go, and come back

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Previous two good inner approximations of $\mathcal{A}(x^*)$ are recurrent sets





Recurrent sets: Letting things go, and come back


A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Property of Recurrent Sets

- $\mathcal{R}$ need **not** be **connected**
- $\mathcal{R}$ does **not** require f to **point inwards** on all $\partial\mathcal{R}$

Recurrent sets, while not invariant,
guarantee that solutions that start in this set,
will come back **infinitely often, forever!**



Recurrent set $\mathcal{R}$: 

A recurrent trajectory: 

Question: Can we use recurrent sets as functional substitutes of invariant sets?

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations



Roy Siegelmann



Yue Shen



Fernando Paganini



Nonparametric Stability Analysis

R. Siegelmann, Y. Shen, F. Paganini, and E. Mallada, “A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions”, CDC 2023



R. Siegelmann, Y. Shen, F. Paganini, and E. Mallada, “Recurrent Lyapunov Functions”, TAC 2025, submitted

Lyapunov's Direct Method

Key idea: Make sub-level sets invariant to trap trajectories

Theorem [Lyapunov '1892]. Given $V: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, with $V(x) > 0, \forall x \in \mathbb{R}^d \setminus \{x^*\}$, then:

- $\dot{V} \leq 0 \rightarrow x^*$ stable
- $\dot{V} < 0 \rightarrow x^*$ as. stable



Challenge: Couples shape of V and vector field f

- Towards decoupling the $V - f$ geometry
 - Controlling regions where $\dot{V} \geq 0$ [Karafyllis '09, Liu et al '20]
 - Higher order conditions: $g(V^{(q)}, \dots, \dot{V}, V) \leq 0$ [Butz '69, Gunderson '71, Ahmadi '06, Meigoli '12]
 - Discretization approach: $V(x(T)) \leq V(x(0))$ [Coron et al '94, Aeyels et. al '98, Karafyllis '12]
 - Multiple Lyapunov Functions: $\{V_j: j \in [k]\}$ [Ahmadi et al '14]

A Butz. Higher order derivatives of Lyapunov functions. IEEE Transactions on automatic control, 1969

Gunderson. A comparison lemma for higher order trajectory derivatives. Proceedings of the American Mathematical Society, 1971

Coron, Lionel Rosier. A relation between continuous time-varying and discontinuous feedback stabilization. J. Math. Syst., Estimation, Control, 1994

Aeyels, Peuteman. A new asymptotic stability criterion for nonlinear time-variant differential equations. IEEE Transactions on automatic control, 1998

Ahmadi. Non-monotonic Lyapunov functions for stability of nonlinear and switched systems: theory and computation, 2008

Karafyllis, Kravaris, Kalogerakis. Relaxed Lyapunov criteria for robust global stabilisation of non-linear systems. International Journal of Control, 2009

Meigoli, Nikravesh. Stability analysis of nonlinear systems using higher order derivatives of Lyapunov function candidates. Systems & Control Letters, 2012

Karafyllis. Can we prove stability by using a positive definite function with non sign-definite derivative? IMA Journal of Mathematical Control and Information, 2012

Ahmadi, Jungers, Parrilo, Roozbehani. Joint spectral radius and path-complete graph Lyapunov functions. SIAM Journal on Control and Optimization, 2014



Liu, Liberzon, Zharnitsky. Almost Lyapunov functions for nonlinear systems. Automatica, 2020

Lyapunov's Direct Method

Key idea: Make sub-level sets invariant to trap trajectories

Theorem [Lyapunov '1892]. Given $V: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, with $V(x) > 0, \forall x \in \mathbb{R}^d \setminus \{x^*\}$, then:

- $\dot{V} \leq 0 \rightarrow x^*$ stable
- $\dot{V} < 0 \rightarrow x^*$ as. stable



Challenge: Couples shape of V and vector field f

- Towards decoupling the $V - f$ geometry
 - Controlling regions where $\dot{V} \geq 0$ [Karafyllis '09, Liu et al '20]
 - Higher order conditions: $g(V^{(q)}, \dots, \dot{V}, V) \leq 0$ [Butz '69, Gunderson '71, Ahmadi '06, Meigoli '12]
 - Discretization approach: $V(x(T)) \leq V(x(0))$ [Coron et al '94, Aeyels et. al '98, Karafyllis '12]
 - Multiple Lyapunov Functions: $\{V_j: j \in [k]\}$ [Ahmadi et al '14]

Question: Can we provide stability conditions based on recurrence?



Recurrent Lyapunov Functions

A continuous function $V: \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a **Recurrent Lyapunov Function** if

$$\mathcal{L}_f^{(0,\tau]} V(x) := \min_{t \in (0,\tau]} V(\phi(t, x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Preliminaries:

- Sub-level sets $\{V(x) \leq c\}$ are τ -recurrent sets.



Definition: A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **τ -recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \in (t, t + \tau]$ s.t. $\phi(t', x_0) \in \mathcal{R}$.


Recurrent Lyapunov Functions

A continuous function $V: \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a **Recurrent Lyapunov Function** if


$$\mathcal{L}_f^{(0,\tau]} V(x) := \min_{t \in (0,\tau]} V(\phi(t, x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Preliminaries:

- Sub-level sets $\{V(x) \leq c\}$ are τ -recurrent sets.
- When f is L -Lipschitz, one can trap trajectories.



$$F_c = \max_{x \in V_{\leq c}} \|f(x)\|$$




Recurrent Lyapunov Functions

A continuous function $V: \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a **recurrent Lyapunov function** if

$$\mathcal{L}_f^{(0,\tau]} V(x) := \min_{t \in (0,\tau]} V(\phi(t, x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Theorem [CDC 23]: Let $V: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be a recurrent Lyapunov function and let f be L -Lipschitz

- Then, the equilibrium x^* is stable.
- Further, if the **inequality is strict**, then x^* is asymptotically stable!



Exponential Stability Analysis


The function $V: \mathbb{R}^d \rightarrow \mathbb{R}_+$ is **α -exponential** recurrent Lyapunov function if

$$\mathcal{L}_{f,\alpha}^{(0,\tau]} V(x) := \min_{t \in (0,\tau]} e^{\alpha t} V(\phi(t, x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Theorem [CDC 23]: Let $V: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ satisfy

$$\alpha_1 \|x - x^*\| \leq V(x) \leq \alpha_2 \|x - x^*\|.$$

Then, if V is **α -exponential** recurrent Lyapunov function, x^* is **α -exponentially stable**.



Norm-based Converse Theorem

Theorem: Assume x^* is λ -exponentially stable: $\exists K, \lambda > 0$ such that:

$$||\phi(t, x) - x^*|| \leq K e^{-\lambda t} ||x - x^*||, \quad \forall x \in \mathbb{R}^d.$$

Then, $V(x) = ||x - x^*||$ is α -exponential recurrent Lyapunov function, i.e.,

$$\min_{t \in (0, \tau]} e^{\alpha t} ||\phi(t, x) - x^*|| - ||x - x^*|| \leq 0, \quad \forall x \in \mathbb{R}^d,$$

whenever $\alpha < \lambda$ and $\tau \geq \frac{1}{\lambda - \alpha} \ln K$.

Remarks:

- The rate α must be strictly smaller than the rate of convergence λ (trading off optimality).
- Any norm is a Lyapunov function!

Question: Is the struggle for its search over?

Nonparametric Verification of Exponential Stability

Proposition [CDC 23*]: Let $||\cdot||$ be any norm and $x^* = 0$. Then, whenever

$$\min_{t \in (0, \tau]} e^{\alpha t} (||\phi(x, t)|| + r e^{Lt}) \leq ||x|| - r$$

for all y with $||y - x|| \leq r$

$$\min_{t \in (0, \tau]} e^{\alpha t} ||\phi(y, t)|| \leq ||y||$$



Remarks:

- Only requires a trajectory of length τ
- Trades off between **radius** r and verified **performance** α
- Amenable for parallel computations **using GPUs**

Nonparametric Stability Verification via GPUs

- **Basic Algorithm:**


- Consider $V(x) = ||x - x^*||_\infty$
- Build a grid of hypercubes surrounding x^*
- Test grid center points:
 - Simulate trajectories of length τ
 - Find α s.t. the verified radius is $r \geq \ell/2$
- Hypercube **not verified**, **split in 3^d** parts
- Repeat testing of new points



Nonparametric Stability Verification via GPUs

- **Basic Algorithm:**

- Consider $V(x) = ||x - x^*||_\infty$
- Build a grid of hypercubes surrounding x^*
- Test grid center points:
 - Simulate trajectories of length τ
 - Find α s.t. the verified radius is $r \geq \ell/2$
- Hypercube **not verified, split in 3^d** parts
- Repeat testing of new points
- **Exponentially expand** to outer layer
- Repeat testing in new layer



Nonparametric Stability Verification via GPUs

- **Basic Algorithm:**


- Consider $V(x) = ||x - x^*||_\infty$
- Build a grid of hypercubes surrounding x^*
- Test grid center points:
 - Simulate trajectories of length τ
 - Find α s.t. the verified radius is $r \geq \ell/2$
- Hypercube **not verified**, split in 3^d parts
- Repeat testing of new points
- Exponentially expand to outer layer
- Repeat testing in new layer

Q: How many samples are needed?

If x^* is λ -exp. stable

$$\mathcal{O} \left(q^{-d} \log \left(\frac{R}{\varepsilon} \right) \right)$$

with $q = \frac{1 - Ke^{(\alpha - \lambda)\tau}}{1 + e^{(L + \alpha)\tau}} < 1$.




Nonparametric Stability Verification via GPUs

- **Basic Algorithm:**

- Consider $V(x) = ||x - x^*||_\infty$
- Build a grid of hypercubes surrounding x^*
- Test grid center points:
 - Simulate trajectories of length τ
 - Find α s.t. the verified radius is $r \geq \ell/2$
- Hypercube **not verified**, **split in 3^d** parts
- Repeat testing of new points
- Exponentially expand to outer layer
- Repeat testing in new layer

- **Two Alg. Variations:**

- Alg. 1: Find largest $\alpha_{\max}$ for region $\mathcal{X}$
- Alg. 2: Find region $\mathcal{X}_\alpha$ for given α





Numerical Illustration – Find Best α

Consider the 2-d non-linear system:
with $B_{ij} \sim \mathcal{N}(0, \sigma^2)$

$$\dot{x} = \begin{bmatrix} 0 & 2 \\ -1 & -1 \end{bmatrix} x + B \begin{bmatrix} x_1^2 \\ x_1 x_2 \\ x_2^2 \end{bmatrix}$$

$$\sigma = 0.3$$

$$\alpha_{\max} = 0.470$$




Numerical Illustration – Find region $\mathcal{X}_\alpha$


Consider the system of n Kuramoto oscillators:

$$\dot{\theta}_i = \frac{k}{n} \sum_{j=1}^n \sin(\theta_j - \theta_i)$$

Parameters: $n = 3$ and $\alpha = 1$



System dimension: $d = n - 1$



Outline


- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations



Yue Shen



Hussein Sibai



Nonparametric Safety Verification using Recurrence


Y. Shen, H. Sibai, E. Mallada, “Generalized Barrier Functions: Integral Conditions and Recurrent Relaxations”, in 60th Allerton Conference on Communication, Control, and Computing 2024

Safety in Dynamical Systems

Consider the continuous-time dynamical system: $\dot{x} = f(x)$

- $\phi(t, x_0)$: solution at time t starting from x_0
- X_u : set of unsafe states

Goal: Find the safe set $\mathcal{X}_s := \{x_0 \in \mathbb{R}^d | \phi(t, x_0) \notin \mathcal{X}_u, \forall t \geq 0\}$



Safety in Dynamical Systems **via Invariant Sets**


Consider the continuous-time dynamical system: $\dot{x} = f(x)$

- $\phi(t, x_0)$: solution at time t starting from x_0
- X_u : set of unsafe states

Goal: Find the safe set $\mathcal{X}_s := \{x_0 \in \mathbb{R}^d \mid \phi(t, x_0) \notin \mathcal{X}_u, \forall t \geq 0\}$

General Approach: Use invariant sets!

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S}, \forall t \geq 0$




Certifying Safety using Barrier Functions

Theorem - Nagumo's Barrier Functions [Nagumo '42] :

Let $h: \mathbb{R}^d \rightarrow \mathbb{R}$ be differentiable, with 0 being a *regular value*.
Then h is a Nagumo's Barrier Function (NBF) satisfying:

$$L_f h(x) := \lim_{t \rightarrow 0} \frac{h(\phi(t, x)) - h(x)}{t} \geq 0, \quad \forall x \in h_{=0},$$


if and only if $h_{\geq 0} := \{x \in \mathbb{R}^d \mid h(x) \geq 0\}$ is invariant.



Then $h_{\geq 0}$ is a safe set whenever $h_{\geq 0} \cap X_u = \emptyset$



Mitio Nagumo



Shaping Safe Behavior using Barrier Functions (BFs)

Barrier functions provide a flexible framework to shape the behavior of trajectories

Nagumo's (NBF)


$$L_f h(x) \geq 0, \quad \forall x \in h=0$$

Exponential Barrier Functions (EBF)

$$L_f h(x) \geq -\alpha h(x), \quad \forall x \in h \geq -c$$

Problem: Finding Barrier Functions is usually a complex undertaking

Key Challenge: The *invariance condition* on $h_{\geq 0}$ couples the geometry of f and the set $h_{\geq 0}$



Other: Zeroing BF (ZBFs), Minimal BF (MBFs), Control BF (CBFs), High Order CBFs (HOCBFs), ...

S. Prajna, A. Jadbabaie. *Safety Verification of Hybrid Systems Using Barrier Certificates*. HSCC 2004

P. Wieland, F. Allgöwer. *Constructive safety using control barrier functions*. IFAC Proceedings Volumes 2007

A. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, P. Tabuada. *Control barrier functions: Theory and applications*. IEEE ECC 2019

R. Konda, A. Ames, S. Coogan. *Characterizing safety: Minimal control barrier functions from scalar comparison systems*. IEEE L-CSS 2020

W. Xiao, C. Belta. *High-order control barrier functions*. IEEE TAC 2021

Integral Nagumo's Barrier Function (INBF)

Nagumo's Barrier Functions [Nagumo '42]:

Let h be differentiable, regular at zero, and

$$L_f h(x) \geq 0, \forall x \in h_{=0}$$

Integral Nagumo's Barrier Function:



Let h be continuous, and

$$h(\phi(t, x)) \geq 0, \forall x \in h_{=0}, t \geq 0$$



The super-level set $h_{\geq 0}$ is invariant.

* requires more conditions on h



Nagumo's Recurrent Barrier Function (RNBF)

Thm: Integral Nagumo's Barrier Function:

Let h be continuous. Then:

$$h(\phi(t, x)) \geq 0, \forall x \in h_{=0}, t \geq 0$$

if and only if $h_{\geq 0}$ is **invariant**



Thm: Nagumo's Recurrent Barrier Function:

Let h be continuous. Then:

$$\max_{t \in (0, \tau]} h(\phi(t, x)) \geq 0, \forall x \in h_{=0}$$

if and only if $h_{\geq 0}$ is **τ -recurrent**

Recall: A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **τ -recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \in (t, t + \tau]$ s.t. $\phi(t', x_0) \in \mathcal{R}$.



Recurrent Nagumo's Barrier Function (RNBF)

Thm: Integral Nagumo's Barrier Function:

Let h be continuous. Then:

$$h(\phi(t, x)) \geq 0, \forall x \in h_{=0}, t \geq 0$$

if and only if $h_{\geq 0}$ is **invariant**

As $\tau \rightarrow 0$
 $\longleftrightarrow$
 By definition



Thm: Recurrent Nagumo's Barrier Function:

Let h be continuous. Then:

$$\max_{t \in (0, \tau]} h(\phi(t, x)) \geq 0, \forall x \in h_{=0}$$

if and only if $h_{\geq 0}$ is **τ -recurrent**

Recall: A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **τ -recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \in (t, t + \tau]$ s.t. $\phi(t', x_0) \in \mathcal{R}$.



Integral Exponential Barrier Function (IEBF)

Exponential Barrier Functions:

Let h be differentiable, and

$$L_f h(x) \geq -\alpha h(x), \quad \forall x \in h_{\geq -c}$$



Integral Exponential Barrier Function:

Let h be continuous, and

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x), \quad \forall x \in h_{\geq -c}$$

for all $t \geq 0$

The super-level set $h_{\geq 0}$ is invariant.



* requires more conditions on h

α – Exponential Recurrent Barrier Function (α –ERBF)

Thm: Integral Exponential Barrier Function:

Let h be continuous. If:

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x), \quad \forall x \in h_{\geq -c}$$



for all $t \geq 0$, then, $h_{\geq 0}$ is **invariant**

Thm: α -Exponential Recurrent Barrier Function:

Let h be continuous. If:

$$\max_{t \in (0, \tau]} e^{\alpha t} h(\phi(t, x)) \geq h(x), \quad \forall x \in h_{\geq -c}$$

then, $h_{\geq 0}$ is **τ -recurrent**



α – Exponential Recurrent Barrier Function (α –ERBF)

Thm: Integral Exponential Barrier Function:

Let h be continuous. If:

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x), \quad \forall x \in h_{\geq -c}$$

for all $t \geq 0$, then, $h_{\geq 0}$ is **invariant**

As $\tau \rightarrow 0$
 $\longleftrightarrow$
 By definition



Thm: α -Exponential Recurrent Barrier Function:

Let h be continuous. If:

$$\max_{t \in (0, \tau]} e^{\alpha t} h(\phi(t, x)) \geq h(x), \quad \forall x \in h_{\geq -c}$$

then, $h_{\geq 0}$ is **τ -recurrent**



Question: Do we gain anything from relaxing the invariance condition in BFs?



α, β – Recurrent Exponential Barrier Functions

We first generalize REBF using different exponential rates $\alpha, \beta > 0$:

$$\max_{t \in (0, \tau]} e^{\alpha t} [h(\phi(t, x))]_{+} + e^{\beta t} [h(\phi(t, x))]_{-} \geq h(x), \quad \forall x \in h_{\geq -c}$$



Signed Norms are Recurrent Barrier Functions!

Theorem: Assume there exists an **Exponential BF (EBF)** h :

$$\dot{h}(x) \geq -\alpha h(x), \quad \forall x \in h_{\geq -c}$$

Then for any set $\mathcal{S}$ with $h_{\geq 0} \subseteq \mathcal{S} \subseteq h_{\geq -c}$, the function


$$\hat{h}(x) := -\text{sd}(x, \mathcal{S})$$

is an $\bar{\alpha}, \underline{\alpha}$ -ERBF, if $\underline{\alpha} < \alpha < \bar{\alpha}$ and $\tau \geq \Omega\left(\frac{1}{\bar{\alpha}-\alpha} + \frac{1}{\alpha-\underline{\alpha}}\right)$

Remarks:

- The rates $\underline{\alpha} < \bar{\alpha}$ must be strictly smaller/bigger than α (giving up optimality).
- *Any signed norm for most sets is a Recurrent Barrier function!*

Question: How to use Recurrent Barriers for safety?



Certifying Safety using Recurrent Sets


Theorem - Consider a closed set S that is τ -recurrent. Then its τ -reachable tube:

$$\mathcal{R}_{[0,\tau]}(S) := \bigcup_{\substack{x \in S \\ t \in [0,\tau]}} \phi(t, x)$$

is **invariant**.


Moreover, S is **safe** when:

1. $\mathcal{R}_{[0,\tau]}(S) \cap \mathcal{X}_u = \emptyset$, or
2. $S \cap \mathcal{R}_{[-\tau,0]}(\mathcal{X}_u) = \emptyset$




Basic Algorithm

- Given unsafe region $\mathcal{X}_u$, precision $r_{\min}$
- Build initial grid of hypercubes $\mathcal{G} = \{g_i := B_{r_i}(x_i)\}$
- **Stage 1:** τ —Backward reachability
 - Find $S_u = \cup_i B(x_i, r_i)$ such that: $\mathcal{R}_{[-\tau, 0]}(\mathcal{X}_u) \subset S_u$
- **Stage 2:** Check RBF on $h(x) = -\text{sd}(x, (S_u)^c)$



Basic Algorithm

- Given unsafe region $\mathcal{X}_u$, precision $r_{\min}$
- Build initial grid of hypercubes $\mathcal{G} = \{g_i := B_{r_i}(x_i)\}$
- **Stage 1:** τ —Backward reachability
 - Find $S_u = \cup_i B(x_i, r_i)$ such that: $\mathcal{R}_{[-\tau, 0]}(\mathcal{X}_u) \subset S_u$
- **Stage 2:** Check RBF on $h(x) = -\text{sd}(x, (S_u)^c)$
 - For $g_i \in \mathcal{G}$, while $\mathcal{G}$ not empty:
 - If: g_i satisfies α —ERBF condition, **continue**
 - Else if: g_i can never satisfy α —ERBF condition, **add g_i to S_u**
 - Else: **refine grid**



$\mathcal{X}_u$:


S_u :

Basic Algorithm

- Given unsafe region $\mathcal{X}_u$, precision $r_{\min}$
- Build initial grid of hypercubes $\mathcal{G} = \{g_i := B_{r_i}(x_i)\}$
- **Stage 1:** τ —Backward reachability
 - Find $S_u = \cup_i B(x_i, r_i)$ such that: $\mathcal{R}_{[-\tau, 0]}(\mathcal{X}_u) \subset S_u$
- **Stage 2:** Check RBF on $h(x) = -\text{sd}(x, (S_u)^c)$
 - For $g_i \in \mathcal{G}$, while $\mathcal{G}$ not empty:
 - If: g_i satisfies RBF condition, **continue**
 - Else if: g_i can never satisfy α —ERBF condition, **add g_i to S_u**
 - Else: **refine grid**
 - **Finish when:**
 - all points satisfy RBF condition, or precision $r_{\min}$ is reached

Claim:

- The set $\mathcal{S} = (S_u)^c$ satisfies: $\mathcal{S} \cap \mathcal{R}_{[-\tau, 0]}(\mathcal{X}_u)$
- The function $h(x) = -\text{sd}(x, \mathcal{S})$ is an α —ERBF



$\mathcal{X}_u$: 

S_u : 



Nonparametric Safety Verification – Stage 1

Stage 1: τ –Backward reachability

- Find S_u with $\mathcal{R}_{[-\tau,0]}(X_u) \subset S_u$

Stage 2: RBF condition

- Check $h(x) = -\text{sd}(x, S)$ is α -ERBF





Nonparametric Safety Verification – Stage 1

Stage 1: τ –Backward reachability


- Find S_u with $\mathcal{R}_{[-\tau,0]}(X_u) \subset S_u$

Stage 2: RBF condition


- Check $h(x) = -\text{sd}(x, S)$ is RBF



Numerical Validation: Reachability vs Recurrence





running time: 354.12s
97.4% safe set covered



running time: 20.68s
99.6% safe set covered



~17x faster
+2.2% more area

Numerical Validation: Reachability vs Recurrence





~6x faster
+6.4% more area

Numerical Validation: Reachability vs Recurrence



~2x faster
+11.2% more area


Numerical Validation: Reachability vs Recurrence




6x faster
+61.9% more area

Numerical Validation: α – Exponential RBF


$\alpha = 4$




$\alpha = 3$



$\alpha = 2$



$\alpha = 1$



Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Outline

- **Relaxing Invariance: Merits and trade offs**
 - *Recurrent Sets*: Letting thing go and come back
- **Nonparametric Analysis via Recurrent Sets**
 - *Stability analysis*: Recurrent Lyapunov Functions (RLFs)
 - *Safety verification*: Recurrent Barrier functions (RBFs)
- **Self-Improving via Nonparametric Control Policies**
 - Policy Improvement using Expert Demonstrations

Classical policy improvement works in discrete spaces

“Policy improvement” is a fundamental building block of classical RL

Policy iteration = Policy evaluation + Policy improvement

Policy evaluation

- Given π , evaluate it to find $Q^\pi(\cdot, \cdot)$

- Can evaluate “separately” for each (s, a)
- Can store Q in a table

Policy improvement

- Given $Q^\pi(\cdot, \cdot)$, define: $\pi' : \mathcal{S} \rightarrow \mathcal{A} : \pi'(s) \in \operatorname{argmax}_{a \in \mathcal{A}} Q^\pi(s, a)$

- Then:

$$V^{\pi'}(s) \geq V^\pi(s) \quad \forall s \in \mathcal{S}$$

- Given a , maximize an array of size $\mathcal{S}$

Rinse and repeat until $V^{\pi'} \equiv V^\pi \implies \pi = \pi' = \pi^*$

Policy iteration in continuous spaces is hard

Policy evaluation

- *Approximately* evaluate $V^\pi(s)$ & $Q^\pi(s, a)$ w/ function approximation.
- E.g. using neural networks and Bellman Residual Minimization:

$$\mathcal{L}_{\text{BRM}}(Q) := \frac{1}{|\mathcal{D}|} \sum_{\mathcal{D}} (Q(s, a) - (r + \gamma Q(s', a')))^2.$$

Policy improvement

- Given $Q^\pi(\cdot, \cdot)$, how can we do $\pi' : \mathcal{S} \rightarrow \mathcal{A} : \pi'(s) \in \operatorname{argmax}_{a \in \mathcal{A}} Q^\pi(s, a)$?
- Only works if $\mathcal{A}$ discrete and finite (e.g. DQN)

Modern RL alg's. guarantee improvement *on average*

E.g. *Trust region* methods (TRPO, PPO).

- Given initial state distribution ρ , the “value” of policy π is:

$$\eta(\pi) = \mathbb{E}_{s_0 \sim \rho, a \sim \pi(\cdot|s)} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]$$

- Guarantees are usually of the form:

$$\eta(\pi') \geq \eta(\pi)$$

provided the new policy is not *too far* away from the old one.

Drawbacks:^{[1],[2]}

- Need to take small steps.
- Sensitivity to hyperparameter tuning.

[1] Wang, Tao, Sylvia Herbert, and Sicun Gao, "Fractal landscapes in policy optimization," NeurIPS, 2023

[2] Wang, Tao, Sylvia Herbert, and Sicun Gao, "Mollification effects of policy gradient methods," arXiv preprint arXiv:2405.17832, 2024



Agustin Castellano



Sohrab Rezaei



Jared Markowitz




Nonparametric policy improvement in continuous action spaces

A. Castellano, S. Rezaei, J. Markovitz, and E. Mallada, Nonparametric Policy Improvement for Continuous Action Spaces via Expert Demonstrations, 2025, submitted to RLC.

Problem Setup

Markov Decision Process:

$$\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, f, \rho, \gamma \rangle$$


Goal: find optimal policy

$$Q^*(s, a) = \max_{\pi} \left(Q^{\pi}(s, a) \triangleq \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right] \right)$$

- **Assumption:** Q^* is L-Lipschitz:

$$|Q^*(s, a) - Q^*(s', a')| \leq L(d_{\mathcal{S}}(s, s') + d_{\mathcal{A}}(a, a'))$$

- This is satisfied, e.g., if $f(\cdot, \cdot)$ and $r(\cdot, \cdot)$ are Lipschitz and $\gamma L_f < 1$

- **Fact:** If Q^* is L-Lipschitz then V^* is L-Lipschitz

Expert data: we have $\mathcal{D} = \{(s_i, a_i, Q_i)\}_{i=1}^{|\mathcal{D}|}$, where $a_i = \pi^*(s_i); Q_i = Q^*(s_i, a_i)$

Expert data: we have $\mathcal{D} = \{(s_i, a_i, Q_i)\}_{i=1}^{|\mathcal{D}|}$, where $a_i = \pi^*(s_i)$; $Q_i = Q^*(s_i, a_i)$

1. **How** can we use these transitions to learn a nonparametric policy?
2. **What** guarantees can we get when we add more transitions?
3. **Where** should we add transitions to improve performance?

Overview of our method




Figure 1: Overview of the proposed method. From left to right: *i*) a dataset containing expert triplets (s_i, a_i, Q_i) is used to *ii*) build a lower bound on the optimal value function; *iii*) acting greedily with respect to it gives our policy; *iv*) if high-probability suboptimality conditions are not met, we collect more expert trajectories and repeat the process.

1. How can we use these transitions to learn a nonparametric policy?

Building bounds & Nonparametric Policy

Expert data: we have $\mathcal{D} = \{(s_i, a_i, Q_i)\}_{i=1}^{|\mathcal{D}|}$, where $a_i = \pi^*(s_i)$; $Q_i = Q^*(s_i, a_i)$

- Use the data to define **lower** bounds on optimal values:

$$V_{\text{lb}}(s) \triangleq \max_{1 \leq i \leq |\mathcal{D}|} \{Q_i - L \cdot d_{\mathcal{S}}(s, s_i)\} \quad Q_{\text{lb}}(s, a) \triangleq \max_{1 \leq i \leq |\mathcal{D}|} \{Q_i - L \cdot (d_{\mathcal{S}}(s, s_i) + d_{\mathcal{A}}(a, a_i))\}$$

- **Define:** *greedy* policy w.r.t. lower bound

$$\pi(s) \triangleq \operatorname{argmax}_{a \in \mathcal{A}} Q_{\text{lb}}(s, a)$$

- Note argmax always gives actions in dataset, and $Q_{\text{lb}}(s, \pi(s)) = V_{\text{lb}}(s)$
- What can we say about $V^{\pi}(s)$?


Nonparametric policy *improves* over lower bound

(Made up of optimal trajectories)

Theorem 3.4 (Policy evaluation). Let $\mathcal{D}$ be a consistent dataset (Definition 3.3) and π as in Definition 3.1. Then, for all $s \in \mathcal{S}$ the following two inequalities hold:

$$V_{\text{lb}}(s) \leq r(s, \pi(s)) + \gamma V_{\text{lb}}(f(s, \pi(s)))$$

$$V_{\text{lb}}(s) \leq V^\pi(s) \leq V^*(s).$$




Theorem 3.5 (Policy improvement). Let $\mathcal{D}, \mathcal{D}'$ be consistent datasets with $\mathcal{D} \subset \mathcal{D}'$. Let V_{lb} and V'_{lb} be the lower bounds constructed with $\mathcal{D}$ and $\mathcal{D}'$ respectively. Then the following **non-deterioration** conditions hold:


More data = better lower bounds

$$V_{\text{lb}}(s) \leq V'_{\text{lb}}(s) \quad \forall s \in \mathcal{S},$$


Improvement on added points

$$V^\pi(s) \leq V^{\pi'}(s) \quad \forall s \in \Pi_{\mathcal{S}}[\mathcal{D}' \setminus \mathcal{D}],$$





1. How to learn a policy?




2. What guarantees with more transitions?

More data = better lower bounds

$V_{\text{lb}}(s) \leq V'_{\text{lb}}(s) \quad \forall s \in \mathcal{S}$, and

Improvement on added points

$V^\pi(s) \leq V^{\pi'}(s) \quad \forall s \in \Pi_{\mathcal{S}}[\mathcal{D}' \setminus \mathcal{D}]$,



3. Where to add transitions?

- Any new trajectory improves the bound...
- However, adding points close to other trajectories won't help much...
- We need a notion of **performance**

Performance metric $\Delta(s)$ guides where to sample data

- Given $\varepsilon > 0$, we want guarantees of the form:

$$V^*(s) - V^\pi(s) \leq \varepsilon \quad \forall s \in \mathcal{S} ? \quad \forall s \in \text{Supp}(\rho) ?$$

Too much to ask!

- More reasonable:

$$\mathbb{E}_{s \sim \rho} [V^*(s) - V^\pi(s)] \leq \varepsilon \quad \text{or} \quad \mathbb{P}_{s \sim \rho} [V^*(s) - V^\pi(s) \leq \varepsilon] \geq p$$

- Sample initial state $s \sim \rho$
- If we know $V^*(s) - V^\pi(s) \leq \varepsilon$, our policy is good enough, don't add data

Many ways to measure this gap

Less conservative



Gap	True	bound-based	distance-based
$\Delta(s)$	$V^*(s) - V^\pi(s)$	$V_{\text{ub}}(s) - V_{\text{lb}}(s)$	$2L \cdot \min_{1 \leq i \leq \mathcal{D} } d_{\mathcal{S}}(s, s_i)$

Easier to compute



Algorithm


Input: Lipschitz constant L

For each episode **do**:

1. Sample $s \sim \rho$
 2. If $\Delta(s) > \varepsilon$:
 - Run optimal trajectory with π^*
 $\tau = (s_0, a_0, Q_0, s_1, a_1, Q_1, \dots)$
 - **Repeat:** add tuples to dataset (from $i = 0$)
 $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_i, a_i, Q_i)\}$
- Until:** $\Delta(s_i) \ll \varepsilon$.


Else:

Continue




Certifying suboptimality with high probability

Theorem 4.1 (Probabilistic Guarantees). *Assume Algorithm 1 ran for E episodes; let Δ_e be defined as in line 7 of the algorithm. Let $\mathcal{S}_0$ denote the support of the initial state distribution ρ .*




Experiments

- We use the `lqr_n_m` environments from DeepMind's Control Suite




- Results on `lqr_2_1`: $\delta = 0.1$, $p = 0.1$ $\epsilon \approx 10\%$




Experiments on lqr_4_2 (stiff links)

Hyper-parameters: $\delta = 0.1$, $p = 0.1$ $\epsilon \approx 5\%$


After 1 episode...




After 1000 episodes...



c.f. optimal control





Results on 1qr_6_2



Research Goals

- To develop analysis and design methods that *trade off complexity and performance*.
- To allow for *continual improvement*, without the need for redesign, retune, or retrain
- To design control policies with controlled sensitivity to parameter changes



Conclusions and Future work

- **Takeaways**

- Proposed a **relaxed notion of invariance: recurrence**.
- **Nonparametric theory for dynamical systems analysis** leading to:
 - General **Lyapunov** and **Barrier Function** conditions **satisfied by any norm!**
 - Algorithms that are **parallelizable and progressive/sequential**.
- **Nonparametric policies:** Guaranteed improvement with each demonstration.

- **Ongoing work**

- **Recurrence:** Information theoretical lower bounds of control recurrence sets
- **Lyapunov/CBF Theory:** Generalize other Lyapunov notions, Control Lyapunov Functions, Control Barrier Functions, Contraction
- **Nonparametric policies (NP):** NP policy iteration, enforcing safety and stability using NP, exploring alternative inductive biases (beyond Lipschitz)

Thanks!

Related Publications:

[CDC 23] Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions*, **CDC 2023, journal under review**

[HSCC 24] Sibai, M, *Recurrence of nonlinear control systems: Entropy and bit rates*, **HSCC, 2024**

[Allerton 24] Shen, Sibai, M, *Generalized Barrier Functions: Integral conditions and recurrent relaxations*, **Allerton 2024, journal submitted**

[RLC 25] Castellano, Rezaei, Markovitz, and M, *Nonparametric Policy Improvement for Continuous Action Spaces via Expert Demonstrations*, 2025, **submitted to RLC**.

Enrique Mallada
mallada@jhu.edu
<http://mallada.ece.jhu.edu>