Mixed Monotonicity Reachability Analysis of Neural ODE: A Trade-Off Between Tightness and Efficiency

Co-funded by the European Union

Code & Paper

Abdelrahman Sayed Sayed   Pierre-Jean Meyer   Mohamed Ghazel
Univ Gustave Eiffel, COSYS-ESTAS, F-59657 Villeneuve d'Ascq, France
{abdelrahman.ibrahim, pierre-jean.meyer, mohamed.ghazel}@univ-eiffel.fr

Idea

- **Challenges in Neural ODE Verification** – neural ODE model complex dynamical systems effectively, but lack efficient reachability analysis tools.
- Existing Methods/Tools – like CORA (zonotopes) and NNV2.0 (star sets) provide tight over-approximations but at high computational cost.
- **Ours** – Novel interval-based reachability analysis method implemented in TIRA using continuous-time mixed monotonicity, prioritizing efficiency.
- **Core Technique** – Decompose neural ODE vector field into a mixed monotone form; exploit initial set geometry and boundaries via homeomorphism for bound propagation.
- **Trade-Off** – Favor efficiency over tightness, enabling real-time verification.

Method

Given a neural ODE

$$\dot{x}(t) = f(x(t)), \quad (1)$$

We compute an over-approximation $\Omega(\mathcal{X}_{in})$ such that the set of the reachable outputs of the neural ODE:

$$\mathcal{R}_{\text{neural ODE}}(\mathcal{X}_{in}) \subseteq \Omega(\mathcal{X}_{in}).$$

neural ODE are naturally invertible and exhibit homeomorphism, allowing over-approximations from initial set to map boundaries to boundaries, and interiors to interiors.

(a) non-homeomorphism   (b) Homeomorphism

The neural ODE (1) is mixed monotone if there exists a decomposition function $g : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ such that for all $x, \hat{x} \in \mathbb{R}^n$, the following conditions hold:

- g is increasing in its first argument (off-diagonally):

$$\forall i, j \in \{1, \ldots, n\},\ j \neq i: \quad \frac{\partial g_i}{\partial x_j}(x, \hat{x}) \geq 0,$$

- g is decreasing in its second argument:

$$\forall i, j \in \{1, \ldots, n\}: \quad \frac{\partial g_i}{\partial \hat{x}_j}(x, \hat{x}) \leq 0,$$

- f is embedded in the diagonal of g:

$$g(x, x) = f(x),$$

This decomposition implies that the embedded dynamical system is evolving in $\mathbb{R}^{2n_x}$:

$$\begin{bmatrix} \dot{x} \\ \dot{\hat{x}} \end{bmatrix} = \begin{bmatrix} g(x, \hat{x}) \\ g(\hat{x}, x) \end{bmatrix} = h(x, \hat{x}),$$

which is monotone with respect to the orthant $\mathbb{R}^{n_x}_{+} \times \mathbb{R}^{n_x}_{-}$ in its state space.

Results

- For **2D spiral system**, CORA and NNV2.0 achieve tighter over-approximations than our TIRA's approaches, with CORA's dashed-boundaries being the tightest. However, CORA's computational time is approximately **25 times** greater, and NNV2.0 is approximately **6 times** greater than TIRA's single-step mixed monotonicity approach.
- For **5D FPA system**, CORA outperformed both NNV2.0 and TIRA, But, TIRA's single-step mixed monotonicity remains the fastest method, with CORA requiring approximately **131 times** more computation time than TIRA's single-step mixed monotonicity approach.
- TIRA's simple rectangular box over-approximation intervals are easier to compute, resulting in shorter computational times compared to CORA and NNV2.0.

References

[1] Meyer, P.J., Devonport, A., Arcak, M.: TIRA: Toolbox for interval reachability analysis. In: 22nd ACM International Conference on Hybrid Systems: Computation and Control. pp. 224–229 (2019)

[2] Meyer, P.J., Devonport, A., Arcak, M.: Interval reachability analysis: Bounding trajectories of uncertain systems with boxes for control and verification. Springer Nature (2021)

[3] Sayed, A.S., Meyer, P.J., Ghazel, M.: Bridging neural ode and resnet: A formal error bound for safety verification. In: International Symposium on AI Verification. pp. 97–114. Springer (2025)

Overview

Contributions:

- **Novel method**: Reachability analysis of neural ODE using mixed monotonicity techniques for continuous-time dynamical systems.
- **Interval-Based Approach** - Exploit interval analysis for sound, efficient over-approximations of neural ODE reachable set.
- **Implementation & Evaluation** – TIRA toolbox with single-step, incremental, and boundary methods; compared to CORA and NNV2.0 on a 2D spiral and 5D FPA systems.

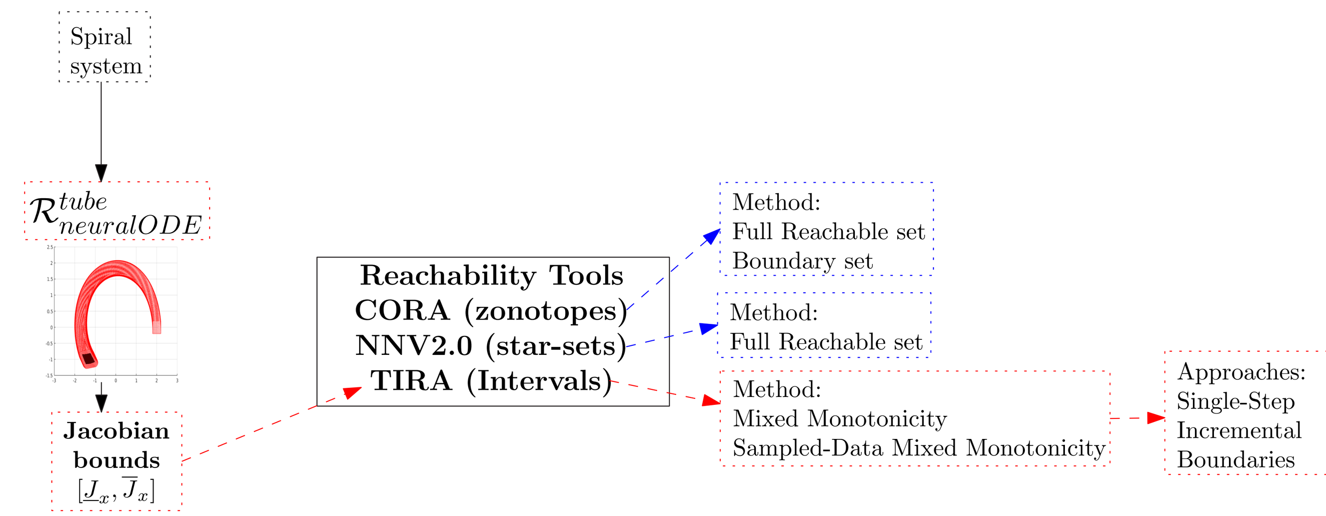

Illustration of the steps for reachability analysis of neural ODE using different tools, methods and mixed monotonicity approaches

Numerical illustration

Comparison with CORA and NNV2.0 Computational times:

Methods	Spiral @ $t = 1$sec.	FPA @ $t = 2$sec.
CORA Full Reachable Set	19.64	13.22
CORA Boundaries only	70.83	109.1
NNV2.0 Full Reachable Set	17.25	11.98
TIRA (single-step) Mixed-Monotonicity	0.66	0.83
TIRA (single-step) Sampled-Data Mixed-Monotonicity	0.95	1.34
TIRA (incremental) Mixed-Monotonicity	63.13	25.41
TIRA (incremental) Sampled-Data Mixed-Monotonicity	111.16	48.06
TIRA (Boundary) Mixed-Monotonicity	2.84	7.06
TIRA (Boundary) Sampled-Data Mixed-Monotonicity	4.35	12.76

Comparison with CORA and NNV2.0 over-approximations Tightness:

Methods	Spiral	FPA		
	$x_1 - x_2$	$x_1 - x_2$	$x_3 - x_4$	$x_4 - x_5$
CORA Full Reachable Set	1.61	1.33	1.11	1.13
CORA Boundaries only	1.15	1.18	0.99	1.08
NNV2.0 Full Reachable Set	1.71	2.52	8.74	2.43
TIRA (single-step) Mixed-Monotonicity	24.59	2.29	2.30	1.79
TIRA (single-step) Sampled-Data Mixed-Monotonicity	12.14	33.57	40.67	8.05
TIRA (incremental) Mixed-Monotonicity	24.59	2.29	2.30	1.79
TIRA (incremental) Sampled-Data Mixed-Monotonicity	23.24	18.92	43.64	5.50
TIRA (Boundary) Mixed-Monotonicity	12.05	2.29	2.30	1.79
TIRA (Boundary) Sampled-Data Mixed-Monotonicity	12.14	33.57	40.67	8.05

Visual result for the 2D spiral system

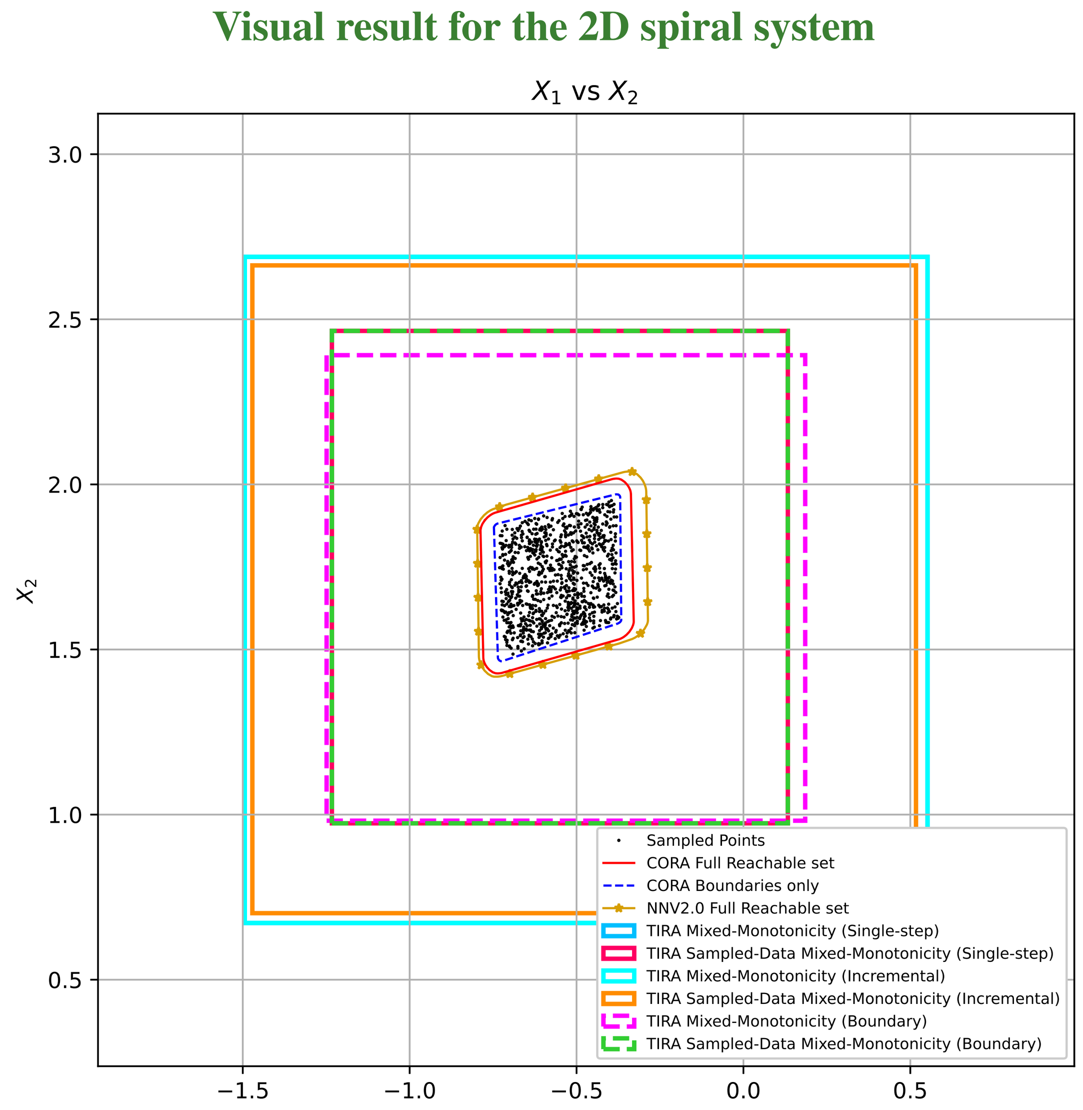

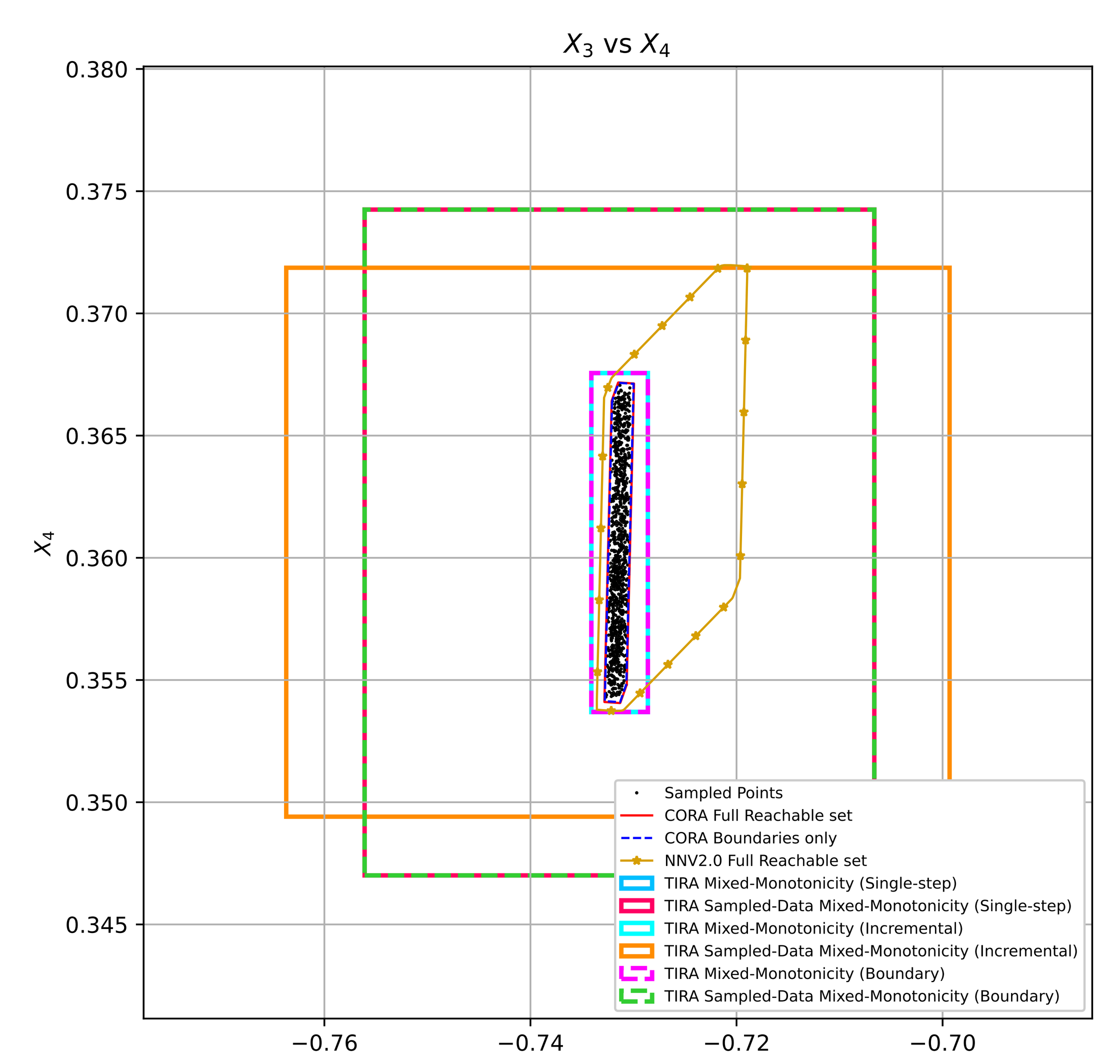

Visual results for the 5D FPA system