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

Abdelrahman Sayed Sayed Pierre-Jean Meyer Mohamed Ghazel

Université Gustave Eiffel, COSYS-ESTAS, F-59657 Villeneuve d'Ascq, France

December 7^{th} 2025, San Diego, California

NEURAL INFORMATION PROCESSING SYSTEMS

Université Gustave Eiffel

Co-funded by the European Union

NeurReps 2025

Motivation

Neural ODE Verification Challenges

- Few verification tools
- Computationally intensive

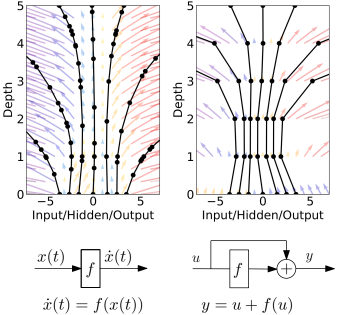

Motivation

Neural ODE Verification Challenges

- Few verification tools
- Computationally intensive

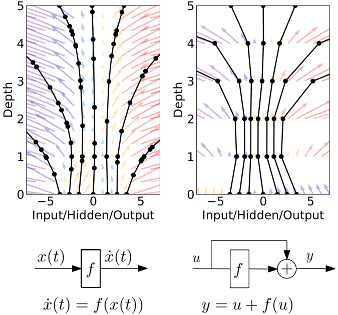

Proposed Approaches

Spiral system

$\mathcal{R}^{tube}_{neuralODE}$

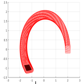

Jacobian bounds $[\underline{J}_x, \overline{J}_x]$

Reachability Tools
CORA (zonotopes)
NNV2.0 (star-sets)
TIRA (Intervals)

Method: Full Reachable set, Boundary set

Method: Full Reachable set

Method: Mixed Monotonicity, Sampled-Data Mixed Monotonicity

Approaches: Single-Step, Incremental, Boundaries

Neural ODE Reachability

- We consider the following neural ODE:

$$\dot{x}(t) = \frac{dx(t)}{dt} = f(x(t)) \tag{1}$$

Definition (neural ODE Reachability)

Given an initial input set: $\mathcal{X}_{in} \subseteq \mathbb{R}^n$ and final time t_f, we define the set of neural ODE reachable outputs as:

$$\mathcal{R}_{\text{neural ODE}}(\mathcal{X}_{in}) = \{y \in \mathbb{R}^n \mid y = \Phi(t_f, u),\ u \in \mathcal{X}_{in}\}$$

- Where Φ corresponds to solution of (1) based on IVP:

$$x(t_f) = \Phi(t_f, x(0)) = \Phi(t_f, u)$$

Homeomorphism

- Homeomorphism preserves topological characteristics

Definition (Homeomorphism)

For two sets $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$, there exists a map $m(.) : \mathcal{X} \to \mathcal{Y}$ which is a homeomorphism w.r.t. $\mathcal{X}$ if it is a continuous bijection and the map inverse $m^{-1}(.) : \mathcal{Y} \to \mathcal{X}$ is also continuous.

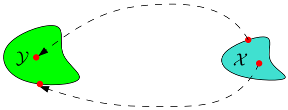

Non-homeomorphic

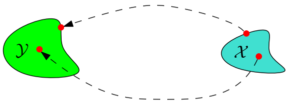

Homeomorphic

Homeomorphism(Cont.)

Lemma

Assuming that the two sets $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$ are closed and bounded. For a homeomorphism map $m(.) : \mathcal{X} \rightarrow \mathcal{Y}$, m maps the boundaries of the set $\mathcal{X}$ to the boundaries of the set $\mathcal{Y}$, and the interior of the set $\mathcal{X}$ to the interior of the set $\mathcal{Y}$.

- Since neural ODE are naturally invertible $\Rightarrow$ Homeomorphic

Continuous time Mixed Monotonicity

Continuous-time neural ODE (1): $\dot{x} = f(x(t))$

Definition (Neural ODE Mixed monotonicity)

neural ODE (1) is mixed-monotone is there exists a decomposition $g : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ such that:

- g is increasing in the first argument: $g(x, \hat{x})$
- g is decreasing in the second argument: $g(x, \hat{x})$
- f is embedded in the diagonal of g: $g(x, x) = f(x)$

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

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

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

Continuous time Mixed Monotonicity (cont.)

Requirements: The only requirement for applicability is to compute the neural ODE Jacobian bounds

$$J(x) = \frac{\partial f}{\partial x}(x)$$

Continuous time Mixed Monotonicity (cont.)

Requirements: The only requirement for applicability is to compute the neural ODE Jacobian bounds

$$J(x) = \frac{\partial f}{\partial x}(x)$$

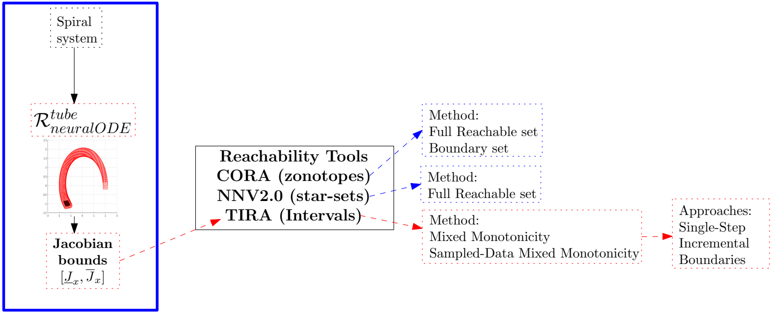

Numerical Examples

Spiral Non-linear System

- 2 dimensional neural ODE
- $\dot{x} = f(x) = W_2 \text{tanh}(W_1 x + b_1) + b_2$,
  where:
  - $\rightarrow$ W_1 and W_2 are the weight matrices
  - $\rightarrow$ b_1 and b_2 are the bias vectors
  - $\rightarrow$ $tanh$: hyperbolic tangent AF

Fixed-Point Attractor System (FPA)

- 5 dimensional neural ODE
- $\dot{x} = f(x) = \tau x + W \text{tanh}(x)$,
  where:
  - $\rightarrow$ $\tau = -10^{-6}$

Single-Step Reachability analysis

- Compute $\mathcal{R}_{\text{neural ODE}}$ from the initial time to final time t_f
- neural ODE embedded system is solved over the full time horizon $[0, t_f]$, with $t_f = 1$ for Spiral system and $t_f = 2$ for FPA system

Incremental Reachability analysis

- Refines the single-step approach, by applying the mixed monotonicity embedding and propagating the bounds sequentially
- The $\mathcal{R}_{\text{neural ODE}}$ output of one step is used as the input for the next step
- This approach can yield tighter over-approximations than single-step, but with a higher computational cost due to repeated numerical integrations of the embedded monotone system

Boundary Reachability analysis

- Compute $\mathcal{R}_{\text{neural ODE}}$ from the boundary of $\mathcal{X}_{in}$ rather than the entire input set
- This approach offers computational efficiency, as it scales linearly with the state dimension

Boundary Reachability analysis

- Compute $\mathcal{R}_{\text{neural ODE}}$ from the boundary of $\mathcal{X}_{in}$ rather than the entire input set
- This approach offers computational efficiency, as it scales linearly with the state dimension

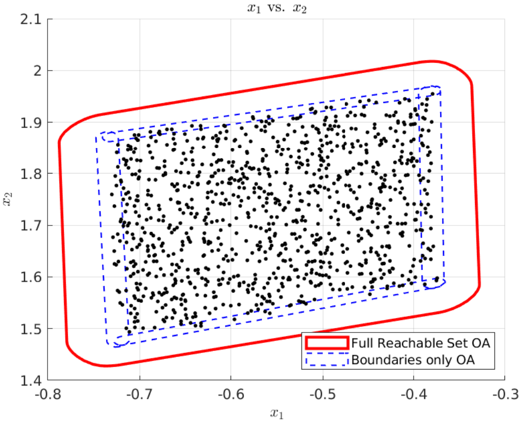

Boundary Reachability analysis

- Compute $\mathcal{R}_{\text{neural ODE}}$ from the boundary of $\mathcal{X}_{in}$ rather than the entire input set
- This approach offers computational efficiency, as it scales linearly with the state dimension

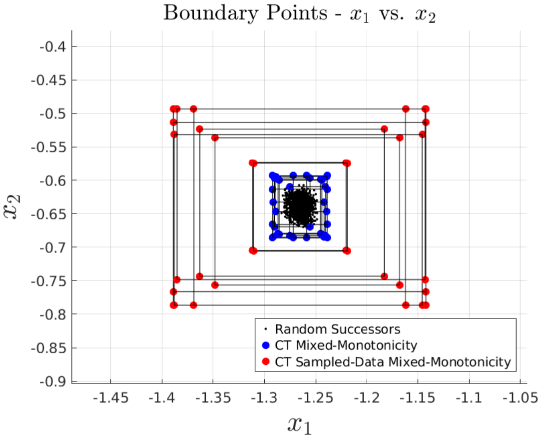

Spiral Comparison over Tools and Approaches

- Similar TIRA over-approximations: single-step and incremental MM & single-step and dashed boundary-based SDMM

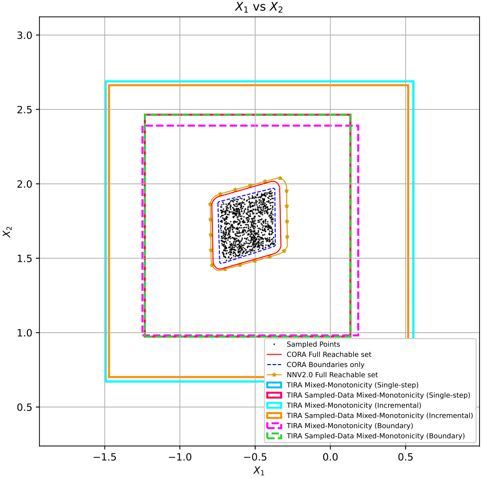

Spiral Comparison over Tools and Approaches (cont.)

- CORA zonotopes and NNV2.0 star-set achieved tighter over-approximations than TIRA's Interval approaches

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

Spiral Comparison over Tools and Approaches (cont.)

- TIRA's single-step MM is **25 times** faster than CORA and **6 times** faster than NNV2.0

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

FPA Comparison over Tools and Approaches

- Similar TIRA over-approximations: single-step, incremental and dashed-boundary-based MM & single-step and dashed boundary-based SDMM

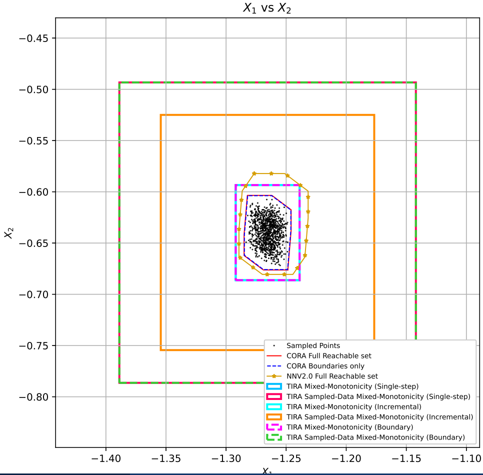

FPA Comparison over Tools and Approaches

- Similar TIRA over-approximations: single-step, incremental and dashed-boundary-based MM & single-step and dashed boundary-based SDMM

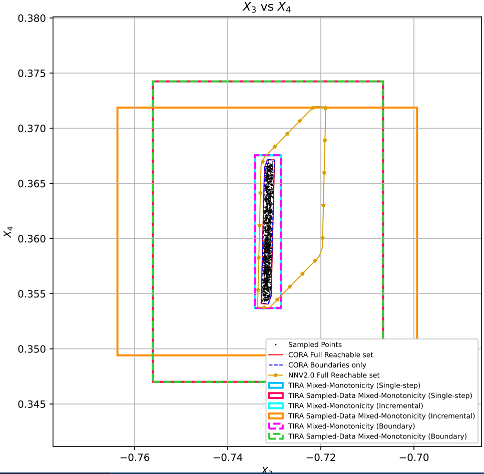

FPA Comparison over Tools and Approaches

- Similar TIRA over-approximations: single-step, incremental and dashed-boundary-based MM & single-step and dashed boundary-based SDMM

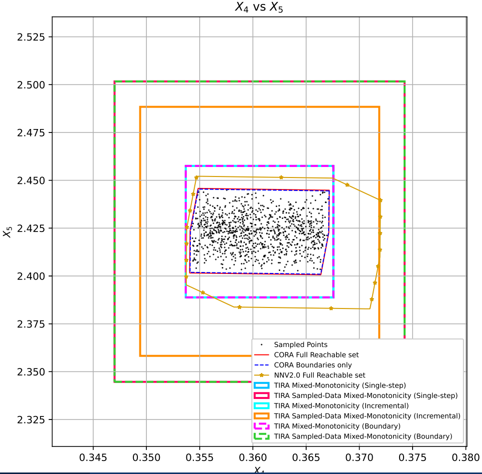

FPA Comparison over Tools and Approaches (cont.)

- CORA zonotopes achieved tighter over-approximations than TIRA's Interval approaches

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

FPA Comparison over Tools and Approaches (cont.)

- TIRA's single-step MM is **131 times** faster than CORA and **14 times** faster than NNV2.0

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

Conclusions

Interval-based reachability method for neural ODE

- Lightweight neural ODE reachability analysis alternative
- Sound over-approximations, albeit at the cost of tightness

Future Work

- Extend boundary-based reachability approach to include incremental method
- Partitioning the initial input set into smaller subsets
- Incorporate the framework into a verifier to check safety properties in neural ODE

Contact: abdelrahman.ibrahim@univ-eiffel.fr