Neural Continuous-Time Supermartingale Certificates

Greg (Grigory) Neustroev[1,2]   Mirco Giacobbe[1,3]   Anna Lukina[2]

[1]Zeroth Research, UK   [2]TU Delft, NL   [3]University of Birmingham, UK

AIPV, May 19, 2026

Zeroth Research

Neural Continuous-Time Supermartingale Certificates

Greg Neustroev

Zeroth Research TU Delft

Mirco Giacobbe

Zeroth Research UNIVERSITY OF BIRMINGHAM

Anna Lukina

TU Delft

Zeroth Research

What Is This Talk About?

Can we certify safety and/or persistence (≈ stability) of a dynamical system that is:

- *neural-controlled*,
- *continuous-time*, and
- *stochastic*.

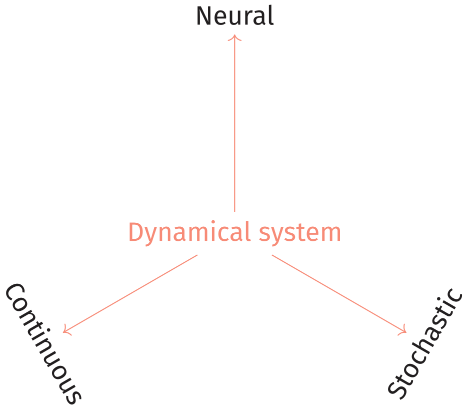

Zeroth Research

What Is This Talk About?

Can we certify safety and/or persistence (≈ stability) of a dynamical system that is:

- *neural-controlled*,
- *continuous-time*, and
- *stochastic*.

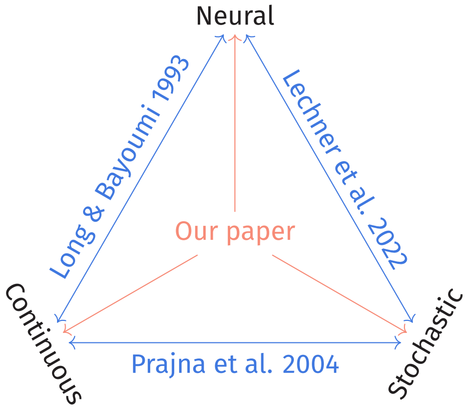

Zeroth Research

What Is This Talk About?

Can we certify safety and/or persistence (≈ stability) of a dynamical system that is:

- *neural-controlled*,
- *continuous-time*, and
- *stochastic*.

As a bonus, we also allow non-stationarity!

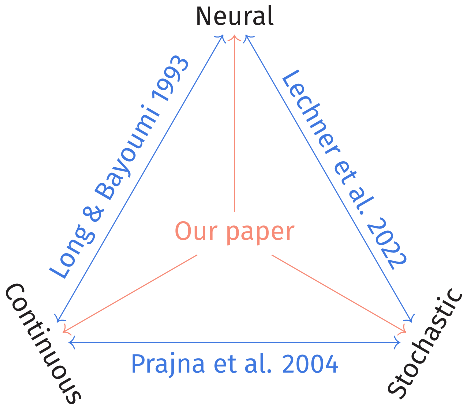

What Exactly Do We Certify?

For any state from the subspace of initial states $\mathbb{X}_0$:

- *Reach-avoidance* — the system is able to reach the subspace of target states $\mathbb{X}_\star$ without entering the unsafe states $\mathbb{X}_\oslash$;
- *Staying* — at some point the system stays in the target states $\mathbb{X}_\star$.

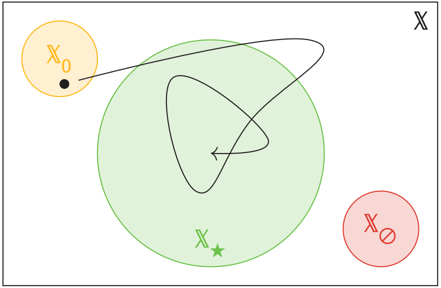

We are interested in formal guarantees of either property (or both).

Zeroth Research

Deterministic Example

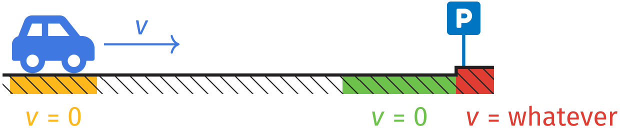

Zeroth Research

Deterministic Example

Control: $a = \pi(p, v)$

System's Dynamics: $\dot{p} = v$

$\dot{v} = a$

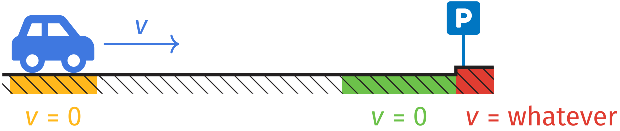

Zeroth Research

Deterministic Example

Control: $a = \pi(p, v)$

System's Dynamics: $\dot{p} = v$

$\dot{v} = a$

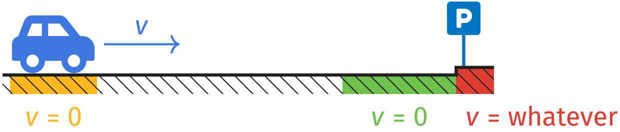

General formulation: $u = \pi(x)$ $\dot{x} = f(x, u)$

Zeroth Research

Deterministic Certificates

How do we certify safety and stability for a deterministic dynamical system?

$$u = \pi(x) \qquad \text{and} \qquad \dot{x} = f(x, u) \qquad \rightarrow \qquad \dot{x} = f\big(x, \pi(x)\big) = f_\pi(x)$$

Zeroth Research

Deterministic Certificates

How do we certify safety and stability for a deterministic dynamical system?

$$u = \pi(x) \qquad \text{and} \qquad \dot{x} = f(x, u) \qquad \rightarrow \qquad \dot{x} = f\big(x, \pi(x)\big) = f_\pi(x)$$

Stability

Lyapunov functions V are stability certificates; they satisfy:

- $V(0) = 0$;
- $V(x) > 0$ for $x \neq 0$;
- $\nabla V \cdot f_\pi(x) < 0$ for $x \neq 0$.

They certify that the system converges to the equilibrium $x = 0$.

Zeroth Research

Deterministic Certificates

How do we certify safety and stability for a deterministic dynamical system?

$$u = \pi(x) \quad \text{and} \quad \dot{x} = f(x, u) \quad \rightarrow \quad \dot{x} = f\big(x, \pi(x)\big) = f_\pi(x)$$

Stability

Lyapunov functions V are stability certificates; they satisfy:

- $V(0) = 0$;
- $V(x) > 0$ for $x \neq 0$;
- $\nabla V \cdot f_\pi(x) < 0$ for $x \neq 0$.

They certify that the system converges to the equilibrium $x = 0$.

Safety

Control barrier functions V are safety certificates; they satisfy:

- $V(x) \leq \alpha$ for $x \in \mathbb{X}_0$;
- $V(x) \geq \beta$ for $x \in \mathbb{X}_\oslash$, where $\alpha < \beta$;
- $\nabla V \cdot f_\pi(x) \leq -\epsilon$ for $x \in \mathbb{X} \setminus \mathbb{X}_\star$ s.t. $V(x) \leq \beta$.

They certify that the system reaches $\mathbb{X}_\star$ from $\mathbb{X}_0$ without entering $\mathbb{X}_\oslash$.

Zeroth Research

Intuition Behind the Certificates

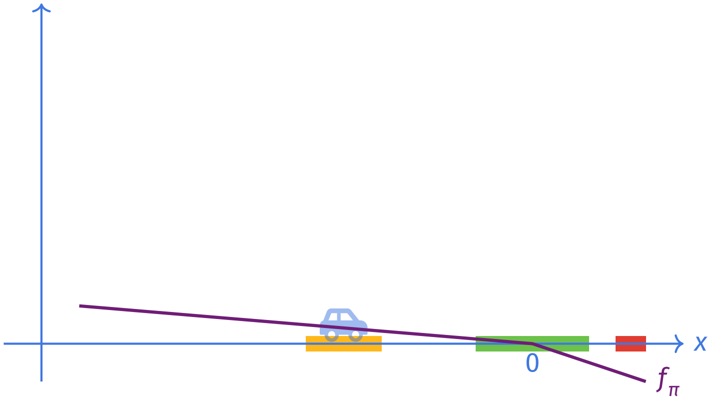

Intuition Behind the Certificates

Stability

- $V(0) = 0$
- $V(x) > 0,\ x \neq 0$
- $\nabla V \cdot f_\pi(x) < 0,\ x \neq 0$

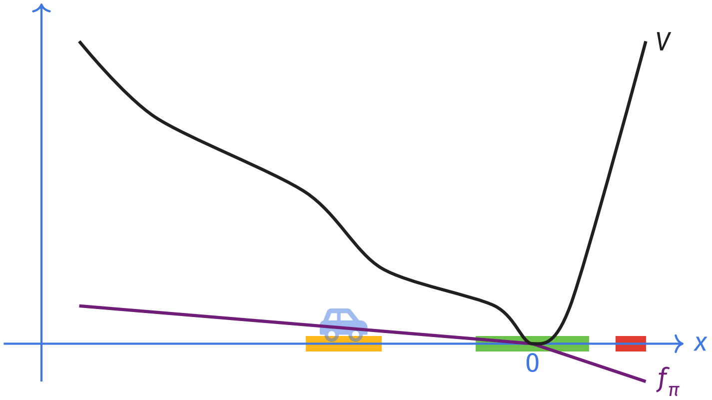

Zeroth Research

Intuition Behind the Certificates

Safety

- $V(x) \leq \alpha,\ x \in \mathbb{X}_0$

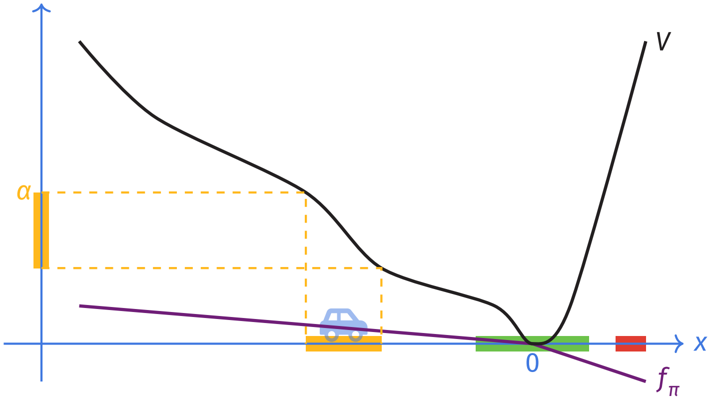

Zeroth Research

Intuition Behind the Certificates

Safety

- $V(x) \leq \alpha,\ x \in \mathbb{X}_0$
- $V(x) \geq \beta,\ x \in \mathbb{X}_{\oslash}$

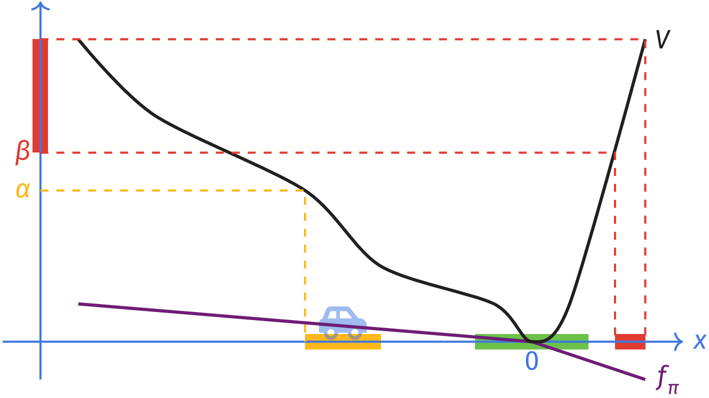

Zeroth Research

Intuition Behind the Certificates

Safety

- $V(x) \le \alpha$, $x \in \mathbb{X}_0$
- $V(x) \ge \beta$, $x \in \mathbb{X}_\oslash$
- $\nabla V \cdot f_\pi(x) \le -\epsilon$, $x \in \mathbb{X} \setminus \mathbb{X}_\star$ and $V(x) \le \alpha$

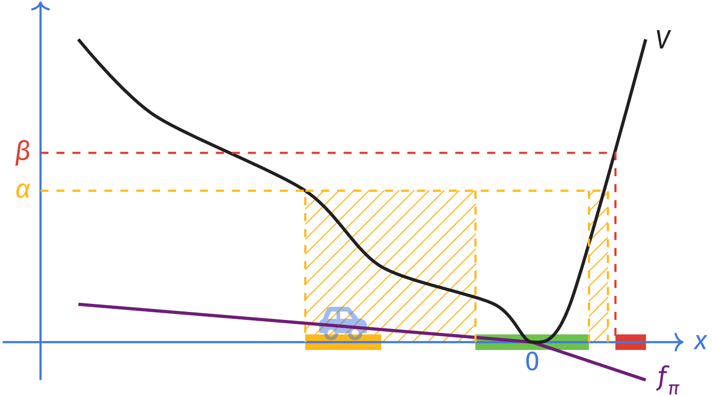

Zeroth Research

Stochastic Differential Equation

$$\mathrm{d}X_t = \underbrace{f(t, X_t)}_{\text{drift}}\,\mathrm{d}t + \underbrace{g(t, X_t)}_{\text{diffusion}}\,\mathrm{d}W_t$$

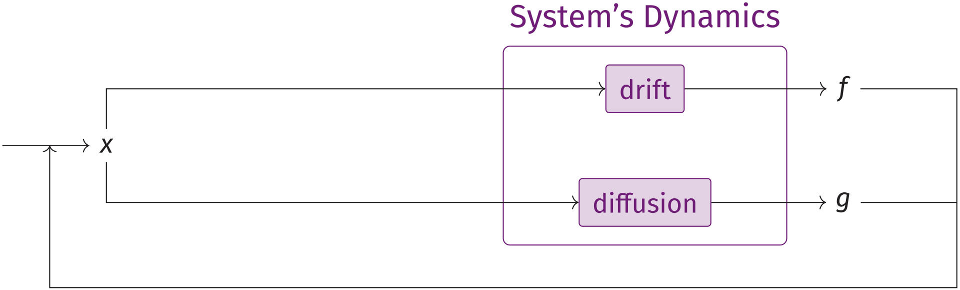

Zeroth Research

Neural Stochastic Differential Equation

$$\underbrace{U_t = \pi(X_t)}_{\text{policy}} \qquad dX_t = \underbrace{f(t, X_t, U_t)}_{\text{drift}}\, dt + \underbrace{g(t, X_t, U_t)}_{\text{diffusion}}\, dW_t$$

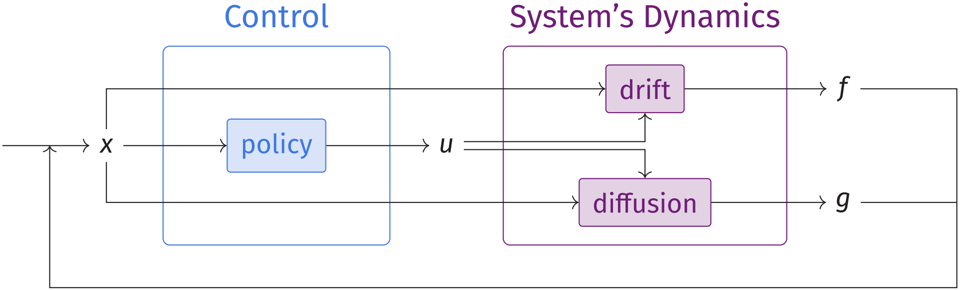

Zeroth Research

Reach-Avoid Specification

For any initial state $x_0 \in \mathbb{X}_0$, certify:

- *reach-avoidance* with probability at least ε.

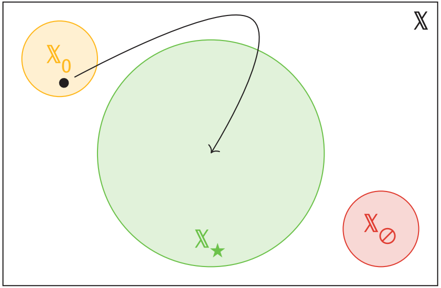

Zeroth Research

Reach-Avoid Specification

For any initial state $x_0 \in \mathbb{X}_0$, certify:

- *reach-avoidance* with probability at least ε.

We prove that we can find a certificate and set barriers so that

$$\varepsilon = 1 - \frac{\alpha}{\beta}$$

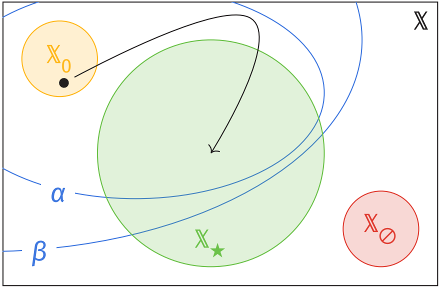

Zeroth Research

Reach-Avoid Certificate

1. $V(t, x) \geq 0$ for $x \in \mathbb{X}$;

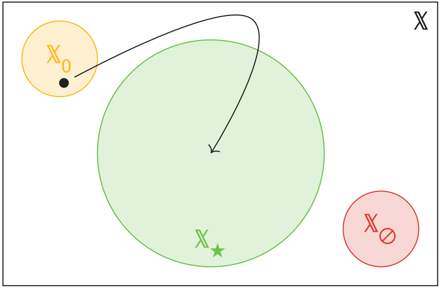

Zeroth Research

Reach-Avoid Certificate

1. $V(t, x) \geq 0$ for $x \in \mathbb{X}$;
2. $V(t, x_0) \leq \alpha$ for $x_0 \in \mathbb{X}_0$;

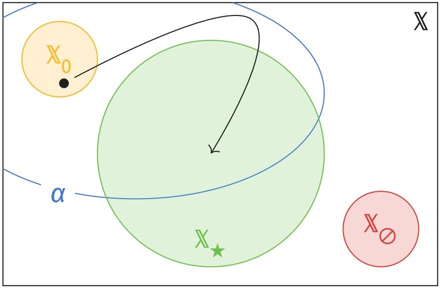

Zeroth Research

Reach-Avoid Certificate

1. $V(t, x) \geq 0$ for $x \in \mathbb{X}$;
2. $V(t, x_0) \leq \alpha$ for $x_0 \in \mathbb{X}_0$;
3. $V(t, x) \geq \beta$ for $x \in \mathbb{X}_{\oslash}$;

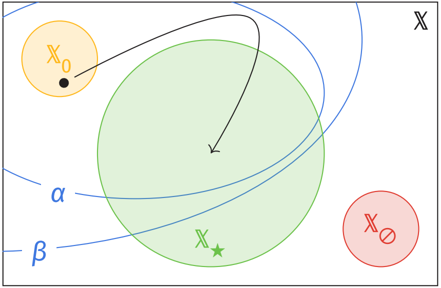

Zeroth Research

Reach-Avoid Certificate

1. $V(t, x) \geq 0$ for $x \in \mathbb{X}$;
2. $V(t, x_0) \leq \alpha$ for $x_0 \in \mathbb{X}_0$;
3. $V(t, x) \geq \beta$ for $x \in \mathbb{X}_\oslash$;
4. $G_\pi V(t, x) < 0$ for $x \notin \text{int}\, \mathbb{X}_\star$ such that $V(t, x) \leq \beta$.

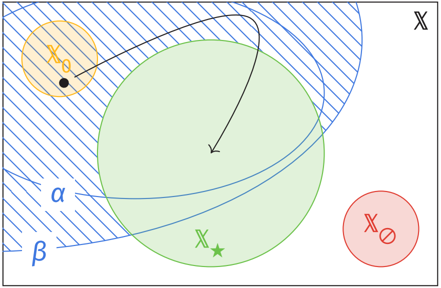

G_π is the infinitesimal generator (≈ stochastic gradient). It shows the expected change of a function along the state trajectory process.

How Do We Find This Certificate?

We train it using a neural network using the following loss:

$$L = L_{\oslash} + L_0 + L_{\star} + L_{\downarrow} + R$$

$$L_{\oslash} = \sum_{x \in B \cap \mathbb{X}_{\oslash}} \left(\beta - V_\theta(x)\right)^+$$

$$L_0 = \sum_{x_0 \in B \cap \mathbb{X}_0} \left(V_\theta(x_0) - \alpha\right)^+$$

$$L_{\star} = \sum_{x \in B \cap \mathbb{X}_{\star}} \left(\beta_S - V_\theta(x)\right)^+$$

$$L_{\downarrow} = \sum_{x \in B \cap (L_{\bar{\beta}} \setminus L_{\bar{\alpha}_S})} \left(\mathrm{G}_\pi V_\theta(x) + \zeta\right)^+$$

$$R = \lambda \prod_{i=0}^{N-1} \|\mathbf{W}^{(i)}\|_{\infty \to \infty}$$

Zeroth Research

The Certificate's Correctness

We use interval bound propagation to verify that the certificate is correct.

The Certificate's Correctness

We use interval bound propagation to verify that the certificate is correct.

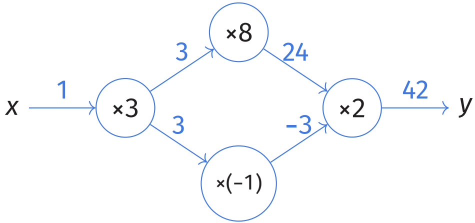

The certificate network

The Certificate's Correctness

We use interval bound propagation to verify that the certificate is correct.

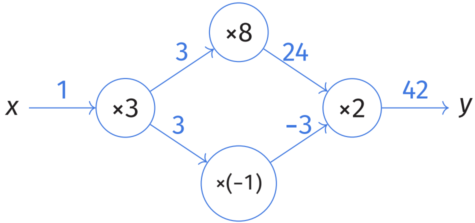

The certificate network

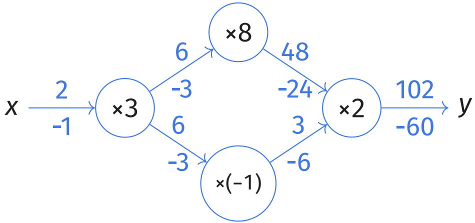

The bounded certificate network

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

$\alpha = 10$

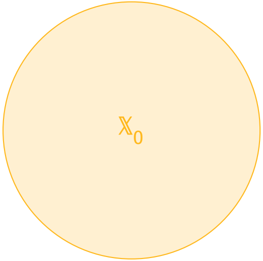

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

$\alpha = 10$

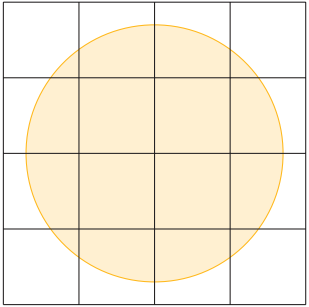

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.

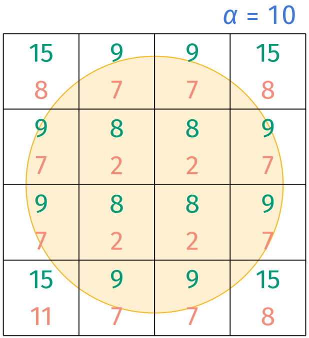

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.
3. A cell is falsified if the lower bound is above the threshold. The considered property does not hold.

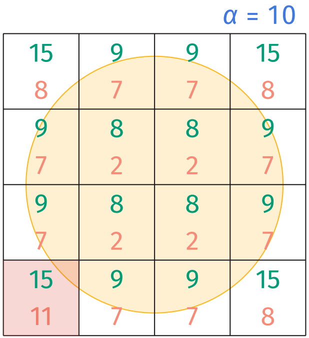

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.
3. A cell is falsified if the lower bound is above the threshold. The considered property does not hold.
4. A cell is verified if the upper bound is below the threshold. The property holds in the cell.

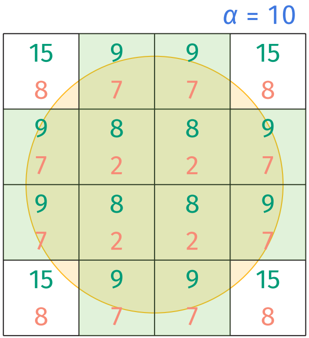

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.
3. A cell is falsified if the lower bound is above the threshold. The considered property does not hold.
4. A cell is verified if the upper bound is below the threshold. The property holds in the cell.
5. If there are unverified cells, split them and repeat.

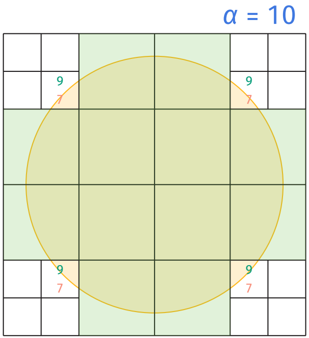

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.
3. A cell is falsified if the lower bound is above the threshold. The considered property does not hold.
4. A cell is verified if the upper bound is below the threshold. The property holds in the cell.
5. If there are unverified cells, split them and repeat.

$\alpha = 10$

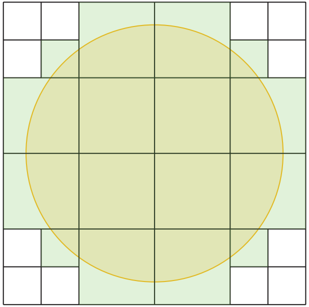

Zeroth Research

Cell-Splitting for Verification

Verify that $V \leq \alpha$ for all $x_0 \in \mathbb{X}_0$

1. Cover $\mathbb{X}_0$ with cells (axis-aligned bounding boxes).
2. Find the lower and upper bound on V at each cell.
3. A cell is falsified if the lower bound is above the threshold. The considered property does not hold.
4. A cell is verified if the upper bound is below the threshold. The property holds in the cell.
5. If there are unverified cells, split them and repeat.

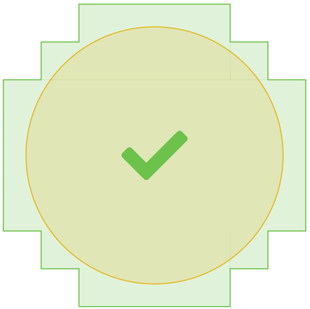

Zeroth Research

Everything Together All at Once

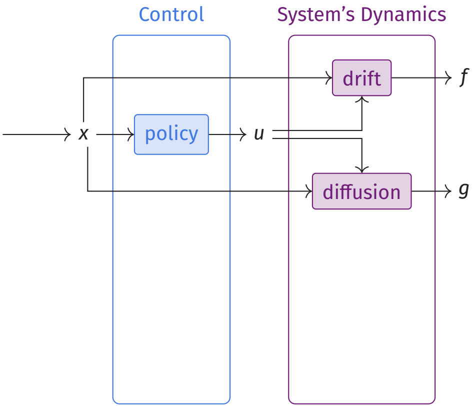

Everything Together All at Once

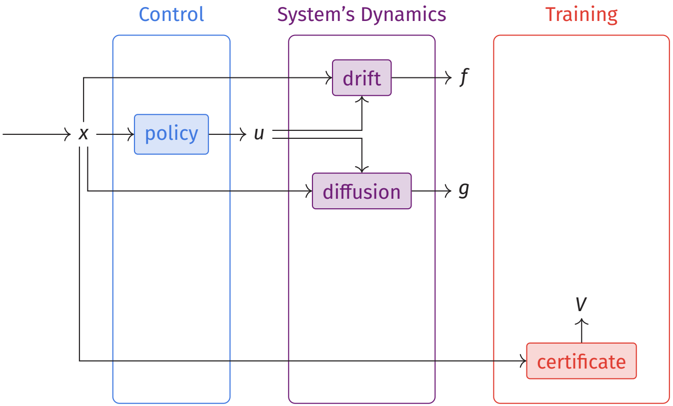

Zeroth Research

Everything Together All at Once

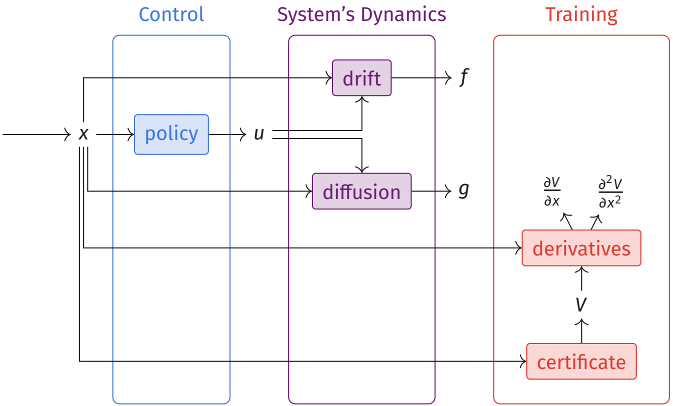

Zeroth Research

Everything Together All at Once

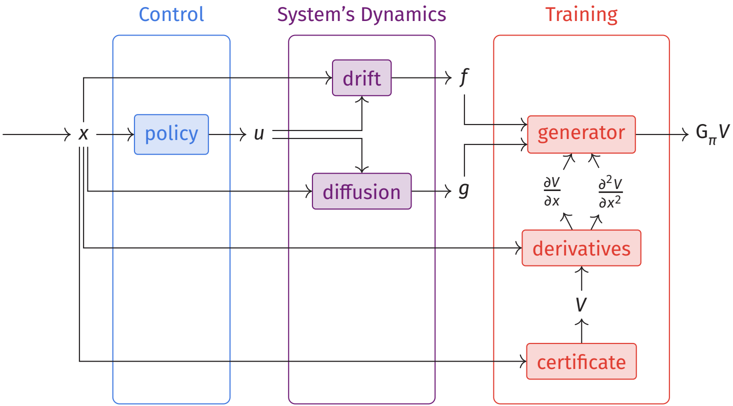

Zeroth Research

Everything Together All at Once

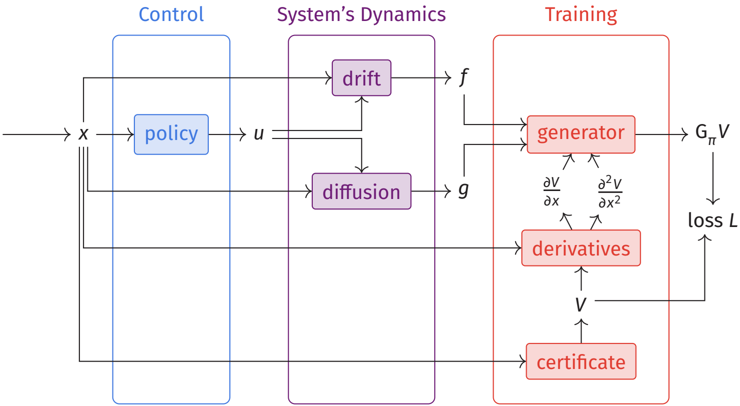

Zeroth Research

Everything Together All at Once

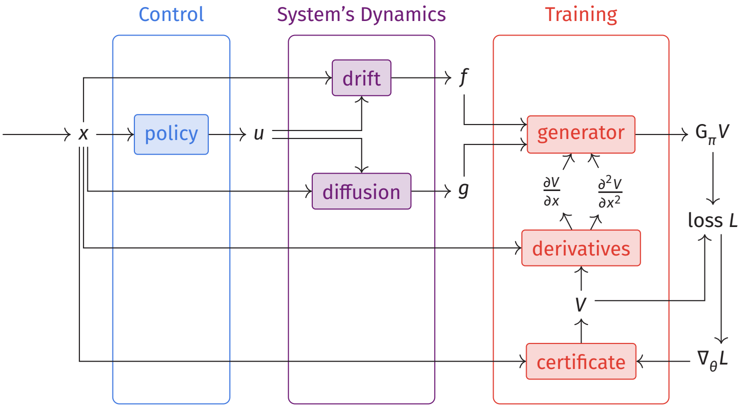

Zeroth Research

Everything Together All at Once

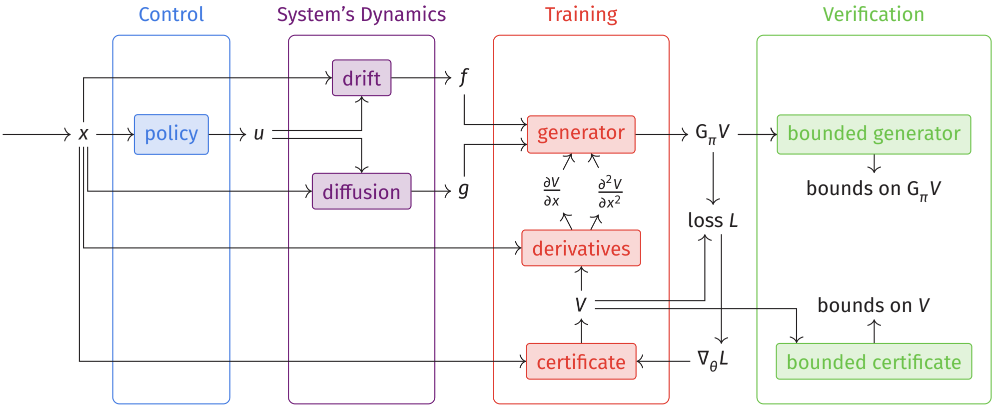

Zeroth Research

Everything Together All at Once

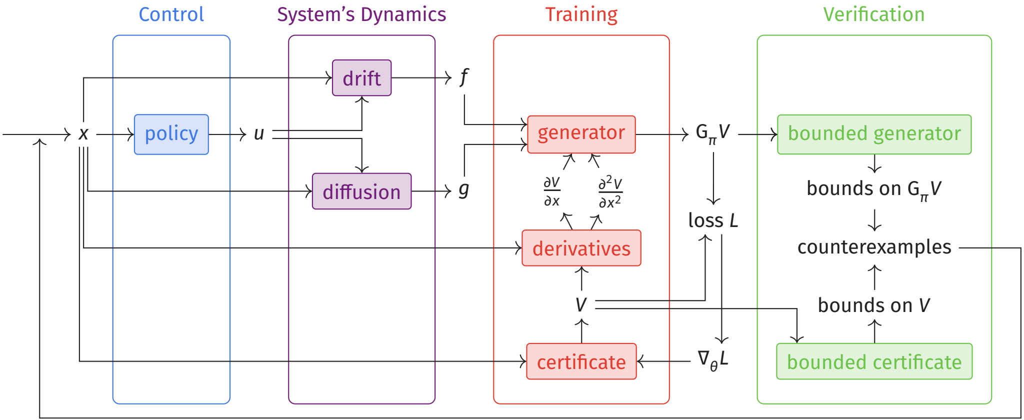

Zeroth Research

Example — Inverted Pendulum

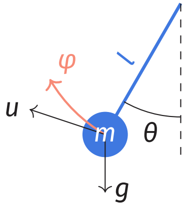

$$\ddot{\theta} = \frac{g}{l}\sin\theta + \frac{Mu - b\dot{\theta}}{ml^2}$$

Zeroth Research

Example — Inverted Pendulum

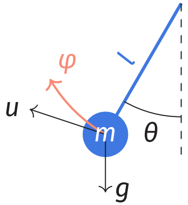

$$d\varphi_t = \left(\frac{g}{l}\sin\theta_t + \frac{Mu_t - b\varphi_t}{ml^2}\right)dt + \sigma\, dW_t$$

$$d\theta_t = \varphi_t\, dt$$

Zeroth Research

Example — Inverted Pendulum

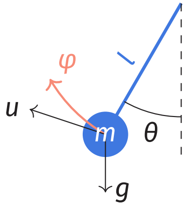

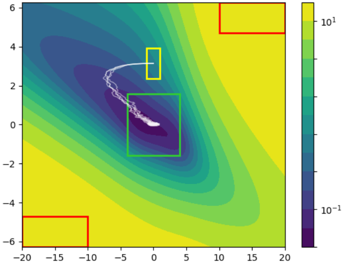

$$d\varphi_t = \left(\frac{g}{l}\sin\theta_t + \frac{Mu_t - b\varphi_t}{ml^2}\right)dt + \sigma\, dW_t$$

$$d\theta_t = \varphi_t\, dt$$

Zeroth Research

Thank you for your attention!

Zeroth Research