Neural Continuous-Time Supermartingale Certificates

Grigory Neustroev, Mirco Giacobbe, Anna Lukina
greg.neustroev@zeroth.org, m.giacobbe@bham.ac.uk, a.lukina@tudelft.nl

1 What Is This Paper About?

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

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

As a bonus, we also allow non-stationarity!

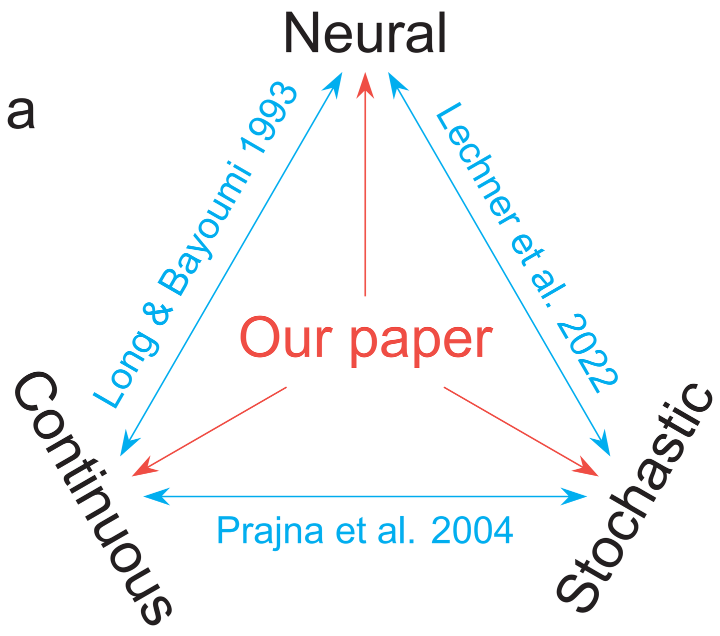

2 How Is the Dynamical System Defined?

The system is described by a stochastic differential equation. Colored functions can be neural networks.

Control: $U_t = \underbrace{\pi(X_t)}_{\text{policy}}$

System's Dynamics: $dX_t = \underbrace{f(t, X_t, U_t)}_{\text{drift}}\, dt + \underbrace{g(t, X_t, U_t)}_{\text{diffusion}}\, dW_t$

3 What Do We Certify?

For any initial state in $\mathbb{X}_0$, with given high probabilities:

Reach-avoidance — the process $(X_t)_{t \geq 0}$ reaches the target $\mathbb{X}_\star$ without becoming unsafe ($\mathbb{X}_\oslash$);

Staying — at some point the system stays in $\mathbb{X}_\star$.

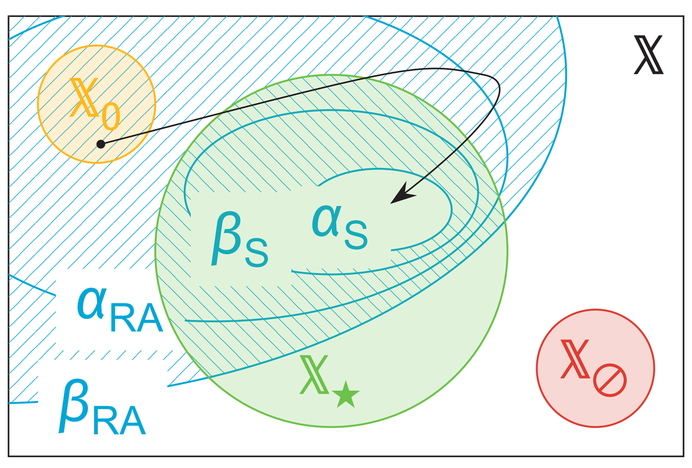

4 How Do We Certify These Properties?

We find a *reach-avoid-stay certificate* $V(t, x)$ such that:

- $V(t, x) \geq 0$ for all states $x \in \mathbb{X}$;
- $V(t, x_0) \leq \alpha_{RA}$ for initial states $x_0 \in \mathbb{X}_0$;
- $V(t, x) \geq \beta_{RA}$ for unsafe states $x \in \mathbb{X}_\oslash$;
- $G_\pi V(t, x) < 0$ for $x \notin \operatorname{int} \mathbb{X}_\star$ such that $V(t, x) \leq \beta_{RA}$;
- Inside $\mathbb{X}_\star$ there exists a subspace where $V(t, x) \leq \beta_S$;
- $G_\pi V(t, x) < 0$ for $x \in \mathbb{X}_\star$ such that $\alpha_S \leq V(t, x) \leq \beta_{RA}$.

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

5 How Do We Find the Certificate?

We train a neural network by minimizing the total violation of the conditions defining the certificate.

6 How Do We Verify the Certificate's Correctness?

After training, we ensure the conditions hold everywhere in the required subspaces. We use interval bound propagation from `autoLiPRA` to do so.

7 How Are the Neural Networks Connected?

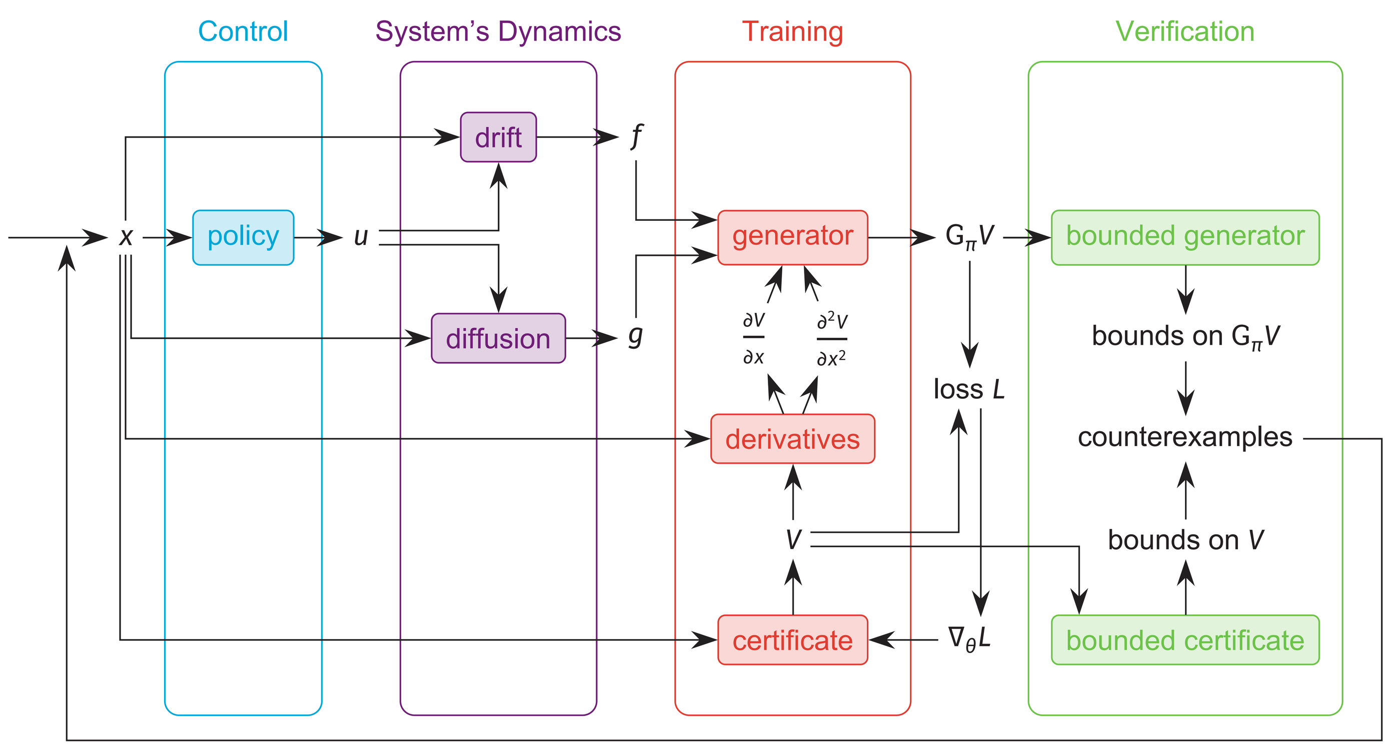

8 Does It Work in Practice?

This is an example of a certificate (right) learned for the neural-controlled inverted pendulum problem (left):

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

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

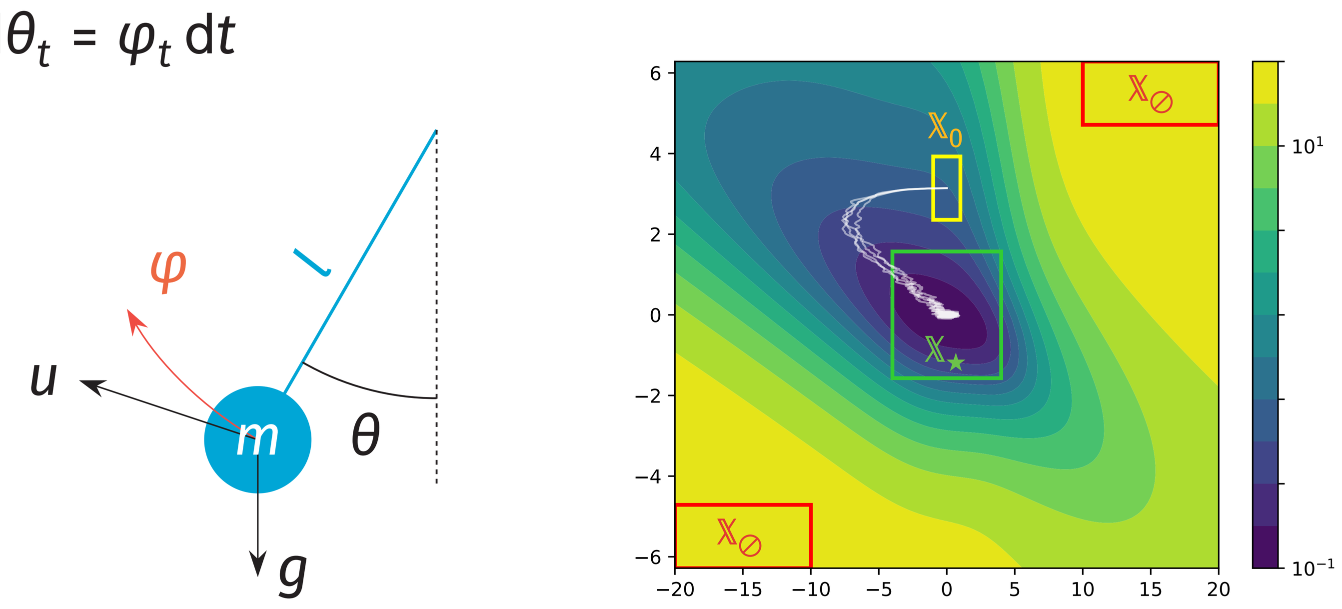

9 How Do I Learn More?

Use the QR codes to access the paper (left) and code (right).

TU Delft
University of Birmingham
SUMI Lab
ellis
NWO