Jye Sawtell-Rickson
m1361019@cgu.edu.tw

Chang Gung University

April 30, 2026

Paper Review: Tolerance of RL Controllers in CPS

Zhang, Kapoor et al.

Outline

m13610190@cgu.edu.tw

Issues

- CPS with RL controllers face challenges in real-world deployment due to "reality gaps."
- Training methods like robust RL lack formal guarantees of tolerance.
- No existing frameworks for post-training analysis of tolerance under deviations.

m1361019@cgu.edu.tw

Contribution

- Formal definition of tolerance using Signal Temporal Logic (STL).
- Introduced the tolerance falsification problem.
- Proposed a two-layer analysis framework for finding tolerance violations.

m1361019@cgu.edu.tw

System Model

- CPS modeled as Markov Decision Processes (MDPs) with RL-based controllers.
- STL used to describe system requirements and deviations.
- Deviations modeled as parameter changes in system dynamics.

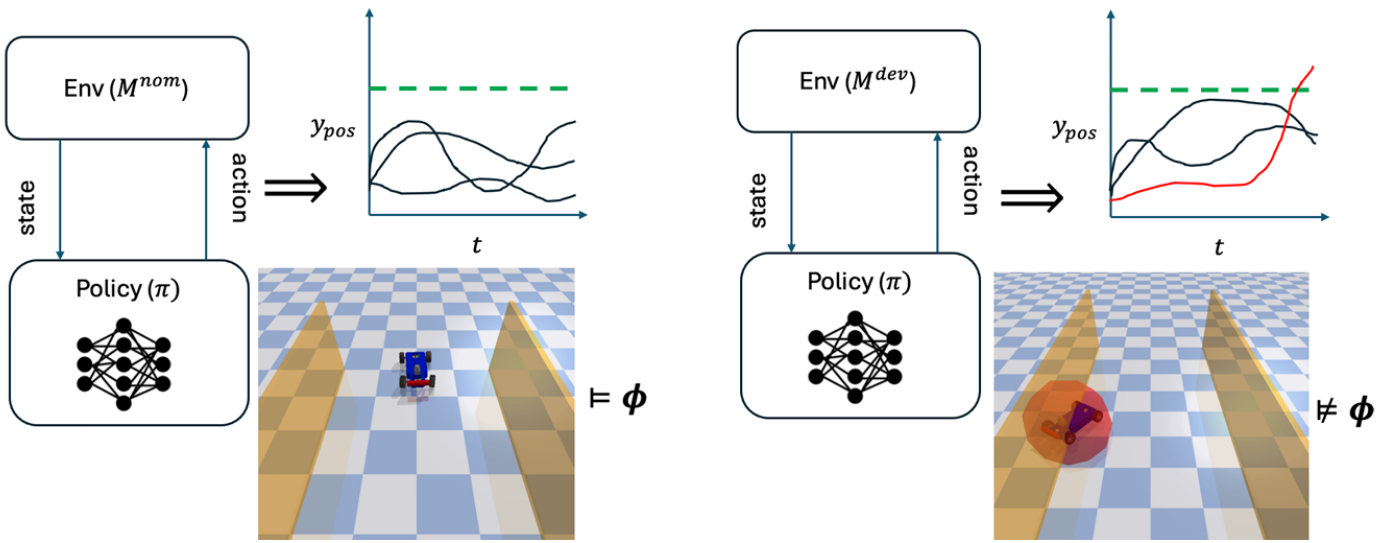

Figure: STL is used to describe system requirements and deviations in the CarRun system.

m1361019@cgu.edu.tw

Proposed Problem

- **Tolerance Falsification Problem**: - Identify system deviations causing violations of STL-defined requirements.
- Focus: Find minimal deviations likely to occur in practice.

m13610190@cgu.edu.tw

Solution

- Two-layer framework:
  - **Upper layer:** Explores deviations in system parameters.
  - **Lower layer:** Finds violating trajectories for a given deviation.
- Heuristic leveraging trajectory similarity to guide the search.

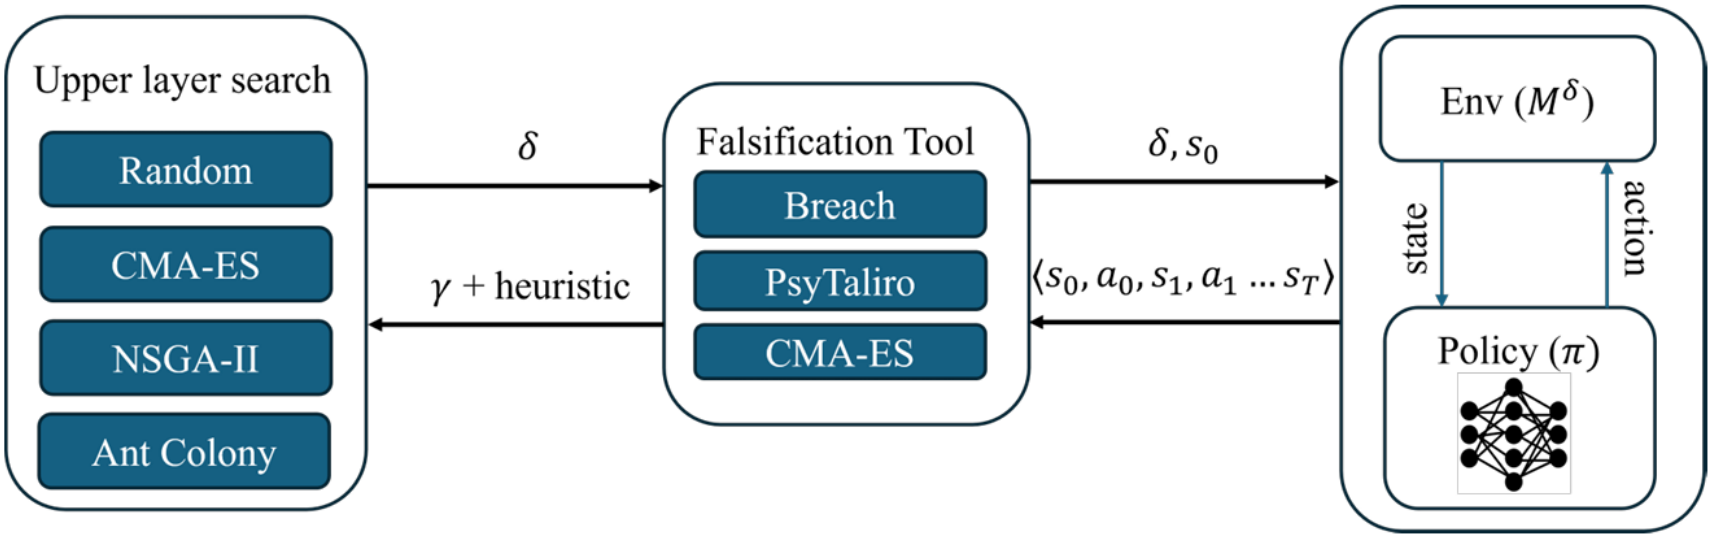

Figure: Two-layer framework

Implementation

- Python-based framework integrating:
  - CPS falsifiers (e.g., Breach, PsyTaLiRo).
  - Simulation platforms (OpenAI-Gym, PyBullet, Simulink).
  - CMA-ES optimization algorithm for both layers.

m1361019@cgu.edu.tw

Evaluation

- Benchmarked on diverse CPS scenarios:
  - Lunar Lander, Adaptive Cruise Control, etc.
- Outperformed one-layer falsifiers in finding smaller, realistic deviations.

	One-layer search			CMA-ES			CMA-ES w/ Heuristic		
	Viol.	Min Dst.	Avg. Dst.	Viol.	Min Dst.	Avg. Dst.	Viol.	Min Dst.	Avg. Dst.
Cartpole	**90**	0.300	**0.399**	69	0.285	0.449	79	**0.256**	0.417
LunarLander	-	-	-	74	0.026	**0.222**	**84**	**0.020**	0.293
CarCircle	11	0.143	0.255	22	0.102	**0.219**	**57**	**0.068**	0.454
CarRun	25	0.191	**0.249**	68	0.161	0.449	**109**	**0.156**	0.399
ACC	N/A	N/A	N/A	43	**0.110**	**0.323**	**110**	0.138	0.415
WTK	**300**	0.299	**0.443**	54	**0.296**	0.454	45	0.319	0.533

m1361019@cgu.edu.tw

Evaluation

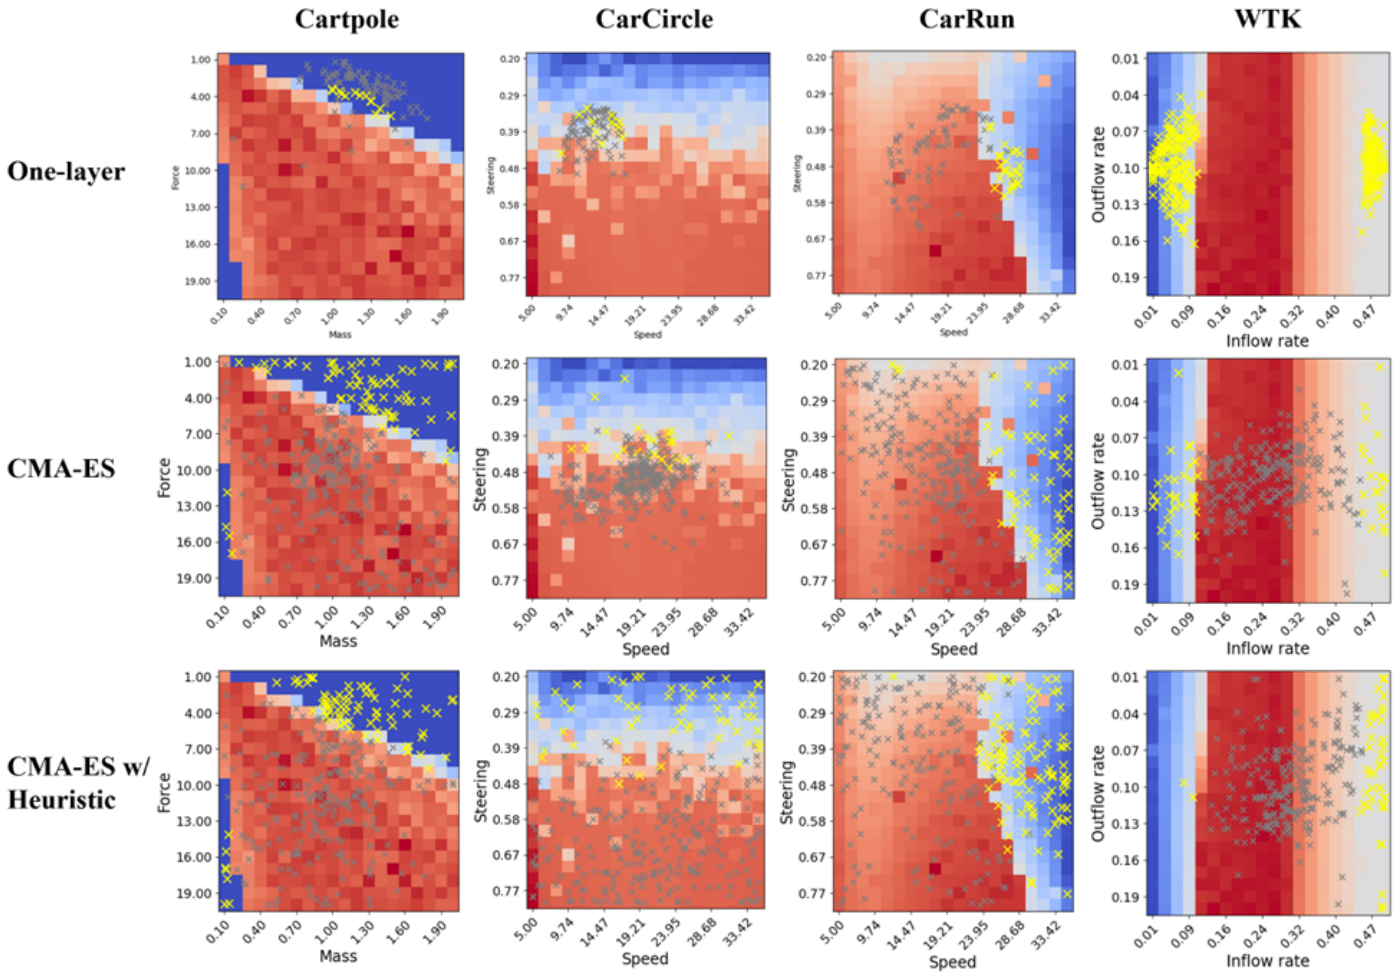

Figure: Visualise deviations with heatmaps

m1361019@cgu.edu.tw

Future Works

- Extend STL evaluation to probabilistic metrics.
- Explore alternative distance metrics (e.g., Wasserstein distance).
- Integrate STL decomposition for complex specifications.

m13610190@cgu.edu.tw

Questions?