

Subrat Prasad Panda

Ph.D. Candidate, Computer Science, NTU Singapore

Trustworthy RL · Neurosymbolic AI

✉ subratpr001@e.ntu.edu.sg

🔗 Google Scholar

📞 +65 ****3021

🐙 github.com/subratpp

🌐 subratpp.github.io

📍 Singapore

🌐 linkedin.com/in/subratpp

Research Statement

I am broadly interested in **Trustworthy Reinforcement Learning**, particularly in designing interpretable, robust, and verifiable policies for safety-critical domains. My current work focuses on **Neurosymbolic** approaches such as programmatic policies, combining explicit knowledge with learning to bridge System 1 (neural) and System 2 (symbolic) reasoning.

Research Areas and Skills

Reinforcement Learning	Hierarchical and goal-conditioned RL, sample-efficient learning in sparse-reward and long-horizon tasks, robust and distributional RL, RL fine-tuning of foundation models
Neurosymbolic AI	Programmatic and interpretable policies (differentiable decision trees), symbolic planning with incremental knowledge, interpretable state abstractions and world models
Safety & Verification	Formal verification of learned controllers, reach-avoid and safety certificates, abstraction-based reasoning for safety-critical systems
Applications	Cyber-physical systems, microgrid energy management, autonomous navigation and robotics, LLM-based planning for embodied agents
Programming	Python, C, PyTorch, TensorFlow, ROS 2, MATLAB, $\LaTeX$
Coursework	Mathematical Statistics, Deep Learning, Machine Learning, Computer Vision

Education

2023 – 2027 expected	Ph.D., Computer Science Nanyang Technological University (NTU) Thesis: <i>Neurosymbolic Techniques for Safe and Efficient Reinforcement Learning</i> Advisors: Arvind Easwaran (NTU), Blaise Genest (IPAL, CNRS France & Singapore)	Singapore
2019 – 2021	M.Tech., Computer Science Indian Statistical Institute (ISI) Thesis: <i>Strategies for User Allocation and Service Placement in Multi-Access Edge Computing</i> Advisor: Ansuman Banerjee CPI: 82.1%	Kolkata, India
2012 – 2016	B.Tech., Electronics and Instrumentation National Institute of Technology (NIT) Silchar Thesis: <i>Real-Time Position Control of DC Motor using PID and Fuzzy Logic</i> Advisor: Munmun Khanra CPI: 7.92/10	Silchar, India

Research Experience

2022 – 2023	Research Associate CNRS@CREATE ▶ Designed a granular control system for microgrids using Deep Reinforcement Learning.	Singapore
2021 – 2022	Machine Learning Engineer InFocusP Innovation Pvt. Ltd., under attachment to Google X (USA) ▶ Developed solutions for the Abstraction and Reasoning Challenge (ARC) using Reinforcement Learning.	Pune, India
2021 – 2021	Research Engineer TiHAN, IIT Hyderabad ▶ Developed braking and throttle control systems using Arduino for a test vehicle. ▶ Deployed a lane detection model on the NVIDIA AGX board.	Hyderabad, India

Publications

Under Review

- 1 S. P. Panda, B. Genest, and A. Easwaran, **Neurosymbolic Learning of Interpretable Abstraction in Reinforcement Learning**, Under review, 2027. [Neurosymbolic](#) [RL](#)
- 2 S. P. Panda, M. Prashant, and A. Easwaran, **Critical Option-Critic via Distributional RL**, Under review, 2027. [RL](#) [Hierarchical RL](#)
- 3 S. P. Panda, L. Yuke, B. Genest, D. Zikelic, and A. Easwaran, **Learning Abstractions for the Safety Verification of RL Policies**, Under review, 2027. [Verification](#) [Neurosymbolic](#) [RL](#)

Journal Articles

- 1 S. P. Panda, B. Genest, and A. Easwaran, **Approximation-Free Differentiable Oblique Decision Trees**, *Journal of Machine Learning Research (JMLR)*, vol. 27, 2026. [Neurosymbolic](#) [RL](#)
- 2 S. P. Panda, B. Genest, A. Easwaran, R. Rigo-Mariani, and P. Lin, **Methods for Mitigating Uncertainty in Real-Time Operations of a Connected Microgrid**, *Sustainable Energy, Grids and Networks*, vol. 38, p. 101334, 2024. [RL](#)

Conference Proceedings

- 1 S. P. Panda, B. Genest, and A. Easwaran, **Neurosymbolic Reasoning with Incremental Knowledge for Sample Efficient Hierarchical Reinforcement Learning**, in *ECML-PKDD*, 2026. [RL](#) [Hierarchical RL](#) [Neurosymbolic](#) [Planning](#) [Knowledge](#)
- 2 S. P. Panda, B. Genest, A. Easwaran, and P. N. Suganthan, **Vanilla Gradient Descent for Oblique Decision Trees**, in *European Conference on Artificial Intelligence (ECAI)*, IOS Press, 2024. [Neurosymbolic](#) [RL](#)
- 3 S. P. Panda, A. Banerjee, and A. Bhattacharya, **User Allocation in Mobile Edge Computing: A Deep Reinforcement Learning Approach**, in *2021 IEEE International Conference on Web Services (ICWS)*, 2021, pp. 447–458. [RL](#) [Edge AI](#)
- 4 S. P. Panda, K. Ray, and A. Banerjee, **Service Allocation/Placement in Multi-Access Edge Computing with Workload Fluctuations**, in *International Conference on Service-Oriented Computing*, 2021, pp. 747–755. [Edge AI](#)
- 5 S. P. Panda, K. Ray, and A. Banerjee, **Dynamic Edge User Allocation with User Specified QoS Preferences**, in *International Conference on Service-Oriented Computing*, 2020, pp. 187–197. [Edge AI](#)

Workshop Papers

- 1 S. P. Panda, B. Genest, and A. Easwaran, **Programmatic Reinforcement Learning for Trustworthy Microgrid Management**, *AI4X International Conference*, 2025. [RL](#) [Neurosymbolic](#)

Talks and Tutorials

- 1 S. P. Panda, **Neurosymbolic Techniques for Safe Reinforcement Learning**, Invited talk, 68NQRT Seminar, ICUR, Rennes, France, Sep. 2026.
- 2 S. P. Panda, **Reinforcement Learning**, Tutorial lecture with hands-on notebooks, *Deep Learning Bootcamp 2026*, NTU Singapore, 2026.

Open-Source Software and Educational Content

Research Code

- DTSemNet**  Differentiable oblique decision trees and tree policies for RL (JMLR 2026, ECAI 2024).
- InK-BWTS**  Neurosymbolic hierarchical RL with Belief World Tree Search (ECML-PKDD 2026).
- FITenth**  ROS 2 race stack (SLAM, particle-filter localisation, Pure Pursuit, etc.) on Jetson AGX.

Educational Content (ongoing effort)

- Deep RL Jutsu**  Notebook-style tutorials implementing deep RL algorithms from scratch (DQN, DDPG, SAC, PPO), with a curated guide to RL libraries and resources.
- GenAI**  Small-scale implementations of generative models (VAE, GAN, diffusion, state-space models, CLIP, VLA) on toy problems to explain how they work.
- RLFT**  Tutorial series on RL fine-tuning of foundation models: RLHF, DPO and GRPO, from a from-scratch GPT to a pretrained LLM.

Awards and Scholarships

2023 – 2027	DesCartes Scholarship: Ph.D. scholarship awarded by CNRS@CREATE under the DesCartes project (research collaboration between Singapore and France).
2021	M.Tech Best Thesis Nomination: ranked among the top 5 theses of the batch.
2019 – 2021	M.Tech Scholarship: scholarship awarded to pursue a master's degree.
2006 – 2012	Jawahar Navodaya Vidyalaya (JNV) Scholarship: fully funded residential education from Class 6 to 12.

Academic Services

Teaching and Mentoring

Final Year Projects	Undergraduate Final Year Projects NTU <i>Project supervision</i> <i>SafeDAG: Precedence-Constrained LLM Steering for Safer Planning in Embodied Agents</i> (ongoing); programmatic policies for interpretable RL using pre-trained models / LLMs; robustness vs. explainability trade-offs in CNNs.
URECA	Undergraduate Research Experience on Campus NTU <i>Research mentoring</i> Mechanistic interpretability of Decision Transformers; ARC reasoning with LLMs; physics-inspired model-based RL; residual reinforcement learning.
F1Tenth	F1Tenth Autonomous Racing Team NTU <i>Mentor</i> Built the ROS 2 race stack on Jetson AGX and advised students on control and navigation strategies.

Reviewing

Reviewer	CAV, AAAI, ECAI, IEEE TSMC, RTSS, NeurIPS, AAMAS, NeSy.
-----------------	---

Leadership

Vice President	CCDS Student Club , NTU: organized events such as the Deep Learning Bootcamp and AI Meetup.
-----------------------	--

Academic References

Ph.D. Advisor	Arvind Easwaran , Associate Professor, CCDS, NTU Singapore	arvinde@ntu.edu.sg
Ph.D. Advisor	Blaise Genest , Research Scientist, CNRS@CREATE, Singapore	blaise.genest@cnrsatcreate.sg
M.Tech Advisor	Ansuman Banerjee , Professor, ISI Kolkata, India	ansuman@isical.ac.in

Academic Collaborators

Verification	Dorđe Žikelić , Assistant Professor, CCDS, NTU Singapore	djordje.zikelic@ntu.edu.sg
Decision Trees	P. N. Suganthan , Research Professor, KINDI, Qatar University	p.n.suganthan@qu.edu.qa
Edge AI	Arani Bhattacharya , Assistant Professor, IIIT Delhi, India	arani@iiitd.ac.in