Nikhilesh Prabhakar

Richardson, TX | +1 (336) 875-9019 | prabhakar.nikhilesh@gmail.com | github.com/nikhileshp | nikhileshp.github.io

PhD Candidate. US Citizen - No sponsorship required.

EDUCATION

University of Texas at Dallas Aug 2020 - Present

Ph.D., M.S. in Computer Science. Advisor: Dr. Sriraam Natarajan (StARLing Lab)

Jonsson School Graduate Fellowship

Vellore Institute of Technology Jul 2016 - May 2020

B.Tech. in Computer Science (CGPA 9.23/10)

PUBLICATIONS

N. Prabhakar, V. Balaji, A. Karanam, K. Kersting, S. Natarajan. **Neurosymbolic Imitation Learning with Human Guidance: A Privileged Information Approach.** *Under Review, ECML, 2026.*

N. Prabhakar, R. Singh, H. Kokel, S. Natarajan, P. Tadepalli. **Exploiting Relational Planning and Task-Specific Abstractions for Multiagent RL.** *AAMAS, 2025.* [paper]

N. Prabhakar, R. Singh, H. Kokel, S. Natarajan, P. Tadepalli. Combining Planning and RL for Solving Relational Multiagent Domains. *AAAI Workshop (PRL), 2024.*

H. Kokel, **N. Prabhakar**, B. Ravindran. Hybrid Deep RePReL: Integrating Relational Planning and RL. *FUSION, 2023.* [paper]

N. Prabhakar, J. Anbarasi L. Analysis of the Complex Air Transport Network: A Global Perspective. *SNAM, 2021.*

RESEARCH INTERESTS

Primary: Reinforcement learning, foundation models for robotics, sample-efficient learning, neuro-symbolic AI

Secondary: Sim-to-real transfer, multi-agent systems, relational planning, imitation learning, interpretable policy learning, reward shaping

SELECTED PROJECTS

Multi-Agent Drone Coordination in Simulation Apr 2026 - Present

Closed-loop neural perception and decentralized multi-agent control in a continuous-action drone swarm, with a relational planner as a structured prior over subtasks. Onboard detectors ground visual observations into symbolic state over spatial predicates (zone priority, adjacency, status) and trigger dynamic replanning; multi-agent RL learns flight policies shaped by planning-derived rewards. Trained in physics simulation and validated in software-in-the-loop flight. Extends AAMAS 2025 relational-planning abstractions to continuous-control robotics with closed-loop neural perception. Testbed: precision-agriculture field survey.

Structured Visual Grounding for Long-Horizon Procedural Tasks Mar 2026 - Present

Two-stage neurosymbolic framework that induces procedural concepts as interpretable logical rules from egocentric video. Stage one fine-tunes an open-vocabulary scene graph generator with a task-adapted relation head on step-boundary keyframes, then links per-step graphs across time by object identity to form a sparse spatio-temporal graph per recording. Stage two grounds graph nodes and relations into symbolic atoms and runs inductive logic programming with predicate invention to learn per-task step rules and cross-task higher-order concepts. Dataset: CaptainCook4D (384 egocentric recordings). First milestone (May 2026): task-graph agreement per target concept against annotated procedural structure.

Relational Planning for Multi-Agent RL (MaRePReL) AAMAS 2025

Decomposes multi-agent tasks into operator hierarchies via a relational planner and learns per-operator state abstractions that drive sample-efficient policy learning and zero-shot generalization to novel object counts. Evaluated on relational Taxi and Office worlds against flat multi-agent baselines. Core contribution: structured planning priors act as a reward-shaping signal for deep RL, cutting the sample complexity that plagues unstructured multi-agent methods. [paper]

Neurosymbolic Imitation Learning with Human Attention Sep 2025 - Feb 2026

End-to-end neurosymbolic framework treating expert gaze as privileged information for imitation learning. A neural perception module grounds raw visual frames into probabilistic first-order atoms, and gaze-derived saliency scores reweight atom valuations to suppress task-irrelevant objects before a differentiable forward-chaining reasoner learns interpretable rule-weighted policies. Achieves 2.5x the score of the unmodulated baseline and 20x behavior cloning on Asterix, matching full-data performance with 10% of training data. Policy distillation yields inspectable rules that surface what the network has learned. [code]

PROFESSIONAL EXPERIENCE

StARLing Lab, University of Texas at Dallas

Aug 2021 - Present

Graduate Research Assistant, Richardson, TX

- Developed neuro-symbolic and relational RL frameworks for multi-agent systems; published at AAMAS, AAAI, and FUSION with reward shaping from structured planning priors as the central technical thread.
- Built sample-efficient policy learning pipelines combining inductive logic programming, relational regression trees, and deep RL across discrete and continuous-control benchmarks.
- Co-developed research with collaborators at Oregon State University and TU Darmstadt.

Other Experience

- Workflow Chair, AAAI-24 (Vancouver). Managed paper workflows for the 38th AAAI Conference.
- ML Consultant, Careband (Summer 2021). Activity recognition from accelerometer data for wearable deployment.

TEACHING

- Graduate TA: Machine Learning (CS6375), Digital Logic (CS4341), UT Dallas, 2021-2023.
- Undergraduate TA: Introduction to Machine Learning (CSE4020), VIT, Fall 2019.

SKILLS

Programming: Python, C++, PDDL, SQL, ASP (Clingo), R

Research Methods: Reinforcement learning, multi-agent RL, sample-efficient learning, imitation learning, relational planning, neurosymbolic reasoning, inductive logic programming, sim-to-real transfer, foundation-model grounding for robotics

Coursework (in progress): ARENA Mechanistic Interpretability Curriculum (Alignment Research Engineer Accelerator); Generative AI with Large Language Models (DeepLearning.AI / AWS).