

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. and M.S. in Computer Science. Advisor: Dr. Sriraam Natarajan. Jonsson School Graduate Fellowship.

Vellore Institute of Technology

Jul 2016 - May 2020

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

RESEARCH INTERESTS

Primary: Sequential decision-making, interpretable policy learning, reinforcement learning, neuro-symbolic AI.

Secondary: Inductive logic programming, sample-efficient learning, planning under uncertainty, multi-agent systems.

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.*

H. Kokel, N. Prabhakar, B. Ravindran. Hybrid Deep RePREL: Integrating Relational Planning and RL. *FUSION, 2023.*

SELECTED PROJECTS

Automated Predicate Discovery for Relational RL

Apr 2026 - Present

Discovers relational symbolic abstractions from reinforcement-learning experience without hand-written domain knowledge. Reward-adjacent trajectory segmentation, breakpoint-context clustering, and per-operator state abstractions yield interpretable predicate vocabularies and policies that zero-shot generalize across task scales, recovered from reward signal alone.

Neurosymbolic Imitation Learning with Privileged Information

Sep 2025 - Feb 2026

Framework treating expert attention as privileged information for sample-efficient imitation. A neural module grounds observations into probabilistic first-order atoms; saliency reweights atom valuations before a differentiable forward-chaining reasoner learns interpretable rule-weighted policies. Achieves 2.5x the unmodulated baseline and 20x behavior cloning at 10% of training data.

Hierarchical Multi-agent Relational RL (MaRePREL)

AAMAS 2025

Hierarchical decision-making framework combining relational planning with deep reinforcement learning across coordinated agents. Decomposes joint tasks via relational subtask abstractions, yielding sample-efficient policies that transfer across task scales without retraining.

PROFESSIONAL EXPERIENCE

StARLing Lab, University of Texas at Dallas

Aug 2021 - Present

Graduate Research Assistant

Richardson, TX

- Developed neuro-symbolic and relational reinforcement-learning frameworks for sequential decision-making and multi-agent coordination, resulting in publications at AAMAS, AAI, and FUSION.
- Co-developed research with collaborators at Oregon State University and TU Darmstadt.

ML Consultant, Careband (Summer 2021): activity recognition from accelerometer data for wearable deployment. Workflow Chair, AAI-24 (Vancouver).

TEACHING

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

SKILLS

Research Areas: Reinforcement learning, neuro-symbolic AI, sequential decision-making, interpretable policy learning, multi-agent systems, planning under uncertainty.

Methods: Inductive logic programming, hierarchical decision processes, sample-efficient learning, predicate invention, differentiable reasoning, answer set programming.

Programming: Python, C++, SQL, R, Java.