Rajesh D. Mangannavar

rajeshdm9@gmail.com | 541-908-9645 | linkedin.com/RajeshDM | github.com/RajeshDM | Oregon, USA

Education

Computer Science and Artificial Intelligence — ***PhD - Dual major*** Sept 2019 - Present
Oregon State University | Advisor: Dr. Prasad Tadepalli GPA : 3.82

Theoretical Computer Science — ***Integrated M.Tech*** Aug 2013 - Jul 2018
International Institute of Information Technology, Bangalore (IIITB), India GPA: 3.64

Work Experience

Graduate Research Assistant — ***Oregon State University*** Sept 2019 - Present

- **Intelligent Agent for Object Retrieval in Interactive Environments**
  - Developed an agent navigating the AI2Thor simulated environment to retrieve soccer balls, demonstrating advanced machine common sense reasoning with 91% success rate across 5000 different scenes.
  - Implemented perception and reasoning algorithms using RGB and depth images for environment mapping, enabling obstacle avoidance and deduction of object locations based on their absence in certain areas, including tool utilization when necessary.
  - Led core algorithm development while mentoring 5 team members working on perception, tool utilization, and system integration.
- **Hierarchical Object-Oriented POMDP Planning for Object Rearrangement** [2]
  - Developed HOO-POMDP framework for multi-object rearrangement problems in partially observable, multi-room environments with robust belief management for handling uncertainty.
  - Designed an abstract planner generating sub-goals paired with low-level policies and abstraction system, handling complex scenarios including blocked paths, object swaps, and goal obstructions.
  - Achieved 49% success rate for full scene rearrangement and 71% for object-level rearrangement, outperforming non-hierarchical baseline (13% and 33% respectively).
  - Demonstrated scalability across environments with 5-20 objects and 1-4 rooms, maintaining performance even when only 20% of objects were initially visible.
- **Graph Neural Network (GNN) Based Action Ranking for Planning** [3]
  - Proposed a novel GNN approach to learn relational policies for classical planning through action ranking.
  - Designed an action-centric graph representation, a custom GNN encoder, and GRU-based decoder that effectively captures relationships between objects, predicates, and actions.
  - Demonstrated strong generalization ability, successfully solving problems significantly larger (upto 8 times larger in some cases) than training instances across domains like BlocksWorld, Gripper, and Miconic (100% success rate) as well as more complex domains such as Grid (92%) and Logistics (71%) while outperforming a leading value-function based method by achieving 13× higher coverage on hard problems (89% vs 6.5%) and 2.5× better plan quality ratio (0.99 vs 0.39).
  - Reduced training data requirements by 75% compared to value-function methods while delivering superior performance, enabling more efficient and scalable approaches to automated planning. Implemented in **PyTorch Geometric**.

Software Engineer — ***CISCO*** AUG 2018 - SEPT 2019

- Strengthened security of routers running on IOS XR (a Cisco proprietary operating system for large-scale routers that support the internet backend at the global level) by developing and implementing secure workflows across a suite of services.
- Accelerated user onboarding by implementing natural language processing to enhance command line interface usability on Cisco's products and leveraged reinforcement learning to continuously improve the model through user feedback.

Technical Projects

Deep Learning — ***Denoising Diffusion Probabilistic Models*** Oct 2023 - Mar 2024

- Developed an innovative solution using diffusion models to efficiently generate paths for objects of various shapes in cluttered environments, overcoming the limitations of traditional methods like Rapidly-exploring Random Trees (RRT).
- Achieved 80% accuracy in path generation and obstained a 3x speed improvement over RRT for paths longer than 50 steps.

Ag Aid Hackathon (Winner) — ***XGBoost*** Jan 2023 - Jan 2023

- Developed a machine learning solution to forecast cool-season precipitation in the Sacramento Valley to help farmers prepare crops effectively.
- Overcame challenges posed by the coastal range rain shadow and Pacific currents, by leveraging 73 years of precipitation data, climate indices, and 20 years of spatial data for soil moisture, ice extent, and sea temperature.
- Conducted extensive data exploration and feature engineering, employing techniques such as Dynamic Time Warping, aggregating spatial features, and correlation analysis to extract meaningful patterns from the chaotic precipitation data.
- Built an XGBoost regression model combined with a classifier to predict precipitation, achieving 75% accuracy on temporal data and 84% accuracy on spatial data.

Integrating Learned Affordance Models and Symbolic Planning [1] — ***Planning*** Jun 2020 - Jul 2021

- Developed a novel approach integrating PDDLStream with learned affordance models, enabling intelligent agents to successfully navigate and manipulate objects in 3D environments.
- Achieved 100% success rate in basic object manipulation tasks and 80% in more complex blocked navigation scenarios, outperforming traditional hand-designed sampling methods by 30-300%.
- Created a framework for learning probabilistic models that sample valid continuous parameters for robot actions, improving performance in task and motion planning problems.
- Demonstrated superior performance in AI2-Thor environments with varying complexity, showing robust generalization across different object configurations.

Deep Learning — *Convolutional Neural Networks (CNN)* Jan 2020 - April 2020
- Implemented 2x2 convolutional kernels in CNNs, improving efficiency and reducing model size while maintaining accuracy. Enhanced robustness to label noise using techniques from 'Learning to Learn from Noisy Labeled Data'.
- Reduced FLOPs by 50% in Resnet-50 C2sp model with minimal impact on parameters and accuracy. **Implemented in PyTorch.**

k-SVRG: Variance Reduction for Large Scale Optimization — *Optimization Techniques* Sept 2019 - Dec 2019
- Analyzed several variance reduction techniques for Stochastic Gradient Descent (SGD), including SVRG and SAGA, identifying key limitations in memory usage and computational efficiency.
- Implemented and evaluated the k-SVRG algorithm which balances memory efficiency and computational speed by storing only a subset of gradients, demonstrating 30% faster execution time compared to standard SVRG.
- Developed an enhanced k-SVRG algorithm with adaptive learning rates that dynamically adjusts the learning rate (n = n * 0.9) after fixed iterations, resulting in 10% faster convergence while maintaining solution quality.

Deep Learning — *Reinforcement Learning in continuous environments (RL) [5]* Jan 2018 - Jul 2018
- Designed a Prioritized Deep Deterministic Policy Gradient (PDDPG) algorithm that outperformed DDPG on 8 out of 10 Mujoco environments through improved exploration strategies.
- Implemented parameter space noise for exploration, demonstrating up to 40% faster learning convergence compared to standard action space noise methods.
- Achieved significantly higher cumulative rewards with PDDPG, showing continued improvement even after 2000 training iterations where standard DDPG performance plateaued.

Game theory and FMEA analysis [4] — *Cyber Physical Systems* Jan 2017 - May 2017
- Developed a novel game-theoretic approach to model unpredictable attacker behavior and optimize defense strategies for cyber-physical systems.
- Applied Failure Mode and Effects Analysis (FMEA) to identify vulnerabilities and assess attack impact, informing cybersecurity planning and demonstrated improved system resilience by determining optimal defensive actions against specific attack vectors.

Scholastic achievements

- Winner of the Ag Aid Hackathon, Water challenge held at OSU 2023.
- Dean's Merit List, IIIT Bangalore – June 2018.
- Second runner-up at Hack.Bangalore by Mercedes-Benz Research and Development India (2016); Idea presented at Mobile World Congress (MWC), Barcelona (2017) representing Daimler.
- Won 1st prize in DreamWorks hackathon conducted at IIITB in July 2015.

Technical Skills

Languages:Python, C/C++, Java, Assembly, MATLAB, R
Frameworks: PyTorch, PyTorch Geometric, AI2Thor, Habitat (simulation environments), TensorFlow, POMDP-py, PDDL
AI Techniques: Graph Neural Networks (GNN), Reinforcement Learning (RL), Diffusion (DDPM), Convolutional Neural Networks (CNN), XGBoost, Partially Observable MDPs (POMDPs), Task and Motion Planning (TAMP)

Publications

[1] Rajesh Mangannavar. Planning with affordances: Integrating learned affordance models and symbolic planning. *arXiv preprint arXiv:2502.02768*, 2025.

[2] Rajesh Mangannavar, Alan Fern, and Prasad Tadepalli. Hierarchical object-oriented pomdp planning for object rearrangement. *arXiv preprint arXiv:2412.01348*, 2024.

[3] Rajesh Mangannavar, Stefan Lee, Alan Fern, and Prasad Tadepalli. Gabar: Graph attention-based action ranking for relational policy learning. *arXiv preprint arXiv:2412.04752*, 2024.

[4] Rajesh Mangannavar and Dugar Pranay. Using fmea and game theory for security models and analysis of cyber-physical systems. In *International Conference of Communications security and Information, Sydney, Australia*, 2018.

[5] Rajesh Mangannavar and Gopalakrishnan Srinivasaraghavan. Learning agents with prioritization and parameter noise in continuous state and action space. In *Advances in Neural Networks–ISNN 2019: 16th International Symposium on Neural Networks, ISNN 2019, Moscow, Russia, July 10–12, 2019, Proceedings, Part I 16*, pages 202–212. Springer, 2019.