

Dixant Mittal

28C, Dover Crescent, Singapore (133028)

www.dixantmittal.com
dixant.mittal@gmail.com

EDUCATION

- National University of Singapore** Singapore
Doctor of Philosophy (Ph.D.) in Computer Science; CAP: 4.66/5.00 Jan 2019 – Present
- National University of Singapore** Singapore
Master of Computing (M.Comp.) in Computer Science; CAP: 4.25/5.00 Aug 2017 – Dec 2018
Master's Thesis: *Active Information Gathering to Disambiguate Referring Expressions*
- National Institute of Technology Kurukshetra** Kurukshetra, India
Bachelor of Technology (B.Tech.) in Information Technology; CGPA: 7.92/10.00 Aug 2011 – Jul 2015

PUBLICATIONS

- Mohit Shridhar, **Dixant Mittal**, David Hsu; "INGRESS: Interactive Visual Grounding of Referring Expressions"; The International Journal of Robotics Research, 20.

PROJECTS

- **Posterior Variance Estimates to Improve Exploration in Monte Carlo Tree Search (Work in progress):** We aim to improve exploration in MCTS by explicitly representing a distribution over the value function and combining it with a memory module.
- **POMCPnet: Learning to Search in Partially Observable Environments:** Online planning algorithms are preferred to solve large POMDP problems. Performance of an online planning algorithms depends on the quality of heuristics used, which requires a considerable amount of domain expertise. We aim to let a planning algorithm learn the optimal heuristics for a problem and how to use it to guide search, on its own, by embedding it inside a neural network architecture and jointly optimising it in an end-to-end approach. We design a differentiable version of POMCP, named POMCPnet, by approximating the simulation, rollout and backup phases of POMCP as different neural networks.
- **Active Information Gathering to Disambiguate Referring Expressions:** In a Human-Robot Interaction environment, the human instructions can be ambiguous sometimes. We propose to interactively gather information about shape, colour or placement of the referred object by formulating the problem as POMDP and solving it using an online planner. Additionally, we formulate a surrogate objective function, which satisfies the adaptive submodular properties. Consequently, a greedy solution of the surrogate objective is near-optimal. Both approaches were 20% more accurate and 33% faster to identify the referred object compared to the baselines.

WORK EXPERIENCE

- Moovita** Singapore
Research Intern Dec 2017 - Dec 2018
- **Intuition-Net:** A high-level behaviour planner for autonomous vehicles to make decisions under complex and partially observable scenarios. It combines model-free learning (neural networks) with model-based planning (UCT search), and uses particle filters to deal with environment uncertainties.
 - **image-to-control:** An end-to-end neural network based steering controller for autonomous vehicles. This system takes RGB image from front camera as input and outputs a steering value.
- Ixigo** Gurgaon, India
Senior Software Engineer Mar 2017 - Jul 2017
- **Trending Searches:** Built using Kafka, this service was used to analyse the search requests for trains/stations and provide a list of most searched entities in the past week.
- Snapdeal** Gurgaon, India
Software Engineer Jul 2015 - Mar 2017
- **Identity Management System:** Provided scalable APIs to handle user login, session management and user information for a user base of 25 million users.

SKILLS

- **Research Interests:** Machine Learning, Planning/Search, Reinforcement Learning, Approximate Inference
- **Programming Languages:** Python, Rust, C++
- **Frameworks:** Pytorch, Kafka