Sumit Kumar

412 315 5110 | sumit.sks4@gmail.com | LinkedIn | Website | GitHub | Google Scholar

Work Experience

Aurora Innovation Mountain View, CA

Staff Tech Lead Manager, Behavior Planning *Mar 2023 – Present*

- Serving as DRI and team lead for Aurora's Intent Prediction team (up to 8 engineers) within the Behavior Planning organization, owning the core ML forecasting component, including architecture decisions, roadmap, and cross-functional alignment across Motion Planning, Perception, Maps, Simulation, and Testing.
- Led the transition of the motion planning architecture from fragmented, heuristic-based decision systems to a unified, learned model architecture, establishing Intent Prediction as the primary decision-making engine.
- Spearheaded the team's contribution to Aurora's planning paradigm shift from discrete decision-based reasoning to continuous trajectory reasoning, delivering accurate actor forecasts required by downstream cost functions.
- Drove the architectural transition of Aurora's core ML forecasting model from GNN to transformer, delivering a Scene Forecaster model that served as the forecasting engine for Aurora's commercial launch and remains in production onboard all fully autonomous Aurora vehicles today.
- Pioneered Aurora's first ML-based Scene Generator, a transformer model that produces realistic human-like actor trajectories from keyframe inputs, enabling directed scenario creation for motion planner training and evaluation.
- Leading an effort to build a large-scale generative scenario creation pipeline combining natural language-driven keyframe injection with the Scene Generator model to augment real driving logs with synthetic actor modifications, generating diverse training and evaluation scenarios for the motion planner.

Software Engineer, Motion Planning *Jan 2021 – Mar 2023*

- Proposed and developed Aurora's first GNN-based scene reasoning model producing lane occupancy, multi-modal trajectories, and interaction likelihoods for the downstream planner — resulting in multiple Aurora patents.
- Deployed the model onboard Aurora's semi-truck and passenger car platforms, producing actor forecasts directly consumed by the downstream engineered ranker system.
- Led projects of up to 10 engineers to expand model capabilities and build supporting infrastructure spanning training pipelines, dataset generation, metrics, visualization and debugging tooling.

UATC LLC (Acquired by Aurora Innovation) Pittsburgh, PA

Autonomy Engineer, Prediction *Oct 2019 – Jan 2021*

- Developed TrafficGraphNet, a semi-supervised GNN-based interaction prediction model for actor behavior forecasting, representing traffic actors and elements (traffic lights, stop signs) as a hybrid graph to explicitly model discrete interaction modes and predict accurate multi-agent trajectories.

Advanced Agent – Robotics Technology Lab, CMU Pittsburgh, PA

Graduate Research Assistant (Advisor: Prof. Katia Sycara) *Aug 2017 – Aug 2019*

- Proposed Recurrent Neural Processes (RNP), a deep latent variable model for spatiotemporal prediction with uncertainty estimation, outperforming GP-based baselines across multiple benchmarks (MS thesis, CMU 2019).
- Developed a GNN-based multi-agent RL model using a shared agent-entity graph for cooperative behavior, achieving state-of-the-art results with strong zero-shot generalization across team sizes (ICML 2019 Workshop).
- Co-developed MAME, a meta-RL framework with decoupled exploration and exploitation policies and self-supervised adaptation, outperforming prior baselines on continuous control benchmarks (CoRL 2019).
- Proposed an active learning algorithm using Gaussian Processes for autonomous informative sample collection in high-throughput crop phenotyping (AAMAS 2019).

Personal Robotics Lab, CMU Pittsburgh, PA

Research Intern (Advisor: Prof. Siddhartha Srinivasa) *May 2016 – Aug 2016*

- Proposed a probabilistic belief model over robot configuration space to reduce collision checks during motion planning, leveraging kNN methods and topological C-space structure to efficiently identify collision-free configurations.

Education

Carnegie Mellon University — Pittsburgh, USA
Master of Science in Robotics (GPA: 4.14 / 4.33) — *Aug 2017 – Aug 2019*

Indian Institute of Technology Kanpur — Kanpur, India
Bachelor of Technology in Mechanical Engineering, Minor in AI (GPA: 9.4 / 10) — *Aug 2013 – May 2017*

Patents

- **Systems and Methods for Interaction-Based Trajectory Prediction** [Patent] — *Issued May 2024*
  US Patent No. 11,975,726
- **Goal-Based Motion Forecasting** [Patent] — *Issued Oct 2023*
  US Patent No. 11,801,871
- **Autonomous Vehicle Motion Planning** — *Under Review*
  US Patent Application No. 63/616,284
- **Training Motion Planning with Hybrid Simulation** — *Under Review*
  US Patent Application No. 19/433,843

Publications

- **Interaction-Based Trajectory Prediction Over a Hybrid Traffic Graph** [Paper] — *IROS 2021*
  Sumit Kumar, Yiming Gu, Jerrick Hoang, Galen Clark Haynes, Micol Marchetti-Bowick
- **MAME: Model Agnostic Meta Exploration** [Paper, Code] — *CoRL 2019*
  Swaminathan Gurumurthy, Sumit Kumar, Katia Sycara — *Oral at ICML 2019 Workshop*
- **Learning Transferable Cooperative Behavior in Multi-Agent Teams** [Paper, Code] — *ICML 2019*
  Akshat Agarwal*, Sumit Kumar*, Katia Sycara — *Spotlight at ICML 2019 Workshop*
- **Active Learning with Gaussian Processes for High Throughput Phenotyping** [Paper, Code] *AAMAS 2019*
  Sumit Kumar, Wenhao Luo, George Kantor, Katia Sycara

Selected Projects

Learning Hierarchical Policies in Dynamic Environments [Report]

Advisor: Prof. Ruslan Salakhutdinov, School of Computer Science, CMU

- Proposed a hierarchical meta-RL framework combining MAML-based skill learning with a high-level master policy to solve sparse reward locomotion tasks in dynamic environments with varying friction and agent morphology.
- Demonstrated that MAML-trained sub-policies achieve up to 2x higher task success rate and significantly fewer steps compared to pre-trained policies across challenging environment configurations.

Deep Reinforcement Learning for Sparse-Reward Manipulation Problems [Report]

Advisor: Prof. Matt Mason, School of Computer Science, Carnegie Mellon University

- Proposed Prioritized Hindsight Experience Replay (PHER), combining HER with TD-error-based transition prioritization and importance sampling to improve sample efficiency for multi-goal robotic manipulation tasks.
- Demonstrated improved learning speed and task success rate over standard HER on 7-DOF Fetch robot manipulation tasks including pick-and-place and push in OpenAI Gym environments.