

Rohan Chitnis

ronuchit@gmail.com | www.rohanchitnis.com | <https://github.com/ronuchit> | (408) 772-2079

Education

Massachusetts Institute of Technology, Cambridge, MA

GPA: 4.8 / 5.0. Minor: Econometrics.

Doctor of Philosophy in Electrical Engineering and Computer Science, May 2022.

Master of Science in Electrical Engineering and Computer Science, May 2018.

Relevant Coursework: Advanced Algorithms, Randomized Algorithms, Advanced Natural Language Processing, Econometrics, Time Series Analysis.

University of California, Berkeley, Berkeley, CA

Graduated with Highest Honors (GPA in top 3%). GPA: 3.95 / 4.0.

Bachelor of Science in Electrical Engineering and Computer Sciences, May 2016.

Relevant Coursework: Advanced Robotics, Machine Learning, Deep Reinforcement Learning, Artificial Intelligence, Computer Vision, Optimization, Graphics, Computational Geometry, Image Processing, Probability and Random Processes, Algorithms, Data Structures.

Research Experience

MIT Learning & Intelligent Sys. Lab (Adviser: Leslie Kaelbling) 09/16 - 05/22

- Conduct research in artificial intelligence for robotics, toward a PhD.
- Focus on integrated learning and planning for solving long-horizon robotic tasks.

Facebook AI Research, Robotics (Research Intern with Abhinav Gupta) 06/19 - 09/19

- Researched formulations of intrinsic motivation for emergence of synergistic behavior in robotics via deep reinforcement learning.
- Built and tested my algorithms in bimanual manipulation environments, both in simulation and on real Sawyer robot arms.

Google Brain, Robotics (Research Intern with Sergey Levine) 05/17 - 09/17

- Researched methods for speeding up deep reinforcement learning for robotics, by regressing on event outcomes to improve exploration.
- Designed and implemented a stand-alone framework for learning-based manipulation.

UC Berkeley Robot Learning Lab (Adviser: Pieter Abbeel) 02/13 - 05/16

- Performed work in hierarchical combined task and motion planning for execution of long-horizon robotic tasks such as laundry.
- Integrated reinforcement learning to improve existing approaches.
- Lead coordinator of lab outreach program, providing tours to visitors of varied ages.

UC Berkeley Oscii Lab (Adviser: John DeNero) 04/15 - 05/16

- Conducted research in Natural Language Processing.
- Improved performance of neural machine translation using Huffman code compression.

Journal Publications

Integrated Task and Motion Planning

Caelan Reed Garrett, **Rohan Chitnis**, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

Annual Review of Control, Robotics, and Autonomous Systems, Volume 4, 2021.

Conference Publications

Learning Efficient Abstract Planning Models that Choose What to Predict

Nishanth Kumar, Willie McClinton, **Rohan Chitnis**, Tom Silver, Tomás Lozano-Pérez, Leslie Pack Kaelbling.

Conference on Robot Learning (CoRL), 2023.

Predicate Invention for Bilevel Planning

Tom Silver*, **Rohan Chitnis***, Nishanth Kumar, Willie McClinton, Tomás Lozano-Pérez, Leslie Pack Kaelbling, Joshua Tenenbaum.

AAAI Conference on Artificial Intelligence (AAAI), 2023.

Learning Neuro-Symbolic Relational Transition Models for Bilevel Planning

Rohan Chitnis*, Tom Silver*, Joshua Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling.

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022.

Reinforcement Learning for Classical Planning: Viewing Heuristics as Dense Reward Generators

Clement Gehring*, Masataro Asai*, **Rohan Chitnis**, Tom Silver, Leslie Pack Kaelbling, Shirin Sohrabi, Michael Katz.

International Conference on Automated Planning and Scheduling (ICAPS), 2022.

Towards Optimal Correlational Object Search

Kaiyu Zheng, **Rohan Chitnis**, Yoonchang Sung, George Konidaris, Stefanie Tellex.

IEEE International Conference on Robotics and Automation (ICRA), 2022.

Learning Symbolic Operators for Task and Motion Planning

Tom Silver*, **Rohan Chitnis***, Joshua Tenenbaum, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.

GLIB: Efficient Exploration for Relational Model-Based Reinforcement Learning via Goal-Literal Babbling

Rohan Chitnis*, Tom Silver*, Joshua Tenenbaum, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

AAAI Conference on Artificial Intelligence (AAAI), 2021.

Planning with Learned Object Importance in Large Problem Instances using Graph Neural Networks

Tom Silver*, **Rohan Chitnis***, Aidan Curtis, Joshua Tenenbaum, Tomás Lozano-Pérez, Leslie Pack Kaelbling.

AAAI Conference on Artificial Intelligence (AAAI), 2021.

CAMPs: Learning Context-Specific Abstractions for Efficient Planning in Factored MDPs

Rohan Chitnis*, Tom Silver*, Beomjoon Kim, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

Conference on Robot Learning (CoRL), 2020.

Efficient Bimanual Manipulation Using Learned Task Schemas

Rohan Chitnis, Shubham Tulsiani, Saurabh Gupta, Abhinav Gupta.

IEEE International Conference on Robotics and Automation (ICRA), 2020.

Intrinsic Motivation for Encouraging Synergistic Behavior

Rohan Chitnis, Shubham Tulsiani, Saurabh Gupta, Abhinav Gupta.

International Conference on Learning Representations (ICLR), 2020.

Learning Compact Models for Planning with Exogenous Processes

Rohan Chitnis, Tomás Lozano-Pérez.

Conference on Robot Learning (CoRL), 2019.

Learning Quickly to Plan Quickly Using Modular Meta-Learning

Rohan Chitnis, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

IEEE International Conference on Robotics and Automation (ICRA), 2019.

Learning What Information to Give in Partially Observed Domains

Rohan Chitnis, Leslie Pack Kaelbling, Tomás Lozano-Pérez.

Conference on Robot Learning (CoRL), 2018.

Integrating Human-Provided Information Into Belief State Representation Using

Dynamic Factorization

Rohan Chitnis, Leslie Pack Kaelbling, Tomás Lozano-Pérez.
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018.

Finding Frequent Entities in Continuous Data

Ferran Alet, **Rohan Chitnis**, Leslie Pack Kaelbling, Tomás Lozano-Pérez.
International Joint Conference on Artificial Intelligence (IJCAI), 2018.

Sequential Quadratic Programming for Task Plan Optimization

Dylan Hadfield-Menell, Christopher Lin, **Rohan Chitnis**, Stuart Russell, Pieter Abbeel.
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2016.

Guided Search for Task and Motion Plans Using Learned Heuristics

Rohan Chitnis, Dylan Hadfield-Menell, Abhishek Gupta, Siddharth Srivastava, Edward Groshev, Christopher Lin, Pieter Abbeel.
IEEE International Conference on Robotics and Automation (ICRA), 2016.

Variable-Length Word Encodings for Neural Translation Models

Rohan Chitnis, John DeNero.
Conference on Empirical Methods in Natural Language Processing (EMNLP), 2015.

Modular Task and Motion Planning in Belief Space

Dylan Hadfield-Menell, Edward Groshev, **Rohan Chitnis**, Pieter Abbeel.
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

Combined Task and Motion Planning Through an Extensible Planner-Independent Interface Layer

Siddharth Srivastava, Eugene Fang, Lorenzo Riano, **Rohan Chitnis**, Stuart Russell, Pieter Abbeel.
IEEE International Conference on Robotics and Automation (ICRA), 2014.

Honors / Awards

Facebook Fellowship Finalist, 2020. Top 4% of applicants for the Facebook Fellowship, Machine Learning track.

NSF GRFP Fellow, 2016. Awarded NDSEG Fellowship (declined).

Hertz Fellowship Finalist, 2016. One of 40 finalists for the Hertz Fellowship, a highly reputable fellowship for student researchers in the physical, biological, and engineering sciences.

Runner-up for the Computing Research Association (CRA) Outstanding Undergraduate Researcher Award (Male, PhD-granting institution), 2016. Highly prestigious award recognizing North American undergraduate students who show outstanding research potential in a field of computing.

Sole recipient of the EECS Mark D. Weiser Excellence in Computing Scholarship, 2015. Merit-based scholarship awarded to one student in the Berkeley EECS department each year for excellence in research.

Member of the EECS Honors Degree Program, concentration in Mathematics. Honors program with 20-30 students. Requirements include research and extended studies in concentration outside EECS.

UC Berkeley Outstanding Graduate Student Instructor (OGSI) Award recipient, 2015. Awarded to top 10% of Teaching Assistants across the university each year.

UC Berkeley Regents' and Chancellor's Scholar. Merit-based scholarship awarded to top 1.5% of applicants to UC Berkeley each year.

National Merit Scholar. Merit-based scholarship awarded to high-achieving high school students for partial college tuition payment.

**Teaching
Assistantships**

CS106A: Code In Place. *Stanford, free online course during COVID-19.* Summer 2020.
6.036: Introduction to Machine Learning. *MIT.* Fall 2018.
CS189: Introduction to Machine Learning. *UC Berkeley.* Spring 2016.
CS188: Introduction to Artificial Intelligence. *UC Berkeley.* Fall 2015.
CS61A: Structure and Interpretation of Computer Programs. *UC Berkeley.* Spring 2015, Spring 2014, Summer 2013.
CS61C: Great Ideas in Computer Architecture. *UC Berkeley.* Fall 2014.

Industry

Nuro Inc., Planning (Software Engineering Intern) 06/21 - 09/21

- Explored integrated data-driven and rule-based reasoning to improve planning stack.
- Developed a synergistic approach that improves behavior in 20-40% of challenge scenes.

Airbnb Inc., Search Ranking (Software Engineering Intern) 06/16 - 08/16

- Worked on incorporating mobile data into machine learning models used by the Search Ranking team. Collected and organized data using tools in Hive, Presto, and Scala.

eBay Inc., Checkout (Software Engineering Intern) 05/14 - 08/14

- Developed an end-to-end pipeline to create a model that classifies checkout transactions.
- Collected data using Hadoop MapReduce under the Apache Pig framework.

**Technical
Skills**

Languages: Python, C++, Java, Scala, LaTeX.
Software: PyTorch, TensorFlow, Theano, Unix/Linux, Robot Operating System (ROS), OpenCV, Apache Pig, Apache Spark, Hadoop MapReduce.