

Ashesh Jain

142 Gates Building, Stanford University, CA 94305
ashesh@cs.cornell.edu, www.cs.cornell.edu/~ashesh

- Interests** My research interest lies at the intersection of machine learning, robotics, and computer vision. Broadly, I build machine learning systems & algorithms for agents – such as robots, cars etc. – to learn from informative human signals at a large-scale. Most of my work has been in multi-modal sensor-rich robotic settings, for which I have developed sensory fusion deep learning architectures. I have developed and deployed algorithms on multiple robotic platforms (PR2, Baxter etc.), on cars, and crowd-sourcing systems.
- Education**
- Cornell University**, New York, USA (2012-2016)
Ph.D. Computer Science
- Specializing in Machine Learning, Robotics, and Computer Vision
 - Learning from large-scale human signals for robots and assistive cars
 - First author papers in NIPS, CVPR, ICCV, SIGKDD, ICRA, ISRR, and IJRR
 - PhD committee: Ashutosh Saxena, Bart Selman, Thorsten Joachims, Doug James, Robert Kleinberg.
- Indian Institute of Technology Delhi**, India (2007-2012)
B.Tech. Electrical Engineering & M.Tech. Information and Communication Technology
- Advisor: Prof. Manik Varma and Prof. S.V.N. Vishwanathan
 - Thesis: Large-scale Algorithms for Multiple Kernel Learning
- Experience**
- Stanford University**, California, USA (2014-2016)
Visiting Ph.D. Computer Science
- Working with Prof. Ashutosh Saxena and Prof. Silvio Savarese
 - Leading Brain4Cars and Machine Learning lead on RoboBrain.
- Purdue University**, Indiana, USA (Summer 2011)
Research Intern, Statistics Department with Prof. S. V. N. Vishwanathan
- Developed Multiple Kernel Learning algorithms to scale to Millions of kernels
- The Royal Bank of Scotland**, India (Summer 2010)
- Software**
- NeuralModels**: A deep learning framework for quick prototyping of structures of Recurrent Neural Networks, Sensory-fusion architectures, and deep learning on graph structured data. Built on top of Theano. ([Github URL](#))
- Brain4Cars**: A sensory-fusion architecture written in ROS for fusing multiple camera streams, tactile sensors, and GPS information. The package has integrated learning algorithms based on Bayesian networks for driver intention understanding and feature extraction modules. ([URL: Data set and Bayesian network code](#)) ([URL: Deep learning architecture](#))
- SPG-GMKL**: Multiple kernel learning (MKL) with a million kernels. This is a generalized MKL tool kit based on spectral projected gradient. It can be used for optimizing arbitrary combination of kernels and regularizer. ([Download page](#))
- Publications**
(in Chronological order)
- A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-RNN: Deep Learning on Spatio-Temporal Graphs
In CVPR, 2016 (**Full Oral**) (**Best Student Paper Award**)

A. Jain, H. S. Koppula, S. Soh, B. Raghavan, A. Singh, and A. Saxena. Brain4Cars: Car That Knows Before You Do via Sensory-Fusion Deep Learning Architecture
In IJRR, 2016 (In review)

A. Jain, A. Singh, H. S. Koppua, S. Soh, and A. Saxena. Recurrent Neural Networks for Driver Activity Anticipation via Sensory-Fusion Architecture.
In ICRA, 2016

A. Jain, H. S. Koppua, B. Raghavan, S. Soh, and A. Saxena. Car That Knows Before You Do: Anticipating Maneuvers via Learning Temporal Driving Models.
In ICCV, 2015

A. Jain and A. Saxena. Brain4Cars: Sensory-Fusion Recurrent Neural Models for Driver Activity Anticipation
In BayLearn Symposium, 2015 (**Full Oral**)

A. Saxena, A. Jain, O. Sener, A. Jami, D. K. Misra and H. S. Koppula. RoboBrain: Large-Scale Knowledge Engine for Robots.
In ISRR, 2015.

A. Jain, S. Sharma, T. Joachims and A. Saxena. Learning Preferences for Manipulation Tasks from Online Coactive Feedback.
In IJRR, 2015.

A. Jain, D. Das, J. K. Gupta and A. Saxena. PlanIt: A Crowdsourcing Approach for Learning to Plan Paths from Large Scale Preference Feedback.
In ICRA, 2015

H. S. Koppula, A. Jain and A. Saxena. Anticipatory Planning for Human-Robot Teams.
In ISER, 2014.

A. Jain, B. Wojcik, T. Joachims and A. Saxena. Learning Trajectory Preferences for Manipulators via Iterative Improvement.
In NIPS, 2013.

A. Jain, S. Sharma and A. Saxena. Beyond Geometric Path Planning: Learning Context-Driven Trajectory Preferences via Sub-optimal Feedback.
In ISRR, 2013.

A. Jain, S. V. N. Vishwanathan and M. Varma. SPG-GMKL: Generalized multiple kernel learning with a million kernels.
In SIGKDD, 2012.

Research in Popular Press

- Brain4Cars appeared on CNN technology front page and Discovery News. (2015)
- Interviewed by MIT Technology Review for Car that knows before the driver makes a mistake. (2015)
- RoboBrain covered by The New York Times, Wired, and MIT Technology Review. (2014)
- Interviewed by BBC World News prime time show *Click* for building robots for supermarkets. Broadcasted to millions of viewers. (2014)
- Research video on Teaching Robots from Human Signals received more than 100,000 hits on YouTube in less than 24 hours. Research covered by Discovery Channel, FOX

News, IEEE Spectrum, Techcrunch and many others. (2013)

- One of two students selected from IIT Delhi to present at the Indo-German Winter Academy on semiconductor device physics and fabrication. (2010)
- Director's Merit Award IIT Delhi, for highest semester grade point (2010)
- Summer Undergraduate Research Award, IIT Delhi, for embedded systems project on: Design and development of low cost optical densitometer for rural paramedics (2009)

**Projects
2012-Present**

Brain4Cars: Deep Learning and Perception for Smart Car Cabins (2014 - Present)

Project lead: <http://www.brain4cars.com>

- Integrated car cabin with an array of cabin sensors (cameras, tactile sensors, wearable devices, etc.) in order to extract valuable statistics about the driver.
- Developed sensory-fusion deep learning architecture for fusing information from multiple sensors in order to predict the driver's future maneuvers.

RoboBrain: Massive Knowledge Graph for Robots (2014 - Present)

Leading PhD student: <http://robobrain.me>

- Developed machine learning algorithms for knowledge graph combining multi-modal data from Natural Language, Vision, 3D-perception and Robot trajectories
- Conceptualized & built the distributed system architecture for managing multi-modal data, learning from crowd-sourcing feedback and visualizing the knowledge graph.
- Leading the project on major fronts: complete software stack, knowledge graph building, crowd-sourcing feedback and probabilistic beliefs on graph

PlanIt: A Crowd-sourcing Approach for Learning to Plan Paths

- Built a machine learning system to teach robots good trajectories from crowd-sourcing feedback. Project webpage: <http://planit.cs.cornell.edu>
- Modeled user feedback as a generative process with latent user intentions

Beyond Geometric path planning: Path planning in context

- Developed an algorithm to learn user preferences over robot trajectories from sub-optimal coactive feedback
- Deployed the algorithm on PR2 and Baxter robotic platforms and trained them for various household-chores and grocery checkout tasks
- Received substantial attention from popular press such as FOX News, IEEE Spectrum, Techcrunch, Discovery Channel and many others

SPG-GMKL: Multiple kernel learning (MKL) with Millions of kernels

- Developed MKL algorithm that scaled to Millions of kernels on UCI data sets
- The released software package is being used by many research groups

**Other
Activities**

Workshop Organizer for Learning Plans with Context from Human Signals, RSS 2014
Reviewer for NIPS, ICML, RSS, ICRA, AURO

Research talks

Keynote at the ICCV workshop on Autonomous driving. Title: Deep Learning for Spatio-Temporal Problems: On Cars, Humans, and Robots, Dec 2015

Invited talk at Zoox Labs, Dec 2015

Talk at University of Washington Seattle Department of Computer Science, Nov 2015

Invited talk at Qualcomm Deep Learning Research Center, Nov 2015

Oral at BayLearn 2015 on Brain4Cars: Sensory-fusion Recurrent Neural Networks

Invited Talk at RSS Workshop on Model Learning for Human-Robot Communication, July 2015

Invited Talk at ICML Workshop on Machine Learning for Interactive Systems, July 2015

Invited Talk at ICRA Tutorial on Planning, Control, and Sensing for Safe Human-Robot Interaction, May 2015

Invited Talk at IIT Kanpur Department of Computer Science. RoboBrain and Learning from Weak Signals, Feb 2015

Stanford Semantics and Geometry Seminar. RoboBrain and Learning from Weak Signals, Feb 2015

Stanford Robotics Seminar. Learning from Weak Signals, Nov 2014

Introductory talk at LPCHS workshop RSS 2014. Learning from Humans.

Cornell AI Seminar and ISRR 2013. Beyond Geometric Path Planning.

ICML Robot Learning workshop 2013.

Oral at SIGKDD 2012.

Invited spotlight at Mysore Park Workshop on Machine Learning 2012.

Lecture at Indo-German Winter Academy 2010.

Mentoring

Led more than 20 Master and Undergraduate students from Cornell and Stanford University. Building the teams for Brain4Cars and RoboBrain. Below are few outstanding students.

Vaibhav Aggarwal (won **Cornell ELI 2014 award** for his work on PlanIt)

Debarghya Das, 2014 - Full year (now at Google NYC)

Jayesh Gupta, Summer 2014 Intern from IIT Kanpur, India (now PhD at Stanford)

Shikhar Sharma, Summer 2013 Intern from IIT Kanpur, India (now PhD at U. of Toronto)

Brian Wojcik, Spring 2013 (now at Microsoft, Redmond)

Bharad Raghavan, 2014 - 2016, MS at Stanford

Shane Soh, 2014 - 2015, MS at Stanford

Avi Singh, 2015 - 2016, Intern from IIT Kanpur (now PhD at U. of Washington)

Siddhant Manocha, 2015 - 2016, Intern from IIT Kanpur (now at Google India)

Arpit Agarwal, 2015 - 2016, Intern from IIT Kanpur (now MS at CMU)