Nathan Ratliff

Curriculum Vitae

Phone: **(412) 983-1336**
Email: **ratliff.nathan@gmail.com**
Web: **www.nathanratliff.com**

6211 Elter Ct
Pittsburgh, PA 15209

Areas of Expertise: Machine learning, robotics, imitation learning, structured prediction, navigational planning, predictive behaviors, heuristic learning, footstep and grasp prediction, motion planning, trajectory optimization, LADAR classification, optical character recognition (OCR), semi-supervised learning.

Education:

Carnegie Mellon University, August 2004 – May 2009
Robotics Institute
Ph.D. awarded May 2009
Masters degree awarded May 2007
Thesis title: **Learning to Search: Structured Prediction Techniques for Imitation Learning**
Advised by Professor J. Andrew Bagnell

University of Washington, October 1999 – June 2003
Bachelor of Science in Computer Engineering
Bachelor of Science in Mathematics
Graduated with Honors
Advised by Professor Dieter Fox

Research Interests:

My work spans a range of applications at the intersection of machine learning and robotics, including **imitation learning, navigational planning, predictive behaviors, heuristic learning, footstep and grasp prediction, motion planning, LADAR classification, and optical character recognition**. My primary goal is to develop intelligent robot behavior, and accordingly my research explores novel theoretical frameworks leveraging tools from mathematical fields such as functional analysis and Riemannian geometry that inspire high performing and practical learning approaches tailored to the structural requirements of real-world robotic platforms. Reinforcement learning is the crème de la crème of robot learning, but it is necessarily slow--- discovering skills, policies, and techniques from scratch takes time for both humans and robots. But learning from others is powerful. Much of my work has revolved around furthering ideas in machine learning to develop novel approaches to learning from demonstration. This focus led me from theoretical and practical advances in max margin structured prediction, such as subgradient and functional gradient optimization approaches that greatly improve the efficacy and efficiency of the technique across a range of structured prediction problems, to new covariant functional gradient algorithms for robot trajectory optimization that place robot motion planning within a setting where we can leverage our imitation learning tools for this higher dimensional computationally complex planning environment. I solve real problems by analyzing the mathematical and theoretical foundations of complex real-world systems.

Key innovations:

1. **Maximum margin planning (MMP):** This framework forms the first general approach of inverse optimal control (IOC). It reduces IOC to convex formulations of structured prediction and results in efficient low-memory algorithms for imitation learning.
2. **LEArning to seaRCH (LEARCH):** This suite of learning techniques is a collection of nonlinear functional gradient algorithms for fast practical implementations of MMP. Within hours LEARCH learned a cost function for the Spinner/Crusher platform (a large scale outdoor robot at the National Robotics Engineering Center (NREC)) that rivaled the performance of the system's previous cost function that took a team of engineers thousands of man hours to meticulously hand tune.
3. **Subgradient methods for structured prediction:** This work generalizes LEARCH to a broader class of structured prediction problems. These methods are now used in state-of-the-art structured learning systems for applications including LADAR classification, navigational planning, and terrain classification. They simultaneously make training faster and more memory efficient by leveraging well engineered combinatorial inference algorithms in the inner loop. They demonstrably converge faster than competing approaches both in theory and in practice.

4. **Covariant Hamiltonian Optimization for Motion Planning (CHOMP):** This algorithm is the first trajectory optimization method to efficiently solve real-world motion planning problems in everyday settings. It effectively merges the planning and optimization stages of traditional systems (i.e. it can start from an initially infeasible trajectory and explicitly push the trajectory out collision during optimization) by leveraging covariant gradient and functional optimization theory from machine learning.
5. **Density based distances for semi-supervised learning:** We developed a new approach to semi-supervised learning that can be seen as a hard-max version of more traditional Laplacian-based expected path length methods on data graphs. This algorithm leverages a novel variant of Dijkstra's algorithm called Dijkstra* that integrates nearest neighbor computations with the shortest path search to avoid explicitly constructing a nearest neighbor graph in advanced with a predetermined number of neighbors. This algorithm solves semi-supervised learning problems in record time while performing shortest-path searches directly through the dense fully-connected data graph, thereby scaling beyond the previous computational limitations of Laplacian-based information flow approaches.

Research and Professional Experience:

Software Engineer – *Google, Pittsburgh*
April 2011 – June 2013

- *Developed large scale learning systems to assess landing page quality.*
- *Ensured smooth 24-7 operation of large scale systems.*

Research Scientist – *Intel Labs, Pittsburgh*
June 2010 – February 2011

- *Led effort aimed at integrating decision technology into Digital Signage.* Initiation of a multiyear effort to leverage decision making tools (e.g. bandit and reinforcement learning algorithms) to deliver customized content to the patrons in retail settings.
- *Theory of manifold projection for nonlinear machine learning.* General class of covariant algorithms building off functional gradient approaches to boosting. Reduce risk minimization to nonlinear least squares regression: utilize fast specialized regression packages to speed learning. Compact representation of nonlinear models (e.g. train finite sized ensemble). Inference efficient form of structured prediction.
- *Constraint manifolds for trajectory optimization.* Proved theorem relating nonlinear constrained CHOMP (covariant gradient trajectory optimization) updates to unconstrained projection. The theory led to a practical and efficient constrained trajectory optimization algorithm.

Research Assistant Professor – *Toyota Technological Institute in Chicago (TTIC)*
September 2009 – June 2010

- *Built wheeled robot for vision research and telepresence.* Researched component tradeoffs culminating in a system featuring an RC truck base, a 5-DOF neck built of Lynxmotion components for flexible manipulation of a webcam, and a control architecture consisting of an onboard laptop interfacing to a Phidget servo controller. Wrote all backend system control software for movement and driving, and implemented teleoperation using a combination of a Java client / onboard C++ server and Skype visualization.
- *Developed Dijkstra* algorithm for very fast semi-supervised learning on large data sets.* Dijkstra* is a lazy version of Dijkstra's algorithm that runs efficiently on graphs with very large degree. In this case, the graph is fully connected and the algorithm utilizes efficient space partitioning data structures in the inner loop to avoid explicitly representing the edges.
- *Lectured for TTIC course on machine learning* both in TTIJ (Nagoya, Japan) and remotely from Chicago: statistics, linear regression, neural networks.

Research Fellowship – *Intel Research*, Pittsburgh under Dr. Siddhartha Srinivasa and Dr. Dave Ferguson
May 2008 – September 2008

- *Covariant Hamiltonian Optimization for Motion Planning (CHOMP)*: Developed efficient covariant and functional gradient based trajectory optimization techniques for motion planning. Applications include robot manipulation, quadrupedal locomotion, simultaneous planning of navigational and manipulation behaviors.
- *Inverse Optimal Heuristic Control (IOHC)*: Developed imitation learning algorithm for training stochastic policies. Applications include taxicab route prediction and pedestrian prediction.

CMU, Robotics Institute, Admissions Committee
October 2006 – March 2008

- Reviewed applications for the Masters and PhD programs.
- Ranked candidates by suitability for the programs.

Research Fellowship – *Intel Research*, Pittsburgh under Dr. Siddhartha Srinivasa
May 2007 – June 2007, August 2007

- Applied newly developed imitation learning algorithm (LEARning to searCH) to learning grasp metrics.

- Initiated research into Rapidly-exploring Randomized Trees (RRT) for grasp planning.

Research Assistant – *Carnegie Mellon University*, Robotics Institute under Dr. J. Andrew Bagnell
August 2004 – May 2009

- *LEARning to searCH (LEARCH), MMPBoost*: Nonlinear generalizations of imitation learning algorithms (MMP) using (exponentiated) functional gradient techniques. Learns nonlinear features of the problem during optimization.
- *Maximum margin planning (MMP)*: Developing subgradient based structured prediction techniques for imitation learning with applications to autonomous mobile robot navigation, quadruped locomotion, grasp prediction, optical character recognition, LADAR classification.
- *Kernel Conjugate Gradient (KCG)*: Speeding up kernel machines by performing a functional form of conjugate gradient optimization directly in the reproducing kernel Hilbert space (RKHS) of the problem.

Teaching Assistant – *Carnegie Mellon University*, Robotics Institute
September 2006 – December 2006

- Select lectures for Professor J. Andrew Bagnell's course on Statistical Techniques in Robotics.
- Led review sessions, held regular office hours, and graded homework.

Software Development Engineer – *Amazon, Inc.*
June 2003 – June 2004

- Worked with team of developers on integration software for third-party merchants.
- Migrated merchant contact point SOAP server to robust third-party Systinet WASP Server for C++ web services framework.

Autonomous Robotic Systems Software Developer – *Intel Research (Seattle)*
July 2002 – September 2002

- Implemented object detection and a potential field based reactive navigation system for the ActivMedia Pioneer2-DXe mobile robot as part of the Intel Research / UW PlantCare project.
- Developed line finding and circle detection algorithms for laser scan vision processing.
- Extended an FLTK / OpenGL based laser range finder display system.

Autonomous Robotics Research Assistant – *University of Washington*, Computer Science and Engineering, under Dr. Dieter Fox
October 2001 – June 2002

- Designed high-level behavioral skills and implemented lower-level behavioral skills such as reactive mobile robot navigation.
- Implemented a dead-reckoning system used as input to a particle filter based localization module.

Teaching Assistant – *University of Washington*, Computer Science and Engineering
January 2002 – March 2002

- Taught a weekly recitation section to review lecture material, and held office hours for individual instruction.
- Proctored and graded midterm and final exams; graded homework.

Software Developer – *Applied Physics Laboratory,* University of Washington under Keith Kerr, MS
June 2000 – September 2001; October 2002 – June 2003

- Implemented multithreaded Java middleware for delivering scientific data.
- Developed a Java Web Start distributed High Seas Warning software package used by the Navy.

Publications:

N. Ratliff. Learning Geometric Reductions for Planning and Control. ICML Workshop on Robot Learning, June 2013.

M. Zucker, N. Ratliff, A. Dragan, M. Pivtoraiko, M. Klingensmith, C. Dellin, J. A. Bagnell, & S. Srinivasa. CHOMP: Covariant Hamiltonian Optimization for Motion Planning. International Journal of Robotics Research, May 2013.

A. Lavoie, M. Otey, N. Ratliff, & D. Sculley. History Dependent Domain Adaptation. Domain Adaptation Workshop at NIPS '11, December 2011.

A. Bijral, N. Ratliff, N. Srebro. Semi-supervised Learning with Density Based Distances. 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011), July, 2011.

A. Dragan, N. Ratliff. S. Srinivasa. Manipulation Planning: Extending Trajectory Optimization to Goal Sets. International Conference on Robotics and Automation (ICRA), Shanghai, China, May 2011 (*oral presentation*).

B. Ziebart, N. Ratliff, G. Gallagher, K. Peterson, J. A. Bagnell, M. Hebert, A. Dey, & S. Srinivasa. "Planning-based Prediction for Pedestrians." Proc. IEEE/RSJ International Conference on Intelligent RObots and Systems (IROS), St. Louis, Missouri, October 2009. (*oral presentation*)

N. Ratliff. "Learning to Search: Structured Prediction Techniques for Imitation Learning." Doctoral dissertation, Carnegie Mellon University, Robotics Institute, August 2009.

Y. Seo, N. Ratliff, & C. Urmson. "Self-supervised Aerial Image Analysis for Extracting Parking Lot Structure." International Joint Conference on Artificial Intelligence (IJCAI). Pasadena, California, July 2009. (*oral presentation*).

N. Ratliff, D. Silver, & J. A. Bagnell. "Learning to Search: Functional Gradient Techniques for Imitation Learning." Autonomous Robotics Special Issue on Robot Learning. Volume 27 Number 1, June 2009.

N. Ratliff, B. Ziebart, K. Peterson, J. A. Bagnell, M. Hebert, A. Dey, S. Srinivasa. "Inverse Optimal Heuristic Control for Imitation Learning." Artificial Intelligence and Statistics (AIStats). Clearwater Beach, Florida, March 2009. (*oral presentation*).

N. Ratliff, M. Zucker, J. A. Bagnell, & S. Srinivasa. "CHOMP: Gradient Optimization Techniques for Efficient Motion Planning." IEEE Int. Conference on Robotics and Automation (ICRA). Kobe, Japan, April 2009.

R. Diankov, N. Ratliff, D. Ferguson, S. Srinivasa, & J. Kuffner. "BiSpace Planning: Concurrent Multi-Space Exploration." Robotics: Science and Systems (RSS), Zurich, Switzerland, June 2008.

N. Ratliff, J. A. Bagnell & S. Srinivasa. "Imitation Learning for Locomotion and Manipulation", IEEE-RAS International Conference on Humanoid Robotics, Pittsburgh, PA, November 2007.

N. Ratliff, J.A. Bagnell & M. Zinkevich. "(Online) Subgradient Methods for Structured Prediction", Artificial Intelligence and Statistics (AIStats), San Juan, Puerto Rico, March 2007.

N. Ratliff & J.A. Bagnell. "Kernel Conjugate Gradient for Fast Kernel Machines", International Joint Conference on Artificial Intelligence, Hyderabad, India, January 2007. (*oral presentation*)

N. Ratliff, D. Bradley, J.A. Bagnell, J. Chestnutt. "Boosting Structured Prediction for Imitation Learning." Neural Information Processing Systems, Vancouver, B.C., Canada, December 2006. (*single track oral presentation, 1 of 25)*

N. Ratliff, J.A. Bagnell & M. Zinkevich. "Subgradient Methods for Maximum Margin Structured Learning", International Conference on Machine Learning Workshop on Learning in Structured Output Spaces, Pittsburgh, PA, June 2006.

N. Ratliff, J.A. Bagnell & M. Zinkevich. "Maximum Margin Planning", International Conference on Machine Learning, Pittsburgh, PA, June 2006. (*oral presentation*)

J.A. Bagnell, N. Ratliff & M. Zinkevich. "Maximum Margin Planning", Neural Information Processing Systems Workshop on Machine Learning Based Robotics in Unstructured Environments, 2005.

N. Ratliff and J.A. Bagnell. "Kernel Conjugate Gradient", Tech Report CMU-RI-TR-05-30, Robotics Institute, Carnegie Mellon University, June 2005.

N. Ratliff. "Autonomous Mobile Robot Reactive Approach Methodologies", Undergraduate senior thesis, University of Washington, Seattle, WA, June 2003

Recognition:

Intel Research Fellowship 2007-2008
Admissions Committee at Carnegie Mellon University, Robotics Institute 2007-2008
Regular conference/journal reviewer: ICML, NIPS, AAAI, IJCAI, IJRR, JAIR, AURO, IEEE journals
Microsoft Endowed Scholarship recipient (2001-2002 academic year)
Member of the Golden Key National Honor Society
Member of the National Society of Collegiate Scholars
Member of the Phi Eta Sigma National Honor Society

References:

Dr. Siddhartha Srinivasa
Intel Research, Pittsburgh
CM2
4720 Forbes Ave, Suite 410
Pittsburgh, PA 15213
(412) 297-4068
ss5@andrew.cmu.edu

Priya Narasimhan
Carnegie Mellon University
Dept. Electrical and Comp. Eng.
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-8801
priya@cs.cmu.edu

Prof. Martial Hebert
Carnegie Mellon University
Robotics Institute
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-2585
hebert@ri.cmu.edu

Prof. Geoffrey Gordon
Carnegie Mellon University
Machine Learning Dept.
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-7399
ggordon@cs.cmu.edu

Prof. Drew Bagnell
Carnegie Mellon University
Robotics Institute,
Machine Learning Dept.
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 681-8669
dbagnell@ri.cmu.edu

Prof. Andrew Ng
Computer Science Dept.
Stanford University Room 156,
Gates Building 1A
Stanford, CA 94305-9010
(650) 725-2593
ang@cs.stanford.edu