

UCL Gatsby Unit – Statement of Research Interests

My research goals are oriented towards the intersection of deep reinforcement learning, inference and decision making under uncertainty, and computational neuroscience that can have implications towards game playing and developing novel reinforcement learning(RL) models. I am interested to explore how inference in the brain works for real time decision making under uncertainty, and how can robots optimize its task performance inspired from computational models of the human brain. I believe, research towards computational neuroscience and reinforcement learning can have impacts towards artificial intelligence research. Working under Professor Peter Dayan and Professor Maneesh Sahani, along with external supervision from Dr. David Silver and my previous supervisor Professor John Shawe-Taylor from UCL CSML, I believe the PhD program at the Gatsby Computational Neuroscience Unit can therefore flourish my research objectives and interests.

Through the PhD program, I am particularly drawn towards research in how deep reinforcement learning, combined with techniques from Bayesian inference, probabilistic programming and uncertainty in models can progress research towards language processing tasks. Currently, RL based policy search techniques are applied for attention mechanisms in computer vision research. I believe, such techniques can further be used for conversational modeling and neural machine translation related large-scale language processing tasks. I believe, through unsupervised learning of large state representations for high dimensional inputs, RL algorithms can be scaled up efficiently. By understanding the decision making model in the human brain, and solving questions towards how the brain computes inference to anticipate upcoming events, can make breakthroughs in research at the intersection of Bayesian inference and deep reinforcement learning. This can further have applications towards RL robotic agents being able to play football or cricket, where such games requires humans being able to predict or anticipate the movement of the ball.

Furthermore, recent work on neural machine translation for conversational modeling, such as being able to interpret stories from movies in a human-like manner, can have applications towards reinforcement learning for optimizing robotic task performance. This would mean AI agents being able to communicate in multiple human languages to learn to improve their behavior policy based on social interactions. Additionally, I am particularly drawn towards research in understanding how the cognitive and inference capabilities in the human brain works, to further develop novel reinforcement learning and decision making computational models for robotics and human computer interaction.

Having taken few online machine learning courses, and curious to explore the field in an immersive setting, after my second year of undergraduate studies, I worked as a research student at the Machine Learning group at Johns Hopkins University (JHU). I worked on a summer project aimed at developing an intelligent healthcare system under the supervision of Professor Suchi Saria. My project focused on developing and implementing a cost sensitive tree of classifiers model that can be applied to large scale ICU patient data for classifying patients with septic shock. I worked towards implementing the decision tree classifier model that would extract features at lowest cost, and would subgroup the patient population data along each branch of the tree based on the symptoms and types of medical test (features).

Following my summer experience, in my final undergraduate year I took more than half of my courses from the graduate level MSc Machine Learning program at UCL. Having taken separate courses on Graphical Models, Reinforcement Learning and fundamental Supervised Learning, I got motivated to further broaden my interests towards recent research advances which encouraged me to regularly attend research talks and PhD reading groups in machine learning at UCL. After attending a talk from Dr. Alex Graves on automatic handwriting generation and neural Turing machines, and from Dr. Demis Hassabis at Google DeepMind on playing Atari games using deep

reinforcement learning, I got interested towards developing research skills in such areas, questioning and solving underlying research problems, which further convinced me to pursue an academic career with the goal to become a Faculty member in the future.

I further got motivated to do my undergraduate thesis in reinforcement learning on convergence of deterministic policy gradient algorithms. My thesis on improving convergence of deterministic policy gradients was supervised by Professor John Shawe-Taylor (UCL Center for Computational Statistics and Machine Learning) and co-supervised by Professor Miguel Rodrigues (UCL EEE) in affiliation with Dr. David Silver based on his recent work at Google DeepMind. I worked towards developing and implementing both stochastic and deterministic policy gradient algorithms on several benchmark RL tasks to analyze convergence rates. I worked towards adaptive learning rates based on recent work from Dr. Tom Schaul (Google DeepMind) and derived approximate Hessians of both stochastic and deterministic gradients to study convergence rates and global optimal convergence of policy gradient algorithms. Results from my work showed that using adaptive learning rates in RL settings, we can ensure elimination of fine-tuning and achieve faster and better local optimal convergence on benchmark RL MDPs; while existence of second order approximate Hessian theorems for both stochastic and deterministic gradients in my work.

In the summer of 2015, I worked in the summer undergraduate research fellowship (SURF) program at Caltech, under the supervision of Professor Richard Murray at Caltech Computing and Mathematical Sciences, Control and Dynamical Systems Lab. My work was part of a larger aim, in collaboration between Caltech, NASA JPL and MIT, towards developing resilient spacecraft executive software architecture, such that Mars Rovers can perform robotic tasks in space taking exploration risks into account. I worked towards integrating real time dynamic mapping capabilities into the popular Pioneer 3-DX simulation robot based on lidar sensors and integrated obstacle avoidance and path planning algorithms into the software architecture. I particularly enjoyed working on such a large research project at Caltech, where Dr. Michel Ingham and Dr. Tara Estlin from NASA JPL, Robotic Systems Estimation, Decision and Control group also supervised my work.

Currently, I am pursuing a MPhil in Machine Learning at University of Cambridge, St John's College, working under the supervision of Professor Zoubin Ghahramani at the Computational and Biological Learning lab. Besides taking regular courses with focus in Bayesian modeling and machine learning, along with practical emphasize on Gaussian Processes, I am further taking courses in speech recognition, language processing and computer vision. Additionally, for my Masters project, I would be working in areas of modeling uncertainty in deep learning, and how uncertainty can account for adversarial examples that can cause deep neural networks to make mistakes with high confidence. My work would additionally focus towards how the top layer in deep networks can be replaced with Gaussian Processes to perform prediction with confidence, taking uncertainty into account.

Gatsby Unit at UCL, having a lot of faculties working towards AI research would mean that, being a graduate student, I would be able to discuss and communicate my research interests effectively with other labs for collaboration. I believe my research interests, combined with my past relevant experiences and my enthusiasm to work with Professor Peter Dayan, along with external supervision from Dr. David Silver, would provide me the ideal platform to further nurture my interests during the PhD program.