Carnegie Mellon University: Statement of Purpose
Name: Riashat Islam UserID: riashat.islam.93@gmail.com

My research goals are oriented at the intersection of Bayesian reasoning and deep reinforcement learning for sequential decision-making and attention-based mechanisms. With interests in approximate inference, deep learning and reinforcement learning, I am interested to explore how such approaches can aid towards data efficient learning for robotics, scene understanding and image or video caption generation tasks. My research aims are towards understanding how memory networks, such as long short-term memory and recurrent neural networks can be applied to reinforcement and transfer learning. I believe my research objectives and interests can flourish intellectually by joining the Machine Learning PhD program, in the Machine Learning Department at Carnegie Mellon University, working under the supervision of Professor Ruslan Salakhutdinov, Professor Emma Brunskill or Professor Eric Xing. I am also interested to work in collaboration with the CMU Robotics Institute working under supervision of Professor Maxim Likhachev and Professor Jeff Schneider. Based on my recent conversations with the above supervisors of interest via email, I understand my research interests are relevant to future directions of work in their labs.

In particular, I am drawn towards understanding how recent achievements in deep reinforcement learning can be applied to video scene understanding, question answering or neural machine translation models in natural language processing and computer vision tasks, for sequentially understanding document paragraphs or video scenes. Furthermore, recent work led by Professor Ruslan Salakhutdinov on image or video caption generation for movie interpretation, or inferring from scenes, can be extended by using novel attention based recurrent network models that uses policy gradient reinforcement learning techniques. I also got fascinated by recent work on deep neural machine translation for conversational modeling, such as being able to interpret stories from movies in a human-like manner, can also have further 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, towards the goal of learning multiple complex high dimensional policies at the same time.

Additionally, based on my recent interests on guided policy search approaches in deep reinforcement learning, I am interested to investigate how such approaches can be data-efficiently applied on a larger scale in domains such as intelligent tutoring systems or robotics that relies on user interactions and learns optimal behavior policy accordingly. Google DeepMind has shown results of playing Atari games using Deep Q-Learning techniques for continuous control. However, such approaches to game playing RL are often not scalable to practical robotic tasks due to the inefficiency and dependency in large of training examples for learning. For my research, I am considering approaches towards how, for example, using approximate inference algorithms (such as stochastic expectation propagation algorithms) or control as inference problem, deep reinforcement learning can be scaled up efficiently for practical tasks. This might also involve learning predictive models of the non-linear dynamics using Bayesian approaches, while inference in deep RL can solve problems towards ensuring exploitation and exploration.

After my second year of undergraduate studies, I worked as a summer research student in the Machine Learning group at Johns Hopkins University (JHU), under supervision of Professor Suchi Saria. The overall goal of research was towards an intelligent healthcare system that can suggest the medical tests to take for the type of disease, based on symptoms when visiting doctors. 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. Furthermore, I worked on implementing a cost-sensitive 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).

In my final undergraduate year I took almost all 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. 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 and co-supervised by Professor Miguel Rodrigues in collaboration with Dr. David Silver and Dr. Guy Lever based on their 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 tasks. Theoretical proof of Hessian of deterministic gradient showed that the second order approximations in model-free settings are in fact dependent on model dynamics, which left room for future work (and delayed submission to ICML EWRL workshop in 2015).

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 Jet Propulsion Lab (JPL) and MIT, towards developing a 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 laser sensors and integrated obstacle avoidance and path planning algorithms into the risk-aware software architecture. Additionally, the project was in collaboration with NASA JPL, where Dr. Michel Ingham and Dr. Tara Estlin (JPL Robotic Systems Estimation, Decision and Control group) further supervised my work.

Currently, in the taught Masters (MPhil) program in machine learning and language processing at University of Cambridge, I am taking courses focused towards approximate inference and Bayesian reasoning. I am participating in three Kaggle competitions for classification, regression and density modeling tasks, where I am investigating the effectiveness of deep neural networks pre-trained with autoencoders and restricted Boltzmann machines using Dropout. Furthermore, using the HTK toolkit for speech recognition, I am working towards both approaches of using DNNs and GMMs at acoustic modeling for speech recognition to compare performance. As part of my Masters thesis, under supervision of Professor Zoubin Ghahramani, I would be working towards using approximate inference based approaches for trajectory optimization in RL, combined with guided policy search methods (based on recent work from Dr. Sergey Levine) for optimal policy search based on optimized trajectories using inference. My research focus is towards understanding how Bayesian reasoning for inference and decision making, combined with deep learning, can provide a key tool for data efficient learning while quantifying uncertainty in exploration of RL agents. Furthermore, I am interested to focus my research towards using memory-based inference and scalable probabilistic decision-making for policy search, aimed at research at the intersection of variational and approximate inference and deep reinforcement learning, using the MuJoCo simulator and Atari games as experimental framework. At Cambridge, I am also working towards DNNs modeled as approximate Gaussian Processes, incorporating uncertainty in deep learning, towards the goal of adversarial training that has gained recent attention in deep learning.

I have already discussed my research interests with Professor Emma Brunskill (via email) regarding my interests in deep reinforcement learning. Furthermore, based on my recent interests in memory based deep learning models for video scene interpretation and generating images from captions, I am also interested to work towards memory based deep learning models under supervision of Professor Ruslan Salakhutdinov. Furthermore, having contacted Professor Maxim Likhachev regarding my research interests, I understand I can also work under his supervision in model-based planning combined with reinforcement learning approaches. I believe I can further nurture my research interests by joining the PhD program at Carnegie Mellon University, working towards practical applications of deep reinforcement learning for computer vision or natural language processing and robotics tasks.