Statement of Purpose: Stanford Computer Science PhD program
Riashat Islam

My research goals are oriented towards scalable deep learning, Bayesian reasoning, approximate inference, and deep reinforcement learning that can be extended for sequential decision making and attention based mechanisms for video scene understanding and question answering tasks. Stanford, being pioneers of large-scale deep learning (based on work from Professor Andrew Ng) has made major contributions in computer vision and NLP using large-scale deep network architectures. I am interested to explore how memory-based inference and memory networks can aid towards computers learning complex arithmetic manipulations based on sequence learning. My research aims are towards understanding how memory networks, such as long short term and recurrent neural networks can further be applied to deep reinforcement learning algorithms for NLP and vision tasks. This further depends on understanding the fundamental questions of how inference and decision-making process in the human brain works, and how humans learn to recall from their previous observations. I believe my research objectives and interests can flourish intellectually through the Computer Science PhD program at Stanford University, working under the supervision of Professor Chris Manning or Professor Fei Fei Li at the NLP and Vision labs respectively. During the last summer, I visited both the Vision and NLP labs to meet Justin Johnson (PhD student in Vision Lab) and Thang Luong Minh and Panupong Pasupat (NLP lab), and had also been in contact with Professor Li and Professor Manning regarding my research interests (via email).

In particular, I am drawn towards understanding how recent achievements in deep reinforcement learning can be applied to question answering or neural machine translation models in natural language processing, for sequentially understanding which part of the text to focus attention on to be able to understand the entire conversations or paragraphs. Furthermore, recent work on image or video caption generation for movie scene interpretation or inference from scenes can further be improved by using novel attention based recurrent network models that uses policy gradient reinforcement learning techniques. 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 further have applications towards reinforcement learning for optimizing robotic task performance. This would mean AI agents being able to communicate in 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.

For my research experience, after my second year of undergraduate studies, I worked as a summer research student at 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. 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. During my summer research, I am partially worked towards Bayesian time series modeling with other interns that fascinated me to pursue related research in machine learning.

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. 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 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 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 while left room for future work (and delayed submission to 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 the 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 laser sensors and integrated obstacle avoidance and path planning algorithms into the 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 also working with both approaches of using DNNs and GMMs for acoustic modeling in speech recognition tasks, using GPUs as computational resource. As part of my Masters thesis at Cambridge, I will be working in the Cambridge Machine Learning group under supervision of Professor Zoubin Ghahramani (alongside PhD student Shane Gu). My Masters thesis research would focus towards adversarial training of deep networks, and how modeling uncertainty in DNNs using Dropouts as a Bayesian approximation, can aid towards having predictive distributions as outputs to take account of adversarial examples. I would also be working towards understanding whether different Bayesian neural networks with smoothness prior can make the predictions more calibrated towards the goal of "calibrated deep learning". Furthermore, on a different direction of research, I would also 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 from UC Berkeley) 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.

At Stanford, I am interested to work in Professor Fei Fei Li's Computer Vision Lab. In addition to recent work on image or video scene understanding and caption generation, I am interested in working towards sequential attention based models using policy gradient RL techniques, for the agent to learn where to look in the scene. Deep reinforcement learning with recurrent neural network attention based models can make major contributions towards scene understanding – I am interested in developing computational models based on how the human perception works by selectively focusing attention and combining information from different fixations over time. I would also be motivated to work in the Natural Language Processing group led by Professor Dan Jurafsky and Professor Chris Manning. Having talked to Professor Manning's current PhD students, I understand my research experience in deep reinforcement learning would be of interest in the group. Deep learning models have made major contributions towards NLP tasks for question answering and using dynamic memory networks to be able to understand meanings of texts, and reason over relevant facts. The ability to understand texts from scratch, combined with read and write capabilities in memory networks would mean AI agents being able to write and understand computer program codes by itself.

Stanford has been a pioneer in machine learning research, with previous faculties such as Professor Andrew Ng and Professor Daphne Koller, working at the intersection of AI, robotics and machine learning. I believe my research interests, combined with my past relevant experiences and my enthusiasm to work with Professor Chris Manning and Professsor Fei Fei Li would provide me the ideal platform to further nurture my interests during the PhD Computer Science program at Stanford University.