**Academic Statement of Purpose**
**Riashat Islam**
**Computer Science and Engineering PhD Program**

My research goals are oriented towards machine learning, at the intersection of deep learning and reinforcement learning (RL) for data efficient robotics and game playing tasks, for making intelligent agents learn to perform complex group tasks. I am primarily interested in how scalable RL approaches combined with feature learning, controller learning and probabilistic models of control dynamics can be learnt efficiently for practical applications based on high dimensional observations. Additionally, I am interested to pursue research towards scalable Bayesian reasoning, approximate inference and deep reinforcement learning (using either neural networks or Gaussian Processes) such that it can be extended towards sequential decision-making and attention based mechanisms such as active sensing. I am fascinated to understand how memory networks, such as long short-term memory and recurrent neural networks can be applied to RL tasks. My research goals depend 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. Deep reinforcement learning research can solve questions of how true artificial intelligence can be achieved through robotic agents remembering from its memory component to improve on their tasks. My research objectives and can flourish intellectually through the CSE PhD program at University of Michigan, Ann Arbor, where I am interested to work in the Artificial Intelligence Laboratory under Professor Satinder Singh Baveja and Professor Honglak Lee.

Additionally, 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 interpretation, or inferring from scenes can further be improved 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. Based on recent work on guided policy search (Professor Pieter Abbeel and Dr. Sergey Levine from UC Berkeley) for deep RL for robotics tasks, I am interested to investigate how such approaches can be data-efficiently applied on a larger scale in other domains such as intelligent tutoring systems that relies on user feedback 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 not often 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 approaches of how, for example, using approximate inference algorithms (such as message passing algorithms), deep reinforcement learning can be scaled up efficiently for practical tasks. After completing my PhD, I hope to join as a Faculty working towards similar directions of reinforcement learning with diverse applications.

For my research experiences, 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 (features).

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 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 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 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 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 lidar sensors and integrated obstacle avoidance and path planning algorithms into the software architecture. Additionally, the project was in collaboration with NASA Jet Propulsion Lab (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 tool for data efficient learning while quantifying uncertainty in exploration.

At the University of Michigan, I am keen to work under the supervision of Professor Satinder Singh Baveja who had been a pioneer working in reinforcement learning. I hope to explore how deep learning for vision and NLP can be used towards collaborative reinforcement learning, in which agents can learn by asking natural language questions, or through observations by being able to infer and interpret the scenes and actions in the observations. Such research can have contributions towards RL for cooperative tasks under a multi-agent RL, or for human robot interaction, such as the example of having an online tutoring RL agent for teaching. Deep neural networks in RL can also be used for problems towards inverse reinforcement learning, where the agent can learn from demonstrations – where essentially, deep architectures can be used to approximate the reward function. Having been in contact with both Professor Singh and Professor Lee at Michigan, I understand my research interests in deep RL are relevant to future directions of work in their respective labs. I have also sent a draft of my research proposal towards data-efficient RL approach to both Professor Singh and Professor Lee.

I believe my research interest combined with my past relevant experiences and my enthusiasm to work with Professor Satinder Singh Baveja and Professor Honglak Lee would provide me the ideal platform to further nurture my interests working in deep learning for reinforcement learning during the Computer Science and Engineering PhD program at University of Michigan, Ann Arbor working towards deep reinforcement learning in the Artificial Intelligence lab.