The University of Edinburgh

Paul Patras
12/14/2016

University Partner:
The University of Edinburgh

Academic Supervisor - Name and Email Address:
Dr Paul Patras, ppatras@inf.ed.ac.uk

Suggested Project Title:
Deep Reinforcement Learning for Dynamic Network Control

Suggested Project Summary:
Research Theme: Numerous real-world network control problems can be modelled as Markov Decision Processes, whereby agents select from a set of actions to interact with the target system. Through these actions, agents change the system's state and receive feedback via some reward functions. For instance, the volume of user traffic and the service demand for computing resources at different time intervals could be seen as possible states of a mobile network and respectively a data centre. Engineers will employ specific tools to optimise certain performance metrics, such as application latency or energy consumption, which can be regarded as reward function instances. A powerful approach to maximising return is Q-learning, which does not require a model of the system in order to find an optimal action policy. The action-state space is however often too large to enumerate and impossible to store in memory. Potential exists to overcome these limitations through deep learning models that approximate Q functions, as demonstrated by recent results in robotics and computer-based board games (e.g. AlphaGo).

Deep Q-Networks (DQN) remain unexplored in the computer, data centre, and mobile networking domains, and will lie at the core of the proposed research. We will design deep learning models for the purpose of balancing mobile/data centre network resources and optimise their performance (in terms of memory, computational speed, throughput, energy, etc.), train the models offline on the Informatics GPU cluster, and apply the techniques developed on real "machines". As traditional Q-learning is known to exhibit slow convergence rates, we will work on designing mechanisms that converge at speeds comparable with those of model-based reinforcement learning methods.

Outcomes: The project will develop a set of DQN algorithms for network control problems, including elastic computing and dynamic allocation of resources in mobile broadband networks. The results obtained will be disseminated in top-tier academic conferences/journals and the code produced patented in collaboration with the industry partner(s) or released into the public domain.

Business Domain: This project should be particularly relevant to companies that

manipulate large volumes of data, who wish to forecast service demand, and aim to dynamically optimise network resources. These may include Amazon, Akamai, British Telecom, and EE.

Innovation Potential: GPU technology and parallel computation techniques enable training neural networks at unprecedented speeds. As deep reinforcement learning is a young area, this project has potential to advance DQN for network control in hyper-connected digital infrastructure.

Data Science Research: The project will fuse novel machine learning techniques, real-time data stream processing, and computer network orchestration protocols. Thus it covers many facets of data science.

Impact for Scotland: The planned research could promote Scotland both as a leader in academia and as a fertile scene for industrial innovation. In particular, we expect i) the results of this project and presentations at reputed fora (e.g. NIPS, NSDI) will attract scientific citations and build esteem, and ii) the technical achievements of the project will attract new talent and stimulate emerging startups.

Collaboration Sought for the Project:
We are seeking an industrial partner that processes vast amounts of information in real-time and who owns their own production network and/or data centre infrastructure. In the absence of realistic data sets such infrastructure routinely gathers, the progress of this project would be limited to synthetic scenarios that may not capture real-life situations. In addition, validating the performance of the proposed DQN models for network control, would only be possible if experimenting on (parts of) a real deployment.

Published or Private?:
Yes