Project 8:
Learning to act with recurrent neural networks

Team members: Natalia Díaz Rodríguez, Yad Faeq, Bogdana Rakova, and others

PRESENTATION BY YAD FAEQ

https://youtu.be/l-V5Iy2oSyg

PROPOSAL

In the general reinforcement learning setting, an agent needs to learn an optimal policy for achieving goals in an environment, while general purpose agents need to exhibit this ability across a diverse range of environments. Our aim is to design a recurrent neural network architecture that uses a hierarchy of LSTM units, an external memory, and learned compositions of modules to achieve transfer learning and avoid catastrophic forgetting. This functionality could be applied to the design and control of nanoscale systems.

Transcript

As the title says, we are learning to act with recurrent neural networks. In the current space for reinforcement learning, an agent learns to adapt to a cyclic environment. So, we have an agent that learns to adapt to a specific task that needs to be achieved. However, a general purpose agent, in most cases cases encountered, needs to generalize long enough to learn over different environments. Our aim is to design an LSTM network, an architecture that avoids catastrophic forgetting, and it uses transfer learning to adapt to different environments. For example, an agent that will learn to adapt to one environment will learn to control, for example, a macroscopic scanning control agent. We want to design an architecture that learns to adapt to the different environments, it will rebalance inputs, and achieve designated tasks using transfer learning.

When you are in a lab, some of the experiments can be very expensive. If you miss something during scanning, this agent will assist the person doing the scanning to avoid such errors. This will save time and the cost of replicating the experiment. Other things that would be nice to have, since are mission learning practitioners, would be others with additional expertise, to look over the real world applications of what we are building, and along the lines of building such products, one of the most important things is the actual GPU, the actual hardware that we use, so we are looking forward in terms of funding to have such hardware provided.

.

References:

"Progressive Neural Networks" AA Rusu, NC Rabinowitz, G Desjardins, H Soyer, J Kirkpatrick, K Kavukcuoglu, R Pascanu, R Hadsell. Google DeepMind, London, UK. (Sep 2016; this version, v3). https://arxiv.org/abs/1606.04671

"Reuse of Neural Modules for General Video Game Playing" A Braylan, M Hollenbeck, E Meyerson, R Miikkulainen. Department of Computer Science, The University of Texas at Austin. (Submitted on 4 Dec 2015). https://arxiv.org/abs/1512.01537

"On Learning to Think: Algorithmic Information Theory for Novel Combinations of Reinforcement Learning Controllers and Recurrent Neural World Models" Juergen Schmidhuber. The Swiss AI Lab. (Submitted on 30 Nov 2015)
https://arxiv.org/abs/1511.09249

"Learning to Communicate to Solve Riddles with Deep Distributed Recurrent Q-Networks" JN Foerster, YM Assael, N de Freitas, S Whiteson. University of Oxford, United Kingdom; Canadian Institute for Advanced Research, CIFAR NCAP Program, Google DeepMind. (Submitted on 8 Feb 2016). https://arxiv.org/abs/1602.02672

Data & Resources

OpenAI gym for data - https://gym.openai.com/docs

"OpenAI Gym is a toolkit for developing and comparing reinforcement learning algorithms. It makes no assumptions about the structure of your agent, and is compatible with any numerical computation library, such as TensorFlow or Theano. You can use it from Python code, and soon from other languages."

Resources: GPU instances from EC2 for research, Titan X GPU

https://www.google.com/#q=GPU+instances+from+EC2+for+research%2C+Titan+X+GPU