

## **Deep Learning Workshop, Edinburgh, 6 May 2014.**

**Organiser: Amos Storkey**

### **Background**

With the recent purchase of DeepMind Technologies by Google, deep learning methods are getting increased attention in many spheres. At the same time a number of members of the SICSA community have been developing or working with deep learning methods for some time.

Deep Learning is a transformative modelling technology in machine learning. Historically, it was well understood that the process of machine learning involved deciding on a representation, and then applying machine learning tools. The process of deciding a representation was generally manual, involving feature construction, variable selection, data reduction, variable rescaling and transformation, and feature selection. The reality of machine learning was that the real performance of a system was more determined by these manual representation steps than by the machine learning methods that were subsequently applied.

Hence, the issue of automated learning of representations was seen as a key process for improving the power and capability of machine learning methods. Deep Learning methods have come to the fore in this process by providing automated unsupervised methods for representation learning. These unsupervised methods work via iterated layerwise refinement of the previous representations to capture multiscale features that are useful across many tasks. This process of using further and further levels of abstraction, while controlling for abstractions that explain the data is what makes deep learning methods so versatile.

### **Workshop**

On the 6 May 2014 we ran a workshop on Deep Learning to bring together those who are working in Machine Learning, Statistics, and the wider community, who are working on, using or interested in Deep Machine Learning Methods. This workshop was designed to be a benefit to those currently utilising or developing such methods, those wanting to know what can be done and understand some of the practical issues hidden behind the hype, or those who are just interested in finding out more. It was also to be of interest to those in early stage businesses who wish to harness this exciting technology.

The workshop was in the School of Informatics, University of Edinburgh. The outline of the event was

- 08.45 Welcome and Coffee
- 09.00 Introductory Talk
- 09.20 Carina Silberer: Learning Grounded Meaning Representations with Autoencoders
- 09.50 Callum Main: Deep Learning for State Tracking in Spoken Dialogue Systems (work with Verena Rieser and Zhoran Wang)
- 10.20 Coffee, Discussion and Ideas Factory
- 11.00 Invited Speaker: Ruslan Salakhutdinov (University of Toronto, confirmed)
- 12.00 Catered Lunch.
- 13:00 Krzysztof Geras: Scheduled denoising autoencoders (work with Charles Sutton)
- 13:30 Benigno Uribe: A Deep and Tractable Density Estimator (work with Iain Murray and Hugo Larochelle)
- 14:00 Coffee, Panel Discussion Period
- 15:00 Invited Speaker: Volodymyr Mnih (Google DeepMind, confirmed)

- 16:00 David Hofmann: Reliable pattern recognition for processing of myoelectric signals
- 16:30 Brief Break.
- 16:40 Spotlights and Open Mike Session
- 17:00ish Summary, and Directions
- 17:30 Evening pub trip for those who don't need to leave

## **Outcome**

The workshop had 170 people attending at least one part of the workshop, and about 140 attendees for the workshop as a whole. There was a considerable split across SICSA and Scottish institutions, with attendees from Aberdeen, Dundee, Edinburgh, Glasgow, Glasgow Caledonian, Heriot Watt, Napier, Robert Gordon, St Andrews, Stirling. Altogether approximately 40 of the full workshop attendees were from Universities outside the University of Edinburgh (the host institution), with another 20 from companies or other institutions. The registration was full two weeks before the event, (the room has a capacity of 180, and we hit our cost limits), and unfortunately there were others who wanted to attend who were not able to. A number of the visitors were international, some coming from far afield just for the workshop.

Feedback from people at the workshop was overwhelmingly positive, with requests that the event be repeated. The format worked well, though with some dropoff in attendance after the second invited talk, which we might mitigate against for future events. The invited talks were very well received, and proved to be good and complementary choices.

## **Follow up**

We will discuss a number of follow up themes, that we will contact workshop attendees about. One specific idea is to set a SICSA wide data analytics competition (probably via Kaggle), that will require deep learning techniques. This will allow collaborative teams to form to enable SICSA members who are new to this area to begin to get some practical experience using these methods, with some forum based feedback.

## **Costs**

The total cost estimate was £1600 for the catering, and about £2000 for the visiting speakers (one speaker was coming from Toronto, the other via Copenhagen). SICSA covered £1100 of those costs, with the School of Informatics contributing the remaining amount (approx £2500).

## **Comments**

This workshop was a workshop in a specific technical area of machine learning, with a high attendance. I believe this is indicative of the demand for data science related events from across SICSA. It appears that well targeted workshops on specific but important research subjects can attract a substantial audience and be beneficial to the research community. We will use this momentum as part of the new SICSA data science theme, but also believe the response to this workshop is indicative of the importance of this area across SICSA.