

**Report on SICSA AI Theme meeting, 30 August 2017, University of Stirling.**

Written up by Prof Leslie Smith, 15 September 2017.

The meeting ran from 10.00 to about 16.30, and was a mix of talks and discussion groups, all oriented around the theme of *How do the many facets of learning in AI relate to each other?* The programme is enclosed.

Statistics: there were 37 attendees, 32 from Scottish Universities (Abertay: 1, Heriot-Watt: 3, Edinburgh Napier: 7, Edinburgh University: 5, Glasgow University: 1, Robert Gordon's University: 5, Stirling University: 10), and 5 from an Edinburgh company, Codeplay.

There were six talks, organized into two sets of three, with Emma Hart chairing the first session and Leslie Smith chairing the second one.

Peter Ross reviewed AI techniques, noting that the learning so far was very much machine training, rather than machine learning. Aims were externally set, and data for training provided, rather than the machine choosing the goals, or finding a deeper understanding of the system. Such systems can also make surprising mistakes, particularly near boundaries of manifolds in high dimensional spaces. We need to think about 0-shot and 1-shot learning as well, and Ross suggested that the use of episodic memory (introspection) might underlie this.

David King discussed the role of learning in games. Initially, games used relatively simple techniques for controlling their actions, FSTs, Decision trees, etc., but in more recent times they are using Gas, n-gram prediction, CBR, and NNs, (the example game used was "Civilization"). The problems in games are like those of the real world, for example, the stability-plasticity dilemma, and the necessity for real-time behaviour makes accuracy not always the best metric of performance. Modular systems seem to be a good way forward.

Iain Murray reviewed the different types of learning used in systems, particularly the ways in which certain companies (e.g. youtube) use learning to recognize features. He discussed the use of self-organised learning as opposed to pure training, but pointed out that there are many problems that simply are not yet solvable within the frameworks of learning currently in use.

Alan Bundy reported on the use of learning in the correction of faulty logical theories. This is a different form of learning from (D)NNs, and is about axiom (belief) revision. These need revised in the light of new evidence, and often take the form of splitting concepts, based on analysis of failure. And this is definitely learning, as opposed to training.

Nick Taylor briefly described the Edinburgh Centre for Robotics (joint between Heriot-Watt and the University of Edinburgh), and its large number of postgraduate students. He then discussed the stages of learning from a Piaget-ian perspective. (sensorimotor/preoperational/concrete operational/formal operational), and tried to put them in the context of robotics. Noted the different types of tasks, some of which are much harder for a robot (go and fill the coal scuttle). He then gave a snapshot of the research at

the Edinburgh Centre for Robotics.

Emma Hart discussed lifelong learning in the context of artificial immune systems. Immunology provides inspiration for building systems that a) continually improve over time b) selectively transfer knowledge from a learned problem to a new problem c) generates novel methods of solving unseen problems. A system for solving continual stream of combinatorial optimization problems that improves with experience and solves instances with diverse characteristics was discussed (NELLI)

The presentations were followed by lunch, at which the delegates could also examine the posters. There were only a small number of posters (4).

Adair: Evolving training sets for improved transfer learning in brain computer interfaces,  
Graham: Contextual Learning and Processing in Humans and Machines,  
Swingler: A Quick Introduction to Mixed Order Hyper Networks, and  
Taou, Corne and Lones: Evolving Boolean networks for the control of gene regulatory networks.

This was followed by the discussion sessions. The discussion topics were:

1. Explain-ability: why does it matter? How might it be achieved.
2. How can we join symbolic and other techniques (Bayesian, statistical, neural) together?
3. What *is* learning (as opposed to training, adaptation, development)?

The original intention had been to have each group discuss two topics, but on the day, each discussion group discussed one topic, and reported back.

There was a rather good and wide ranging final discussion covering the following topics:

- Knowledge representation: something that's been of importance to AI researchers from the start: to what extent (and how) do (e.g.) DNNs represent knowledge?
- Piagetian development: do the different types of AI techniques show stages? (for example, coding the data initially in DNNs, before going on to use these codes in interpretation.
- Axioms and reliability: there had been discussion of the possibility of using DNN-type systems to produce the predicates that might then be used in logical systems. But often these predicates come with probabilities – or at the least, are not certainties. Can they be used in logical theory systems? Could we use probabilistic or fuzzy-logics instead of the 2-valued logic? It was noted that some of the statements are definite (“Sweden in Europe”, whereas others may be altered.
- What's the role of biology in understanding (different types of) learning? From insects to humans? Kane noted that we need to realise that our learning is related to our interaction with the world (of which we are part), and that the dualistic view of separating “us” from the rest of the world is not always appropriate. This might also mean that entities with quite different senses (and perhaps distributed communicating entities) might learn quite different things (this came from Peter Ross's talk).

Finally we discussed what the topics of future meetings of the AI Theme might be, and came

up with the following:

- Social Impact of AI
- The connection between AI and Human Intelligence
- Perception and AI
- AI in safety-critical environments
- Autonomous Systems (including robotics, and ethical issues)
- Barriers to the adoption of AI (in industry, commerce, healthcare, ...)

In addition, there was a suggestion of holding a public event<sup>1</sup>, making the AI theme more generally visible.

The meeting the adjourned to the bar of the MacRobert Centre for further discussions.

---

<sup>1</sup> Since then, I have had a discussion with Julie Allan, of the MacRobert Centre here, and they would be interested in helping us to hold such an event. I was thinking of a something with a title like *AI: threat or saviour*, with a couple of short accessible presentations, followed by a more general chaired discussion. This could be held in one place initially, and then repeated a few times in different Scottish locations.

SICSA AI THEME MEETING

**How do the many facets of learning in AI relate to each other?  
Programme**

30 August 2017, 1000-1700, Lecture room B3 (and 2X4/6), Cottrell Building,  
University of Stirling, Stirling FK9 4LA.

**0930-1000: Arrival and coffee (2X4/6)**

**1000-10.10:** Lecture Theatre B3: Introduction to the day: aim of the meeting,  
and outline of the day: Leslie Smith and Emma Hart

**1010-1245:** Lecture Theatre B3: Short talks on different aspects of learning in  
Artificial Intelligence:

Chair: *Emma Hart.*

1010-1030: Peter Ross: *Some future directions for machine learning and AI*

1030-1050: David King: *Learning in games*

1050-1110: Iain Murray: *Density estimation: machine learning meets statistics*

**1110-1140 Coffee (2X4/6)**

(Short talks continue: Lecture Theatre B3)

Chair: *Leslie Smith*

1140-1200: Alan Bundy: *Reformation: A Generic Algorithm for Repairing Faulty  
Logical Theories*

1200-12.20: Nick Taylor: *Machine learning at the Edinburgh Centre for  
Robotics*

1220-1240: Emma Hart: *Lifelong learning in Optimisation*

**1240 – 1345: Buffet lunch (2X4/6)**

1320 – 1430: *Circulation & Poster session*

1430 - 1515: *Discussions session 1 (groups):*

Discussion topic 1: What are the differences in approach to learning in different areas  
of AI? How might one bring these together?

Discussion Topic 2: What are the possibilities/obstacles to bringing different types of  
learning together?

**1515: Coffee arrives and discussions continue**

1515-16:00: *Discussions session 2 (groups):*

(groups swap over)

1600-17:00: Groups report back, and concluding remarks.

**1700 onwards: informal discussions in a local hostelry.**