Imperial College London

HLC Newsletter

ISSUE 03

July 2020

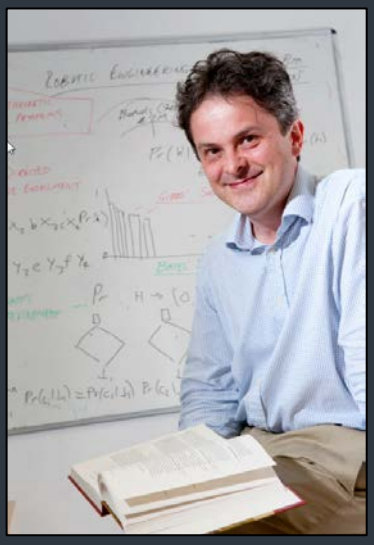

The Machine Intelligence 21 workshop, June/July 2019

Stephen Muggleton, Imperial College London

The Machine Intelligence 21 (MI21-HLC) was held at Cumberland Lodge in Windsor from 30th June to 3rd July 2019 (see http://mi21-hlc.doc.ic.ac.uk/). The workshop was the first of its kind to be supported by the EPSRC Human-Like Computing Network (2018-2023) and built on successful talks and discussions between internationally leading researchers in Artificial Intelligence and Cognitive Science at the MI20-HLC workshop in 2016. The research described aims to address present incompatibilities of Human and Machine reasoning and learning approaches. According to the influential US funding agency DARPA (originator of the Internet and Self-Driving Cars) this new area represents the Third Wave of Artificial Intelligence (3AI, 2020s-2030s), and is being actively investigated in the US, Europe and China.

The EPSRC's UK network on Human-Like Computing (HLC) is one of the first internationally to initiate and support research specifically in this area. Starting activities in 2018, the network represents around sixty leading UK groups Artificial Intelligence and Cognitive Scientists involved in the development of the inter-disciplinary area of HLC. The research of network groups aims to address key unsolved problems at the interface between Psychology and Computer Science.

Papers from the MI20 and MI21 workshops and the Third Wave Artificial Intelligence workshop are presently being compiled and edited as chapters for a book to be published in 2020 by Oxford University Press. The chapters of this book are authored by a mixture of UK and other international specialists. Some of the key questions addressed in the book include how AI systems might 1) explain their decisions effectively, 2) interact with human beings in natural language, 3) learn from small numbers of examples and 4) learn with minimal supervision. Solving such fundamental problems involves new foundational research in both the Psychology of perception and interaction as well as the development of novel algorithmic approaches in Artificial Intelligence.

Stuart Russell

The key benefit of the network will be the development and maintenance of an inter-discplinary community of HLC researchers at the forefront of research in Artificial Intelligence and Cognitive Science

Alan Bundy

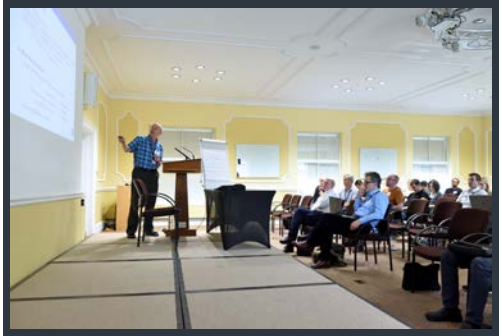

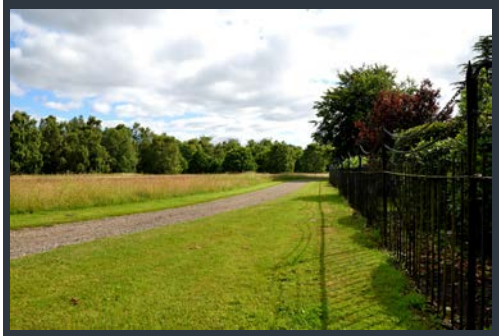

The grounds around Cumberland Lodge

Alan Bundy, University of Edinburgh

The Third Human-Like Computing Network+ Workshop was held at Cumberland Lodge in the middle of Windsor Great Park from 30th June to 3rd July. This was also the site of some of the second workshop and, once again, proved an excellent location. The Park is beautiful and, although close to Central London, feels remote, which encourages engagement with the Workshop and promotes interaction. The food is also very good. Like its predecessor, it was also co-badged as a Machine Intelligence Workshop, the 21st, ensuring that its proceedings would be published by Oxford University Press. As I write, the camera-ready copy is being sent to OUP for printing. The Workshop also followed our tradition of a mixture of long talks and short talks with posters. There were several non-UK speakers, some of whom joined by video.

A prime aim of the HLC Network+ is to promote interaction between the Artificial Intelligence and Cognitive Science communities. The talks confirmed the success of this aim. Many proposed computational models of human or other animal cognition – and these came from both the AI and the Cogsci participants, motivated by the goals of both communities. These models ranged over understanding language & gesture, visual perception, dialogue & argumentation and learning. Some were applied to practical problems, such as human/machine interaction including explanation, computer-aided tutoring, choosing & repairing knowledge representations and automating data-science. Some were aimed at understanding specific aspects of human cognition, such as, early language learning, cooperation, virtual bargaining and the conjunction fallacy. None of the talks could be exclusively classified as straight AI or Cogsci – nor could the discussions arising from these talks. New collaborations were seeded from these discussions.

We also had talks about potential sources of funding for these collaborations. These talks were from Rhys Perry from the Engineering and Physical Sciences Research Council (EPSRC), Ali Anjomshoaa from the Knowledge Transfer Network (KTN), which is part of Innovate UK (IUK), and from Kenneth Kwok from the Agency for Science, Technology and Research (A*STAR) in Singapore about the human-centric AI research programme.

Stuart Russell

Nick Chater

Stephen Muggleton

Ute Schmid

Luc De Raedt

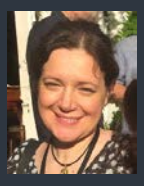

Katya Tentori

Alan Bundy, University of Edinburgh

There isn't space to describe all the talks in detail, so I'll just select a few that caught my eye.

- Stuart Russell summarised the proposals, in his recent book "Human Compatible", to ensure that Artificial Intelligence is beneficial to humans. AI systems will learn the real preferences of their human users by playing game-theoretic assistance games with them.

- Nick Chater described *Virtual Bargaining*, the amazing ability that humans have to cooperate with other humans in conditions where their ability to communicate is severely constrained. They can imagine themselves in the shoes of their partner and 'bargain' with them virtually to arrive a mutually optimal strategy. This ability is especially interesting to me because we are working with Nick to model it, as I described in my talk.

- Stephen Muggleton has applied a noise-tolerant version of his Meta-Level Learning (MIL) algorithm to construct logic program explanations of visual images. For instance, the learned program can predict the location of light sources illuminating an image. One of the notable features of MIL is its ability to invent new intermediate predicates, e.g., for the concept of an individual step in a description of a staircase.

- Ute Schmid showed how logic-based learning was used to repair faulty rules in her Dare2Del tool, which assists users to identify files that can be deleted. When it makes inappropriate suggestions, Dare2Del uses user interaction to identify and repair the rules responsible for the error.

- Luc De Raedt describe his SYNTH system, which automates aspects of data science by combining predictive autocompletion; predictive modelling and auto-transformation.

- Katya Tentori described a logic account of the very common, Conjunction Fallacy, in which, contrary to the laws of probability a subject will assign a higher probability to A & B than to A (or B). For instance, suppose that Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more likely?
    - A: Linda is a bank teller.
    - A&B: Linda is a bank teller and is active in the feminist movement.

  Katya observed that this happens when B is a hypothesis that would confirm the background information.

The Machine Intelligence 21 workshop, June/July 2019






The Machine Intelligence 21 workshop, June/July 2019






EPSRC Engineering and Physical Sciences Research Council

WARWICK THE UNIVERSITY OF WARWICK

Imperial College London