Imperial College London

HLC Newsletter

ISSUE 05 March 2021

this issue

HLC Calls for Funding 1

Awards and Funding 2

Journals and Publications 3

Media 4

Upcoming Meetings and Events 5

HLC Year 1 Kick-start Final Report 6

Meetings and Events 7

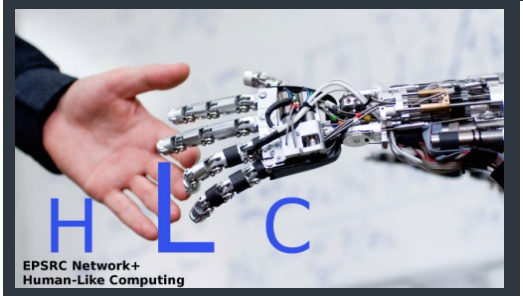

HLC Calls for Funding

Year 4 calls for the EPSRC Network+ on Human-Like Computing are now open and we invite you all to apply.

The calls are available on our website, on this link:

http://hlc.doc.ic.ac.uk/calls.html

The calls include 2 x 80K Kick-start grants, an Industrial engagement grant, 2 x Travel grants, EPSRC grant submission endorsements and Innovate UK grant submission endorsements. Full details of the format and deadlines for these are included on our website.

The closing date for all calls is **Friday 7th May 2021** and the result announcement is expected by 9th July 2021. Notification of funding will be provided shortly afterwards.

Whilst we appreciate that no travel is currently permissible, we encourage you to apply for the 2021 travel grants and if you are successful then the dates for your actual travel can be discussed at this point. Due to the major disruption caused by COVID-19 we have extended the period in which the travel funding can be spent. Any unexpected delays to your planned travel, should be discussed with the HLC Director at the time.

Awards & Funding

Marko Tesic

Congratulations to Marko Tesic, Birkbeck, University of London who has been awarded a UK Intelligence Community Postdoctoral Research Fellowship from the Royal Academy of Engineering.

The purpose of the award is to support research that advances national security. The work that he is conducting is in the area of explainable artificial intelligence, where he will be looking into the role of explanation in (re)building trust in AI systems.

More specifically, he will aim to explore (i) what a human user takes to be explanatory in the AI context and (ii) what types of explanations are most conducive to build trust in an AI system's outputs

The fellowship is in collaboration with some members of both the UK and Australian intelligence communities.

Marko has started research on his fellowship in March 2021 and it will span over the next 2 years. Marko has been granted £200,000 to support his research.The research will be conducted at Birkbeck, University of London. Ulrike Hahn is his research advisor at Birkbeck and he has multiple advisors from the intelligence community as well.

Read more: https://www.raeng.org.uk/grants-prizes/grants/support-for-research/ic-postdoctoral/current-and-recent-awards

Ulrike Hahn

Top journals, conference papers and publications

We know of the following papers recently published by members of the network:

- Dany Varghese and Alireza Tamaddoni-Nezhad both from the University of Surrey have had a conference paper titled "One-shot rule learning for challenging character recognition" accepted in the proceedings of the 14th Intl. Rule Challenge (RuleML) 2644:10-27, 2020

- Nir Oren, University of Aberdeen and Louise Dennis, University of Manchester have had a paper based on human-like computing intuitions (namely argumentation and dialogue) titled "Explaining BDI Agent Behaviour Through Dialogue" accepted for AAMAS.

- FACE: Feasible and Actionable Counterfactual Explanations. Rafael Poyiadzi, Kacper Sokol, Raúl Santos-Rodríguez, Tijl De Bie and Peter Flach. AIES '20: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society https://dl.acm.org/doi/10.1145/3375627.3375850

- Explainability Fact Sheets: A Framework for Systematic Assessment of Explainable Approaches. Kacper Sokol and Peter Flach. FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. https://dl.acm.org/doi/abs/10.1145/3351095.3372870

- L. Ai, S.H. Muggleton, C. Hocquette, M. Gromowski, and U. Schmid. Beneficial and harmful explanatory machine learning. *Machine Learning*, 2021. https://doi.org/10.1007/s10994-020-05941-0.

- S. Patsantzis and S.H. Muggleton. Top program construction and reduction for polynomial time meta-interpretive learning. *Machine Learning*, 2021. https://doi.org/10.1007/s10994-020-05945-w

- Konstantina Spanaki, Erisa Karafili, Stella Despoudi “AI Applications of Data Sharing in Agriculture 4.0: A Framework for Role-based Data Access Control” to appear at the *International Journal of Information Management*, Elsevier, 2021.

- Konstantina Spanaki, Erisa Karafili, Uthayasankar Sivarajah, Stella Despoudi, Zahir Irani “Artificial intelligence and food security: swarm intelligence of AgriTech drones for smart AgriFood operations” in *Journal of Production Planning & Control*, Taylor & Francis, 2021. https://doi.org/10.1080/09537287.2021.1882688

- Erisa Karafili, Fulvio Valenza “Automatic Firewalls’ Configuration Using Argumentation Reasoning” in *Workshop on Emerging Technologies for Authorization and Authentication (ETAA@ESORICS),* 2020. https://doi.org/10.1007/978-3-030-64455-0_8

Media

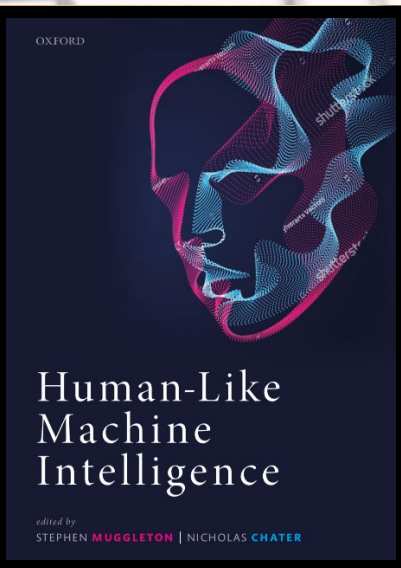

We are pleased to announce that a new book, Human-Like Machine Intelligence, edited by Stephen Muggleton and Nick Chater may now be pre-ordered for delivery when available.

This book will be published as a hardback book.

Notify Me When In Stock

Estimated publication date: 3rd June 2021

560 Pages | 92 line art, 18 combo, and 13 halftones

246x171mm

ISBN: 9780198862536

Read more:
https://global.oup.com/academic/product/human-like-machine-intelligence-9780198862536?facet_narrowbybinding_facet=Hardback&facet_narrowbyprice_facet=50to100&facet_narrowbyproducttype_facet=Digital&lang=en&cc=ru

Stephen Muggleton

Nick Chater

Upcoming Meetings and Events

A reminder that a Dagstuhl Seminar on **"Approaches and Applications of Inductive Programming"** is due to be held from 9th to 12th May 2021. Due to COVID-19 this event is now being held virtually.

The organizers of this event are:
Andrew Cropper (University of Oxford, GB)
Luc De Raedt (KU Leuven, BE)
Richard Evans (DeepMind – London, GB)
Ute Schmid (Universität Bamberg, DE)

It is the fifth time that the bi-annual AAIP event will be held as a Dagstuhl Seminar and the 9th time over all. AAIP started as a workshop at ICML 2005 in Bonn with Stephen Muggleton as keynote speaker.

Read more:
https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=21192

Peter Flach, University of Bristol is giving a talk and Kacper Sokol, University of Bristol is giving a short demonstration at the upcoming The Turing Presents: AI UK conference being held on the 23rd and 24th March 2021. Peter's talk is titled "Computer says 'I don't know' – the case for Honest AI". Kacper's demonstration is titled "Did You Get That? The Intelligibility of the State-of-the-art Explanations in Artificial Intelligence and Machine Learning".
Read more: https://www.turing.ac.uk/ai-uk

Kacper Sokol, University of Bristol is giving an invited talk at the "Fair and Explainable Models" session of the 31st European Conference on Operational Research. Kacper's talk is titled "Making Machine Learning Explanations Truthful and Intelligible".
Read more: https://euro2021athens.com/

Andrew Cropper

Luc De Raedt

Richard Evans

Ute Schmid

Peter Flach

Kacper Sokol

Professor Francecsa Toni, Imperial College London

Final Report – HLC Year 1 Kick-start

Final Report:
Evaluating dialectical explanations for recommendations

Partners:
- Francesca Toni (Imperial College London, Computing) – PI
- Dave Lagnado and Christos Bechlivanidis (UCL, Experimental Psychology) – CoIs
- Antonio Rago (Imperial College London, Computing) – PostDoc

The project aimed at addressing the lack of transparency of AI techniques, e.g. machine learning algorithms or recommender systems, one of the most pressing issues in the field, especially given the ever-increasing integration of AI into everyday systems used by experts and non-experts alike, and the need to explain how and/or why these systems compute outputs, for any or for specific inputs. The need for explainability arises for a number of reasons: an expert may require more transparency to justify outputs of an AI system, especially in safety-critical situations, while a non-expert may place more trust in an AI system providing basic (rather than no) explanations, regarding, for example, films suggested by a recommender system.

The main aim of this project was to conduct experiments to determine whether and which computed dialectical explanations, extracted from argumentation graphs for explaining recommendations, are useful to humans and whether human feedback can improve the outputs of the recommender system. The planned experiments were identified as useful to confirm or falsify the hypothesis that argumentation can serve as a paradigm for human-machine interaction, in the specific setting of recommender systems and argumentative explanations
The project resulted in a number of publications, including:

Argumentation as a Framework for Interactive Explanations for Recommendations. Antonio Rago, Oana Cocarascu, Christos Bechlivanidis and Francesca Toni. KR 2020. https://proceedings.kr.org/2020/83/

Mining Property-driven Graphical Explanations for Data-centric AI from Argumentation Frameworks. Oana Cocarascu, Kristijonas Cyras, Antonio Rago, Francesca Toni. Human-Like Machine Intelligence. Estimated publication date: 3rd June 2021. ISBN: 9780198862536

Meetings and Events

Peter Flach, University of Bristol and Kacper Sokol, University of Bristol organised a hands-on ML explainability tutorial at **ECML-PKDD 2020.**

Title: What and How of Machine Learning Transparency: Building Bespoke Explainability Tools with Interoperable Algorithmic Components

Read more: https://events.fat-forensics.org/2020_ecml-pkdd (slides, computational notebooks and tutorial recordings are linked from this page)

In addition, Peter Flach, University of Bristol led a tutorial on classifier calibration at **ECML-PKDD 2020.**

Title: Classifier Calibration: How to assess and improve classifier confidence and uncertainty

Read more: https://classifier-calibration.github.io/

Funding

Peter Flach, University of Bristol and Kacper Sokol, University of Bristol recently started a new European project where Bristol is one of the partners: "TAILOR is one of the four H2020 ICT-48 European network of AI excellence centres". They are leading a work package on Collaboration, which is tasked with:

- developing AI-powered collaboration tools;

- building training platforms and materials;

- composing a PhD curriculum and training programme in AI; and

- organising AI summer/winter schools.

They will also collaborate with other work packages, working on Trustworthy AI, and Paradigms and Representations in AI. Peter and Kacper will be hosting an event to disseminate our work; and will later share the details.

Read more: https://tailor-network.eu/

Peter Flach, University of Bristol and Kacper Sokol, University of Bristol received the 2020/21 Alan Turing grant for the development of online training materials. By the end of the year they will publish an interactive online book about hands-on machine learning explainability. The tentative title is: "Hands-on Machine Learning Explainability: From Theory to Practice".