

AISB QUARTERLY

THE NEWSLETTER OF THE SOCIETY FOR THE STUDY OF
ARTIFICIAL INTELLIGENCE AND SIMULATION OF BEHAVIOUR



Celebrating 150th Issue of the Q

About the Cover

What is better than celebrating the 150th issue of the Q than throwing the swarms at the celebratory cake and let them *have the cake and eat it*, using their abundant Artificial Intelligence to recreate the memory of their journey and illustrate the turns of events while taking the journey. This work used Dispersive Flies Optimisation¹, a global optimiser and a Swarm Intelligence algorithm, which has been applied in various domains including medical imaging, optimising machine learning algorithms, deep neuroevolution, feature selection in data mining, as well as digital art.

Feeling geeky and arty?

If you are interested in designing a cover with the help of your off-the-shelf AI-boosted algorithms, feel free to contact the editor on aisbq@aisb.org.uk with your cover design (taking into account the already “set on stone” orange shade of the cover) along with a blurb on how you managed to get to the final results.

¹https://en.wikipedia.org/wiki/Dispersive_flies_optimisation

Editorial

Since my role as the editor of the Q, many issues have been raised in the previous editorials, several of which are still strongly persisting in the minds and hearts of many (sometimes not crossing the threshold of “enough”). Given it is the **150th issue** of the Q, while maintaining the upbeat tone inspired by the Q’s milestone achievement and longevity, some of these concerns are alluded to.

Some of these issues, which we are well aware of and deal with on a daily basis are the uncertainty of Brexit and its aftermath; the academics dispute and fight for fair pay and equal pay in higher education sector; and the future of AI in the UK, both in education sector (with the possible reduction in the number of EU students and academics thinking of relocating) and the industry (with many business entertaining the thoughts of move out of the UK either partially or altogether).

Despite all this, life goes on for Artificial Intelligence (and perhaps slightly less explicitly for the Simulation of Behaviour) and more planning and ambitious plans are crystallised regularly by policy makers through the encouragements of their (academic and business) advisers. One of the the latest ones, released on 6th November 2018, is the government’s announcement of five new centres of excellence for digital pathology and imaging, including radiology, using AI medical advances².

²<https://www.gov.uk/government/news/artificial-intelligence-to-help-save-lives-at-five-new-technology-centres>

³<https://www.theguardian.com/uk-news/2018/oct/15/prince-harry-and-meghan-markle-pregnant-baby-spring>

⁴Note: spring as season, that is, and not Spring Framework!

These £50m funded centres in Leeds, Oxford, Coventry, Glasgow and London aim to use: AI to diagnose diseases at an earlier stage; speed up the diagnosis process; develop more intelligent analysis of medical imaging, leading to better clinical decisions, by bringing together medical doctors, businesses and academics.

On a different (Natural Intelligence) yet related (to Artificial Intelligence) topic, Google recently stated that searches for “when is spring?” had reached its peak after the announcement of a pregnancy with the child expected during the season³. While we, intentionally or unintentionally, immerse our minds, hearts and evening discussions to the cause of the digital revolution and AISB, it is advisable to encourage ourselves and others not to lose touch with our surrounding and nature, and focus less on AI to untangle all the “complex” questions, including the aforementioned one⁴.

This issue allocates some of the pieces to celebrate the well-deserved 150th issue of the Q and the society’s history. Wishing the Q and all its members a happy celebration!

Mohammad Majid al-Rifaie

Editor

@mohmaj

University of Greenwich

Old Royal Naval College, London, U.K.

November 2018

The AISB in 2018 and beyond

Bertie Müller, AISB Chair

Swansea University, berndt.muller@swansea.ac.uk

We often hear questions like “What is the role of a learned society like AISB in the 21st century?”, “What value does it give to its members?”, and “Why should I join a society?” Most answers are quite straightforward, but the value is sometimes hard to convey. As a member of a learned society, you are a part of a wider research ecology that provides intellectual stimulation, public outreach, and reputational value for the individual and profession. As the world’s oldest AI society, AISB works both publicly and behind the scenes to advance research in AI and cognitive science to further the common good.

As a relatively small not-for-profit organisation, AISB has always prioritised supporting early-career researchers, such as PhD students, but this can only work with a healthy number of full-paying members. Recently, we have been confronted with members cancelling their membership because of the lack of apparent value of being a member. The perception of only receiving weekly email bulletins and quarterly printed newsletters does not seem to be sufficient added value. This perception neglects the fact that the events and opportunities published in the bulletins are carefully curated, unlike on other mailing lists. More importantly, it neglects the role AISB is playing in the national and international AI community that has seen some dramatic changes in recent

years. AISB values its members’ opinions. AISB is proud to have sixteen distinguished Fellows and a predominantly research-active membership. AISB is likewise proud to have representatives from think tanks and from businesses as members. We will continue to represent the diversity of our membership as a reflection of the AI landscape in the UK.

Geographical diversity has become rare in the AI economy with most UK activity clustered around London. AISB has traditionally followed a policy of geographical inclusivity, evidenced by holding the annual AISB Convention at venues across the UK, and we are proud to announce that Falmouth University will be hosting the 2019 AISB Convention with the theme “Artificial Intelligence, Imagination and Invention – A[I]³”. Apart from our own convention, AISB will continue to actively support government initiatives (e.g., Select Committee on AI, All-Party Parliamentary Group on AI - APPG-AI, AI Global Governance initiative), seek to influence policy through think tanks (e.g., Reform.UK, Future Care Capital, Future Advocacy), and seek dialogue with industry (e.g., CBI, Techerati, AI Summit, AI Europe, Marketforce). For all of these activities, the main aim for AISB is to ensure that higher-education institutions, businesses, and governments adopt a responsible approach to AI that values

transparency and ensures accountability, i.e., AI for the good of humanity.

At the European level, AISB was the only representative for the UK at the European Commission (EC) Workshop on “The European AI Landscape” in Brussels in January 2018 that led to the publication of an EC report and a Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions entitled “Artificial Intelligence for Europe”. The establishment of a network of co-ordinated AI initiatives is key to avoiding a concentration of social and economic benefits in the primary clusters of AI research and AI startup activity. As stated above, AISB has been working towards this by holding research meetings across all geographic regions in the UK and by giving the UK’s AI sector a voice in Europe. AISB acknowledges the importance of the vision set out by the CLAIRE⁵ initiative and will continue to actively engage in the discussion. In fact, I have personally had the pleasure to attend the inaugural CLAIRE symposium in Brussels in August. We publish the vision set out by the initiators of CLAIRE elsewhere in this issue and will keep you updated on future activities.

Finally, there will be an exciting addition to the portfolio of events run by AISB from 2019. In addition to the scientific conference AISB holds every year in the Spring, we are going to establish a second annual event in the Autumn. This event will be targeted at the general public and at businesses working with AI. The idea is to showcase in exhibitions how AI has become part of everyday life and affects diverse areas like art, leisure and business alike. We hope to facilitate this event to demystify AI, showcase good examples of AI, and stimulate discussion. Since 2014, AISB has organised the longest running Turing-Test competition, the Loebner Prize. Due to various reasons, we will no longer run this as a standalone event, but intend, instead, to use the Autumn event to exhibit chatbots in a modified context and format, divided into two categories. One category will continue the tradition of the Loebner Prize, while a second category will be open to less constrained conversational systems. Watch our web site and this space for further announcements.

We are always open to feedback and suggestions from our members. Please direct these to chair@aisb.org.uk.

Best regards,
Dr Bertie Müller, October 2018
(AISB Chair)

⁵Confederation Of Laboratories For Artificial Intelligence Research In Europe (claire-ai.org)

A European Vision for AI

Call for the Establishment of a Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE)

Prepared by Holger Hoos (Universiteit Leiden, The Netherlands)

Morten Irgens (Oslo Metropolitan University, Norway)

Philipp Slusallek (German Research Center for Artificial Intelligence, Germany)

Based on discussions with many key members of the European AI Community⁷

The Need for a European AI Strategy

Artificial intelligence (AI) methods and technologies are posed to bring transformative change to societies and industries world-wide. The game-changing nature of AI and its role as a major driver of innovation, future growth, and competitiveness are internationally recognised. As a result, AI is at the top of national and international policy agendas around the globe.

In the United States of America, huge investments in AI are made by the private sector, and a substantial governmental plan was launched in 2016, which includes significant long-term in-

vestments in AI research⁸. Similarly, in 2017, the Canadian government has started making major investments in AI research, focusing mostly on existing strength in deep learning⁹. In 2017, China released its Next Generation AI Development Plan, with the explicit goal of attaining AI supremacy by 2030¹⁰.

However, in terms of investment in talent, research, technology and innovation in AI, Europe lags far behind its competitors. As a result, the EU and associated countries are increasingly losing talent to academia and industry elsewhere¹¹. Europe needs to play a key role in shaping how AI

⁷Find additional information about this initiative at claire-ai.org

⁸Preparing for the Future of Artificial Intelligence”, Executive Office Of The President, National Science and Technology Council, Committee on Technology, the Office of Science and Technology Policy, 12 October 2016, ISBN 1544643136, 9781544643137. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

⁹Canada funds \$125 million Pan-Canadian Artificial Intelligence Strategy, Canadian Institute for Advanced Research, 22 March 2017, <https://www.newswire.ca/news-releases/canada-funds-125-million-pan-canadian-artificial-intelligence-strategy-616876434.html>

¹⁰New Generation Artificial Intelligence Development Plan, State Council of China 2017, http://www.gov.cn/zhengce/content/2017-07/20/content_5211996.htm, translation: <https://www.newamerica.org/documents/1959/translation-fulltext-8.1.17.pdf>

¹¹Big tech firms’ AI hiring frenzy leads to brain drain at UK universities, the Guardian, 2 November 2017, <https://www.theguardian.com/science/2017/nov/02/big-tech-firms-google-ai-hiring-frenzy-brain-drain-uk-universities>

changes the world, and, of course, benefit from the results of AI research. The reason is obvious: AI is crucial for meeting Europe's needs to address complex challenges as well as for positioning Europe and its nations in the global market.

Europe has started to react: In April 2018, 25 countries pledged to increase national research funding for AI as part of a common "European approach"¹². In parallel, the European Commission laid out a preliminary plan for strengthening AI across Europe¹³, realising that more focussed instruments are needed beyond those planned in H2020 to turn the tide and achieve the research and innovations we need, and on the scale we need. This urgent sense of need for action was also clearly expressed in a recent open letter by a number of AI researchers, who proposed a European research centre in machine learning and related areas of AI¹⁴.

In the following, we outline a proposal that builds on and expands on these initiatives, and, we believe, is necessary to meet their objectives. In particular, we strongly support the ambition and vision articulated in the recent EC Communication (see footnote 13), and we endeavour to present a specific approach to realising it.

All of AI, all of Europe, with a Human-Centred Focus

There is a pressing need for increasing Europe's strength and position in the area of AI research. Based on extensive discussions within the community of European AI researchers, following the recent EC Communication on AI (see footnote 13), a strong consensus has emerged on key aspects of a coordinated European research effort.

In particular, a broad and ambitious vision is needed for European AI research to thrive and for Europe to stay competitive with other major players. The research and innovation efforts required in this context should encompass all of AI, and include all of Europe. Furthermore, by building on our existing strength in AI and commitment to European values, Europe should take a human-centred approach to AI.

We call for a vision that aims to

1. have European research and innovation in artificial intelligence be amongst the best in the world, that
2. encompasses all of AI and all of Europe
3. has a strong focus on human-centred AI

¹²Declaration, Cooperation on Artificial Intelligence, Brussels, 10 April 2018, http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=50951, See also <https://ec.europa.eu/digital-single-market/en/news/eu-member-states-sign-cooperate-artificial-intelligence>

¹³Artificial Intelligence for Europe. Communication From The Commission To The European Parliament, The European Council, The Council, The European Economic And Social Committee And The Committee Of The Regions, COM(2018) 237, SWD(2018) 137, the European Commission, April 25, 2018. Link, as of May 2018: http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=51625

¹⁴Initiative to establish a European Lab for Learning & Intelligent Systems (The ELLIS Initiative), Bach et al. 2018, Link as of May 2018:<https://ellis-open-letter.eu/>.

The best in the world. In order to meet Europe's challenges and to secure markets, European research and innovation needs to be among the best in the world. The good news is that Europe is very well positioned: We have a strong tradition of excellence in all areas of AI, and many of the top European researchers are recognised as leading figures within the AI community and related fields of research and technology.

All of AI. Artificial intelligence comprises a broad spectrum of methods and techniques, each with their own important applications. While recently, advances in machine learning techniques have enabled rapid progress across many areas of AI, future applications of AI will increasingly leverage combinations of AI techniques. It is therefore crucial that Europe builds on its existing strengths across the full spectrum of AI, covering all of machine learning, knowledge representation and reasoning, search and optimisation, planning & scheduling, multi-agent systems, natural language processing, robotics, computer vision, and other areas.

A broad view of AI that includes all areas within the field is essential to meet the challenges that lie ahead of us, especially in human-centred, ethical AI, where explanations and deep understanding (of natural language, images, etc.) are essential to achieve trust between humans and machines, and to

thus obtain the best solutions to the problems we face as individuals and societies. Moreover, AI researchers need to adopt a multi-disciplinary approach and work with experts from other areas, not only from mathematics, engineering and the natural sciences, but also with social scientists.

All of Europe. Human talent is already a limiting factor in AI research and development in Europe. For a European initiative to succeed, it needs to attract, educate, and harness talent, and drive innovation across the continent, leveraging the strength in AI currently found in many European countries, and ensuring diversity and inclusion across languages, cultures and gender. It is therefore of key importance to foster AI excellence across Europe.

Human-centred AI. Artificial intelligence increasingly enables new forms of production, services, and medical treatments, but may also lead to increased bias, inequity, manipulation, invasion of privacy, and job loss¹⁵. We believe that responsible AI research and deployment should be strategically focussed on augmenting human capabilities, rather than replacing them, on compensating for human bias and limitations, and on serving and protecting the human and ethical values that are of core importance to European societies¹⁶. Research on AI in Europe thus

¹⁵The Future of Employment: How susceptible are Jobs to Computerisation? Carl Benedikt Frey and Michael A. Osborne, the Oxford Martin Programme on Technology and Employment, University of Oxford, Link: www.oxfordmartin.ox.ac.uk/publications/view/1314

¹⁶The Age of Artificial Intelligence. Towards a European Strategy for Human-Centric Machines", EPSC Strategic Notes, European Political Strategy Centre, Issue 29, 27

needs to understand, anticipate, and address ethical, legal and social aspects (also known as, and in EU's Framework Programmes usually referred to as, Responsible Research and Innovation, or RRI).

As AI scientists, we are keenly aware that AI is a disruptive set of technologies. Consequently, we need to act at the European level and issue a set of principles and guidelines regarding the responsible use of AI - similar to what physicists did in 1955 with the "Russell-Einstein manifesto". This **"AI manifesto" should stipulate limits of responsible use and anticipate the consequences of deploying specialised AI systems as well as of creating general, human-level AI.** It should also define how to quantitatively and qualitatively assess whether AI systems or agents comply with those limits. We believe that European AI researchers are in an ideal position to play a leading role in an ambitious, global effort to address these issues and have a responsibility to exercise leadership in this area.

A Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE)

The discussions within the community of European AI researchers have also led to a clear understanding that Europe needs not only to increase its research activity level, but also to coordinate better and collaborate more closely. This requires investment in both outstanding AI research and in structures that allow effective collaboration and transfer of results. In partic-

ular, major actions are required to develop and retain key talent and expertise in AI, and existing strength needs to be leveraged and expanded.

Specifically, we call for the establishment of a **Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE)**, comprising a network of centres of excellence, strategically located throughout Europe, and a new, central facility that serves as a hub, providing state-of-the-art infrastructure, and fostering the exchange of ideas and expertise.

To be effective in meeting the above vision, CLAIRE should consist of the following key elements:

- A collaborative network of relevant existing and new research labs and organisations across Europe. Under the leadership of some of the top researchers in the field, this network should jointly identify fundamental research questions, discuss the most promising approaches, and help organise collaborative efforts to address them.
- A selection of some of these research labs, located strategically throughout the European Economic Area and EFTA, to be designated "Centres of Excellence in AI", should play strong regional or national roles as hubs for the members of the collaborative network in their region.

March 2018: https://ec.europa.eu/epsc/sites/epsc/files/epsc_strategicnote_ai.pdf

- A new facility that serves as a highly visible and vibrant focal point for the collaborative network, the “CLAIRE Hub”. Here, excellent scientific personnel at all levels and from all partners would find an outstanding research environment for AI, where they can work together, face-to-face, for periods of time (e.g., an extended version of the highly successful Leibniz Centre for Informatics in Dagstuhl, Germany). This hub should provide cutting-edge infrastructure and support, but would not have permanent scientific staff.

This is a model that builds on existing strengths, brings together the still fragmented AI research activities and expertise in Europe, while at the same time creating centres of excellence and a structure that can efficiently focus research and distribute results.

Our vision for CLAIRE is in part inspired by the extremely successful model of CERN¹⁷. CERN’s research activities rely on distributed, collaborative efforts of many physics laboratories across Europe that jointly define the research questions, discuss how to address them, and then collaboratively develop the experiments and publish the results. This network is supported by a central facility, a joint working environment, and a distributed research infrastructure for collaboration and data sharing.

In other aspects, CLAIRE will differ from CERN: Despite the central facility, its structure will be more distributed, as there is less need for reliance on a single experimental facility. It will also have much closer collaboration with industry, to quickly and efficiently transfer new results and insights. Similar to CERN, the suggested structure will allow for the establishment of a common, well-recognised “trademark” for high-quality European AI research. As can be seen with CERN, research gets much wider exposure by being associated with a “trademark” like CERN, without diminishing the scientific reward for the individual researchers and their laboratories.

How CLAIRE can Ensure the Success of European AI

While deliberately refraining from defining details of the organisation and financing of the proposed Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE) at this early stage, we believe that the following ideas and concepts are important to ensure its success:

Wide range of applications. CLAIRE should support AI research that is expected to have major short-, medium- and long-term impact across a wide range of application areas, including efficient and safer transportation, advanced healthcare, smart industry, effective and sustainable agriculture, accelerated scientific research, and others. In order to have a significant impact on applications, fund-

¹⁷Position paper on “CERN for AI”, OECD, Oct 2017, see also: <https://www.oecd-forum.org/users/71431-philipp-slusallek/posts/28452-artificial-intelligence-and-digital-reality-do-we-need-a-cern-for-ai>

ing should be targeted towards existing scientific strengths, novel research opportunities and key European interests. CLAIRE should also put in place suitable mechanisms to engage with industries and collaborate with them on defining and tackling applications in various sectors.

Social impact. CLAIRE should also conduct AI research that aims at resolving some of the open issues regarding the social impact of AI, such as fairness, transparency, explainability and value alignment, with a clear focus on building trustworthy AI that is beneficial to people and aligned to European values.

Attracting talent. Funding should be focussed on existing strengths and support attractive fellowships for Master students, PhD candidates and post-doctoral researchers, as well as exceptional junior, mid-career, and senior researchers, using efficient, light-weight but solid scientific reviewing and allocation processes. The aim should be to attract the best talent from all over the world.

The CLAIRE Hub. The CLAIRE Hub should be created to provide a visible, vibrant center for AI research in Europe. This facility should comprise a large, state-of-the-art data and computer centre, cutting edge robotics laboratories, test facilities for key application areas, such as autonomous transportation, advanced agriculture and automated scientific experimentation, usability labs, and oth-

ers. It should have outstanding support staff, including programmers, usability and interface experts, and hardware technicians. The centre should also maintain a repository of datasets open for researchers across Europe. We do not envision the CLAIRE Hub including permanent research staff, but rather as hosting affiliated researchers and visitors (including researchers on sabbatical / study leave) at all levels of seniority for limited periods of time to exchange ideas, work on projects, and jointly use infrastructure only available there.

Strong infrastructure. The CLAIRE network should be supported by strong infrastructure in terms of computing, big data storage (including long-term storage and secure storage for sensitive data), and networking as well as infrastructure for maintaining joint AI platforms and services. It needs to be able to support large-scale AI research that can compete at the level of large private entities, while focusing on areas specifically relevant for Europe. Collaborations with existing initiatives, such as GÉANT¹⁸ or the new EuroHPC¹⁹, will be essential.

Fostering talent. The central facility should run summer schools, seminars, public outreach activities and workshops at the highest scientific level and from all areas of AI. A special focus should be on identifying and supporting as early as possible the best emerging AI talent across Europe.

¹⁸<https://www.geant.org/>

¹⁹<https://ec.europa.eu/digital-single-market/en/eurohpc-joint-undertaking>

Centres of Excellence. The distributed Centres of Excellence in AI should become highly visible and vibrant regional environments. Excellent scientific personnel and students receive secondary appointments and spend part of their time at the central hub or in other centres of the network. The Centres of Excellence would benefit from (i) outstanding research infrastructure that encourages collaboration across many areas of AI, (ii) state-of-the-art collaboration infrastructure, such as conference and working environments, (iii) an innovation infrastructure that facilitates industry collaboration and entrepreneurship, including pre-incubators, innovation advisors, and well-developed relationships with governmental innovation support and investor organisations.

Supporting collaboration. To stop the current level of brain-drain and attract the best talent, CLAIRE should provide considerable support for exchange and interaction of researchers at all levels of seniority, across all areas of AI. The CLAIRE Hub should provide a focal point for such exchange and interaction. There should be considerable incentives to establish joint research projects among researchers in the network, under the lead of top researchers in the field. CLAIRE should provide support for developing and running high-quality study programmes focussed on AI.

Because top researchers produce excellent results with high consistency, a substantial part of the funding should

be allocated based on track record (e.g., similar to the Reinhart Koselleck funding²⁰ in Germany for excellent researchers, which is based on 5-page proposals for up to 1.25 M€) rather than the need for extensive research proposals. Of course, substantial funding opportunities also need to be provided for talented junior researchers without a long-standing track record.

While the individual researcher needs to enjoy full academic freedom, the community as a whole would greatly benefit from a more coordinated approach to guiding European AI research. Other disciplines, e.g., physics, have long provided excellent examples for this. To that end, CLAIRE would establish a scientific process through which fundamental research questions, the most promising approaches to their solution, and key steps to implement them would be identified and regularly updated. The process would combine top-down and bottom-up mechanisms, leveraging guidance from top European scientists that are highly trusted by the community, as well as new ideas from outstanding members of the community, including rising stars.

The scientific discussions that form the basis of this process will offer a unique opportunity to strengthen, focus and coordinate European AI research, while offering a solid basis for developing guidelines for industry, politics and the general public. In parallel, CLAIRE would organise and drive similar discussions regarding both the requirements of European indus-

²⁰http://www.dfg.de/en/research_funding/programmes/individual/reinhart_koselleck_projects/index.html

tries as well as the benefits, social consequences, and key European values that should define and drive a European approach to AI. The open and well-founded discussion between science, industry, and society will be a key element of CLAIRE.

Realising the Vision

Obviously, many details of our vision for CLAIRE remain to be specified, and we strongly believe that this should be done based on further discussions with members of the European AI community and other stakeholders. Yet, given the global competition, CLAIRE (the distributed “CERN for AI”) should be created as quickly as possible, to maximise retention of critical AI talent in Europe and to start defining and focusing a joint European AI agenda.

A number of activities in that direction have already been started: The “Humane AI” proposal for an EU Flagship project on AI has just progressed to the second phase, together with two other proposals in the AI context: ro-

botics and language technology. Several proposals for a European AI-on-demand platform are currently under review (ICT-26), with the goal of starting work on a software platform supporting AI research and development throughout Europe later this year. The EC Joint Research Centre has also started to devote significant attention to AI. In addition, many member states have ramped up their AI research programmes and platforms, including large cross-national AI initiatives, such as the planned French-German collaboration on AI. Again, most of them highlight the need for human-centred AI and share the core vision of CLAIRE. Finally, we note that the previously mentioned open letter (see footnote 12) calling for substantial investment in machine learning research in Europe, with a focus on excellence in fundamental research, is well-aligned with our vision. All these initiatives form a strong basis for the larger vision of CLAIRE, which unites and strengthens AI across Europe²¹.

²¹List of supporters can be found at <https://claire-ai.org/>

Holding Our Brea(d)(th)

John A Barnden, AISB Fellow and AISB Patron

University of Birmingham, UK, J.A.Barnden@cs.bham.ac.uk

The multi-/cross-/interdisciplinary eclecticism of AISB, hinted at though hardly described by the SB part of the society name, always appealed to me. It was a main quality that I sought to support and promote during my time on the committee, especially as chair of it. The society is notable in the way that it welcomes, at its Conventions for instance, but also in the more-recently instituted series of AISB Workshops, any topic from abstruse philosophy to artistic, bodily expression such as dance ...not leaving out the philosophy of dance. Somewhat in line with this breadth, in 2012 I enjoyed organising, together with the President of IACAP (International Association for Computing and Philosophy), a joint congress that served as that year's AISB Convention. I was touched by the fact that many philosophers who came felt that it would be beneficial to repeat the experience. So that's a type of event that we should keep on our society's to-do list. And perhaps with the populace's and government's increasing general awareness of AI, combined with dispiriting narrowness of understanding of what it is (even senior academics I have recently interacted with outside AI imagining that it simply equates to machine learning), it is ever more important to stress the intertwinedness of the concerns of AI with those of many other disciplines or walks of life.

Writing the present piece, and thinking about conference organisation, has reminded me of a disruptive question that was asked by a variety of people during my years on the committee: why does academia in general think that invited plenary speakers at conferences should have their expenses covered? Such speakers tend after all to be more senior people with successful careers who could often afford to pay for themselves, at least when still employed but sometimes also when retired. Wouldn't the money be better spent on having more bursaries for students or others who have greater difficulty in paying, or on generally reducing the registration fees? Or, in the special case of AISB Conventions, paying the expenses of invitees to the component Symposia, who are often earlier in their careers? Sure, an invited plenary speaker is providing a valuable service to a conference, but so is everyone who brings their ideas and questions along: and we don't systematically pay back expenses to people who give good talks or ask good questions. I think we should disentangle money from the question of the honour paid to speakers by inviting them and the honour they pay a conference by contributing to it. I'm not sure of the practical details of how to make this work, but I could imagine some sort of honour system (in a different sense) where an invitee is free to tell

the organising committee whether they can afford to pay for themselves or not.

Perhaps the AISB could be at the forefront of a new spirit of academic in-

itation ... not forgetting that before too long some invitees will be robots. Our money now needs robots, but will robots need our money?

It's Simulation, Jim, but Not as We Know It.

Hugh David

When I joined the Society for the Study of Artificial Intelligence and Simulation of Behaviour about 1967, it was for two reasons. I had been, briefly, the only university lecturer in Cybernetics in England a subject about which I knew very little - there was very little to know. I was also very interested in Simulation. I was completing a Ph. D. on the ability of air traffic controllers to predict conflicts on radar screens, using what was for its time a very advanced Radar Simulator [2]. I later realised that the basic premise of my study was completely wrong, but it was a learning experience for all concerned [4]

After completing my Ph.D. at Loughborough, I joined the European Organisation for the Safety of Air Navigation (EUROCONTROL) in 1970, two years before the U.K. joined the Common Market. (Three new handles and two new blades, but just as good as new.) From then on I supported the AISB as the Duke of Wellington supported the Church of England - not as a pillar, but a flying buttress - from outside. EUROCONTROL was building the first large scale digital ATC simu-

lator in the world, as part of building the most advanced Air Traffic Control centre in the world at Maastricht in the Netherlands. It was also the first international centre - it now covers Belgium, Luxemburg, the Netherlands and half of Germany. The simulator and the centre are still the most advanced in the world, having been updated regularly - like Jack Hobbes's cricket bat - three new blades and two new handles, nut as good as new.)

Real-Time Simulation is interesting because it is atheoretical - even anti-theoretical. (On arriving at the EUROCONTROL Experimental Centre I was told "Here we do not do Science - we do experiments!") While 'real scientists' try to reason from the general to the particular, engineers tend to study the particular, and view any generalisation with deep suspicion. Publication is regarded as vanity, and conference attendance as an unearned holiday. The EUROCONTROL Real-Time simulator is reconfigured for each simulation. It can have as many as forty controllers and assistants, twenty-odd (some very odd) simulator pilots, and ten 'feed controllers',

who simulate the adjacent sectors from which traffic is received and to which it is sent. Radio, landline and intercom communications are recorded, as are all orders given to simulated aircraft. A vast mass of data is collected, most of which is filed and forgotten. Several alternative organisations are simulated, and controllers' opinions are collected using standard forms and 'de-briefings' after exercises. It is usually clear after a few weeks simulation which is the preferred alternative. In some cases, a new organisation is developed half way through the simulation, a textbook problem in confounding, learning and more or less unconscious bias.

Over the years, some conclusions have become clear. Controllers cannot give reliable estimates of times for tasks [3]. They are not computers, and do not apply reason except on very rare occasions. They recognise situations which are usually partly repetitive. Most flights take place from the same place each day to the same destination, using the same type of aircraft. Although the take-off time varies, the time in sector is fairly constant. Controllers recognise situations, and remember how they were solved. They do not extrapolate aircraft tracks. They can predict possible future conflicts where they have 'learned the sector' with an accuracy that is higher than the data shown allows. They plan their future workload, taking account of possible future problems. In essence, they are constantly simulating the future of the simulation. They do not wait for problems to appear, but adopt strategies to prevent problems happen-

ing. They ignore aircraft which have no potential problem, and cannot remember them when questioned immediately after work [1, 9]. Controllers judge systems as a whole, and make no allowance for minor errors - either it works or it doesn't. Measuring mental strain -, the effect of doing the job on the controller -, is very difficult. Electrophysiology and electroencephalography are difficult to apply to more than a few of the forty controllers, and very vulnerable to confounding factors, particularly any physical work. Secondary tasks, assuming that 'spare capacity' can be measured at all, assume that controllers have a constant maximum capacity and can switch tasks without cost. Neither of these is true [7]. Eye movement can now be measured fairly easily, if very expensively, but we still do not know why controllers look at a particular point. (It may be useful to find that they never do look at a particular display.) Analyses of hormones and other chemicals in blood or saliva are usually upset by unexpected interference. For instance, controllers appeared to secrete more catecholamine when not stressed. It turned out they had more time to drink coffee - which produced breakdown products confused with catecholamine [8].

Various models have been made of the Air Traffic Control process. Most of these are 'event-driven' assuming that controllers react to external events [5]. Controllers say that this state only occurs when they have 'lost the picture' and are on the verge of complete breakdown.

The lesson I have learned from thirty-odd years of looking at Air Traffic Con-

trol is that human operators are more complex, more fallible, more adaptable and more goal-oriented than most people believe, and that they can make faulty systems work - mainly by ignoring silly general rules [10].

Finally. I have learned to respect the men and women, controllers, pilots, engineers, air and ground crew and others who by their sustained and conscientious efforts, make a potentially dangerous system into the safest mode of transport ever known. If you do not have thirty years to spare try my book [6].

References

- [1] André Bisseret, *Représentation et décision experte: Psychologie cognitive de la décision chez les aiguilleurs du ciel*, Octares éd., 1995.
- [2] H David, *Human factors in air traffic control: a study of the ability of the human operator to predict dangerously close approaches between aircraft on simulated radar displays*, Ph.D. dissertation, Ph. D. Thesis, Loughborough Univ. of Technology, 1969.
- [3] H David, 'Data collection for the euro-control ground model.eec', in *Note No.5. Bretigny-sur-Orge, France: Eurocontrol*, (1972).
- [4] H David, 'The radar air traffic controller - a paradigm shift paper presented at contemporary ergonomics 1984', in *Proceedings of Ergonomics Society Annual Conference, 2 5 April 1984*, (1984).
- [5] H David, 'Use of microsaint for the simulation of atc activities in multi-sector evaluations', *Contemporary Ergonomics*, 213-218, (1997).
- [6] H David, *Control Room Simulation*, R+D Hastings: Amazon Books, 2017.
- [7] H David and C Noonan, 'Potential measures of strain on controllers', *EEC Report: EUROCONTROL Experimental Centre, Bretigny sur Orge, France*, (164), (1983).
- [8] F. Klimmer, H. M. Aulmann, and J. Rutenfranz, 'Urinary catecholamine elimination in air traffic control workers under occupationally induced emotional and mental stress', *International Archives of Occupational Health*, 30(1), 16, (1972).
- [9] MT Lafon-Millet, 'Observations en trafic réel de la résolution des conflits entre avions évolutives', *Rapport INRIA No CO/R/55*, (1978).
- [10] JT Reason, *The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries*, Ashgate Publishing, Ltd., 2008.

AISB 2019 Convention, Falmouth

Swen E. Gaudl (Falmouth University, swen.gaudl@gmail.com)

The 2019 Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, AISB 2019 will be held from April 16th to April 18th 2019 at Falmouth University, Cornwall. The convention is chaired by Swen E. Gaudl and Edward Powley. Additionally, the organisation committee consists of Tanya Krzywinska, Rob Saunders, Heidi Ball, Kamran Harandy and Michael A. Scott.

As in the past years, AISB 2019 will provide a unique forum for presenting cutting-edge research and burning issues around all areas of AI.

The theme for this year is “Artificial Intelligence, Imagination and Invention”.

The annual convention is a unique place for presenting and discussing research as well. The convention provides a forum fostering trans-disciplinary exchange and bringing together researchers, practitioners, artists and scholars from various backgrounds and countries. After this year’s selection process, we accepted a total of nine symposia and two workshops with topics ranging from machine learning and language acquisition to robotic dance performance, creativity and games. We are also excited that the proposals came both from national and international groups, demonstrating that research is not and should not be bound by borders.

Further on, we give a brief introduction to the events both work-

shops & symposia that will be held during AISB2019. For more information and for details on keynotes, exhibitions, travel and other social events which will complement the programme, as well as the full call for papers of the various symposia, please visit our website at <http://aisb2019.falmouthgamesacademy.com/>

6th Computational Creativity Symposium

Organised by Maximilian Droog-Hayes and Juan Manuel Alvarado López (Queen Mary University of London, UK)

Over the last few decades, computational creativity has attracted an increasing number of researchers from both arts and science backgrounds.

This symposium aims at bringing together researchers to discuss recent technical and philosophical developments in the field, and the impact of this research on the future of our relationship with computers and the way we perceive them: at the individual level where we interact with the machines, the social level where we interact with each other via computers, or even with machines interacting with each other. We invite submissions related to both theoretical and technical work on modelling creative systems which produce musical, pictorial or linguistic works that represent imaginative concepts. Relevant topics might include discussions about creativity as a source of imagination in relation to

multiple domains, or how a creative process can produce a novel and interesting result.

More general topics of interest:

- Novel systems and theories in computational creativity
- Evaluating computational creative systems, processes & artefacts
- Computational aesthetics theory
- Representational issues in creativity, including visual & perceptual representations
- Social aspects of computational creativity & IP issues
- Creative autonomy & constraint
- Computational appreciation of artifacts, including human artwork

We welcome participants to submit their work as a paper of up to 8 pages, or as a demo. As is the tradition of the Computational Creativity symposium, the best paper and the best presentation would each be awarded a prize.

10th AI & Games Symposium

Organised by Daniela Romano, David Moffat & Swen E. Gaudl (University College London, Glasgow Caledonian University, Falmouth University, UK)

The AISB AI & Games symposium celebrates 11 years since it first started in 2008 as AI & Narrative Games for Education. In the past years the symposium has been acting as a meeting place for researchers and practitioners from academia and industry who

are involved with the design, development and evaluation of AI in the context of games. In particular, the Symposium focuses on the application of artificial intelligence or intelligent-like techniques, frameworks and theories to the creation of intelligent games. AI can be used in any manner suitable in a game, from algorithms to making it more engaging, personalised, and/or interactive.

Example topics (research & practice) of the symposium; these can be applied to X (a game, or VR, or design process, or any form of experience):

- Use of AI techniques (e.g. planning, learning, evolution etc.)
- Design & engineering of AI components
- (Semi-)Automatic PCG
- Intelligent or adaptive player interaction
- AI for user analytics and/or player-modelling
- Agent path-finding and/or decision-making
- Games (or simulations) as a platform for building agents
- Environmental simulations
- Interactive narrative generation
- Intelligent Narrative Technologies
- Experimental AI

Authors could be specialised in: AI, machine learning, planning, narrative, education and training, media, multi-media, virtual reality and virtual experiences, game design and development, game interaction design, charac-

ters design, interaction design and evaluation for children and/or adults, and any other relevant area.

A poster and demo session is also held usually over coffee time. In all past years we had one invited speaker.

Decolonising the computational imagination: Human & machine creativity as situated practice

Organised by Eleanor Dare, Eleni Ikoniadou, Claudia Dutton and Laura Ferrarello (Royal College of Art London, UK)

Projects such as the Leverhulme's CFI (Centre for the Future of Intelligence) Value Alignment (2018) aim to prevent AI from acting in ways which are irreconcilable with our explicit moral values, but what about our unconscious values, our implicit assumptions and inherited ideas about what constitutes imagination and invention?

This four-part workshop will address the complex theme of how we can we model an artificial, or even a natural inventive intelligence, until we've understood our own assumptions about what that means – to imagine something new, if, indeed, "newness" (Lim & Oyama, 2014) is an inevitable component of human and machine creativity. The practical goal(s) of this workshop will be to generate collaborative insights, creative performances and design methodologies addressing the theme of decolonising AI, the workshops are a positive opportunity to re-think commonly held assumptions about human and machine creativity, to engage with broader practices and to understand

the situated nature of all computation. We will make the case for decolonising human-machine creativity - as both an ethical and a pragmatic imperative.

The 4-part workshop will use practical exercise to uncover the assumptions, processes and structures embedded in all computational processes which claim to emulate human inventiveness. The workshop will explore the extent to which those models and processes are culturally specific, bringing to the surface our historically situated notions of what constitutes human intelligence, and, by extension, what constitutes human subjectivity and creativity.

This workshop will present a series of calibration exercises, performative engagements, dialogue and speculation, designed to reveal to participants their implicit and explicit notions of creative intelligence and inventive insight. Each RCA academic (who has a research specialism addressing aspects of AI), will present a 45-minute interactive/collaborative exercises.

Philosophy after AI: language, imagination and creativity Symposium

Organised by Giusy Gallo and Claudia Stancati (University of Calabria, Italy)

This symposium aims for a philosophical approach to the latest issues about the study of human mind developed in the field of Artificial Intelligence. The goal of the second symposium Philosophy after AI is to investigate the philosophical roots of imagination and creativity and the role they play in AI researches and/or, con-

versely, how much they have changed (and are changing) after AI developments.

The first edition of the symposium has been devoted to themes such as mind and knowledge. Among the issues to be investigated the mind-body problem and the category of subject need to be deeply explored, from the ethical perspective encouraged last year till a new light given by the philosophical stance on imagination and creativity. The mind-body problem can be read again involving the neuroscientific research, including the provocative theory of the extended mind. The second question is about the philosophical category of subject: how to set out the boundaries of the self? How is the concept of subject together with the concept of imagination involved in AI researches? Is human imagination due to be replicated? Are technological developments affecting human creativity in long term? The third issue concerns the nature of learning and creativity and the current research in the field of machine learning. The development of AI asks for the role of the advancement in such field plays in studies devoted to language, including the helpful effect on people with disabilities. Moreover, we should follow the way machines implement human language (e.g. Siri, Cortana, ...): could machines and human beings understand each other?

Recently AI researchers are developing autonomous machines which can exhibit behaviours such as prejudice, identifying, copying and learning them from a big amount of data. This kind of “understanding” and “simulation” requires a philosophical attitude on the

power of imaginative dimension of prefiguration of behaviour.

We invite talks on the following topics (non-exclusive):

- Linguistics and AI
- AI research on language
- The role of imagination in AI research
- Philosophy, science and AI
- Mind-body problem and AI
- Truth, post-truth and AI
- Language & cognition
- Learning, creativity & AI
- Creativity, machine-learning & language
- Social media, devices & human sociality

Social Interactions in Complex Intelligent Systems (SICIS)

Organised by Stefania Monica and Federico Bergenti (UNIVERSITÀ DI PARMA, Italy)

A complex intelligent system (CIS) is a large network of interacting agents where non-trivial global patterns and behaviours emerge, normally without a central control, from the combination of simple behaviours of individual units. Social interactions in CIS give rise to collective properties that hold at the macroscopic level, such as the formation of polarised opinions or the appearance of trends and subcultures, whose emergence cannot be easily inferred from the analysis of the behaviour of single agents at the micro-

scopic level. The study of complex intelligent systems represents a novel approach to investigate how social interactions among agents lead to emergent behaviours which exhibit some sort of intelligence.

Methodologically, social interactions in CIS accommodate both local and global phenomena, and therefore they are a key concept to understand the behaviour of a complex intelligent system. Social interactions are also the key concept to explain how real and artificial societies behave. The models used to study interactions among agents may derive from various fields, such as statistical physics, information theory, and non-linear dynamics. They are used to describe the effects of interactions among agents from a microscopic point of view, and the derivation of observable behaviours of the system may be addressed using various approaches, such as analytic and simulative tools, statistical methods, and empirical observations. Sophisticated research methodologies are being developed and used in the analysis of social complex intelligent systems, including graph theory, bifurcation diagrams, network analysis, agent-based modelling, theoretical physics, non-linear modelling, and computational models including cellular automata, and multi-agent systems.

Complex intelligent systems (CIS) and models of social interactions are used to describe processes in various fields, such as Artificial Intelligence, Computer Science, Mathematics, Biology, Economics, Physics, Sociology, Economy, and many others. Hence, they represent a promising

multi-disciplinary research field. The symposium is meant to offer an interdisciplinary forum on all aspects related to social interactions in complex intelligent systems. It aims to stimulate discussions and synergies among participants, which are expected to have diverse and complementary research backgrounds.

This will be the third edition of the symposium. The structure of the symposium involves talks given by authors of accepted papers.

Topics of interest to the symposium include, but are not limited to:

- Collective intelligence and co-operation
- Interacting agents and emergent behaviours
- Self-adaptiveness and self-organization in CIS
- Opinion dynamics in CIS
- Cultural dynamics
- Social consensus and agreement
- Social simulation and agent-based simulation
- Social networks analysis and simulation
- Mathematical analysis of CIS
- Mathematical and physical models of CIS
- Bio-inspired analysis of CIS
- Game theoretic and economical models of CIS
- Emergent properties and behaviours in CIS

Symposium on Analytics-based (Cognitively-enabled) Social Systems

Organised by Aladdin Ayesb, Indrani Lahiri, Miguel Arevalillo-Herráez and John Bishop (De Montfort University, Goldsmiths, UK, University of Valencia, Spain)

Several technological advances in recent years make the presence of systems in our life persistent and ubiquitous, whether in hardware form such as phones or software form such as social media that is embedded in the various devices we use, e.g. TV. This was driven by and equally drove further the advances in cognitive systems research. This enabled two things to happen: greater amount of data to be generated and new ways of socialising. These two events feed back into the development of cognitive systems in the forms of Data Analytics and Social Systems, providing a full loop of interaction between technological development and daily societal impact. Terms such as cyberbullying, personalised ads, analytics, and many others are becoming part of the daily language and news bulletins. Whilst technology gives rise to social issues and anxieties, e.g. meddling into elections, these social issues impose challenges and expectations especially from Artificial Intelligence, to resolve these issues, e.g. identifying fake news or unlawful content in social media platforms.

This symposium aims to look at this new breed of systems from a multidisciplinary viewpoint, examining, amongst others, the social impact

and its effects on pushing the boundaries of developing such systems.

Topics include but are not limited to:

- Personalisation, Machine learning and AI
- Quantified self and data cultures
- Mobile and locative services
- Social media, politics & big data
- Cyberpsychology
- Ethics & Privacy in Social Systems
- Data-system sustainability
- Cybersurveillance & IoT
- Social & cognitive theories tested
- Social bots and the management of sociality
- Social data collection and mining
- Social recommender systems and social robots
- Analytics applications
- Text analytics and language development
- Smart cities and smart services

Explainable Artificial Intelligence

Organised by Serge Thill and Maria Riveiro (Radboud University, NL, University of Skövde, SE)

AI systems are increasingly present in everyday society, from simple computer systems to agents such as autonomous vehicles or social robots. In this context, several researchers have noted that it is critical to understand

how human users perceive such systems - in particular, the degree to which they understand how the system works, and what mental models they build of the underlying algorithms. "Explainable AI" (XAI) thus refers to AI systems that behave or provide the necessary information so that their working becomes comprehensible to the human user. The need for transparency and explanations in AI-based systems to support interpretability, understandability, and trust has been highlighted recently by multiple authors from disparate disciplines in both AI and HCI communities, and is, for example, the focus of Google's People & AI research initiative.

Given the breadth of fields in which the notion of explainable AI turns up in one form or another, there are varying interpretations of what the concept really entails. To a traditional AI researcher, for example, explainable AI is often the opposite of a black box system. Similarly, in the situation awareness literature, the term "system awareness" (or sometimes system transparency) is used to describe to what degree a user understands, that is to say, has a reasonable model, of the inner working of the decision support system they interact with.

In the cognitive sciences, meanwhile, the interest is two-fold, covering both how humans understand other humans, and how these mechanisms can be mapped onto machines. In particular, Theory of Mind (ToM) refers to the human ability to infer the (internal) mental states of other human beings. Insights from the cognitive sciences have then be used in, for example, robotics,

to argue that it is necessary for machine systems to tap into the human ability to understand other agents through such simulation mechanisms. This is similar to long traditions of research in, for example, Human-Computer Interaction, which asks how humans perceive their environment in order to provide effective interfaces for computer programs.

Similarly, the application range is varied. In the automotive domain, it can lead to appropriately calibrated trust in an autonomous vehicle's abilities, similar to how appropriate system awareness can manage trust and also cognitive load in the domain of decision support systems. For a cognitive scientist, understanding how to design intuitively understandable AI systems can lead to a better understanding of the human mind itself. It is clear from the description above, that XAI research is a very varied field whose precise meaning and core missions are interpreted differently by different researchers.

The purpose of the present symposium is to bring together researchers from all aspects of XAI, and to foster an exchange of the current state of the art while facilitating the development of synergetic connections between different sub-fields of XAI.

We are interested in a diverse symposium that can cover several aspects of XAI outlined above. We particularly encourage submissions around the themes of human expectations of intelligent systems, and, conversely, how such systems can create explanations that are relevant, and in line with those expectations. Additionally, work concerned with the evaluation of such systems is also relevant to this symposium.

Movement that Shapes Behaviour: Rethinking how we can form relationships with non-humanlike embodied agents

Organised by Petra Gemeinboeck, Rob Saunders and Elizabeth Jochum (Falmouth University, UK, University of NSW, The University of Sydney, AUS, Aalborg University, DK)

This symposium is a transdisciplinary forum for exploring the potential of movement for shaping the expressive and relational capacities of non-humanlike robots and how we perceive them as social agents. Social robots are expected to affect every aspect of our lives in the near future. Currently, the design of social robots in research labs often mimic humanlike or animal-like features, both in terms of how they look and how they behave. We believe, however, that movement and its expressive, relation-making qualities hold the key to widening the spectrum of how we can interact with robots, without relying on a human- or animal-like veneer.

The importance of movement in the simulation of behaviour can be traced back to early cybernetic experiments and artworks, such as, Grey Walter's tortoises and Gordon Pask's conversational systems. Similarly, Heider and Simmel's classic cognitive psychology experiments using simple animated geometric figures, demonstrated the potential of movement to generate social meaning. This symposium emphasises the importance of methods and practices from the fields of animation, choreography, dance, pup-

petry and theatre. Grounded in embodied knowledge, they offer valuable insights for embodied AI, e.g., working with movement as a material, embodying 'bodies', relation-making through movement dynamics, embodied perception, and kinaesthetic empathy.

This symposium will bring together researchers and practitioners to explore how movement qualities can enable an embodied agent to communicate non-verbally, take on a social presence, make connections or enact an identity without mimicking living creatures. The topic opens up a number of important questions and challenges for embodied AI: how can we access, apply or learn from the embodied, often tacit knowledge of movement experts? How can we effectively study people's subjective experiences and ability to connect or interact with such machine-like agents? How does a robot's movement abilities integrate with its perceptual and cognitive processes, to make sense of other agents and its environment? How could this embodied emphasis lead to an integrated enactive approach to human-robot interaction? We are particularly interested in contributions from researchers and practitioners developing interdisciplinary theories, concepts and/or approaches that can inform or directly tackle embodied, interactive experiences with machine-like agents.

Language Learning for Artificial Agents (L2A2)

Organised by Stephen McGregor, Katrien Beuls, Stephanie Gross, Brigitte Krenn, Friedrich Neubarth

and Thierry Poibeau (*Austrian Research Institute for Artificial Intelligence (OFAI), AU, Vrije Universiteit Brussel, BE, CNRS/École normale supérieure, FR*)

Our symposium will be a venue for presenting and discussing recent and ongoing work in the various areas related to the development of artificial linguistic agents. At the heart of this event is the topic of grounded language learning and the host of computational techniques that are being explored as solutions to this socially significant area. As such, we will solicit submissions detailing results, ongoing work, projections, descriptions of data, model learning procedures, and ideas that pertain to artificial linguistic agents and their emerging role in the world. Beyond the core artificial intelligence approaches involved in this area of research, a range of fields, including robotics, computer vision, cognitive science, developmental linguistics, and philosophy are entailed by this ambitious programme.

As members of the artificial intelligence community, we feel we are a part of the spirit of optimism and advancement currently prevalent amongst computational linguists. At the same time, as researchers with an awareness of the theoretical and philosophical issues surrounding computational approaches to language, we are sensitive to concerns regarding the way that some forthcoming technological developments might overlook important questions about the grounded, human aspects of language, and the unforeseen consequences that might arise from pursuing information engineering projects without due regard

for the social or environmentally situated aspects of natural language.

Paper topics may include, but are not limited to, the following:

- Representation learning
- Ontology construction
- Construction grammar
- Learning language from multimodal data
- Multi-lingual approaches to grounded language learning
- DL for grounded language learning
- Embodied approaches to NLP
- Modelling non-linguistic components of language learning
- Symbol grounding/ungrounding problems
- Language games
- Embodied conversational agents
- Human interactions with artificial agents
- Computational models of developmental linguistics
- Modelling language on multiple timescales
- Evolutionary computational linguistics
- Social considerations in developing artificial linguistic agents

With this in mind, we intend to offer this symposium as both a platform for the presentation of exciting new results and a forum for engaging with some of the hard questions that emerge at the boundaries of language technology and life in the world.

We anticipate our symposium to be a showcase for research from international authors researching a variety of related topics.

Intelligent Machines & Human Behaviour

Organised by Simon Wells, Kate Pangbourne and Hannah Bowden (Edinburgh Napier University, University of Leeds, University College London, UK)

This workshop will primarily take the form of paper presentations around thematic topics in the area of Intelligent machines and human behaviour. Accepted papers will be grouped into thematic sessions that incorporate extensive time for questions and discussion. The session will close following a town-hall discussion session, designed to facilitate the mapping of papers to the interdisciplinary landscape, and the development of future collaborations between participants.

We will solicit contributions from a wide range of relevant topics related to how AI can affect human behaviour. These can include, but are not limited to the use of AI in Captology, digital persuasion, behaviour change, gamification. We are interested not only in focused reports concerning research into the applications of these techniques to specific problems, such as within healthcare and transport behaviour, but also in more general consideration of the risks posed and benefits gained from application of these techniques within human society. Of specific interest are contributions addressing the dark side of these interactions, examining how techniques can be mis-

used and how such misuse can be defended against.

Our rationale for focusing upon this topic is as follows: Artificially intelligent machines are becoming increasingly prevalent in modern society and are likely to play an important, even ubiquitous, role in future everyday decision making. This is a trend that is likely to accelerate as new techniques for automated-reasoning and machine-learning are applied to decision making within real-world domains. That these machines will have a great impact upon human society is beyond doubt. There is the potential for such machines to improve nearly every aspect of human life, particularly when artificial intelligence can overcome the well known shortcomings in human decision making such as those identified by behavioural economists. Insights from behavioural economics are behind the rise of "nudge" initiatives, and are in themselves subject to a critique of their ethics. However there is also the potential for AI machines to act to the detriment of people. For every cancer successfully detected at an early stage, there could be a bank computer denying (or approving) a mortgage, or the consequences of an autonomous vehicle that makes a poor decision about whether to evade an obstacle or emergency brake. This is not to ascribe explicitly malicious intent, but merely to recognise that most current, and likely future, machine systems will be as imperfect as those who have created them. Additional complexities can stem from the interplay between intelligent machines and human society. A further layer of risk and complexity is added once humans with

malicious intent are included. Whilst a machine can be used to help recognise poor behaviours, for example eating excess junk food, and can in turn help manage that person's behaviour in order to form better habits, such an approach could be used in the absence of informed consent.

Thus the study of how machines, in particular intelligent machines that can learn to recognise behaviours and respond accordingly, interact with humans, and how the behaviour of humans can be directly or indirectly affected as a result, is a topic of timely and deep importance.

AI and Robotics Normative Spheres: Towards a Sustainable Society and Technology.

Organised by Dr Aurora Voiculescu, Dr. Jack Stilgoe, Prof. Alan Winfield, Prof. Susumu Hirato, and Prof. Norihiro Hagita (University of Westminster, University College London, UWE Bristol, UK, Chuo University Tokyo, Advanced Telecommunications Research Institute International, JP)

In the past decades, an increasing number of human intellectual activities, such as perception, recognition, decision-making, inference have been replicated through Artificial Intelligence (AI) technologies. AI actions, based on these intellectual processes, have led to such technologies being used in a multitude of support activities in businesses and services throughout the economy and society. Big data and machine learning have led to increased progress in machines offering "cognitive in-

sight", classifying information, identifying patterns, processing of natural language, to mention just a few. Moreover, intelligent machines share now physically, more and more, the same space as humans, with automated vehicles, care robots, surgical robots, hotel receptionists, becoming a common encounter. While the support that such AI and robotics technologies can bring to human activities is expanding at an ever-increasing rate, the normative – ethical and regulatory – environment needed for welcoming such technologies is evolving at a much slower pace and, with few exceptions, mostly in a reactive rather than a proactive manner.

The symposium sets out to create a platform of debate, as a regular recurring feature of the AISB Convention, inviting AI and robotics scientists as well as social scientists to engage critically in dialogue within a multidisciplinary environment.

The symposium organisers invite submissions on all aspect of the ethical and/or regulatory issues encountered within or outside AI and robotics labs, on issues including but not limited to:

- ethical codes and guidelines in AI R&D environments
- signals of the emergence of a new sense of ethics at the human/machine interface
- relevant dimensions in considering the balance between human decisions and AI-based decisions
- technologies for cyber security and privacy protection
- ELS implications of autonomous robots

- challenges derived from the AI enhancing of human senses and abilities
- the challenges of re-distributing responsibility for accidents involving physical or digital AIs
- access and exploitation of big data in relation to privacy and/or social justice
- rights and incentives in developing socially-mindful AIs
- from moral, to legal, to electronic personhood: ethical, legal and conceptual challenges
- types of responsibility at the human/machine interface
- machine ethics and law-abiding algorithms
- transparency – the white box/black box conundrum – normative implications

- ethical dimensions of stakeholder consultation in AI and robotics
- creating the social space for discussing alternative visions of the society–technology dynamics
- governance and regulation of new technologies – who does the shaping, owning, benefiting from it, accessing it and making decisions about it
- reflecting on the impact of AI technologies on society – where should we start? A top-down or bottom-up approach?

Participants in the workshop will be invited to submit contributions to a special issue of an interdisciplinary journal (Ethics and Information Technology; AI and Society and Connection Science are under consideration for SI publication in 2019).

Dear Aloysius. . .

Agony Uncle Aloysius, will answer your most intimate AI questions or hear your most embarrassing confessions. Please address your questions to fr.hacker@yahoo.co.uk.

Note that we are unable to engage in email correspondence and reserve the right to select those questions to which we will respond. All correspondence will be anonymised before publication.

Dear Fr. Hacker,

Celebrating the 150th issue, how would you retrospectively view your column for the Q over the years, with all the old and new

readers, the loyal and cynical ones, and the witty and/or geeky ones?

Yours, QED

Dear QED,

Congratulations on the 150th issue of AISBQ. Its longevity is largely due, I'm sure you'll agree, to my excellent columns.

The first **A.HACKER™** (Aloysius Has AI, Cognition & Knowledge: an Exceptional Robot!) article for Q was in issue 23, where I gave the catechism for the Logo programming language. My regular diary columns started, in issue 50, so I am also celebrating: this is my 100th

column! In issue 107 the diaries were replaced by my guide for young AI researchers and, then, from issue 122 by 6 episodes of my life history, briefly interrupted by an additional 'Guide' in issue 127. The current 'Dear Aloysius' Agony Uncle columns started in issue 129. Each column has described wittily-named, examples of our Institute's stunning, AI-based products. Each of these products has transformed society. The loyalty of our readers has been regularly rewarded - even the cynical ones are gratified to have their expectations confirmed. New readers have a treat in store in the Q back issues.

During the history of my columns, AI has been transformed from a little-known craft cooperative, where the World-wide research community could all meet in same conference room, to a multi-billion £ industry, requiring each of the World's governments to have an AI strategy. The Institute has always been well ahead of the AI frontier. It has now, though, become impossible to further exaggerate the potential impact of AI. So, my work is done.


My Institute's work has been inspired by its religion: The Church of God the Programmer. Its tenet is that we live in a simulated universe: **COSMOS**TM (Computer-Oriented Simulation of Mass-energy, Opinion and Space-time). People often ask me why I am so certain that The Programmer exists. I don't have a mathematical theorem QED, but the scientific evidence is overwhelming.

- Physicists have identified the very high resolution of the **COSMOS**TM simulation. The side of each voxel is 1.616229×10^{-35} metres (the Planck length), and the frame rate is 1.854884×10^{43} frames/sec (the speed of light divided by that length).

- The rules of Physics keep changing. Just as physicists think they have found a theory of everything, The Programmer 'refactors' **COSMOS**TM to reveal novel experimental evidence that contradicts the old theory and requires a new one of greater depth and complexity. This keeps the scientists busy and out of trouble.
- From relativity via quantum mechanics to dark energy, 'reality' is now revealed to be ineffable - beyond human understanding.

Our Institute's religion has repaid our faith by enabling our fantastic inventions to flourish. The Programmer achieves this via **MIRACLE**TM (Magical Intervention into Reality Achieves Contradictions of the Laws of Everything). This gives us a unique selling point - even our most geeky competitors find it impossible to duplicate our results without Her assistance.

Yours, Aloysius



Fr. Aloysius Hacker
Cognitive Divinity Programme
Institute of Applied Epistemology

Back matter

Articles may be reproduced as long as the copyright notice is included. The item should be attributed to the *AISB Quarterly* and contact information should be listed. *Quarterly* articles do not necessarily reflect the official AISB position on issues.

Editor

Dr. Mohammad Majid al-Rifaie
School of Computing and Mathematical Sciences
University of Greenwich, Old Royal Naval College
aisbq@aisb.org.uk

Advertising and Administration

Dr. Rob Wortham (AISB Executive Office)
Dept of Electronic & Electrical Engineering
University of Bath, Claverton Down
Bath, BA2 7AY, United Kingdom
admin@aisb.org.uk

AISB Patron

Prof John Barnden (*University of Birmingham*)

AISB Fellows

Prof Harry Barrow (*Schlumberger*)
Prof. John Barnden (*University of Birmingham*)
Prof. Margaret Boden (*University Sussex*)
Prof. Mike Brady (*University of Oxford*)
Prof. Alan Bundy (*University of Edinburgh*)
Prof. Tony Cohn (*University of Leeds*)
Prof. Luciano Floridi (*University of Oxford*)
Prof. Kerstin Dautenhahn (*University of Waterloo*)
Prof. John Fox (*Cancer Research UK*)
Prof. Jim Howe (*University of Edinburgh*)
Prof. Nick Jennings (*Imperial College London*)
Prof. Aaron Sloman (*University of Birmingham*)
Prof. Mark Steedman (*University of Edinburgh*)
Prof. Austin Tate (*University of Edinburgh*)
Prof. Mike Wooldridge (*University of Oxford*)
Dr. Richard Young (*University College London*)

AISB Committee

Chair: Dr. Bertie Müller (*Swansea University*)
Vice Chair: Dr. Yasemin J Erden (*St. Mary's University College*)
Secretary: Andrew Martin (*Goldsmiths, University of London*)
Treasurer: Dr. Rob Wortham (*University of Bath*)
Editor: Dr. Mohammad Majid al Rifaie (*University of Greenwich*)
Publications: Dr. Floriana Grasso (*University of Liverpool*)
Membership: Stephen McGregor (*Queen Mary University of London*)
Equality, Diversity and Inclusivity: Dr Anna Jordanous (*University of Kent*)
Industry Liaison Officer: Dr. Ed Keedwell (*University of Exeter*)
Loebner-Prize Officer: Dr. Nir Oren (*University of Aberdeen*)

Contents

Editorial	3
The AISB in 2018 and beyond <i>B. Müller</i>	4
A European Vision for AI <i>H. Hoos, et al.</i>	6
Holding Our Brea(d)(th) <i>J. A. Barnden</i>	14
It's Simulation, Jim, but Not as We Know It <i>H. David</i>	15
AISB 2019 Convnetion <i>S. E. Gaudl</i>	18
Dear Aloysius . . .	29

The AISB Quarterly is published by the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). AISB is the UK's largest and foremost Artificial Intelligence society. It is also one of the oldest-established such organisations in the world. The society has an international membership of hundreds drawn from academia and industry. We invite anyone with interests in artificial intelligence or cognitive science to become a member.

AISB membership includes the following benefits:

- Quarterly newsletter.
- Electronic subscription to *Connection Science* published by Taylor & Francis.
- Student travel grants to attend conferences.
- Discounted rates at AISB events and conventions.
- Free attendance of Members Workshops.
- Discounted rates on various publications.
- A weekly email bulletin and web search engine for AI-related events and opportunities.

You can join the AISB online via: <http://aisb.org.uk>