AISB QUARTERLY

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



No. 148

November 2017

About the Cover

Cover artwork by Rebecca Johnson, @becmjo (Goldsmiths, University of London)

This cover utilises Dispersive Flies Optimisation [1] to perform a simple style transfer between two input images of equal dimensions. The first image features a centred 'Q' and the second image features a form of vector field [2], a technique to simulate fluid motion, to serve as the underlying background.

Upon initialisation, each fly was assigned a candidate solution in the form of random coordinates within the dimensions of the background image. An individual fitness was then calculated from the pixel value within the 'Q' image correlating to the fly's position. A threshold of the fitness was calculated to determine whether the fly's coordinates resided within the area of the 'Q' in the first input. If this fitness exceeded the threshold, a second individual fitness was calculated for the specific fly. This fitness was established by the pixel value correlating to it's coordinates within the background input. Another threshold was set to determine whether the fly's location corresponded to an area within the vector field. If this second fitness exceeded the threshold, the fly then converted the vector field pixel to a darker colour based upon the strength of it's fitness.

Similar to nature, the swarm faced a dynamically changing environment where the background image pixels were altered dependent certain conditions. By using a high disturbance threshold, the swarm was provoked to continually explore all areas of the search space. As a result, the swarm was able to progressively alter all desired pixels of the vector field to create the shaded 'Q' within.

Check the link below to see the process through which the cover is generated: https://youtu.be/l8RkwJcpfCY

References

[1] al-Rifaie, M.M., 2014, Dispersive Flies Optimisation. In Computer Science and Information Systems (FedCSIS), 2014 Federated Conference on (pp. 529-538). IEEE.

[2] Aminov, Y., 2000, Geometry of Vector Fields, Gordon and Breach Science Publishers.

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.

Editorial

After the reading week, in the beginning of one of my Natural Computing lectures, a student raised his hand and asked: "Would you think swarm intelligence and evolutionary computing techniques are any better than machine learning algorithms in giving us insights on how a solution is reached?"

The question was triggered following my recommendation of a recent article in the Guardian "*Computer says no:* why making AIs fair, accountable and transparent is crucial¹", which I had suggested as a bed time reading during the week.

Students and, to some extent, researchers and academics are often more interested in finding solutions rather than necessarily the reason behind how solutions are identified. There is clearly a shift in paradigm which is seeing more computer scientists (and, as always, psychologists) from various backgrounds updating their approach and priorities to adhere to the much needed change.

Another relevant matter, which could impact the long term role of the United Kingdom and its relevance in the ever growing world of artificial intelligence, is the House of Lords' hearing of the Artificial Intelligence Committee that was held on Tuesday 10 October 2017². It is useful to hear the sometimes contradicting and at times complementary accounts of the AI scientists and technology journalists in the two consequent hearings.

We are likely to see some of these debates arise in the symposia of the AISB2018 convention, whose summary is presented in this issue, along with a rich set of articles. One of them highlights the positive implications AI could have on tackling modern slavery; in another piece, the long standing "battle" between humans and machines is discussed through playing poker. Additionally, a personal perspective on chatbots is presented, discussing their origins. Finally, some conference reports detail the highlights of the event, along with a take on the World Summit AI.

As always, the concluding words are reserved for our very Fr. Aloysius Hacker.

Mohammad Majid al-Rifaie Editor @mohmaj

Goldsmiths, University of London London, United Kingdom 20 November 2017

¹https://www.theguardian.com/science/2017/nov/05/computer-says-no-whymaking-ais-fair-accountable-and-transparent-is-crucial

²http://www.parliamentlive.tv/Event/Index/073717ca-484b-4015-bd10f847cea3f249

The heart inside the machine: Using AI to combat modern slavery

James Ollerenshaw, james.ollerenshaw@digitalreasoning.com

The incredible story of how artificial intelligence has connected a nonprofit, governments, and the banking industry in the global fight against human trafficking.

Introduction

"I've seen video content of a child that's the same age as mine being raped by an American man that was a sex tourist in Cambodia. And this child was so conditioned by her environment that she thought she was engaging in play."

With palpable emotion, Ashton Kutcher delivered these words to the Senate Foreign Relations Committee hearing on human trafficking and slavery³. His impassioned testimony generated headlines around a world where, hidden from view, 27 million adults and 13 million children⁴ are victims of modern slavery. Any notion that slavery was disgrace of history is crushed by the realization that there are more slaves today than at any point in our past⁵.

The dark side of Internet

Kutcher may be best known as an actor, but for anyone doubting his

reason to speak on this issue he had a formidable rebuke. As chairman and co-founder of Thorn, he told the committee, his "day job" is to lead a nonprofit organization committed to using technology to fight human trafficking and the sexual exploitation of children.

Thorn advances developments in technology to tackle the problem of modern slavery because, in common with many other industries, much of the market has moved online. Demand for prostitutes means almost 60⁶ percent of trafficked humans are sold into the sex trade. Most are under the age of 24^7 , but some (as Kutcher's example underlines) are considerably younger. Online advertising is the contemporary equivalent of a seedy back street or phone box stickers, but reaches a much larger audience. When Thorn surveyed underage sex trafficking victims, 3 out of 4 said they had been advertised online.

The scale of the market is hard to fathom. Every day, more than 100,000 escort ads are posted to the internet in the United States alone. Law enforcement faces a massive challenge in identifying both victims and perpetrators within this enormous and ever-

```
^{3} \rm http://www.digitalreasoning.com/partnering-to-combat-human-trafficking <math display="inline">^{4} \rm http://www.acrimesomonstrous.com/
```

 $^{^{5}}$ See footnote 4.

⁶https://www.unodc.org/documents/data-and-analysis/glotip/Trafficking_in_ Persons_2012_web.pdf

⁷http://www.newsmax.com/Pre-2008/Millions-Suffer-Sex-Slavery/2001/04/23/id/ 661571/

changing pile of data. Finding ads that indicate a trafficked victim among those which are legitimate requires trained officers. To say that the volume of the data has overwhelmed the available human resources is an egregious understatement.

Hidden in plain sight

Law enforcement had already adopted technology to help mitigate the task, applying keyword analytics to sift through escort ads and highlight dubious language. The problem with this technique is that human traffickers deliberately avoid using words or terms that might trigger an alert. Meanings are implied, spellings are distorted, and ads are continually edited and reposted to make it hard for officers to keep track.

Keyword analytics solves the problem of churning through huge volumes of data, but without any real understanding of meaning or context it frequently reports false positives while all too easily missing what is significant. This state of affairs leaves trafficked minors hidden in plain sight and law enforcement with no option but to commit to a complex and highly manual investigation task.

This type of computing problem had long irritated Tim Estes. Frustrated with the inability of computers to learn from data, he'd been inspired to build software that could understand the semantics of human language and how that are influenced by context. His company, Digital Reasoning, was formed with a mission to use this ability to help humans solve the world's toughest problems.

Digital Reasoning found its first customers in the shadow of the 9/11 attacks on the United States. Working with the defense and intelligence communities, the company proved that machines could learn human communication and flexibly extract relevant details without relying on keywords or lexicons. Ongoing development led to software able to read vast amounts of communications data, learn what's important, and cross-compare between multiple sources. These capabilities enabled it to report insights, patterns and relationships not apparent to the human eve.

From the battlefield to banks

Success in intelligence and on the battlefield attracted interest from other quarters. Rocked by scandals of market fixing and rogue trading, the financial services industry was keen to find better ways to identify corrupt employees. "Both Wall Street and the intelligence world want the same thing: to find unknown unknowns in the data," said Roger Hockenberry, the former chief technology officer of the Central Intelligence Agency's Directorate of Operations.

Pressure from regulators, billions in fines, and substantial reputational damage had caused banks to swell their compliance organizations. Despite this, finding the bad apples in businesses with many thousands of employees was a major challenge. Compliance officers knew that colluding employees used email and chat messages to plan and enact their crimes, but they faced similar challenges to their peers in law enforcement. Inferred meanings and coded terms made it difficult for keyword analytics to spot the messages that mattered. Compliance teams had little option but to manually review thousands of alerts, most of which would be false positives.

By this stage, Digital Reasoning had honed its technology into an artificial intelligence platform, Synthesys, that could be flexibly applied to analytics tasks where human communication was the main data source. Unlike keyword based solutions, Synthesys emulates a human-like ability to make sense of human communication in context. It learns as it works, allowing it to cope with nuance, aliases, coded terms, and ambiguity. What had proven to be effective for national security also worked for banks, with Synthesys piecing together insights that could not be extracted using conventional technology.

Artificial intelligence helped in two key ways. First, deeper insights and greater accuracy proved to be much more effective at disclosing which employees posed a threat, protecting the banks from future regulatory censure. Second, the technology was a huge time saver for compliance officers. Using the intelligent assistants built into Synthesys, insights were automatically resolved to specific entities (people, places, objects, events), focusing attention on the problem's source and visualizing the hidden connections with other employees. This combined impact was just what compliance organizations needed. Within the space of two years, Synthesys became the most widely adopted artificial intelligence technology among leading investment banks.

Finding finance's felons

Miscreant employees were not the only headache for compliance chiefs. Fines for failures to prevent money laundering were another source of huge cost and embarrassment for the industry. Virtually every major bank has been hit by millions, sometimes billions in fines⁸. Many were quick to see the potential for using artificial intelligence to root out financial crime.

This was another situation where human communication in immense volumes was an impediment to pro-Despite acting in accordance gress. with know-your-customer or KYC regulations that were created to prevent criminals exploiting the financial system, banks were struggling to confirm customers' identities and intent, and to reveal any hidden relationships or undeclared interests. Inadequate knowledge allowed money launderers to circumvent banks' defenses and disguise the source of their funds by using false identities, financial "smurfs" (people who use their account to make many small deposits and transfers on behalf of the money's true owner), and by establishing shell corporations.

Financial crime and compliance chiefs believed that that data analytics might hold the answer. Barry Koch, formerly Senior Vice President and Chief Compliance Officer for The West-

⁸http://www.int-comp.com/ict-views/posts/2016/07/22/top-5-money-launderingcases-of-the-last-30-years/

⁹http://www.osce.org/secretariat/115618?download=true

ern Union Company, described a process⁹ for using financial traits that indicate risk to identify potential criminality. While many of the attributes and characteristics he defined can be gleaned from transactional data, Koch highlights the importance of trawling public source information to improve the quality of the analysis. Useful sources for evidence of aliases and criminal interests include news sites, companies records, and court reports.

With a limited number of trained officers to call on, compliance organizations had focused their resources on customers deemed to be a higher risk. However, being based on flawed knowledge, risk scores were not reliable. Moreover, the time taken to complete investigations meant that illegal activity could continue for months, sometimes years before being spotted. KYC was a thoroughly broken process, but this time the banks could turn to artificial intelligence to solve the analytics problem.

The link to human trafficking

Human trafficking is typically understood as a human rights issue, but this assessment tends to overlook the financial incentives that drive it. Estimates of its profits range from \$9 billion to \$31.6 billion¹⁰. A study in the Netherlands found that a pimp's average earnings from a single sex slave were \$250,000 per year¹¹. The Panama Papers exposed the magnitude of financial crime and the truth about how illegal funds are generated, among them the profits from human trafficking¹². Money launderers give these profits a veneer of legitimacy, allowing them to be moved into the global banking system.

Governments and intergovernmental organizations had enacted legislation to target human trafficking, but many nations have failed to implement these laws and among those which have the number of convictions remains low^{13} . Baroness Goudie, a member of the United Kingdom's House of Lords and chair of the United Nation's Women Leaders' Council to Fight Human Trafficking, argues¹⁴: "Ultimately, no number of declarations will end the business of human trafficking. We need real investment from governments, and the involvement of people who are trained not only to spot human trafficking but also to follow the money."

Tougher legislation is starting to emerge. The UK Prime Minister, Theresa May, has promised to "lead the way in defeating modern slavery." Focusing on the profits of human trafficking is a key part of the obligations set out in the country's Modern

¹⁰https://www.unglobalcompact.org/docs/issues_doc/labour/Forced_labour/HUMAN_ TRAFFICKING_-_THE_FACTS_-_final.pdf

¹¹See footnote 4.

¹²https://panamapapers.icij.org/20161201-global-impact.html

¹³http://psm.du.edu/media/documents/us_research_and_oversight/csr_reports/us_ crs_rl34317.pdf

¹⁴https://www.theguardian.com/global-development/poverty-matters/2013/jul/ 15/slavery-industry-money-human-trafficking

Slavery Act. This stipulates that any firm with a turnover above £36 (\$44) million must publish an annual statement¹⁵ of the steps it has taken to ensure that modern slavery is not occurring in its business. Crucially, this extends beyond the UK and includes all parts of its overseas supply chain. The threat of fines and reputational damage has sent bosses scrambling. As Edward Naish of PwC warns, "Imagine waking up to news reports that your business is taking no steps to combat slavery. It doesn't bear thinking about."

As Goudie points out, targeting the money is the most effective means of undermining modern slavery, but when that money is laundered it requires the analytical capabilities of artificial intelligence to expose the true identities of people and organizations linked to these crimes. Mounting pressure on businesses and banks to make more thorough checks on their customers and suppliers is helping to frustrate human traffickers, at least at the more senior levels of criminal networks. However, in a heartening twist of providence, the very artificial intelligence being used to root out financial crime in banks and business was about to fight human traffickers head on.

A solution for law enforcement

With Synthesys being used to defend national security and bring integrity back to banking, Digital Reasoning was making progress on its mission to use technology to help solve the world's toughest problems. However, when the firm's executives met Ashton Kutcher's organization, Thorn, and learned how law enforcement was being engulfed by data volumes, they could not ignore the potential of Synthesys to deliver a solution.

Experiments with artificial intelligence were not in law enforcement budgets, but Digital Reasoning and Thorn formed a partnership to explore how Synthesys could be used to respond to the challenge. Supported by donated time and funds from sponsors, the team worked to create an intelligent assistant that could take on much of the rote work and complexity of analyzing escort ads.

The result was Spotlight: a webbased tool that provides law enforcement with intelligence about suspected human trafficking networks. Using the artificial intelligence capabilities of Synthesys, it uses text analytics and computer vision to automate the process of identifying which escort ads are for trafficked individuals.

An outsized impact

To date, Spotlight has assisted in more than 8,300 investigations conducted by 780 law enforcement agencies in all 50 US states. It has contributed to the identification of 6,625 victims and, by revealing hidden relationships in the data, has also help officers unmask 2,255 pimps. By bringing a human-like acuity to big data analytics, Spotlight

¹⁵http://pwc.blogs.com/industry_perspectives/2016/09/uk-business-with-aturnover-over-36m-have-you-prepared-your-modern-slavery-statement-not-anobligati.html

has accelerated investigation times by 60%.

"Our mission matches the magnitude of the internet, but by working with Digital Reasoning we found a way to empower law enforcement to identify and rescue trafficked children by turning huge volumes of data into an asset," said Julie Cordua, CEO of Thorn.

Working with banks, first to tackle internal rule breaking and later to identify and disrupt the activities of money launderers, has enabled Digital Reasoning to continue to support the development and maintenance of Spotlight. Now both Thorn and Digital Reasoning are eager to expand the use of Spotlight beyond the US. The UK government is looking at the technology and this would provide an anchor point to expand its use across Europe.

The fight to stop modern slavery and human trafficking will not be easily won, but artificial intelligence is empowering law enforcement, governments, and the banking industry to find the sources of these crimes and bring their perpetrators to justice.

Conference Report: International Joint Conference on Artificial Intelligence (IJCAI) 2017

Kacper Sokol (University of Bristol, k.sokol@bristol.ac.uk)

This year the 26th International Joint Conference on Artificial Intelligence was held in Melbourne, Australia between August 19th and 25th. The conference attracted over 2540 submissions and nearly 1800 attendees. The programme was split into four days of workshops, tutorials and special sessions; and three days of the main conference track.

The range of tutorials was broad and covered topics such as Argumentation in Artificial Intelligence: From Theory to Practice and Interactive Machine Learning: From Classifiers to Robotics. Workshops offered insights into interesting research areas, for example: Fourth International Workshop on Theory and Applications of Formal Argumentation and Explainable AI Workshop. The latter has proven to be one of the most attended. Finally, special sessions included a number of collocated conferences and Doctoral Consortium, where I published my 2-page extended abstract titled "The Role of Textualisation and Argumentation in Understanding the Machine Learning Process". The Doctoral Consortium participants were given an opportunity to present their work in a 3-minute elevator pitch and later discuss their ideas during one hour poster session.

The main conference started with an evening reception at the Melbourne Cricket Ground. The event was a great networking opportunity and featured live entertainment and animals native to the land Down Under.

This year the main conference theme: Autonomy and AI, aimed at encouraging "debate and analysis of the limits and safeguards that must be established when giving AI systems more autonomy". The conference was opened with a warm welcome from a local indigenous Elder and the organisers. The scientific programme was divided into multiple thematic tracks e.g. Natural Language Processing, Data Mining and Personalization, which allowed attendees to focus on their favourite topic.

Keynote and Panels

In-between the variety of events happening during the conference, all the attendees gathered for the plenary sessions. The keynote, which opened the scientific programme of the conference. was given by Stuart Russell - Professor of Computer Science and Smith-Zadeh Professor in Engineering at UC Berkeley – probably best known for his textbook in AI: "Artificial Intelligence: A Modern Approach". He addressed the conference theme with his talk on "Provably Beneficial AI" presenting recent achievements in the field and showing important directions for future research that will benefit deployment of safe AI. Furthermore, he talked about his work in probabilistic programming and its application to global seismic monitoring with NET-VISA model developed by his research group. They have used NET-VISA as a tool for nuclear treaty verification, outperforming currently deployed system¹⁶.

The conference programme was also rich in panel discussions on wide range of topics:

- AI and Autonomy: Current Opportunity or Future Threat?
- AI and Societal Challenges
- AI in 2027

All of them were set to discover views from different angles on the main conference theme. They provided an opportunity for both the panellists and the attendees to discuss current limitations of AI systems and speculate about the future development possibilities. All in all, the conference programme was broad and rich in technical sessions, invited talks and AI competitions.

Textualisation and Argumentation in AI

My paper – "The Role of Textualisation and Argumentation in Understanding the Machine Learning Process" – was published in the Doctoral Consortium track and was given a short oral and poster presentations. The talk highlighted the importance of transparency and explainability of machine learning systems and introduced the concept of textualisation – a natural language narrative describing selected aspects of

 $^{16}\mathrm{NET}\text{-}\mathrm{VISA}:$ Network Processing Vertically Integrated Seismic Analysis, Arora et al.

the machine learning process (data, model, predictions). Moreover, textualisation put within a human-machine message exchange system powered by argumentation theory could lead to machine learning models and predictions arguing persuasively for their choices. Therefore, allowing human operator to learn selected aspects of the system by querying its implicit logic.

In a broader perspective, attending IJCAI helped me to network with researchers, otherwise difficult to reach. I met people who are the main contributors to the research areas that I am most interested in. These interactions allowed me to share my research vision and receive important feedback and directions for future work. The conference as a whole helped me to establish links for possible future collaborations, refine my research plan and make it more concrete. To sum up, I came back full of new ideas, energy and passion to continue my research and excel at it. Overall, IJCAI 2017 was an excellent conference and a valuable learning experience.

Acknowledgements: Attending IJ-CAI'17 would be impossible without a Travel Award kindly provided by AISB, for which I am grateful. Furthermore, I would like to express my gratitude for the financial support provided by the Artificial Intelligence Journal and the University of Bristol and its Alumni Foundation.

Report: World Summit AI

Patrick Dhondt, patrick@artistdetective.com

I found myself wading through a crowded silo in Amsterdam, its dark interior only interrupted by neon strobes and an overbearing sound of music that undermined any attempt at verbal communication. The attendees, a balanced mix of young adult men and women, were heavily invested in their smartphones, or were otherwise engaged in private conversations in dimly red-lit corners. One might think I had walked into a disco club, but no, this was the setting of the World Summit AI, held on the 11th and 12th of October on the grounds of a repurposed gas factory. The other half of the silo stood in stark

contrast to the exhibition area, with a spacious centre stage where the most prominent talks were featured against the backdrop of huge overhead projector screens, and offering seats for a good half of the 2200+ visitors. Outside, barracks featured four "streams" of lectures, categorised in practical application, startups, tech talks, and workshops.

The summit started off with a panel discussion with representatives of four powerhouses: Ralf Herbrich of Amazon, Robert High of IBM Watson, Nocila Morini-Bianzino of Accenture and Zhang Tong of Tencent, on how they were pushing the frontiers with AI.

Xian-Sheng Hua of Alibaba Cloud explained how they were using their Deep Learning platform in combination with video and image recognition to perform real-time traffic predictions and recommendations. They have tested their system in two Chinese cities in 2016, tying into roadside cameras and traffic lights, with significant improvement in traffic flow made. They intend to eventually expand their services globally.

Gary Marcus, known AI critic and nonetheless founder of the machine learning startup Geometric Intelligence that was acquired by Uber, gave a pleasantly down-to-earth talk about the current abilities and inabilities of AI. The advancements that he acknowledged were speech recognition, image recognition, natural language understanding and targeted advertising, although each in narrow domains (e.g. speech recognition in quiet environments). The areas that he considered little progress made in were conversational interfaces, automated scene comprehension for the blind, domestic robots and reliable driverless cars. The main takeaway was that we are not as close to strong AI as some people think, particularly in reference to Deep Learning. Marcus illustrated the point that Deep Learning can not handle much deviation from its training data, showing examples of image recognition gone wrong, like a bottle mistaken for a pen just because it had an uncommonly shaped cap. He believes that we should look more at how children learn and know things, and look into psychology.

Marcus' view on the limitations of Deep Learning were recently echoed by Geoffrey Hinton, one of the pioneers in that field, who has suggested that we should turn our backs on backpropagation in Neural Networks and invent entirely new methods.

Steve Chen of NASA's Jet Propulsion Lab talked about the various possible applications of AI in space exploration that NASA was working Autonomous space probes and on. rovers could explore farther reaches of space without waiting hours for instructions from Earth, also allowing them to adjust for local weather conditions like storms. Other daring plans included pinpoint-accurate harpooning and hitchhiking on comets, and employing self-replicating robots to 3D-print more robots from locally found materials.

Andy Peart and Lucas Willering of Artificial Solutions introduced and demonstrated in tutorial-like fashion how to use their chatbot platform to build "intelligent user interfaces that talk like humans". As a former participant in the AISB's Loebner Prize Turing Test, this subject interested me in particular. Their demonstration guided us through training the system with a few variants of the questions "Tell me about the movie X." or "Who was the lead character in X?", after which it could instantly handle other questions of similar meaning vet different words. due to a vast collection of word association data behind the scenes. These questions then triggered responses that had to be manually set up, like retrieving the name of an actor and inserting it in a prewritten answer. If

one then asked a follow-up question like "Did he win an Oscar?" (again trained with variants), one had to set that "he" or "she", at this point in the conversation flow. referred to actor names. While I was impressed with the flexibility of the machine learning of variant questions, I was puzzled by how much still had to be manually rigged in branching paths of conversation. Perhaps this provides ideal accuracy for enterprises in limited domains, but pronouns can be handled more generally for such short-range contexts using current coreference resolvers. The answering of questions still seemed to be up to the developer. All in all not quite what I would have expected from the title, which could be said for much on the programme. A workshop titled "Start your own AI experiment today" was more of a guided thought experiment into the resources and organisation of an AI project, than actually tinkering with AI technology.

I did not attend the second day of the summit, however there was one centre stage event that I should like to note:

The Afghan all girls robotics group told a touching story about the opposition they met in entering a robot-

ics contest in the US, hailing from a country where few girls are even allowed to study, and being twice barred from the US until international outrage against Muslim bans got them through. With limited resources and one eight of the time that other contestants had they had built a robot for sorting differently coloured balls. While they did not reach top ranks in the contest, they received a silver award for courageous achievement. These girls are now seen as role models by other girls in Afghanistan who want to grow beyond the restrictions of tradition. They came to the World Summit AI in spite of death threats from extremist groups, and got a standing ovation from the audience. This wakening story just goes to show that AI and robotics can not only improve technology, but also human welfare. It makes me think that we too, should have a critical look at still-present discouragement of women in STEM in our own cultures.

www.artistdetective.wordpress.com

AISB 2018 Convention, Liverpool

Floriana Grasso and Louise Dennis

(University of Liverpool, *floriana,louised@liverpool.ac.uk*)

The 2018 Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, AISB 2018 will be held at the University of Liverpool, chaired by Floriana Grasso and Louise Dennis, from April 4th to April 6th 2018. Following a tradition of many years, AISB 2018 is set to provide a unique forum for presenting cutting edge research and burning issues around all areas of AI. A special theme has been chosen for this year, namely "AI for the Digital Society", and contributions are especially welcomed to all symposia which also address the main theme.

After the selection process, a total of thirteen symposia have been accepted, which provide a broad and diverse range of topics, for what promises to be a truly interdisciplinary and exciting event. A concise introduction to the symposia is presented below, but for more information and for details on keynotes, exhibitions and other social events which will complement the programme, as well as the full call for papers of the various symposia, please visit the website at http://aisb2018.csc.liv.ac.uk/

AI, Games & Virtual Reality

Organised by Daniela Romano (UCL Interaction Centre), and David Moffat, (Glasgow Caledonian University).

The AI, Games & VR Symposium, acts as a meeting place for researchers and practitioners from academia, education and industry who are involved with the design, development and evaluation of AI in the context of games or virtual reality and any other form of immersive experience (e.g. 360 videos) It focuses on the application of artificial intelligence or intelligent-like techniques, frameworks and theories to the creation of interactive engaging intelligent games. Papers connecting games to all areas of computational intelligence and traditional AI are considered.

AI, robots and public engagement

Organised by Michael Szollosy (University of Sheffield

A great deal of attention has been paid in recent years to the question of ethics and responsible innovation in researching artificial intelligence and robotics. A key aspect of this has been improved attempts to create a full and frank dialogue between those who build these systems and that nefarious entity known as "the public": the policymakers, politicians, legal experts, and the rest of the "stakeholders". Despite improved efforts, however, most of the public seems to remain largely misinformed as to the development of robotics and AI, causing some wild expectations, unfounded optimism, and fears, both founded and unfounded, that will have dramatic impacts upon research. So while there seems to be a (near) universal acknowledgement that there is a need to engage with the public in our research, it seems that this is something that we can still do much better than we presently are.

This symposium will bring together researchers, those with expertise in public engagement and those with a specific interest in helping guide conversations on robots and AI in terms of ethics, responsible research and the social impacts of new research. This may also include experts public engagement offering workshops in how to organise events, how to talk to the public and specific stakeholder groups, and how best to make meaningful impacts (e.g. on policy).

A public event will accompany the symposium, to be announced.

Artificial Intelligence in Technologies for Health

Organised by Kirsten Smith (University of Southampton), Matt Dennis (University of Portsmouth), and Rosemary J. Thomas (University of Aberdeen) There is a clear potential for Artificial Intelligence to improve healthcare services; for example, clinician decision support, patient behaviour change, analysing patterns in health data, intelligent personalised health support and gamification. Artificial Intelligence in Technologies for Health (AITH) will showcase innovative Artificial Intelligence research that focuses on promoting access, improving efficiency and enhancing quality within healthcare. This workshop aims to connect the current research in AI with the needs of patients, healthcare workers and manufacturers to promote research that is

both practical and innovative. We invite contributions that address theoretical or applied research of AI for health and wellbeing.

Assessing Agency, Moral and Otherwise: Beyond the Machine Question

Organised by Joel Parthemore (Technical University of Eindhoven, NL)

Robotization and other forms of automation increasingly find themselves among the most heard buzz words throughout the manufacturing sector

Bevond the mundane and beyond. assembly-line robots, one hears about self-driving cars, "killer" battlefield robots, sex robots, prototype care robots. Lavpersons and researchers alike talk about the more sophisticated examples in a way that appears to ascribe them agency and, in some cases, stops little short of personhood: describing them as having feelings, weighing choices, making decisions wrong and right. How much is hype and how much substance? To what extent are people speaking metaphorically - and aware of doing so - to what extent do they really mean what they are saying? Are existing artefacts - or, if not, can potential future artefacts be - agents in any sub-Can they be moral stantial sense? agents, capable of making moral decisions and being held responsible for the consequences? Most importantly, how do the answers to these questions shape our ethical interactions with machines that, in some important ways at least, remind us of ourselves? How do they inform our assignments of moral responsibility?

This symposium takes as its starting point that questions of artefactual agency and machine ethics are red herrings. What matters is what qualifies any purported agent as an agent and what qualifies certain agents - whatever their origins - as moral agents.

The 18th Worskhop on Computational Models of Natural Argument

Organised by Floris Bex (University of Utrecht, NL), Floriana Grasso (University of Liverpool, UK), and Nancy

The series of workshops on Computational Models of Natural Argument (CMNA18) is continuing to attract high quality submissions from researchers around the world since its incep-Like the past edition in 2001. tions, CMNA XVIII acts to nurture and provide succor to the ever growing community working on Argument and Computation, a field developed in recent years overlapping Argumentation Theory and Artificial Intelligence. AI has witnessed a prodigious growth in uses of argumentation throughout many of its subdisciplines: agent system negotiation protocols that demonstrate higher levels of sophistication and robustness; argumentation-based models of evidential relations and legal processes that are more expressive; groupwork tools that use argument to structure interaction and debate; computer-based learning tools that exploit monological and dialogical argument structures in designing pedagogic environments; decision support systems that build upon argumentation theoretic models of deliberation to better integrate with human reasoning; and models of knowledge engineering structured around core concepts of argument to simplify knowledge elicitation and representation problems. Furthermore, benefits have not been unilateral for AI, as demonstrated by the increasing presence of AI scholars in classical argumentation theory events and journals. and AI implementations of argument finding application in both research and pedagogic practice within philosophy and argumentation theory.

The workshop focuses on the issue of modelling "natural" argumentation. Naturalness may involve the use of means which are more visual than linguistic to illustrate a point, such as graphics or multimedia. Or to the use of more sophisticated rhetorical devices, interacting at various layers of abstraction. Or the exploitation of "extra-rational" characteristics of the audience, taking into account emotions and affective factors.

The 5th Computational Creativity Symposium

Organised by Maximilian Droog Hayes (Queen Mary, University of London), Mohammad Majid al-Rifaie (Goldsmiths, University of London), and Stephen McGregor (Queen Mary, University of London)

Over the last few decades, computational creativity has attracted an increasing number of researchers from both arts and science backgrounds. Philosophers, cognitive psychologists, computer scientists and artists have all contributed to and enriched the literature.

Many argue a machine is creative if it simulates or replicates human creativity (e.g. evaluation of AI systems via a Turing-style test), while others have conceived of computational creativity as an inherently different discipline, where computer generated (art)work should not be judged on the same terms, i.e. as being necessarily producible by a human artist, or having similar attributes, etc.

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.

Cybernetic Serendipity Reimagined

Organised by Joseph Corneli (University of Edinburgh), Colin Johnson (University of Kent), Anna Jordanous (University of Kent), and Christian Guckelsberger, (Goldsmiths, University of London)

Following the successful AISB Member Workshop VII: Serendipity Symposium held at St Mary's University in June 2017, with keynotes by leading serendiptologists and nine contributed talks, we invite full papers, short papers, and demos for a Symposium on the theme of "Cybernetic Serendipity" – broadly understood.

We take the 50th Anniversary of the famous Cybernetic Serendipity exhibition, curated by Jasia Reichardt at the London's Institute for Contemporary Arts, as our inspiration. The show introduced a broader audience for the first time to emerging work on the intersection of art and technology. "Cybernetic Serendipity" lends itself beautifully (perhaps even serendipitously) to a sensory and emotionally evocative exploration of "AI for the Digital Society", the theme of this year's AISB convention.

Serendipity has been studied extensively to theorise human discovery. However, there has been little consideration of serendipity in AI – it simply wasn't needed, because systems were designed to run in a perfectly predictable manner within highly constrained environments. In recent years, anticipation of AI systems with increasingly unpredictable behaviour leads us to reconsider the role of serendipity may play in a computational context. Serendipity has been addressed in a variety of adjacent fields such as recommender systems, machine ethics, information retrieval, information science, planning and computational creativity. With this symposium, we want to encourage a mutually beneficial exchange between these and other disciplines beyond computing.

Digital Behaviour Intervention for Cyber-Security

Organised by Judith Masthoff, Matthew Collinson, and John Paul Vargheese (University of Aberdeen)

This symposium focuses on how digital technology can motivate and influence people to behave more cyber-securely. It brings together researchers, designers, developers and cyber-security experts interested in computers designed to change cyber-security attitudes and behaviours. The symposium covers a wide range of topics on persuasion, from behaviour intervention methods to persuasive argumentation and persuasive user interfaces. Digital behaviour interventions have a great practical potential. They have been applied in many domains, for instance to improve health and to move towards sustainable living. There has been much progress in the research community on digital behaviour interventions, as shown a.o. by the successful Persuasive conference series, a special issue of the UMUAI journal, and a successful series of workshops on Computational Models of Natural Argument (an area overlapping with persuasion). Recent developments within cyber security have emphasised the need for greater consideration and acknowledgement of human behaviour. The action or inaction of individual users has been demonstrated to have a severely negative impact on the security of organisations across all sectors of the economy. Hence, there has been a lot of interest in cyber-security behaviour, awareness and policy compliance. However, most of this work has not been routed within the behaviour change literature. There is currently an emergence of work that is beginning to combine these two strands of research, and this symposium will help to further build this community.

Emotion Modelling and Detection in Social Media and Online Interaction

Organised by Francesca D'Errico (University of Roma-Tre, Italy), Floriana Grasso (University of Liverpool, UK), Malvina Nissim (University of Groningen, NL), Nicole Novielli (University of Bari, Italy), and Viviana Patti (University of Torino, Italy)

The worldwide diffusion of social media has profoundly changed the way we communicate and access information. Social media is changing the way people interact with each other and share information, personal messages, and opinions about situations, objects and past experiences.

On one hand user-generated content comprise an invaluable wealth of data, ready to be mined for training predictive models. On the other hand, the pervasive use of online social media in computer-mediated communication, is opening new challenges for social sciences and human-computer studies as one of the biggest drawbacks of communication through social media is to appropriately convey and recognize sentiment through text. Furthermore, the sentiment analysis and emotion recognition in online user-generated contents presents its own specificities and challenges due to their characteristics. language use, and to the huge available volume of data.

The aims of this symposium include: presenting the state of the art in emotion modelling and tools for online interaction; fostering discussion around interdisciplinary research area at the intersection between cognitive sciences, computational linguistics, and social computing; enhancing the state of the art in affect recognition in social media: discuss challenges and opportunities of research ethical concerns and applications addressing the role of sentiment and emotions in computer-supported cooperative work and online interaction on social media, with a special focus on education, entertainment, health, egovernment, games, hate speech monitoring.

Philosophy after AI: mind, language and action

Organised by Giusy Gallo, and Claudia Stancati (University of Calabria, Italy) The goal of this one-day symposium is to claim for a philosophical approach to the latest issues about the study of human mind developed in the field of Artificial Intelligence. First, the mind-body problem can be read again involving the neuroscientific research (e.g. mirror neurons, researches on memory and the relationship between brain and action, the reductionist approach), including the provocative theory of the extended mind which enlivened the debate placing the mind-body-scaffolding problem. The second question is about the philosophical category of subject: how to set out the boundaries of the self? The third issue concerns the nature of learning and creativity and the current researches in the field of machine learning. Since learning and creativity have been connected to the acquisition of language and the linguistic change, they have been investigated by linguists and philosophers. The fourth issue is the philosophical theme of teleology, which has been widely debated in cybernetics, crosses all the issues above mentioned. Functionalism, simulation, representationalism, mentalism, identity of explanatory principles are the philosophical milestones which follow the development of cognitive sciences and AI since the last century. Finally, the main issue, in the shadow of H. Simon, concerns what does "artificial" mean.

The 2nd Symposium on Social Interactions in Complex Intelligent Systems

Organised by Federico Bergenti (University of Parma), Stefania Monica (University of Parma), and Paolo Petta (Austrian Research Institute for Artificial Intelligence)

A complex intelligent system 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 complex intelligent systems (SICIS) give rise to collective properties that hold at the macroscopic level, whose emergence cannot be easily inferred from the analysis of the behaviour of single agents at the microscopic 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 concern both local and global phenomena, and therefore they are a key concept to understand the behaviour of a complex intelligent system. The models used to study interactions among agents describe the effects of interactions from a microscopic point of view. Therefore, the derivation of observable behaviours from such interactions may be addressed using various approaches, such as statistical methods, empirical observations, analytic approaches and simulative tools.

Complex intelligent systems are used to describe processes in various fields, such as Artificial Intelligence, Computer Science, Mathematics, Biology, Economics, Physics, Sociology, and Economy. 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. The aim of the symposium is to stimulate discussions and synergies among participants, which are expected to have diverse and complementary research backgrounds.

Swarm Intelligence & Evolutionary Computation

Organised by Mohammad Majid al-Rifaie, Tim Blackwell, and Haya Alhakbani (Goldsmiths, University of London)

Swarm intelligence (SI) and evolutionary computation (EC) techniques have been thriving research topics, specially with the dominating presence of big data in all aspects of technology and and their importance in policy making for institutions, governments and international bodies. Self-organising nature of swarm intelligence in both nature and computational models is key to the attractiveness of such techniques; several such techniques are already proposed, not only explaining and reflecting on the natural-and-social phenomena but also their application to solve complex problems in many fields is an ongoing observation.

Additionally, noisy environments and/or incomplete data are often at the heart of hard real-world data where search and optimisation-related problems are amongst the core issues. Ever since the inception of SI and EC techniques, researchers have been attracted to the complex emergent behaviour, robustness and easy-to-understand architecture of nature-inspired swarm intelligence algorithms; and, particularly in challenging search environments, these algorithms have often proved more useful than conventional approaches.

This symposium would be facilitating the discussion of emerging topics in this context and would encourage early-career researchers, enthusiasts as well as senior academics to engage in a dialogue surrounding the applications and theories based on swarm intelligence and evolutionary computation techniques.

Verification of Systems that Learn

Organised by Louise Dennis (University of Liverpool) and Alice Miller (University of Glasgow)

Machine learning is of particular value in areas where developing a precise specification of desired behaviour is outside the scope of our current understanding of the world. For instance machine learning is widely deployed for image classification tasks. In these cases the specification is that the classifier should match the perception ability of a human. This is a difficult property to formally specify. Even when properties can be formally specified, the results of many machine learning systems (e.g. a set of weights in a neural network) are difficult to map onto these or to reason about in appropriate terms. The aim of this symposium is to bring together researchers interested in the question of how systems that learn may be verified. It will take the form of a number of scientific presentations and posters.

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,

Our company, Big Data Unlimited, sells statistical machine learning solutions to the World's major companies and governments. The General Data Protection Regulation (GDPR) legislation now threatens our business plan. It will require us to provide explanations of those solutions. This is impossible. Our algorithms are inherently black boxes. Can you help our company meet its legal obligations?

Yours, Daddy Data

Dear Daddy Data,

The black box nature of your company's algorithms are both the cause of your GDPR problems and their solution. Just as no one can tell how those ML algorithms arrived at their solution, so no one can tell whether a purported explanation actually explains their inner workings. You need **COP OUT**TM (Causes of Outputs Proclaimed; Offers Undeniable Tales). Given your algorithm's solution, **COP OUT**TM will generate a plausible explanation for it. Whether it is an accurate explanation is both unknowable and irrelevant. Yours, Aloysius

Dear Fr. Hacker,

As you know, my parish consists of half a dozen, small, remote islands in the Inner Hebrides. A few years ago you kindly provided me with a team of telepresence robots: one based in each of my churches. These robots enable me to conduct simultaneous services in each church.

Recently, when a storm prevented me from travelling to one of the islands to conduct a wedding service in person, I decided to do so remotely. All went smoothly until the storm became so severe that communication with the robot was lost. I had assumed that it would then go silent, but it seems that the robot was more intelligent than I had imagined. Its machine learning software had memorised the marriage service, which it continued to conduct it in my absence. So it was the robot, and not myself, that declared the happy couple to be man and wife. Since it did not have the necessary authority to declare them married. I'm concerned that they are not ecclesiastically married. Should I leave them in blissful ignorance of their state of sin? Yours. Switched-on Priest

Dear Switched-on Priest,

Do not be concerned. Before shipment, all our $\mathbf{RITE}^{\text{TM}}$ (Robots: Intelligent and Theologically Educated) ministerial robots are ordained into the Church of God the Programmer. They do, therefore, have the authority to officiate at the declaration of consent. Your couple are ecclesiastically married.

Yours, Aloysius

Dear Fr. Hacker,

I'm concerned that quantum computers will endanger our use of encryption to protect state secrets as well as the financial information used in legitimate online trading. How soon, if at all, will quantum computers become a reality and how can we then restore confidentiality in our communications.

Yours, S Pook

Dear S Pook,

The Institute's code breaking quantum computer has been very successful for several years. Indeed, our decryption of your own communications reveal that you already have code-breaking quantum computers yourselves, so your question is a little disingenuous. Nor do you need me to tell you how to use quantum cryptography to protect your secrets. I think you are phishing.

What you may not have yet invented is quantum ontologies. These ontologies exist in a state of superposition, so that they can simultaneously represent opposing World views. We find these very useful for presenting different 'facts' to rival groups. For a very reasonable consideration, we can supply **QUACK**TM (Quantum Usance of Alternative and Contrasting Knowledge), a tool you will find very useful in your diverse propaganda work. 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.

$\mathbf{Editor}-\text{aisbq@aisb.org.uk}$

Dr. Mohammad Majid al-Rifaie (Goldsmiths, University of London)

Advertising and Administration

Rob Wortham (AISB Executive Office) Dept of Computer Science 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. 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 (University of South Wales); Vice Chair: Prof. John Barnden (University of Birmingham); Secretary: Andrew Martin (Goldsmiths, University of London); Treasurer: Rob Wortham (University of Bath); Editor: Dr. Mohammad Majid al Rifaie (Goldsmiths, University of London); Publications: Dr. Floriana Grasso (University of Liverpool); Publicity & Media Officer: Dr. Colin Johnson (University of Kent); Workshop Officer: Dr. Yasemin J Erden (St. Mary's University College); Public Understanding Officer & Schools Liaison: Janet Gibbs (Kings College London); Equality, Diversity and Inclusivity: Dr Anna Jordanous (University of Kent) Loebner-Prize Officers: Janet Gibbs (Kings College London), Dr. Nir Oren (University of Aberdeen); Dr. Ed Keedwell (University of Exeter).

Contents

Editorial	3
The heart inside the machine J. Ollerenshaw	4
Conference Report: International Joint Conference on AI (IJCAI) K. Sokol	9
Report: World Summit AI P. Dhondt	11
AISB 2018 Convention F. Grasso and L. Dennis	13
Dear Aloysius	21

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

ISSN 0268-4179 © the contributors, 2017