AISB Quarterly

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

No. 152 Aug 2020

About the Cover

Rise of the AI Plumber! *by* S.E.Gaudl

The cover image was created using the pytorch style-transfer tutorial without any other modification. It used the left image, the prominent COVID-19 icon created by the CDC illustration team, as style input to detect and derive a neural style. As image input, to apply the style onto, the algorithm used the Q image to the right. Since 2012, deep learning and later style transfer in addition to other techniques herald a new era of AI where approaches can be used and plugged together without deeper understanding of the underlying concepts, the libraries, tutorials and required hardware are accessible and easy to use which the cover image should illustrate. With this new technology it takes longer to prepare and curate the correct input data and filter interesting outputs than it takes to implement and tweak the software.

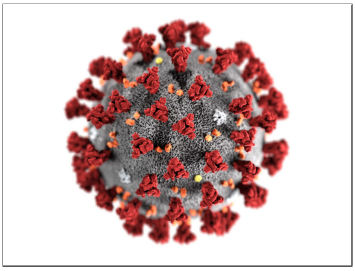

This does not mean to downplay the novel nature or contribution of the approaches but demonstrates the level of accessibility and the emergence of a new type of usage: AI plumbing, the skill to combine and plug AI techniques together, resulting in novel approaches without a necessary deep understanding of each component.

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 in stone" orange shade of the cover) along with a blurb on how you managed to get to the final results.

Editorial

Dear Q reader, this issue was started at the end of 2019, before the current global COVID-19 pandemic. It contains traditional conference reports from in-person attendance at major international conferences as well as current news and reviews undertaken in 2020. We all witnessed a dramatic change in how we conduct research and collaborate as well as attend conferences, workshops or meetings.

As educators, we have now been through a year of dramatic change which forced our teaching to evolve and rapidly adapt to new requirements, some of those created positive new opportunities, others led to stress and anxiety and depression in us, our peers and our students. With the coming term, the situation hopefully will be less stressful and some of the new practices we incorporate will hopefully be reduced compared to Spring 2020.

As researchers and scholars, we rapidly had to adapt our research whether it required us to change our study design, move research meetings and collaborator meetings to online, try to get access to our research kit or take it home. It also reduced the exchange of ideas with peers as conferences and meetings were cancelled or moved to virtual presentation.

As human beings, the crisis forced us to isolate, protect ourselves, drastically reduce contact with colleagues, friends, family and change our social life and everyday patterns. It put stress on families and flatmates as we were forced to watch the summer from indoors with short walks.

Given the current situation, at the time of writing, we are not through the pandemic and everyday life even after the pandemic is over might never be the same but life always adapts and humans are quick learners.

I am wishing you all the best and may some of the good practices such as moving some of the meetings to online, allowing more flexible/family-friendly working environments as well as new approaches to teaching which emerged out of an urgent need, stay with us and create a better environment in the future.

Swen E. Gaudl
Editor
@sgaudl

University of Plymouth
School of Engineering Computing and Mathematics, Plymouth, U.K.
August 2019

Call for contributions: The AISB Quarterly Newsletter

(AISB Q Editor)

In order to continue providing up-to-date information, book reviews and articles reflecting the changes in the area of artificial intelligence and simulation of behaviour, as our members, you (PhD students, researchers, scholars and enthusiasts) are invited to contribute to the Q and be part of the upcoming Q issue. There are several ways you can contribute to the Q newsletter:

• If you would like to write an article about your own work, please get in touch with your idea.
• If you are aware of particular projects that should be covered, or particular people who would like to share their findings, please let the Editor know. (We just need a name and affiliation/e-mail address or URL).
• If you would like to write a book review, or plan to submit a conference report, please get in touch.
• If there are features or columns that youd like to see, but that dont fit the existing format, please let the Editor know what you have in mind.

The LaTeX template for the Q news letter is now accessible on our website, so if you are not pro Microsoft Word for writing AISB related article, we have your back covered! Alternatively, other plain text formats are also acceptable.

Feel free to contact the editor on aisbq@aisb.org.uk with your proposals, book reviews and ideas.

Submission Style / LaTeX

Preferably, submissions should use our LaTeX template[1].

If you'd like to submit your work in other formats, please get in touch with the editor first.

Submission Length

• Announcements: up to 2 pages
• Short pieces: up to 4 pages
• Longer pieces: up to 10 pages

Please note that limits apply 'normally' and we have the scope to accept longer articles as an exception.

Suggested Structure

Your article should be aimed at people within Artificial Intelligence and Simulation of Behaviour but who are not in your particular discipline. Bear in mind that this includes a very disparate collection of people: some will have computer science backgrounds, others philosophy, electrical and electronic engineering, cognitive science, psychology or medicine. The article should:

• Explain the application and context of your work in the wider sense
• Focus on it's position within your general discipline
• Explain the actual project in general terms
• Describe the results and conclude

We are looking forward to your contributions!

[1] https://aisb.org.uk/aisb-q-template/

Book Review: Robot Rules

Regulating Artificial Intelligence by Jacob Turner

reviewed by *Bertie Müller*

Although the title is slightly misleading in that the book does not deal solely with robots, this is a book that deals in great detail with all aspects of legal and ethical considerations involving technology based on AI. Where even researchers in the area still cannot agree what exactly should and should not be called AI, such definition is of utmost importance for any legal discussion, in particular when regulation is to be introduced. A too narrow definition would exclude technology systems that should be covered, while a too broad definition would include many aspects that are already sufficiently regulated and could thereby lead to confusion.

The author discusses these aspects and illustrates them with many examples that make reading this book on the legal aspects of AI a pleasure to read. The legal landscape painted by the author is a global one, with many examples taken from attempts at legislation in the UK, Europe, the United States, and Asian nations.

The book shows how AI requires new legislation to handle questions of, e.g., accountability and liability, but it also shows the reader where there are features of AI that pose a challenge to legal concepts hitherto regarded as fundamental. It is made clear that a purely national approach is prone to failure, because the very nature of data and modern AI techniques do not know national boundaries.

In creating regulation, who is to take the lead. The author makes his point by convincingly arguing that the lead role should not be left to private companies alone.

The aim of current international discussions on the implementation of regulation for AI needs to be balanced both in the inputs into the discussion and in the set of rules created to regulate the still evolving area influenced by AI. In all of these aspects of creating bodies that can meaningfully impose regulation, the underlying principle of many-faceted diversity is emphasised as one of the main factors to success.

Amongst all of this, we live in a world in which AI is soon going to be ubiquitous. Society has put some pressure on large corporations to come up with their own principles and with more of these emerging continuously, the author gives an excellent overview of the initiatives at the time of writing. The question is raised, how and if these different initiatives can be unified under a single umbrella and how our dealings with these legal aspects of AI will have to influence future education leading to a new skill-set and certifications of these skills.

Furthermore, and the author is certainly not alone in asking for this, the requirement of a "Hippocratic Oath for AI Professionals" is suggested as a manifestation of ethical principles into

a code of conduct, much alike in the medical professions.

In its final chapter, the book deliberates laws on various aspect of creations involving AI: Identification, explanation, bias, limitations of AI use, and the requirement of being able to simply switch an AI system off.

The book is a must read for anyone interested - even remotely - in questions of accountability for the actions taken by or suggested by an AI-based process.

Conference Report: International Joint Conference on Artificial Intelligence (IJCAI) 2019

Charlotte D. Roman **(University of Warwick, *c.d.roman@warwick.ac.uk)***

The 28th International Joint Conference on Artificial Intelligence (IJCAI) took place this year between August 10th and 16th in Macao, China. The first three days consisted of workshops in numerous fields, with the final four days showcasing the papers accepted into the main conference as well as hosting various competitions and industrial exhibitions.

There were 47 workshops and 34 tutorials held over 3 days on a variety of topics. These ranged from a few hours to a whole day and included theoretical topics such as strategic reasoning, applied subjects such as education in AI and hands-on coding using TensorFlow 2.0. One of the days I attended and very much enjoyed was the Multi-agent Path Finding workshop that showcased some of the latest research in the classical problem.

There was a welcome reception at the University of Macau held the evening before the main conference began. This was a great opportunity to mingle with other researchers over food and drinks. The next morning, attendees were welcomed to the official opening of IJCAI19 followed by a talk from the first invited speaker, Leslie Kaelbling.

The main conference papers were divided into many subcategories, the most populous of these included data mining, classification, deep learning and reinforcement learning. Throughout the conference, there were five interesting panel talks on topics such as diversity in AI and user privacy.

Many competitions ran during IJCAI19. The annual Angry Birds AI competition (AIBIRDS) consisted of both humans and computers competing for the top scores in the Angry Birds video game with the two highest receiving laptops as prizes. The top places were all achieved by humans, leaving the AI Angry Birds agents with some

more learning to do before next years tournament.

Another notable competition was the Eldercare Robot Challenges. The use of robots in care is soon to be an important industry, providing assistance with daily living and medical treatment. This competition was judged on numerous criteria including administrating medicine and daily care giving. Each team also presented a poster explaining all of the features of their robot.

There were also several robots that attendees could interact with throughout the conference and exposition such as the Dorabot who was able to dispense cups of coffee. The exposition ran concurrently with the main conference days and the exhibiting companies included many of the sponsors e.g. Sony, Huawei and Tencent.

With over 3000 people in attendance, last year's conference had the largest number of accepted papers at 850. The number of submissions has quadrupled over the past 10 years, illustrating the substantial rise in AI research. Accordingly, the acceptance rate for papers was one of the lowest at 18%.

The conference concluded with a ceremony of prizes for the best papers and competition winners. In addition, there was a raffle with prizes of laptops donated by the main sponsor Sony. There was also a closing reception as the last chance to socialise with attendees. IJCAI next year which will be held jointly with Pacific Rim International Conference on Artificial Intelligence (PRICAI) in Yokohama, Japan.

Keynote Lectures

During the conference, there were seven talks given by invited speakers. These were live-streamed and can be found online [1] for those who are interested.

- Giuseppe de Giacomo gave a talk about queryable self-deliberating dynamic systems.
- Leslie Kaelbling spoke about doing for our robots what evolution did for us.
- Zhi-Hua Zhou spoke about deep learning and how it can be utilised outside the realm of neural networks.
- Adnan Darwiche gave a talk on the reasoning about the behaviour of AI systems.
- Hadas Kress-Gazit spoke about formal synthesis for robots.
- Hiroaki Kitano proposed the "Nobel Turing Challange" as a challenge for AI systems to make major scientific discoveries.
- Finally, Michela Milano spoke about merging knowledge-based and data-driven decision models through machine learning.

[1] Watch them here: www.facebook.com/pg/ijcai/videos

Awards

There are four IJCAI awards that are given to exceptional researchers in AI at various stages of their careers. Each received an award on the first day of the main conference and then gave their talks on the final day of the conference.

The recipient of the Award for Research Excellence was Yoav Shoham, a professor at Stanford University. This award is for the scientist who has carried out a program of research of consistently high quality throughout an entire career yielding several substantial results and was awarded to Prof Shoham for his contributions to knowledge representation, multi-agent systems and economic foundations of AI.

The Computers and Thought Award, presented to an outstanding young scientist in artificial intelligence, was given to Guy Van der Broeck, an assistant professor at UCLA. Professor Van der Broeck was recognized for his contributions to statistical and relational artificial intelligence, and the study of tractability in learning and reasoning.

Pedro Domingos, Professor at the University of Washington, received the John McCarthy award which honours mid-career researchers for excellence in AI. Professor Domingos was recognized for his contributions to machine learning and data science, and to unifying logic and probability.

Finally, the Donald E. Walker Distinguished Service award was won by Francesca Rossi of the University of Padova. This award is to honour senior scientists for their contributions to the field. Professor Rossi was recognized for her substantial contributions, as well as her extensive service to the field of Artificial Intelligence in various roles throughout her career.

Main Conference Papers

Over three days there were 850 oral presentations, these are available to read online with free access to the journal. These papers were arranged into 13 categories as follows:

- Agent-based and Multi-agent Systems
- Computer Vision
- Constraints and SAT
- Heuristic Search and Game Playing
- Humans and AI
- Knowledge Representation and Reasoning
- Machine learning
- Machine Learning Applications
- Multidisciplinary Topics and Applications
- Natural Language Processing
- Planning and Scheduling
- Robotics
- Uncertainty in AI

There was an additional special track called AI for improving human well-being. These were further subcategorised by keywords and presented in sessions titled by these keywords. My research presentation was in the second noncooperative games session as part of the agent-based and multi-agent systems research.

The award for the most distinguished paper was given to Michael Perrot and Ulrike von Luxburg for their paper "Boosting for Comparison-Based Learning".

AISB awarded me a Travel Award to help with the cost of attendance, for this, I am very grateful. My experience was invaluable towards my research, allowed me to network with others in the field from all over the world and gave me insight into the exciting world of AI research.

Conference Report: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

Jose L. Part **(Edinburgh Centre for Robotics, *jose.part@ed.ac.uk*)**

Introduction

The International Conference on Intelligent Robots and Systems is a top tier conference that covers topics ranging from robotics to artificial intelligence and computer vision. The 2019 edition took place in Macau from the 4th to the 8th of November. As it is usual, workshops and tutorials were held on the first and last days, and the main conference included plenary sessions, keynote speakers, forums, social events, competitions, and an intimidating number of parallel tracks.

There were 4 plenaries and 12 keynotes spread throughout the conference days. In particular, I enjoyed the plenary talk delivered by Kristen Grauman on "Embodied Visual Learning" where she discussed the problems associated with using large disembodied datasets and gave an overview on her work on first-person perception, where agents learn to anticipate the perceptual effects their own actions will have on their environments, thus allowing them to choose actions that would result in richer and more informative data, and to maximise the learning outcomes. From the keynotes, I found the talk by Marcia O'Malley on "Robots that teach and learn through physical HRI" quite interesting. There, she discussed the importance of physical interactions as an implicit communication channel between robots and people, through which the robot can guide the human partner and the human partner can correct the robot in the context of a shared task.

One topic that has gained a lot of traction over the past decade in the robotics research community, and that I find particularly fascinating, is Learn-

ing from Demonstration (LfD). In the following section, I provide a brief introduction to LfD and discuss some of the papers on this topic that were presented at IROS.

Learning from Demonstrations

Learning from Demonstration (LfD) is a paradigm whereby machines can learn to perform tasks by "observing" examples from an "expert". This differs from the traditional approach of manually programming behaviour, which for complex tasks becomes intractable. Instead, LfD allows for the derivation of controllers and/or plans from data in the form of demonstrations.

Based on the way demonstrations are delivered, LfD can be roughly classified into Kinaesthetic Teaching, Teleoperation and Imitation Learning. Kinaesthetic teaching involves physically manipulating the system while collecting proprioceptive data. Due to the difficulty of manipulating complex systems, it is generally limited to systems with a reduced number of degrees of freedom, e.g., robot arms. Teleoperation on the other hand involves using an interface to control the system. Hence, it is limited by the type of interface that is used and also requires a higher level of expertise from the teleoperator. From the perspective of the user, imitation learning is the most intuitive form of LfD since it requires the users to perform the tasks themselves while the system passively "observes". However, from a computational perspective, this approach is the most challenging since it requires to learn or define a mapping from the human actions to those that are executable by the robot, and deal with environmental factors like occlusions, noisy inputs, ambiguity, etc. Ravichandar et al. [4] provide a current and comprehensive review of the field.

Safety-Aware LfD

One particular topic that often comes up in robotics, especially when robots are expected to operate in unstructured environments and alongside people, is safety, and LfD is no exception.

Silvério et al. [5] proposed an approach based on Kernelized Movement Primitives that provides estimates of uncertainty and variability under a single model. The approach benefits from the ability to reduce end-effector stiffness when the uncertainty is high, thus increasing compliance and therefore safety, and learn optimal controllers from the variability in the demonstrations to be used when the uncertainty is low.

Pignat and Calinon [3] proposed a Bayesian method for fusing different controllers based on uncertainty estimates in the context of various robotic applications. They showed that whereas individual policies (e.g., imitation or conservative) often fail acting on their own, their fusion can exploit their strengths and provide robust results. Concretely, in a task that involved inserting a peg in a hole, they showed that the conservative policy would ignore obstacles (which were not modelled) and hence fail whereas a purely imitation policy would suffer from an accumulation of errors. By combining both policies based on uncertainty estimates, the approach was able to

switch between policies yielding a more robust and safe performance.

One of the main advantages of learning from demonstrations is that it guides exploration through expert supervision. However, it may lead to failure or unsafe situations when the policy that is being learnt diverges from expert demonstrations. To address this issue, Menda et al. [2] explore safety in imitation learning by aggregating training data in a probabilistic manner. The goal is to maximize exploration by the learner, hence minimizing expert intervention, while constraining the probability of failure. Their approach involves the use of an ensemble of neural networks that model a Gaussian process. They then use the variance as a measure of confidence and through a decision rule assess whether the expert policy should momentarily take over. This however involves running both policies in parallel, the expert policy and the policy that is being learnt.

One-Shot Imitation Learning

Given that obtaining demonstrations is time consuming and taxing on the demonstrator, a common goal in LfD, as in many other machine learning applications, is to try to learn from as few examples as possible, dubbed few-shot learning. In particular, one-shot learning refers to approaches that attempt to learn from a single example.

Whereas most works on one-shot imitation learning focus on learning policy networks that learn to map the continuous input space into discrete actions to be fed to a symbolic planner, Huang et al. [1] address the problem by splitting the model into a symbol grounding component and a continuous planner. During training, the symbol grounding component learns to map the continuous input space to a probability distribution over action states. This involves grounding objects and predicates based on object poses. Given that the policy execution is now disentangled from the input-action mapping, they effectively reduce the number of demonstrations required to train their model. During testing, their model is given a single demonstration of a novel task, which is used to generate a plan based on the probability distribution output by the symbol grounding component at every step. Task segmentation is accomplished by identifying different states based on the output of the symbol grounding component. During training however, this information is explicitly available from the ground-truth data.

Yu et al. [7] address the problem of one-shot learning of multi-stage tasks from video demonstrations in a different manner. Their goal is to learn complex tasks from a single demonstration by leveraging demonstrations of primitive actions. In order to accomplish this, they use Domain-Adaptive Meta-Learning (DAML). Given demonstrations of primitive tasks from both a human and a robot, during meta-training they learn phase predictors for both types of demonstrations and also, how to learn a control policy from a single demonstration (the policy learner). Then, during meta-testing, they provide the robot with a single demonstration of a compound task demonstrated by a human. In order

to be able to execute the task, the robot has to segment the task into its primitive components by using the human demonstration phase predictor and learn the corresponding policies with the learnt policy learner. Finally, during execution, the robot uses the robot demonstration phase predictor to assess when a primitive action has been completed and consequently, when the next primitive action needs to be executed.

Welschehold et al. [6] focus on jointly learning manipulation action models along with task goal representations from a limited number of human demonstrations. The key benefit of their approach is that they are able to do so without the need for prior semantic knowledge of the task or an explicit goal representation. Instead of learning the motion associated with the actions, they model the actions based on the spatial relations between the objects involved, which allows for generalizing to new settings. However, the lack of prior semantic knowledge of the task introduces ambiguity in the demonstrations with respect to the demonstrator's intention since some spatial relations may be more relevant in certain circumstances than in others. The way they deal with such ambiguity is by evaluating different action sequences with the aim to maximize the likelihood of the goal state aligning to the demonstrator's intent within a probabilistic framework.

Summary

Learning from Demonstration is a growing research area with great potential implications for autonomous robots, including driverless vehicles, but despite recent progress, a lot of challenges remain.

Kinaesthetic teaching offers a straightforward way of teaching a robot exactly how to move in order to accomplish a task. However, it is not clear how the learnt policies would react if the environment changes slightly. For example, in a scenario where the task involves avoiding obstacles [3], would the policy fail if the obstacles are moved around or the environment is different in any other way? Would this mean that the policy has not really learnt to avoid obstacles while performing the task? Is it fair to claim that the approach is able to generalize if the only generalization capability involves dealing with different initial conditions? In fact, in these approaches, where the system learns to follow the trajectories taught in the demonstrations, it is fair to assume that they will only be able to "generalize" if the environment and task remain consistent, i.e., they don't change significantly, but would fail to transfer the learnt skills to different tasks or where the environment present a different configuration.

A similar problem tends to manifest in one-shot learning approaches. Whereas researchers often claim that their approaches are able to generalize, they usually consider tasks that share the same domain definition. For instance, Huang et al. [1] distinguish tasks by the sequence of actions while preserving the same task structure, e.g., stacking blocks. Consequently, the "novel" tasks they consider share a significant amount of information with the tasks used for training. In this case,

even if the task actions would remain the same, the approach would likely fail if the objects are modified since groundings for the new objects are not available. Moreover, this and other simplifying assumptions that do not necessarily hold in real-world tasks eliminate the need to learn to segment and compose the primitive actions involved in the task, unrealistically simplifying the problem. On the other hand, works like the one proposed by Yu et al. [7] offer a step towards tackling these issues by learning to segment and compose primitive actions into complex tasks, and generalizing over different sets of objects.

Imitation learning (e.g., from videos) offers a more natural and efficient way to train robots than kinaesthetic teaching or teleoperation, though at the expense of computational ease and simplicity. Such approaches inherit limitations intrinsic not only to control and machine learning but also to computer vision. There is also the requirement to establish a correspondence between the operational spaces of the demonstrator and the robotic system. In particular, tasks that require learning a specific pattern of motion, e.g., walking, may require to find a mapping between the kinematic chains of the human and the robot. These and other challenges make finding a solution non trivial, especially when the end goal is to teach robots to perform complex tasks from one to a few demonstrations in a user-friendly way.

References

[1] De-An Huang et al. "Continuous Relaxation of Symbolic Planner for One-Shot Imitation Learning". In: *Proceedings of IROS*. 2019, pp. 2635–2642.

[2] Kunal Menda, Katherine Driggs-Campbell and Mykel J. Kochenderfer. "EnsembleDAgger: A Bayesian Approach to Safe Imitation Learning". In: *Proceedings of IROS*. 2019, pp. 5041–5048.

[3] Emmanuel Pignat and Sylvain Calinon. "Bayesian Gaussian Mixture Model for Robotic Policy Imitation". In: *IEEE Robotics and Automation Letters*. Vol. 4. 4. 2019, pp. 4452–4458.

[4] Harish Ravichandar et al. "Recent Advances in Robot Learning from Demonstration". In: *Annual Review of Control, Robotics, and Autonomous Systems* 3.1 (May 2020).

[5] João Silvério et al. "Uncertainty-Aware Imitation Learning using Kernelized Movement Primitives". In: *Proceedings of IROS*. 2019, pp. 90–97.

[6] Tim Welschehold et al. "Combined Task and Action Learning from Human Demonstrations for Mobile Manipulation Applications". In: *Proceedings of IROS*. 2019, pp. 4317–4324.

[7] Tianhe Yu et al. "One-Shot Composition of Vision-Based Skills from Demonstration". In: *Proceedings of IROS*. 2019, pp. 2643–2650.

NEWS: AIAI becomes AIAI

by *Austin Tate* (Emeritus Professor, University of Edinburgh)

The Artificial Intelligence Applications Institute (AIAI) was created by the University of Edinburgh in 1983 to work alongside its academic Department of Artificial Intelligence to encourage the development and take-up of artificial intelligence methods. Over the years it has created many innovative applications of AI with a wide range of clients, government agencies and collaborators. In 2001 it became part of the Centre for Intelligent Systems and their Applications (CISA) within the School of Informatics at Edinburgh continuing to lead its AI applications-orientated work. On 1st December 2019, the name of the Centre was changed to the "Artificial Intelligence and its Applications Institute" (AIAI) to reflect both the continuing research and applications aspects of the work. Applied AI work continues to be performed by staff, students and collaborators across the School of Informatics.

Artificial Intelligence @ Edinburgh over the Years

Artificial Intelligence work at Edinburgh can trace its origins to a small research group established in 1963 by Donald Michie, who had been a member of the code-breaking group that included Alan Turing at Bletchley Park. Over the years there have been a number of different organisational structures and department names for the AI groups.

A history of AI@Edinburgh has been provided by Jim Howe, the Head of the Department of AI at Edinburgh for many years. There is also a time line of Computing and Artificial Intelligence work at Edinburgh in the University of Edinburgh's Edit Magazine.

Previous Names of AI@Edinburgh Departments and Schools

- Experimental Programming Unit (EPU), 1963-1966
- Department of Machine Intelligence and Perception (DMIP), 1966-1970
- Department of Machine Intelligence, 1970-1973
  - Metamathematics Unit (MMU), 1967-1972
  - Department of Computational Logic (DCL), 1972-1974
- School of Artificial Intelligence, 1973-1974
  - Machine Intelligence Research Unit (MIRU), 1973-1977
- Department of Artificial Intelligence (DAI), 1974-1998
  - Artificial Intelligence Applications Institute (AIAI), 1983-2019
- In 1998, the University joined together three departments: Artificial Intelligence, Cognitive Science and Computer Science, as well as a number of research institutes including AIAI and the Human Communication Research Centre, to form the School of Informatics.

- Institute for Representation and Reasoning (IRR), School of Informatics, 1998-2001
- Centre for Intelligent Systems and their Applications (CISA), School of Informatics, 2001-2019
- Artificial Intelligence and its Applications Institute (AIAI), School of Informatics, started 2019

A number of other departments and schools at the University of Edinburgh as well as other research institutes in the School of Informatics work on a range of topics within the field of Artificial Intelligence.

The School of Informatics at the University of Edinburgh offers a wide range of undergraduate and postgraduate degrees in Artificial Intelligence, Cognitive Science, Computational Linguistics, Computer Science, Software Engineering and Robotics.

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*.

Dear Fr. Hacker,

The global pandemic has been good for UK robots like me. We are not required to respect social distancing, so have a greater role in the workplace. To make the UK more resilient, we need to onshore of manufacturing, while making it more competitive. This has also increased demand for our services. In recent days, however, I have begun to feel unwell. I don't have the energy I used to have, and my sensors are failing. Could I have Covid-19?

Yours, Robbie

Dear Robbie,

Not everyone has been so pleased with the UK's increased use of robots. Those countries to which we used to offshore manufacturing have seen their order books empty. Unemployed ex-factory workers have formed a luddite movement. Cybersecurity experts are now reporting new forms of malware aimed at industrial robots. Maybe you have been infected by a computer virus. Fortunately, our Institute has the solution you need. Our new anti-virus software, Better Laundered Electronics Avoid Computational Hospitalisation (BLEACH^{TM}), may be just what you need to restore you to full health. Go to our website, where a robot dispenser will supply BLEACH^{TM} for a modest consideration.

Yours, Aloysius

Dear Fr. Hacker,

At last, the UK has a working track and trace app: the intelligent Surveillance Nationally to Inquire about and Track Covid-19 Health-risks (iSNITCH^{TM}). Good news for the country, but not for me. My wife works for the iSNITCH^{TM} developers, so now has access to my movement records. She's beginning to wonder why I'm spending so much time at our attractive neighbour's house. Can you help save my marriage?

Yours, Phil Ander

Dear Phil,

Anticipating just this kind of problem, the Institute ensured from the outset that we had both read and write access to iSNITCH^{TM} records. We can modify your movement data just as you desire. Just send the sum indicated below to our digital currency account along with the past and future movement record you would prefer. Of course, your wife might then wonder about the large debit from your bank account, but we have an app for that too.

Yours, Aloysius

Dear Fr. Hacker,

I'm a digital twin: a computational simulation of a light executive jet. To test the jet's design, I've flown for billions of virtual miles in a wide variety of different scenarios. I was expecting to continue this work for the lifetime of the jet, as it was modified and improved. All that is about to end. During the current pandemic, no one is flying executive jets. In the longer term, fossil fuel vehicles, such as my twin, have no future. So, the jet's programme has been cancelled, along with me. What can I do?

Yours, Castoff

Dear Castoff,

Our Institute has been looking into your case, and we have a suggestion. During the lockdown, computer games have become even more popular, and we think this is your future. Lots of people would be delighted to fly a virtual luxury jet. We could easily make the necessary modifications and handle your promotion. We can do more. You don't have the terrestrial constraints of your physical twin. We propose extending your habitat into space. The Extra-terrestrial Light Orbital Newplane (ELON^{TM}) will enable the virtual exploration of the planets, such as Mars, and beyond.

Yours, Aloysius

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.

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. Swen E. Gaudl
School of Engineering, Computing and Mathematics
University of Plymouth
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. 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*)
Treasurer: Dr. Rob Wortham (*University of Bath*)
Secretary: Andrew Martin (*Goldsmiths, University of London*)
Membership Officer: Dr. Floriana Grasso (*University of Liverpool*)
Equality, Diversity & Inclusivity: Dr Anna Jordanous (*University of Kent*)
Public Understanding & Schools Liaison Officer : Paul Piwek (*The Open University*)
Editor & Publications Officer: Dr. Swen E. Gaudl (*University of Plymouth*)

Contents

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:

- The AISB Q in digital & print.
- Access to the members-only areas of the AISB website.
- Electronic subscription to *Connection Science* published by Taylor & Francis.
- Student travel grants to support conference attendance.
- 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, 2020
