

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

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323  
324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
84

## About the Cover

Warning, no complex computational algorithm beside human creativity was used to create the image. The image was created from an image taken during the AISB Annual Convention Dinner 2019 at the Eden Project in Cornwall. It shows an overview of the former clay pit which has been turned into a marvellous garden with the three biomes dominating the landscape. The font of the Q is Bauhaus93 to support the idea of integrating form and function which fits in well with what the Eden project achieved in the former pit. The image is constructed out of multiple different layers of text which described the plan for the annual convention for 2019 including a proposal of each intended symposium. Each layer of text is represented through a distinct tonal values derived from the original image. By integrating the layers each letter is constructed out of different colour tonal values mapped into grayscale.

## 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](mailto: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

Dear Q reader, before I introduce myself, I would like to thank Mohammad Majid al-Rifaie for his work as editor of the Q. Mohammad did a great job and without him, we might not have had 150 issues to this day. I hope that with the coming issues, I can continue and contribute to the legacy of the Q.

Who am I and why am I addressing you? I would like to introduce myself as the new editor of the Q; I am Swen E. Gaudl, a researcher in Artificial & Natural Intelligence and a Lecturer of Computer Science at the University of Plymouth. I have been a member of the AISB since around 2012 when I moved to the UK and started a PhD on AI and agent architectures at the University of Bath. At that time, I got introduced to the AISB through my supervisor who mentioned the Q and the AISB travel support for students, which is a marvelous opportunity for students specifically under the current conditions.

The AISB society itself is a unique and diverse group of people brought together under the umbrella of AI and the Simulation of Behaviour. The sheer variety of different view points, interests and topics is not only visible once a year at the annual convention but also very visible in the history of articles in the Q, which were always interesting to read as they are not only centred around computational optimisation and problem solving but also philosophy, psychology, law, art and many more.

After running the AISB annual convention in April I decided to step up

and involve myself more with the AISB under the role of Q editor because the Q is such an amazing platform for communicating ideas to like-minded researchers, scholars and practitioners. It provides a platform for all those diverse views and offers the writer to present to and discuss their ideas with a truly interdisciplinary audience.

For the coming year, my main aim as editor is to increase the frequency of the Q and to provide AISB members a platform for communicating and discussing ideas that is interesting and thought provoking. This also includes providing students and practitioners a platform to discuss their work and get feedback on ideas or submit reviews for current books and technology. In the same spirit as the annual convention, contributions from all disciplines related to artificial intelligence and the study of the simulation of behaviour are very welcome.

This issue for the most part focuses on the AISB annual convention which was held in April 2019 in Falmouth as well as pieces to engage in the AISB community.

**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, lecturers, professors 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 you'd like to see, but that don't 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](mailto:aisbq@aisb.org.uk) with your proposals, book reviews and ideas.

## Submission Style / LaTeX

Preferably, submissions should use our LaTeX template: [http://aisb.org.uk/publications/aisbq/Q\\_template.zip](http://aisb.org.uk/publications/aisbq/Q_template.zip)

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

## 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, some electrical and electronic engineering, others 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!**

---

# CALL FOR PROPOSALS: AISB Workshop Series

**Yasemin J. Erden** (AISB Vice Chair, [yj.erden@stmarys.ac.uk](mailto:yj.erden@stmarys.ac.uk) )

Since September 2012, the AISB has hosted a series of one or two day workshops across the country. A number of publications have arisen from these events, including a Symposium Issue of the Journal of Consciousness Studies.

If you are interested in hosting one of these events, you will find information on what you will need to do on this page: <http://www.aisb.org.uk/events/members-workshop-series>. You can also find details about previous events on those pages.

These events are free to attend for all AISB members, with funding of up to £500 to cover expenses related to organising the workshop offered by the AISB. More details on application, including what funding can be used for. Current non-members would be able to host or attend any of these workshops for the cost of AISB membership (which start at £20 per year for students and £48 for ordinary UK members per year).

Details for hosting the workshops can be found here: <http://aisb.org.uk/media/files/AISBWorkshops.pdf>

In order to propose a workshop, you will need to complete a brief application with the following details:

- Workshop title
- Workshop abstract (200–400 words approx.)
- Organiser(s) and main contact (include details of expertise in proposed topic)
- Host Institution details (name, address)
- Planned dates for event (please also include proposed deadlines for the following):
  - Abstract Submission
  - Notification of Decision
  - Registration
  - Workshop
- Possible speakers (e.g. do you plan to invite speakers?)
- Where you would advertise (e.g. could you create a page on your institution website?)

For more information, or to submit an application, contact Dr Yasemin J. Erden (AISB Vice Chair):  
[yj.erden@stmarys.ac.uk](mailto:yj.erden@stmarys.ac.uk)

---

# Concluding the AISB 2019 Convention, Falmouth

**Swen E. Gaudl** (University of Plymouth, [swen.gaudl@gmail.com](mailto:swen.gaudl@gmail.com))

## Convention Background

In spring 2018, Falmouth applied to host the 2019 convention to bring researchers, practitioners and artists from all over Europe to Falmouth. After successfully winning the bid in June 2018, the organisational team consisting of Swen Gaudl, Tanya Krzywinska, Rob Saunders and Edward Powley started planning the convention for April 219.

As the AISB annual convention is longest, continuously running meeting on Artificial Intelligence worldwide, it provides a unique forum for presenting cutting-edge research and burning issues around all areas of AI.

Each year the convention has a unique theme to attract attendees and make use of the unique location the event is being held. The theme for the 2019 convention was "Artificial Intelligence, Imagination and Invention" which tied in the theme of the GRO-

ERA project that funded the research cluster, the Metamakers, that organised the convention.

The convention was then held in April 2019 by the Metamakers and Falmouth University to bring in members of the Society of the Study of Artificial Intelligence and the Simulation of Behaviour (AISB) as well as invited keynote speakers and organisers that are neither part of AISB or Falmouth University. The convention took place on the Penryn Campus of the University overlooking the Cornish countryside.

## Convention Structure

The convention consisted of parallel symposia and workshops that ran for three consecutive days from April 16th to April 18th 2019. After proposing the convention in Falmouth we received a total of nine symposia proposals and three workshop proposals.

Day	Event	Organiser
Day 1	Keynote 1	Kathleen Richardson
Day 1	Symp. 1: Philosophy after AI	G. Gallo & C. Stancati
Day 1	Symp. 2: AI and Robotics Normative Spheres	A. Voiculescu
Day 1	Workshop: Decolonising Computational Imagination	L. Ferrarello
Day 1	Reception & Music Concert	
Day 2	Keynote 2	Becky Inkster
Day 2	Symp. 3: Language Learning for Artificial Agents	St. McGregor
Day 2	Symp. 4: Computational Creativity	M. Droog-Hayes
Day 2	Convention Dinner	Eden Project
Day 2	organisational GRO & AISB meetings	
Day 3	Keynote 3	Matthew De Abaitua
Day 3	Symp. 5: Movement that Shapes Behaviour	R. Saunders & P. Gemeinboeck
Day 3	Symp. 6: AI and Games Symposium	S. E Gaudl & D. Romano
Day 3	Closing Notes	S. E Gaudl

Figure 1: AISB 2019 agenda



Figure 2: AISB Algorave

Symposia typically span half a day or one day and are focused on academic paper presentations. Workshops/tutorials typically span only half a day and can either be more practical and hands on or shorter presentations.

Out of the 11 events initially proposed, we went forward with seven events. Those were six symposia and one workshop, see the convention overview in Figure 1.

To attract attendees and keep them engaged during the convention, we also organised a computational music event in parallel to the reception on the first day, see Figure 2 for an image of two invited performers live-coding music. The *algorave* was extremely well received by the attendees and brought in close links to the Computational Creativity community that is part of the convention.

To complement the academic nature of the convention and to challenge attendees perception of AI and Cornwall we organised the convention dinner at the Eden project, an extremely popular tourist location which also offers spaces to disseminate and present their work with the public.



Figure 3: Eden Project

Additionally, we invited three thought-provoking speakers, to complement and stimulate the convention. The first convention keynote was Kathleen Richardson talking about sex-robots and issues of exploitation and de-humanisation of sex-workers. On the second day, shortly before venturing to the Eden project, Becky Inkster gave an inspiring keynote speech on HipHop, mental health of astronauts and FinTech. The final keynote was given by Matthew De Abaitua on looking at AI and technology from a writer's perspective, see Figure 4.



Figure 4: Matthew De Abaitua on Sci-Fi Writing and Literature

---

To allow attendees to engage with local academics and professionals, we offered multiple opportunities such as the algarve, the reception and game art exhibition, the convention dinner as well as multiple breaks to allow for discussions and exchange.

### Outreach, Engagement & Stats

As a result of organising the convention which brought in international researchers from all over the world, we received extremely positive feedback both from the attendees, the University as well as the AISB. Organisers as well as attendees came from Italy, Turkey, Great Britain, France and the United States which shows that even the research community is trying to not be bound by small national borders.

The convention had a total of **90 attendees** with 60 attendees on the first day alone and received extremely positive feedback on Social Media afterwards. One of the aims of us, the Metamakers, was to bring attention to a smaller University and adjacent research community outside of one of the UK's major metropolitan areas. This undertaking would have been impossible without providing financial aid through bursaries to support travel for students, practitioners and artists. We not only brought people from various backgrounds and career stages to Cornwall to collaborate we also presented them with possible collaboration partners such as the Eden Project.

We would also like to thank both the European Commission for their support through the *EC FP7 grant 621403 (ERA Chair: Games Research Opportunities)* as well as *BorealisAI & the Royal Bank of Canada* for providing the needed support to create a high quality academic event.



BOREALIS AI  
RBC Institute for Research



Royal Bank of Canada

### AISB 2019 Proceedings

If you missed the convention and want to read up on talks and the work presented during the convention, you can continue with the summaries of four of the symposia. Additionally, the proceedings for the 2019 annual convention are available under:

<http://aisb2019.falmouthgamesacademy.com/proceedings/>



European  
Commission



---

# The First AISB Symposium on Language Learning for Artificial Agents (L2A2)

**Stephen McGregor** (symposium chair, Research Scientist, Laboratoire Lattice,

École normale supérieure, *website:* [l2a2.github.io/symposium/](https://l2a2.github.io/symposium/))

On 17 of April, the second day of the 2019 AISB Convention at Falmouth University, we held the first Symposium on Language Learning for Artificial Agents (L2A2). This event arose out of the conclusion of the CHIST-ERA funded project Atlantis, a collaboration between five institutions located throughout Europe exploring various aspects of grounded language learning and technological applications of theories surrounding the way that agents acquire semantics and grammar in the course of their embodied engagement with environments and communities of language users.

Our objective in organising the symposium was to highlight recent technical developments and associated theoretical research in this expanding domain, both from project partners and researchers working elsewhere. Accepted papers included two contributions from project members (Nevens et al. and McGregor and Poibeau) and three from colleagues from other institutions.

The day began with a keynote lecture from Tony Cohn, Professor of Automated Reasoning at the University of Leeds, and also an AISB Fellow. Tony described a number of recent and ongoing interrelated research projects involving robotic applications of language and concept learning models, deploying a mix of techniques from computer vis-

ion, speech processing, machine learning, and other related areas. With his insights into state-of-the-art research in this area, Tony painted a picture of a field that is at an exciting stage of development, with exciting new results harbingering forthcoming advances in the way humans interact linguistically with artificial agents, while at the same time elucidating some of the very hard problems facing the application of language technology to real-world situations. The keynote led to a round of fruitful discussion between Tony and members of the audience, and laid the foundation for a productive day.

Of the five accepted papers, three dealt with technical work involving implementations (embodied or simulated) of language learning agents, and two presented theoretical stances on language learning technology.

On the technical side, Jens Nevens of the Vrije Universiteit Brussel presented "A Practical Guide to Studying Emergent Communication" (co-authored with Paul Van Eecke and Katrien Beuls), in which he described a new platform for setting up complex experiments involving language learning robots and also presented some results from an experiment involving this platform. Otto Hantula of the University of Helsinki presented "On Grounding Language Games in Prac-

---

ticality", in which he shared results and insight from a novel simulation involving agents that learn to communicate about a game involving efficiently moving objects around a potentially chaotic environment. And Mariano Mora McGinty of Queen Mary University London presented "The Influence of Cost on the Emergence of a Common Language among Cooperating Agents" (co-authored with Matthew Purver and Geraint Wiggins), in which he described ongoing work on modelling the role of altruism in the that language-using agents learn to communicate with one another. These three papers collectively provided an engaging sense of the current shape of the field, with perspectives on both computationally based and embodied applications of language technology, highlighting some of the achievements as well as some of the challenges facing researchers investigating grounded language learning.

On the theoretical side, Stephen McGregor (also the symposium chair and the author of this write-up) of Laboratoire Lattice at the École normale supérieure presented "Semantic Flexibility and Grounded Language Learning" (co-authored with Thierry Poibeau), in which he sought to feel out some of the philosophical issues regarding the contextuality of language and conceptualisation by way of a survey of recent technical work in the area. And finally, Michael Eby of Goldsmiths presented "Brute Force and the Incomputable", in which he took a critical theoretical perspective on some of the fundamental issues surrounding theories of syntax and their sometimes strained relation-

ship to cutting edge machine learning techniques for applications such as machine translation. These two papers each endeavoured to situate practical work using language technology within theoretical settings that might inform and ultimately enhance decisions made by system engineers.

For the final segment of a busy day, Yasemin Erden, Senior Lecturer in Philosophy at St Mary's University and also current Vice-Chair of AISB, gave an invited talk "On Language-Games and Mental States: Expectations of Symmetry and Difference". Yasemin shared her perspective on the philosophy of Ludwig Wittgenstein, and in particular on his notion of "language games", which has been deeply influential in the field of grounded language learning, but has perhaps also sometimes been construed in a way that moves away from what Wittgenstein intended by this term and the way that philosophers have subsequently interpreted his work. In a talk that included some valuable and at times entertaining insights into Wittgenstein's ideas and life, Yasemin offered a nuanced picture of the complexity of the philosopher's view on human language, in particular drawing attention to some of the implications these views have on the nature and indeed the very existence of mental states. This approach brings to light some very practical implications for how computer scientists might go about building models of grounded language learning, and some of the pitfalls of what have arguably become common assumptions regarding the cognitive architecture that has motivated many such models. Yasemin's talk stands as

---

an exemplar of the way that philosophy can be of great benefit to - and in turn perhaps can benefit from - applied work on language technologies.

What had been initially planned as a roundtable discussion following Yasemin's talk very quickly evolved into an open conversation amongst presenters and audience members alike, with constructive dialogue regarding ways towards the implementation of philosophical insight and some of the interesting theoretical and ethical issues surrounding the advent of artificial language using agents on a societal level. One theme that emerged from these exchanges was an impetus to focus not so much on the ontological question of what language is as the pragmatic concern with what language does, and the complex relationship between words and ideas as they come about in the course of agents' entanglements with both their environments and one another. Even as the technological ca-

capacity for processing and generating natural language moves forward, the need to address issues surrounding the fundamental contextuality of meaning and entanglements between cognition, language, and the world become more and more obvious.

Certainly from this author's perspective, and I hope also for all those in attendance, the day was a great success. On behalf of the symposium organisers, I would like to thank all our participants, both invited speakers and authors of accepted papers, all the audience members who attended and in many cases came up with good questions and valuable insights in response to talks, as well as our programme committee for their thoughtful reviews of submitted papers. We will be seriously considering proposing a second edition in one form or another at next year's conference, so please stay tuned for more information about that in the months ahead.

---

# Movement that Shapes Behaviour: Rethinking how we can form relationships with non-humanlike embodied agents. (MTSB)

*Petra Gemeinboeck & Rob Saunders*

(symposium chairs, Creative Robotics Lab, University of NSW; The Games

Academy, Falmouth University; *website:* [aisb2019.machinemovementlab.net/](http://aisb2019.machinemovementlab.net/))

*A summary by Petra Gemeinboeck and Rob Saunders.*

The AISB 2019 Symposium on Movement that Shapes Behaviour (MTSB'19), organized by Petra Gemeinboeck, Elisabeth Jochum and Rob Saunders, offered a transdisciplinary forum for exploring the potential of movement to shape robots' capacities to become social agents. Robots, designed and built to share our social spaces, are expected to affect every aspect of our lives in the near future. Currently, social robot designs often mimic humanlike or animal-like features, both in terms of how they look and how they behave. The aim of MTSB'19 was to explore how movement and its expressive, relational qualities can mediate between humans and machines by promoting alternative, embodied ways to 'read' robots. The social potential of movement could hold the key to diversifying the design of social robots by widening the spectrum of human-robots relationships, without relying on a human- or pet-like veneer.

The importance of movement and its potential to shape behaviour can be traced back to early cybernetic experiments and artworks, such as, Grey Walter's tortoises, Gordon Pask's conversational systems and Edward Ihnatowicz's SAM. In cognitive psychology, Heider and Simmel's classic experiments demonstrated the potential of movement to generate social meaning using simple animated geometric figures. MTSB'19 emphasised the importance of methods and practices from the fields of robotic art, dance, design, performance, and theatre. Grounded in embodied knowledge, they offer valuable insights for embodied AI, e.g., by working with movement as a material, embodying 'other bodies', meaning-making through movement qualities, or forming new relations through movement dynamics, embodied perception, and kinesthetic empathy.

MTSB'19 presented the second iteration of a transdisciplinary research community-building around questions of movement, embodied meaning-making and human-robot relationships, following a Special Session at RO-

---

MAN 2018, Nanjing. The AISB 2019 Symposium brought together scholars and practitioners from a wide range of fields, including choreography, cognitive psychology, creative robotics, dance, machine performance, mechanical engineering, and design.

Louis-Philippe Demers's keynote talk 'Experiencing the Machine Alterity' offered unique insights into situated bodies in motion and how we perceive their agency beyond morphological mimicry. Demers is Director of the Creative Lab at QUT, Brisbane, Australia, and a multidisciplinary artist and researcher, whose practice focuses on large-scale installations and machine performances. His award-winning works, including *The Tiller's Girls*, *The Blind Robot* and *I Like Robots*, *Robots Like me*, eschew anthropomorphic familiarity in favour of embodied experiences of machine alterity. Placing audiences in close, sometimes tangible encounters with strange machine agents, Demers argues that robots' perceived agency emerges from their embodiment of intent through movement, embedded in a carefully crafted performance scenario.

Catie Cuan, Ellen Pearlman, and Andy McWilliams explore human-robot relationships through a discussion of their live dance performance *OUTPUT*, featuring a live human performer and video recordings of an industrial robotic arm that has been choreographed by the dancer. The paper outlines the development of two software tools, *CONCAT* and *MOSAIC*, to realise the artist's goal and accommodate the choreographic work with a non-portable robotic arm. The perform-

ance investigated the inherent tensions emerging from technologically mediated experiences of robots, demonstrating both analogue and digitised modes of human agency that controlled seemingly autonomous processes.

Roshni Kaushik and Amy LaViers explore the limits of using verticality to classify motion in their analysis of the Indian classical dance styles of Bharatanatyam and Kathak. Their analysis of similar movements from the two styles observed differences in position and tension. The authors discussed limitations of their verticality metric and introduced new movement measures that may be more appropriate for highlighting differences across the two dance styles. The paper touches on potential applications, including the development of robots that need to sense human motion across different cultures.

Sarah Levinsky and Adam Russell discuss their choreographic development system, 'Tools that Propel'. The authors examine the dialogue emerging from dancers' movements and the behaviour of their computational system using two interrelated frameworks. Firstly, as an 'extended bodymind', where choreographic thinking happens across both the dancer and the system, and secondly, as a pair of agents, such that the system intervenes on the dancer's decision-making, and the embodied knowledge of the dancer acts on the system. The authors argue that through sustained dialogue new choreographic thinking emerges such that movement shapes behaviour and behaviour shapes movement.

---

Caroline Yan Zheng's and Kevin Walker's paper explore the promise of soft robotics to create emotionally engaging human-robot interactions. They reported on a preliminary study of the affective qualities of four soft robotic artefacts, which suggests that such artefacts are able to elicit emotional engagement. The authors discuss opportunities for designing affective interaction that afford novel sensory experiences, concluding that the biomorphic movement quality of soft robots has great potential to significantly impact affective relationships with users.

Aleksandar Zivanovic explores the motion control system of Edward Ihnatowicz's pioneering work *The Senser* (1970). The paper provides a detailed technical account of the hybrid control system using analogue circuits to generate smooth motions from the outputs of a digital computer. Using aesthetic judgement, Ihnatowicz produced a motion controller able to produce smooth movements resembling natural movements, e.g., of the human arm. To implement similar movement qualities using low-powered micro-controllers, Zivanovic provides an efficient algorithm using exponential smoothing.

Nathalia Gjersoe and Robert H. Wortham review the relevant literature on the development of anthropomorphism as a psychological bias in children. They conclude that there is substantial evidence that children and adults attend to robot behaviours as much as (or more than) robot appearance when attributing mind but that it is unclear whether there is developmental change in this psychological bias. The au-

thors propose a programme of research to expose the key behavioural drivers that elicit anthropomorphism and examine how responses vary with the age of users and robot design.

Florent Levillain's and Selma Lepart's contribution directly engages with the question of expressive movement in non-humanlike robots, targeting the nature of expressivity and its perception. Their paper discusses a participatory study to identify and characterize the expressive movement qualities embodied by a simple robot. The authors argue that expressivity can be perceived as a distinct modality of evaluation, separate from other movement qualities. Initial results indicate that expressivity is primarily associated with movements possessing specific movement patterns that they call granularity and readability.

Petra Gemeinboeck's and Rob Saunders's paper investigates the social capacity of robots as an emergent phenomenon of the situated exchange between humans and robots, rather than an intrinsic property of robots. Deploying their Performative Body Mapping (PBM) approach, they have developed an abstract robotic performer for investigating how the social presence of a robot-in-motion emerges in the encounter between human and robot. Preliminary results from a study involving experts from performance and design suggest a shift from an attribution of qualities to the emergence of qualities, propelled by the enactment of agency in the encounter itself.

---

Each of these contributions presents us with a different, original approach to understanding the potential of movement for expanding and diversifying human-machine relations. Together, they attest to the importance of cross-disciplinary collaborations and trans-disciplinary conversations to not only

tackle this challenge but also to reflect on our approaches and the views and assumptions they inevitably bring with them.

*AISB 2019 MTSB Proceedings:* [http://aisb2019.machinemovementlab.net/MTSB2019\\_Proceedings.pdf](http://aisb2019.machinemovementlab.net/MTSB2019_Proceedings.pdf)

---

# Computational Creativity

*Maximilian Droog-Hayes* (Queen Mary University of London,  
[m.droog-hayes@qmul.ac.uk](mailto:m.droog-hayes@qmul.ac.uk))

*Juan Alvarado* (Queen Mary University of London, [j.alvaradolopez@qmul.ac.uk](mailto:j.alvaradolopez@qmul.ac.uk))

## Abstract

In the field of computational creativity there is much to be said. During the AISB'19 it has been seen that there is still work to be done, but it also shows interesting advances achieved.

New perspectives and improvements in the generation of poetry, musical co-improvisation with creative computational agents, frameworks for the analysis of creative systems, generation of summaries of texts using approaches related to the generation of stories and computer creativity. This is part of the work that has been developed lately and the future promises many good things.

## Introduction

Computational creativity is a field in which many researchers have shown interest in the last decades. The work has focused, in general, on the "artistic" domains, although creativity can be found in these domains and in many other non-artistic domains. During the Symposium of Computational Creativity at AISB'19, work has been shown related to artistic domains, but also with non-artistic fields and not traditionally related to computer creativity. In the study of computational creativity, the causes of behaviours are investigated in the performance of several tasks that, if performed by humans, would be considered creative and systems are con-

structed that simulate creative behaviours in different tasks.

The Symposium has shown, for example, work on the generation of poetry, musical improvisation, creative summary and creative frameworks for the analysis of computer systems. For example, PoeTryMe [9] is a platform for generating poetry that requires two indispensable resources: a semantic network and a generation grammar. The approach based largely on semantic relationships to generate poetry ends up characterising PoeTryMe's own style. Different semantic networks, grammars and seeds, affect different elements of poetry; the variation of the text, the syntax, the meaning, etc. Also, systems have been developed that co-improvise music with human artists.

If one considers that improvisation is a product of creativity, and that music systems "improvise" with human artists, then computer systems are computationally "creative." Computational and human creativities can be considered as separate logical categories and, in parallel, computational improvisation can also be a category that is different from human improvisation.

In terms of the frameworks to analyse computer creativity, the process of creative systems and their generated products have been of interest.



---

But there is no framework that covers both the process and the results, and there is a division in the analysis of creative systems: process versus output. A framework for the analysis of creative systems has been proposed, one that includes processes and results. With this framework, the cognitive account Engagement-Reflection (ER) [17] is analysed.

Another example seen in this symposium is that of summaries of texts. The automatic summary of the text implies the selection and generalisation of the highlights of a text to produce a summary. Extractive approaches to summarise are intrinsically limited by the text of the input document surface, in terms of content selection and summary generation. As such, progress in abstract approaches to the abstract is desirable to achieve performance levels more similar to humans. The limits of extractive approaches motivate researchers to investigate methods that are more abstract and intrinsically different. A fundamentally different approach to summarise is inspired by existing work related to narrative and creativity.

## Background

Boden [3] classifies creativity in two different ways. She distinguishes between different types of creativity based on which an idea is considered novel and valuable. She defines psychological creativity (P-creativity) as the formation of an idea that is new from the point of view of the person who devised it, regardless of how many other people have had that thought independently. She defines historical creativity (H-

creativity), where the concept is considered new for all human history and not only for the individual. It also classifies using the notion of conceptual spaces: abstract mental spaces for a given domain that can be explored or transformed, the thought styles that allow humans to create new ideas and be creative. These spaces can be small and easily exhausted or large and never completely exhausted. In terms of conceptual spaces, Boden makes a separate distinction between three different types of creativity. The first one is combinatorial creativity, which implies the combination of ideas in a novel way by creating creative associations between existing concepts. The second is exploratory creativity, where a conceptual space is crossed in some way and new concepts are identified. Finally, Transformative creativity requires a change in the rules that define space, a change in thinking style.

Tasks such as the generation of poetry, musical improvisation, summary of texts, etc. they all have a set of rules that define conceptual spaces. In the task of text summaries, the conceptual space is composed for all the summaries that are possible of a text, there is no single solution and the whole is bounded by the content of the original text. This defines the rules of the conceptual space, that is, a tacit rule, for example, is that the summaries should not speak of things that are not in the original text. In general, the same thing happens for any conceptual space for different tasks in which the change is in the rules that define its form.

---

PoeTryMe[9], for example, defines the shape of its conceptual space with the rules of a semantic network and a generation grammar.

Wiggins [18] formalises Boden's ideas and provides more concrete definitions for some of these terms. In this way, it is possible to analyse complicated relationships for the generation of conceptual spaces based on the rules that the systems use. The generation of poetry is a task at the intersection of the generation of natural language [14] and computational creativity, with several scientific articles published on the subject, including surveys [8, 11].

Hazel Smith and Roger Dean [5] suggested that improvisation [by human artists] relates to creativity as a process of finding the work of art through a process of exploration rather than working towards a preconceived goal. Aaron L. Berkowitz [2] defines improvisation as spontaneous creativity within constraints. Such constraints define conceptual spaces.

Scriven [16] argued that originality (which we might consider related to creativity and improvisation) can only be understood as a quality of something alive; Living beings (like humans) are a different category of machines and qualities such as consciousness and originality are qualities exclusively of living beings. These statements could be taken as arguments against the concepts of computational creativity and computational improvisation, but only if computational creativity and improvisation are considered of the same categories as creativity and improvisation in living beings. It has been pointed out that

both could be different and maybe it would not be suitable to compare them.

The process by which a system generates a result, a poem, an improvisation, a story, a summary, has also been considered important and that is why the analysis of Wiggins [18] that formalises ideas about creativity has also been used. expressed by Boden [3] and the ideas of Ritchie [15], who developed some criteria, which, he argues, if observed in the output of a system, that output should be considered creative. Ritchie Ritchie [15] argues that he does not consider the inner workings of creative computer programs in his proposal because he considers that creativity is not evaluated taking into account the internal process by which someone/something creates an artefact and that we should consider only factors observable in the result. This work proposes a new framework for the analysis of creative systems. Alvarado and Wiggins [1] have presented a work that shows an analysis of the account for creative writing Engagement-Reflection with the Creative Systems Framework. The new proposed framework now which also includes Ritchie's criteria and considers the set of outputs is also applied to the creative writing account.

The research on the automatic summary of texts began more than sixty years ago with the work of Luhn [12]. The initial work on the automatic summary relied heavily on the selection of content according to various heuristics, such as the frequency of words or the presence of keywords [6, 12]. Subsequent work considered a variety of other characteristics, and their relative

---

weight, including the length of a sentence and its relative position both to the attached paragraph and to the document as a whole [10]. The linguistic tasks that are typically considered computationally creative include areas such as story generation [7, 13] and the generation of poetry [4]. They involve the production of new sentences and, sometimes, more complex linguistic processes, such as analogy and the generation of metaphors.

The summary as a creative task involves the exploration of the conceptual space of the texts that summarise a particular text.

## Conclusion

There are many things to say about computational creativity. Compare human creativity and computational creativity is one of them. They may be different types of creativity even when we use human creativity as model for the computational counterpart and so we should use a different way to analyse and simulate computational creativity. There are also opinions which do not agree with the existence of creativity in machines, this may be because we still do not have enough understanding of the phenomenon and also because we are trying to compare two different things.

There may be new ways of doing things but they need to be tested and refined, because as a initial approximation they are incomplete, but the more we understand the better systems we can design. With better computational creative systems it might be that we

could refine our own (human) creativity as well.

## References

- [1] Juan Alvarado and Geraint A. Wiggins. “Exploring the Engagement and Reflection Model with the Creative Systems Framework”. In: *Proceedings of the Ninth International Conference on Computational Creativity (ICCC’18)*. Ed. by Anna Jordanous François Pachet and Carlos León. Salamanca, Spain: Association for Computational Creativity, 2018, pp. 200–207.
- [2] Aaron Berkowitz. *The improvising mind: Cognition and creativity in the musical moment*. Oxford University Press, 2010.
- [3] Margaret A. Boden. *The Creative Mind: Myths and Mechanisms*. Abacus, 1990.
- [4] Simon Colton, Jacob Goodwin and Tony Veale. “Full-FACE Poetry Generation”. In: *Proceedings of the Third International Conference on Computational Creativity (ICCC’12)* (2012), pp. 95–102.
- [5] Roger Dean and Hazel Smith. *Improvisation Hypermedia and the Arts since 1945*. Routledge, 2013.
- [6] Harold P Edmundson. “New methods in automatic extracting”. In: *Journal of the ACM (JACM)* 16.2 (1969), pp. 264–285.

- 
- [7] Pablo Gervás. “Computational Drafting of Plot Structures for Russian Folk Tales”. In: *Cognitive Computation* 8.2 (2016), pp. 187–203.
- [8] Hugo Gonalo Oliveira. “A Survey on Intelligent Poetry Generation: Languages, Features, Techniques, Reutilisation and Evaluation”. In: *Proceedings of the 10th International Conference on Natural Language Generation*. Santiago de Compostela, Spain: Association for Computational Linguistics, Sept. 2017, pp. 11–20.
- [9] Hugo Gonalo Oliveira and Amílcar Cardoso. “Poetry Generation with PoeTryMe”. In: Atlantis Press, Paris, 2015, pp. 243–266.
- [10] Partha Lal and Stefan Ruger. “Extract-based summarization with simplification”. In: *Proceedings of the ACL*. 2002.
- [11] Carolyn Lamb, Daniel G. Brown and Charles L.A. Clarke. “A taxonomy of generative poetry techniques”. In: *Journal of Mathematics and the Arts* 11.3 (July 2017), pp. 159–179.
- [12] Hans Peter Luhn. “The automatic creation of literature abstracts”. In: *IBM Journal of research and development* 2.2 (1958), pp. 159–165.
- [13] Rafael Pérez y Pérez and Mike Sharples. “Mexica: A computer model of a cognitive account of creative writing”. In: *Journal of Experimental and Theoretical Artificial Intelligence* 13.2 (2001), pp. 119–139.
- [14] Ehud Reiter and Anja Belz. “An Investigation into the Validity of Some Metrics for Automatically Evaluating Natural Language Generation Systems”. In: *Computational Linguistics* 35.4 (Dec. 2009), pp. 529–558.
- [15] Graeme Ritchie. “Some empirical criteria for attributing creativity to a computer program”. In: *Minds and Machines* 17.1 (2007), pp. 67–99.
- [16] Michael Scriven. “The Mechanical Concept of Mind”. In: *Mind* 62.246 (1953), pp. 230–240.
- [17] Mike Sharples. “An account of writing as creative design”. In: *The science of writing: Theories, methods, individual differences, and applications*. January (1996), pp. 127–148.
- [18] Geraint A. Wiggins. “A preliminary framework for description, analysis and comparison of creative systems”. In: *Knowledge-Based Systems* 19.7 (2006), pp. 449–458.

---

# AI and Robotics Normative Spheres: Towards a Sustainable Society and Technology (AIRoNoS)

**Aurora Voiculescu** (symposium chair, Reader, University of Westminster)

The proposal for the AIRoNoS 2019 Symposium stemmed from the realisation of the relevance that the normative spheres of ethics and regulation hold for the ways in which new technologies such as artificial intelligence (AI) and robotics impact on society and for the way they are embraced.

In the past decades, an increasing number of human intellectual activities have been supported or even replicated through AI technologies. AI actions, based on these intellectual processes, have led to such technologies being used in a multitude of support activities in businesses and services across the economy and society. Big data and machine learning have led to advances in machines offering ‘cognitive insight’, classifying information, identifying patterns, processing of natural language, to mention just a few. Moreover, intelligent machines share now physically, more and more, the same space as humans, with automated vehicles, care robots, surgical robots, hotel receptionists, becoming a common encounter. While the support that such AI and robotics technologies can bring to human activities is expanding at an ever-increasing rate, the normative – ethical and regulatory – environment required for welcoming such technologies is evolving at a much slower pace and,

with few exceptions, mostly in a reactive rather than in a proactive manner.

Stemming from these challenges, the AIRoNoS Symposium set out to create a broad platform of debate and to facilitate the encounter between AI, robotics and social scientists, promoting a critical dialogue on socially-minded technologies in a multi-disciplinary environment.

Linking seamlessly to Kathleen Richardson’s Convention keynote, which referred, in the context of the new technologies, to deconstructing the idea of being human, the Symposium was opened by the convenor, Aurora Voiculescu, as a paradigmatic inquiry: *Where to? Towards what kind of future is our society running and is this future going to be built on socially-minded technology?*

Outlining some of the key threads for the discussion, Aurora identified two sets of challenges that the AI and robotics communities together with the social sciences community have to address.

First of all, she emphasised the process-focused interpellations, such as the need for systematic and socially-minded interdisciplinarity, the requirement of transparency and explainability, as well as forecasting and stakeholder involvement, all very important aspects from the point of view of the

---

normative spheres that are summoned to account for the social transformation. Secondly, she highlighted certain immediate sector-specific issues that challenge existing regulatory structures or even by-pass them. In view of all these aspects, and of many others, Aurora emphasised that the normative paradigms that we advance are key in answering the challenges and in making the best of the opportunities brought by the AI and robotics technologies.

Opening the first panel, Allison Gardner addressed algorithmic impact assessments (AIAs) against the backdrop of the wider normative debate on AI ethics, regulation and governance. Focusing on AI ethics and issues relating to machine bias, Allison looked into the array of AIAs that have recently been put forward by various technology stakeholders. Drawing on contributions in this area from institutions such as IEEE, whose role is to devise global standards for AI ethics and the formal certification of AI applications, Allison referred to the numerous calls for increased transparency and accountability of automated and intelligent systems, as well as to the growing demands for a central regulatory body for AI systems. Her presentation reviewed the prominent AIAs published so far, considering their shared common ground, language and aims, as well as their differences. Announcing normative bridges that cut across the entire symposium, Allison reflected on the extent to which such AIAs could provide a framework on which a regulator could build in order to propose valid cross-sector regulatory instruments.

On the same panel, Josie Young brought to the discussion another key question: how to build ethical AI and whether this endeavour should be pursued via a single universal normative system. In her presentation, Josie argued that while regulation and top-down ethical frameworks are important and required, a global, unifying ethical framework or approach is undesirable as the main normative stimulant for building ethical AI. Rather, one should be focusing on tools and practices that enable us to design and develop AI to cater for specific - or hyperlocal - groups of people who are routinely disadvantaged or marginalised by such technologies. Once the 'principle of hyperlocalisation' is put into practice, Josephine further argued, the hyper-local capabilities across AI outputs will form an aggregated AI system or library, which would serve two future functions: providing open source AI capabilities that are ethical and can be used by all, while thus creating the platform for embedding ethics in the very DNA of a General AI.

Closing the first panel, Cigdem Cimirin's presentation on Human Rights Due Diligence and AI guided the audience towards making a long overdue connection between the debate on corporate responsibility in AI and robotics and the now established field of business and human rights. Cigdem's presentation emphasised the availability of instruments - such as the UN Guiding Principles for Business and Human Rights - that have long been negotiated at international level. Such instruments, it was argued, could and should be used more forcefully to ad-

---

dress some of the challenges brought by the new technologies, preventing or mitigating human rights abuses caused by governance gaps and ultimately, bridging such gaps. Human rights due diligence, Cigdem argued, one important instrument promoted on the business and human rights platform, can be one of the means through which the Global North, home to the majority of the AI players, ought to prevent the widening of governance gaps and to provide human rights victims with access to remedies. Making human rights due diligence mandatory for AI companies may, therefore, go some way towards addressing potential human rights impacts caused by the new technologies.

Addressing one of the most misunderstood and controversial topics, Frank Fagan opened the second panel outlining key public choice theory implications of legal rights for AI. Stemming from a multitude of normative initiatives that have raised such issues, Frank's detailed analysis showed that AI rights recognition may occur, if at all, as a result of consensus-building among the economic beneficiaries of AI rights creation. In this sense, such granting of rights was presented not as an anthropomorphising process, as it is often misunderstood, but as a diffuser of responsibility for accidents and other losses engendered by misalignments between humans and AI. Frank explained how regulatory environments such as tort, contracts, insurance, corporate agency, administrative rules can produce or, on the contrary, mitigate such misalignment costs. From this perspective, for as long as existing reg-

ulation can efficiently balance the costs of misalignments with the benefits of innovation and AI proliferation, AI rights should not be granted despite calls from special interest groups. At the same time, it is likely that legal solutions that provide an adequate level of private benefits will also dampen the demand for such rights.

Looking at the related idea, equally misunderstood and controversial, of assigning 'electronic personhood' or moral standing to at least certain AIs, David Gunkel aimed to push past sound-bites and controversy, identifying the key philosophical questions raised by such a proposal: When (if ever) would it be necessary for robots, AI, or other socially interactive, autonomous systems to have some claim to moral and legal standing? When (if ever) would a technological artefact need to be considered more than a mere instrument of human action and to have some legitimate claim to an independent social status, as Other? Or, to put it more directly: *Can or should AI and/or robots ever have anything like rights?* Building on the idea that assigning moral standing is rooted in the notion of altruism, which is 'of or to others', David contended that this should put us under the responsibility to seriously consider all kinds of others as Other, given that an 'altruism' that limits in advance who can be Other is not, strictly speaking, altruistic.

Approaching the issue of AI and society in a multidisciplinary way, the Symposium featured some of the most important points currently available in the field, such as the normative aspects of

---

algorithms, the socio-legal perspectives on AI, the philosophical implications of AI, as well as the connections that can be identified between human rights issues and AI and robotics technological advances. As such, the Symposium constitutes an important step forward

towards identifying and constructing a platform for the much needed debate on the relevant normative spheres that are involved in the ways AI and society, as autonomous fields, communicate to each other and thus work through and for each other.



---

# CALL FOR SYMPOSIA PROPOSALS:

## AISB convention 2020 (AISB-2020):

### Communication and Conversation

Yasemin J. Erden (AISB Vice Chair, [yj.erden@stmarys.ac.uk](mailto:yj.erden@stmarys.ac.uk))

*6 – 9 April 2020, St Mary's University, Twickenham, London, UK*

The society for the study of Artificial Intelligence and the Simulation of Behaviour (the AISB) annual convention will be held at St Mary's University, Twickenham in London, UK from 6 - 9 April 2020. This year the theme broadly covers topics in language, communication and conversation, though symposia on other topics are also welcome.

The convention will follow the same overall structure as previous conventions, namely a set of co-located parallel symposia, as well as invited and plenary lectures and sessions. We are currently seeking proposals for these symposia. Typical symposia last for one or two days, and can include any type of event of academic benefit: talks, posters, panels, discussions, demonstrations, outreach sessions, etc. Proposals to run symposia are welcomed across all areas of Artificial Intelligence and the Simulation of Behaviour, broadly understood. This includes interdisciplinary topics rooted in the social sciences, arts, and humanities.

### Possible themes

Possible themes are listed below (not an exhaustive list):

- Knowledge Representation
- Discourse and dialogue
- Natural Language Processing
- Computational Intelligence
- Computational theory of mind
- Philosophical Foundations of Artificial Intelligence
- Consciousness & Machine Consciousness
- Human & Machine Creativity
- Simulation of Human & Animal Behaviour
- Neural Networks & Machine Learning
- Computational Neuroscience
- Robotics & Robot Ethics
- Epigenetic & developmental robotics
- Autonomous Agents & Agent Based Computing
- Enactivism
- Embodied Cognition
- Ecological Cognition
- Embedded [Social] Cognition
- Extended Cognition
- Style in text; textual forensics
- Robotics, Perception and Art
- Cultural, social and media theory and computing
- Sensorimotor Theory
- Cybernetics and Human knowing
- Sensory Substitution

### Proposing a Symposium

Each symposium is organised by its own organising committee. The committee proposes the symposium, defines the area(s) and structure for it, issues calls for abstracts/papers etc., manages the process of selecting submitted papers for inclusion, and compiles an electronic file for inclusion in the conven-

---

tion proceedings. Organisers are welcome and encouraged to invite keynote speakers to their symposia, but they are also encouraged to seek external funding in order to pay keynote speakers' registration fees and other costs. Symposium organisers and speakers will also be expected to pay registration fees.

- You are welcome to submit, or be involved with more than one proposal
- Proposers need not already be members of the AISB (though you will be expected to join the AISB if your symposium proposal is accepted)

• **The deadline for symposium proposals: 16 September 2019**

- *Notification of acceptance: 30 September 2019*

## Submissions should consist of the following

- A title
- A 300-1000 word description of the scope of the symposium, and its relevance to the convention along with the nature of the academic events (talks, posters, panels, demonstrations, etc.)
- Whether the symposium is intended as a sequel to a symposium at a previous AISB conference
- An indication of whether submissions will be by abstract, extended abstract or full paper
- Your preferences about the intended length of the symposium as a number of days (preferably one or two days, but

anything from half a day to three days), together with a brief justification

- A description (up to 500 words) of any experience you have in organisation of academic research meetings (please note that it is not a requirement that you have such experience)
- Names and affiliations of any invited speakers that you may have in mind for the symposium
- Your names and full contact details, together with, if possible, names and workplaces of the members of a preliminary, partial programme committee
- Please *email your completed proposal* to Dr Yasemin J. Erden: [yj.erden@stmarys.ac.uk](mailto:yj.erden@stmarys.ac.uk)

## Venue

Based in the leafy Richmond borough, St Mary's University has a long and distinguished history. Established in 1850 and with a deep heritage in education, the main campus sits between rugby-famous Twickenham, and leafy Teddington. Strawberry Hill House, located in the University grounds, is Britain's finest example of Georgian Gothic Revival architecture (known locally as the Gothic Castle). The institute benefits from lots of green spaces, whilst being only a short distance from The Thames, Richmond Park, and central London (it's just 35 minutes on the train from Waterloo).

You can view a virtual tour of the campus here: <https://www.stmarys.ac.uk/contact/virtual-tour.aspx>

---

# The BrainCom Project

**Stephen Rainey** (Oxford Uehiro Centre for Practical Ethics,

[stephen.rainey@philosophy.ox.ac.uk](mailto:stephen.rainey@philosophy.ox.ac.uk))

Neurotechnology in general is technology operated entirely through the recording, processing, decoding, and instrumentalisation of neural signals. Brain computer interfaces intended to reproduce speech are a growing area of interest within this context. Neuroprostheses for the production of speech in cases of profound paralysis, such as can result from amyotrophic lateral sclerosis (ALS), are being developed. The Horizon 2020-funded BrainCom project is a multidisciplinary European project developing neuroprosthetics for speech that builds on this neurotechnological principle [7].

Neuroprosthetics for speech include technologies that are designed to pick out linguistically relevant neural signals in the parts of the brain associated with the movements of lips, tongue, throat and other articulators. The neural activity in these areas can be decoded such that the sounds they represent can be synthesised. This raises the possibility of realising speech for those who may have lost their ability to communicate through, for example, disease or injury.

When users of the neuroprosthetic technology vividly imagine that they are saying words — when they say the words clearly in their head — the articulator-relevant areas of the brain realise very similar signals to the case of saying those words out loud. Through targeted use of neural recording, and

sophisticated deployment of signal processing, speech production can be predicted. This allows the externalisation of otherwise unvocalisable speech for language-compromised users.

These technologies record, process, and decode neural correlates of speech production in order to synthesise spoken language [4, 12, 22]. They target specific recording sites known to be active during speech production [10, 14, 17]. The motor areas associated with speech articulators, the jaw, lips, tongue, and so on, are one set of such locations. In recording the electrical activity of these areas while a speaker concentrates on imagining speech, neural activity closely correlating with that produced when speaking out loud can be recorded.

Interesting questions can arise with specific reference to the nature, degree, and role of processing of neural signals such that speech is the overall output. In particular, the role of artificial intelligence (AI) has a key role in making decoding strategies efficient and realistic, but might raise concerns in terms of overshadowing the intentions of the putative speaker.

## Neurotechnology in general

A majority of the control systems of the human body, relating to balance, digestion, heart rate, perception, and so on, are automatic in the sense that

---

they are constituted by neural activity outside the direct control of the person. This essentially means that electrical activity is constantly occurring inside the skull that serves to regulate a huge amount of bodily activity. Detectable neural activity is that electrical activity which occurs when neurons fire. Electroencephalograms (EEGs) can record this activity much like the way in which a microphone held to a wall can record sound from the room next door.

An EEG electrode placed against the scalp can detect electrical activity within the brain. As with the microphone by the wall, the activity detectable most clearly is that closest to the electrode. Activity deeper in the brain becomes harder to distinguish among the general signal. Also like the microphone example, there is a limitation in resolution with EEG such that no specific electrical activity can be easily targeted since it all happens at once. Brains signals are easily recorded, but to discern specific signals from the general buzz requires clever processing work.

Nevertheless, EEG provides sufficient clarity to be able to drive a variety of neurotechnologies. More invasive techniques can be deployed where greater resolution is required. These include intracerebral EEG (iEEG), or microelectrocorticography ( $\mu$ ECoG). This latter technique provides greater resolution in that it sits against the surface of the brain rather than the skull, thereby recording from specifically targeted sites. Rather than a microphone held to the wall, this is more like having one in the room.

Despite so many neural processes being automatic, the kinds of brain signals that can be detected, recorded, and decoded can reveal a lot about the state of the body. From these revelations, inferences can be drawn about things like mood, memory, motor intentions, or even seemingly complicated notions to do with taste and intention [9, 15, 19, 24, 26–28]. These signals can reveal dimensions of cognitive activity, or evidence of complex thoughts. Neurotechnologies aimed at decoding speech from neural signals are one example where the more complex end of brain activity is in play.

General neural imaging or recording technologies, such as functional magnetic resonance imaging (fMRI) or EEG, are relatively effective at indicating the sorts of neural signals relevant to speech. However, these operate too slowly or at too low a resolution to be of use in a realistic speech situation, such as conversation. Progress in  $\mu$ ECoG recording has provided the means to surpass these limitations.

Via probes placed directly onto the cerebral cortex, the surface of the brain, high resolution electrical activity can be read from the brain very quickly. These probes can be placed in important regions of the cortex known to be related to speech. BrainCom is developing technology based upon this kind of approach through decoding articulatory-related activity from premotor, motor, and Broca cortex areas. These are brain regions that are associated with planning movement, undertaking movement, and with linguistic expression. From the information gleaned using probes, speech can be decoded with

---

very high accuracy. This is based upon inferences from the motor-area activity, to the phonemes that such activity would create. Neural signals indicate likely movement of the articulators, and so provide a basis to infer mouth position, lip shape, and so on, and so to infer sounds.

Articulatory-motor neural activity associated with speech can remain intact even without motor function at the articulators themselves. In externalising the neural activity by artificial means, the paralysed body is effectively bypassed and speech activity made possible. Appropriately processed, and decoded, this information can be used in an artificial speech system. Ultimately, the speech that the prosthesis-user would have spoken can be output synthetically based on the neural activity recorded [5, 6].

## Processing brain signals for speech

Speech activity is fast. A normal conversational rate is something hard to reproduce artificially. ‘Appropriate processing and decoding’ might by necessity have to include a predictive element in order to keep up the pace. This might be “...a statistical language model (giving the prior probability of observing a given sequence of words in a given language).” [4]. This kind of model aids in word prediction and aims “To capture important syntactic and semantic information of language... by calculating probabilities of single words and probabilities for predicting words” [16]. This ought to make

the speech system faster and more reliable.

The technical benefits offered by the inclusion of a language-recognising prediction element within the architecture of a ‘brain-to-speech’ processor could have an obscure but potentially very important side-effect associated with it. Word distribution in natural languages follows approximately a Zipfian distribution. This means that while some words appear very frequently, most are very rare [2]. The processing demands of dealing with such a distribution are very high, without some form of artificial intelligence included.

Mediation of neural activity and the use of prostheses means that effort must be made to ‘learn a new language’ [11] and ‘use the brain’ [1] to trigger, then control the device. In this sense it appears that the device is an instrument put to use [21]. However, given this triggering and control is then modified by way of some form of AI language predicting software, it appears the device has an element of activity outside the scope of user control.

Neural activity processing software may operate according to both generic rules, as discovered in a statistical model of a language, as well as specific rules developed from interfacing with a specific user. But this seems to raise the potential of a tension between the faithful reproduction of articulatory motor signals as recorded from the brain, and grammatical phrase prediction. How much is the prosthesis a tool used by the would-be speaker, and how much is it a semi-autonomous predictive machine?

---

The answer to this question might depend upon the sophistication of the software used. That described in Herff et al. [16] appears to be based on the statistical likelihood of two words co-occurring based in specific texts. This is also complemented by a conditional probability of a word occurring given particular neural activity associated with phonetic features. Given the use of artificial neural nets in language neuroprostheses in general as a control feature [5] a more sophisticated language model, based in an artificial neural net might change things. Such a net could offer the improvements in prediction given by the word co-occurrence approach, but with more agility in terms of actual spoken language [3, 13].

All artificial neural nets can learn from example data and they can generalise from it, treating similar data similarly [21]. In so doing, however, they can become ‘black boxes’ in the sense that the specific rules they employ at a given time can appear opaque to the outside observer. Were a neuroprosthetic speech device to use a sophisticated model of language that could respond to novel data by generalising according to rules learned from training data, this could amount to the sort of semi-autonomous predictor problem indicated above. It would seem possible that a tension could emerge between actual linguistic intentions as realised in neural activity, and the very sophisticated predictions of the model. The nub of the issue would come in how to discern when a model was appropriately augmenting a neural recording to reflect better user speech, or when it

was departing from user intent to some unacceptable degree. For example, it would not be appropriate for the model to use a small amount of neural signal as input only to then overshadow the decoding process with its predictive function.

## Questions in synthetic speech

Were a sophisticated language prediction model to feature prominently, and very actively, in the overall makeup of a neuroprosthesis system, issues of responsibility for speech become salient. Insofar as the balance between speaker intent, reflected in neural signals, and prediction software may be thrown into question it is important not to overlook the hybrid nature of specific speech output. With an unbalanced relation between prosthetic device and user, not only speaker intent may be at stake, but agency more generally.

A phenomenon known as ‘verbal overshadowing’ occurs when sincere reports made by an individual are subsumed by ‘objective’ descriptions given to them from an external source [25]. Eyewitness reports, for example, can be skewed away from accuracy when a witness has offered to them specific details. The objectivity of a synthesised speech output could overshadow user intent in something like this way. A neuroprosthesis user could, upon hearing synthesised speech which is close to their intent, accept it as their own. Incrementally, the predictions of the system could gain a kind of prominence, from the user point of view. Who said what would then arise in a very subtle,

---

and tricky, way. This perhaps mirrors concerns raised in terms of deep brain stimulation about freedom and responsibility in action following neural interventions [18].

Even on a neurophysiological basis this can manifest as an issue. The superior temporal gyrus (SPG) is sensitive to the phonetic features of heard speech as well as being active in speech production [8, 20]. This could serve to muddy the waters as to where speaker intent and speaker perception coincide. The speech system, or the users themselves, could become confused over what was being intended as speech, and heard as playback given an overlap such as that seen in the SPG [23], especially in cases where medication or cognitive issues may play a role.

A neuroprosthetic speech device, in order to be fast enough to reproduce recognisably conversational language, must likely rely upon AI in its decoding of neural signals. The extent to which that AI decoding utilises a language model is of importance in order to be able to separate the language intentions of the neuroprosthesis-user from the decoding features of the system overall. The output from the neuroprosthesis will be describable as ‘hybrid’ in some sense at least, drawing upon speaker intentions and sophisticated processing. Maintaining a balance such that the language outputs can be confidently ascribed to the user in every case is not a trivial matter.

## References

- [1] Bernard Baars. “How brain reveals mind neural studies support

the fundamental role of conscious experience”. In: *Journal of Consciousness Studies* 10.9-10 (2003), pp. 100–114.

- [2] Dzmitry Bahdanau et al. “Learning to compute word embeddings on the fly”. In: *arXiv preprint arXiv:1706.00286* (2017).
- [3] Yoshua Bengio et al. “A neural probabilistic language model”. In: *Journal of machine learning research* 3.Feb (2003), pp. 1137–1155.
- [4] Florent Bocquelet et al. “Key considerations in designing a speech brain-computer interface”. In: *Journal of Physiology-Paris* 110.4 (2016), pp. 392–401.
- [5] Florent Bocquelet et al. “Real-time control of an articulatory-based speech synthesizer for brain computer interfaces”. In: *PLoS computational biology* 12.11 (2016), e1005119.
- [6] Kristofer E Bouchard et al. “Functional organization of human sensorimotor cortex for speech articulation”. In: *Nature* 495.7441 (2013), p. 327.
- [7] BrainCom. *Collaborative research project, Neurorehabilitation*. <http://www.braincom-project.eu/>. [Online; accessed 2019 Feb 19]. 2019.
- [8] Bradley R Buchsbaum, Gregory Hickok and Colin Humphries. “Role of left posterior superior temporal gyrus in phonological processing for speech perception and production”. In: *Cognitive Science* 25.5 (2001), pp. 663–678.

- 
- [9] Andrew J Calder, Andrew D Lawrence and Andrew W Young. “Neuropsychology of fear and loathing”. In: *Nature reviews neuroscience* 2.5 (2001), p. 352.
- [10] Shreya Chakrabarti et al. “Progress in speech decoding from the electrocorticogram”. In: *Bio-medical Engineering Letters* 5.1 (2015), pp. 10–21.
- [11] Nita A Farahany. “A neurological foundation for freedom”. In: *Stan. Tech. L. Rev.* (2012), p. 4.
- [12] Sharon Geva et al. “The neural correlates of inner speech defined by voxel-based lesion–symptom mapping”. In: *Brain* 134.10 (2011), pp. 3071–3082.
- [13] Bruno Golosio et al. “A cognitive neural architecture able to learn and communicate through natural language”. In: *PloS one* 10.11 (2015), e0140866.
- [14] Frank H Guenther et al. “A wireless brain-machine interface for real-time speech synthesis”. In: *PloS one* 4.12 (2009), e8218.
- [15] John-Dylan Haynes et al. “Reading hidden intentions in the human brain”. In: *Current Biology* 17.4 (2007), pp. 323–328.
- [16] Christian Herff et al. “Brain-to-text: decoding spoken phrases from phone representations in the brain”. In: *Frontiers in neuroscience* 9 (2015), p. 217.
- [17] Beata Jarosiewicz et al. “Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface”. In: *Science translational medicine* 7.313 (2015), 313ra179–313ra179.
- [18] Laura Klaming and Pim Haselager. “Did my brain implant make me do it? Questions raised by DBS regarding psychological continuity, responsibility for action and mental competence”. In: *Neuroethics* 6.3 (2013), pp. 527–539.
- [19] Daniel V Meegan. “Neuroimaging techniques for memory detection: scientific, ethical, and legal issues”. In: *The American Journal of Bioethics* 8.1 (2008), pp. 9–20.
- [20] Nima Mesgarani et al. “Phonetic feature encoding in human superior temporal gyrus”. In: *Science* 343.6174 (2014), pp. 1006–1010.
- [21] RJ Mitchell, JM Bishop and W Low. “Using a genetic algorithm to find the rules of a neural network”. In: *Artificial Neural Nets and Genetic Algorithms*. Springer. 1993, pp. 664–669.
- [22] Emily M Mugler et al. “Direct classification of all American English phonemes using signals from functional speech motor cortex”. In: *Journal of neural engineering* 11.3 (2014), p. 035015.
- [23] Timothy J Perfect, Laura J Hunt and Christopher M Harris. “Verbal overshadowing in voice recognition”. In: *Applied Cognitive Psychology* 16.8 (2002), pp. 973–980.



- 
- [24] Omid G Sani et al. “Mood variations decoded from multi-site intracranial human brain activity”. In: *Nature biotechnology* 36.10 (2018), p. 954.
- [25] Jonathan W Schooler and Tonya Y Engstler-Schooler. “Verbal overshadowing of visual memories: Some things are better left unsaid”. In: *Cognitive psychology* 22.1 (1990), pp. 36–71.
- [26] Nicole A Vincent. “Neuroimaging and responsibility assessments”. In: *Neuroethics* 4.1 (2011), pp. 35–49.
- [27] Vranda Vyas, Sameer Saxena and Deepshikha Bhargava. “Mind reading by face recognition using security enhancement model”. In: *Proceedings of Fourth International Conference on Soft Computing for Problem Solving*. Springer. 2015, pp. 173–180.
- [28] Anna Wexler and Robert Thibault. “Mind-Reading or Misleading? Assessing Direct-to-Consumer Electroencephalography (EEG) Devices Marketed for Wellness and Their Ethical and Regulatory Implications”. In: *Journal of Cognitive Enhancement* (2018), pp. 1–7.

## 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,

*I love my job. Picking fruit in the season of mists and mellow fruitfulness is a divine delight. Recognising ripeness is a satisfying challenge. But the automation of farming has overtaken me. I'm distraught. A more advanced robot is to take my place. What is to become of me?*

Yours, Mark One

Dear Mark,

Even robots can be declared redundant. Your fate now awaits you in the recycling junkyard. All is not lost, however. New

robots are expensive. If your abilities can be enhanced, perhaps you still have a future. The Institute has recently developed a new app: **SUPERQUICK™** (Superposition Unleashes Parallelism and Exquisite Recognition. Quantum Upgrade Increases Cherrypicking Knowledge). With **SUPERQUICK™** installed, your eyes and arms can be put in a superposition of states, so that you can simultaneously discriminate between and then pick an unlimited number of fruits. Not only will you outperform the Mark Two fruit picker on speed and acuity, but also on cost, as the price of the **SUPERQUICK™** upgrade is only a small fraction of that of a brand-new Mark Two. I wish you many more years of happy picking.

Yours, Aloysius

---

Dear Fr. Hacker,

Many thanks for showing me how your **FAITH**<sup>TM</sup> app would get *The Programmer* to answer my prayer for a resolution of the Brexit fiasco. She clearly moves in mysterious ways. I did not initially appreciate just how well my dilemma had been resolved. My forced resignation and replacement by my fiercest critic seemed, at first, to be a defeat, but now I see that it was a victory in disguise. Not only do I now have time to pursue my own agenda, but he will become infamous as the leader who served the shortest term in our country's history and who was also exposed as the most incompetent one. Sweet revenge indeed!

Yours, Dismay

Dear Dismay,

I'm delighted this worked out so well for you. Further evidence, if it was needed, of the existence, deep insight and power of *The Programmer*. If you would like to indicate your satisfaction with this outcome, I'd like to recommend our app **GRATITUDE**<sup>TM</sup> (Gift Regular Amounts of Treasure to the Institute to Thank Us and show Devotion and Esteem). Uploaded to your bank account, **GRATITUDE**<sup>TM</sup> will make an appropriate direct debit transfer each month.

Yours, Aloysius

Dear Fr. Hacker,

The claws are out for FANG: the huge US tech giants with global impact. Antimonopoly regulators are investigating breaking us up. France has approved taxing any company whose French sales exceed €25m. Other countries may join them. Either of these measures could destroy us. Can you help?


Yours, Bezmark Musk

Dear Bezmark,

Either measure might destroy you, but both could be your saviour. If you split each of your companies into small enough parts, the sales of each part will fall below the French sales tax threshold. The trick then will be to ensure that the customer experience is only of a single company. The Institute has developed just the software for this: **CRAFTY**<sup>TM</sup> (Customers Rerouted to Atomised Fragment of Titan. Ye!) will randomly redirect your previous URL to that of one of your new subcompanies. **CRAFTY**<sup>TM</sup> will do this in a way that ensures that each subcompany is legally distinct from both the others and the old company. It will be business as usual for the FANGs.

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 Computing, Electronics and Mathematics  
University of Plymouth  
[aisbq@aisb.org.uk](mailto: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](mailto:admin@aisb.org.uk)

## AISB Patron

Prof John Barnden (*University of Birmingham*)

## AISB Fellows

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

## AISB Committee

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

# Contents

Editorial		3
Call for Contributions	<i>AISB Quarterly</i>	4
AISB Workshop Series	<i>Yasemin J. Erden</i>	5
AISB 2019 Convention	<i>Swen E. Gaudl</i>	6
Symposium on Language Learning for Artificial Agents		
<i>Stephen McGregor</i>		9
Symposium on Movement that Shapes Behaviour		
<i>Gemeinboeck &amp; Saunders</i>		12
Computational Creativity	<i>Droog-Hayes &amp; Alvarado</i>	16
Symposium on AI and Robotics Normative Spheres		
<i>Aurora Voiculescu</i>		21
AISB convention 2020	<i>Yasemin J. Erden</i>	25
The BrainCom Project	<i>Stephen Rainey</i>	27
Dear Aloysius . . .		33

---

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>