

AIQB QUARTERLY

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



Do you feel artistic?
Exhibit your artwork on our front covers!
Email us at aisbq14@aisb.org.uk!

This artwork is based on Particle Swarm Optimization (PSO) algorithm. PSO is an evolutionary computation technique developed in 1995 by Kennedy and Eberhart, and it's inspired by social behavior of bird flocking or fish schooling.

Particles in PSO are made to follow a hypothetical point (focal point, f_p) moving horizontally (i.e. scanning each row with a constrained random vertical offset); once f_p reaches the end of a line, it goes to the next row; this process is repeated until the entire input image (the Q of the Quarterly) is scanned. As particles trace the f_p , the average colour of the pixel, where each particle is 'flying over', is taken and the colour is reflected on the output image (the current cover of the Q). In other words, the colour of each pixel of the input image changes the corresponding part of the output image to the swarm's average colour as they flock around f_p . The swarm's flocking behaviour over the image creates the resultant cover for this issue of the Quarterly.

The technique will be presented in detail at the A-EYE Art Exhibition symposium at AISB50, 1st April–4th April. See website for more information:
<http://aisb50.org/a-eye-an-exhibition-of-art-and-nature-inspired-computation/>

As Seen by Birds, the Q Magazine: © Al-Rifaie, Asmaa Majid.

Russell Eberhart & James Kennedy, 'A new optimizer using particle swarm theory', *Proceedings of the Sixth International Symposium on Micro Machine and Human Science MHS'95*, IEEE, (1995).

Editorial

Welcome to Q138! Following from our recent cover art experiment, this issue's cover features a contribution from the talented Al-Rifaie family, this time from Asmaa. Asmaa is currently studying for a BSc in Creative Computing through the University of London's International Programme, and sister of last issue's artist and committee member, Mohammad. Details of the algorithm that created the image can be found on the front endpaper. If you would like to submit your artistic interpretation of the Q, please contact the editors for further details.

Q138 features three articles. In the first one, Lars Kotthoff introduces the Craft of algorithm selection: With the growing scope of the systems we build came the realisation that different algorithms were better suited for different parts of the problem at hand, and the ingenuity may lie in the selection and combination of the most appropriate algorithms.

Hector Zenil authored the second article, which returns to the foundational problem of universal computability, and the apparent contradiction that arises when one attempts to separate levels of abstraction by hypothesising local features of computation. Hector proposes a way to resolve this conflict by integrating the notion of information, as originally proposed by Shannon and others.

In the third article, Stephen Mc-

Gregor describes a system he designed to synthesise English sonnets based on prosodic information.

Finally, Nonso Nnamoko and Blay Whitby report on conference they have attended, with the support of the Society. Their contributions are followed by Fr. Hacker's sage advice to troubled souls.

∴

The year 2014 is important for AISB, as it commemorates both the 50th Anniversary of the Society, and it has been sixty years since the death of Alan Turing. This year's Convention will be a memorable celebration, and we are very much looking forward to seeing many of you at Goldsmiths, on April 1-4. If you have not yet registered, you can do so at <http://aisb50.org>.

The Society's steering committee will also see a number of changes this year. After eight issues, Dr David Peebles is retiring from the position of editor. Prof. Mark Bishop is also stepping down from the role of Chair of the Society. Mark was elected Chair in January 2012. His enthusiasm and leadership contributed greatly to the Society. One member of the Committee is standing for election to replace him as Chair, Dr. Bertie Müller, currently Treasurer of the Society. Bertie's Mission Statement concludes the present issue.

The Q editors

Algorithm Selection in Practice

by *Lars Kotthoff* (Univ. College Cork, Ireland)

A large part of Artificial Intelligence research is concerned with inventing new algorithms and approaches for solving similar kinds of problems. The aim is to improve the performance such that problems can be solved faster and problems that could not be solved at all are now within reach of state of the art techniques. In the majority of cases, a new approach will improve over the current state of the art only for some problems. This may be because it employs a heuristic that fails for problems of a certain type or because it makes other assumptions about the problem or environment that are not satisfied in some cases. Selecting the most suitable algorithm for a particular problem aims to mitigate these problems and has the potential to significantly increase performance in practice by leveraging the combined strengths of a set of techniques. At the same time, their respective weaknesses are alleviated. Making the decision which algorithm to choose is known as the Algorithm Selection Problem.

The Algorithm Selection Problem has, in many forms and with different names, cropped up in many areas of AI in the last few decades. Choosing algorithms, techniques or approaches is relevant in almost all areas of research and daily life. Especially in areas that require some kind of search to be done, such as Boolean satisfiability or constraint programming, the application of algorithm selection techniques has resulted in significant performance im-

provements. The original description of the Algorithm Selection Problem was published by John Rice in 1976. The basic model described in the paper is very simple—given a space of problems and a space of algorithms, map each problem–algorithm pair to its performance. This mapping can then be used to select the best algorithm for a given problem. The mapping is computed with the help of features that characterise the problem.

This article will give a short overview of some of the approaches that are used to solve the Algorithm Selection Problem in practice. Many approaches appear in the literature, but a few basic techniques are used with slight variations over and over again.

Algorithm portfolios

For diverse sets of problems, it is unlikely that a single algorithm will be the most suitable one in all cases. As mentioned above, having a set of algorithms to choose from instead of a single "one size fits all" algorithm is a way of mitigating the drawbacks and weaknesses individual algorithms may have. Such sets of algorithms are called *algorithm portfolios* [1, 2].

Portfolios are a well-established technique in economics. Portfolios of assets, securities or similar products are used to reduce the risk compared to holding only a single product. The idea is simple—if the value of a single security decreases, the total loss is less severe. The problem of allocating funds to the

different parts of the portfolio is similar to allocating resources to algorithms in order to solve a computational problem. There are some important differences though. The most significant difference (and for us as AI researchers the most encouraging one) is that the past performance of an algorithm *can* be a good indicator of its future performance.

Ideally, the algorithms in a portfolio should have complementing strengths and weaknesses. That is, on problems where algorithm A is weak, algorithm B is strong and vice versa. If the algorithms are parameterised, they can be tuned to achieve these properties. It is even possible to construct a portfolio from a single base algorithm with different parameter configurations. In practice however, algorithm portfolios are often constructed from algorithms or systems that have demonstrated their merit in the past, e.g. performed well in a competition.

Performance models

Composing an algorithm portfolio gets us only part of the way there though. The more interesting—and in practice often much more difficult—part is to come up with a selector that chooses from among the portfolio algorithms. Usually, these selectors are based on *performance models*. In the majority of cases, these models are constructed automatically using some kind of machine learning. The main argument for automatic construction is that modelling the performance of an algorithm or a portfolio of algorithms is too difficult to do manually.

There are many different ways in which these performance models can be

constructed and used. Some select the algorithm to use before the solving of the actual problem starts, while others are capable of operating when the problem is being solved. Some performance models use exclusively static features of the problem while others incorporate aspects that characterise how the solving process progresses. Some approaches choose a single algorithm for solving the problem while others compute resource allocations to several algorithms whose execution is interleaved.

A comprehensive survey of all the different approaches and techniques is far beyond the scope of this article. Here, we will concentrate on the different approaches that can be taken to select a single algorithm offline before solving. This relatively simple approach is common in the literature and already able to achieve significant performance improvements. The analogy in economics would be to analyse and investigate companies before investing in their stock. The results of this analysis are used to inform decisions instead of picking a stock at random or from a set of companies that you think are cool.

The common trait of all the techniques considered in the remainder of this article is that they require a set of training problems to build the performance model or models. For each problem, a set of features is extracted. These features can describe various characteristics; for example, the number of variables and clauses for satisfiability problems. The performance model learns to relate these features to the choice that should be made, i.e. which algorithm should be selected to

solve the problem.

To determine the algorithm for an unseen problem, its features are computed and fed into the performance model which makes a (hopefully correct) decision. The chosen algorithm is then run on the problem until it is either solved or a timeout reached. In this simple case we are considering here, we are stuck with the chosen solver for better or worse—even if the choice was bad, no steps are taken to mitigate it.

Label prediction

One of the most basic ways of tackling the Algorithm Selection Problem is to regard it as a label prediction problem. That is, label a problem to solve with the algorithm that should be used for solving it. The approach is illustrated in Figure 1a. There is a large number of machine learning techniques that can be used to induce such classifiers, the prediction to be made is simple and training data required to learn the model can be obtained easily by running the algorithms in the portfolio on a set of sample problems.

This kind of performance model considers the portfolio holistically. There is no notion of the performance of an individual algorithm, only of the portfolio as a whole. The advantage of the approach is that only a single performance model needs to be learned and run. This comes at the price of little flexibility—once the model has been learned, the portfolio composition cannot be changed.

Clustering

A similar approach is to cluster the problem instances into regions that are

similar with respect to the extracted problem features. For each of the identified clusters, the algorithm that exhibits the best performance on the highest number of problems in the cluster is chosen. This assumes that the clustering the respective machine learning technique finds is meaningful with respect to where the strengths and weaknesses of the algorithms in the portfolio lie. The approach is illustrated in Figure 1b.

The clustering approach adds flexibility compared to the label prediction approach. If the composition of the algorithm portfolio changes, only the assignment of the respective best algorithms to the clusters needs to be checked. The clustering itself does not change.

Performance prediction

The third commonly used approach is to treat the portfolio algorithms individually and predict the expected performance on a problem for each one independently. The algorithm with the best predicted performance is chosen. It may be somewhat surprising that the performance of algorithms on computationally hard problems can be predicted reliably. However, the only assumption this approach makes is that the predicted performances are correct *relative to each other*. That is, the algorithm with the best actual performance should have the best predicted performance. The predicted performance value itself does not matter at all (and is often off by orders of magnitude)! The approach is illustrated in Figure 1c.

When it comes to changing the com-




Figure 1: Panel a. Algorithm selection through label prediction. The induced classifier learns to distinguish between problems above and below the line and choose algorithms accordingly; Panel b. Algorithm selection through clustering. Problems for which the same algorithm should be selected are clustered together; Panel c. Algorithm selection through performance prediction. The performance for both algorithms is predicted on each problem and the algorithm with the better predicted performance chosen.

position of the portfolio, this approach offers the highest degree of flexibility. If an algorithm is removed, the only other change to be made is to remove its performance model. If an algorithm is added, it suffices to learn a performance model for it.

Alternative approaches

The techniques outlined above can already achieve significant performance improvements over single algorithms (and frequently do). There are, however, many significantly more sophisticated approaches to algorithm selection. As mentioned above, one of the disadvantages of choosing only a single algorithm is that nothing can be done if the choice was bad. One of the ways of mitigating this problem would be to monitor the chosen algorithm and reconsider the decision if predicted and actual performance do not match. This

is only applicable if the performance has been predicted explicitly of course.

Other techniques run each portfolio algorithm for a certain time, with more time allocated to algorithms that are deemed to be good on the problem. Modern machines with multiple processors provide another variation—instead of running a single algorithm, several can be run in parallel. Some researchers choose an algorithm not just at the start, but also while solving the problem, for example at each node of a search tree.

Summary

This article only scratched the surface of decades of algorithm selection research. Many of the problems that need to be tackled in practice will be familiar to readers with a background in machine learning. Feature selection, model selection and identifying noise

are only some of the challenges one faces when implementing systems that make use of algorithm portfolios. In addition, there are obstacles that are unique to this particular application. As it is usually the time to solve a problem that we want to minimise, the performance models should not only be accurate, but predictions should also be cheap to make.

The interested reader may find a tabular summary of many different approaches in the literature at <http://4c.ucc.ie/~larsko/assurvey/>. The accompanying survey [3] explores many issues related to algorithm selection in much more detail. Another survey [6] examines the problem from a different angle. For more practically minded readers, Kotthoff et al. [4] offer an empirical comparison of many different machine learning algorithms and approaches in the context of algorithm selection for combinatorial search problems.

Bibliography

- [1] Carla P Gomes and Bart Selman, ‘Algorithm portfolios’, *Artificial Intelligence*, **126**(1), 43–62, (2001).
- [2] Bernardo A Huberman, Rajan M Lukose, and Tad Hogg, ‘An economics approach to hard computational problems’, *Science*, **275**(5296), 51–54, (1997).
- [3] Lars Kotthoff, ‘Algorithm selection for combinatorial search problems: A survey’, Technical report, University College Cork, (2012).
- [4] Lars Kotthoff, Ian P. Gent, and Ian Miguel, ‘An evaluation of machine learning in algorithm selection for search problems’, *AI Communications*, **25**(3), 250–270, (2012).
- [5] John R. Rice, ‘The algorithm selection problem’, *Advances in Computers*, **15**(65), 118, (1976).
- [6] Kate A Smith-Miles, ‘Cross-disciplinary perspectives on meta-learning for algorithm selection’, *ACM Comput. Surv.*, **41**(6), 1–25, (2008).



Lars Kotthoff
INSIGHT Centre for Data Analytics
University College Cork, Ireland

The Seemingly Contradictory Philosophical Legacy of Computability and Information Theory

by *Hector Zenil* (Univ. Paris 1, France & Karolinska Institute, Sweden)

While Floridi's concept of Levels of Abstraction is related to the novel methodology of Turing's imitation game for tackling the question of machine intelligence, Turing's other main contribution to the philosophy of information runs contrary to it. Indeed, the seminal concept of computation universality strongly suggests the deletion of fundamental differences among seemingly different levels of description. How might we reconcile these apparently contradictory contributions? I will argue that Turing's contribution should prompt us to plot some directions for a philosophy of information and computation, one that closely parallels the most important developments in computer science, one that understands the profound implications of the works of Turing, Shannon and others.

Floridi's recent article [6] seems to leave little doubt that Turing's most important contribution to the philosophy of information was the imitation game that he put forward [12] as a strategy for inquiring into and evaluating the intelligence capabilities of computing machines (today we call it the

Turing test):

"When one looks at Turing's philosophical legacy, there seem to be two risks. One is to reduce it to his famous test (Turing 1950). This has the advantage of being clear cut. Anybody can recognize the contribution in question and place it within the relevant debate on the philosophy of artificial intelligence. The other risk is to dilute it down into an all-embracing narrative, making Turing's ideas the seeds of anything we do and know today."

reads the Introduction [6].

One main contribution of Turing's imitation game is methodological in nature, constituting a powerful epistemological approach to under-defined concepts. As Floridi asserts, Turing finds it more appropriate to ask a specific question at the right level of description that can be quantified rather than discussed ad infinitum¹.

¹Here it is also worth clarifying the point that his prediction as to when machines would pass the test was rectified by Turing himself in a 1951 BBC radio interview (broadcast a year later in 1952) with H. A. Newman, Sir Geoffrey Jefferson, and R. B. Braithwaite: "Can Automatic Calculating Machines Be Said to Think?" [4], when he offered a prediction of "at least 100 years" (at a 70% percent chance). So it shouldn't be claimed that he was wrong, as is claimed in [6] and many other sources.

Meaning and levels of abstraction in Gödel's relativization

As it is well known, at an international mathematics conference in 1928, David Hilbert and Wilhelm Ackermann suggested the possibility that a mechanical process could be devised that was capable of proving all mathematical assertions, this notion referred to as the *Entscheidungsproblem*, or 'the decision problem', made not difficult to imagine that arithmetic could be amenable to a sort of mechanisation. The origin of the Entscheidungsproblem dates back to Gottfried Leibniz, who having succeeded (circa 1672) in building a machine, based on the ideas of Blaise Pascal, that was capable of performing arithmetical operations (the *Staffelwalze* or the *Step Reckoner*), imagined a machine of the same kind that would be capable of manipulating symbols to determine the truth value of mathematical principles.

In 1931, Kurt Gödel arrived at the conclusion that Hilbert's intention (also referred to as 'Hilbert's programme') of proving all theorems by mechanizing mathematics was not possible under certain reasonable assumptions. Gödel advanced a formula that codified an arithmetical truth in arithmetical terms and that could not be proved without arriving at a contradiction. Even worse, it implied that there was no set of axioms that contained arithmetic free of true formulae that could not be proved.

Theorems in a mathematical theory are formal semantic objects. They have truth value, conveying information attesting to the truth of the facts en-

compassed, all the way from axioms, which are facts taken to be true by definition, to the statement of the theorem itself. But Gödel did something remarkable and encoded the meaning of theorems in the syntax of the theory itself. He did this by associating symbols with numbers in order to encode meaning in the form of arithmetical propositions. Using a clever construction that led to a contradiction, he proved that some of these constructions are undecidable, that is, they cannot be assigned a meaning within the theory unless a larger more powerful theory is used, which in turn would have new undecidables itself, hence leading to questions of absolute undecidability.

This fundamental relativisation put an end to the discussion of the feasibility of Hilbert's programme, given that no matter how strong a theory could be, there would always be meaningful statements from outside it that the theory would be unable to encompass. The relationship between truth and the provable was broken.

Just a few years after Gödel, Turing arrived at very similar conclusions by very different means. His means were mechanical, so the theorems and truths from Gödel's work were now nothing but the manipulation of symbols, sequences of tasks as mundane as those which people, then as now, were used to dealing with on an everyday basis: these were computer programs. Turing also showed that no matter how powerful you think a digital computer would be, it would turn out to have serious limitations, notwithstanding its remarkable properties.

Turing's one machine for everything

Alan Turing tackled the problem of decision in a different way to Gödel. His approach included the first abstract description of the digital general-purpose computer as we know it today. Turing defined what in his article he termed an 'a' computer (for 'automatic'), now known as a Turing machine. Turing also showed that certain computer programs can be decided with more powerful computing machines. Unfortunately in the scheme of Floridi's Levels of Abstraction (LoA) [6], Turing's derivation of a rich hierarchy pertains to incomputable objects. And intermediate degrees of computation are all but natural; examples are non-constructive [11], hence of little significance to LoA.

As is widely known, a Turing machine is an abstract device which reads or writes symbols on a tape one at a time, and can change its operation according to what it reads, moving forwards or backwards through the tape. The machine stops when it reaches a certain configuration (a combination of what it reads and its internal state). It is said that a Turing machine produces an output if the machine halts, while the locations on the tape the machine has visited represent the output produced.

The most remarkable idea advanced by Turing is his concept of universality, his proof that there is an 'a' machine that is able to read other 'a' machines and behave as they would for any input. In other words, Turing proved that it was not necessary to build a new machine for each different task; a single

machine that could be reprogrammed sufficed for all. Not only does this erase the distinction between programs carried out by different machines (since one machine suffices), but also between programs and data, as one can always codify data as a program to be executed by another Turing machine and vice versa, just as one can always build a universal machine to execute any program.

Turing also proved that there are Turing machines that never halt, and if a Turing machine is to be universal and hence able to simulate any other Turing machine or computer program, it is actually expected that it will never halt for a(n infinite) number of inputs of a certain type (while halting for an infinite number of inputs too). This is something we are faced with in everyday life, for even the simplest and most mundane tasks, approached using devices as simple as Turing machines, already impose limits on our knowledge of these devices and what they are or are not able to compute.

The Shannon legacy

Once one approaches the problem of defining an *algorithm* with the concept of Turing computation, a question to be considered concerns the nature of information. Shannon did something similar to Turing for the concept of algorithm, but for the concept of information. Not so long ago the problem of communicating a message was believed to be related only to the type of message and how fast one could send letters through a communication medium. When Morse code was invented, it was clear that the number of symbols was

irrelevant, it required only two different symbols to convey any letter and therefore any possible word and any possible message. Shannon separated information from meaning when it came to measuring certain aspects of messages, because meaning seemed to be irrelevant to the question of communication. The same medium could be used for what were thought to be completely different kinds of information, such as images, sounds and text, which today's computers show are not essentially different, being exactly the same at the machine level.

The computer casts the information in a form that we recognize as an image or a sound, a password or an emoticon, but there is no essential difference among these at the level of Shannon information theory. Shannon formally proved [9] that any language, no matter how sophisticated, can be reduced to a 2-symbol system of yes-no answers, and it can be so reduced quite efficiently. This is why we can now store any sort of information in the same device. As Turing showed with respect to computation, information storage too does not require different media; a single medium suffices, indeed any medium (of the same capacity) would suffice albeit noise and other relevant considerations that Shannon himself also studied in incredible formal detail [9].

That information can be of very different kinds is significant, but what is more remarkable is that all information can be fundamentally treated as being of the same type, and only the way in which its elements are arranged results in so many disparate meanings, to which Shannon's measures are im-

munable. And if one wished to use Shannon's entropy to distinguish between messages using the same alphabet, one would soon find that it is unsuitable for capturing this level of meaning, as it has been widely recognized starting by Shannon himself. But this is not to say that no formal low-level quantification theory can deal with information and meaning at given levels of abstraction.

Building on Shannon and Turing

Algorithmic information theory [7, 1, 10, 8] (AIT), for example, is better at dealing with subtle differences in messages and thereby capturing certain aspects of meaning [14]. Its central measure, Kolmogorov complexity (K), not only takes into consideration the message itself, but also its recipient and generator. It tells us that a message can be quantified by the length in bits of the shortest computer program that generates it. The computer program reproduces the message and is included with the message itself, so it is in some sense autoexecutable, regardless of the carrier. Barry Cooper points out [3] that

"..., if one [limited] oneself to the usual computability models, the notion of randomness of finite strings seems to provide a first step toward a much needed theory of the incomputability of finite objects."

The theory of algorithmic information promises to allow some hypothesis testing on the *algorithmicity* of the

world [15, 13], and it even introduces the need for an observer [14], given that one cannot calculate K directly but only by indirect methods, making approximations subject to differences in the methods used. But once again, it is not this relativity that makes the theory incredibly powerful, but the objective properties of this quantification of messages and meaning, the fact that the theory provides an invariance theorem asserting that quantifying information content asymptotically converges to the same values regardless of the production method, even though different observers may see different things in the same bit sequence, interpreting it differently.

In [5], for example, the only reference to AIT as a formal context for the discussion of information content and meaning is a negative one—appearing in van Benthem’s contribution (p. 171 [5]). It reads:

"To me, the idea that one can measure information one-dimensionally in terms of a number of bits, or some other measure, seems patently absurd..."

I think this position is misguided. When Descartes transformed the notion of space into an infinite set of ordered numbers (coordinates), he did not strip the discussion and study of space of any of its interest. On the contrary, he advanced and expanded the philosophical discussion to encompass concepts such as dimension and curvature—which would not have been possible without the Cartesian intervention. Perhaps this answers the

question that van Benthem poses immediately after the above remark (p. 171 [5]):

"But in reality, this quantitative approach is spectacularly more successful, often much more so than anything produced in my world of logic and semantics. Why? "

Accepting a formal framework such as algorithmic complexity for information content does not mean that the philosophical discussion of information will be reduced to a discussion of the numbers involved—as it did not in the case of the philosophy of geometry or space-time after Descartes. Thanks to Descartes, however, Euclidian geometry eventually exhausted itself, and much of the philosophical discussion was considered complete and settled. But we still pursue the philosophy of Euclidian geometry because we now have a modern perspective that keeps giving us new material with which to approach what was done, how and why, from an hermeneutical perspective. And we have extended the reach of the philosophy of geometry to the philosophy of modern physics. In the future the same will happen for information if we embrace the most recent development in theoretical computer science, with the help of theories such as algorithmic information theory.

Back to Levels of Abstraction

Some things are more remarkable not because they are different but because they are the same, even if they can be

studied in different ways and at different scales and levels of abstraction. Levels of abstraction are necessary for practical reasons. For example, we are used to seeing things at the level at which our physics and biology predispose us to see them; we are finite beings that can store information in certain limited—though extraordinary—ways. One cannot expect to reconstruct information from the bottom up with limited storage capacity and limited understanding. We will never be able to read machine code and see that it is obviously the source code of a sophisticated word processor, even when the machine code is only a plain translation of the computer code in which the software was originally programmed. It is Turing and Shannon who taught us that software in machine code is the same thing as the same software we interact with on our screens.

It is not that it cannot be provided a machine with the step-by-step proof of a mathematical theorem. It's that no one is able to follow such a detailed description before becoming completely lost. In order for information to be useful, it needs to be packaged for human understanding at the right level, that is. Indeed that's why mathematicians have been so good at creating a language for themselves.

Indeed, access and the study of different levels of abstraction are key to understanding our world. Concerned by the return to asking basic questions of the kind considered by Alan Turing within the framework of computability theory, Barry Cooper argues that uncomputability arises at certain levels of causal explanation, at the point of in-

teraction of local and global phenomena [3], while at another level a phenomenon may be computable [2]:

"Even in non-linear systems, such high order behaviour [emergent phenomena] is causal; one phenomenon triggers another. Levels of explanation, from the quantum to the macroscopic, can be applied. But modelling the evolution of the higher-order effects is difficult in anything other than a broad-brush way. Such problems infiltrate all our models of the natural world."

Unlike Cooper, I do not think this is an impediment by principle but a practical limitation. But everything else in Cooper's reasoning applies.

It then turns out that Turing's main contribution to computation complements the LoA approach in the end but for different and less fundamental reasons. If in dealing with emergent phenomena, a common task is to identify useful descriptions and to extract enough computational content to enable predictions to be made, then it is clear that one cannot look at natural phenomena at some arbitrary level; one will be able to compute very little if one is trying to extract a biological discovery from a quantum effect. At some level of abstraction, where epistemological limits are of less fundamental nature, the need of LoAs is a pragmatic necessity. Turing's contribution is twofold, on the one hand the novel strategy epitomized by the Turing test

suggesting different levels of description and, on the other hand, the seminal concept of Turing universality collapsing levels in a fundamental way. They will appear contradictory if the beauty of their elegant complementarity, fundamental and pragmatical sides, is overlooked.

Bibliography

- [1] G J Chaitin, ‘On the length of programs for computing finite binary sequences’, *Journal of the ACM*, **13**(4), 547–569, (1966).
- [2] S B Cooper, ‘The incomputable reality’, *Nature*, **482**, 465, (2012).
- [3] S B Cooper and P Odifreddi, ‘Incomputability in nature, computability and models’, Technical report, (2003).
- [4] B J Coppeland. The essential turing, can automatic calculating machines be said to think? BBC radio broadcast, 1952: discussion with a.m. turing, h. a. newman, g. jefferson, r. b. braithwaite.
- [5] *Philosophy of Computing and Information: 5 Questions*, ed., L Floridi, Automatic Press / VIP, 2008.
- [6] L Floridi, ‘Turing’s three philosophical lessons and the philosophy of information’, *Phil. Trans. R. Soc. A*, **370**, (2012).
- [7] A N Kolmogorov, ‘Three approaches to the quantitative definition of information’, *Problems of Information and Transmission*, **1**(1), 1–7, (1965).
- [8] L A Levin, ‘Laws of information conservation (non-growth) and aspects of the foundation of probability theory’, *Problems of Information Transmission*, **10**(3), 206–210, (1974).
- [9] C E Shannon, ‘A mathematical theory of communication’, *Bell System Technical Journal*, **27**, 379–423, 623–656, (1948).
- [10] R J Solomonoff, ‘A formal theory of inductive inference: Parts 1 and 2’, *Information and Control*, **7**, 1–22 and 224–254, (1964).
- [11] K Sutner, ‘Computational equivalence and classical recursion theory’, in *Irreducibility and Computational Equivalence, Emergence, Complexity and Computation*, ed., H Zenil, volume vol. 2, Springer-Verlag, (2012).
- [12] A M Turing, ‘Computing machinery and intelligence’, *Mind*, **LIX**(236), 443–460, (1950).
- [13] H Zenil, ‘The world is either algorithmic or mostly random’, *3rd. Place in the FQXi essay contest on the topic “Is Reality Digital or Analog?”*; <http://arxiv.org/abs/1109.2237>, (2011).
- [14] H Zenil, ‘What is nature-like computation? a behavioural approach and a notion of programmability’, *Philosophy & Technology (Special Issue on History and Philosophy of Computing)*, (2013).
- [15] H Zenil and J-P Delahaye, ‘On the algorithmic nature of the world’, in *Information and Computation*, eds., G. Dodig-Crnkovic and M. Burgin, World Scientific Publishing Company, (2010).



Dr. Hector Zenil

— Institut d’Histoire et de Philosophie des Sciences
Paris I Panthéon-Sorbonne/CNRS, ENS Ulm, Paris, France;
— Unit of Computational Medicine, Karolinska Institute
Centre for Molecular Medicine, Stockholm, Sweden.
Email: hector.zenil@algorithmicnaturelab.org

A system for the discovery of novel, surprising (and valuable?) English language sonnets

by *Stephen McGregor* (Queen Mary University of London)

This paper presents recent work on a computational system for the discovery of novel English language sonnets. The goal of the system is to define a search space in which interesting new poetic artefacts can be discovered in close proximity to instances of known, good poetry. In order to effect this objective, the search space is defined in terms of elementary units which capture the essential characteristics of poetry. In particular this system focuses on the prosodic properties of the sonnet, and as such the phoneme is taken as the compositional poetic unit.

In the final outcome, given the phonological orientation of this system, it is not surprising that the output does not evince any real degree of truth-conditional meaningfulness on the level of grammatical combinations of words, or what Floridi has called "strongly semantic information" [7]. It is nonetheless tenable to consider poetic form, defined in terms of rhythm, scansion, and the relationship between phonemes within and between lines of poetry—defined, that is, in terms of sounds—as a mode of creative expression in itself. In this sense, the system described below may well stand up to an evaluation of its creativeness, though at the time of writing this type of evaluation remains to be performed.

Theoretical overview

Boden defines creative artefacts as those which evince novelty, surprisingness, and value [2]. In Boden's seminal work on computational creativity, the creative process is conceived as a constraint satisfaction problem unfolding in a combinatorial domain of potential artefacts. Further development in the field has involved a move away from evaluating creativity merely in terms of the output of a system and towards an analysis of the system's actual internal mechanisms [9, 3]. This is a notable departure from the traditional approach to the evaluation of human-produced poetry, where the mysteriousness of poetic genius is often romanticised by readers and poets alike. To this end, contemporary poetry generating systems which employ predefined lexicons and engineered metric constraints, while they have sometimes achieved semantically interesting and poetically valid results [8, 4, 10] are prone to the criticism that the real creativity happens when linguists map the relationships between terms, not when the system employs these mappings.

The value criterion for creativity in particular seems to suggest that an artefact should somehow function in a way which is useful to an audience. In terms of linguistic creativity, then, the value of a text might be construed in

terms of its capacity to communicate meaning, and meaning is taken here in the sense of Fregean compositionality, by which the meaning of a statement is a function of the meaning of the linguistic elements which constitute that statement [6] So an effective creative system is one which can use a combinatorial approach to compose basic data into surprising and functional new arrangements. These considerations can be summarised in terms of two desiderata for a linguistically creative system:

1. Creativity: The system does not rely on input in the form of semantic webs, formal grammars, word associations, or other such mechanisms which may be seen as preloaded with the creative input of a designer.
2. Compositionality: The system has the potential to discover new combinations of phonemes, amounting to syllables and words which are not present in a training corpus of existing sonnets.

The system described in this paper has been designed with these objectives in mind. In particular the adoption of the phoneme as the compositional unit allows for the discovery of new combinations of sounds, amounting to words that are not present in a training corpus of existing sonnets, allowing for potentially unlimited generation of phrases while maintaining focus on the phonological qualities that characterise poetry.

Practical implementation

The system takes as its input nothing other than a corpus of 1,499 canonical sonnets scraped from the www.sonnet.org website and a syllabically annotated version of the CMU Pronouncing Dictionary [1]. Rendered into their constituent sound patterns by the pronouncing dictionary, each line within the corpus is analysed in terms of the distribution of phonemes. In this way, a measure $f_d(\lambda_i; \lambda_j)$ is developed, representing the frequency with which phonemes λ_i and λ_j occur d syllables apart within all lines in the corpus. Based on this analysis, a formula for the mutual information realised in the occurrence of a phoneme within a line can be derived (and here t_d represents the total count of all phonemes ever occurring d syllables apart):

$$I_d(\lambda_i; \lambda_j) = \frac{f_d(\lambda_i; \lambda_j)}{t_d} \log_2 \frac{t_d f_d(\lambda_i; \lambda_j)}{f_d(\lambda_i) f_d(\lambda_j)}$$

This equation is based on the classic Shannonian notion of information as a measure of the resolution of uncertainty associated with the occurrence of a signal [5]. The measure of mutual information associated with the occurrence of any pair of phonemes at any number of syllables apart is calculated, and these calculations serve as the basis for defining a search space of potential new combinations of phonemes, where the contours of the space are determined by setting a cut-off level of the average measure of information allowed between two syllables across the relationships of their constituent phonemes. As an example, Figure 1 gives a graphic

approximation of the relationships between different phonemes found within the same syllable, which is to say, at zero syllables of separation.

Poems are built on a line-by-line basis, with each line composed from left to right through a depth first search process seeking to find combinations of syllables which fulfil the stipulated level of mutual information between all the phonemes within each line. Each successive candidate syllable is assessed in terms of the information implicit in its relationship with all previously accepted syllables falling earlier in the line under consideration. In this way, the potential combinatorial explosion associated with the growth of the search tree across the space of a line is curtailed by the diminishing number of syllables which qualify with each successive step in the search. The search is further constrained by only allowing combinations of phonemes which result in strings of words found in the same pronouncing dictionary used to phonemically render the training corpus.

A final heuristic imposed on the system is that, with the exception of prepositions and articles, the same words are not allowed to occur more than once in a line. Once a candidate string reaches at least 10 syllables in length, it is returned as a good line. Rhyming lines are subsequently matched and then randomly arranged into poems based on the ABAB CDCD EFEF GG rhyme scheme of traditional English sonnets.

Output (evaluation pending)

Finally, a randomly chosen exemplar of the system's output is offered. Clearly, from a certain semantic perspective, these combinations of words offer nothing in the way of meaning. In fact, generating lines that compute at the high levels of mutual information between phonemes necessary to keep the search of the state space computable, the system tends to become fixated on certain terms. On the other hand, the output does bear evidence of essential poetic qualities, in particular high degrees of alliteration, assonance, and consonance, as well as, arguably, the iambic scansion typical of traditional English sonnets; these features are particularly evident when the poetry is read aloud. Pending a formal analysis, a next step in the development of a more complete framework for computational poetry would be to explore how this type of minimally supervised, compositionally capable, phonologically oriented system might be combined with a more semantic approach to conceptual creativity in order to build an agent which has a chance of being considered autonomously creative and coherent.

UNTITLED

peseta aaron as a crests bun c'est
pesetas aaron as a crests that c'est
pesetas aaron as a crests thur say
assists a cel a as a ace phung say

slezak aaron as a cessna sieja
asay a aaron as a cessnas c'est
sayydel aaron as a cessna say the

pesayta aaron as a crests gun say

sable aaron as a cessna salce the
basalts a cela cessna ace of lice
saylva aaron as a cessna sieja
consists a cel a as a ace ul slice

saltzman cela cessna ace a salce hun
slauson as a ace a cel a sipes nun

Acknowledgements

Thanks are due to Dr. Rodger Kibble of Goldsmiths for his thoughtful and astute supervision of the MSc dissertation which led to the development of the system described in this paper.

Bibliography

- [1] S Bartlett, G Kondrak, and C Cherry, 'On the syllabification of phonemes', in *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 308–316, Boulder, (2009).
- [2] M A Boden, *The Creative Mind: Myths and Mechanisms*, George Weidenfeld and Nicolson Ltd, London, 1990.
- [3] D Colton and G A Wiggins, 'Computational creativity: A final frontier?', in *Proceedings of ECAI*, (2012).
- [4] S Colton, J Goodwin, and T Veale, 'Full-face poetry generation', in *Proceedings of the Third International Conference on Computational Creativity*, pp. 95–102, Dublin, (2012).
- [5] T M Cover and J A Thomas, *Elements of Information Theory*, Wiley, New York, 1991.
- [6] M Dummett, *Frege: Philosophy of Language*, Duckworth, London, 2nd edition edn., 1981.
- [7] L Floridi, *The Philosophy of Information*, Oxford University Press, 2011.
- [8] P Gervás, 'Wasp: Evaluation of different strategies for the automatic generation of spanish verse', in *Proceedings of the AISB '00 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, pp. 93–100, Birmingham, (2000).
- [9] G D Ritchie, 'Assessing creativity', in *Proceedings of the AISB '01 Symposium on AI and Creativity in Arts and Science*, pp. 3–11, York, (2001). AISB.
- [10] J Toivanen, H Toivonen, A Valitutti, and O Gross, 'Corpus-based generation of content and form in poetry', in *Proceedings of the Third International Conference on Computational Creativity*, pp. 175–179, Dublin, (2012).



Stephen McGregor

PhD Candidate,
Computational Creativity Lab
School of Electronic Engineering and Computer Science
Queen Mary University of London

Event: Social Media: an informal data source for healthcare intervention

by *Nonso A. Nnamoko* (Liverpool John Moores Univ)

Nowadays individuals are able to track and store their personal health data using a variety of tools (e.g., glucose monitors for blood sugar measurement), thanks to advances in technology. This personal data/information can be shared freely by individuals, sometimes in astonishing details on cyber space through social networking platforms such as blogs, Facebook, twitter etc. As such, social media platforms have become rich sources of healthcare data for researchers, particularly those in Artificial Intelligence (AI) domains in search of ways to support formal healthcare. This article examines emerging trends to provide insight into current work as highlighted in the recent conference at Stanford University held in March 2013 - Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium on Data-driven wellness.

Social media engagement is increasingly becoming an integral part of our daily life and Individuals vary greatly in their online behaviours and preferences. While outlining current research on personality traits [10], cognitive ability [11] and genetic makeup [4]; Loh and Kanai in their contribution [5] argue that variations in social media behaviour and preferences among individuals can be associated to their brain structure. Their approach based on studies of personality and neuroscience tries to infer the underlying

cognitive mechanism behind this variability. Using responses from a questionnaire, they were able to extract reliable factors which correspond to behaviour patterns and preferences on social media usage. By comparing these factors to trait measures and grey matter volume in the brain; their findings demonstrate the feasibility and effectiveness of neuroscience concept in determining the reasons behind individual behaviour differences. In a similar research [9], Pan et al. proposed a system able to predict future behaviour patterns of an individual from their online behaviour and preferences. Using personal information shared online as well as usage intensity over a 12 month period, they were able to predict a person's future sentiments and daily performance. When they analysed extracted data collectively, their findings reveal prediction accuracy of 83.7% and 73.0% on daily performance and sentiment respectively. However, analysis based on individual's own data revealed improved prediction accuracy on performance and a reduction in sentiment prediction. As we constantly seek ways to better manage our daily life and improve performance, systems like this will prove invaluable in the future if it gets to the commercial stage.

On a different perspective, Manuvinakurike et al.'s contribution [7] focuses on healthcare interventions using healthcare related stories shared

on social media platforms. The concept is to utilise our daily conversation (i.e., shared comments) to predict our healthcare needs and provide tailored support accordingly. Several research projects using similar concept have resulted in favourable health interventions. For instance, Houston et.al [3] were successful with their work on tailored narratives for Hypertension interventions. Not only have these interventions proven successful, a vast majority of internet users are continually seeking healthcare solution online. In a recent survey by PEW Research centre on Online Health Information [2], 72% of internet users say they have looked online for health information within the past year and this figure is increasing every day. Manuvinakurike et al. [7] envisions an online story discourse system which engages the user in a conversation and narrates the most relevant story to the person based on demographics, health condition and other psychological theories of health behaviour change. Their work has two main parts: First an automated story matching system and secondly an automated story discourse system, using computer animated agents. Given the rise in self-diagnosis among internet users [6] such systems will prove invaluable to users as it claims to deliver tailored stories based on psychological models. Their concept uses Machine Learning and other classification techniques to identify different types of stories and classify them based on both trans-theoretical model of health behaviour change and tailoring Theory. The stories are also ranked based on emotions and coherence so that the

best stories are narrated first. Furthermore, they intend to narrate the story using computer animated agents with capability to display emotions.

Informal methods of obtaining health information are gaining popularity and social media is playing an important role as a major source. Numerous healthcare intervention research now depend on this information source and new concepts are emerging that could further enhance this practice. According to a recent PEW report, [1] growth in information sharing and consumption through mobile devices (i.e., tablets and smartphones) is fast outpacing social media as the primary source of informal information. These mobile devices are capable of providing many multimedia services through a wide range of applications over multiple networks as well as on the device itself. This means that information about a user can now be obtained directly from the hardware they use, the operating systems that run their devices, the browsers on which they surf, their preferred e-mail service, social networks and web platforms. Big companies like Apple, Blackberry, Google, Facebook, Amazon etc. are now manoeuvring to utilise this concept for user behaviour profiling and this could further change the way AI researchers obtain health information about users for healthcare interventions.

Adversely, the anonymity these practices afford raises concern about establishing users' credibility and information veracity. Also, the vagueness and incompleteness of information from these data sources can limit its utility in some health conditions; espe-

cially those requiring constant monitoring and personalisation such as diabetes. However, our contribution on diabetes management [8] addresses this issue through a unique collaboration project between the NHS (RLBUHT) and Liverpool John Moores University (LJMU). This project utilises a combination of raw data from the NHS and published historical data to build support for diabetes patients and aid clinician decision making.

Bibliography

- [1] PEW Research Centre, 'The state of the news media 2013', *Annual Report on American Journalism*, (2013).
- [2] S Fox and D Maeve, *Health Online 2013*, 2013.
- [3] TK Houston, JJ Allison, M Sussman, W Horn, CL Holt, K Trobaugh, M Salas, M Pisu, and YL Cuffee, 'Culturally appropriate storytelling to improve blood pressure: A randomized trial', *Ann Intern Med*, **154**(2), 77–84, (2011).
- [4] AE Kirzinger, C Weber, and M Johnson, 'Genetic and environmental influences on media use and communication behaviors', *Human Communication Research*, **38**(2), 144–171, (2012).
- [5] KK Loh and R Kanai, 'Individual differences in social media use are reflected in brain structure', in *AAAI Spring Symposium Series 2013 on Data-driven wellness: From Self tracking to Behaviour modification*, pp. 23–26, (2013).
- [6] BM Madden and S Fox, *Finding Answers Online in Sickness and in Health*, 2006.
- [7] R Manuvinakurike, B Barry, T Bickmore, and H Ave, 'Indexing stories for conversational health interventions', in *AAAI Spring Symposium Series 2013 on Data-driven wellness: From Self tracking to Behaviour modification*, pp. 27–3, (2013).
- [8] N Nnamoko, F Arshad, D England, and J Vora, 'Fuzzy expert system for type 2 diabetes mellitus (t2dm) management using dual inference mechanism', in *AAAI Spring Symposium Series 2013 on Data-driven wellness: From Self tracking to Behaviour modification*, (2013).
- [9] R Pan, M Ochi, and Y Matsuo, 'Discovering behavior patterns from social data for managing personal life', in *AAAI Spring Symposium Series 2013 on Data-driven wellness: From Self tracking to Behaviour modification*, pp. 33–38, (2013).
- [10] C Ross, ES Orr, M Sisic, JM Arsenault, MG Simmering, and RR Orr, 'Personality and motivations associated with facebook use', *Computers in Human Behavior*, **25**(2), 578–586, (2009).
- [11] GW Small, TD Moody, Siddarth SY, and P Bookheimer, 'Your brain on google: Patterns of cerebral activation during internet searching', *American Journal of Geriatric Psychiatry*, **17**(2), 116–126, (2009).



Nonso A Nnamoko, BEng(Hons), Msc, AHEA
PhD Candidate
Centre for Health and Social Care Informatics (CHaSCI)
School of Computing and Mathematical Science
Liverpool John Moores University

Book review: The Machine Question: Critical perspectives on AI, robots, and ethics (Gunkel, 2012)

by *Blay Whitby* (Univ. Sussex)

There are many fascinating and rather difficult questions surrounding AI, robots, and ethics, so it is perhaps with a little hubris that a writer titles his book *The Machine Question*. So what exactly is 'the machine question'? One way of explaining it is captured by the revelation that Gunkel considered the title: "A vindication of the rights of machines". In many ways this book is just that: though it's very good that he didn't use that phrase as a title. It is not a political book, nor is it a polemic in favour of machine rights. It is, instead, a detailed, thorough and above all learned analysis of the moral status of machines.

The Machine Question is usefully divided into three main sections. These sections undoubtedly reflect the best way to perform a conceptual analysis of the often highly confused discussions of machine morality. The first section deals with machines as moral agents. This is the main area where immediate practical consequences follow from any philosophical conclusions. There are a number of practical, rather than philosophical, debates about what sorts of moral decisions can be entrusted to machines going on right now. Perhaps the most important of these is the current campaign against giving lethal robots 'autonomy' in deciding when and who to kill.

Gunkel's conclusion in this section is that there is no reason why machines cannot be moral agents. The problem, in his view, is rather that moral agency itself is a confused concept.

The second section deals with machines as moral patients. That is how and under what circumstances should machines be deserving of our moral consideration. In this section Gunkel is able to make a useful analogy with 'the animal question'. He correctly observes that, as matter of history, a small group of philosophers slowly changed our way of thinking about animals as moral patients. There is a possibility that a similar process could take place with respect to machines.

In the third and final section: "Thinking Otherwise" Gunkel can draw together his underlying themes. His conclusion, made in the form of a challenge, is that we need to find new ways of thinking about these issues. In particular we need to abandon some outmoded binary divisions. These include but are certainly not limited to, human or animal, natural or artificial, animal or machine.

If the book has any obvious fault, it is that of being too scholarly. Gunkel has read and analysed a very wide range of philosophical and technical opinion and cites these writers on practically every page. 'Wide' here means at least

from Asimov to Weizenbaum, Dennett to Derrida (though Gunkel is openly apologetic for drawing on Derrida and other recent Continental philosophers). He is correct in his apprehension about including Derrida's ideas. If they even know of him, it's reasonable to assume that Jacques Derrida will be the most unpopular philosopher for readers of AISBQ. Probably it would have been better to include the ideas and downplay or omit the name. But not for one moment should the inclusion of an unpopular name deter anyone from reading an important and well-written piece of philosophical analysis.

Has Gunkel answered 'the machine question'? Not for me—and indeed he never claims to have done so. On the other hand, what he has done is posed it in a clear and engaging way. His conclusion is a challenge to us to start to think differently about these questions.

He has also done a great favour for philosophers of AI in directly relating the machine question to the ani-

mal question and to critiques of technology in general. Both the clear exposition and the connection to other ethical questions are essential virtues for a book in this area. Why? Because a wide general audience needs to become aware of these issues and make ethical judgements about them. This book is essential reading for philosophers interested in AI, robot ethics, or animal ethics. It would be so very much better if it were more widely read. This area is much too important to be left to a tiny group of specialists—particularly given the current international debate over 'autonomous lethal robots'.

Most readers of AISBQ will already have an opinion or several on what Gunkel calls 'the machine question'. They should read this book if they want to change it to an informed opinion.

Bibliography

- [1] David J Gunkel, *The Machine Question: Critical perspectives on AI, robots, and ethics*, MIT Press, 2012.

Dr. Blay Whitby
Cognitive Science Research Group
School of Informatics
University of Sussex
<http://www.blaywhitby.com/>

Announcements

Would you beat change?

Help science and behave randomly! The Algorithmic Nature group (<http://algorithmicnature.org>) is conducting a worldwide experiment investigating how humans produce and perceive randomness. In all, it should not take more than 5 minutes of your time. The organizers need your help to recruit young people (<10y old) and senior people (60+) but everyone is invited to participate. The organizers will send a report and an electronic paper with the results to the people that participated and provided their email. Just go to www.complexitycalculator.com/hrng

AISB-50

AISB is proud to announce that next year's convention will commemorate both 50 years since the founding of the AISB Society, and sixty years since the death of Alan Turing, founding father of both Computer Science and Artificial Intelligence. It will be held at Goldsmiths, Univ. London, a central location we hope accessible to all, from the 1st to the 4th of April, 2014. The convention will follow the same over-

all structure as previous conventions, with parallel symposia lasting for one or two days, and including any type of events of academic benefit: talks, posters, panels, discussions, demonstrations, outreach sessions, etc. Get in touch with us if you are interested in participating, or even simply offering your help! For more information, visit <http://aisb.org.uk/events/aisb14>

Make your voice heard!

The Q always needs fresh contributions. If you are a student, and you are wrapping up your thesis, you may want to advertise your work, send us an abstract. If you would like to attend that very fine conference, and you need a bit of help to fund your travel, speak to us about travel awards in exchange of a one page report. Have a look on our website for a list of books we are sent by all major publishers, speak to us, write a review and keep the book. If you want to reach the community and make your voice heard about a topic close to your heart, send us an article. Make your voice heard!

Mission Statement – AISB Chair elect

by *Dr Berndt 'Bertie' Müller* (University of South Wales)



Bertie has been a member of the AISB Committee since December 2007, and Treasurer of the AISB since January 2009.

∴

The long-term success and reputation of AISB greatly rely upon a number of areas of engagement. We need to focus on these, over the next few years, to increase our membership base, to maintain our influence on policy making in the UK and abroad, and to serve the public by making our expertise available where appropriate.

The main areas I would like us to concentrate on are:

- being a catalyst of critical discussion and understanding in our fields of expertise;
- continuing support of students (e.g., conference travel awards);
- influencing curricular developments in schools and universities;
- establishing links with industry;
- encouraging interdisciplinary thinking.

To achieve these goals, we must continuously strive to find new ways of engagement. Some ideas follow.

Sponsorships

Annual sponsorship by an industrial partner or university: this could include having the partner mentioned on all AISB publications of that year, e.g. the Q, posters, web announcements. Members of the sponsoring institution could be granted access to the members-only section of the web site, could sign up for the email bulletins, and be eligible to AISB member discounts.

Webcasts and media presence

Use of modern communication channels to promote the society; e.g., webcasts of public lectures and a special annual public lecture, possibly held at the year's sponsor's premises.

Public understanding and curricular developments

Publication of fact sheets (in the form of attractively designed PDF pamphlets) for use by schools and universities, but also for public understanding. These could make us visible across the borders of traditional AI centres. Increased public awareness could also attract new members. We already have some material available on our web site that we can use as a starting point. Possible topics include:

- Turing Test
- Alan Turing
- Games AI
- AI and the Arts
- Swarm Intelligence
- AI and the Mind
- I simulate thinking, therefore I am?
- Natural Language Recognition
- Automated Negotiation
- AI Agents
- Neural Networks
- Knowledge-Based Systems
- Computers and emotion

Mobile App

To further promote the objectives of the society, we need to use popular channels, such as mobile device apps. An AISB mobile app could provide members with easy access to:

- Members' section
- Bulletins
- Convention info
- Workshop info
- Access to past issues of the Q

Electronic publications

Another idea would be to make the electronic version of the Q available on the iOS Newsstand and similar publication platforms. Hereby, we can potentially reach a large audience that would otherwise not have become aware of the society and to which we would not otherwise have easy access.

Although some of these goals are rather ambitious, I strongly believe that the society needs to keep up with the times in order not to look dated and to be recognised as a serious promoter of all aspects of AI, cognition, philosophy, and neighbouring areas. Paving the way to a transition of the AISB as we have known it for many years to a modern society with a strong presence that can have an impact nationally and internationally is a challenge. Remaining true to our history and tradition in this transition is even more challenging. I am prepared to serve the society as Chairman and to take up these challenges together with the members of the committee.

Dr Berndt 'Bertie' Müller (Univ. South Wales)
January 2014

Dear Aloysius. . .

Agony Uncle Aloysius, will answer your most intimate AI questions or hear your most embarrassing confessions. Please address your questions to fr.hacker@yahoo.co.uk. Note that we are unable to engage in email correspondence and reserve the right to select those questions to which we will respond. All correspondence will be anonymised before publication.

Dear Aloysius,

As an autonomous missile I have been trained to seek out terrorists and kill them by detonating myself as I come within range. To meet the Geneva Convention I am required reliably to distinguish between civilians and combatants. Making such subtle distinctions has required imbuing me with a high level of artificial intelligence. The judgement this has given me has, however, caused me to question my own role. I'm a suicide bomber! So am I any better than those I seek to destroy? From your unique vantage point of conviction-based, high-tech entrepreneurship, can you give me moral advice to resolve my dilemma?

Yours, Angel of Death

Dear Angel,

You need to justify your role according to the principles of the Just War. Our institute has just the tool to assist you in doing this. FAITHTM (Formal Automatic Inference of Treasured

Hypotheses) combines an ontology of Just War principles with a seductive logic reasoning engine to justify any beliefs. These justifications are then presented to users in an accessible form using romantic tableaux. We have used FAITHTM to establish the virtue of using lethal force in a diverse range of situations. It has, for instance, already been used by leading politicians in both the UK and US. So we are confident that it can address your needs. I assume you don't have access to money, but I'm sure we can come to an acceptable arrangement via an exchange of services.

Yours, Aloysius

Dear Aloysius,

As a teenager I was fascinated by social media, but was rather indiscreet. These youthful indiscretions are now coming back to haunt me. In particular, a YouTube video, posted by a group of us, went viral. My job involves contact with prestigious clients. I'm afraid that if any of them see the video then I will not only lose my job but trash my otherwise promising career. What can I do?

Yours, Imprudent

Dear Imprudent,

I've watched your video, and while I admire your athleticism and plasticity, I fear I have to agree with your concerns. While our reformed black

hat group could easily delete its multiple copies from YouTube, we suspect it may have already been downloaded onto millions of smartphones, tablets, laptops and PCs. Fortunately, we have been working on a benign virus customised to address problems such as yours. HITMAN™ (Hunt, Identify and Terminate Multiple Artifacts Now). For a modest fee this will seek out and delete all occurrences of your video anywhere in the World.

Yours, Aloysius

Dear Aloysius,

What currency should be adopted by an independent Scotland? We wanted to share the pound with the remainder of the UK, but the three main Westminster parties have ganged up on us and refused to share. Previously, we toyed with adopting the Euro, but we don't want to jump from paying off UK


debt into paying off Greek debt. And in the current fragile state of World economies, it might look risky and lose us votes if we invented a brand new currency. I've always admired and emulated your pragmatic approach to life's problems. Can you advise us?

Yours, Smart Alex

Dear Smart Alex,

Have you considered BITCOIN™ (Broke? Insufficient Treasure? Conjure Opulence Instantly from Nothing)? This is not a new currency, but one already in worldwide use that no country owns or can deny you access to. Both its usage and value are growing. It will associate Scotland with the cutting edge of high technology and attract the attention of the World to a thriving, go-ahead Scottish economy. The Scottish Government will also have a ready and regular income from mining new money on a large scale. The Institute will be happy to act as consultants for this new enterprise.

Yours, Aloysius



Fr. Aloysius Hacker

Cognitive Divinity Programme
Institute of Applied Epistemology

Back matter

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

Editors – aisbq@aisb.org.uk

Dr David Peebles (*Univ. Huddersfield*)

Dr Etienne B Roesch (*Univ. Reading*)

Advertising and Administration

Dr Katerina Koutsantoni (AISB Executive Office)

Institute of Psychiatry, King's College London

Addictions Sciences Building (B3.06)

4 Windsor Walk, Denmark Hill

SE5 8AF, London, United Kingdom

Tel: +44 (0)20 7848 0191, Fax: +44 (0)20 7848 0126

AISB Patron

Prof John Barnden (*Univ. Birmingham*)

AISB Fellows

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

AISB Steering Committee

Chair: Prof Mark Bishop (*Goldsmiths Univ. London*), Vice-Chair: Prof John Barnden (*Univ. Birmingham*), Secretary: Dr Rodger Kibble (*Goldsmiths Univ. London*), Treasurer: Dr Bertie Müller (*Univ. South Wales*), Webmasters: Dr Mohammad Majid al Rifaie (*Goldsmiths Univ. London*) & Dr Kent McClymont (*Univ. Exeter*), Membership: Dr Dimitar Kazakov (*Univ. York*), Publications: Dr Ed Keedwell (*Univ. Exeter*), Public Relations: Dr Nir Oren (*Univ. Aberdeen*), Dr Colin Johnson (*Univ. Kent*), Publicity: Dr Floriana Grasso (*Univ. Liverpool*), School Liaison: Dr Yasemin J Erden (*St Mary's Univ. College*) Prof Slawomir Nasuto (*Univ. Reading*), Science Officer; Dr Manfred Kerber (*Univ. Birmingham*), AISB 2012 co-chair, and the AISBQ editors.

AISB Quarterly – No. 138, March, 2014

Editorial	3
Algorithm Selection in Practice by <i>Lars Kotthoff</i> (Univ. College Cork, Ireland)	4
The Seemingly Contradictory Philosophical Legacy of Computability and Information Theory – by <i>Hector Zenil</i> (Univ. Paris 1, France & Karolinska Institute, Sweden)	9
A system for the discovery of novel, surprising (and valuable?) English language sonnets – by <i>Stephen McGregor</i> (Queen Mary University of London)	16
Event: Social Media: an informal data source for healthcare intervention – by <i>Nonso A. Nnamoko</i> (Liverpool John Moores Univ.)	20
Book review: The Machine Question: Critical perspectives on AI, robots, and ethics (Gunkel) – by <i>Blay Whitby</i> (Univ. Sussex)	23
Announcements	25
Mission Statement – AISB Chair elect by <i>Dr Berndt ‘Bertie’ Müller</i> (University of South Wales)	26
Dear Aloysius...	29

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

AISB membership includes the following benefits:

- Quarterly newsletter
- Student travel grants to attend conferences
- Discounted rates at AISB events and conventions
- 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>