SIGEVOlution

newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation

Volume 5
Issue 2

in this issue

Editorial

Welcome back my friends to SIGEVOlution! We disappeared for a while, but we are now back on the saddle again. GECCO-2011 ended less than two weeks ago, but it feels such a long time since we were at the opening reception and at the Guinness Storehouse. I want to take this opportunity to thank once more Michael O'Neill, Anthony Brabazon, Irene Ward, and the UCD team for their hard work. They have been the best local team a general chair could dream of! We all owe them an immense thank you!

This issue continues the fifth volume of SIGEVOlution. Another issue should be out soon after this one to clear the backlog. In the first paper, Margaret Boden discusses the reasons both for believing and for doubting that evolutionary art could be wholly free from personal signatures. In the second paper, Gerard Howard and his colleagues present a spiking neuro-evolutionary system which implements memristors as neuromodulatory connections. The paper was one of the best paper nominees at GECCO-2011 and it is reprinted here with the permission of ACM (and the authors of course!). The issue ends with the usual calendar of forthcoming events.

The cover is a shot by Dave Fagan, the UCD/GECCO-2011 official photographer; more photos are available on the GECCO-2011 local arrangements blog. We need your help to collect all the photos and the videos of GECCO-2011. So, please send all your GECCO-2011 shots to geccopictures@gmail.com or to me. The best photo and the best video of GECCO-2011 will be awarded with an Amazon gift card!

This new issue comes to you thanks to Margaret A. Boden, Gerard Howard, Ella Gale, Larry Bull, Ben de Lacy Costello, Andy Adamatzky, Cristiana Bolchini, and board members Dave Davis and Martin Pelikan. And once again, I thank the UCD team, Michael O'Neill, Anthony Brabazon, Irene Ward, Eoin Murphy, Miguel Nicolau, Alex Agapitos, John Mark Swafford, Nguyen Quang Uy, Wei Cui, Michael Fenton, Jonathan Byrne, Clíodhna Tuite, Erik Hemberg, and Kerrie Sheehan.

Pier Luca
July 27, 2011

SIGEVOlution Volume 5, Issue 2

Newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation.

Contributors to this Issue

Margaret Boden
Gerard Howard
Ella Gale
Larry Bull
Ben de Lacy Costello
Andy Adamatzky

Contents

 ISSN: 1931-8499

Anonymity and Evolutionary Art

Margaret Boden, University of Sussex, Falmer, Brighton, BN1 9QJ, UK, M.A.Boden@sussex.ac.uk

Human artists typically have a personal signature, by which their individual authorship can be recognized. Modernist artists tried to avoid such idiosyncracies, focussing on abstract structure instead–and welcomed computers, accordingly. But even those computer artists who have deliberately tried to lose their signature have not managed to do so. Perhaps evolutionary methods might help? Reasons are discussed both for believing and for doubting that evolutionary art could be wholly free from personal signatures.

1 The Quest for Anonymity in Art

Artworks are typically attributable, by art historians and connoisseurs, to a particular person. Indeed, Romantic views of art value the fact that the individual artist's 'personal signature' enables one to recognize the authorship of the work. This personal signature is not literally a signature. Rather, it is a set of subtle features of the work, of which the actual artist may not even be consciously aware [4].

Modernist artists, reacting against Romanticism, down-played the role of the individual person in art. They stressed formal (often minimalist) structure, not perceptible idiosyncracies. Typically, the art-object was no longer celebrated as a unique artefact, nor the human artist as an individual person.

This attitude was epitomized in an influential statement by the modernist Sol LeWitt: "the idea becomes a machine that makes the art, [where] all of the planning and decisions are made beforehand and the execution is a perfunctory affair" [8, pag. 824]. Once the plan has been chosen, LeWitt said, "The artist's will is secondary to the [artmaking] process he initiates from idea to completion" [9, item 7]. Indeed, he produced many 'remote' artworks, where he faxed instructions intended to be followed by anonymous people who, by following these instructions, would make the work using standard off-the-shelf materials such as 2-inch by 2-inch wooden strips. The Romantic ideal, of art as the expression of human individuality, had been abandoned.

2 The Impersonality of Computers

It's not surprising, given the sentiments quoted above, that when computers appeared on the scene many artists with modernist sympathies welcomed them specifically for their impersonal, non-human, nature. (Romantics, by contrast, recoiled from them in horror.)

At base, the reason for the existence of personal signatures lies in factors concerning the economy of information processing in human minds [4]. Computers are only indirectly affected by such factors. And, of course, they are immune to the motor habits of the programmer, and normally cannot develop any motor habits of their own (as we'll see in Section 3, certain sorts of robot may be exceptions to that).

The psychological basis for the personal signature therefore disappears. Or, more accurately, it is pushed into the background. The aims and imagination (and programming skills) of the computer artist will always have idiosyncratic features, which may or may not be reflected in the computer output. But for those mid-century artists who already wished to obscure, or even escape from, their human individuality, it seemed that the very *impersonality* of computers might help.

Today, that is still a very natural assumption. So much so, that three leading computer artists have recently felt the need to reassure newcomers to the genre that *if* they want to set their individual stamp on the computer's behaviour, then they can. As they put it: "As a designer working with generative processes [i.e. computer art/design] one may still wish to leave a recognizable mark on a creation. This may be achieved statically using fixed components with a trademark style.

A more interesting way to achieve this is to ensure either that the organization of the artefact bears the stamp of its designer, or that its behaviour falls within the gamut of work typically produced by the designer. Of course the designer may not be interested in producing a recognizable style, however the utilization of generative techniques does not preclude this option" [10, 6.1]. We'll return to the issue of "the organization of the artefact [bearing] the stamp of its designer" in Section 4.

One of the first artists to welcome computers for their very impersonality was the young Paul Brown. Visiting the "Cybernetic Serendipity" exhibition in London in 1969, he was inspired by the hope that this new methodology would enable him to do something he was already trying to do: namely, to lose his personal signature. Now, some forty years later, he is an internationally known computer artist. But his artworks are still recognizable, to those familiar with his oeuvre, as Brown's. Even his very earliest pieces [5] have an evident visual kinship with his recent/current work. In other words, it turned out that losing his individual artistic stamp, as his modernist sympathies inclined him to do, was easier said than done.

One reason is that Brown himself, after forty years as a professional artist, still cannot say just what his personal signature is (i.e. just what needs to be avoided). In general, recognising a particular artist's signature and describing it explicitly are two very different things [4, sectn. iii]. Whatever it is in Brown's case, it certainly is not a matter of a specific mark (such as a particular form of ear-lobe) recurring in his work. It is more a matter of an overall stylistic 'feel' that he cannot pin down in words.

He had hoped as a young man that the clarity with which art-making has to be defined if computers are involved might help him both to identify his signature and (by changing the generative rules as a result) to lose it. Reasonable enough hopes, one might think. But no: his computer-generated work still betrays its human author's individual hand. And this, even though he has deliberately aimed for aesthetic anonymity.

It appears, then, that if someone wishes to use computers so as to lose their personal signature, deliberate self-effacement in the hands-on practice of one's art is not the way to do it. Can some other way of achieving self-effacement be found?

3 Could Anonymity be Evolved?

Recently, Brown has been using computers in a new way in trying to achieve his long-standing artistic goal. An interdisciplinary team, with Brown as a leading member, has tried to evolve line-drawing robots whose products are of some aesthetic interest (no more than that!), but which do not carry the telltale traces of a work by Brown himself.

In evolutionary art in general, the selection at each generation can be done interactively, by a human being making the comparisons, or automatically by the program itself. In this particular case, interactive selection is best avoided, because it is likely to carry the personal mark of the human artist. Even automatic selection, however, requires that a 'fitness function' be defined, which the program can use to make its selections. (The fitness function itself may evolve, again either interactively or automatically.) As we'll see in Section 4, this fact is the Achilles' heel of Brown's current research.

The first obvious question to ask about this project–which is named *Drawbots*– is "Why evolve line-drawing gizmos, as opposed to simply designing (programming/building) them?" The second is "Why use robots, as opposed to computer graphics (i.e. programs for drawing images on paper or virtual images in cyberspace)?"

The answer to the first question is that if the line-drawing computer system has been evolved then, thanks to the many random mutations that will have taken place, it has *not* been prespecified in detail by the artist-programmer. Accordingly, there may (sic) be a chance of avoiding that individual's personal signature. Whether that "may" can, in practice or even in principle, be replaced by a "will" is the key point at issue here.

As for the second question, the answer is that a robot, being a material object functioning in the physical world, can be affected not only by its program and/or internal design but also by unexpected–and perhaps serendipitous–events in the physical environment. Again, this offers a means by which the programmer's personal signature may be bypassed, or anyway diluted. (An early example of this sort of thing occurred in the 1970s, when the moving 'legs' of a kinetic sculpture–alias a robot–happened to scratch the wooden floor of London's Royal Academy. Although the RA was doubtless incensed, the sculptor, Darrell Viner, was intrigued. He was so "fascinated by the structure of the repetitive scratches and their relationship to cross-hatching" that he went on to make artworks produced by comparable, though simulated, means–[6, 283].

The "serendipity" in the physical events involved can even include cases where a radically new feature appears in the robot's behaviour. In a previous experiment done by a member of the *Drawbots* team, a population of robots evolved a new sensory capacity–not merely an improved sensory capacity–as a result of contingent, and previously unremark-ed, facts about the physical environment [1]. That suggests the possibility that a fundamentally transformative change in the *Drawbots'* drawing-style might occur. If so, then presumably the new style would not bear Brown's individual mark, even if the previous style had done so.

The Drawbots themselves are small wheeled vehicles carrying a retractable pen. And the task in the team's minds is line-drawing. By that is meant not drawing pictures that represent real things (as both stick-men and Renaissance cartoons do), nor even drawing geometrical designs, but simply drawing *lines* . . . which can curve, cross, stop, and approach each other in myriad ways–and which may sometimes change in thickness too. Brown's hope is that robots can be evolved which will draw aesthetically acceptable lines *that do not exhibit his personal signature*. In other words, the fitness function/s to be followed by the robot should guarantee aesthetic acceptability but should not be so 'rich' as to express his personal style.

In principle, that would not preclude there being a telltale identifier, or quasi-signature (one can hardly say a "personal" signature), *produced by an evolved robot itself*. This would be a pattern that distinguishes its drawings from those of its siblings and close cousins. The evolution of such patterns is in principle possible because new performance details will follow from random mutations, and these details can be perpetuated provided that they do not compromise fitness.

Such details could include drawn patterns or line-features discriminated by the gizmo's visual sensors. Indeed, a robot might even develop particular motor habits, driven by motor circuits conserved in its 'brain' (see Section 2). Suppose that a sudden movement, caused by a recently mutated motor circuit, led to a mark that was then selected (along with the rest of the drawing) by Brown. This might lead the motor circuit to endure, forming the basis of a future motor habit. That habit could be involved either in many different stylistic choices, or only in one (think of an overall stylistic 'feel' and of tell-tale ear-lobes, respectively). In short, the general style that is selected via the fitness function could allow for idiosyncratic expression (alias signatures) by different robots within the same generation or lineage.

If the fitness function were to include measures of computational economy, the different robots might even develop quasi-signatures for much the same (psychological) reasons that human beings do. However, it is hardly likely that such patterns would arise as a matter of course, as they do in the work of human artists. For the root of the personal signature, as remarked above (see also [4, sectn. iii]), is the need for economy in information processing within a highly complex system–a criterion that does not apply in robots as simple as those being considered here.

Whether it is actually possible for the drawbots to lose the stamp of Brown's individual artistry depends on a number of things. One is the extent to which Brown, or anyone else, can say just what his personal signature consists in. If he knew that, he would be in a much better position to try to avoid it. However, as explained in Section 2, he does not.

Possibly, the *Drawbots* research may help him towards a better–if still incomplete–understanding of this. For in examining the various drawings made by the drawbots, he will have to ask himself two questions: *Is it aesthetically acceptable?* and *Is it evidently a 'Brown'?* In answering that second question over and over again, as the drawing style mutates across the generations, and in posing it to colleagues with an appropriately practised critical eye, he may achieve a more explicit understanding of just what his own style is. (Then again, he may not.) But that could happen without his ever answering *No* to the second question. In that case, he still would not have 'lost' his signature, despite understanding it more deeply. Whether the increased understanding would enable him to dilute it, if not to shed it, in his (non-evolutionary) future work is an interesting question.

Another factor that will affect the likelihood of success in the project is the extent to which aesthetic acceptability can rest on relatively primitive visual features. “Primitive”, here, means both *simple* and *naturally salient*. For example, shininess (of satin, silver, polished ivory, lurex, chromium . . .) is relatively simple to discriminate, and naturally salient too. That’s so for good evolutionary reasons, involving the fitness-enhancing nature of reflective expanses of water [3, 8.iv.a]. In other words, it’s no accident that shininess is aesthetically appealing to a very wide range of individuals and cultures. Are there any features of line-drawings such as those the drawbots could produce which are naturally attractive (and easily discriminable) in a comparable way?

For example, if the drawbots were able to change pens, might they evolve a preference for the shiny lines left by a silver-ink pen? They could do so, if their visual apparatus could discriminate shininess. To be sure, the robotics team would have to build reflectance into the fitness function: no robot ’naturally’ prefers it. But reflectance is such an easily discriminable property, and so near-universally liked by human beings, that the team could not be accused of cheating were they to do that. (Some cultural groups positively avoid shininess, regarding it as vulgar; but that is irrelevant here, since this discriminatory attitude has developed *precisely because* the liking for shininess is so very common.) Nor would putting silveriness into the fitness function result in drawings that display Brown’s personal signature, for that (whatever it is) is not a matter of shininess.

It’s easy to see that Brown’s authorial mark does not involve shininess. What it *does* involve is less clear. Suppose it were to turn out that all the perceptible features favoured (via the fitness function) by ’aesthetically competent’ drawbots were relatively high-level and/or complex, with no ’natural’ attractiveness for human beings in general. In that case, their drawings would probably be more specific to Brown’s personal style. His project would have failed. However, “success” and “failure” here admit of several levels. In the language used above, Brown’s signature may become more or less *diluted*, even if it cannot be entirely lost.

Among the naturally discriminable features that are already being considered by the *Drawbots* team are holes, line-crossings, and fractals (of varying complexity or depth). But why should one expect any of these things to be ’naturally’ attractive?

Well, consider fractals, for instance. These are ubiquitious in Nature, both in living things and in environmental features such as rocks and coastlines. According to the ’biophilia’ hypothesis [16], *Homo sapiens* has evolved to respond favourably not only to conspecifics and other aspects of our original ecological niche (the African Savannah) but also to living things and natural environments in general. If that’s so, then fractals might well have some natural attraction for us. That’s merely an argument for plausibility. But there is also some evidence that fractals of a certain kind are spontaneously favoured in art as in nature–and even, as William Congreve said of music, that they can soothe the savage breast. Richard Taylor claimed, in the late-1990s, that Jackson Pollock’s canvasses, far from being random splashes of paint, have specific fractal properties to which most viewers respond in a positive way, and by which his paintings can be distinguished from fakes [13, 14]. Specifically, people prefer those Pollock paintings which have a fractal dimension of 1.5 (his later paintings reach 1.8+). By comparison, people asked to choose between natural images (or between simulated coastlines) prefer a fractal dimension of 1.3. Taylor’s claim aroused huge interest (e.g. Spehar et al. 2003), and was later followed by experiments showing that viewing Pollock’s images can actually reduce stress [15].

Taylor’s early remarks about how to discriminate genuine Pollocks from fakes, have recently been challenged [7]. One aspect of that challenge is especially intriguing here: Katherine Jones-Smith reported that a careless doodle done by her showed the same fractal properties as those found in Pollock’s work. She didn’t ask whether the doodle had any aesthetic value. To the contrary, she implied that, being a thoughtless scribble, it did not. But if she had asked people whether they “liked” it, or whether they preferred it to some other mark (maybe one produced accidentally), she might have found that people ascribed some–albeit small–degree of aesthetic merit to it. If that were so, it suggests that a suitably fractal-favouring drawbot might make aesthetically acceptable (’natural’) drawings that don’t show anyone’s individual mark: not hers, not Pollock’s, and not Brown’s either.

4 The Likelihood of Success–and What it Would Mean

The discussion in Section 3 suggested that it is *in principle* possible for Brown's personal signature to be lost by evolved robots (even though it is also possible for those robots to develop individual 'signatures' of their own). But what of the likelihood of this happening *in practice?* Are there any specific reasons (beyond those mentioned in Section 3) to suspect that the *Drawbots* project will succeed, or fail? And if it succeeds, would it follow that the creativity exhibited in the drawings of the newly-evolved drawbots must be attributed to the drawbots themselves, rather than to Brown? *'No signature, no creative authorship'*, perhaps?

As remarked above, the Achilles' heel of the project lies in the fitness function. This is true in two related senses, one philosophical and one psychological.

First, if it is Brown who is continually deciding on the fitness function as the research proceeds then perhaps it is *his* aesthetic judgement, and also *his* artistic creativity, which is really responsible for the final drawings? (For shorthand purposes, let's ignore the creative role of the other human beings on the team.) Many philosophers would say that there is no "perhaps" about it, that *of course* Brown's creativity lies behind whatever aesthetic interest the drawbots' drawings happen to have. For they believe that it is in principle absurd to ascribe creativity, or aesthetic judgment, to any computer system–no matter how superficially impressive its performance may be.

Their belief typically rests on assumptions about one or more of four highly controversial issues, including intentionality and consciousness [2, ch.11]. Accordingly, it can be challenged–though not definitively refuted. However, even if one were happy to reject their claim as a general philosophical position, that would not settle the question at issue here. For in the specific case of the *Drawbots* research, the largely human source of the fitness function is a distinct embarrassment for anyone wanting to grant all the creative credit to the computer.

This embarrassment would persist whether or not the project succeeded in its own terms–that is, irrespective of whether Brown's signature had been lost. For if the final fitness function were to exploit only what in Section 3 were called "primitive" aesthetic properties, so that Brown as an individual artist had become invisible in the final-stage drawings, it would still be true that the aesthetic decisions involved in developing the fitness function were *such as are naturally made by human beings*. Brown's hand (judgment) would still be there–but functioning as the hand of a generic human being, not of a particular individual. (In other words, the fitness function would describe the general style, without imposing any detailed 'authorial' implementation.)

That argument would apply even if the robots' drawing style had shown a truly fundamental change: a new style (presumably, a 'non-Brown' style), as opposed to an improved style. We saw in Section 3 that the physical 'embodiment' of the drawbots makes it in principle possible for such serendipitous change to occur. By definition, the stylistic change would have been caused by some unconsidered and/or contingent feature of the robots' physical environment.

So Brown couldn't be credited with initiating it. But he could, perhaps, be credited with 'causing' it, since the incipient change will be maintained (and perhaps developed) only if it is approved/selected by his personal decision or by the fitness function already evolved under his direction. In such a case, Brown might be regarded as the creative spirit behind the final drawings *even though* he never foresaw them, and *even though* they are free of his personal mark.

What of the psychological question? Are there any psychological reasons to expect that Brown *will not* be able to decide on a fitness function that entirely avoids his personal signature?

One psychological consideration that is important in aesthetic judgments (see [4]: sectn. iii) is relevance–considered in terms of computational closeness and/or efficiency (Sperber and Wilson 1986). This issue is less obviously crucial here than it would be if Brown were trying to evolve robots capable of realistic representational drawings. If the drawbots were intended to draw human faces, for instance, they had better include depictions of eyes, mouth, and even (the relatively less relevant) ear-lobes.

And they had better *not* add horns, or wings. But if a tinge of surrealism were to be favoured (by Brown), then a horn-like protuberance appearing in generation 1,000 might be selected and 'shaped' so that recognizable devilish/goat-like horns were visible at generation 9,000. The same might occur if Brown felt that familiar myths about the Devil were relevant to the 'topic' of the drawings. In either case, Brown's own judgements about relevance would be reflected in the robots' behaviour, and–to the extent that these are idiosyncratic–so would his personal mark.

In fact, Brown has always been an abstract artist, so is not aiming to evolve 'representational' robots. Even so, issues of relevance–or rather, issues of what he deems to be relevant–may arise.

Aesthetic acceptability depends in part on intelligibility. To be sure, intelligibility may be more or less easy to achieve in differing artistic styles. But utter chaos will satisfy nobody. In other words, one factor underlying judgments of aesthetic acceptability is the computational effort that is involved in comprehension. A 'messy' line-drawing (or doodle), for instance, may be unacceptable largely because its components *do not* appear to be mutually relevant. That is, they do not appear to be 'coherent', or to 'make sense'. (Perhaps there are no closed curves, suggesting bounded physical objects? And/or perhaps there are no T-junctions where one line stops as it meets another, suggesting occlusion of a line/edge by some other physical thing?)

These judgments are not usually conscious–and it may not be possible to make them fully conscious. It follows that it may not be possible for Brown to avoid them deliberately.

A closely related issue is the extent to which Brown can banish his own preferred schemas from the fitness function. (Compare: evolving robots to draw faces *without* eyes.) If he cannot, because these schemas are so deeply entrenched in his mind and experience, they will inevitably be reflected in the fitness function and therefore in the final drawings.

At that point, we come full circle to the issue discussed in Section 3 in terms of "simplicity" and "naturalness". The more that the features favoured in the fitness function are complex, culture-based, and idiosyncratic to Brown, the less will the final-generation drawbots be free of his personal stamp.

If the Brown signature is preserved, despite all his efforts, that will be because he has found it necessary to build relatively 'rich' criteria into the fitness function. As we've seen, it is still an open question as to how rich the final criteria of aesthetic fitness will need to be. If they are all relatively simple, then Brown's creative inspiration may seem less important. At most, the fact that he is a human being will be relevant, not the fact that he is Paul Brown. (Any idiosyncratic 'signature' visible in the drawings might be attributable to the evolutionary vicissitudes of the robots themselves, as explained above.)

What if, contrary to all his hopes, Brown's personal signature remains still visible to experts (dare we say connoisseurs?) looking at the robots' drawings? In such a case, and *even if* one were willing in principle to grant creativity to some computer systems, it would seem bizarre to attribute creativity to the drawbot. For the concept of the personal signature arose specifically in order to attribute a given work of art to one creative source–normally, one human individual–rather than another [4, sectn. ii]. The signature, in short, points to the person. This was recognized by the computer artists (quoted in Section 2) who spoke of "the organization of the artefact [bearing] the stamp of its designer". Whether that telltale organization were deliberately designed, as they were assuming, or gradually evolved, as in the *Drawbots* project ('failure' here being supposed), it would point to one person: Brown.

References

[1] Bird, J., and Layzell, P. (2002), 'The Evolved Radio and its Implications for Modelling the Evolution of Novel Sensors', *Proceedings of Congress on Evolutionary Computation*, CEC-2002, 1836-1841.

[2] Boden, M. A. (2004). *The Creative Mind: Myths and Mechanisms*, 2nd edn., expanded/revised (London: Routledge).

[3] Boden, M. A. (2006). *Mind as Machine: A History of Cognitive Science* (Oxford: Oxford University Press).

[4] Boden, M. A. (2010) 'Personal Signatures in Art', in M. A. Boden, *Creativity and Art: Three Roads to Surprise* (Oxford: Oxford University Press), 92-124.

[5] Brown, P. (1977), 'The CBI North West Export Award'. First published in *Page Sixty Two–Special Terminate CACHe Issue: Bulletin of the Computer Arts Society, Northern Hemisphere, Autumn 2005:* 12-13.

[6] Brown, P. (2008), ’From Systems Art to Artificial Life: Early Generative Art at the Slade School of Fine Art’, in C. Gere, P. Brown, N. Lambert, and C. Mason (eds.), *White Heat and Cold Logic: British Computer Arts 1960-1980, An Historical and Critical Analysis* (Cambridge, Mass.: MIT Press): 275-289.

[7] Jones-Smith, K., and Mathur, H. (2006), ’Revisiting Pollock’s Drip Paintings’, *Nature*, 444 (Nov. 30):E9-E10 (published online 29 Nov.).

[8] LeWitt, S. (1967), ’Paragraphs on Conceptual Art’, Artforum, 5(10): 79-83. Reprinted in K. Stiles and P. Selz (eds.), *Theories and Documents of Contemporary Art: A Sourcebook of Artists’ Writings* (London: University of California Press), 822-826.

[9] LeWitt, S. (1969), ’Sentences on Conceptual Art’, Art-Language, 1: 11-13. Reprinted in K. Stiles and P. Selz (eds.), Theories and Documents of Contemporary Art: A Sourcebook of Artists’ Writings (London: University of California Press), 826-827.

[10] McCormack, J., Dorin, A., and Innocent, T. (2004). ’Generative Design: A Paradigm for Design Research’, in J. Redmond, D. Durling, and A. de Bono (eds.), Futureground, vol. 1 (Melbourne: Design Research Society). Available as a pdf-file from McCormack’s home page.

[11] Spehar, B., Clifford, C., Newell, B., and Taylor, R. (2003), ’Univeersal Aesthetic of Fractals’, *Computers and Graphics*, 27: 813-820.

[12] Sperber, D., and Wilson, D. (1986). *Relevance: Communication and Cognition* (Oxford: Blackwell).

[13] Taylor, R., Micolich, A. P., and Jonas, D. (1999a), ’Fractal Analysis of Pollock’s Dripped Paintings’, *Nature*, 399: 422 (one page only).

[14] Taylor, R., Micolich, A. P., and Jonas, D. (1999b), ’Fractal Expressionism’, *Physics World*, October.

[15] Taylor, R. P., Spehar, B., Wise, J. A., Clifford, C. W. G., Newell, B. R., Hagerhall, C. M., Purcell, T., and Martin, T. P. (2005), ’Perceptual and Physiological Responses to the Visual Complexity of Pollock’s Dripped Fractal Patterns’, *Journal of Non-Linear Dynamics, Psychology and Life Sciences*, 9: 89-114.

[16] Wilson, E. O. (1984), *Biophilia* (London: Harvard University Press).

About the author

Margaret A. Boden is Research Professor of Cognitive Science at the University of Sussex. She is a member of the Academia Europaea, and a Fellow of the British Academy and of the American Association for Artificial Intelligence. In 2002 she was awarded an OBE “for services to cognitive science.” Her writing has been translated into 20 foreign languages. Her latest books are Mind as Machine: A History of Cognitive Science, 2 vols. (2006) and Creativity and Art: Three Roads to Surprise (O.U.P. 2010).

Email: M.A.Boden@sussex.ac.uk

Evolving Spiking Networks with Variable Memristors

Gerard Howard, Ella Gale, Larry Bull, Ben de Lacy Costello, and Andy Adamatzky
University of the West of England, Coldharbour lane, Bristol, UK
david4.howard@uwe.ac.uk, ella.gale@uwe.ac.uk, larry.bull@uwe.ac.uk,
ben.delacycostello@uwe.ac.uk, andrew.adamatzky@uwe.ac.uk

This paper presents a spiking neuro-evolutionary system which implements memristors as neuromodulatory connections, i.e. whose weights can vary during a trial. The evolutionary design process exploits parameter self-adaptation and a constructionist approach, allowing the number of neurons, connection weights, and inter-neural connectivity pattern to be evolved for each network. Additionally, each memristor has its own conductance profile, which alters the neuromodulatory behaviour of the memristor and may be altered during the application of the GA. We demonstrate that this approach allows the evolutionary process to discover beneficial memristive behaviours at specific points in the networks. We evaluate our approach against two phenomenological real-world memristive implementations, a theoretical "linear memristor", and a system containing standard connections only. Performance is evaluated on a simulated robotic navigation task.

1 Introduction

The memory-resistor, first implemented as a "memistor" by Widrow [39], then theoretically characterized by Chua [3] and renamed a "memristor", has recently enjoyed a resurgence of interest from the research community after being manufactured *in silico* by HP labs [36]. A memristor is a fundamental passive two-terminal device whose state (memristance) is both nonvolative and dependent on past activity. Nonvolative memory [13] is perfect for low-power storage, and the device's dynamic internal state facilitates information processing.

These properties make the memristor an ideal candidate for use in nanoscale neural architectures [22], where the memristor can function as a synapse between - for example - Complementary Metal-Oxide Semiconductor (CMOS) neurons [27].

In this paper, we introduce the notion of a *variable memristor* e.g. a memristor whose conductance profile can vary as a result of the evolutionary process. As the conductance profile of the memristor is responsible for its behaviour, variable memristors can potentially impart a variety of adaptive behaviours to the networks. We analyse the computational properties of variable memristors when cast as synaptic connections in evolutionary Spiking Neural Networks (SNNs [10]).

Neural architectures require some form of learning mechanism in order to harness their computational power. A typical approach involves utilising a form of Hebbian learning [12] to realise Spike Time Dependant Plastictiy (STDP) [18], so that memristors between a presynaptic and postsynaptic neuron can alter their efficiacy dependant on the spike timings of those neurons. Here, the memristive element of the network allows the weight of the connections to vary during a trial and provides a neuromodulatory learning architecture which is shown to be beneficial to the evolutionary design process.

We couple this neuromodulatory process with a constructive model of neuro-evolution, whereby the evolutionary process can add or remove both neurons and connections (which may be memristors) during the application of the Genetic Algorithm (GA) [14].

Our hypothesis is that the additional degrees of functional freedom afforded to the variable memristors can be benficially harnessed by the evolutionary process. To test this hypothesis we compare the variable memristor networks to a number of alternates, (i) PEO-PANI networks [8], (ii) HP networks, [36], (iii) idealised"linear memristor" networks, (iv) networks comprised of static (e.g. non-memristive) connections. A simulated robotics navigation task is selected for this purpose. To our knowledge, this is the first approach that allows for the self-adaptation of the characteristic performance of the memristors alongside neuroevolution of both neurons and connection structure.

The remainder of the article is ordered as follows: Section 2 introduces background research, Section 3 introduces the system, Section 4 details SNN implementation, and Section 5 outlines the static memristors. Section 6 details the discovery component and Section 7 details the topology mechanisms used. Following this, Section 8 details the variable memristor implementaiton, Section 9 gives the environment and Section 10 shows the experimental setup and analyses the results of the experiments that were carried out. Section 11 concludes with a discussion and future research directions.

2 Background

2.1 Spiking Networks

Spiking Neural Networks (SNNs) are a relatively recent phenomenological model of neural activity in the brain. In an SNN, a number of neurons are linked via unidirectional, weighted connections that provide a method of intra-network communication. The medium of communication is the action potential, or spike, which is emitted from a neuron and received by all connected neurons that the given neuron is presynaptic to. Each neuron has a measure of excitation, known as "membrane potential". A spike is emitted after an arbitrary neuron surpasses a certain level of excitation within a given window of time. This build-up of membrane potentials and release of postsynaptic current within a network is able to produce dynamic activation patterns through time, providing increased computing power [20] [31] when compared to other network models, such as the Multi Layer Perceptron (MLP) [30]. Two well-known formal SNN implementations are the Leaky Integrate and Fire (LIF) model [10] and the Spike Response Model (SRM) [10].

Neuro-evolution involves the use of evolutionary techniques to alter the topology or weights of neural networks. A survey of various methods for evolving both weights and architectures is presented in [9], similarly the evolution of networks for robotics tasks is covered by Nolfi and Floreano [24].

2.2 Memristors

Memristors (memory-resistors) are the fourth fundamental circuit element, joining the capacitor, inductor and resistor. A memristor can be defined as a resistor whose current resistance value (a) depends on the previous charge that has passed through it (b) is nonvolatile. Formally, a memristor is a passive two-terminal electronic device that is described by the non-linear relation between the device terminal voltage, v, terminal current, i (which is related to the charge q transferred onto the device), and magnetic flux, φ, as (1) shows. Resistance increases or decreases depending on the direction of the current.

$$v = M(q)i \text{ or } i = W(\varphi)v \tag{1}$$

The memristance (M) and memductance (W) properties are both nonlinear functions, defined in (2) as:

$$M(q) = d\varphi(q)/dq \text{ and } W(\varphi) = dq(\varphi)/d\varphi \tag{2}$$

Previous applications of memristors within neural paradigms include HP memristors [36], which have been used in the manufacture of nanoscale neural crossbars [32], and AgSi memristors, which have been shown to function in neural architectures [17].

2.3 Synaptic Plasticity

Hebbian learning [12] is thought to account for adaptation and learning in the brain. Briefly, Hebbian learning states that "neurons that fire together, wire together" - in other words in the event that a presynaptic neuron causes a postsynaptic neuron to fire, the synaptic strength between those two neurons is increased so that such an event is more likely to happen in the future. Such a mechanism allows for self-organising, correlated activation patterns.

Integration of neuroevolution with heterogeneous neuromoduation rules (which are similar to the different memristor types used here) is investigated by Soltoggio (e.g., [34]), and has been extended to robot controllers [7]. Probabilistic spike emission, governed by modulatory Hebbian rules, have also been investigated [21]. Urzelai and Floreano [38] present a nodes-only encoding scheme where synapses are affected by four versions of the Hebb rule.

Previous studies have presented methods to implement STDP learning using memristive synapses, notably [2][17][19][33]. Consistent between the papers is the use of a "two-part spike". The temporal coincidence between a spike sent backwards from a postsynaptic neuron and one sent forwards from a presynaptic neuron is used to alter the voltage across the memristor between those neurons; if it surpasses a threshold voltage the weight of the synapse is altered. The main difference between [17][33] and [2][19] is that in [17][33], the two-part spike is implemented as a discrete-time stepwise waveform approximation, whereas [2][19] use values calculated from continuous waveform equations, allowing them to operate in continuous time.

3 The System

The system consists of a population of SNNs which are evaluated on a robotics test problem, and altered via GA operation which is detailed in Section 6. To introduce the terminology to be used throughout this paper: each experiment lasts for 1000 evolutionary *generations*; each generation involves new networks in the population being evaluated on the test problem (a *trial*). Each trial consists of a number of *timesteps*, which begin with the reading of sensory information and calculation of action, and end with the agent performing that action. Every timestep consists of 21 *steps* of SNN processing, at the end of which the action is calculated.

4 Spiking Network Implementation

We base our spiking implementation on the LIF model, although it must be stressed that our model is heavily simplified in terms of the number of simulation steps used per action calculation. Neurons can be stimulated either by an external current or by connections from presynaptic neurons. Each neuron has a membrane potential, y, where $y > 0$, which slowly degrades over time.

As spikes are received by the neuron, the value of y is increased in the case of an excitory spike, or decreased if the spike is inhibitory. If y surpasses a positive threshold, $ythresh$, the neuron spikes and transmits an action potential to every neuron to which it is presynaptic, with strength relative to the connection weight between those two neurons. The neuron then resets its membrane potential to some low number, given in Section 10. At time t, the membrane potential of a neuron is given as (3):

$$y(t+1) = y(t) + (I + a - by(t)) \tag{3}$$

$$If(y > ythresh) y = c \tag{4}$$

Equation (4) shows the reset formula. Here, $y(t)$ is the membrane potential at time t, I is the input current to the neuron, a is a positive constant, b is the degradation (leak) constant and c is the reset membrane potential of the neuron. A model of temporal delays is used so that, in the single hidden layer only, a spike sent to a neuron not immediately neighbouring the sending neuron is received x steps after it is sent, where x is the number of neurons between the sending neuron and receiving neuron.

4.1 Action Calculation

Action calculation involves the current input state being repeatedly processed 21 times by each network (an experimentally determined number of steps). For the purposes of this paper, each network was initialised with 6 input neurons (used to pass sensor values to the network), nine hidden neurons, and 2 output neurons that are used to calculate the action. Neurons are arranged into layers based on this classification. Each output neuron had an activation window that recorded the number of spikes produced by that neuron over the last 21 steps. To calculate the three possible actions that a network could advocate, we classified the spike trains at the two output neurons as being either *low* or *high* activated. A neuron was said to be *highly* activated if it had spikes in over half ($>$11) of the positions in the sampling window after 21 steps; otherwise it was said to have *low* activation. The combined spike trains at the two output neurons translated to a discrete movement according to the output activation strengths; (high, high) or (low, low) = forwards, (high, low) = left turn, (low, high) = right turn, which were calculated once at the end of each timestep. See Section 9.1 for precise details of sensory state generation and possible actions.

5 Memristive Connections

From the description of SNNs in Section 2.1 and that of memristors in Section 2.2, the strength of a connection weight in a neural network can be intuitively seen as the inverse memristance of that connection. The impact of the memristor on the functionality of the network depends on the memristive equations used. Next, we describe the equations governing the three comparitive memristors; details on creating variable memristors are given in Section 8.

5.1 HP Memristor

The HP memristor is comprised of thin-film TiO_2 and TiO_{2-x}, which have different resistances. The boundary between the two compounds moves in response to the charge on the memristor, which in turn alters the resistance of the device. To allow for future self-adaptation, memristance equations are refactored from the original, as given in [36]. The original profiles, used for the static memristors, are recreated using $\beta = 1$.

In the following equations, W is the scaled weight (conductance) of the connection, G is the unscaled weight of the connection, M is the memristance, $sf1$ and $sf2$ are scale factors, R_{off} is the resistance of the TiO_2, R_{on} is the resistance of the TiO_{2-x} , q is the charge on the device, $qmin$ is the minimum allowed charge, and β encompasses physical properties of the device. Equations are presented in order of calculation.

$$sf1 = 0.99/1 - \left(\frac{1}{-R_{off}R_{on}\beta q_{min} + R_{off}}\right) \tag{5}$$

$$sf2 = 1/\left(\frac{-R_{off}R_{on}\beta(R_{on} - R_{off})}{-R_{on}R_{off}\beta + R_{off}} sf1\right) - 1 \tag{6}$$

$$q = \left(\frac{1}{-R_{off}R_{on}\beta}\right)\left(\frac{sf1}{(W+sf2)} - R_{off}\right) \tag{7}$$

$$M = -R_{off}R_{on}\beta q + R_{off} \tag{8}$$

$$G = \frac{1}{M} \tag{9}$$

$$W = Gsf1 - sf2 \tag{10}$$

5.2 PEO-PANI Memristor

The PEO-PANI memristor consists of layers of PANI, onto which Li^{+}-doped PEO is added [8]. We have phenomenologically recreated the performance characteristics of the PEO-PANI memristor in terms of the HP memristor, creating a memristance curve similar to that seen in [8]. Two additional parameters, $G_{q_{min}}$ and $G_{q_{max}}$, are the values of G when q is at its minimum (q_{min}) and maximum (q_{max}) values respectively. As with the HP equations, $\beta = 1$ will produce the static PEO-PANI profile.

$$q_{max} = R_{on} - R_{off} / - R_{on}R_{off}\beta \tag{11}$$

$$G_{q_{min}} = 1/(-R_{off}R_{on}\beta q_{min} - R_{on}) + R_{on} \tag{12}$$

$$G_{q_{max}} = 1/(-R_{off}R_{on}\beta q_{max} - R_{on}) + R_{on} \tag{13}$$

The two scale factors are recalculated as

$$sf1 = 0.99/(G_{q_{max}} - G_{q_{min}}) \tag{14}$$

$$sf2 = (G_{q_{min}} sf1) - 0.01 \tag{15}$$

$$q = \left(\frac{1}{((W+sf2)/sf1) - R_{on}} + R_{on}\right)\left(\frac{1}{-R_{off}R_{on}\beta}\right) \tag{16}$$

$$M = (-R_{off}R_{on}\beta q - R_{on}) + (1/R_{on}) \tag{17}$$

G is calculated as in (9), and W as in (10).

5.3 Linear Memristor

The variable memristor alters W by $1/mem\_lifetime$, therefore it takes $mem\_lifetime$ memristance events to linearly increase W from R_{off} to R_{on}.

5.4 STDP

In Section 2.3 a number of STDP implementations were reviewed. We have elected to follow [33][17] in using discrete-time stepwise waveforms, as our SNNs operate in discrete time.

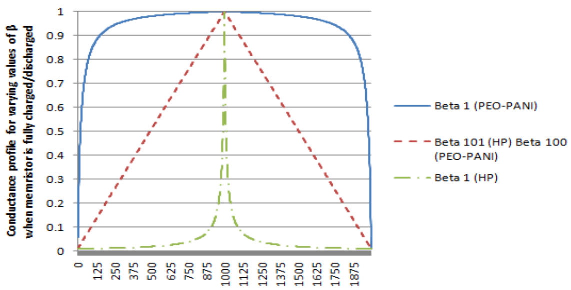

Fig. 1: Displaying memristive profiles attained with different values of β when fully charging, then fully discharging, the memristor. The x axis shows 1000 positive STDP events, followed by 1000 negative STDP events, assuming $mem\_lifetime$=1000. Static HP and PEO-PANI memristors have β=1

We augment each neuron in the network with a $last\_spike$ variable, which is initially 0. When a neuron spikes, this value is set to some positive number. At the end of each of the 21 steps that make up a single timestep, each memristive connection is analysed by checking the $last\_spike$ values of its presynaptic and postsynaptic neurons. If the calculated value exceeds some positive $spike\_threshold$, memristance of the synapse occurs. Whether the connection increases or decreases depends on which neuron has the highest $last\_spike$ value, providing pre- to postsynaptic temporal coincidence. If the $last\_spike$ values are identical, memristance does not occur. At the end of each step, each $last\_spike$ value is then decreased by 1 to a minimum of 0, creating a discrete stepwise waveform through time. Each STDP event either increases or decreases q, by Δq, as detailed in (18), which is then used to calculate the weight as detailed in Sections 5.1 - 5.3. The memory of the system is therefore contained in q.

$$\Delta q = (q_{max} - q_{min})/mem\_lifetime \tag{18}$$

From Fig. 1 it can be seen that the amount of change in connection weight depends heavily on the memristors value of β and current connection weight. Both HP and PEO-PANI memristor types display increased sensitivity (larger ΔW per STDP event) when β is a low number and either $W > 0.1$ for HP networks, or $W < 0.9$ for PEO-PANI networks. Linear memristors display constant sensitivity.

6 Discovery Component

Having described the component parts of our networks, we now detail the implementation of the GA that acts on them. In our GA, two parents are selected fitness-proportionately, mutated, and used to create two offspring. We use only mutation to explore weight space; crossover is omitted as sufficient solution space exploration can be obtained via a combination of self-adaptive weight and topology mutations; a view that is reinforced in the literature, e.g. [29]. The offspring are inserted into the population and two networks with the lowest fitness deleted. Parents stay in the population competing with their offspring.

6.1 Self-adaptive Mutation

We utilise self-adaptive mutation rates as in Evolutionary Strategies (ES) [28], to dynamically control the frequency and magnitude of mutation events taking place in each network. This allows increased structural stability in highly fit networks whilst allowing less fit networks to vary more strongly per GA application. Here, the μ $(0 < \mu \leq 1)$ value (rate of mutation per allele) of each network is initialized uniformly randomly in the range [0,0.25]. During a GA cycle, a parent's μ value is modified as in (19), the offspring then adopts this new μ, and mutates itself by this value, before being inserted into the population.

$$\mu -> \mu \exp^{N(0,1)} \tag{19}$$

Only non-memristive networks can alter their connection weights via the GA. Connection weights in this case are initially set during network creation, node addition, and connection addition randomly uniformly in the range [0,1]. Memristive network connections are always set to 0.5, and cannot be mutated from this value. This aims to force the memristive networks to harness the plasticity of their connections during a trial to successfully solve the problem.

7 Topology Mechanisms

In addition to self-adaptive mutation, we apply two evolutionary topology morphology schemes to allow the modification of the spiking networks in two ways; by adding/removing hidden layer nodes, and adding/removing inter-neural connections.

The effect of this self-adaptive, constructivist framework is to tailor the evolution of the network to the complexity of the environment explicitly. This allows each network to control its own architecture autonomously in terms of (i) amount of mutation that takes place in a given network at a given time (ii) adapting the hidden layer topology of the neural networks to reflect the complexity of the problem considered by the network. Both memristive and non-memristive networks use these topology mechanisms.

7.1 Constructivism

Constructivist learning postulates that neural structures are initially small and sparsely connected [26]. Learning acts to add appropriate structure (neurons/connections) until some satisfactory level of computing power is attained; suitable specialized neural structures emerge as a result of the learner's interaction with its environment. Implementations include Synapsing Variable Length Crossover (SVLC) [16], and its inspiration, the Species Adaptive Genetic Algorithm (SAGA) [11]. Also relevant is Neuro Evolution of Augmenting Topologies (NEAT) [35], which combines neurons from a predetermined number of subpopulations to encourage diverse neural utility and enforce niche-based evolutionary pressure.

The implementation of constructivism in this system is based on that used to evolve constructive SNNs in neural Learning Classifier Systems (LCS) [15], due to the demonstrated utility of this approach when using spiking neurons. Each network has a varying number of hidden layer neurons (initially 9, and always > 0); additional neurons can be added or removed from the single hidden layer. Constructivism takes place during a GA cycle, after mutation. Two new self-adaptive parameters, ψ $(0 < \psi \leq 1)$ and ω $(0 < \omega \leq 1)$, are incorporated into the model. Here, ψ is the probability of performing a constructivism event and ω is the probability of adding a neuron; removal occurs with probability $1-\omega$. Both have initial values taken from a random uniform distribution, with ranges [0,0.5] for ψ and [0,1] for ω. Offspring networks have their parents ψ and ω values modified using (19) as with μ. Nodes created during constructivism are initially excitatory with 50% probability, otherwise they are inhibitory.

7.2 Connection Selection

Automatic feature selection is a method of reducing the dimensionality of the data input to a process by using computational techniques to select and operate exclusively on a subset of inputs taken from the entire set. In the context of neural networks, the connection structure - as opposed to the connection weights - of artificial neural networks was first evolved by Dolan and Dyer [5]. In this paper we allow each connection to be individually enabled/disabled, a mechanism termed "Connection Selection". During a GA cycle a connection can be enabled or disabled based on a new self-adaptive parameter τ (which is initialized and self-adapted in the same manner as μ and ψ). If a connection is enabled for a non-memristive network, its connection weight is randomly initialised uniformly in the range [0,1], memristive connections are always set to 0.5. During a node addition event, new connections are set probabilistically, with $P(connectionenabled) = 0.5$. Connection selection is particularly important to the memristive networks. As they cannot alter connection weights via the GA, variance induced in network connectivity patterns plays a large role in the generation of STDP in the networks.

8 Variable memristors

Despite being a field in its infancy, a number of different memristors have been manufactured from a variety of different materials, including $Ti0_2$ [36], conductive polymer [8], AgSi [17], and crystalline oxides [6]. Unique memristive profiles are being discovered regularly; for this reason it is assumed that any evolved memristors will have an approximate physical analogue and thus any results will be (eventually) physically replicable. The notion that varied memristive behaviours could be combined in a single network is an attractive one from a computing perspective, as more functional degrees of freedom are afforded to the synapses. Additionally, certain memristive behaviours may be more beneficial than others in certain positions within the network. Mixing different types of synaptic plasticity has been previously investigated by Soltoggio [34] , Maass and Zador [21], and Urzelai and Floreano[38].

To allow the memristive profiles of the connections to change, we self-adapt β, which in reality is derived from the thickness of the device and the mobility of the oxygen vacancies in TiO_2 and TiO_{2-x}. The self-adaptive memristor profiles are allowed to range from HP-like to PEO-PANI-like profiles, each of which are governed by different equations, outlined in Sections 5.1 and 5.2. Because of this, we augment each memristor with a *type*, which is set to either 0 and 1 on memristor initialisation, with P=0.5 of each *type* being selected based on a uniform distribution. β is then initialised randomly uniformly in the range [β_{min}, β_{max}]. If *type = 0*, the refactored HP equations are used to calculate the profile of the device; otherwise the PEO-PANI equations are used. During GA activity, on satisfaction of a new self-adaptive parameter ι, which is self-adapted as with μ, a memristors β changes by $+/-10\%$ of the total range of β. If a memristors new value of β surpasses a threshold β_{max}, the *type* of the memristor is switched and a new β calculated as Δβ - β_{max}. In this way, a smooth transition between the different profile types is provided.

9 Test Environment

Our chosen robotics simulator was the Webots platform [23], a test bed that is popular amongst the research community; alternatives are summarised [4]. Previous applications of Webots in the field of evolutionary robotics include the application of incremental neuro-evolution to generate complex behaviours [37], and investigations into hierarchical neural control [25].

9.1 The Agent

The agent was a simulated Khepera II robot with 8 light sensors and 8 distance sensors. At each timestep (64ms in simulation time), the agent sampled its light sensors, whose values ranged from 8 (fully illuminated) to 500 (no light) and IR distance sensors, whose response values ranged from 0 (no object detected) to 1023 (object very close). All sensor readings were scaled to the range [0,1] for computational reasons (0 being unactivated, 1 being highly activated). Six sensors were used to comprise the input state for the SNN, three IR and three light sensors at positions 0, 2 and 5 as shown in figure 2(a). Additionally, two bump sensors were added to the front-left and front-right of the agent to prevent it from becoming stuck against an object.

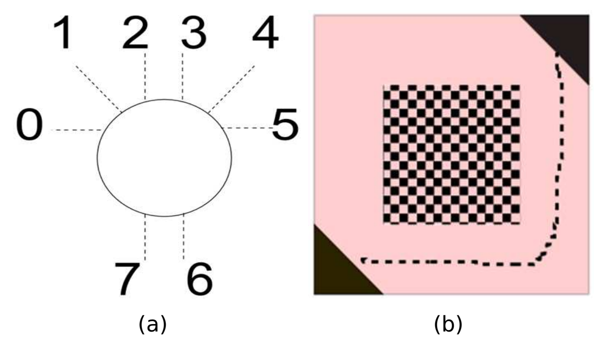

Fig. 2: (a)The sensory configuration of the simulated Khepera agent (b)The test environment; the agent starts in the lower left triangle and must navigate to the upper-right triangle, avoiding the central obstacle.

If either bump sensor was activated, an interrupt was sent causing the agent to reverse 10cm and the agent to be penalised by 10 timesteps. Movement values and sensory update delays were constrained by accurate modelling of physical Khepera agent. Three actions were possible: forward, (maximum movement on both left and right wheels based on physical Khepera data) and continuous turns to both the left and right (caused by halving the left/right motor outputs respectively).

9.2 Environment

The agent was randomly located within a walled arena - the environment - which it could not leave, with coordinates ranging from [-1,1] in both x and y directions. We constrained the agents initial starting position to a triangle in the lower left-hand corner of the environment, where $x+y<-1.5$. Adding to the complexity of the environment, a three-dimensional box was placed centrally in the arena, with vertices on “ground level” at ($x=-0.4$, $y=-0.4$), $(-0.4, 0.4)$, $(0.4, 0.4)$, and $(0.4, -0.4)$, and raised to a height of $z=0.15$. A light source, modelled on a 15 Watt bulb, was placed at the top-right hand corner of the arena ($x=1$, $y=1$), which the agent must approach. The environment is shown in figure 2(b).

Tab. 1: Detailing performance t-test results (p values) for all systems in the experiment

	Performance	High fitness	Neurons	Connectivity	ψ	ω	τ
SA-HP	**0.036**	0.146	0.198	0.316	0.08	**0.002**	0.09
SA-PEO	0.128	0.421	0.232	0.79	0.058	**0.005**	0.061
SA-LIN	0.132	**0.048**	0.520	0.227	0.865	**0.008**	0.915
SA-GA	0.077	0.532	**0.046**	**0.019**	0.478	**0.005**	**0.017**

When the agent reached the goal state (where $x+y>1.6$), the responsible network received a constant fitness bonus of 2500, which was added to the fitness function f outlined in (20). The denominator in the equation expresses the difference between the position of the goal state (1.6) and the current agent position ($posx$ and $posy$), and st is the number of timesteps taken to solve. The minimum value of this function is capped so that $f>0$. The fitness of an agent is calculated at the end of every timestep, with the highest attained value of f during the trial kept as the fitness value for that network. Optimal performance gives $f=11800$, which corresponds to 700 timesteps from start to goal state with no collisions.

$$f=(1/(1.6-(|posx+posy|)))\times 1000-st \tag{20}$$

10 Experimental Setup

In the following experiments we gauged the impact of variable memristor connections, comparing to benchmark systems comprised of homogeneous static HP, PEO-PANI and linear memristors, and constant connections. We refer to the various network types as follows: variable memristor = SA, static HP memristor = HP, static PEO-PANI memristor = PEO, static linear memristor = LIN, non-memristive network = GA. All experiments had a population size of 100 networks and were evolved for 1000 generations, with a maximum of 4000 timesteps per trial. SNN parameters were *initial hidden layer nodes*=9, $a=0.3$, $b=0.05$, $c=0.0$, $c_ini=0.5$, $ythresh=1.0$, *output window size*=21, $last\_spike=3$, $spike\_threshold=4$. In memristive networks, all connections were memristive. Memristive parameters were $R_{on}=0.01$, $R_{off}=1$, static β=1, $\beta_{min}=1$, $\beta_{maxHP}=100$, $\beta_{maxPEO-PANI}=100$ $qmin=0.0098$, $mem\_lifetime=1000$.

The experiment was repeated ten times, the statistics recorded were the averages of these ten runs. Every 20 trials, the current state of the system was stored and used to create the results that follow. To facilitate useful comparisons, we defined a notion of "performance". As the start location was tightly constrained, we say the performance of the system is equal to the first trial in which the goal state is found, so that a lower value indicated higher performance. This measure allowed us to perform t-tests to compare the respective performances of the four systems. In the following tables, "Performance" was the average performance as outlined above. "High fitness" refers to the average fitness of the highest-fitness network in each run. "Neurons" were the average final connected neurons per network in the population and "Connectivity" was the average percentage of enabled connections in the population. During a trial, some of the memristive connections in the networks may experience STDP, altering their weights. After every trial, memristive connections were reset to their original values of 0.5.

10.1 Results

In solving the test problem, two general high-fitness strategies were employed by the SA networks. The first involved a chain of "forwards" actions, a number of "turn right" actions as the agent circumvented the obstacle, and finished with successive "forwards" actions until the goal state was reached. The second strategy was a mirror of the first, but passing below the obstacle and turning left. In either case, STDP was harnessed to turn the agent. HP-governed SA beta profiles were found to quickly reduce synaptic efficiacy to the left (right) motor, causing peturbation of calculated action during turn by bringing that motor below the "high activited" threshold.

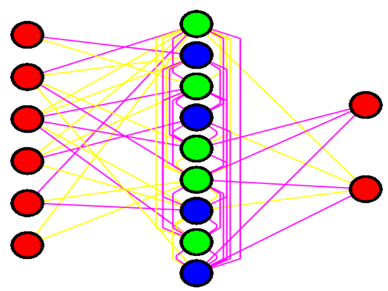

Fig. 3: An example evolved network. Light coloured neurons are excitory, dark neurons are inhibitory

PEO-PANI-governed SA profiles to the same motor were used to swiftly increase the level of spiking activity (usually in response to a light sensor surpassing some threshold) until a "forwards" action was calculated after the turn was completed. A typical evolved network is shown in Figure 3.

10.1.1 Performance

T-tests (data columns 1-4, Table 1) show a number of promising results. SA networks achieved statistically higher "performance" than HP networks (p=0.036), and higher "high fitness" than networks LIN networks (p=0.048). Figure 4(a) shows SA networks trail slightly behind PEO and LIN networks in terms of performance, but are better than HP and GA. This is an encouraging result considering the vastly increased search/behaviour space the SA discovery component has to deal with due to β, as it indicates that the variable memristor induces no significant performance overhead. Similar results can be seen for average fitness, Figure 4(b).

Figure 4(c) reveals that SA networks finish with the fewest number of required neurons, significantly less than GA networks (Table 1 showns p=0.046). However, Figure 4(d) shows that GA networks have statistically sparser connectivity (p=0.019) than SA networks. Connections were expected to be more prolific in SA networks as they are computationally more powerful than static connections, a notion echoed in recent literature [1].

Considering possible hardware implementations, CMOS neurons are larger and more complex than the synapses that connect them. As neuron numbers are more likely to be a constraint, a reduction in neurons despite increased connectivity can be said to be beneficial. In the case of both Figure 4(c) and 4(d), the profile of the HP memristor was due to one anomalous result, and as such did not infer statistical significance in either case.

10.1.2 Self-adaptive parameters

Self-adaptive parameter results can be seen in the final 3 columns of Table 1 (figures not shown). μ was not compared as it was only used in GA networks (final average value 0.022). Similarly ι was only used in SA networks (final average value 0.026). Two main results were of interest (i) SA networks had a universally lower ω, which governed the rate of neuron addition. This result indicated more parsimonious SA network evolution, allowing significantly less neurons than GA networks (ii) GA networks had lower τ than SA networks, allowing sparser connectivity (Figure 4(d)).

10.1.3 Evolution of β

As β varied between 1-101 in the case of SA HP-governed profiles and 1-100 in the case of SA PEO-PANI-goverened profiles, the total range of β is 199, where a value between 1-101 is considered to be a HP-governed profile and anything $>$101 is a PEO-PANI-goverened profile.

Analysis of the SA networks revealed that the connections to the motor on the side that made the turn evolved less linear profiles, allowing for quicker action switching behaviour. In addition, connections between the input and hidden layer had a lower average maximum (145.98 vs. 150.11) and higher average minimum (50.945 vs. 27.533) than those between the hidden and output layer. This suggests that connections to motors in general were evolutionarily preferred to have more nonlinear conductance profiles. Connections in both of these areas had higher average maximum and lower average minimum β values than connections within the hidden layer (116.16 and 84.1 respectively), suggesting that more steady memristance profiles were preferred there.

11 Conclusions

In this paper we have introduced the notion of a variable memristor and analysed its synaptic performance when compared to three static memristor types and non-memristive connections in a simulated robotics environment. The main benefit of memristive STDP over other STDP implementations lies in hardware implementations, as the efficiacy and past history of the synapse is stored in the nonvolatile physical state of the device and thus does not require simulation.

Our hypothesis was that the additional degrees of functional freedom afforded to the variable memristors allowed them to outperform these other networks. Numerous findings were found to support this hypothesis, including higher performance than HP networks, higher quality solutions than linear memristor networks, and fewer required neurons than GA networks. These findings suggest that self-adapation of β is harnessed by the evolutionary process to provide flexible plastic networks with more implicit degrees of freedom than the other network types. Variable plasticity was harnessed via STDP to achieve more expedient goal-finding behaviour with reduced topological complexity when compared to certain other network types. Importantly, the self-adaptation process itself was found to be non-disruptive with respect to network performance. Possible future research directions include hardware and mixed-media implementations.

References

[1] L. Abbott and W. Regehr. Synaptic computation. *Nature*, (431):796–803, 2004.

[2] A. Afifi, A. Ayatollahi, and F. Raissi. Stdp implementation using memristive nanodevice in cmos-nano neuromorphic networks. *IEICE Electronics Express*, 6(3):148–153, 2009.

[3] L. Chua. Memristor-the missing circuit element. *Circuit Theory, IEEE Transactions on*, 18(5):507 – 519, Sept. 1971.

[4] J. Craighead, R. Murphy, J. Burke, and B. Goldiez. A survey of commercial open source unmanned vehicle simulators. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 852 – 857, apr 2007.

[5] C. P. Dolan and M. G. Dyer. Toward the evolution of symbols. In J. J. Grefenstette, editor, *Genetic Algorithms and their Applications (ICGA'87)*, pages 123–131, Hillsdale, New Jersey, 1987. Lawrence Erlbaum Associates.

[6] W. Doolittle, W. Calley, and W. Henderson. Complementary oxide memristor technology facilitating both inhibitory and excitatory synapses for potential neuromorphic computing applications. In *Semiconductor Device Research Symposium, 2009. ISDRS '09. International*, pages 1 –2, 2009.

[7] P. Durr, C. Mattiussi, A. Soltoggio, and D. Floreano. Evolvability of Neuromodulated Learning for Robots. In *Proceedings of the 2008 ECSIS Symposium on Learning and Adaptive Behavior in Robotic Systems*, pages 41–46, Los Alamitos, CA, 2008. IEEE Computer Society.

[8] V. Erokhin and M. P. Fontana. Electrochemically controlled polymeric device: a memristor (and more) found two years ago. *ArXiv e-prints*, July 2008.

[9] D. Floreano, P. Dürr, and C. Mattiussi. Neuroevolution: from architectures to learning. 2008.

[10] W. Gerstner and W. Kistler. *Spiking Neuron Models - Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.

[11] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: artificial evolution, real vision. In *Proceedings of the third international conference on Simulation of adaptive behavior : from animals to animats 3: from animals to animats 3*, pages 392–401, Cambridge, MA, USA, 1994. MIT Press.

[12] D. O. Hebb. *The organisation of behavior*. Wiley, New York, 1949.

[13] Y. Ho, G. M. Huang, and P. Li. Nonvolatile memristor memory: Device characteristics and design implications. In *ICCAD*, pages 485–490. IEEE, 2009.

[14] J. H. Holland. Adaptation. In R. Rosen and F. M. Snell, editors, *Progress in theoretical biology IV*, pages 263–293. Academic Press, New York, 1976.

[15] G. Howard, L. Bull, and P.-L. Lanzi. A spiking neural representation for xcsf. In *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pages 1 –8, jul 2010.

[16] B. Hutt and K. Warwick. Synapsing variable-length crossover: Meaningful crossover for variable-length genomes. *Evolutionary Computation, IEEE Transactions on*, 11(1):118 –131, 2007.

[17] S. H. Jo, T. Chang, I. Ebong, B. B. Bhadviya, P. Mazumder, and W. Lu. Nanoscale memristor device as synapse in neuromorphic systems. *Nano Letters*, 10(4):1297–1301, 2010. PMID: 20192230.

[18] W. M. Kistler. Spike-timing dependent synaptic plasticity: a phenomenological framework. *Biological Cybernetics*, 87(5-6):416–427, 2002.

[19] B. Linares-barranco and T. Serrano-gotarredona. Memristance can explain spike-time- dependent-plasticity in neural synapses, 2009.

[20] W. Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 1997.

[21] W. Maass and A. M. Zador. Dynamic stochastic synapses as computational units. *Neural Computation*, 11(4):903–917, 1999.

[22] C. Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.

[23] O. Michel. WebotsTM: Professional mobile robot simulation. *International Journal of Advanced Robotic Systems*, 1(1):39–42, 2004.

[24] S. Nolfi and D. Floriano. *Evolutionary Robotics*. The MIT Press, Cambridge, Massachusetts, 2000.

[25] R. W. Paine, R. W. Paine, J. Tani, and J. Tani. How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior*, 13:211–225, 2005.

[26] S. R. Quartz and T. J. Sejnowski. The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*, Jan. 01 1997.

[27] J. M. Rabaey. *Digital integrated circuits: a design perspective*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.

[28] I. Rechenberg. *Evolutionsstrategie: optimierung technischer systeme nach prinzipien der biologischen evolution*. Frommann-Holzboog, 1973.

[29] M. Rocha, P. Cortez, and J. Neves. Evolutionary neural network learning. In *Progress in Artificial Intelligence*, volume 2902 of *Lecture Notes in Computer Science*, pages 24–28. Springer Berlin / Heidelberg, 2003.

[30] D. Rumelhart and J. McClelland. *Parallel Distributed Processing*, volume 1 & 2. MIT Press, Cambridge, MA, 1986.

[31] K. Saggie-Wexler, A. Keinan, and E. Ruppin. Neural processing of counting in evolved spiking and mcculloch-pitts agents. *Artificial Life*, 12(1):1–16, 2006.

[32] G. Snider. Computing with hysteretic resistor crossbars. *Appl. Phys. A.*, 80:1165–1172, 2005.

[33] G. Snider. Spike-timing-dependent learning in memristive nanodevices. In *Nanoscale Architectures, 2008. NANOARCH 2008. IEEE International Symposium on*, pages 85 –92, jun 2008.

[34] A. Soltoggio. Neural plasticity and minimal topologies for reward-based learning. In *Proceedings of the 2008 8th International Conference on Hybrid Intelligent Systems*, pages 637–642, Washington, DC, USA, 2008. IEEE Computer Society.

[35] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.

[36] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams. The missing memristor found. *Nature*, 453:80–83, 2008.

[37] R. A. Téllez and C. Angulo. Progressive design through staged evolution. *Frontiers in Evolutionary Robotics*, 2008.

[38] J. Urzelai and D. Floreano. Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evol. Comput.*, 9:495–524, December 2001.

[39] B. Widrow. An adaptive 'adaline' neuron using chemical 'memistors'. Technical Report 1533-2, Stanford Electronics Laboratories, Stanford, CA, oct 1960.

About the authors

Gerard Howard is a research associate based in the Department of Computer Science and Creative Technologies at UWE. His research to this point has been focused on the application of evolutionary approaches to neural control systems.

Ella Gale is a research associate with the Unconventional Computing Group at UWE. She obtained her M.Sci. and Ph.D. degrees from Imperial College London in 2003 and 2007 respectively.

Larry Bull is a Professor of Artificial Intelligence and based in the Department of Computer Science and Creative Technologies at UWE. His research interests are in intelligent and unconventional systems, with an emphasis on evolution. He is the founding Editor-in-Chief of the Springer journal Evolutionary Intelligence.

Ben de Lacy Costello received his B.Sc. in Applied Chemical Sciences and Ph.D. in Materials Science from the University of the West of England, Bristol. He is now a Senior Research Fellow in the Faculty of Health and Life Sciences and his research is split between producing sensors and systems for early disease diagnosis and nature inspired unconventional computing.

Andrew Adamatzky is a Professor in Unconventional Computing in the Department of Computer Science, ranger of the Unconventional Computing Centre, and a member of Bristol Robotics Lab. He does research in reaction-diffusion computing, cellular automata, physarum computing, massive parallel computation, applied mathematics, collective intelligence and robotics.

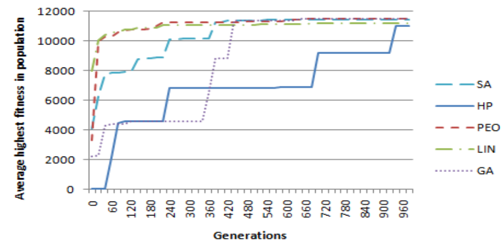

(a)

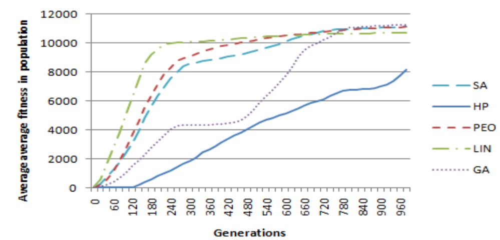

(b)

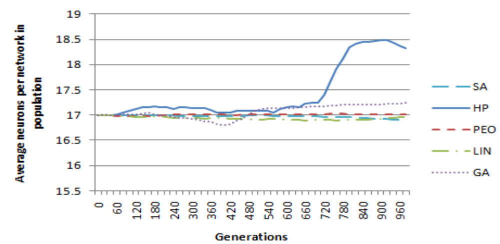

(c)

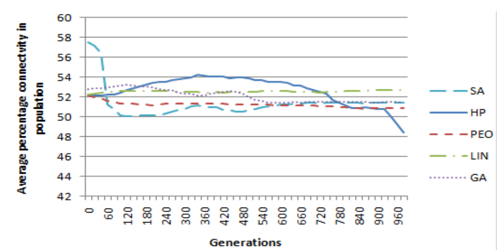

(d)

Fig. 4: (a) average highest fitness (b) average average fitness (c) average connected hidden layer nodes, (d) average enabled connections for the experiment

April 2012

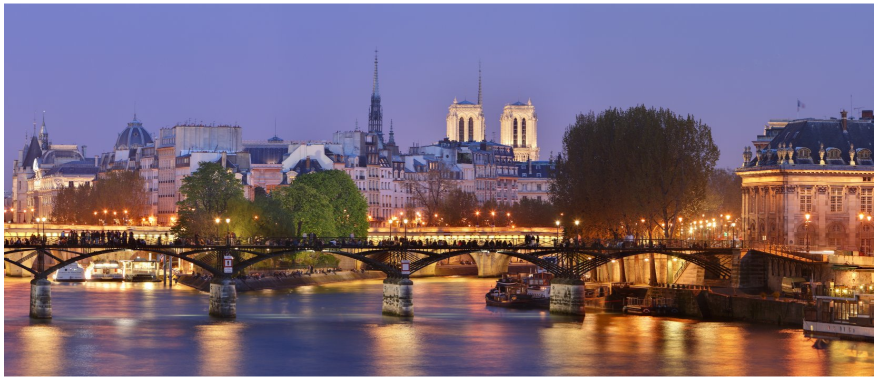

Learning and Intelligent OptimizatioN Conference - LION 6

January 16-20, 2012, Paris, France

Homepage: http://www.intelligent-optimization.org/LION6/

Call for Papers: www

Deadline October 14, 2011

Notification to authors: November 28, 2011

Conference dates: January 16-20, 2012

Camera ready for post-proceedings: February 24, 2012

The LION conference aims at exploring the intersections between machine learning, artificial intelligence, mathematical programming and algorithms for hard optimization problems. The main purpose of the event is to bring together experts from all these areas to present and discuss new ideas, new methods, general trends, challenges and opportunities in applications as well as in research aiming at algorithmic advances. The conference program will consist of plenary presentations, introductory and advanced tutorials, technical presentations, and it will give ample time for discussions.

Relevant Research Areas

LION 6 solicits contributions dealing with all aspects of learning and intelligent optimization. Topics of interest include, but are not limited to:

- Metaheuristics such as tabu search, iterated local search, evolutionary algorithms, ant colony optimization, particle swarm optimization, and memetic algorithms
- Hybridizations of metaheuristics with other techniques for optimization
- Hyperheuristics and automatic design of heuristics
- Machine learning-aided search and optimization
- Algorithm portfolios and off-line tuning methods
- Reactive search optimization, autonomous search, adaptive and self-adaptive algorithms
- Specific adaptive metaheuristic techniques applied to propositional satisfiability, scheduling and planning, routing and logistics problems
- Interface(s) between discrete and continuous optimization
- Algorithms for dynamic, stochastic and multi-objective problems
- Multiscale and multilevel methods

For all the previous approaches:

- Experimental analysis and modeling
- Parallelization techniques
- Theoretical foundations
- Innovative applications

High-quality scientific contributions to these topics are solicited, in addition to advanced case studies from interesting, high-impact application areas.

Submission Details

LION 6 accepts the following three submission types:

- Long paper: original novel and unpublished work (max. 15 pages in Springer LNCS format);
- Short paper: an extended abstract of novel work (max. 4 pages in Springer LNCS format);
- Work for oral presentation only (in any Latex format, no page restriction). For example work already published elsewhere, which is relevant and which may solicit fruitful discussion at the conference.

Further Information

Up-to-date information will be published on the web site www.intelligent-optimization.org/LION6. For information about local arrangements, registration forms, etc., please refer to the above-mentioned web site or contact the organizers.

LION 6 Conference and Technical co-chairs

- Youssef Hamadi, Microsoft Research, UK (youssefh@microsoft.com)
- Marc Schoenauer, INRIA, France (Marc.Schoenauer@inria.com)

April 2012

Evostar 2012 - EuroGP, EvoCOP, EvoBIO, EvoMusart and EvoApplications
April 11-13, 2012, Malaga, Spain
Homepage: http://www.evostar.org
Flyer: pdf
Deadline November 30, 2011
Notification to authors: January 14, 2012
Camera-ready deadline: February 5, 2012

evo* comprises the premier co-located conferences in the field of Evolutionary Computing: **eurogp**, **evocop**, **evobio**, **evomusart** and **evoapplications**.

Featuring the latest in theoretical and applied research, evo* topics include recent genetic programming challenges, evolutionary and other meta-heuristic approaches for combinatorial optimization, evolutionary algorithms, machine learning and data mining techniques in the biosciences, in numerical optimization, in music and art domains, in image analysis and signal processing, in hardware optimization and in a wide range of applications to scientific, industrial, financial and other real-world problems.

eurogp (flyer)

15th European Conference on Genetic Programming Papers are sought on topics strongly related to the evolution of computer programs, ranging from theoretical work to innovative applications.

evocop (flyer)

12th European Conference on Evolutionary Computation in Combinatorial Optimization Practical and theoretical contributions are invited, related to evolutionary computation techniques and other meta-heuristics for solving combinatorial optimization problems.

evobio (flyer)

10th European Conference on Evolutionary Computation, Machine Learning and Data Mining in Computational Biology Emphasis is on evolutionary computation and other advanced techniques addressing important problems in molecular biology, proteomics, genomics and genetics, that have been implemented and tested in simulations and on real-life datasets.

evomusart (flyer)

1st International Conference and 10th European Event on Evolutionary and Biologically Inspired Music, Sound, Art and Design

evoapplications (flyer)

European Conference on the Applications of Evolutionary Computation

evocomnet

9th European event on nature-inspired techniques for telecommunication networks and other parallel and distributed systems

evocomplex

3rd European event on algorithms and complex systems

evofin

6th European event on evolutionary and natural computation in finance and economics

evogames

4th European event on bio-inspired algorithms in games

evohot

7th European event on bio-inspired heuristics for design automation

evoiasp

14th European event on evolutionary computation in image analysis and signal processing

evonum

5th European event on bio-inspired algorithms for continuous parameter optimisation

evopar

1st European event on parallel and distributed Infrastructures

evorisk

1st European event on computational intelligence for risk management, security and defence applications

evostim

7th European event on nature-inspired techniques in scheduling, planning and timetabling

evostoc

9th European event on evolutionary algorithms in stochastic and dynamic environments

evotranslog

6th European event on evolutionary computation in transportation and logistics

July 2012

GECCO 2012 - Genetic and Evolutionary Computation Conference

July 7-11, 2012, Philadelphia, PA, USA

Homepage: http://www.sigevo.org/gecco-2012

Deadline January 13, 2012

Author notification: March 13, 2012

Workshop and tutorial proposals submission: November 07, 2011

Notification of workshop and tutorial acceptance: November 28, 2011

The Genetic and Evolutionary Computation Conference (GECCO-2012) will present the latest high-quality results in the growing field of genetic and evolutionary computation.

Topics include: genetic algorithms, genetic programming, evolution strategies, evolutionary programming, real-world applications, learning classifier systems and other genetics-based machine learning, evolvable hardware, artificial life, adaptive behavior, ant colony optimization, swarm intelligence, biological applications, evolutionary robotics, coevolution, artificial immune systems, and more.

Organizers

General Chair:	Jason Moore
Editor-in-Chief:	Terence Soule
Publicity Chair:	Xavier Llorá
Tutorials Chair:	Gabriela Ochoa
Students Chair:	Josh Bongard
Workshops Chair:	Bill Rand
Competitions Chairs:	Daniele Loiacono
Business Committee:	Wolfgang Banzhaf
	Marc Schoenauer
EC in Practice Chairs:	Jörn Mehnen
	Thomas Bartz-Beielstein,
	David Davis

Important Dates

Paper Submission Deadline	January 13, 2012
Decision Notification	March 13, 2012
Camera-ready Submission	April 9, 2012

To Propose a Tutorial or Workshop

A detailed call for workhop and tutorial proposals will be posted later so stay tuned! Meanwhile, for enquiries regarding tutorials contact gecco2012tutorials@sigevolution.org while for enquiries about workshops contact gecco2012workshops@sigevolution.org.

More Information

Visit www.sigevo.org/gecco-2012 for information about electronic submission procedures, formatting details, student travel grants, the latest list of tutorials and workshop, late-breaking papers, and more.

Contact

For general help and administrative matters contact GECCO support at gecco2012@sigevolution.org

GECCO is sponsored by the Association for Computing Machinery Special Interest Group for Genetic and Evolutionary Computation.

September 2012

PPSN 2012 – International Conference on Parallel Problem Solving From Nature

September 1-5, 2012, Taormina, Italy
Homepage: http://www.dmi.unict.it/ppsn2012/
Call for paper: www
Email: ppsn2012@dmi.unict.it
Paper Submission Deadline: March 15, 2012
Author Notification: June 1, 2012
Workshop Proposals Submission: October 15, 2011

PPSN XII will showcase a wide range of topics in Natural Computing including, but not restricted to: Evolutionary Computation, Quantum Computation, Molecular Computation, Neural Computation, Artificial Life, Swarm Intelligence, Artificial Ant Systems, Artificial Immune Systems, Self-Organizing Systems, Emergent Behaviors, and Applications to Real-World Problems.

Paper Presentation

Following the now well-established tradition of PPSN conferences, all accepted papers will be presented during small poster sessions of about 16 papers. Each session will contain papers from a wide variety of topics, and will begin by a plenary quick overview of all papers in that session by a major researcher in the field. Past experiences have shown that such presentation format led to more interactions between participants and to a deeper understanding of the papers. All accepted papers will be published in the LNCS Proceedings.

Paper Submission

Researchers are invited to submit original work in the field of natural computing as papers of not more than 10 pages. Authors are encouraged to submit their papers in LaTeX. Papers must be submitted in Springer Verlag's LNCS style through the conference homepage, here.

About the Newsletter

SIGEVOlution is the newsletter of SIGEVO, the ACM Special Interest Group on Genetic and Evolutionary Computation.

To join SIGEVO, please follow this link [WWW]

Contributing to SIGEVOlution

We solicit contributions in the following categories:

Art: Are you working with Evolutionary Art? We are always looking for nice evolutionary art for the cover page of the newsletter.

Short surveys and position papers: We invite short surveys and position papers in EC and EC related areas. We are also interested in applications of EC technologies that have solved interesting and important problems.

Software: Are you are a developer of an EC software and you wish to tell us about it? Then, send us a short summary or a short tutorial of your software.

Lost Gems: Did you read an interesting EC paper that, in your opinion, did not receive enough attention or should be rediscovered? Then send us a page about it.

Dissertations: We invite short summaries, around a page, of theses in EC-related areas that have been recently discussed and are available online.

Meetings Reports: Did you participate in an interesting EC-related event? Would you be willing to tell us about it? Then, send us a short summary, around half a page, about the event.

Forthcoming Events: If you have an EC event you wish to announce, this is the place.

News and Announcements: Is there anything you wish to announce? This is the place.

Letters: If you want to ask or to say something to SIGEVO members, please write us a letter!

Suggestions: If you have a suggestion about how to improve the newsletter, please send us an email.

Contributions will be reviewed by members of the newsletter board.

We accept contributions in LaTeX, MS Word, and plain text.

Enquiries about submissions and contributions can be emailed to editor@sigevolution.org.

All the issues of SIGEVOlution are also available online at www.sigevolution.org.

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in the Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library
- to allow users to copy and distribute the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will make every effort to refer requests for commercial use directly to you.