Converging on the Divergent: The history (and future) of the International Joint Workshops in Computational Creativity

Amílcar Cardoso, Tony Veale and Geraint A. Wiggins

1 A matter of Emphasis

When heckled, professional comedians frequently lament that "everyone's a comedian!". It's easy to see why: professional comedians don't possess different kinds of brains from others, or engage in radically different kinds of behaviours from others, and, moreover, the success of their acts is predicated on others' shared ability to understand and reason about comic situations they describe. The difference between comedians and their audience is a matter not of kind, but of degree, a difference that is reflected in the vocational emphasis they place on humour.

Researchers in the field of computational creativity find themselves in a similar situation. As a sub-discipline of Artificial Intelligence, computational creativity explores theories and practices that give rise to a phenomenon, creativity, that all intelligent systems, human or machine, can legitimately lay claim to. Who is to say that a given AI system is not creative, insofar as it solves non-trivial problems or generates useful outputs that are not hard-wired into its programming? As with comedians' being funny, the difference between studying computational creativity and studying artificial intelligence is one of emphasis rather than one of kind: the field of computational creativity, as typified by a long-running series of workshops at AI-related conferences, places a vocational emphasis on creativity, and attempts to draw together the commonalities of what human observers are willing to call "creative" behaviours. The study of creativity in AI is not new, but it is unusual. When Margaret Boden included a chapter on creativity in her text-book, *Artificial Intelligence and Natural Man* (Boden, 1977), colleagues asked "Why on earth are you doing that?" (Boden, 1999). Sometimes, it seems that creativity is, for AI-believers, that place beyond the pale, where lies intelligence itself for AI-skeptics.

Since the mid-1990s, interest in creativity from an AI perspective has begun to blossom. Workshops dedicated to computational creativity now occur yearly or more, the foremost being the International Joint Workshop on Computational Creativity (IJWCC). This series grew out of a number of events in the 1990s, including the International Workshop on Computational Humor at the university of Twente in 1996, the Mind II conference on Creative Computation at Dublin City University in 1997, and the Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB) in Edinburgh in 1999, whose central theme was creativity in AI. This last was probably the largest AI/creativity-focused conference to be held to date, attracting 225 participants. Subsequent to these developments, in 2000, the AISB Convention hosted the first of several workshops on various aspects of AI and creativity; one year later, a workshop series on Creative Systems began holding workshops jointly with major AI-related conferences, namely ICCBR 2001, ECAI 2002 and IJCAI 2003; in 2008, AAAI held its first Spring Symposium on Computational Creativity. It seems clear that a relatively small and new community would struggle to sustain so much activity, and therefore, in 2004, the Creative Systems series merged with the AISB workshops to form the International Joint

Workshop series, which proceeds to date. 2009 will be the tenth year of the workshops, in one form or another, and a corresponding event is planned.

Another indication of the surge of interest in the scientific study of creativity was the triennial Creativity and Cognition conference, originally motivated from Loughborough University, UK, but now floating internationally free. It has, as one might expect from its title, a less computational slant, but nevertheless resides under the banner of the Association for Computing Machinery. Creativity, of both biological organisms and machines, is becoming a hot scientific topic.

2 The Great Creativity Debate

Creativity is an elusive phenomenon to study, or even to define, made all the more vexing by our fundamental inability to pin it down in formal terms. Ask most people the question "what is creativity?" and you are more likely to elicit an anecdote, an aphorism, or a metaphor, than you are a literal definition, least of all a definition that can contribute to the construction of a convincing computational model. It's not surprising, then, that a formal definition of creativity—and our inability to find one that satisfies everybody—has been the elephant in the room at all of the computational creativity workshops to date. Many hours have been spent in argument about what does and does not constitute creativity. Fortunately, these arguments have always been philosophically and socially engaging, and they expose the claims of workshop contributors to the deepest possible scrutiny. Of course, the history of AI records a similar debate on the search for a consensual definition for intelligence that might be useful for building computer systems. In that debate, we have so far only agreed to disagree.

One key reason for the degree of debate on this topic at creativity workshops is that this is a field defined by a word, "creativity", rather than a concept, CREATIVITY. The word that has, historically, undergone several shifts in meaning, and it continues to mean different things to different people. One of the most beguiling aspects of language is the illusion of certainty it can grant to a speaker. The possession of a word for a given concept often implies possession of that concept itself, and when this word is both familiar and commonplace, like "creativity", we can easily fall prey to the belief that the underlying concept is itself familiar, coherent and easy to grasp. The slipperiness of CREATIVITY famously led Newell et al. (1963) to despair of an essentialist account and propose instead a multi-partite definition. They suggest four intermingling criteria for categorizing a solution as creative:

1. The answer has novelty and usefulness (either for the individual or society).
2. The answer demands that we reject ideas we had previously accepted.
3. The answer results from intense motivation and persistence.
4. The answer comes from clarifying a problem that was originally vague.

This approach is, of course, a classic AI formulation: there is at its base the implicit assumption that the created artefact is an "answer", and therefore that there must have been a question. In the creative arts and in the less empirically motivated sciences and mathematics, this need not be the case: creative motivation may be altogether less well-defined. Nevertheless, we have seen theories and models that embody each of these criteria at the creativity workshops, which we catalogue below. Most papers emphasise the first criterion, presenting computational models that are capable of generating outputs that are novel (to themselves, at least) and demonstrably useful (either aesthetically or analytically).

An alternative view of the definition of *computational* creativity is proposed by Wiggins (2006):

> "The performance of tasks [by a computer] which, if performed by a human, would be deemed creative."

This is a different kind of definition, because it casts creativity as a relation between the creator and an observer. In some contexts, this style of definition would be a cop-out, merely postponing the problem; in this context, however, it is appropriate, since creativity really is in the eye of the beholder.

3 The Eye of the Beholder

The IJWCC workshops thus provide a good forum for papers that focus on practical concerns in the development of computational systems that might, at some level, exhibit creativity. In the history of the workshop series, story-generation is a strongly represented theme, as are musical composition and improvisation, humour, metaphor, analogy and other clever uses of language. Most papers view creativity as an additional element of a system that authors would be building anyway, regardless of their interest in creativity, because their interests lie in stories, music, art or language and because these phenomena are all the more appealing when they exhibit creativity. But to be truly successful, the workshops must succeed in aligning this interest in creativity with the public's perception of creative behaviour, which is naturally inclined toward prototypical cases of human creativity and instinctively biased against anything that is artificial. This bias is further intensified by the act of creative appreciation itself, since it is the audience that imbues a product with much of its meaning and creative value. This is a point at which the analogy between computational creativity and artificial intelligence breaks down: the concept of intelligence does not entail a corresponding requirement. In general, it is often enough for a mechanism to perform a useful function with some degree of autonomy to earn the label "intelligent".

However, it bodes well for our computational efforts that artefacts have to be neither very complex nor obviously labour-intensive to attract this kind of creativity-enhancing attention. Consider, for example, a simple but memorable piece of word-play by the artist Marcel Duchamp. At a party in Paris in 1953, the tin-foil wrappers on the candy given to guests was graced with a Duchamp pun:

"A Guest + A Host = A Ghost".

The mechanics of the pun are easy to appreciate, while the knowledge needed to construct it is readily found in most dictionaries (i.e., that "host" and "guest" are related terms). Yet the scientist Stephen Jay Gould (2000) finds creativity at many different levels of this simple pun, describing it as Duchamp's "deepest and richest play on words". It's hard to read Gould's account without concluding that much of this creativity resides in his own analysis, and not in the (literally) throw-away pun itself, but we can begin to see that even the simplest computational outputs can be accorded a high level of creativity if viewed with the benevolent eye of an audience that openly (and without bias) expects creativity. Even the simplest combinations can yield large payoffs in terms of creative appreciation. Consider another example, this time a visual pun, and one that is wholly unintentional. Commissioned from a professional design company at a cost of £14,000, the bland logo for the UK Office of Government Commerce in Figure 1 (upper) gives rise to a more creative (and altogether more suggestive) interpretation when rotated 90° to the right; enough so to draw the attention of the UK national press (Simpson, 2008).

A system for generating puns like Duchamp's may not be an obvious topic for an IJCAI paper, but JAPE (Binsted et al., 1997) which plays the same game has achieved considerable success at engaging children with no speech to be linguistically creative, in the STANDUP project in Scotland (Black et al., 2007). To achieve human levels of computational creativity, we do not necessarily need to start big, at the level of whole poems, songs stories or paintings; we are more likely to succeed if we are allowed to start small, at the level of simple but creative phrases, fragments and images.

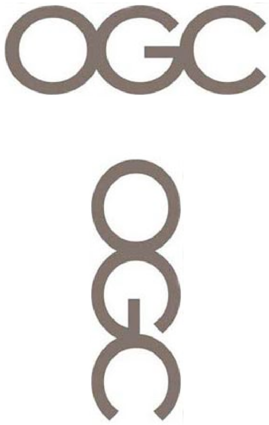

Figure 1: The logo above (for the Office of Government Commerce) reveals the rather more suggestive image below when creatively rotated 90° clockwise.

In contrast with the example of Figure 1, the same perceptual principle is at work in the much more labour-intensive and artful image-pair in Figure 2, which depicts a single frame from the 1904 comic-strip *The Upside-Downs of Gustav Verbeek*. In these one-page strips of six panels, Verbeek manages to tell a story that is twelve panels long, by asking the reader to physically invert the page once the first six panels have been read. As shown in Figure 2, in which an old man in a canoe, catching a large fish, becomes a mythical *roc* eating a defenceless maiden, each panel must be artfully constructed so that it yields an equally well-formed panel when inverted. The artfulness of Figure 2 is so far beyond the current abilities of AI systems that it is almost painful to contemplate how we might ever reach this level of creative generation. However, the representational and processing demands of Figure 1 are well within our current computational grasp, and allow us to begin to identify principles that can be exploited in much more complicated and unexpected ways.

One might ask why artefacts of this kind are worth generating in the first place. The answer lies in the kind of flexible, dual-purpose representations and emergent processes that are necessary to enable their generation. By and large, we are not going to understand these processes and representations by studying problems that are already in the AI mainstream, or by studying problems that have an apparent commercial or industrial dimension. The IJWCC workshops do not focus on pressing problems in AI (especially not problems with particular "correct" solutions), nor even on systems with immediate applications, but on problems that might illuminate the nature of human and machine creativity, and thus, one day, find their way into the AI mainstream. The fact that much of the work reported sits squarely within a given creative domain (e.g., music, mathematics, etc.) does not undermine the potential for general mechanisms—after all, there is no reason to suppose that a different mechanism for creativity is required for each possible domain.

An example of the generality of computational creativity technology can be found in the Curious Agents of Saunders and Gero (2004), where agents exhibiting certain creatively exploratory behaviours cooperate in design—but the same technology has proved successful in simulating the biology of human stem cells (d'Inverno and Saunders, 2005). So we can argue that these systems, at least, have creative behaviours that are capable of novelty, and which are directly useful to humans—fulfilling at least two of Newell et al.'s criteria.

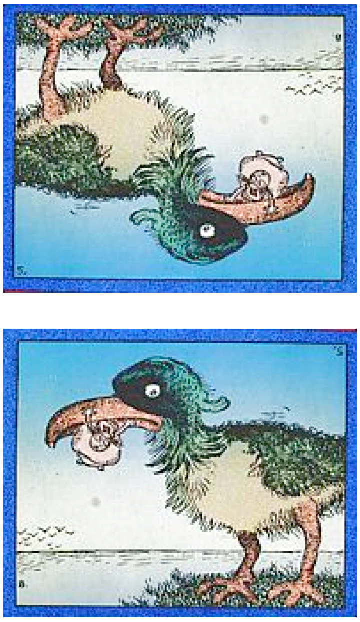

Figure 2: The lower image (panel 5 of 6 in the original), is a 180° rotation of the upper image, panel 8. Cover one of the images to see the other most clearly.

This issue of evaluation is perennial in computational creativity (Boden, 1998): how can a computational system know when its outputs are worthy of the term "creative"? In fact, of course, there is another aspect, in the scientific context, of evaluation: how do we empirically and rigorously evaluate the systems we build and decide whether or not they can genuinely be called "creative"? Ritchie (2007) presents some criteria which may be applied to a creative system formulated in a certain way, to begin an argument that it is (or is not) creative. Ritchie's criteria are cleverly couched in terms of what the system knows about its own domain, and thus can potentially be applied without a general model of AI in the background; but Boden's problem of evaluation, not of the system itself, but of the artefacts it produces as part of the creative process, is *much* harder and probably needs to be deferred until we are substantially more capable in general automated reasoning and knowledge representation.

4 Show and Tell

Philosophical arguments about the nature of creativity are just as unlikely as the evaluation problems to be resolved to universal satisfaction in the short or medium term. Though useful and engaging, they don't play to the core strengths of a computational perspective, which sees the construction of working models as the most convincing way to drive home a point; this is especially so because the identification of creativity seems so much to be a relation between artefact and observer, and not just an abstract property of the observer itself. Therefore, computational creativity workshops have begun to phase out these often circular arguments in favour of more straightforward discussions and demonstrations of what can be achieved computationally and whether and why an audience might eventually dub this "creative". Researchers therefore come to IJWCC workshops with laptops primed to give demos of what their systems can do, ready to show off features that have been added since their papers were first accepted. No abstract insight can compare with the ability to show a real creative system in full flow. Interactive "show and tell" sessions have thus become an irreplaceable—and very enjoyable—feature of the workshops.

5 Towards the Grand Challenges

As a forum for computational research in creativity, the IJWCC attracts both top-down and bottom-up approaches to creative behaviour. Top-down approaches are those that tackle a complex problem such as art generation or music composition in its entirety, albeit at a level of achievement that leaves much room for improvement. The goal of top-down development in creativity research is to establish a framework or architecture in which individual modules can be successively developed or plugged in; and some top-down approaches are expressed entirely as frameworks or as sets of characteristics of systems. Bottom-up approaches are those that isolate some more or less specific module of a larger problem—such as analogical mapping, pun generation or plot organization—that can be individually evaluated and improved. Though collaboration is frequent among IJWCC contributors, no one task or grand challenge has so far allowed for a synthesis for these different approaches on a significant scale. Nonetheless, the IJWCC provides an ideal forum for the development and management of such Grand Challenges.

It may well transpire that a Grand Challenge for computational creativity is not a solution to a particular problem, like many of the current agreed Grand Challenges in AI. Rather, it is likely to be the *way* a system does what it does, and how well, that constitutes the real challenge. For example, one less-than-Grand challenge to overcome is the very common accusation that a rule-based system cannot be creative: "But you just programmed it to do what it does!". There is, of course, an argument that the output of an complex production system is not predictable by its author, but experience shows that this does not wash with most audiences. The obvious solution is to build systems which *learn* to do what they do, before attempting to do it creatively; and there is at least one such creative system (albeit not a very creative one) now reported in the mainstream computer music literature (Wiggins et al., 2009). However, like most systems that involve multiple technologies (in this case, learning and then generation), the difficulties in designing and building the systems increase exponentially with that multiplicity. Therefore, we must follow the paths of the early AI pioneers, using relatively simple, readily comprehensible models, before we proceed to the greater complexities of realistic ones.

6 Progress

There is not space here to give a detailed survey of all the work presented over the past ten years. But there is some value in listing some of the sustained contributions (some of which have already fed into journal publications), grouped into areas of interest.

On the applied side, we have seen papers covering creative systems working in various domains. Linguistic creativity is embodied in work on forms traditionally seen as "creative", narrative and poetry (Gervás, 2000, 2001; Gervas et al., 2007; Levy, 2001), and on more common usage, much of which takes the reasoning to a more conceptual level, linking in to the more abstract conceptual work discussed below (Gervás, 2002; Hayes et al., 2004; Hervas et al., 2006; Veale, 2003). Visual art, too, has appeared, and this work has been notable for interesting attempts to characterise artistic value in perceptual terms, again linking it to the more theoretic work, below (Hull and Colton, 2007; Machado and Cardoso, 2000; P. Machado and Cardoso, 2002; Saunders and Gero, 2001). The third conventionally-understood creative domain, music, has been disproportionately well-represented, perhaps reflecting a broader interest in the scientific or empirical study of music that has been developing since the mid 1908s (Chuan and Chew, 2007; Forth et al., 2008; Iliopoulos et al., 2002; Ribeiro et al., 2001; Whorley et al., 2007). And, perhaps most importantly, because it is so often overlooked, the domain of mathematics and science is represented, too (Colton, 2001; Steel et al., 2000); this domain is as creative as any of the others, and we often do ourselves a disservice by forgetting the fact! What is more, a mathematically-creative program, HR (Colton et al., 2000) was the subject of an award-winning paper at AAAI 2000, and is the first creative program, of which we are aware, to have contributed to a reference text in its field: the Encyclopedia of Integer Sequences (`www.research.att.com/~njas/sequences/`).

There are several groups of papers concerned with more general issues which apply across domains, in various different ways. Saunders and Gero (2001), Machado et al. (2003) and Sosa and Gero (2004) are all concerned with simulations of creativity in social contexts and how creativity can emerge from inter-agent interaction. Pereira and Cardoso (2002), Veale (2003), Hao and Veale (2006) and Hervas et al. (2006) study analogy, metaphor and conceptual blending as potential mechanisms for creative reasoning in a symbolic style. Finally, de Figueiredo and Campos (2001) and O'Donoghue and Crean (2002) are interested in serendipity, and the problem of noticing an accidental creative act.

Without a theoretical literature, however, computational creativity could not dignify itself as a scientific study. The workshops have featured various proposals for models of creative processes, at various levels of abstraction (Colton, 2003; Machado et al., 2004; Magnani et al., 2002; McCormack, 2007; Pease et al., 2002; Pereira and Cardoso, 2002; Sosa and Gero, 2003; Wiggins, 2001, 2003) and also various approaches to the assessment problem, outlined above (Colton et al., 2001; Pease et al., 2001; Ritchie, 2001; Whorley et al., 2007).

7 Prospects

While most people can claim an intuitive feeling for the concept of CREATIVITY, such intuitions do not always facilitate formalization and may even obscure it. Of course, it seems clear that computational creativity in particular is a sub-branch of Artificial Intelligence, but few AI researchers who are not ostensibly working on creativity would deny that their outputs lack creativity. Creativity is thus a rather amorphous concept that has thus far resisted the development of either a foundational set of techniques or key papers, despite the growing numbers of researchers who dedicate their efforts to the enlargement of the field.

Indeed, until recently, the field of creativity research has been as much influenced by the historical anecdotes of famously creative individuals, such as Poincaré, Kekulé and Einstein, as it has been by a dedicated academic literature. Fortunately, this situation is beginning to change, in large part because of the computational emphasis of the yearly workshops that we as a community have organized. In this past decade, we have watched the community grow and solidify, and we have also watched it strive for a mutually acceptable definition of itself and its goals. We believe it is finally close to reaching a level of formalization that will act as a solid foundation for future scientific work.

One of the keystones that a growing research field requires is a canonical literature, a set of papers that guides future research and lays down markers as to the importance of key ideas and mechanisms. Our yearly workshops have so far produced two journal special issues which bring together papers that we expect will stand the test of time (Cardoso and Bento, 2006; Veale et al., 2006).

A tenth anniversary event is currently being planned, to take place in Lisbon in early 2010. We expect this event to be something of a coming-of-age for the field, as we move towards a creative future.

References

Binsted, K., Pain, H., and Ritchie, G. (1997). Children's evaluation of computer-generated punning riddles. *Pragmatics and Cognition*, 5(2):309–358.

Black, R., Waller, A., Ritchie, G., Pain, H., and Manurung, R. (2007). Evaluation of joke-creation software with children with complex communication needs. *Communication Matters*, 21(1):23–28.

Boden, M. (1977). *Artificial Intelligence and Natural Man*. Harvester Press.

Boden, M. (1998). Creativity and artificial intelligence. *Journal of Artificial Intelligence*, 103(2):347–356.

Boden, M. (1999). Preface to special issue on creativity in the arts and sciences. *AISB Quarterly*, 102.

Cardoso and Bento, 2006 (2006). Special issue on computational creativity. *Journal of Knowledge-Based Systems*, 19(7).

Chuan, C.-H. and Chew, E. (2007). A hybrid system for automatic generation of style-specific accompaniment. In Wiggins, G. and Cardoso, A., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity, London, UK*.

Colton, S. (2001). Experiments in meta-theory formation. In Wiggins, G., editor, *Proceedings of the AISB'01 Sympo- sium on Artificial Intelligence and Creativity in Arts and Science*.

Colton, S. (2003). Creative logic programming. In C. Bento, A. C. and Gero, J., editors, *Proceedings of the 3rd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, IJCAI 2003, Acapulco, Mexico*.

Colton, S., Bundy, A., and Walsh, T. (2000). Automatic invention of integer sequences. In *Proceedings of AAAI 2000*, pages 558–563. AAAI Press/MIT Press.

Colton, S., Pease, A., and Ritchie, G. (2001). The effect of input knowledge on creativity. In Cardoso, A., Bento, C., and Wiggins, G., editors, *Proceedings of the First Workshop on Creative Systems, International Conference of Case-Based Reasoning*. ICCBR-01.

de Figueiredo, A. D. and Campos, J. (2001). The serendipity equations. In Bento, C. and Cardoso, A., editors, *Proceedings of the ICCBR-01 Workshop on Creative Systems, Vancouver, Canada*, Washington, DC. Naval Research Laboratory.

d'Inverno, M. and Saunders, R. (2005). Agent-based modelling of stem cell organisation in a niche. In *Engineering Self-Organising Systems*, number 3464 in LNAI, pages 52–68. Springer, Heidelberg.

Forth, J., Mclean, A., and Wiggins, G. (2008). Musical creativity on the conceptual level. In *IJWCC 2008*.

Gervás, P. (2000). Wasp: Evaluation of different strategies for the automatic generation of spanish verse. In *AISB-00 Symposium on Artificial Intelligence and Creativite & Cultural Aspects of AI*.

Gervás, P. (2001). Generating poetry form a prose text: Creativity versus faithfulness. In *AISB-01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, University of York, England.

Gervás, P. (2002). Linguistic creativity at different levels of decision in sentence production. In *Proceedings of the AISB 02 Symposium on AI and Creativity in Arts and Science, 3rd-5th April 2002, Imperial College. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour*, pages 79–88. AISB.

Gervas, P., y Perez, R. P., and Sosa, R. (2007). On the fly collaborative story-telling: Revising contributions to match a shared partial story line. In Wiggins, G. and Cardoso, A., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity, London, UK*.

Gould, S. J. (2000). The substantial ghost: Towards a general exegesis of duchamps artful wordplays. *Tout-Fait, The Marcel Duchamp Online Journal*, 1(2).

Hao, Y. and Veale, T. (2006). A fluid category structure for metaphor processing. In *Procs. 3rd International Joint Workshop on Computational Creativity*.

Hayes, J., Seco, N., and Veale, T. (2004). Creative discovery in the lexical validation gap. In *Proceedings of the 1st International Joint Workshop on Computational Creativity, Madrid, Spain*.

Hervas, R., Pereira, F., Gervas, P., and Cardoso, A. (2006). Cross-domain analogy in automated text generation. In *Procs. 3rd International Joint Workshop on Computational Creativity*.

Hull, M. and Colton, S. (2007). Towards a general framework for program generation in creative domains. In Wiggins, G. and Cardoso, A., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity, London, UK*.

Iliopoulos, C. S., Lemstrom, K., Niyad, M., and Pinzon, Y. J. (2002). Evolution of musical motifs in polyphonic passages. In *Symposium on AI and Creativity in Arts and Science, Proceedings of AISB'02*, pages 67–76.

Levy, R. (2001). A computational model of poetic creativity with neural network as measure of adaptive fitness. In Bento, C. and Cardoso, A., editors, *Proceedings of the ICCBR-01 Workshop on Creative Systems, Vancouver, Canada*, Washington, DC. Naval Research Laboratory.

Machado, P. and Cardoso, A. (2000). NEvAr - the assessment of an evolutionary art tool. In Wiggins, G., editor, *Proceedings of the AISB'00 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, Birmingham, UK.

Machado, P., Romero, J., Manaris, B., Cardoso, A., and Santos, A. (2003). Power to the critics a framework for the development of artificial art critics. In C. Bento, A. C. and Gero, J., editors, *Proceedings of the 3rd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, IJCAI 2003, Acapulco, Mexico*.

Machado, P., Tavares, J., Cardoso, A., Pereira, F. B., and Costa, E. (2004). Evolving creativity. In *Proceedings of the 1st International Joint Workshop on Computational Creativity, Madrid, Spain*.

Magnani, L., Piazza, M., and Dossena, R. (2002). Abduction as a rational means to creativity, unexpressed knowledge and scientific discovery. In C. Bento, A. C. and Wiggins, G., editors, *Proceedings of the 2nd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, ECAI 2002, Lyon, France*.

McCormack, J. (2007). Creative ecosystems. In Wiggins, G. and Cardoso, A., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity, London, UK*.

Newell, A., Shaw, J. G., and Simon, H. A. (1963). The process of creative thinking. In Gruber, H. E., Terrell, G., and Wertheimer, M., editors, *Contemporary Approaches to Creative Thinking*, pages 63–119. Atherton, New York.

O'Donoghue, D. and Crean, B. (2002). Searching for serendipitous analogies. In C. Bento, A. C. and Wiggins, G., editors, *Proceedings of the 2nd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, ECAI 2002, Lyon, France*.

P. Machado, A. D. and Cardoso, A. (2002). Learning to colour greyscale images. *Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour, AISB Journal*, 1(2):209–219. Special issue with updated versions of selected papers from the AISB'02 Convention.

Pease, A., Colton, S., Smaill, A., and Lee, J. (2002). Lakatos and machine creativity. In C. Bento, A. C. and Wiggins, G., editors, *Proceedings of the 2nd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, ECAI 2002, Lyon, France*.

Pease, A., Winterstein, D., and Colton, S. (2001). Evaluating machine creativity. In Cardoso, A., Bento, C., and Wiggins, G., editors, *Proceedings of the First Workshop on Creative Systems, International Conference of Case-Based Reasoning*. ICCBR-01.

Pereira, F. C. and Cardoso, A. (2002). Conceptual blending and the quest for the holy creative process. In C. Bento, A. C. and Wiggins, G., editors, *Proceedings of the 2nd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, ECAI 2002, Lyon, France*.

Ribeiro, P., Pereira, F. C., Ferrand, M., and Cardoso, A. (2001). Case-based melody generation with muzacazuza. In Wiggins, G., editor, *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, pages 67–74.

Ritchie, G. (2001). Assessing creativity. In Wiggins, G., editor, *Proceedings of the AISB'01 Symposiym on Artificial Intelligence and Creativity in Arts and Science*, pages 3–11.

Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, 17(1):67–99.

Saunders, R. and Gero, J. (2001). The digital clockwork muse: A computational model of aesthetic evolution. In Wiggins, G., editor, *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, pages 12–21.

Saunders, R. and Gero, J. S. (2004). Situated design simulations using curious agents. *AIEDAM*, 18(2):153–161.

Simpson, A. (2008). OGC reveals new logo to red faces. *Daily Telegraph. April 29*.

Sosa, R. and Gero, J. (2003). Design and change: A model of situated creativity. In C. Bento, A. C. and Gero, J., editors, *Proceedings of the 3rd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, IJCAI 2003, Acapulco, Mexico*.

Sosa, R. and Gero, J. (2004). The role of gatekeepers in creativity. In *Proceedings of the 1st International Joint Workshop on Computational Creativity, Madrid, Spain*.

Steel, G., Colton, S., Bundy, A., and Walsh, T. (2000). Cross domain mathematical concept formation. In *In Proceedings of the AISB-00 Symposium on Creative & Cultural Aspects and Applications of AI & Cognitive Science*.

Veale, T. (2003). Qualia extraction and creative metaphor in wordnet. In C. Bento, A. C. and Gero, J., editors, *Proceedings of the 3rd. Workshop on Creative Systems, Approaches to Creativity in Artificial Intelligence and Cognitive Science, IJCAI 2003, Acapulco, Mexico*.

Veale et al., 2006 (2006). Special issue on computational creativity. *New Generation Computing*, 24(3).

Whorley, R., Wiggins, G. A., and Pearce, M. T. (2007). Systematic evaluation and improvement of statistical models of harmony. In Wiggins, G. and Cardoso, A., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity, London, UK*.

Wiggins, G. A. (2001). Towards a more precise characterisation of creativity in AI. In Bento, C. and Cardoso, A., editors, *Proceedings of the ICCBR2001 workshop on Creative Systems*, Naval Research Laboratory, Code 5510, Washington, DC. Navy Center for Applied Research in Artificial Intelligence.

Wiggins, G. A. (2003). Categorising creative systems. In Bento, C., Cardoso, A., and Gero, J., editors, *Proceedings of the IJCAI'03 workshop on Creative Systems*. IJCAI.

Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Journal of Knowledge Based Systems*, 19(7):449–458. `doi:10.1016/j.knosys.2006.04.009`.

Wiggins, G. A., Pearce, M. T., and Müllensiefen, D. (2009). Computational modelling of music cognition and musical creativity. In Dean, R., editor, *Oxford Handbook of Computer Music and Digital Sound Culture*. Oxford University Press. In press.