

Computational Creativity: A Continuing Journey

Tony Veale, Pablo Gervás, Rafael Pérez y Pérez

Introduction to the Special Issue of Computational Creativity

Minds and Machines Vol. 20 No. 4 Winter 2010 pp. 483–650

As the final touches are being put to this very special issue of *Minds and Machines*, which showcases a selection of the finest contributions to the 2008 International Joint Workshop on Computational Creativity, Tony Blair has just unveiled his memoirs to a sharply divided public. As befitting this controversial figure’s polarizing legacy, reactions to his autobiography – entitled *A Journey* – have ranged from the painfully predictable to the creatively original. The jeering crowds that awaited Mr. Blair outside book signings and other public events were entirely predictable, as were the cries of “war criminal” and other lazy categorizations that have long since lost their creative bite. Slightly less predictable was the hurling of shoes at Mr. Blair, an analogical echo of another time when footwear was used as a cultural missile against George W. Bush. At a Dublin book signing, this minor act of analogical creativity was both enhanced and undermined by the cost-effective and very ironic choice of shoes – flip-flops! – since Mr. Blair is most often criticized for being unwilling to recant (or “flip-flop”, in political parlance) for his involvement in the Iraq war. However, the most creative protest to greet the publication of *A Journey* was also the least noisome. Through a campaign launched via Facebook, protesters were encouraged to quietly and mischievously re-categorize Mr. Blair’s book in whatever bookshop they happened to find it. As reported in the British newspaper *The Telegraph*, employees in bookstores across the country were puzzled to find copies of the book migrating from its “official” home on the Biography shelves to sections such as Fiction and even True Crime. Euan Booth, the student who originated the Facebook campaign, explained himself thus to *The Telegraph*:

“This is a peaceful and mischievous way of making your point if you feel the same way. It’s a non-violent display of anger using the materials given to me – his book and the crime section – they’re both there, I just put them together”

When it comes to theories of creativity, computational researchers have an embarrassment of riches to choose from, not least because no single theory ever comes close to capturing creativity in all of its glorious guises. The view most favored by computationalists is that which imagines creativity as arising from the exploration (and, occasionally, the transformation) of a conceptual space. This view is especially attractive because it builds on the foundational beliefs in AI that intelligence arises out of the use of effective search techniques. However, it should be clear that in the case of Euan Booth above, such a view is supremely unhelpful. Are we to imagine the young protester exploring a conceptual space of different protest strategies, with an evaluation function that assigns greater value to non-violent solutions? Though undoubtedly useful in some contexts, the exploration view offers little computational insight in others, and is often applied post-hoc as a useful rationalization of what a heuristic-laden program is trying to achieve through other means. In this particular case, the *investment theory* of Robert J. Sternberg and Todd I. Lubart might seem to offer a slightly better explanation of how Booth’s creative ends are achieved. Sternberg and Lubart argue that creative thinkers “buy low and sell high” in the realm of ideas – that is, they take on simplistic, unpopular and seemingly worthless ideas, adapt them to their current goals, and then derive enhanced value from the adapted ideas before moving on to other unfashionable territory. So rather than loudly and crassly asserting his target to be a “criminal”, Booth recognized the untapped value in playfully implying this categorization instead. His idea is bought cheaply – it does not even necessitate the purchase of the offending book – and its easy execution yields creative dividends that capture more column-inches in the media than a whole barrage of footwear ever could. Ultimately, however, this investment theory is just as lacking in technical details about how its creative investments are to computationally realized.

Each computational theorist in our field has his or her own favored theory of creativity. Yet, when it comes to building experimental systems that are designed to test

our ideas and impress our peers, theory inevitably takes a back-seat. For theories still serve a largely meta-level role in our field, providing a guiding philosophy but little real technical detail. So while they often serve as useful glosses for what we are trying to achieve, they rarely provide us with the computational tools to get the job done. As a field, we have more theories to choose from than real, large-scale, knowledge-based resources. If a computationalist were to candidly describe the methodology that is actually employed in practical creativity research, the description would sound very much like Booth's admission to *The Telegraph*: "using the materials given to me ... I just put them together". Fittingly, this observation is nicely borne out by the papers selected for this current special issue on computational creativity. Since these papers are also representative of the submissions to the 2008 International Joint Workshop on Computational Creativity (an event which has now assumed full conference status), we hope you see them as signs that this yearly event is much more than a talking shop for theoretical musings, and that the field of computational creativity is in rude health.

This special issue contains seven papers in all. While we have sought to tip the balance of our selection toward practical attempts at concrete solutions, high-level considerations about the nature of creativity and the optimal direction of enquiry within our field are a necessary part of any well-balanced discussion. Thus, philosophical and methodological issues are the primary focus of two of the seven papers, while the remaining five focus on empirical questions in computational implementations of one kind of creative act or another. Yet each paper, in its own way, manages to marry the philosophical with the computational to provide a worthwhile perspective on creativity.

We start with an intriguing paper from Kyle Jennings, which addresses a key rhetorical challenge at the heart of computational creativity head on: how do we persuade our peers in other fields, as well as eventual end-users, that the apparent creativity manifest in our software's actions are indeed the creative outputs of the software itself, and not the hard-coded unfolding of the programmer's own creativity? Jennings provides a two-pronged response to this constant and inevitable challenge: our programs must not only evaluate their own creations for themselves (rather than push this tedious task onto the unfortunate user), but be responsible for establishing their own set of evolving

standards for what is good and what is not over time, much like a human innovator.

Our second paper, by Jamie Forth, Geraint Wiggins and Alex McLean, attempts to bridge the divide between a formalized model of the popular view of creativity as the exploration (and, sometimes, transformation) of conceptual spaces with a geometric account of conceptual spaces that has been developed outside the field of creativity research, namely Gärdenfor's theory of conceptual spaces. The authors show that the match between the two accounts is based on more than a fortuitous overlap in the names given to each. Indeed, choosing Music as their domain of inquiry, the authors show how Gärdenfor's theory can be used to practically unite the symbolic and sub-symbolic levels of meaning and representation that must be negotiated by any creative musical agent.

As anyone who has endured the tedious process of buying furniture can attest, the factors influencing a purchase go well beyond price, colour and material. Our next paper, by Rafael Pérez y Pérez, Alfredo Aguilar and Santiago Negrete, describes a computer model for the arrangement of furniture. While this is a mundane human task where one might see little scope for creativity, one can also live in a house or apartment for years and never be entirely happy with the layout of the furniture. Indeed, it may well be that this kind of unassuming task – which requires a non-showy and non-obvious micro-creativity – is precisely the kind of task at which creative computers may achieve their most lasting impact. Interestingly, these authors show that the arrangement task benefits from the same engagement-reflection model of creative thinking that has been used in more obvious forms of machine creativity, such as the generation of novel stories.

In fact, the next papers in this issue are squarely on the topic of creative story generation. The first of these, by Federico Peinado, Virginia Francisco, Raquel Hervás and Pablo Gervás, focuses on the issue of novelty in creative stories and how best to evaluate it when the creator is a computational system. Humans, of course, have an intuitive feel for novelty and value, but our challenge as computationalists is to tap into these intuitions and articulate them in an explicit form that is computationally felicitous. These authors study the domain of computer-generated folk-tales, and distill their empirical findings into a set of metrics that can be used to meaningfully assess the originality of the outputs of any kind of computational story-telling system.

We earlier observed that the view of creativity as an opportunistic exploration of a conceptual space is only truly useful when the space of concepts can be formally defined in real, computational terms. Our next story-oriented paper, by Mark O. Riedl, takes this observation to heart, and shows how real computational value can be squeezed from the exploration perspective in the automated generation of stories. Riedl's implemented system is capable of searching a space of partial story solutions because it is guided by combinations of existing story fragments called *vignettes*. To avoid generating simplistic patchworks of canned elements, or *Frankenstories*, Riedl shows how this stock of pre-packaged structures is adapted and transformed anew in each story-telling context.

The exploration of conceptual space is also the theme of our third story-oriented paper in this issue, by Carlos León and Pablo Gervás. Creative computational systems do not accidentally happen upon good stories in a conceptual space, but are guided in their exploration by a pre-defined sense of quality that allows them to choose one path over another, or to favor one area of a space while dismissively avoiding another. Once again, the value of the exploration perspective is only as good as the quality metrics that are used to drive it in a real computational system. To generate stories of quality, a system must be capable of ruthlessly rejecting those candidates in which quality is sorely lacking, so that it may concentrate its limited energies on more promising candidates. These authors tease apart the various dimensions that comprise story quality, from tension to empathy, and employ them as variables in concocting different evaluation functions for the exploration process.

Thus, whether an agent is human or computational, one cannot be creative unless one knows how to appreciate creativity. The last paper in this special issue diverges from the topic of story generation to offer a final ironic sting in this vein. Focusing on what makes an ironic comparison creative, Yanfen Hao and Tony Veale provide a computational analysis of the irony evident in novel self-contained similes that have been automatically harvested from the world-wide-web. Irony is a difficult problem for even humans to tackle, and these authors show that the phenomenon is more pervasive in online comparisons than one might think. Veale and Hao identify markers of varying reliability for sardonic and ironic intent, and show how these can be used to classify similes into ironic and non-ironic categories with high levels of accuracy. Few would argue with the

authors' contention that the ability to recognize irony is a necessary precursor to the meaningful generation of irony by a computer that truly appreciates what it is doing.

To summarize, philosophical considerations are an important ingredient of any study of human and machine creativity, but in computational creativity they must be more than mere window dressing: they must be forced to earn their keep in ways that contribute to the computational bottom-line. The papers in this special issue do not espouse a common philosophical line, which is to be expected (indeed, desired) when so many philosophical approaches are available to the theorist-practitioner. All papers are united by a common *meta*-philosophy, however. To borrow the title from a recent Woody Allen film, the key to making progress in computational creativity is a belief in “Whatever Works”. Philosophy is important, but the ultimate test of success in our field is a computational one. In the words of the 17th-century philosopher G. W. Leibniz, let us compute!

Tony Veale, Pablo Gervás, Rafael Pérez y Pérez

Dublin, Madrid, Mexico City

September, 2010