

Leaps and Bounds:

An Introduction to the Field of Computational Creativity

Tony Veale,

School of Computer Science, University College Dublin, Ireland.

Tony.Veale@UCD.ie

Rafael Pérez y Pérez,

Universidad Autónoma Metropolitana, Cuajimalpa, Mexico.

Rperez@cua.uam.mx

Abstract

Computers have enhanced productivity and cost-effectiveness in all of the creative industries, and their value as tools is rarely doubted. But can machines serve as more than mere tools, and assume the role and responsibilities of a co-creative partner, or even become goal-setting, autonomous creators in their own right? These are the questions that define the discipline of computational creativity. The answers require an algorithmic understanding of how humans give meaning to form, but a transformation in the way we think about creativity is unlikely to occur in a single bound. Rather, interdisciplinary insights from diverse fields must first inform our models, and shape a narrative of creativity in which machines are both tools and creators. To set the stage for the newest work, this introduction to the special issue on computational creativity shows where the field is going, and where it has come from.

1. Introduction

The historical sweep of human innovation has rarely been captured so concisely as in Stanley Kubrick's film, '2001: A Space Odyssey'. In what it is perhaps the most dramatic jump cut in all of cinema, Kubrick cuts from a tumbling bone, recently weaponized by an ancient hominid at the vanguard of an evolutionary leap, to a gently orbiting space station of the future. In doing so, he captures an essential truth about creativity: creators rely crucially on their tools, and so the steady evolution of human creativity has occurred in lockstep with a parallel evolution in human tools. Whether our tools are used for good or ill, they make something possible that was previously difficult or impractical to achieve. As we embed ever more of our creative processes into our tools, these processes – as well as our notions of creativity itself – are also free to grow and evolve. So we now find ourselves, albeit decades behind schedule, at a similarly transformative juncture in human creativity as imagined in Kubrick's film: for we now transfer so much of the creative burden onto our computers and software systems that our tools may soon deserve to be viewed as creative agents in their own right.

We must thus consider the creative role of computers from a dual perspective: the machine as an instrument of human creativity; and the machine as an agent of its own creativity. Despite advances in Artificial Intelligence, the instrumental view remains dominant, as, ironically, such advances can lead us to prize those areas in which we humans still excel, such as art and science, and to see them as a distinctly human preserve. However, Computational Creativity is an emerging field at the nexus of artificial intelligence (AI), cognitive science, art, design and philosophy, that sees creativity as a fundamentally algorithmic activity, one that computers can pursue autonomously or semi-autonomously, not as tools but as creators and co-creative partners. In this way, this new field, called CC for short, aims to bridge the competing viewpoints by showing how tools may evolve into tool users, and thus realize Countess Ada Lovelace's hopes for a poetical science.

The collected papers of this special issue provide a recent snapshot of work in the field. However, no work in this field can be entirely unpicked from any other, nor easily detached from the history of what has gone before. So, from the issue's most strongly represented theme, story generation, to its most expansive, a

general framework for design creativity, readers will find interwoven strands of the field's most pressing concerns and contemporary perspectives, such as conceptual blending, visual semantics, scientific discovery in very large data sets, embodied performance, and a methodology for studying and applying automated creativity in the wild. We shall return to the specifics of each paper in section 6, after having first laid a solid foundation for our discussion. We begin our whistle stop tour of the field in section 2, where a consideration of computers as tools gives way, in section 3, to a consideration of machines as creative agents in their own right. The need for inspiration, and a look at how machines might find it for themselves, is the subject of section 4, while section 5 provides some illustrative examples of the real systems that embody these various concepts and principles. The papers of the special issue are situated relative to those key ideas in section 6, before this introduction concludes with a prognosis for the field in section 7.

2. The Computer as Creative Instrument

Human creativity in the digital age is increasingly facilitated by richly-featured software tools that support rapid experimentation via trial and error, and the growing market for creativity support tools has not gone unnoticed by the field of creativity research (see Lubart, 2005; Shneiderman, 2007). Where once an artist might doodle in a notebook, painters like David Hockney now extoll the virtues of the digital tablet. For decades designers have used professional CAD (computer-aided design) tools to design everything from cars to buildings, but the distinction between professional- and consumer-grade tools is now blurring with significant increases in the power, and dramatic decreases in the cost, of personal devices. Smartphones with all the functionality of a laptop computer come complete with in-built cameras and audio capabilities as standard, creating a market for casual creativity on the go. When users are empowered to create and share content via a multitude of apps, the market incentivizes developers to provide more advanced tools for filtering, editing and generally transforming reality as mediated by our digital devices. Yet not all modalities are equal. Most smartphones offer tools for the manipulation of sounds, images and video, with capabilities that were once the preserve of professional studios. Yet it says

something about the brittle nature of language that comparable consumer-grade tools have not yet been developed for the creation and co-creation, of novel texts.

Indeed, creativity support tools for writers have made little headway beyond digitalized versions of old reliables such as the thesaurus, the style guide and the grammar book. These tools suggest word substitutions or changes in word-order but rarely help writers to be more creative, by suggesting edits that would make a text wittier, more persuasive, or more memorable. To be sure, image-editing tools are just that, tools, and have no creativity of their own. Nonetheless, such tools – comparable applications exist for audio and video too – offer a broad palette of features to encourage creative experimentation via cycles of rapid generation and facile manipulation of novel forms. By comparison, our productivity tools for text creation offer a much-impoverished slate of features. This is changing, with the advent of web-scale language models (Radford et al., 2019) that can be repurposed for a wide range of language generation tasks. Yet the tool perspective, or what we might call the ‘weak CC’ stance, is still capable of generating insights into the nature of human creativity. The availability of tools and their relative capabilities across modalities tells us something of value about what, how and why we humans create. To go deeper still, we will need to explore truly autonomous machine creativity, or what we can call the ‘strong CC’ stance.

3. The Computer as Creative Agent

Like AI, CC is both an engineering discipline and a philosophical stance. It aims to shed light on human creativity by treating the mind as an embodied information-processing machine, but it also sets out to give algorithmic form to psychological findings about the workings of creativity in humans. For practical purposes, it does not seek to define ‘creativity’ itself, as moving targets rarely offer reliable directions, but assumes the folk, human-centric notion of common parlance, with all its cultural and historical baggage. As AI researchers once did for intelligence, strong CC researchers seek to endow machines with the ability to achieve ends that would be deemed creative if achieved by humans (Colton & Wiggins, 2012), perhaps in ways that would be equally to the credit of a human. These ends are more generative than analytical, as CC emphasizes the creation of novel digital

artefacts in various modalities, whether text (e.g., poems, stories, jokes, lyrics), image (e.g., paintings, designs, fonts, logos) or sound (music, speech), as well as multimodal artefacts that integrate all three (e.g., computer games, interactive fiction). Human assessments of these artefacts offer an empirical yardstick for the methods and theories that guide their construction, though CC researchers are careful to resist calls for a creative Turing Test in which humans are asked to discern the synthetic from the handmade. Turing's (1950) test, now a staple of pop culture, promotes a race to the surface, pushing engineers to use guile rather than theory to fool the eye and the ear. This point deserves italics: *Strong CC does not set out to fool humans but to better understand human creativity, by turning formal ideas of novelty and value into working algorithms and new artefacts.*

Following Guilford (1950), CC offers an algorithmic view of divergent thinking that models decision-making as a search in a conceptual space. Where CC differs from AI's use of the metaphor (Newell, Shaw & Simon, 1959) is the specification of the goal state that is sought in the space. AI typically adopts a convergent, problem-solving view of search in which just a few states represent optimal or near-optimal solutions to a problem, so a good AI system must converge on these goal states as quickly and reliably as possible. In contrast, CC imagines a space to contain a great many goal candidates of varying suitability relative to our quality metric or objective function (Boden, 1990). CC systems value diversity as much as quality, so they may not yield the same results every time. A CC system may seek out areas of a space that have yet to be explored by humans or machines.

Search in a state space is a flexible abstraction, as different configurations of a space or an objective function can model competing views of the creative act. A CC system might, for instance, search two spaces in parallel: the space of possible artefacts and the space of objective functions, using the best candidate from the latter to rank the preferred candidates in the former. A system might also search the space of possible formulations of spaces, to search at the meta-level for the optimal space (a style, genre or domain) in which to search for the most novel or interesting artefacts. A search can explore a space of possible artefacts or a space of combinations of artefacts. Alternately, each state may correspond to one step in a creative solution, so the path followed by a search thru the space is itself the

solution. Paulos (1982) used catastrophe theory to visualize the shape of search spaces most suited to joke analysis and generation, noting that such spaces have deliberate dead-ends or discontinuities to trip up an audience. Only a joke teller knows the safest paths through this warped space. We can also model this kind of adversarial search with two competing objective functions – one reflecting the status quo, the other the system’s own insights – so as to implement Sternberg & Lubart’s (1996) investment theory of creativity. By finding states that maximize the differences between both functions, a CC system can ‘buy low and sell high.’

Search offers a low-level computational perspective that accommodates high-level views of various creative strategies. For instance, Koestler’s (1964) theory of bisociation requires an agent to search the overlap of two spaces, while Lakoff & Johnson’s (1980) theory of metaphor, much like Gentner’s (1983) theory of structural analogy, necessitates the search for a mapping between two spaces, representing the source and target domains. Fauconnier & Turner’s (1998, 2002) theory of conceptual blending may posit additional spaces, but the optimality constraints and principles that guide the construction of a well-formed blend likewise require a search in the space of all possible blend configurations. It may seem reductive to view everything as search, but researchers do this not just for philosophical reasons, but for the real traction it offers in building CC systems.

Search in a state-space supports a multitude of algorithm types, from genetic algorithms to artificial neural networks (whose training process exploits search) to *good old fashioned AI* (GOFAI) methods. Each instantiates the basic metaphor in its own way. A greedy search starts from a single, often random, artefact that a CC system works to improve in small strides relative to its objective function. A non-greedy search instead creates a pool of random artefacts that are developed in tandem, perhaps using crossover to generate new candidates from old. These approaches embody Boden’s (1990) concepts of exploratory and combinatorial creativity if the makeup of the space and its objective function remain constant. If the shape of a space can shift during the search, or if the objective function is allowed to evolve, then what results is closer to Boden’s idea of transformational creativity. While transformation is simple in theory, it is complicated in practice by the need to align a CC system’s evolving view of a space with the aesthetics of

the users for whom it is constructed. For, although a machine need not pursue a cognitively-plausible theory of human creativity, the artefacts that it generates must ultimately conform to a human view of what is novel and useful to a user.

4. Sources and Kinds of Computational Inspiration

Even a machine that is largely autonomous in generation will create artefacts for human consideration and consumption, for we have, as yet, no need of machines that create only for each other. The values of the human audience are reflected in a CC system's objective function, which a developer might specify explicitly as a multivariate formula or, alternately, require the machine to learn for itself from an 'inspiring set' of human exemplars (Ritchie, 2007). In fact, a machine will need a number of such sets if it is to adequately model novelty and usefulness. For, to determine whether a digital artefact is original, or to measure its novelty and surprise value, a machine will need the cultural memory of past acts of creativity that we humans take for granted. So Boden (1990,1998) distinguishes P-Creative from H-Creative generation: a result is P- or Psychologically-Creative if it is novel to the agent that generates it, regardless of its originality in any historical sense; in contrast, the first-ever generation of a result is H- or Historically-Creative. A CC system is demonstrably P-Creative to the extent that it reproduces known results for itself – e.g., if it can derive established theorems or laws – but even P-not-H creativity can serve as a promissory note for a system's generative reach and future abilities. Naturally, we prefer it when our CC systems are capable of truly H-Creative results too, even when these amount to novel reformulations of prior results (e.g. a new way of proving theorem T or of achieving result R).

The P/H distinction is a useful, and oft-cited, idea in computational creativity, but it is no less reductive than the creativity-as-search perspective. It lacks, for instance, a socio-cultural dimension, and assumes a simplistic, one-track view of history. Moreover, while P vs. H feels like a binary distinction, each must exist on a gradient. Tiny changes to an exemplar can yield a result that is strictly new, yet we are reluctant to label even those that are deemed patentable as H-Creative. A generative system such as David Cope's *Experiments in Musical Intelligence*, or EMI (Cope, 2006) must have Bach in its bones, in the very definition of its search

spaces and objective functions, before it can produce new Bach-like cantatas of its own. But as impressive as a machine's outputs may be, they will fall short of the exemplars in its inspiring set, at least with regards to originality, if those exemplars are merely used as templates for future, recombinant generation. In the ideal case, a CC system can extract the unifying essence of its inspiring set, and transfer it into novel results of its own that chime with, but do not overly resemble, its inspirations. In effect, a CC system capable of deliberate originality must possess different objective functions for similarity and dissimilarity. If a single distance function is used for both purposes, a system will find it hard to identify outputs that both resemble its exemplars and significantly diverge from them too. All too often, what results is a machine pastiche of a human exemplar that is strictly new in a binary sense but unduly similar to what has gone before.

CC theorists refer to a related point of criticism as *mere generation*. A system that uses templates or scripts to generate artefacts with the formal properties of a given class – to e.g. produce poem-shaped texts or cantata-shaped sounds – is merely generative if it lacks a capacity for self-critique and the ability to rank and filter its own well-formed outputs by quality. Such systems may be autonomous generators of valid forms, but they require a human in the loop to attach value to their outputs. Mere generation and pastiche often work hand in hand, as each exploits a shallow model of the domain. A pastiche generator limits its search to the locality of an inspiring exemplar, using formal criteria to explore neighboring states. Lacking any sense of the value of these states, it relies on closeness to the exemplar to guarantee the meaningfulness of any new variants. Mere generation hollows out or abstracts templates from its inspiring exemplars, in the hope that much of the meaning of the original attaches to the template and to the variants produced with it. But it is important to note that these criticisms should not only be levelled at our machines. Pastiche and mere generation are often exploited by humans too, to produce artefacts with just a veneer of originality. Our 'creative' machines usefully expose the tacit assumptions we make about human creators. So when we criticize them, it is not so much for falling short of the best human creators but for aiming for, and reaching, the low bar set by the worst.

A generative system needs a sense of form and a sense of how form conveys

meaning if it is to yield more than pastiche or blind variations on past successes. To achieve a common type of human creativity that Giora *et al.* (2004) label an ‘optimal innovation’, an agent must be able to predict the change in meaning that arises from a change in form. An optimal innovation in language is any modified phrase that allows an agent to retrieve the familiar form from which it diverges, such as an idiom, cliché, title or slogan, and to calculate the difference in meaning induced by the modification. Puns and related modes of wordplay fall into this category, such as Giora’s example of ‘weapons of mass distraction.’ The familiar phrase, ‘weapons of mass destruction’ (or WMD) was a constant refrain during the second Gulf war, and the change, ‘destruction’ → ‘distraction’, achieves a large shift in meaning (military → political) for a relatively small shift in form. An innovation is optimal if it transcends mere generation and knowingly leverages form to magnify meaning (Veale, 2012). Optimal Innovation theory is thus an attractive foundation for CC systems that aim to create novel, meaningful and witty variants of their inspiring sets with a modest set of change operators.

5. CC In Thought and Action

For a relatively young discipline, CC has amassed a significant body of research literature and system-building knowledge. A recent collected volume, *Readings In Computational Creativity* (Veale & Cardoso, 2019), samples from two decades of CC’s philosophical, engineering and scientific developments. Readers may also find the proceedings of the annual ICCO event – the International Conference on Computational Creativity – and its associated workshops to be a useful source of information about the methods and emphases of this rapidly evolving field.

For those seeking a formalization of the metaphor that underpins search in a conceptual space, Wiggins (2006) offers a formal synthesis of work ranging from Newell, Shaw & Simon (1959) to Boden (1998). Wiggins’ formalism remains popular with CC theorists seeking a common framework in which to situate their own uses of search, state spaces and objective functions. Another staple of the CC literature is the issue of evaluation: how exactly are we to judge the utility and value of a system’s outputs, or measure its capacity for human-scale creativity, or tease apart what a developer puts into a system from what a system puts into its

outputs? Jordanous (2012) proposes a standardized approach to evaluation for CC systems, while Jordanous & Keller (2016) conduct a corpus analysis of the creativity literature to identify the principal dimensions that such systems and evaluations should emphasize. Ritchie (2007) explores the most durable criteria for assessment, whilst Colton (2008) argues that we humans apply a double standard to machine-generated artefacts, making it necessary to discriminate between criteria that impart a perception of creativity from those that actually distinguish the creative act. Colton, Charnley & Pease (2012) go further, and offer two complementary frameworks, IDEA and FACE. IDEA describes creativity in terms of generative acts, while FACE attempts to capture the impact of those acts. Taken together, machine creativity is defined on a scale of generation and impact. Grace *et al.* (2014) affirm the importance of domain data to evaluation, and focus on what it is that makes the design of a new product novel, useful and surprising.

The CC literature spans all of the obvious domains of human creativity, from art and science to music to language. Veale (2012) explores linguistic creativity from a computational perspective, by modeling creativity as a lever that yields disproportionate changes in meaning for relatively minor changes in form. Stock & Strapparava (2003) also explore the value of small surface changes with big impacts on meaning in their *HAHAcronym* humorous acronym generator. Pun generation similarly leverages new meanings from old forms, and the first CC pun generator to be robustly evaluated by true pun lovers, children, is described in Binsted, Pain & Ritchie (1997). More complex textual forms, from metaphors to stories and poems, have also been the focus of CC research and development. Veale & Cook (2018) provide a comprehensive tour of tools and techniques for the building of creative linguistic systems that operate on social media as “bots.” Gervás (2001) used an established AI technology, expert systems, to encode the knowledge needed to generate formal poetry, while Manurung *et al.* (2012) used genetic algorithms to search the space of poetic possibilities (in a take on the Oulipo movement (1988) and its goals) so as to evolve strong poems from weak beginnings. Veale, Shutova & Beigman Klebanov (2016) conduct a wide-ranging survey of computational systems that analyze, interpret, paraphrase or generate metaphors of all kinds, whether in everyday language, newspapers, or poetry.

CC work on narrative takes its cue from structuralists such as Propp (1968). His ideas also inspired human writers seeking a systematic means of production, such as Cook (1928), who, in retrospect, now look like proto-CC thinkers. Turner (1994) explicitly bridged AI and creativity via story-telling, while Riedl & Young (2010) used an AI model of planning to generate satisfying plots, thus taking an AI forte, convergent problem solving, to the meta-level. Pérez y Pérez & Sharples (2001) turned a cognitive model of the writer into an algorithm for narrative generation that cycles through modes of reflection and generation, while Pérez y Pérez (2015) outlines a related model for the collaborative generation of stories. Aguilar & Pérez y Pérez (2015) abstract this model of reflection and generation to view our early cognitive development as an inherently creative process. Veale & Valitutti (2017) generate comic narratives by forcing familiar but ill-paired characters into familiar plot forms, while Victor Raskin, one of the first to bring AI methods to bear on narrative jokes, collects much relevant research in Raskin (2008). Following also in the footsteps of the early structuralists, Gervás (2009) surveys different lines of CC work on narrative and their driving forces.

In seeking to build domain-specific generators of human-scale creativity, CC theorists often build concrete algorithmic realizations of theories from the wider creativity literature. Koestler's Bisociation, for instance, is the basis of a scientific discovery system from Petrič *et al.* (2012), who look for a bisociative overlap in the statistical similarities of very different scientific texts. Fauconnier & Turner's theory of conceptual blending, an evolution of Koestler's position, has proven to be especially popular with CC theorists and system builders. Veale & O'Donoghue (2000) synthesized a range of AI methods, from spreading activation to structure mapping, to build an efficient implementation of blending in conceptual spaces. Pereira & Cardoso (2006) and Pereira (2007) drew on additional AI techniques, including genetic algorithms, to faithfully model the constraints and principles of Fauconnier & Turner's theory. Other CC systems create blends in more implicit ways, such as via the learning mechanisms of an artificial neural network that is designed to fuse multiple sources and kinds of data (e.g. see Heath *et al.*, 2014).

For a sampling of the diversity of work in CC on musical and artistic creativity, readers are directed to Veale & Cardoso (2019). Worthy of special mention are

the Aaron system of Harold Cohen (see McCorduck, 1991), the NEvAr system of Machado & Cardoso (2002), and the Painting Fool of Colton (2011). Each define different objective functions and search spaces, to produce paintings that range from traditional figures in scenes to mathematical landscapes to portraiture. The Painting Fool incorporates a model of affect to give its digital painter an artificial personality of a sort, while NEvAr uses genetic search at two different levels simultaneously, evaluation and generation, to evolve both an image and an objective function that best suits the tastes expressed by its users. With regards to music, in addition to Cope's EMI (2006), readers are directed to Pasquier *et al.* (2016) for an overview of a branch of CC known as musical meta-creativity. On the spectrum of weak to strong computational creativity, meta-creativity floats freely in the middle. It allows humans to express themselves creatively through the building of artificial systems that may themselves shoulder some non-trivial creative responsibilities. Insofar as humans wish to remain firmly in the loop during machine forays into creativity, a process of meta-creation that allows humans and machines to work as co-creative partners seems the most promising vector for theoretical and commercial progress in computational creativity.

6. On The Spectrum: The collected papers of this special issue

As a discipline, computational creativity fosters work that runs the gamut from philosophical inquiry and psychological theorizing to the practical concerns of AI system-building. A definition that embraces all this and more will be necessarily vague. Consider that of Colton & Wiggins (2012), for whom CC is "the study and simulation, by computational means, of behaviour natural and artificial, which would, if observed in humans, be deemed creative." This is a definition of CC that piggybacks on the prevailing social view of human creativity, so that it can evolve with it, yet it astutely avoids any commitment to what it might actually mean. For as much as CC seeks to understand the workings of human creativity, it cannot dictate what creativity means to *us*. Rather, it can explicate the workings of what *we* take it to mean, while showing that certain myths are overly-simple or naïve. In this spirit, Pérez y Pérez (2015) defines CC as the interdisciplinary study of the creative process that uses computational means to reflect on existing knowledge, so as to generate new insights and knowledge on which future work can reflect.

Pérez y Pérez (2018) maps this explicatory role onto a continuum of research: at one pole we find formal or mathematical models of creative behaviour, usually engineered as real systems; at the other pole sits our cognitive and social models of human creativity, and views of what they imply for a generative machine. No work in CC is so singular of purpose that it sits at one pole or another. Rather, each has its own emphases that nudge it this way or that. So, those closer to the cognitive pole seek general insights and social understanding, while those closer to the engineering pole seek to generate real outputs that a human might deem “creative.” Between the two poles we find CC systems that seek to do both: to generate their own outputs while shedding light on human processes. These systems strive for novelty, with outputs that deviate from the norm, but they also strive for value, with outputs that demonstrate utility for an audience. In the end, this what it means for any agent, computational or otherwise, to be creative. It falls to individual researchers, in their specialist domains, to explore their own computational correlates of novelty and value within their chosen frameworks.

The seven papers that make up this special issue exemplify different points on the CC continuum. Consider the rump of three papers that concern themselves with narrative. Each explores much the same mode of human creativity, yet each situates itself on a slightly different part of the continuum. The contribution of Wicke and Veale, for instance, focuses on the performance of a novel story by embodied artificial agents. Although their system, named *Scéalability*, necessarily contains a generative component for the plotting, characterization and rendering of stories, their focus is on physical embodiment. We humans bring our stories to life when we relate them, and this work taps into the same cognitive structures – image schemas and conceptual metaphors, to allow robots to do the same. In the contribution of Concepcion, Méndez, Gervás and Tapscott, it is the engineering of new plots that takes centre stage. This work draws its inspiration from a study of the literary process, to give machines the same basic ability as human writers to weave multiple plot lines into a single, multifocal narrative. León, de la Torre and Gervás complete this triad by taking a step back, to derive some quantifiable insights about the narrative process by looking at how we humans write our own short stories. These three papers represent different stages of a CC investigation: empirical inquiry, exploratory development guided by practice, and the framing

of the results in ways that appeal to the cognitive predispositions of an audience.

Two other papers in the special issue reside to the right of this triad on the CC continuum, as it extends toward its cognitive/cultural pole. The contribution of Pérez-Romero and Aguilar concerns itself with the architecture of CC systems that engage in design tasks. They propose the CCDSF, or *Computational Creative Design System Framework*, as a schematic means of understanding any CC model in this space. Their proposed schema is layered, expansive, and responsive to the needs of systems in diverse domains. Responsiveness is key, for it is all too easy to fall into the trap of dictating how other systems, or other practitioners, should approach their work as a CC task. Rather, researchers must immerse themselves in the concerns of the field they aim to model, to adapt their system from within rather than imposing constraints from without. So that researchers can approach a domain in the spirit of those doing field work, a methodology for doing just that is offered by Pérez y Pérez and Ackerman in their contribution to this issue.

The last two papers in this issue sit much closer to the engineering pole of the continuum, although these too draw some general insights from specific systems. The first, a contribution from Cunha, Lourenço, Martins and Machado, presents a case-study in conceptual blending with a specific application to visual creativity. These authors explore the space of compositional emoji, to generate new images from old for words and ideas that fall outside the scope of the Unicode standard. Conceptual blending is a compositional framework that lends itself particularly well to combinatorial creativity, and owes much to the theory of bisociation first introduced by Koestler (1964). It is upon this foundation that the contribution of Lavrač, Martinc, Pollak, Novak and Cestnik builds, to support a model of scientific discovery in the academic literature. A bisociation, much like a creative blend, is an integration of unlikely parts that, at first blush, seem to have little to do with one another. It is the unlikeliness of the combination that yields novelty, while it is the meaning of the whole – and the new insight it brings – that yields its value.

With the partial exception of this last contribution, the papers of this special issue exemplify a largely symbolic approach to computational creativity. While statistical reasoning has a role to play in any system that quantifies novelty and value, the symbolic approach is characterized by its ontological commitment to

explicit schematic structures, often conceptual in nature, and to its imposition of top-down interpretative processes that accord with our understanding of what is happening under the hood. As AI raises its stake in large-scale, data-driven, and often opaque bottom-up approaches, encouraged by the impressive performance of statistical models – for instance, Christie's auction house in New York recently sold an imaginary portrait created by an artificial neural network, titled *Edmond de Belamy*, for \$432,500 – CC is surely following suit. This CC special issue may be the last in which symbols dominate, yet strands of this old-school AI approach will persist as long as they continue to add value, to CC theories and CC systems. Only a blend of old- and new-school approaches will allow the next generation of CC systems to denote more than a point on our continuum, and capture instead a sweeping interval that jump cuts from engineering reality to theoretical insight.

7. Prognosis and Conclusions

The usefulness of computers as tools for the furtherance of human innovation is well attested by their widespread adoption by the creative industries. Modern software tools do much more than the uninspired drudgework of creativity, and work hand-in-glove with creators to drive innovations that might not be feasible or cost-effective by either working alone (Shneiderman, 2007). Where machines excel is in the exploration of large state spaces, such as those that correspond to the most divergent of problems (Guilford, 1950), which is to say, problems that benefit from casting the widest possible net in the search for solutions, or from using multiple metrics to assess novelty and value. We are already halfway to viewing computers not as drudges and mere tools but as partners in the creative act (Lubart, 2005). At present, this partnership assigns creative responsibilities on the basis of comparative advantage, with each partner doing what each does best. Machines are facile manipulators of form, and can follow the rubrics that have been set for them, or, indeed, that they may have learnt for themselves, to explore the combinatorial possibilities of a given set of rules and constraints.

Humans, in contrast, excel at meaning, and at placing form in its semantic, cultural and historical context. Working together, a partnership of human and machine can reflect on convention to depart from familiar norms, and deviate

from familiar forms, to deliver an unexpected jolt of fresh meaning. For machines to evolve into full partners in the creative act – to become equal co-creators with humans – they will need to be as facile in the attribution of meaning and context as they are in the generation of surface content. This is the key challenge that confronts the algorithmization of human-level creativity. It is a one that readers will see played out in each of the collected papers of this special issue.

References

- Aguilar, W. & Pérez y Pérez, R. (2015). Dev E-R: A computational model of early cognitive development as a creative process. *Cognitive Systems Research*, 33:17–41.
- Boden, M. (1990). *The creative mind: Myths and mechanisms (second edition)*. Routledge, London, UK.
- Boden, M. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, 103, 347–356.
- Binsted, K., Pain, H., and Ritchie, G. (1997). Children’s evaluation of computer-generated punning riddles. *Pragmatics and Cognition*, 52, 309-358.
- Colton, S. (2008). Creativity versus the perception of creativity in computational systems. In Papers of the 2008 AAAI Spring Symposium on Creative Intelligent Systems, Technical Report SS-08-03, Stanford, California, USA, March 26-28.
- Colton, S. (2011). The Painting Fool: Stories from building an automated painter. In McCormack, J. & d’Inverno, M. (eds.) *Computers and Creativity*. Springer Verlag, Berlin.
- Colton, S. and Wiggins, G. (2012). Computational creativity: The final frontier? In de Raedt, L., Bessiere, C., Dubois, D., & Doherty, P. (eds.), *Proceedings of ECAI Frontiers in Artificial Intelligence and Applications*, Paris, France, pp 21-26.
- Colton, S., Charnley, J., and Pease, A. (2012). Computational creativity theory: The FACE and IDEA descriptive models. In Proceedings of the 2nd International Conference on Computational Creativity, ICC 2012.
- Cook, W.W. (1928). *PLOTTO: The master book of all plots*. Battle Creek, Michigan:

Ellis Publishing Company.

Cope, D. (2006). *Computer models of musical creativity*. Cambridge: MIT Press.

Fauconnier, G. and Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, 22(2):133–187.

Fauconnier, G. and Turner, M. (2002). *The way we think. Conceptual blending and the mind's hidden complexities*. New York, NY: Basic Books.

Gentner, D. (1983). Structure-mapping: A theoretical framework. *Cognitive Science*, 7(2), 155-170.

Gervás, P. (2001). An expert system for the composition of formal Spanish poetry. *Journal of Knowledge-Based Systems*, 14(3–4), 181–188.

Gervás, P. (2009). Computational approaches to storytelling and creativity. *AI Magazine*, 30(49–62).

Giora, R., Fein, O., Kronrod, A., Elnatan, I., Shuval, N. and Zur, A. (2004). Weapons of mass distraction: Optimal innovation and pleasure ratings. *Metaphor and Symbol* 19(2):115-141.

Jordanous, A. (2012). A standardized procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4(3), 246–279.

Jordanous, A. and Keller, B. (2016). Modeling creativity: Identifying key components through a corpus-based approach. *PLOS ONE*, 11(10).

Grace, K., Maher, M. L., Fisher, D. and Brady, K. (2014). Data-intensive evaluation of design creativity using novelty, value, and surprise. *International Journal of Design Creativity and Innovation*, 3(3–4):125–147.

Guilford, J.P. (1950). Creativity. *American Psychologist*, 5(9):444–454.

Heath, D., Norton, D. and Ventura, D. (2014). Conveying semantics through visual metaphor. *ACM Transactions on Intelligent Systems and Technology*, 5(2), article 31.

Koestler, A. (1964). *The act of creation*. New York: Macmillan.

Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. Chicago, Illinois:

University of Chicago Press.

Lubart, T. I. (2005). How can computers be partners in the creative process: Classification and commentary on the special issue. *International Journal of Human-Computer Studies*, 63(4-5), 365–369.

Machado, P. and Cardoso, A. (2002). All the truth about NEvAr. *Applied Intelligence, Special Issue on Creative Systems*, 16(2):101-119.

Manurung, R., Ritchie, G. and Thompson, H. (2012). Using genetic algorithms to create meaningful poetic text. *J. Exp. Theor. Artif. Intell.*, 24(1):43-64.

McCorduck, P. (1991). *Aaron's code: Meta-art, artificial intelligence, and the work of Harold Cohen*. W. H. Freeman & Co.

Allen Newell, John Shaw, and Herbert A. Simon. (1959). *The processes of creative thinking*. Rand Corporation.

Oulipo. (1988). *Atlas de littérature potentielle*. Gallimard.

Pasquier, P., Eigenfeldt, A., Bown, O., and Dubnov, S. (2016). An introduction to musical metacreation. *Computers in Entertainment*, 14(2), article 2.

Paulos, J. A. (1982). *Mathematics and Humor: A study of the logic of humor*. Chicago, Illinois: University of Chicago Press.

Pereira, F. C. (2007). *Creativity and Artificial Intelligence: A conceptual blending approach*. Berlin: Mouton de Gruyter.

Pereira, F. C. and Cardoso, A. (2006). Experiments with free concept generation in Divago. *Knowledge Based Systems*, 19(7):459–471

Pérez y Pérez, R. and Sharples, M. (2001). MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental & Theoretical Artificial Intelligence*, 13(2):119-139

Pérez y Pérez, R. (2015). A Computer-based Model for Collaborative Narrative Generation. *Cognitive Systems Research*, 36-37, pp. 30-48.

Pérez y Pérez, R. (2018). The Computational Creativity Continuum. *Proceedings of the Ninth International Conference on Computational Creativity*, pp. 177-184.

Petrič, I., Cestnik, B., Lavrač, B. and Urbančič, T. (2012). Outlier detection in

- cross-context link discovery for creative literature mining. *The Computer Journal*, 55(1):47–61.
- Propp, V. (1968). *Morphology Of The Folk Tale* (second edition). University of Texas Press.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskeve, I. Language Models are Unsupervised Multitask Learners. *Open AI Technical Paper*, 2019.
- Raskin, R. (2008). *The primer of humor research*. Berlin: Mouton de Gruyter.
- Riedl, M. and Young, R. M. (2010). Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research (JAIR)* 39:217– 268.
- Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, 17(1), 67–99.
- Shneiderman, B. (2007). Creativity support tools: Accelerating discovery and innovation. *Communications of the ACM*, 50(12), 20–32.
- Sternberg, R.J. and Lubart, T.I. (1996). Investing in creativity. *American Psychologist*, 51(7):677-688.
- Stock, O. and Strapparava, C. (2003). HAHAcronym: Humorous agents for humorous acronyms. *HUMOR, the International Journal of Humor Research*, 16(3):297–314.
- Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59:433-460.
- Turner, S.R. (1994). *The Creative Process: A Computer Model of Storytelling*, Hillsdale, NJ: Lawrence Erlbaum.
- Veale, T. and O'Donoghue, D. (2000). Computation and blending. *Cognitive Linguistics*, 11(3-4), 253-281.
- Veale, T. (2012). *Exploding the creativity myth: The computational foundations of linguistic creativity*. London, UK: Bloomsbury.
- Veale, T., Shutova, E. and Beigman Klebanov, B. (2016). *Metaphor: A Computational Perspective*. Morgan Claypool, Synthesis Lectures on Human Language Technologies.
- Veale, T. and Valitutti, A. (2017). Tweet dreams are made of this: Appropriate

incongruity in the dreamwork of language. *LINGUA* 197, 141--153.

Veale, T. and Cook, M. (2018). *Twitterbots: Making Machines that Make Meaning*. Cambridge, MA: MIT Press.

Veale, T. and Cardoso, A. (eds.) (2019). *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*. Springer.

Wiggins, G. (2006). Searching for computational creativity. *New Generation Computing*, 24(3):209–222.