Systematizing Creativity: A Computational View

Tony Veale, Amílcar Cardoso and Rafael Pérez y Pérez

Abstract Creativity is a long-cherished and widely-studied aspect of human behavior that allows us to re-invent the familiar and to imagine the new. *Computational Creativity* (or CC) is a recent but burgeoning area of creativity research that brings together academics and practitioners from diverse disciplines, genres and modalities, to explore the potential of our machines to be creative in their own right. As a scientific endeavor, CC proposes that computational modeling can yield important insights into the fundamental capabilities of both humans and machines. As an engineering endeavor, CC claims that it is possible to construct autonomous systems that produce novel and useful outputs that are deserving of the label "creative". The CC field seeks to establish a symbiotic relationship between these scientific and engineering endeavors, wherein the artifacts that are produced also serve as empirical tests of the adequacy of scientific theories of creativity. We argue that, if sufficiently nurtured with volumes such as this, the products of CC research can have a significant impact on many aspects of modern life, with real consequences for the worlds of entertainment, culture, science, education, design and art.

1 From C to CC

Creativity is a multi-faceted phenomenon that manifests itself in different guises in different domains. So creativity in the domain of sports (e.g. as manifest in a team sport like *soccer*, or an intellectual game like *chess* or *Go*) is clearly different to

---

Tony Veale
School of Computer Science, University College Dublin, Ireland.
e-mail: `tony.veale@gmail.com`

Amílcar Cardoso
DEI / CISUC, University of Coimbra, Portugal. e-mail: `amilcar@dei.uc.pt`

Rafael Pérez y Pérez
Universidad Autónoma Metropolitana, Cuajimalpa, Mexico. e-mail: `rpyp@unam.mx`

creativity in the arts domain (e.g. consider *painting* or *poetry*), yet there are enough similarities for exemplary outcomes in each domain to be deserving of the same label, "creative". This heterogeneity makes creativity a notoriously difficult concept to pin down in formal terms, and definitions that favor one area of human activity (such as art) are unlikely to do justice to other areas (such as science, engineering, cooking). Our definitions of creativity – and a great many have been considered in the scientific literature – are no more than accepted conventions, and it is in the very nature of creativity to bend and subvert these conventions.

Computational Creativity (CC) is an emerging branch of AI that studies and exploits the potential of computers to be more than feature-rich tools, and to act as autonomous creators and co-creators in their own right. In a CC system, the creative impetus should come from the machine, not the human, though in a hybrid CC system a joint impetus may come from both together. As a discipline, CC draws on research in Artificial Intelligence, Cognitive Science, Psychology and Social Anthropology to explore the following questions:

- What does it mean to be "creative"? Does creativity reside in the producer, in the process, in the product, or in a combination of all three together?
- How does creativity relate to expertise and to what extent does it necessitate specialized domain knowledge?
- How does creativity exploit and subvert norms and expectations?
- How are the outputs of creativity judged and evaluated? How can we meaningfully measure creativity? What knowledge is needed (of the creator or process) before we can label a work "creative"?
- What constitutes creativity in different domains and modalities?
- How does creativity emerge from group behavior and collective action?
- What cognitive paradigms offer the most insightful explanatory theories of creativity (e.g., search in a conceptual space, conceptual blending, etc.)?

Each of these questions is just as valid to the study of human creativity as they are to the study of machine creativity. What makes CC different is that it adopts an explicitly algorithmic perspective on creativity, and seeks to tie down the study of creative behavior to specific processes, algorithms and knowledge structures. The goal of CC is not just to theorize about the generative capabilities of humans and their machines, but to build working systems that embody these theoretical insights in engineering reality. So CC is both an engineering discipline and an experimental science. in which progress is made by constantly turning insights into applications that can be experimentally tested and evaluated. The purpose of these applications is to create novel artifacts – stories, poems, metaphors, riddles, jokes, paintings, musical compositions, games, etc. – in which a large measure of the perceived creativity is credited directly to the machine. We believe that the future of intelligent computers lies in transforming our computers from passive tools into active co-creators, and that CC is the field that can make this transformation a reality.

CC researchers tend not to trade in definitions of creativity *per se*, but to focus on those aspects of behavior – in both humans and computers – that produce outputs that are novel or surprising and which yield unexpected value. It was in this vein that

Newell, Shaw, and Simon (1963) suggested four different criteria for categorizing an answer to a question, or a solution to a problem, as “creative”:

1. The answer has novelty and usefulness (for the individual or for 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

Though obviously incomplete, each criterion is instinctively appealing because each expresses in literal language the meaning of a conventional metaphor of creativity. For instance, (1) simply reflects the folk view that creative solutions should be “fresh” and “innovative”, perhaps even “ground breaking”; (2) suggests that one must “think outside the box” and reject conventional categories and labels; (3) suggests that to be creative, one must expend copious amounts of “mental energy” in tenaciously exploring the avenues of a wide-ranging conceptual space; and (4) espouses the common belief that creativity requires “illumination” and “insight”.

Given the obvious difficulties in distilling a pure definition of creativity – *pure*, at least, in the sense of being metaphor-free and grounded in objective fact rather than in human intuition – CC researchers pursue one or all of the following approaches:

1. They ignore the need to define the phenomenon objectively, to perhaps employ instead an ad-hoc definition of convenience; this allows practical work on creative systems to continue, perhaps even to an extent that practical results can eventually inform a fuller and more satisfying definition of creativity.
2. They embrace the metaphorical foundations of creativity, to identify processes and mechanisms within our repertoire of computational algorithms and representations that best seem to embody these folk metaphors.
3. They identify an archetypal area of creative endeavor and attempt to model that area computationally. In such work, a formal definition is not needed to underwrite the research as “creative”. However, as in (1) above, the outputs of this research may then feed back into a later formalization of creativity.

These three alternatives summarize, more or less, the research assumptions made by contemporary CC researchers. Because the field is anchored in engineering and experimentation, CC systems produce concrete outputs whose novelty and value can be assessed by human judges in the absence of any formal definition of creativity. Though many CC researchers believe that machines can exhibit creativity on their own terms, perhaps even by using algorithms and knowledge structures that are different to those used by humans, a principal goal of CC is for machines to exhibit human-level creativity that humans will also perceive as “creative”. In striving for this technical goal, CC researchers and their systems can illuminate the processes and biases of human creativity too.

While humans and computers can be creative in the absence of a formal definition of *how* they are being creative, both still need a level of self-understanding and critical awareness to justify the use of the label “creative”. Computers which generate outputs for an external user to evaluate are merely generative in their behavior, and *mere generation* does not rise to the level of human creativity. Rather, the generation of outputs must be coupled with an awareness of the value of the output in terms of its novelty and its utility. A creative computer must embody a particular view of creativity that the computer itself understands, so that the computer can justify its outputs much as a human creator would do. Such a computer cannot be a dumb savant that naïvely flings outputs at an audience. Crucially, it must exhibit an ability to filter its outputs for quality, so that any outputs presented to a user show intentionality and discernment, and just as importantly, it must exhibit an ability to articulate why its outputs may have interesting and unexpected value for its audience. Thus, according to the *investment theory* of creativity (Sternberg & Lubart, 1995, 1996), a creative computer must be able to articulate its sense of how a particular product or idea can be "bought low and sold high".

Though developments in the field of AI have become fixtures of the technological landscape (e.g., machine translation, natural language question answering, driverless cars, grandmaster-level chess and Go), humans still instinctively cling to the idea that creativity is a uniquely human (or uniquely biological) preserve. In this view, when computers apparently exhibit some measure of creativity, this mere appearance of creativity is due to some specifiable slice of the programmer's own creativity having been imprinted onto the algorithmic workings of the system. In CC research this idea is known as *pastiche*, since such computers unknowingly resort to the same kinds of stylistic mimicry that is knowingly exploited by uncreative human artists. For instance, careful musicological analysis of the structure of Bach cantatas can allow a programmer to write software that generates its own novel cantatas in the style of Bach. Though these outputs may fool the human listener, and even delight the unsophisticated ear, they are the product of a system that mimics rather than creates. Such a system has no awareness of its inherent limitations, nor does it have any conceptual input into the hardwired (albeit pseudorandom) processes that it follows. Such systems are more like skilled forgers than creative artists; while they can expertly mimic and recycle, they cannot innovate and nor can they surprise. Moreover, because they explore a pre-defined sweet-spot in the space of possible outputs, pastiche systems take no risks, always produce well-formed outputs, and have no need to self-critique or to ever learn from their failures.

Of course, pastiche has its place, both in human and machine creativity. One can learn from pastiche, and even good creators occasionally lapse into pastiche (recognizing this tendency in himself, Picasso once noted of his own paintings "Sometimes I paint fakes"). Pastiche thus serves as a useful boundary case for computational creativity. Indeed, there are cases where pastiche is precisely what the human co-creator desires (e.g. "let's explore more variations on this theme."). Pastiche-based systems are a useful starting point for the computational exploration of creativity, but the goal of CC as a field is to actively transcend pastiche, to demonstrate that computers are capable of true, human-level creativity.

CC is an interdisciplinary research field that sits at the intersection of the fields of AI, Psychology, Cognitive Science, Linguistics, Anthropology and other human-centered sciences. Given its focus on system-building, the field has most in common with AI, and builds on many of the same foundations, such as intelligent search in a conceptual/problem/state space. Nonetheless, the field has a distinctive character

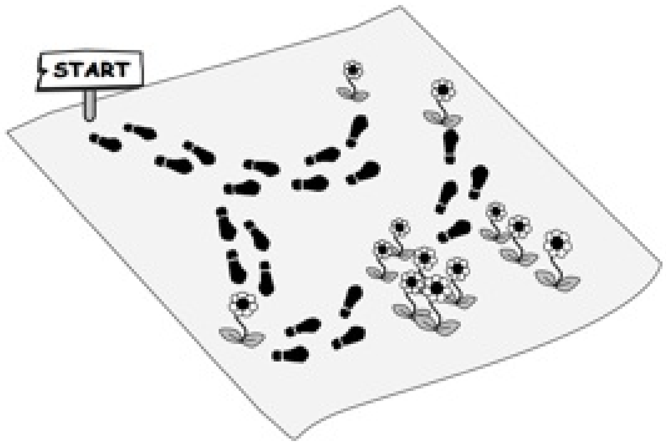

Fig. 1 Search in a state-space (Veale, 2012). Flowers represent acceptable goal-states ors solutions, while footprints illustrate the paths pursued via various cognitive agents.

of its own, which shapes its use of ideas and techniques from other fields. For instance, CC views creativity as arising from more than a merely systematic search of a conceptual space of possibilities. Rather, it recognizes that these spaces are deeply-rutted with conventional pathways, and that creativity arises from how an intelligent agent knowingly exploits or subverts these conventions. Thus, Boden (1990) suggests ways in which creativity might arise from the exploration of such a space, while G. Wiggins (2006) has formalized the CC components of this perspective.

A visual representation of search in a conceptual space is rendered in Fig. 1. Here, flowers depict acceptable solutions – goal states at which a search can profitably terminate – while footprints illustrate the paths taken by a cognitive agent as it explores the space. Since this model projects physical search into mental spaces, we can understand "mental agility" as the cognitive equivalent of those qualities that are desirable for an agile physical search. For instance, one often needs to backtrack gracefully when at a dead-end, and shift smoothly to an alternate avenue of search. Note that the search metaphor is just that, a metaphor, though it one that some CC researchers nonetheless resent as overly reductive. However, alternate metaphors for creative choice-making may yet be reducible to the non-deterministic exploration of an abstract space.

Adaptability, in particular, seems to be a salient aspect of creative behavior that can be formalized in terms of search spaces. Boden (1990) offers an intriguing view of adaptive creativity, of a kind that not only delivers surprising solutions to a problem, but that also changes the way we view the problem itself. Boden argues that one should distinguish *exploratory* creativity – of the kind visualized in Figs. 1 and 2 – from *transformational creativity*. While the former explores the space as it is defined by the problem, looking for previously undiscovered or unappreciated states of unexpectedly high value, the latter actively transforms the space. As illustrated

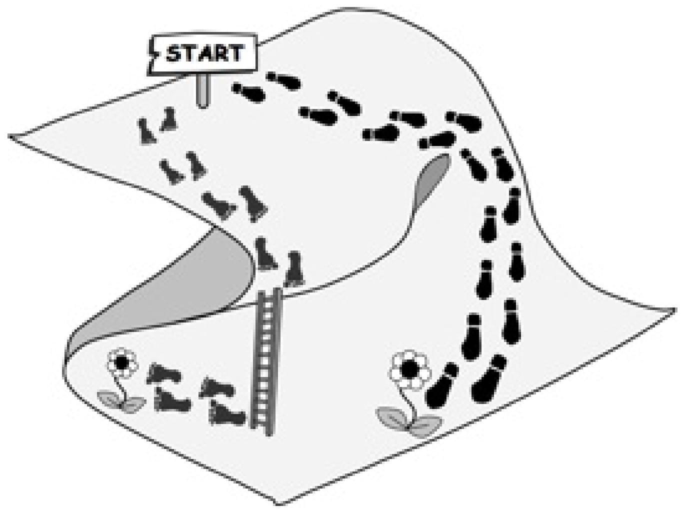

Fig. 2 A creative searcher (shown here as a bare-footed explorer) finds novel ways to navigate a search space, by e.g., looking in hard-to-reach areas, or by identifying unconventional connections between states that previously did not appear connected. (also from Veale (2012))

metaphorically in Fig. 3 this transformation re-defines the criteria of value that gave shape to the space and which drive the search for value in that space.

Boden cites the development of atonal music as a dramatic example of transformational creativity, and one can also point to key developments in science, such as the transformational shift from a Newtonian (absolute) to Einsteinian (relativistic) world-view, or from a classical (determinate) to quantum-mechanical (indeterminate) conception of reality. When searching through a space, whether that space is physical or abstract, a searcher can either contort itself to fit the constraints of the space, or contort the space to fit the needs and values of the searcher.

Transformations of the kind analyzed by Boden are the exception rather than the rule in creativity, in either its *small-C* (everyday creativity on a mundane scale) or *big-C* (exemplary creativity on a historical scale) guises. One finds a more commonplace form of agile exploration of a state space in the narrative jokes that are the common currency of social interaction. Jokes exploit the fact that we all navigate through shared state spaces in our everyday lives, to explain the events in the world around us, and to understand the behaviors of our friends and colleagues. These shared spaces have well-trodden pathways that correspond to the commonsense norms of conventional thought processes, but these rutted paths do not always offer the quickest or surest routes to a solution. In cases when the best path to a solution is circuitous and non-obvious, mental agility is not a matter of speed but of sure-footedness. The shortest path can sometimes lead to incongruity and failure.

Jokes employ state-spaces that have been deliberately warped, so as to fool the unsuspecting explorer into believing that the quickest and most conventional route is also the most intelligent route. In other words, jokes subvert the logic of intelli-

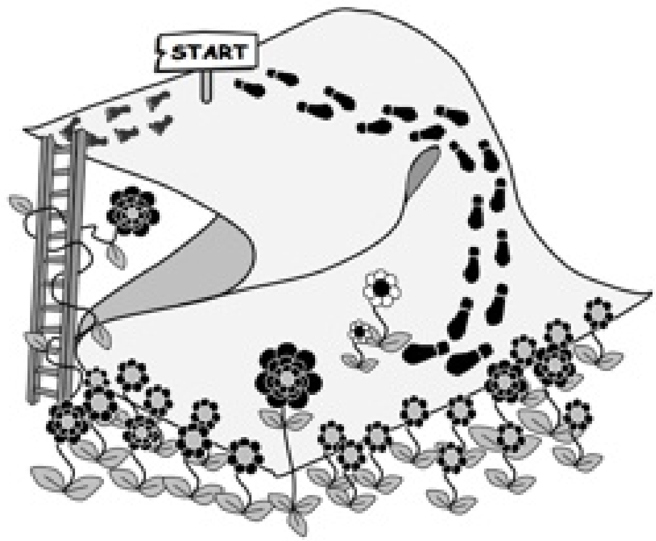

Fig. 3 Transformational thinkers alter the space that they are exploring, to identify high-value targets that lie outside the original space, and which would not have been considered in the original formulation of the problem. Of course, a transformation may also place states that were previously accessible out of bounds to the creative agent. (from Veale (2012))

gent search in a state space, and thereby demonstrate the limits of conventionalized thought processes (Minsky, 1980). The mathematician John Allen Paulos uses the framework of *catastrophe theory* to characterize the kinds of warped spaces that are most used in narrative jokes: as shown in Fig. 4, these typically contain an unexpected "kink" or discontinuity that corresponds to a surprising gap in the logic of the narrative (Paulos, 1982; Veale, 2012). Explorers who jump to conclusions by pursuing the path of the discontinuity can be humbled and surprised by their unthinking use of conventional logic.

It is on the computational treatment of discontinuity, incongruity and contradiction that CC most distinguishes itself from AI as a discipline. In a conventional state-space search, contradictions are viewed as dead-ends from which a computational agent must backtrack. AI makes an assumption that search is important but the avoidance of search is more important still, so contradictions serve as useful boundaries to limit an otherwise costly search. CC, however, views incongruity and contradiction as opportunities for further search, to explore whether anomalies can be resolved on another level of representation to yield results that are surprisingly meaningful. Resolvable contradictions of this kind underpin not just the incongruity of jokes, but the absurdity of surrealist paintings, the semantic tension of metaphors, the pragmatic insincerity of ironic statements, the plot twists of mystery stories, and even the unexpected discoveries of mathematics and science. In his wide-ranging theory of "*Bisociation*", Koestler (1964) argued that the creativity of these diverse phenomena emerges from the collision of two seemingly incompatible frames of ref-

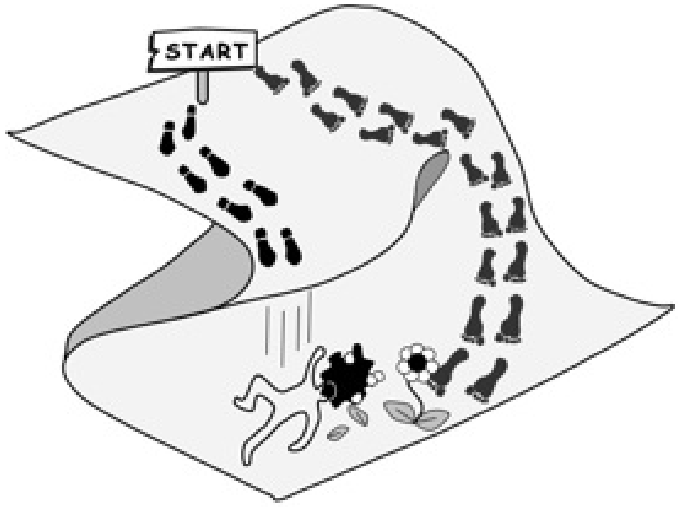

Fig. 4 Some state-spaces are deliberately constructed to be misleading, and the most obvious or conventional path to the solution can lead to a surprising dead end. A sure-footed explorer who knows the space takes a more circuitous route. (from Veale (2012))

erence (see also Lavrač et al. (2017) in this volume). Koestler's ideas form the basis of Fauconnier and Turner's influential theory of *Conceptual Blending* (cf. Fauconnier (1994); Fauconnier and Turner (2002)), though it falls mainly to researchers in CC to anchor these ideas in the algorithmic specificity that can only come from computational model-building (e.g., see Martins, Pereira, and Cardoso (2017); Pereira (2007); Veale (2017); Veale and Li (2011); Veale and O'Donoghue (2000)).

2 The Association for Computational Creativity

A key development in the history of the field has been the establishment of an international Association for Computational Creativity to promote further research in CC and to foster public engagement on the societal issues surrounding CC technologies.

The Association, or ACC, exists to promote the scientific study of human and machine creativity via computational means. The roots of the Association were put down in a series of early workshops and symposia that were explicitly dedicated to the issues of creativity in computers, such as the 2^{nd} Mind conference (1997) and events co-located at AISB (1999-2003), ICCBR (2001), ECAI (2002), EuroGP (2003 & 2004), IJCAI (2003), ECCBR (2004), LREC (2004), IJCAI (2005) and ECAI (2006). Originally, these events were organized by a small cadre of researchers who first coalesced as a working group through EU COST action 282 (*Knowledge Exploration in Science and Technology*), though this group quickly expanded to include scientists from diverse parts of the world. In 2007 the community re-launched the *International Joint Workshop on Computational Creativity* (or IJWCC) as a stand-alone event and an international steering committee for the workshop and its kindred events was formally established (the history of the IJWCC series is described in a special issue of *AI Magazine on CC* (Cardoso, Veale, & Wiggins, 2009).

In 2008 the Steering Committee took the decision to transform the workshop into a conference, thus establishing the International Conference on Computational Creativity (ICCC). The first ICCC was held in 2010 in Lisbon, the second in 2011 in México City, the third in 2012 in Dublin, the fourth in 2013 in Sydney, the fifth in 2014 in Ljubljana, the sixth in 2015 in Park City, the seventh in 2016 in Paris, the eighth in 2017 in Atlanta, and the ninth is planned for 2018 in Salamanca. During the first year of the ICCC in 2010, the members of the Steering Committee formally recognized the necessity of creating an international Association for Computational Creativity. Thus, in 2011 the Association was officially founded, and Geraint Wiggins was elected as its first Chair. In 2015 Rafael Pérez y Pérez was elected as its next chair, and that same year the ACC's official constitution was ratified by its steering committee.

From its origins to the present day, the Association has pursued a range of community-building activities, from the annual organization of the ICCC conference to the creation and maintenance of a comprehensive Wikipedia entry on Computational Creativity to the publication of special journal issues on CC (such as *Knowledge-Based Systems*, 9(7), 2006; *New Generation Computing*, 24(3), 2006; and *Minds and Machines*, 20(4), 2010). The main event organized by the Association is its annual conference, the International Conference on Computational Creativity (ICCC). The conference's mail goals are to provide a space where researchers from across the world can meet to debate ideas, hear about novel approaches to the study of creativity, build partnerships and start collaborations that explore interdisciplinary opportunities. The participation of students and young researchers has always been a priority for the ACC. From the beginning, the members of the steering committee has made it its mission to make publically available all the materials generated by the Association. The proceedings of its past conferences can be downloaded from the association's web page: *www.computationalcreativity.net*.

The support of the Association has provided fertile ground for new CC initiatives. In 2013, for example, seven European CC researchers (and members of the Association) obtained support from the European Commission to organize, under the aegis of the PROSECCO coordination action, a range of activities to stimulate enhanced CC outreach and education, including tutorials, summer schools, code camps, workshops and contact fora. PROSECCO has grown within the environment cultivated by the ACC, just as the Association has itself been shaped and advanced by PROSECCO. For example, the charter of the ACC was a specific deliverable of the PROSECCO project in its first year of operation. Another outcome of this symbiotic relationship is the volume you are now reading, which has always been a planned effort of the ACC but which is now made a reality as a PROSECCO deliverable. Though PROSECCO's funded lifetime as an EC project ended in late 2016, its legacy will live on in the Association.

The Association faces important challenges in the coming years, the most notable of which is its consolidation as an international society for all CC researchers that will continue to promote the goals, the philosophy and the technological vision of CC. This consolidation and growth will be only achieved through the committed participation of all of its members.

3 The PROSECCO Vision

PROSECCO is an international coordination action that was funded by the European Commission (2013-2016) to *Promote the Scientific Exploration of Computational Creativity*. The action was anchored in the belief that our computers can be more than mere “tools” of human creativity, but can actually rise to the level of co-creators that proactively share the creative responsibility with a human peer. As co-creators, our computers will be capable both of generating their own ideas and of framing those ideas in the appropriate modalities (e.g. language, image, sound). The PROSECCO vision of a co-creator is more than a mere facilitator or enabler for human creativity (in the sense e.g. that *Microsoft Word* or *Adobe Photoshop* facilitates content creation, or in the sense that Facebook facilitates collaborative creation), but envisions a largely autonomous agent that explores its own conceptual spaces and expresses its own ideas in its own terms. CC conducts application-driven research into this notion of a computational co-creator in two guises: autonomous systems that receive little or no human input; and semi-autonomous hybrid systems that interact with humans as peers.

This view is shaped not by a desire to replace humans with machines, nor by a perceived lack of human creativity in modern society, but by the belief that large amounts of human creativity remain untapped because users lack the appropriate co-creation software. No matter how richly featured a conventional software tool may be, users are still forced to start from a blank page or an empty screen, or a pre-determined template that simply encourages recycling and pastiche. Future CC systems must not only suggest ideas to users, but articulate, demonstrate and critique those ideas as would a human teammate. By providing humans with partners that can share the creative responsibilities and the creative credit, the goal is not to replace human creativity, but to engage and foster human creativity as only a creative equal can.

This is an ambitious vision that will take decades to fully realize, though in the interim, the CC field continues to build systems that serve useful (and steadily improving) creation and co-creation roles. To meet these challenges, CC must go from being a growing area of niche interest to being a true scientific discipline in its own right. The challenges are both organizational and research-based. As a field, we must coalesce around a clear set of principles, an unambiguous and comprehensive terminology, and a canonical set of techniques, metrics and approaches to evaluation (e.g. see Ritchie (2007)). We must consolidate our own identify as a field while actively engaging with neighboring disciplines. Computational Creativity has long been an implicit element of AI research, one that comes to the fore when AI addresses topics of an obviously creative bent, such as painting (e.g. Harold Cohen's Aaron), analogical reasoning (e.g. see Gentner (1983); Goel (2017); Hofstadter (1995); Veale (2006); Veale and Keane (1997); Winston (1980)), music generation (e.g. the EMI of Cope (2006)), story telling (e.g. the TALE-SPIN system of Meehan (1981) and the MINSTREL system of Turner (1994)), joke generation (e.g. see Binsted, Pain, and Ritchie (1997); Gatti, Ozbal, Guerini, Stock, and Strapparava (2017); Hempelmann (2008)), or metaphor processing (e.g. Fass and Wilks (1983); Veale and Keane (1992); Veale, Shutova, and Klebanov (2016); Wilks (1978)). However, this work was generally seen as AI work, and not as a product of a specific movement toward the realization of true computational creativity.

A key pillar of the PROSECCO coordination action has thus been its educational programme, which has sought to inform and shape the next generation of Computational Creativity (CC) researchers. To this end, the project has organized a major tent-pole educational event in each of its three years. Beginning with an Autumn School in 2013, the project organized a code-camp in both 2015 and 2016, with a variety of smaller events (such as targeted tutorials) spread between these tent-poles. It was of the utmost importance that student participants at these educational events not fall into the beguiling trap of *mere generation* – the alluring belief that machines can be programmed to generate creative outputs without being able to appreciate those outputs for themselves – but to instead build generative systems that would be accepted as creative by the CC community. Our machines cannot appreciate their own outputs if they lack knowledge about the semantic components of their outputs, hence there is a need to provide CC students and researchers with a comprehensive knowledge-base of interconnected and semantically-grounded beliefs. An important outcome of the PROSECCO project has been the development of large-scale semantic resources for use in teaching CC principles and fostering future CC research. Specifically, PROSECCO has developed two complementary semantic knowledge bases to support these goals. The first is the *NOC list*, or *Non-Official Characterization* list, which provides vivid semantic detail on a large cast of famous personalities (800 at last count) and their many attributes. The second is the *Scéalextric* knowledge-base of plot structures and idiomatic renderings for story actions, which significantly lowers the otherwise formidable barriers to entry to the CC domain of automated story generation. Each of these resources, which collectively run to over 60,000 high-quality semantic triples, can be accessed on the dedicated PROSECCO GitHub site: *github.com/prosecconetwork*. In addition, readers may be interested in reading about the experiences of PROSECCO code-camp participants on a blog dedicated to these ongoing educational efforts: *bestofbotworlds.com*.

This volume of canonical papers constitutes another key part of PROSECCO's and the ACC's efforts to reach future and emerging researchers in CC while they are still in the development stages of their education. As an emerging field, CC needs to reach graduate students in a variety of fields and disciplines to create the next wave of active researchers. By reaching these students at a time where their Ph.D plans are still at a formative stage, this book can bring the necessary knowledge on mathematics, psychology, anthropology, sociology, art, language, music, science and design into the field, and demonstrate that CC is a research area in which its students can pursue a truly cross-disciplinary exploration of creativity at the intersection of experimental science, system-building engineering and the humanities.

4 A Thematic Overview

This volume brings together a diversity of papers on a diversity of themes, to collectively chart the terrain that is Computational Creativity. This section provides a brief introduction to each of the chapters and themes that await the reader of the volume.

Novelty is a pillar of many operational definitions of creativity, but what exactly do we mean by "novel"? Is an artefact novel to the extent that it differs from others that we have experienced in the past? But difference is itself a contextual notion, since the dimension along which two things can differ will be primed by our expectations of how they should be the same. Kazjon Grace and Mary Lou Maher thus argue here that it more meaningful to say that an artefact is novel to the extent that it violates our expectations of "more of the same" (Grace & Maher, 2017). Expectations shape our perception of novelty and creativity, but these authors also argue that expectations can crucially shape the generation process too.

Novelty is just one dimension along which the "creativity" of an artificial generative system may be evaluated. Anna Jordanous takes a wide-angle look at the issue of evaluation in this volume, to consider the issue from a historical, a strategic and a methodological perspective (Jordanous, 2017). What does it mean to say that a CC system has undergone an evaluation, how might we interrogate the results of an evaluation, and how might we compare two systems that putatively aim to generate the same kinds of artefact? As Jordanous notes, the aim of CC as a field should be to provide a sound and systematic basis for the rigorous evaluation of our automated systems.

This is an aim that also provides the guiding theme for Graeme Ritchie's contribution to this volume. Ritchie argues that a growth in the engineering sophistication of CC systems must be matched by a comparable growth in the objective rigor and sophistication of our methods for evaluating these systems (Ritchie, 2017). For Ritchie, a proper evaluation shows an understanding of the goals of the work being evaluated, requiring an evaluator to tease apart the theoretical agenda from its engineering application. While noting that CC systems operate in a realm that supports little in the way of objectively-defined and widely-agreed criteria, Ritchie uses an analysis of past evaluations to establish a solid foundation for the evaluation of future CC systems.

Creative systems can operate in the various modalities that we associate with human creativity, from the visual (e.g. painting, design, video games) to the musical to the linguistic. Each modality is associated with its own sense of what constitutes creative "genius". For language, this sense integrates notions of wit, concision and persuasive power. In their contribution to this volume, Lorenzo Gatti, Gözde Ozbal, Marco Guerini, Oliviero Stock and Carlo Strapparava explore linguistic creativity as it inheres in puns and in advertising slogans (Gatti et al., 2017). While the former allows CC to build models of the whimsical creativity that we associate with children, the latter – relying on many of the same techniques – allows CC to investigate an area of human activity with a compelling commercial use-case.

Story-telling is a capability that defines us as human beings. We use stories in every aspect of our social and intellectual lives, to explain the world both to ourselves and to others. What would otherwise be a mere sequence of events, one occurring after the other, becomes a coherent narrative in the hands of a story-teller. Though some make a livelihood from this ability, we are all natural story-tellers, placing narrative at the heart of computational creativity too, not least because so many of our definitions of creativity are little more than compressed narratives of what creativity *should* be. Insofar as we can define creativity at all, and define the social expectations of how a creative person should act, it is because we can tell a good story that draws these expectations into a recognizable narrative. The contribution of Rafael Pérez y Pérez to this narrative in this volume is a model of story-telling (his MEXICA system) that aims to capture how real people tell real stories in social situations (Pérez y Pérez, 2017). As with the best gossip, these stories hinge on how others both obey and subvert social expectations, so Pérez y Pérez sets out in this chapter to capture the kinds of common-sense knowledge that influence our understanding of social norms and those who obey, bend or break them.

As a CC researcher, Pablo Gervás is as much known for poetry generation as he is for story-generation. This duality is not a coincidence, but arises from his conviction that memorable poems typically have a memorable story to tell too. Gervás has thus studied the norms of story-telling and the norms of poetry side be side. So while his contribution to his volume principally concerns the latter, we invite readers to keep the former in mind when enjoying his chapter (Gervás, 2017). There are other points of similarity between this chapter and others in the volume too, since Gervás sets out to evaluate just *how* we should evaluate our computer poets before returning to the subject of his own automated Spanish poet.

Social convention is the invisible hand that guides both the generation and our appreciation of creative artefacts, whether or not we explicitly set out to obey or subvert this conventionality. As CC researchers we often aim to codify the governing conventions for our genre or domain into our systems so that they might satisfy the tastes of a user-base entrained to those conventions. Rob Saunders, is his contribution to this volume, uses a multi-agent framework to explore how those social conventions might arise in the first place (Saunders, 2017). His artificial agents provide a revealing sandbox in which we can observe the emergence, via self-organization, of norms in creative fields. For Saunders, creativity is not a quality of a lone system, but of a social agent interacting in a dynamic world.

If creativity resides as much in social interactions as it does it algorithmic action, perhaps other complex dynamical systems that exhibit an equivalent sensitivity to interaction and context might also be usefully labelled "creative"? Though we naturally take a human-centric view of creativity even in the context of CC, the core ideas of CC can be observed at much longer and much shorter time scales than the human lifespan, and at biological levels that are much higher and much lower than the human organism. In this contribution to this volume, Jon McCormack concerns himself with how ideas of a "biological" creativity can guide and inspire – via appropriate acts of abstraction, simplification and generalization – work in computational creativity (McCormack, 2017).

Abstraction and generalization are the guiding themes of Geraint Wiggins's contribution to this volume (G. A. Wiggins, 2017). Wiggins sets out to formalize the intuitions that hold sway in many discussions of creativity that implicitly see it as an exploration in a space of conceptual possibilities. These intuitions are often given a metaphorical form in layman's language, as when we speak of "exploring all avenues" and "coming up empty", of "hitting a brick wall" or "going around an obstacle", of "reaching a dead-end" or of "finding a goldmine of opportunities". Wiggins does more than recast these intuitions in formal language for the sake of formalization – he demonstrates that when expressed in wholly formal terms the ideas allow themselves to be manipulated in ways that are both enlightening and productive.

An unspoken aspect of the exploration view of creativity is that the explorer is an intentional agent, one that explores a space of possibilities with a specific goal (or at least a specific meta-goal) in mind. In his contribution to this volume, Dan Ventura considers the related questions of system *intentionality* and system *autonomy*, both at a philosophical level and at a practical level afforded by his CC system DARCI (Ventura, 2017). Papers such as Ventura's show CC to be a discipline that sits comfortably at the cross-roads of philosophy of engineering, where profound questions can not just be asked but answered in practical implementation terms.

Ofttimes CC allows these profound questions to be rephrased in simpler terms so that they might give rise to robust, scalable and useful implementations. In his contribution to this volume, Tony Veale explores the phenomenon of conceptual blending (Veale, 2017), but finds it too powerful and too operationally vague to be properly and faithfully implemented by any real CC system. He identifies a subspecies of conceptual blend that it is more amenable to robust computational modeling, coining the term "conceptual mash-up" to distinguish this related notion from its complex forebear. Importantly, Veale shows that such mash-ups are more than curtailed blends; they capture an important aspect of conceptual blending in a form that allows blends to be more than intellectual curiosities or mere playthings, so that mash-ups can actively fill the gaps in a CC system's representation of the world.

A fuller treatment of the merits of blending is offered by the contribution of Pedro Martins, Francisco C. Pereira and Amilcar Cardoso to this volume (Martins et al., 2017). Taking a historical perspective on the development and subsequent revision of an early CC implementation of conceptual blending – named *Divago* – the authors show how those early systems can provide a sound foundation for building a robust modern approach to the most challenging aspects of human creative behavior. As a field that prizes practical and continuous implementation it can be tempting to view the CC terrain as a junk-yard of once-promising but now-abandoned systems and approaches. However, as this paper shows, CC is a field that gives rise to families of related systems and approaches to knowledge representation that grow and evolve in interesting ways over time.

Indeed, a practical CC system will rely as much (if not more) on its domain knowledge and on a felicitous representation of such as it will on any special algorithms for manipulating this knowledge. For instance, if it is the goal of a CC system to find novel insights at the boundaries of two disparate domains, any effective search will crucially hinge on the construction and representation of those domains. In their contribution to this volume, Nada Lavrač, Matjaz Juršič, Borut Sluban, Matic Perovšek, Senja Pollak, Tanja Urbančič and Bojan Cestnik focus on scientific knowledge discovery at the overlap of domains that are constituted by different textual subsets of the scientific literature. In this way these authors give a robust statistical form to the intuitions of Arthur Koestler that creativity arises from the *bisociation* (i.e. simultaneous *bi-association* with two domains of knowledge) of two different perspectives or frames of mind (Lavrač et al., 2017).

Last but not least, the domain of choice for Ashok Goel in his contribution to this volume is a domain that encompasses many others, the *design* domain (Goel, 2017). In particular, Goel applies model-based analogy to the solution of creative problems in design, showing how past engineering solutions can be retrieved and adapted to suit new needs in new contexts. But those past solutions need not be the solutions of human engineers, and Goel shows how model-based analogy can turn *all* of nature into a case-base for biologically-inspired design.

5 Conclusion: Baby Steps in the Right Direction

The pioneering 19^{th}-century scientist Michael Faraday was once pointedly asked by Benjamin Disraeli about the practical uses of research in the nascent field of electricity. Faraday retorted "What use is a baby?" though in another equally apocryphal telling, Faraday replied "Why, one day you will tax it, sir". Disraeli's question seems rather short-sighted with the benefit of hindsight, knowing what we know now about the utility of electricity in modern society. Nonetheless, Faraday had a responsibility as a scientist to educate the general public about his ambitious vision for this startling new phenomenon. The same holds true for the champions of any transformational discipline, and so a dual impact of PROSECCO has been a raised public awareness of the benefits of creative computers and the shaping of realistic public expectations of progress in the field.

Consider the following response to a debate about the value or otherwise of computer-generated paintings in a BBC documentary about AI and CC. The response was published in a British broadsheet newspaper, *The Telegraph*, after an airing of the documentary :

> "(...) one man is trying to teach computers to paint. One picture with colourful dancers was lauded as a creative breakthrough but was actually atrocious. Which proves that as long as computer scientists have no artistic taste, it's unlikely computers ever will."
> (London Telegraph, April 4^{th}, 2013)

The key point is not that the journalist above was *wrong* – the above response is a valid *subjective* response to the output of a CC system – but that the critique is made in an inappropriate frame of reference. CC is a developing discipline and needs to be nurtured as such; expectations must be realistically shaped so that incremental breakthroughs are not cynically strangled at birth, and so that the public can appreciate the merits of computer-generated artifacts as the results of research-in-progress rather than finished research. The CC community must continue to engage with the public about the merits and possibilities of its research, to refute misconceptions and to respond to genuine concerns. By impacting directly on public expectations, our research can foster an environment in which less-than-human computational creativity can make its way into steadily improving software that is aimed at the general public. This volume, with its diverse collection of papers, is intended to help future CC researchers to make this transition a practical reality.

Acknowledgements This work was supported by the EC project PROSECCO, a coordination action funded to *PROmote the Scientific Exploration of Computational Creativity. The official PROSECCO website is:* PROSECCO-network.eu. *For introductory teaching material on Computational Creativity, visit the PROSECCO-sponsored website* RobotComix.com

References

Binsted, K., Pain, H., & Ritchie, G. (1997). Children's evaluation of computer-generated punning riddles. *Pragmatics and Cognition*, *52*, 309358.

Boden, M. (1990). *The creative mind: Myths and mechanisms (second edition)*. Routledge.

Cardoso, A., Veale, T., & Wiggins, G. (2009). Converging on the divergent: the history (and future) of the international joint workshops in computational creativity. *AI Magazine*, *30*(3), 15-22.

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

Fass, D., & Wilks, Y. (1983). Preference semantics, ill-formedness, and metaphor. *Computational Linguistics*, *9*(3-4), 178–187.

Fauconnier, G. (1994). *Mental spaces: aspects of meaning construction in natural language*. Cambridge University Press.

Fauconnier, G., & Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.

Gatti, L., Ozbal, G., Guerini, M., Stock, O., & Strapparava, C. (2017). Computer-supported human creativity and human-supported computer creativity. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Gentner, D. (1983). Structure-mapping: A theoretical framework. *Cognitive Science*, *7*(2), 155170.

Gervás, P. (2017). Exploring quantitative evaluations of the creativity of automatic poets. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity:*

The philosophy and engineering of autonomously creative systems. Springer International Publishing.

Goel, A. (2017). Revisiting design, analogy, and creativity. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Grace, K., & Maher, M. L. (2017). Placing expectation at the centre of computational creativity evaluation. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Hempelmann, C. (2008). Computational humor: Beyond the pun. In V. Raskin (Ed.), *The primer of humor research.* Berlin: Mouton de Gruyter.

Hofstadter, D. (1995). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. HarperCollins Publishers.

Jordanous, A. (2017). Evaluating evaluation: Assessing progress and practices in computational creativity research. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Koestler, A. (1964). *The act of creation*. London, UK: Penguin Books.

Lavrač, N., Juršič, M., Sluban, B., Perovšek, M., Urbančič, T., & Cestnik, B. (2017). Bisociative knowledge discovery for cross-domain literature mining. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Martins, P., Pereira, F. C., & Cardoso, F. A. (2017). The nuts and bolts of conceptual blending: Multi-domain concept creation with divago. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

McCormack, J. (2017). Creative systems: a biological perspective. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Meehan, J. (1981). Tale-spin. In R. Schank & C. Riesbeck (Eds.), *Contemporary approaches to creative thinking.* Hillsdale, NJ: Lawrence Erlbaum.

Minsky, M. (1980). *Jokes and the logic of the cognitive unconscious (a.i. memo no. 603)* (Tech. Rep.). M.I.T. A.I. Laboratory.

Newell, A., Shaw, J., & Simon, H. (1963). The process of creative thinking. In G. T. H. E. Gruber & M. Wertheimer (Eds.), *Contemporary approaches to creative thinking.* New York: Atherton.

Paulos, J. A. (1982). *Mathematics and humor*. University of Chicago Press.

Pereira, F. C. (2007). *Creativity and artificial intelligence: a conceptual blending approach*. Berlin: Walter de Gruyter.

Pérez y Pérez, R. (2017). Representing social common sense knowledge in MEXICA. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, *17*(1), 6799.

Ritchie, G. (2017). The evaluation of creative systems. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Saunders, R. (2017). Multi-agent based models of social creativity. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Sternberg, R. J., & Lubart, I. T. (1995). *Defying the crowd: Cultivating creativity in a culture of conformity*. Hillsdale, NJ: Lawrence Erlbaum.

Sternberg, R. J., & Lubart, T. I. (1996). Investing in creativity. *American Psychologist*, *51*(7), 677-688.

Turner, S. R. (1994). *The creative process: A computer model of storytelling*. Hillsdale, NJ: Lawrence Erlbaum.

Veale, T. (2006). Re-representation and creative analogy: A lexico-semantic perspective. *New Generation Computing*, *24*, 223-240.

Veale, T. (2012). *Exploding the creativity myth: The computational foundations of linguistic creativity*. London, UK: Bloomsbury.

Veale, T. (2017). From conceptual "mash-ups" to "bad-ass" blends: A robust computational model of conceptual blending. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

Veale, T., & Keane, M. (1992). Conceptual scaffolding: A spatially founded meaning representation for metaphor comprehension. *Computational Intelligence*, *8*(3), 494-519.

Veale, T., & Keane, M. (1997). The competence of sub-optimal structure mapping on hard analogies. In *Proceedings of ijcai97, the 15th international joint conference on artificial intelligence.* San Mateo, California: Morgan Kaufmann.

Veale, T., & Li, G. (2011). Creative introspection and knowledge acquisition: Learning about the world thru introspective questions and exploratory metaphors. In *Proceedings of aaai2011, the 25th conference of the association for the advancement of artificial intelligence.* AAAI press.

Veale, T., & O'Donoghue, D. (2000). Computation and blending. *Cognitive Linguistics*, *11*(3-4), 253-281.

Veale, T., Shutova, E., & Klebanov, B. B. (2016). *Metaphor: A computational perspective*. USA: Morgan Claypool: Synthesis Lectures on Human Language Technologies.

Ventura, D. (2017). Autonomous intentionality in computationally creative systems. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.* Springer International Publishing.

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

Wiggins, G. A. (2017). A framework for the description, analysis and comparison of creative systems. In T. Veale & F. A. Cardoso (Eds.), *Computational creativity: The philosophy and engineering of autonomously creative systems.*

Springer International Publishing.
Wilks, Y. (1978). Making preferences more active. *Artificial Intelligence*, *11*(3), 197–223.
Winston, P. (1980). Learning and reasoning by analogy. *Communications of the ACM*, *23*(12), 689-703.