Understanding Creativity: A Computational Perspective

Tony Veale

School of Computer Science and Informatics
University College Dublin
Belfield, Dublin 4, Ireland.
tony.veale@UCD.ie

Pablo Gervás

Facultad de Informática
Universidad Complutense de Madrid
28040 MADRID, Spain.
pgervas@sip.ucm.es

Alison Pease

School of Informatics
University of Edinburgh
Edinburgh, Scotland.
alisonp@dai.ed.ac.uk

§1 Introduction

What is creativity? This question is far more likely to elicit an anecdote, an aphorism or a metaphor than it is a literal definition. Creativity is an elusive phenomenon to study, made all the more vexing by our fundamental inability to pin it down in formal terms. One contributing factor to this general perplexity is a telling lack of a literal vocabulary for describing creativity and its mechanics, for when speaking of creativity, one is generally forced to speak metaphorically. Given this recalcitrance, it is tempting to conclude that the idea of creativity as a coherent concept is simply another bewitching effect of language, of the kind bemoaned by Wittgenstein[5].

As a category of human behaviour, creativity appears to resist all attempts at classical categorization via necessary and sufficient features. Indeed,

it may well the case be that there exists no single creativity mechanism, and that all instances of creative behaviour are best corralled into a meaningful synthesis only by a system of family-resemblances. This realization has lead researchers such as Newell, Shaw and Simon[4)] to attempt a multi-pronged definition of creativity. They suggest four different criteria for categorizing a given answer or solution as creative:

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

None of these criteria appear neither necessary or sufficient for creativity, which is perhaps just as well, since no criterion seems to hit the mark. For instance, consider (1) in the context of linguistic creativity. In the course of an average day, most speakers of a language will utter some statements that are so specific to their ephemeral contexts that they will never have been uttered before (or likely to be ever again); since these utterances serve some communicative function, they are thus both novel and useful, yet they can hardly be considered creative (except in the unconventional, generative-linguistics sense of Chomsky[3)], which describes to the ability of language users to generate an infinitude of different sentences). It would seem then that the meaning of "novelty" in (1) must already presuppose a notion of creativity. Likewise, the rejection of Habeas Corpus as a legal concept in time of war (or terrorist threat) is not as creative as (2) would suggest, for true creativity in such a context would find a way of preserving a treasured belief in troubling situations. Meanwhile, the intense motivation and persistence of (3) is as much a characteristic of tenacious plodders as it is of creative individuals, while for some problems, pedants are every bit as capable as insightful thinkers of providing the clarity called for in (4).

Though obviously flawed, each criterion is instinctively appealing because each recapitulates in more literal terms the meaning of a conventional metaphor of creativity. For instance, (1) simply reflects the commonplace 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 a particular av-

enue or wide-ranging conceptual space; and (4) espouses the common belief that creativity requires "illumination" and "insight". Yet, because of the systematic nature of conceptual metaphor, these metaphors allude to a wealth of intuitions that cannot be summarized in the neat bullet points of (1) thru (4), and one cannot but feel that Newell, Shaw and Simon's attempt at a literal exposition of the facts seriously short-changes these metaphors. A more considered exploration of the underlying metaphors, to perhaps achieve a synthesis at both the source and target levels of description, may yield a tighter and less fragmentary perspective on creativity.

Given these 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 human intuition - researchers can still respond in one of several ways:

1. Ignore the need to define the phenomenon objectively, to perhaps employ instead an ad-hoc definition of convenience; this will allow practical work on creative systems to continue, perhaps even to an extent that practical results will eventually inform a fuller and more satisfying definition.
2. Embrace the metaphoric foundations of creativity, to identify processes and mechanisms within our repertoire of computational algorithms and representations that best seem to embody these metaphors.
3. Identify an archetypal area of creative endeavour and attempt to model that area computationally. In such work, a formal definition is not needed to underwrite the research as "creative". 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 taken by the authors in this special issue on creative computation. These authors describe work originally submitted to the IJCAI'2005 workshop on Creative Systems, though the versions presented here have been elaborated at length. Before we turn you over to the authors themselves, let us consider first a brief capsule account of each paper and how each addresses, in its own way, the vexing problem of defining creativity.

The first, by Geraint Wiggins, presents a formalization of ideas first presented by Margaret Boden[1)2)]. While Boden's exploration of creativity relies

on some key intuitions about the nature of creativity, and results in ideas that have an obvious basis in metaphor (not least the beliefs that creativity can be "exploratory" and "transformational"), Wiggins formalizes these intuitions in a framework of computational search. He argues that this framework is not reducible to the traditional GOFAI framework (of "good old fashioned AI"), but is one that facilitates a more complex search space and a richer search experience. His starting point is an ad-hoc definition of computational creativity that is behaviourist in nature, as "the study and support, through computational means and methods, of behaviour exhibited by natural and artificial systems, which would be deemed creative if exhibited by humans". If this definition seems circular and tangential, this appears to be Wiggins' point: we need a working definition of creativity not to tell us what creativity is (since this is the goal, rather than an input, of creativity research), but to license the research that will ultimately allow us to replace such a definition with a formal specification of computational systems capable of creative activity.

The second paper, by Tony Veale, chooses to explicitly embrace the apparently metaphoric foundations of creativity as it is generally conceived and expressed in language. In this view, the metaphors we employ most to talk of creativity are sufficiently systematic to suggest that language is a path, rather than an obstacle, to understanding creativity. In particular, Veale argues that a metaphoric notion of figure-ground reversal underpins many of the metaphors of creativity that are evident in language: we speak of "thinking outside the box", "colouring outside the lines", effecting a "paradigm shift", a "breakthrough" or a "revolution". To give this argument computational substance, Veale looks at the problem of analogy generation using lexical ontologies like WordNet and HowNet. In this task-specific context, figure-ground reversal can be given a very real embodiment in representational terms. Veale demonstrates that greater analogical coverage can be achieved in a system that is sufficiently flexible to structurally invert logical definitions to derive new definitions for concepts that would otherwise be structurally impoverished.

The third paper, by Graeme Ritchie, explores in considerable detail a key metaphor arising from Boden's perspective on creativity, which Ritchie dubs the "Transformational Creativity Hypothesis". In brief, Boden articulates the view that creativity occurs within a conceptual space, where much of what passes for creativity arises from an exploration of this space, but where the most dramatic forms of creativity necessitate a transformation of this space[1)2)]. This

metaphor/hypothesis is an appealing one that has proven highly influential in the study of computational creativity. In his paper, Ritchie investigates the hypothesis in greater depth, to identify how the metaphors of "space" and "transformation" can be given formal, computational substance. He concludes that, for the time being at least, the metaphors must remain exactly that, metaphors, but that this inherent vagueness does not preclude the construction of genuinely creative computer programs.

The fourth paper, by Diarmuid O'Donoghue, Amy Bohan and Mark Keane, concerns itself more with the practical application of computational creativity than with its philosophical underpinnings. Like Veale, the creative focus is on analogical reasoning, but O'Donoghue, Bohan and Keane are not interested in linguistic analogies constructed from lexical resources, but in visual and geometric analogies. They bring the machinery of analogical reasoning to bear on the problem of interpreting topographic (land-cover) maps, and argue that in understanding such maps, some degree of creative "imagination" must be employed. Their paper amply demonstrates that the role of creative computation is not limited to pyrotechnical exemplars such as figurative language, story telling and music composition, but has a very real utility in everyday software applications.

Nonetheless, to conclude on a pyrotechnical note, the final two papers look to the processes of story generation as an archetype of creative human behaviour, and attempt to formalize the processes of story generation in computational terms. Federico Peinado and Pablo Gervás describe a case-based reasoning approach, in which a case-base of existing plot-lines is leveraged, through the use of a mediating ontology of story elements, to generate new narrative variations. The stories of interest to Peinado and Gervás adhere to the structural form of classical folk-tales as schematized by Vladimir Propp in this now famous analysis of Russian fairy tales. These authors give Propp's ideas about the morphology of the folk-tale computational form via a description logic, and their paper furthermore presents an initial empirical analysis of the stories generated using their approach.

Authors Mark Riedl and Michael Young also look to the process of story generation as an exemplar of creative behaviour, and use their paper to describe a new story-telling system they call "Fabulist". Riedl and Young note that while acceptable stories can be generated using a rigid or closed representation of world knowledge, creative stories require the teller to transcend the limitations of this

representation and effectively "think outside the box". Fabulist achieves this transcendence by providing limited mechanisms for enhancing the closed-world planning process used by the application of traditional to story-telling. These mechanisms enable partial transgressions of the closed world assumption and the stated restrictions on character behaviour, and they allow the planner to obtain better stories and to make the actions of these characters seem all the more credible. Riedl and Young observe that story-telling is inherently more creative when the initial specification of the world state is open and under-specified, thereby granting the story planner and plot generator the freedom to allocate story-resources (such as props and secondary characters) in the service of the plot.

It will be noted that none of these authors either attempt to provide a formal definition of creativity, or view such a definition of creativity as the ultimate goal of their work. While such a definition would certainly be valuable, these papers demonstrate that it is not essential to research in creative computational systems. This reflects an entirely pragmatic view that such a definition is not useful in itself, but as a foundation upon which useful creative systems can be built. The work described here thus represents an evolutionary view to understanding creativity: new systems built on old intuitions can yield new insights on which progressively newer and better systems can be built, to ultimately reveal, in computational terms, the essence of creativity. In effect, these papers perform an important figure-ground reversal in their own right. They shift the focus from defining creativity to achieving creativity; once achieved, even in a meagre form, we will all be better positioned to answer the original question: what is creativity?

Tony Veale, Pablo Gervás, Alison Pease
Dublin, Madrid, Edinburgh

References

1) Boden, M. *The Creative Mind: Myths and Mechanisms*, London: Weidenfeld and Nicholson, 1990.

2) Boden, M. "Computational models of creativity," Handbook of Creativity: 351–373, 1999.

3) Chomsky, N. *Syntactic Structures*, The Hague: Mouton, 1957.

4) Newell, A, Shaw, J. G, and Simon, H. A. "The process of creative thinking," in *H. E. Gruber, G. Terrell and M. Wertheimer (Eds.), Contemporary Approaches to Creative Thinking, pp 63 - 119*, New York: Atherton, 1963.

5) Wittgenstein, L. *Philosophical Investigations. Translated by G.E.M. Anscombe*, London: Blackwell, 1958.