

Kent Academic Repository

Full text document (pdf)

Citation for published version

Johnson, Colin G. (2012) Connotation in Computational Creativity. *Cognitive Computation*, 4 (3). pp. 280-291. ISSN 1866-9956.

DOI

<https://doi.org/10.1007/s12559-012-9141-8>

Link to record in KAR

<https://kar.kent.ac.uk/70809/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Connotation in Computational Creativity

Colin G. Johnson, School of Computing, University of Kent
Email: C.G.Johnson@kent.ac.uk

Author's submitted manuscript.

Definitive version published in *Cognitive Computation*, 4(3):280–291, 2012.

<https://link.springer.com/article/10.1007/s12559-012-9141-8>

Abstract

Computational creativity is the application of computers to perform tasks that would be regarded as creative if performed by humans. One approach to computational creativity is to regard it as a search process, where some conceptual space is searched, and perhaps transformed, to find an outcome that would be regarded as creative. Typically, such search processes have been guided by one or more objective functions that judge how creative each solution is on one or more dimensions.

This paper introduces a contrasting approach, which is search based on the idea of *connotations*. Rather than exploring a space constructed solely of potential outcomes, a larger space is explored consisting of such outcomes together with other relevant information. This allows us to define search processes that include a more exploratory process, out of which an outcome emerges via density of connotations. Both the general principles behind this and some specific ideas are explored.

Keywords

Creativity; connotation and denotation; search

1 Introduction: Search and Creativity

A common stance in AI research is to cast various problems associated with the action of mind as search problems¹. That is, it is asserted that the important aspects of intelligent action can be carried out just as well by a search algorithm as any other putative mechanism.

This can be seen as an aspect of the functionalist stance towards understanding mind [2]. A key part of functionalism is substrate-agnosticism, which hypothesizes that the phenomenon of mind can be equally well realised on any substrate that has sufficient computational richness. A second part of functionalism is a quasi-behaviourist process-agnosticism, which asserts that we can ascribe mentality to any process which, compared with something that we are certain has mentality such as a human, has similar input-output behaviour. That is, the internal detail of the process is regarded as irrelevant to the ascription of mentality. Taken together, these underpin the idea in AI that mental action can be seen as a computational search problem: the search process is similar enough to the action of mind, acting on a substrate (the computer) that has a similar capacity, appropriately scaled, as a mind in terms of perception and action.

This idea of search applies readily to many kinds of mental action. In limited domains, such as chess-playing, the search process produces a similar kind and level of behaviour as the human mind. Even in a more open domain, such as planning a route around a building or choosing what to eat for lunch, the problem can be reduced to a search problem, guided by certain constraints and goals, until a single point, or an equally-good set of points, remains.

This idea has been applied in “creative” domains too, where the aim of the search is to find an outcome that would be considered a creative product if produced by a human. Search

¹For a general discussion of the search stance in AI see any standard AI text such as Russell and Norvig [1].

processes in these domains are often guided by one or more heuristic measures that provide approximations to various aspects of creativity, such as novelty and value, or else guided by interaction with a human.

In this paper we explore search processes that are guided by a web of connotations of potential outcomes rather than by denoting the creative value of specific outcomes. The general idea of a space of connotations and its links to the kinds of search spaces used in computational creativity are defined, and some ideas for how these spaces could be automatically constructed based, for example, on link-analysis in web searches are outlined. A number of potential applications of these ideas are sketched, though no specific implementations are given.

The paper is structured as follows. Section 2 reviews the key ideas of computational creativity from the point of view of search processes, and discusses the structure of spaces used in computational creativity research. Section 3 defines what is meant by denotational search, and gives a number of examples of current computational creativity systems that are based on this idea; we then introduce the contrasting, novel concept of connotational search and explain why it is desirable. Section 4 describes the three main components that would be needed for a connotational search algorithm in general terms: a connotation space, a way of exploring that space, and a way of extracting outcomes at the end of that exploration. We then give more specific details in the form of four domains in which connotational search could be used: finding or building a key concept based on a number of seed concepts, free exploration of connotation space, using the idea of connotational search as an adjunct to problem-directed search, and using connotational search as a supplement to a human-directed process. Finally, Section 5 provides some final conclusions, critical evaluations and directions for future work.

2 Creativity and the Search Process

A number of definitions have been offered for computational creativity. One kind of definition that is found in much recent literature is one that uses human creativity as a defining behaviour: for example, “the study of building software that exhibits behavior that would be deemed creative in humans” [3]. Other definitions, from the very early discussion by Newell et al. [4] to very recent papers such as that of Maher [5] and Duch [6], have focused on identifying a number of characteristics of the outcomes of creativity such as novelty, value, surprise, unexpectedness and utility. Maher [5] also emphasises another contrasting view, which focuses systems that reproduce the content of the creative process rather than the outcome, e.g. combination of ideas, exploring new parts of a conceptual space, transforming a conceptual space, and identifying new spaces of possibilities.

Computational creativity systems have been built in many domains, including art [7, 8], music [9, 10], literature [11, 12], design [13], science [14], and mathematics [15]. Some such systems are built with the idea that they will be creative in their own right; others, that they will act as meaningful, albeit perhaps specialised, partners in the creative process alongside humans.

Search-based accounts of creativity [16, 17, 18] typically begin with the definition of a space of possible outcomes from the creative process. For example, in a system for musical creativity, the conceptual space might consist of melodies. Rather than referring to these as *potential solutions*, as would be conventional in AI search, the more neutral term *outcomes* will be adopted in this paper. The set of possible outcomes for a particular search process will be called the *outcome space* for a space consisting of the possible end-points of the creative process, and the points in this space will be referred to as *outcomes*. Wiggins [17] calls these “artefacts”, but, that has an implication of something built from subcomponents, rather than potentially discovered as a whole.

Boden [19, 20] discusses the related idea of a *conceptual space* that creative activity works within. However, this idea is not given a clear definition. Wiggins [16, 17] has presented a formalisation based on Boden’s ideas. In this formalisation, conceptual spaces are defined as proper subsets of a *universe* of concepts, containing both abstract and concrete concepts, and the notion of evaluating a point in a conceptual space.

In practice, whilst Wiggins’s framework allows the space to contain concepts other than complete works or components of works, attempts to implement computational creative systems in a search based framework have largely involved the exploration of an outcome space. As an example, one large category of such systems are the *evolutionary* systems, where the search heuristic is an evolutionary algorithm (see the overview in [21]). In the vast majority of such systems (see [22] for a detailed enumeration) the space being explored is the space of complete works.

In particular, Boden discusses different concepts of creativity: one of these is *exploratory* creativity, where the creative act consists of exploring a well-defined conceptual space for a novel outcome, but without changing the underlying space being explored. For example, writing a song in a well-known style, or painting a picture of a popular subject in a traditional style.

One variant on exploratory creativity is *combinational creativity*, where aspects of two or more points, either in the same or in different, conceptual spaces, are brought together to form a novel outcome. This can be seen as a form of exploratory creativity even where the concepts that are combined together are in different conceptual spaces—this can be seen as exploring the outer product of those spaces.

Boden’s second main concept is that of *transformational creativity*. This is where the underlying conceptual space is expanded or otherwise transformed to give rise to a new conceptual space for exploration. Canonical examples are the moves from representation to abstract art, the move from natural numbers to integers, the move from rhyming verse to free verse, and so on. This is usually seen as needing some kind of ability to reflect on the nature of the spaces and “self-criticize” [19]. Perhaps surprisingly, given the dominance of Boden’s work in discussions of computational creativity, the number of computational creativity systems that actually attempt to explicitly model this idea of transforming the space is very small.

It is implicit in this discussion that transformational creativity is typically a “more creative” act than exploratory creativity. However, this is difficult to pin down. Conceptual spaces can be on very different scales, and it is not obvious that exploring a large, complex space is necessarily less creative than making a minor transformation to a small, simple space.

Furthermore, creativity happens in multiple conceptual spaces at the same time. Consider the (oversimplified) idea that the conceptual space of poetry moved from rhyming verse to free verse. Is this transformational or exploratory? It depends on what underlying space we are considering. Clearly, in both cases, poets were working with (broadly) grammatical lists of sentences in the, say, English language. In that conceptual space, the search is exploratory. However, if the original conceptual space was lists of rhyming sentences, then the creativity is transformational.

Yet, it would seem that the poet would be aware of operating in *both* of these levels simultaneously. The creative search seems to operate in a partially ordered set of conceptual spaces, where the order is defined by subset inclusion of concepts. For example, all rhyming poems in English (call this C_1) are examples of formatted sets of English words (call this C_2), and so $C_1 \leq C_2$.

Transformational creativity can be seen as moving up this ordering to a broader set of concepts. This is not to deny the scope for genuinely transformational work: in this example, the development of the conceptual space of concrete poetry is an example of a conceptual space that cannot just be seen as exploring further within the more abstract conceptual space of lists of English sentences; but, we can repeat the argument again, and suggest that this is just exploratory within the placement of characters on a page.

Clearly, there are search spaces in which exploratory creativity is, nonetheless, something with vast scope for novel, valuable creative outcomes. Consider a space as large as all syntactically-valid paragraphs in a particular natural language, all arrangements of pixels on a particular size of computer monitor, or the space of all Turing-complete programs. It is unclear whether these are apprehensible enough to be regarded as conceptual spaces.

Wiggins [16, 17] has formalized these ideas of exploratory and transformational creativity.

This is based around the idea of a universe of possible artefacts that could be generated by the creative process, together with a means of exploring this universe, and an underlying background set of generative rules which constrain the set of points in the universe that can be explored. The search process can both explore within these rules, and change the rules themselves. This formalisation is the first stage of creating an implementation of this; however, creating an implementation is still challenging, particularly in terms of recognising the value of transformed spaces.

3 Connotation and Denotation

In this section we will explore what search could mean when the desired outcome from the process is to produce a creative outcome. In particular, we will explore past attempts at applying the concept of search to creative processes and show that many of them can be characterized as being driven by *denoting* how creative each point in the search space is; by contrast, we will explore how systems could instead use the concept of *connotation* as a way of annotating points in the search space, and how a connotation-based search process could be built.

3.1 Search, Objective Functions, and Denotation

Traditional search and optimization algorithms are driven by an *objective function*, which allocates some quality-based score to the points in the search space, or at least allows points to be ranked relative to each other [1]. This allows a search process—deterministic or stochastic—to be defined on the search space, where the choices of points to visit is based on the objective function. In standard search algorithms it is assumed that there is a single objective function, allowing any two points in the space to be unambiguously compare, and that this objective function is static with time.

One variant on this is *multi-criterion* search, where the objective function allocates multiple scores along different dimensions of quality, or allows points to be ranked along multiple dimensions [23]. There are a number of approaches to searching such spaces, from a simple weighted sum of the scores [24, 25] at each point to approaches based on Pareto-dominance [23, 26].

A second variant on this idea of an objective function is a *dynamic* objective function, i.e. one that varies with time [27, 28]. However, in such a system, the dynamics of the search space are independent of the search process as such, the objective function is a function of time, and not of the current or past positions visited by the search process.

We shall adopt the term *denotational* for an objective function where at any point in the search process we can take any point in the search space and assign a (perhaps multi-dimensional, perhaps time-dependent) objective function value for it. Importantly for the subsequent discussion, a denotational objective function *cannot* depend on the previous points visited by the search process, or by other points being stored as a result of that search (e.g. the population in a population-based search algorithm).

A second feature of traditional search and optimization is that it is based on some idea of *steady progress* towards an optimal solution. For example, genetic algorithms work with the idea of biasing the choice of parents from the current population towards the fittest parents [29]; swarm algorithms are guided by the local and global best solutions found [30]; and, deterministic search algorithms such as A^* search use the idea of a heuristic function to determine the current best next point to explore [1].

That is not to say that for each step in the search process, we make the same amount of progress towards a solution, nor that there aren't points where the search process continues with little progress. But, the general trend across the search process is from a very low-quality solution towards a high-quality solution; we can call this as a *staircase* process. Such a process begins with a starting point or set of points that are potential outcomes, and it is possible to identify the current best solution. At each step in the search process, either a better solution has been found (we can visualise this as ascending a step on a staircase) or not (which we can visualise as remaining on the same step). If the search is making successful progress, then the

quality of the solution increases. At any point it is possible to stop and output a solution that is the current best.

With regard to search processes that are intended to support creative processes, these two assumptions are usually taken as read, and some way of measuring progress within the creative process used. One approach is to use *interaction* with a human as a substitute for a fixed objective function (see e.g. [31, 32, 33, 34, 35]). The argument here is that the creative process is ill-defined, multi-faceted and subjective to be summarized by a constantly defined function, but that the judgement and/or emotional reaction of a human will be able to provide an appropriate ranking or scoring (a “subjective function”!). In particular, this draws upon the idea of *post hoc*, oracular recognition of creativity—we might not be able to define in advance of the act what we are *going to mean* by creativity, but we can recognise after the act that creativity has happened (to a greater or lesser degree).

For example, in the *GenJam* system by Biles [36, 37], a human performer interacts with a computer system for melody generation. The format of the interaction is that of a jazz improvisation where the human and computer take turns to play short melodies over an underlying, fixed rhythm section. The computer’s phrases are generated by analysis of and variation of the human’s phrase; the rules used to carry out this variation are learned during the performance via a machine learning algorithm. Feedback is given to the system via two kinds of interaction—the performer can rate the most recent phrase by pressing a key on the computer a certain number of times, or this feedback can come from the audience, via the display of feedback via coloured paddles (a contemporary version of this system would probably use direct audience interaction via mobile phones).

This is successful up-to-a-point, but has its limitations. Firstly, it can be criticized as offloading the creative role onto the human—in terms of machine creativity, this offers at best the human-machine interaction being creative, at worse the computer as a minor supporting player in the human’s creativity. Secondly, there are many practical problems, most notably the time required for the human to make many judgement, particularly towards the beginning (when there are many far-from-interesting solutions being offered) and towards the end (when many near-identical solutions are being explored).

Various approaches have been taken to close this loop and not require human intervention. One is to substitute some well-defined measure as a proxy for an ill-defined concept such as creativity. Most commonly, this has been some measure of the *novelty* of and outcome, sometimes with regard to other outcomes explored or with regard to a corpus of examples (see Peinado et al. [38] for an overview of ways of measuring novelty). Another common measure has been some measure of the *value* of the outcome [17, 39], whilst Ritchie [40] has also discussed proxies such as *typicality* and *untypicality*, and Colton [41] has discussed a “tripod” of appreciation, imagination, and skill.

One attempt to extend these ideas has been to use multiple different objectives in order to capture the multifaceted nature of creativity. This idea can be traced to very early (non-implemented) work by Newell et al. [4], who suggest that evaluating something creative can be achieved by measuring various factors—novelty, usefulness, rejection of previous ideas, persistence and removal of vagueness. More concretely, this has been realised in several recent projects. One example is the work of Rahman and Manurung [11], which uses multi-criterion optimization to search a space of texts for those that score highly on well-defined measures of *meaningfulness*, *grammaticality*, and *poeticness*. In a project by Sorenson and Pasquier, an evolutionary algorithm for game level design is guided by a measure of the amount of *fun* that the level will produce, using ideas from Csikszentmihalyi’s idea of *flow* [42].

Along similar lines, researchers in data mining have explored ideas of *interestingness* of solutions, to address the problem that most of the true facts that can be discovered in a large database are not of interest to the user. In an extensive 2006 survey of the of the field, Geng and Hamilton [43] summarize this work by giving a listing of nine broad factors that have been

used in defining such measures of interestingness: “conciseness, coverage, reliability, peculiarity, diversity, novelty, surprisingness, utility, and actionability”. Some systems have used just one of these; others have attempted to combine multiple ones into a composite interestingness measure.

A rather different approach has been to use the idea of *computerised critics*, which use a separate machine learning process to learn how to make a measurement of creativity [44, 45]. The basis of such an approach is similar to that of using humans in-the-loop, that is to capture the multi-faceted concept of creativity rather than have a single or multiple proxies for creativity. Typically these are supervised learning systems, which use a number of training examples that are given by the user to the critic system as exemplars of creativity. These are still denotational systems; points in the search space are being scored or ranked, albeit by another machine system rather than a human. Furthermore, it is difficult to get critics to make their own criteria for selecting or preferring one work over another, rather than just weighting a set of pre-provided criteria; nor do the tasks approached by critics go beyond the recognition of existing styles.

There has been some success with this approach, but there is a potential danger here that such critics learn *what has been regarded as creative in the past* (as illustrated through the training examples) rather than *what might be considered creative in the future*. One of the more slippery aspects of creativity is that determining what is regarded as creative (at least, *H*-creative [19]) is dependent on what has happened before with the system and its context. This depends to some extent on the level of abstraction of the patterns extracted by the critics.

3.2 Connotation

By contrast with this *denotational* aspect of search, creative search often has a *connotational* aspect to it. A particular point in the search space is valued because it provokes in the viewer/listener/reader a large number of connections, or a particularly interesting or rich set of connections.

This is observable, for example, in art criticism, where one aspect of the depth of a piece of art is its ability to link a number of different ideas in a non-trivial way. For example, consider the following excerpt from a recent exhibition review:

“Flanagan would stitch irregular hessian bags, fill them with sand and wait to see how the substance would behave when contained. The resulting sculptures—irregular, tubular, bulging like thighs or nearly splitting their sides—are nameless but full of character: hapless, toppling, cheery, distended. You can see a good deal of art prefigured there, from Ernesto Neto’s stuffed muslin biomorphs to Sarah Lucas’s stuffed-stocking bunnies.” [46].

This is rich with connotation—*affective connotations* (“cheery”), metaphor (“bulging like thighs”) and showing connections with other art trends (“a good deal of art prefigured there”).

Another visual example is in the creation of logos in advertising and branding. This is a very focused design challenge: the logo needs to be a simple image that will stand for many aspects of an, often rather complex, organisation. For example, consider the logo for the British Conservative Party, which is a simple child-like drawing of an oak tree. This might be seen as being a symbol that has many connotations, which are somewhat contradictory, and thus speaking to different audiences—the connotation with the concept of *growth* suggesting economic growth and the pro-business aspect of the party, whereas the oak as a symbol of Englishness might speak to a traditionalist anti-Europe aspect. This is a kind of *ostensive* creativity. Instead of constructing a complex artefact, the creative act consists of pointing at a single, often rather simple, object, and the creativity is in choosing the right object.

The core of a text, whether an advertising strapline or a fragment of poetry, can lie in the connotations of a small number of words. For example, Paul Celan’s poem *Todesfuge* [47] revolves around the two-word idea of “black milk”, bringing together radically conflicting imagery such as life (a baby drinking milk) and death (black as the colour of mourning garments in many cultures) to create an image powerful enough not to trivialise its holocaust theme.

Clearly not *all* of the creativity of a piece of art or design can be captured in its connotations. For example, part of the power of “black milk” is just in the unsettling visual image that it brings to mind (but, we might further ask, why *is* the image unsettling?). Certainly, this idea speaks more to those aspects of the creative process that are about *idea generation* than the detailed construction of a creative artefact. This is not a negative point, merely a note of the scope of its applicability. Furthermore, sometimes the value of a creative act lies in recognising the connotations of a particular outcome but then *going beyond* them.

Many accounts of the creative process in different domains (e.g. Hadamard in mathematics [48], Webb Young in advertising [49], Lawson in design [50], as well as general accounts such as those of Wallas [51]) begin with some idea of *exploration* or *ruminatio*n. This is one aspect that is lacking from the staircase, denotational models of creative search described earlier—in those, the focus is solely on exploring a space of possible outcomes.

By contrast, connotation provides a convivial framework for this exploratory process. Beginning from some (problem-specific or random) points in the connotation space, an algorithm could explore neighbouring points to build up a network of connotations. However, to prevent this just being arbitrary wandering, we would need to have some means of controlling which links to maintain. Determining which of these connotations are most relevant/meaningful/important will be a difficult challenge in this area.

Also, we would need some means of *forgetting*. Sets of connotations that led nowhere would need to be removed from the exploration. Eventually, we need to remove lots of the things that we have explored and just leave a few key concepts behind.

A related idea is the process of *distillation* that occurs in many kinds of artistic creativity. In such a process the artist works with a set of materials, producing prototype works and then discarding and reducing down aspects of those prototypes, producing simpler pieces which are then combined and elaborated again before being reduced again. The end result of this process can be a simple item or gesture. Nonetheless such distilled works can be very powerful to the audience, and will be considered to be works of great creativity, despite their simplicity. This may be because the distillation process is very good at picking out points in the search space that connote many other points; or it may be that the audience can (intuitively) perceive that there is not a straightforward route through the search space to the discovery of such a simple item.

When discussing denotational search we noted that it was staircase-like: the search progresses by moving from bad outcomes to better ones. The connotational approach is more *scaffolding*-like: the things that we explore in the course of the search algorithm are not imperfect examples of the final outcome, but are pieces of structure that help to build up to that outcome. Points visited during the search are not only valued when they represent mediocre versions of desired outcomes, but also when they might lead towards a desired outcome.

We should note that this is *not* just about biasing the exploration-exploitation tradeoff to spend more time in the exploratory phase. This is a different kind of exploration, in which items that could never be outcomes are explored. Furthermore, this is not just the process of combining aspects of two or more potential outcomes, e.g. through a crossover process such as that in evolutionary algorithms, because the parents need to be viable outcomes, whereas connotations can be other concepts linked with the outcomes.

A small number of systems in the computational creativity literature discuss something akin to connotation-based search. Veale presents a system, the *Jigsaw Bard*, which uses Google searches for phrases with “resonant variations on cultural stereotypes” [52]. Berthold and colleagues [53, 54] have taken Koestler’s idea of *bisociation* [55] as a way of finding missing concepts, both in general knowledge domains (using Wikipedia as a data source) and more specialised domains, such as bioinformatics. Norton et al. [56] automatically associates words with images as the beginning of the idea of providing a system for *appreciation* of an artistic creation.

Meanwhile, Jennings [57] discusses the idea of *affinities* as a way of evaluating the links between

producer-agents and critic-agents; affinities include concepts such as propinquity, similarity and familiarity, which have some similarities to the kinds of connotative links discussed in this paper. Finally, work on *metaphor* (e.g. [58]) uses similar search processes to those used by connotational search processes, however, the focus of such work is on finding and reporting the metaphors, rather than using them as the basis for further exploration.

4 Connotational Search: General Considerations

In this section we would like to explore more practical issues in defining a connotational search process. We consider three main aspects: the definition of a connotation space, the development of a search algorithm for that space and associated ways of assessing which links to explore, and the final selection of outcomes from the space.

4.1 Connotation Space

The first requirement is some idea of connotation itself. To begin formalising this, consider a *connotation space* consisting both of potential outcomes from the creative process (call these *o*-points) and a set of concepts that might be connotations with those outcomes (call these *c*-points). This space can be seen as a graph, where links in the graph represent connotation relations; this could be a weighted graph if we want to have some measure of the strength of the connotation.

Clearly, for any nontrivial domain it will be necessary to be able to construct this connotation space on-the-fly, rather than it being stored complete in some kind of database.

A somewhat related concept has been introduced by Wiggins [16], of a conceptual space where moves within that space can depend on multiple points in the space. This has been noted as being a richer representation than those typically used in AI search: "...note that this formulation is more powerful than the standard formulation of AI state space search because the function is allowed to select arbitrarily many members of [the space of concepts] and is not limited to the head(s) of the sequence. This is a key feature, because it admits search strategies which rely on the combination of or comparison between agenda items." [16]. However, this differs from the concept of connotation space as there is an assumption that a value can be given to each point in the space right away; indeed, all examples in the paper are either potential outcomes or components thereof. In connotation space *c*-points are also included in the graph. For example, in a connotational space where the *o*-points were images, the connotation space might have a *c*-point that is the colour "green". This is not a potential outcome, but finding an outcome associated with that colour might be important.

What is a connotation? Given a point *c* in a connotation space *S*, a connotation is something that (a particular group of) people will associate with that concept (to a greater or lesser extent: there is clearly some fuzziness here), something that people will "bring to mind" when *c* is present in their consciousness, or else that they would recognise as being connected when the two objects are presented. However, in order to carry this out, we will not be able to use this description directly, as assessing all connotations by human judgement using a large number of people is, in practice, impossible. Instead, we will have to use some kind of proxy for this, drawn from a corpus of information.

What kind of relation is connotation? We will use the symbol \xrightarrow{c} for *connotes*. It is *reflexive*: this requires little argument, it is clear that a concept is associated with itself. It is also *symmetric*: if c_1 connotes c_2 , then c_2 connotes c_1 . However, there may be exceptions to this: for example if c_1 is a very specific concept, and c_2 a very general one, then $c_1 \xrightarrow{c} c_2$, but not $c_2 \xrightarrow{c} c_1$. It is not transitive: after a certain number of levels of indirection, connotation doesn't continue to hold. But connotations will often have a small *radius of transitivity*, with concepts two or three connotative links away still connoting the original concept. This idea of a radius of transitivity might be useful in practical applications, e.g. in determining whether to hold a particular concept in a set of indirect connotations of a particular concept. Determining the size of this radius is an empirical question.

There are many kinds of connotations. We can break these down into a number of classes, each of which can be constructed using one or more sources of information as a computational proxy:

- The first is *similarity*, or *syntactic* connotation. That is, a point p in S could look or sound similar to c . This could be realised in some computational system via a pattern-matching system.
- p could have some *semantic similarity* to c , for example words that have the same meaning, or images that have been tagged with the same tag. Sources for this include explicitly curated databases such as WordNet, TrueKnowledge and OpenCyc, and collaboratively constructed semantic resources [59] such as Wikipedia.
- p could be commonly found where c is commonly found; we might term this *pragmatic* or *phenomenological connotation*. Again, sources such as Wikipedia, or just general web searches and autocompletion databases offer examples of this.
- p might be a *part* of c (or vice versa). Again, systems such as WordNet and OpenCyc attempt to curate these kinds of links.

We should note that doing this well would require empirical backup. Before we could be confident about the ability of these proxies to stand in for connotations, we would need to do some studies with users to see how good the proxies are at creating viable connotations.

Another issue is that of *normalisation*. In a space of connotations there are likely to be very generic basins of attraction. This is illustrated by the well-known property of the structure of the link-digraph in Wikipedia. It is said that if you start from a random page in Wikipedia, follow the first substantive link on the page, then follow the first substantive link on the following page, and so on, then eventually you will reach the entry for *philosophy*. Another example from Wikipedia is the category-entry for *living persons*, which is clearly connected to many entries and carries little useful connotative knowledge. The danger with this is that if a connotation space is built on such resources, then these dominate the search. Ideally, we want to search through points that are broadly at the same level of abstraction, and not go through chains that lead us into very generic classes and out again into a more-or-less arbitrary part of the space. One possibility for doing this would be to eliminate from the connotation space links with large in-degree, or some refinement thereof.

4.2 Exploring a Connotation Space

Secondly, we need some way of exploring this space. There are many approaches to this, we will sketch one here, which takes some ideas from the graph-search algorithms known as *Ant Colony Optimization* (ACO) [60, 30], and from *Stochastic Diffusion Search* (SDS) [61].

We start from a number of points in the connotation space—call these seed points. These reflect aspects of the final artefact that we would like to search for. In the example discussed earlier of designing a logo for an organisation, these points could be words or images that capture the values of an organisation, gained e.g. via focus groups, values-elicitation exercises, and mood boards.

The next stage is a process of exploring out from these by following edges in the connotation space that are adjacent to those seed points towards new nodes, and then looking at the edges that are adjacent to those new nodes. At each timestep in the search, a random selection of nodes would be explored, perhaps with some bias towards edges that are closer to the seed concepts to reflect the “radius of transitivity” of the connotation relation as discussed earlier (here reinterpreted as a fuzzy set rather than a strict radius). Each edge in the connotation space has a weight, initially zero. As this process progresses we increase these weights every time an edge is traversed—this is similar to the pheromone deposition process in ACO [60], or to the diffusion of hypotheses in SDS [61]. This process is illustrated in Figure 1.

Just exploring these will rapidly lead to exploring the whole of the connotation space, so we need some way of removing links (i.e. setting their weight to zero) or reducing their weight. One

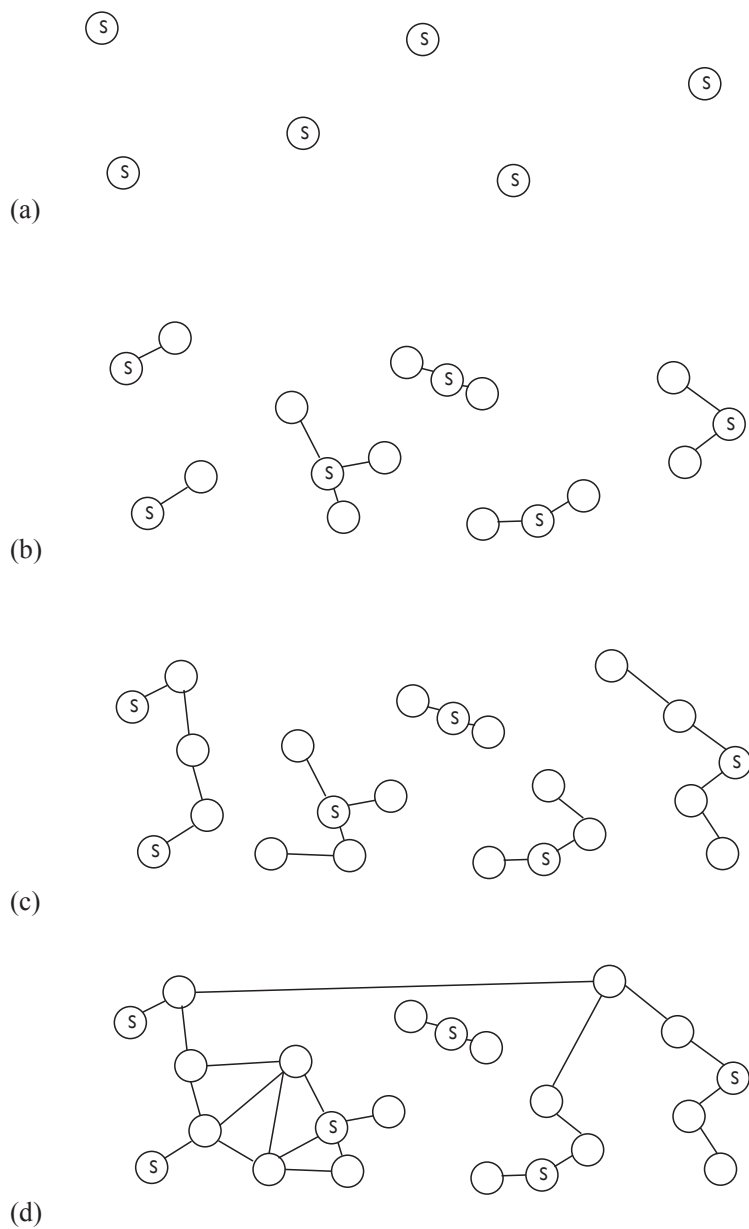


Figure 1: Sketch of a connotational search algorithm. (a) Starting with the seed points; (b) Sampling immediate connotations of those seed points; (c) Exploring further into the connotation graph; and, (d) links forming between different areas of the connotation graph.




Figure 2: Dense links and multiple paths. Seed points are doubly-circled. *Oak tree* is a densely connected point, being connected to various other concepts. *Economic strength* and *tree* are multiply linked, having two independent paths.

way would be to adopt the *evaporation* process used in ACO, where at each timestep edges with positive weights have a small weight reduction—i.e. in order for edges to be active in the long run, they must be being reinforced by another process. A alternative would be to remove specific edges, such as those that are only linked to one other point, after a certain number of timesteps.

Furthermore, we want to reward links that have some pattern that is suggestive of value. One way to do this is to identify *dense points* in the graph, i.e. nodes with several positively-weighted adjacent edges in the connotation graph. This has been used in the analysis of static *concept graphs* by Berthold and colleagues [53, 54] as an interpretation of the idea of *bisociation* of concepts discussed by Koestler in his account of the creative search process [55]. Edges adjacent to these could be strengthened. This is illustrated in Figure 2. This idea of testing particular points is found in SDS—but, in that algorithm this is with regard to an extrinsic problem, whereas in this idea it is the intrinsic structure of the graph locality that is being explored.

A second way to strengthen links would be to look for patterns where there are multiple links (i.e. edge-sequences) between a seed concept and another concept in the space, capturing the idea that the second concept is connected to the seed concept in multiple different ways. This is also illustrated in Figure 2.

A third way takes its inspiration from the lateral thinking techniques of de Bono [62]. One of his techniques for exploring an idea is to pick a random word and work back to the problem at hand from that word, with the idea that connections that are made *en route* might provoke a solution to the original problem. One way of realizing this in this framework would be to pick

a random point in the connotation space, and follow a list of connections back to the original seed points that define the problem at hand. Perhaps this could be done a number of times, and overlaps in the paths further strengthened. This would be computationally challenging as there is no *intuition* for the algorithm to follow about how to get back to the seed concept, by contrast with a human following the same process.

These three examples are brief sketches—they would need to be verified against existing examples to be of validity. A very different approach would be to apply machine learning to extract patterns in connotation graphs that exemplify existing human creativity, or that have been artificially constructed to show the kind of creativity that we are trying to automatically identify.

One potential criticism of this whole approach is that we are just doing what is done in the denotational systems, but that our denotation is just at a higher level of abstraction. Instead of using some computational proxy to denote the creativity of a particular outcome, we are instead creating a computational proxy that denotes the creativity of patterns in the connotation graph.

Another criticism is that all this exploration can be wrapped up into an objective function—the function could statically look for connotational links with the current outcome candidate and assign a score based on that connotational strength. However, this seems a less rich process than the one suggested above, as it is essentially *memoryless*—it relies on a static network of connotative connections, whereas in the search framework discussed above connotations themselves can develop their strength of links as the search unfolds.

4.3 Selecting the Final Outcome

Finally, we need some way of picking the best *o*-points out from the final constructed structure. One approach would be to stop the exploratory phase after a certain point—this could happen when weight-changes have stabilised, or after a certain number of timesteps—and to extract the most strongly-connected *o*-points in the connotation graph.

Another approach would be to increase the strength of the forgetting process until very few links remained. Then, the *o*-points adjacent to those links would be selected. A similar idea would be to use some kind of competitive-learning between *o*-points, similar to the *k*-winner-take-all concept in neural networks [63], where output neurons have competitive interactions to encourage the reduction of activity in more weakly-represented outputs and strengthen the *k* strongest outputs.

4.4 Applying Connotational Search

Algorithms such as the one sketched in this section could be applied in a number of different ways

The first is essentially the idea that we have sketched above, which is to use connotational search to work outwards from a number of seed points that are provided as a starting point. These could be provided by a human, or via some automated process. The example given above of automated extraction of a strongly-connotative key concept for a logo is an example of humans providing the seed points.

Another example illustrates how this could be applied in a more automated fashion. This starts from the work of Krzeczowska et al. [64], who have described an extension to Colton’s *Painting Fool* system that attempts to auto-generate artworks that are inspired by the day’s news. Their system works by extracting keywords from news stories, using internet image search to find images related to those words, and then applying the *Painting Fool* to generate a visually coherent collage of those images. One way to make this richer might be to use a connotational search where the seed points are, say, keywords from news stories, and the pictures used in the collage include images that are linked via connotations to two or more of the news stories.

A greater challenge would be to explore the connotation space without a specific objective in mind, just to find something with some degree of *depth*. This could involve finding pairs of concepts that are well-connected by a certain kind of connection, and then seeing if those can be

connected via another set of unrelated connections.

A third context for the application of these ideas is as an *adjunct* to a problem-focused search. Consider a system such as *genetic programming* [65], which uses an evolutionary algorithm to evolve a program or circuit to solve a problem. Typically, the problem would be specified via a training set of desired input/output behaviour. The algorithm progresses by taking an initially random population of programs or circuits, and then moving towards a solution via a large number of discrete generations in which the best solutions so far, evaluated by running them on the training set, are selected and then mutated and crossed over to create a new solution.

This is purely denotational—as appropriate for a well-specified problem for this. However, connotational search might add something to this. Consider, for example, the idea of creating a set of additional problems that are similar to the original problem, and trying to solve them alongside the original problem. Alternatively, sub-problems that appear to be trying to be solved could be extracted and added to a *connotational set* of problems that are being tackled via the search, and which could guide the problem-solving of the GP system along new lines to get out of local minima. There is a danger, of course, that by creating a *set* of problems rather than a single problem, the search becomes unfocused. Therefore, this might be better suited to the kind of GP system that evolves a set of interlocked modules, such as the *systemic computation* system [66].

The fourth domain for the exploration of these ideas would be to use this alongside a human exploration of the connotation space. We can imagine an interface (perhaps a tabletop interface to allow various people to interact with and discuss the ideas being generated) where various points in the connotation space are being displayed, and users can touch those points to strengthen or weaken their interest in them. Alongside this, a connotation-exploration system along the lines sketched above is running, and further points in the connotation space are displayed; if they are of interest, the user(s) can click on them; if not, then after a while they will just fade away and be replaced.

5 Conclusions and Future Work

In this paper we have explored the idea of *connotation* in the context of computational creativity, and contrasting it with the idea of the more common use of *denotation* in such systems. A three-stage process for connotational search has been discussed, consisting of constructing a space of connotations, exploring that space, and then choosing outcomes from that space. Ideas for how such an algorithm could be implemented have been discussed, as have potential applications.

Clearly, an important piece of future work is to implement and evaluate algorithms based on connotation. Another direction is to explore how humans use various forms of connotation in their creative activities, and how well computational proxies for connotation are at capturing humans' connotational networks. It would also be interesting to explore this idea in specialised domains, for example in graphs of information about, say, bioinformatics (as constructed by Nagel et al. [54]). Finally, a different direction would be to explore the application of this as a *supplement* to other systems, both fully automated problem solving, and humans exploring a creative space, whether individual or collaborative.

References

- [1] Russell S, Norvig P. Artificial Intelligence: A Modern Approach. Prentice Hall; 2009. Third Edition.
- [2] Levin J. Functionalism. In: Zalta E, editor. The Stanford Encyclopedia of Philosophy; 2004 (revised 2009). Available at <http://plato.stanford.edu/>.
- [3] Colton S, López de Mántaras R, Stock O. Computational Creativity: Coming of Age. AI Magazine. 2009;30(3):11–14.

- [4] Newell A, Shaw JG, Simon HA. The Process of Creative Thinking. In: Gruber HE, Terrell G, Wertheimer M, editors. *Contemporary Approaches to Creative Thinking*. New York: Atherton; 1963. p. 63–119.
- [5] Maher ML. Computation, Creativity and Interaction; 2009. Background paper for CHI'09 workshop on Computational Creativity Support, available at <http://computationalcreativity.org/cameraready/maher.pdf>.
- [6] Duch W. Computational Creativity. In: *Proceedings of the 2006 World Congress on Computational Intelligence*; 2006. p. 1162–1169.
- [7] Cook M, Colton S. Automated Collage Generation—with More Intent. In: Ventura D, Gervás P, Harrell DF, Maher ML, Pease A, Wiggins G, editors. *Proceedings of the Second International Conference on Computational Creativity*; 2011. p. 1–3. Available at <http://iccc11.cua.uam.mx/proceedings/index.html>.
- [8] Cohen H. The Further Exploits of AARON, Painter. *Stanford Humanities Review*. 1995;4(2).
- [9] Eigenfeldt A, Pasquier P. Negotiated Content: Generative Soundscape Composition by Autonomous Musical Agents in *Coming Together: Freesound*. In: Ventura D, Gervás P, Harrell DF, Maher ML, Pease A, Wiggins G, editors. *Proceedings of the Second International Conference on Computational Creativity*; 2011. p. 27–32. Available at <http://iccc11.cua.uam.mx/proceedings/index.html>.
- [10] Whorley R, Wiggins G, Rhodes C, Pearce M. Development of Techniques for the Computational Modelling of Harmony. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 11–15. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [11] Rahman F, Manurung R. Multiobjective Optimization for Meaningful Metrical Poetry. In: Ventura D, Gervás P, Harrell DF, Maher ML, Pease A, Wiggins G, editors. *Proceedings of the Second International Conference on Computational Creativity*; 2011. p. 4–9. Available at <http://iccc11.cua.uam.mx/proceedings/index.html>.
- [12] Gervás P. Computational Approaches to Storytelling and Creativity. *AI Magazine*. 2009;30(3):49–62.
- [13] de Melo CM, Gratch J. Evolving Expression of Emotions Through Color in Virtual Humans Using Genetic Algorithms. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 248–257. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [14] Mizetič I, et al. Bisociative Knowledge Discovery for Microarray Data Analysis. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 190–199. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [15] Colton S. Automated Conjecture Making in Number Theory using HR, Otter and Maple. *Journal of Symbolic Computation*. 2005;39(5):593–615.
- [16] Wiggins GA. A Preliminary Framework for Description, Analysis and Comparison of Creative Systems. *Knowledge-based Systems*. 2006;19:449–458.
- [17] Wiggins GA. Searching for Computational Creativity. *New Generation Computing*. 2006;24:209–222.

- [18] Jennings KE. Search Strategies and the Creative Process. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. Proceedings of the International Conference on Computational Creativity; 2010. p. 130–139. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [19] Boden M. The Creative Mind. Abacus; 1990.
- [20] Boden MA. Creativity and artificial intelligence. *Artificial Intelligence*. 1998;103(1–2):347–356. Available from: <http://www.sciencedirect.com/science/article/pii/S0004370298000551>.
- [21] Romero JJ, Machado P, editors. The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music. Springer; 2007.
- [22] Johnson CG. Fitness in Evolutionary Art: What Has Been Used and What Could Be used? In: Machado P, Romero J, Carballal A, editors. Proceedings of EvoMUSART 2012. Springer; 2012. p. 129–140. Lecture Notes in Computer Science 7247.
- [23] Deb K. Multi-objective Optimization using Evolutionary Algorithms. Wiley; 2001.
- [24] Zadeh L. Optimality and Non-Scalar-Valued Performance Criteria. *IEEE Transactions on Automatic Control*. 1963;AC-8(59).
- [25] Kim IY, de Weck O. Adaptive Weighted Sum Method for Multiobjective Optimization. In: 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference; 2004. .
- [26] Mathur VK. How Well Do We Know Pareto Optimality? *Journal of Economic Education*. 1991;22(2):172–178.
- [27] Li C, Yang S. A General Framework of Multi-population Methods with Clustering in Undetectable Dynamic Environments. *IEEE Transactions on Evolutionary Computation*. 2012;In press.
- [28] Yang S, Ong YS, Jin Y, editors. Evolutionary Computation in Dynamic and Uncertain Environments. Springer; 2007. Studies in Computational Intelligence, Vol. 51.
- [29] Mitchell M. An Introduction to Genetic Algorithms. MIT Press; 1998.
- [30] Bonabeau E, Dorigo M, Theraulaz G. Swarm Intelligence. Oxford University Press; 1999.
- [31] Kowaliw T, McCormack J, Dorin A. An Interactive Electronic Art System based on Artificial Ecosystemics. In: Proceedings of the 2011 IEEE Symposium on Artificial Life. IEEE Press; 2011. p. 162–169.
- [32] Yu H, Riedl MO. A Sequential Recommendation Approach for Interactive Personalized Story Generation. In: Proceedings of the 11th International Conference on Autonomous Agents and Multi Agent Systems; 2012. .
- [33] Tokui N, Iba H. Music Composition with Interactive Evolutionary Computation. In: Proceedings of the 2000 Conference on Generative Art; 2000. .
- [34] Marques VM, Reis C, Tenreiro Machado JA. Interactive Evolutionary Computation in Music. In: Proceedings of the 2010 IEEE International Conference on Systems, Man and Cybernetics. IEEE Press; 2010. .
- [35] Kowaliw T, Dorin A, McCormack J. Promoting Creative Design in Interactive Evolutionary Computation. *IEEE Transactions on Evolutionary Computation*. 2012;(in press).

- [36] Biles JA. Life with GenJam: Interacting with a Musical IGA. In: Proceedings of the 1999 IEEE Conference on Systems, Man and Cybernetics; 1999. p. 652–656.
- [37] Biles JA, Eign W. GenJam Populi: Training an IGA via Audience-Mediated Performance. In: Proceedings of the 1995 International Computer Music Conference; 1995. .
- [38] Peinado F, Francisco V, Hervás R, Gervás P. Assessing the Novelty of Computer-Generated Narratives Using Empirical Metrics. *Minds and Machines*. 2010;20:565–588.
- [39] Bundy A. What is the difference between real creativity and mere novelty? *Behavioural and Brain Sciences*. 1994;17(3):533–34.
- [40] Ritchie G. The Transformational Creativity Hypothesis. *New Generation Computing*. 2006;24:241–266.
- [41] Colton S. Creativity Versus the Perception of Creativity in Computational Systems. In: *Creative Intelligent Systems: Papers from the AAAI Spring Symposium*; 2008. p. 14–20.
- [42] Csikszentmihalyi M. *Flow: The Psychology of Optimal Experience*. HarperPerennial; 1991.
- [43] Geng L, Hamilton HJ. Interestingness Measures for Data Mining: A Survey. *ACM Computing Surveys*. 2006 September;38. Available from: <http://doi.acm.org/10.1145/1132960.1132963>.
- [44] Machado P, Romero J, Santos ML, Cardoso A, Manaris BZ. Adaptive Critics for Evolutionary Artists. In: Raidl GR, Cagnoni S, Branke J, Corne D, Drechsler R, Jin Y, et al., editors. *EvoWorkshops*. vol. 3005 of *Lecture Notes in Computer Science*. Springer; 2004. p. 437–446.
- [45] Greenfield G, Machado P. Simulating Artist and Critic Dynamics—An Agent-based Application of an Evolutionary Art System. In: *Proceedings of the 2009 International Joint Conference on Computational Intelligence*; 2009. p. 190–197.
- [46] Cumming L. Pipilotti Rist: *EyeBall Massage*; Barry Flanagan: *Early Works 1965–82—Review*. *The Guardian*. 2011; Available at <http://www.guardian.co.uk/artanddesign/2011/oct/02/pipilotti-rist-barry-flanagan-review>.
- [47] Celan P. *Collected Poems*. Penguin; 1996.
- [48] Hadamard J. *The Psychology of Invention in the Mathematical Field*. Princeton University Press; 1945.
- [49] Webb Young J. *A Technique for Producing Ideas*. McGraw-Hill; 2003; original edition 1943.
- [50] Lawson B. *How Designers Think*. Architectural Press; 2005. 4th Edition.
- [51] Wallas G. *The Art of Thought*. Jonathan Cape; 1926.
- [52] Veale T. We Can Re-Use It For You Wholesale: Serendipity and Objets Trouvés in Linguistic Creativity. In: Ventura D, Gervás P, Harrell DF, Maher ML, Pease A, Wiggins G, editors. *Proceedings of the Second International Conference on Computational Creativity*; 2011. p. 22–26. Available at <http://iccc11.cua.uam.mx/proceedings/index.html>.
- [53] Kötter T, Berthold MR. (Missing) Concept Discovery in Heterogeneous Information Networks. In: *Proceedings of the Second International Conference on Computational Creativity*; 2011. p. 135–140. Available from: <http://www.inf.uni-konstanz.de/bioml2/publications/Papers2011/KoBe11.pdf>.

- [54] Nagel U, Thiel K, Kötter T, Piatek D, Berthold MR. Bisociative Discovery of Interesting Relations Between Domains. In: *Advances in Intelligent Data Analysis X: 10th International Symposium. Lecture Notes in Computer Science Volume 7014*; 2011. p. 306–317. Available from: <http://www.inf.uni-konstanz.de/bioml2/publications/Papers2011/NTKP+11.pdf>.
- [55] Koestler A. *The Act of Creation*. Macmillan; 1964.
- [56] Norton D, Heath D, Ventura D. Establishing Appreciation in a Creative System. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 26–35. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [57] Jennings K. Developing Creativity: Artificial Barriers in Artificial Intelligence. *Minds and Machines*. 2010;20:489–501.
- [58] Indurkha B. On the Role of Metaphor in Creative Cognition. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 51–59. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [59] Gurevych I, Wolf E. Expert-Built and Collaboratively Constructed Lexical Semantic Resources. *Language and Linguistics Compass*. 2010 November;4(11):1074–1090.
- [60] Dorigo M, Stützle T. *Ant Colony Optimization*. MIT Press; 2004.
- [61] Bishop JM. Stochastic Searching Networks. In: *Proceedings of the First IEE Conference on Artificial Neural Networks*; 1989. p. 329–331.
- [62] de Bono E. *de Bono’s Thinking Course*. Crown Publications/BBC; 1982.
- [63] Maass W. On the Computational Power of Winner-take-all. *Neural Computation*. 2000;12:2519–2535.
- [64] Krzeczowska A, El-Hage J, Colton S, Clark S. Automated Collage Generation—with Intent. In: Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. p. 36–40. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [65] Poli R, Langdon WB, McPhee NF. *A Field Guide to Genetic Programming*. Published via <http://lulu.com> and freely available at <http://www.gp-field-guide.org.uk>; 2008. (With contributions by J. R. Koza). Available from: <http://www.gp-field-guide.org.uk>.
- [66] Bentley PJ. Systemic Computation: A Model of Interacting Systems with Natural Characteristics. *International Journal of Parallel, Emergent and Distributed Systems*. 2007 April;22(2):103–121.
- [67] Ventura D, Pease A, Pérez y Pérez R, Ritchie G, Veale T, editors. *Proceedings of the International Conference on Computational Creativity*; 2010. Available at <http://creative-systems.dei.uc.pt/icccx>.
- [68] Ventura D, Gervás P, Harrell DF, Maher ML, Pease A, Wiggins G, editors. *Proceedings of the Second International Conference on Computational Creativity*; 2011. Available at <http://iccc11.cua.uam.mx/proceedings/index.html>.