

Search and notions of creativity

Colin G. Johnson

Computing Laboratory

University of Kent at Canterbury

Canterbury, Kent, CT2 7NF, England

C.G.Johnson@kent.ac.uk

Abstract

In this paper we consider a number of issues concerned with how creativity can exist within a paradigm of computational search. We look both at creativity within the search process and the capacity of the search space to facilitate creativity. Within the latter we consider notions of compressibility and openness in search spaces; the influence of modularity and substructures on creativity; and the distinction between items in a search space *denoting* an external reference versus *connoting* links within and beyond the search space. We go on to discuss whether there really is a true distinction between transformational and exploratory notions of creativity within the context of search. Finally we offer some open questions in this area.

1 Overview

The tensions between *search* and *creativity* is important in understanding the extent to which machine creativity is possible. The aim of this paper is to clarify a number of issues and raise a number of questions concerning this. In the first substantive section of the paper we give an overview of the connections and conflicts between these two notions. We then consider several issues concerning the structure of the search space, centred around a notion of *openness* and the extent to which such a notion is similar to and different from ideas of finiteness, structuredness and compressibility. We move on to consider the idea that many aspects of creativity in search may be linked to the idea of points in the search space being treated in a *connotative* fashion by contrast to the *denotative* style associated with traditional search. Finally we consider the well-known notions of exploratory and transformational creativity in the light of this discussion.

2 Introduction

2.1 The search stance in AI

A common stance in AI research is that of casting various problems associated with the action of mind as *search* problems. That is, we assert that the problem at hand “might as well be” a problem of searching some space.

This can be seen as an aspect of the functionalist stance in understanding mind. The core aspect of functionalism is a substrate-agnosticism which hypothesises that the phenomena that we term “mind” are capable of being realised equally well on any substrate that has sufficient computational richness. A second aspect of functionalism, which will be important in this paper, is a kind of quasi-behaviourist *process-agnosticism*, viz. that we can ascribe (an aspect of) mind to any process that produces similar¹ outputs/behaviour from similar inputs/environment. The details of the process that produces those outputs is considered irrelevant. Hence the stance of approaching mentality as a computational search problem: regardless of the actual process taken by human minds, we can explain the important features as a search problem.

There are a number of reasons why we might want to make this replacement/simplification. Firstly, we might be interested in pragmatic AI, i.e. producing mind-like phenomena for some practical purpose. In these cases the process is irrelevant; as long as we can achieve the results, it doesn't matter how we get there. Secondly, we might be interested in understanding the actions of the human mind. In this case the abstraction of mentality to a search process needs to be argued more carefully. For example we might validly make arguments that are based solely on the structure of the input-output map, and use the search model as a mapping between these, purely because it is the easiest way to make the explanation, whilst accepting that any other equivalent process of producing that map is valid. Nonetheless it seems that we must be more careful with this process-agnosticism than we should be with substrate-agnosticism.

For the remainder of this paper we shall take it as a premiss that relevant aspects of mind can be reproduced with all essential features intact as a search process.

This model works well as an explanation of many mental phenomena. For example we can model simple decision-making processes such as finding a route around a building or choosing what to eat for lunch as a process of reducing the state space (deterministically or stochastically) of possible routes/lunches until we are left with a single point.

¹the details of what is meant by “similar” are rather obscure, particular when we start to consider creative processes.

2.2 Search and creativity

Nonetheless, when it comes to providing an explanation for *creative* mental processes this notion of search breaks down somewhat. Indeed there seems to be a tension between the two notions of “search” and “creativity”.

There are a number of different concepts wrapped up in the idea of machine creativity. The first of these is that of computational support for human creativity, for example in systems that use human feedback to guide the direction of a search process. This is not what concerns us here, and we shall not pay any further attention to this in this paper. The second is the idea of “pure” creativity, the creation of new things which in some way go beyond simple variations of what is currently known. This will be the main focus of this paper. A third notion is that of creatively responding to some external stimulus, i.e. using that stimulus as an inspiration for a creative act. This has been less well studied in the world of machine creativity, and whilst deserving of attention will not be the focus here.

3 Creativity in the search process

Typically we talk about “being creative”; that is, we suggest that it is the *process* that is where the creativity happens. One of the arguments in this paper is that search-based creative processes depend on the structure of the search *space* as well as the search *process*. Nonetheless we shall begin with a brief discussion of creative search processes.

One characteristic of creative search processes is that the criteria for evaluation of where to make the moves in the search space are not easy to capture in rulebound form. Consider for example searching a space of melodies for “interesting” or “tuneful” melodies. Whilst we might be able to suggest a number of heuristics which might happen (for a while) to generate such melodies, such heuristics will eventually run out of steam. To continue to generate new melodies, we need to be generating new heuristics.

One approach to creativity has been to create meta-level heuristics which evaluate potential search moves in terms of general characteristics of interesting results. This has been used for example in systems such as AM (Davis and Lenat, 1982) and HR (Colton *et al.*, 2000; Colton, 2002) that search for “interesting” patterns of numbers. For example in Colton’s HR one way of assessing the interestingness of mathematical conjectures is to ascertain whether the conjecture makes use of concepts that do not otherwise appear together in other conjectures—*viz.* that the concept is *surprising*. A number of similar ideas have been explored in the data mining literature (see (Freitas, 1999) for a survey), for example assessing the interestingness of a rule that summarizes data by looking both at its performance on the dataset and its similarity to a database of rules/expert system that represent “common knowledge”.

Another approach is search strategies that are driven by diversity rather than optimisation (Johnson, 2001b). An example of this is work on diversity-driven genetic algorithms (Johnson, 2001a), which use a kind of tabu search where the tabu list is not simply as a record of where the algorithm has visited previously, but where it has visited and found items

in the search space which are similar to items that have already been found. Thus the algorithm is driven on to constantly explore areas of the search space by negative pressure against returning to well-explored areas, rather than by a positive pressure to find a particular goal.

4 Spaces that facilitate creative search

We turn now aspects of the search *space* rather than *process*; how can we more carefully characterise the seeming *openness* of creative search?

One characteristic of open search spaces is that they have some complexity which belies “easy” search, as exemplified by exploration of the space of designs for mechanical devices or electrical circuits. Even though an exhaustive search would turn up the same result as a “creative” search, both the size of the search space and the complex structure thereof. For example it is not possible for a “naive” thinker to conceive of how to specify and order the “all possible” designs.

Another is because the search space is seen as being extensible. Boden (1990; 1998) (see also (Wiggins, 2001, 2003) for an elaboration and formalization of these ideas) has made a distinction between *transformational* and *exploratory* notions of creativity. Exploratory creativity consists of exploring an existing set of possibilities in a creative fashion; whereas transformational creativity consists of transforming the space that is explored.

As an example of this consider the idea of searching a space of melodies as discussed above. In order to search this space, we will need to give a description of what a “melody” is—e.g. a sequence of notes in a particular key. However this definition has limitations: what about a melody that changes key half way through? So we expand the search space to include such melodies, then The search space can (seemingly) always be extended.

5 Is this just about finite vs. infinite spaces?

This appears to be a more subtle distinction than merely saying that openness is just exploring infinite spaces. Plenty of infinite spaces are not open in this sense: a traditional optimization problem (e.g. maximizing a function defined over all integers) requires the exploration of an infinite search space; however this is not an open, creative process. However the exploration of a finite space can sometimes have this character of creative openness. For example the space of all melodies consisting of a certain number of notes, or all sentences of a certain length drawn from a finite dictionary are both finite processes: nonetheless there is considerable scope for creative search in both of these spaces.

6 Size and structure of search space

What, then, characterises search spaces that are ripe for creative exploration?

Perhaps in some cases this is just a matter of performance of search relative to human cognitive capacity to perform the creative act. It has been suggested (Miller *et al.*, 2002) that the creation of complex circuits is a creative act (though this does have a definite problem solving outcome). However this

may just be a human-parochial point of view. If we were able to explore circuit space much more rapidly in our minds, then perhaps the creative aspect of this would diminish. A related example is that in a typical function optimization problem we do not see the algorithm as being creative: is this because we are (at least theoretically) able to “step back” from the function and “see” the optimum, so finding the optimum doesn’t seem a creative act?

6.1 Structure of the search space: compressibility

One idea is to explore connections with information-theoretic notions of **compressibility**. Non-open search spaces admit a simple description, both in terms of describing the space and describing the effect of members of the space on the “fitness” of that member of the space. For example in the optimization problem each member of the search space translates into a value of the function being optimized: there is nothing rich about the structure of the fitness space, even for a complex function.

However for an open space the space generated by the translation of the search space values into “fitness” space has a richness redolent of incompressibility in information theory. We are not surprised to find areas of this space which are unlike any other areas in the space, and the structure of such space is typically complex, with regions of very different sizes and shapes having particular qualitative characters.

6.2 Structure of the search space: modularity

Another idea is that spaces that are open for creative search have an implicit **modularity** in their structure, and that exploratory creativity consists not just of finding interesting/surprising/whatever points in the space, but in discovering how to exploit the modular building blocks in the space in an intelligent way.

For example one reason that we may ascribe creativity to the process of circuit design is that we see this as a process of breaking down a big problem into subproblems, and working out how those subproblems fit together, and then tackling those subproblems; as we go through this process we also acquire knowledge about common subproblems and solutions thereto, so we accumulate a knowledge-base about the problem. Is it this richness that we term “creativity” in these kinds of problems, as opposed to e.g. a function optimization problem, which seems just to have one level of structure?

Similarly with a more open-ended creative process like writing melodies there is a search space that has a rich structure, in particular there is a rich structure of different kinds of *similarities* between melodies. One melody might have a similar overall arc to others; it might have a similarities in terms of its patterns of tension and release to others; it might conclude in ways that are similar to others. This contrasts with the way in which items in the search space of a traditional non-creative search space such as combinatorial or function optimization have a single measure of distance, e.g. in optimizing a function the difference between the fitness measure of the two items.

6.3 Structure of the search space: connotation

Another, related, aspect of a search space where creative search can occur is that there are many **connotations** between items in the search space. Whilst only one search space item might be presented as the end result of the creative process, the rich structure of the search space means that the person presented with that item is reminded of many other items in the search space.

An example of this is in the process of **distillation** that occurs in many kinds of artistic creativity². 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.

Much of the discussion in this section can be summarised by the distinction between *denotation* and *connotation*. In a traditional AI search processes, evaluation of the moves through the search space is purely based on what the item *denotes*. Perhaps in (some kinds of) creative search processes the search is driven primarily by what each item in the search space *connotes*. What might a fitness function based on this notion of connotation look like?

7 Is there really a difference between exploratory and transformational creativity?

In this final section we explore two challenges that attempt to undermine the traditional difference between exploratory and transformational creativity.

7.1 Constraints on transformation and grades of possibility

The first is the simple assertion that all creativity takes place within a system that is subject to some constraints. For example all human creative actions-in-the-world take place within the “search space” of all possible physical actions. We can conceive of many creative actions which are not realizable due to the physical constraints of the real world, many of which are artistically appealing.

This suggests that there can be limits to the scope of the possibility of the transformation in an act of transformational creativity. Yet perhaps this is not a problem. Perhaps it is not necessary to have *all possible* transformations available in order to validly have a theory of transformational creativity; indeed what is the scope of “all possible” search spaces that a transformation of search space could take us to?

²Many thanks to Michael Finnissy for discussions on this topic

Can we really meaningfully comprehend the space of “all physically possible actions” as a search space, or have we reached a stage by this point where being able to imagine actually searching this space is impossible, so the usefulness of the search space analogy breaks down? This has been noted by Dennett (1991): “sometimes an *impossibility in fact* is theoretically more interesting than a *possibility in principle*”.

All creative action can be seen as the exploration of *possibilities*. Can we learn anything, therefore, from theories of possibility such as Dennett’s (1995) theory of *grades of possibility*? This theory suggests that possibilities can be placed in a hierarchy. The original hierarchy is this: actual, historical, biological, physical and logical; additional intermediate levels can be introduced. The idea is that any possible action or object in the world can be placed at one of these levels: the action/object being impossible at lower levels and the higher levels being unnecessary. As an example Dennett discusses the notion of a flying horse: this is physically possible in some form; however it is biologically impossible due to constraints of energy expenditure, bone strength, et cetera. A carnivorous horse, however, is biologically possible; it is just not *historically* possible, as horse populations never found themselves in a situation where meat-eating was necessary or advantageous.

We can see all creative actions (which are realised in the world) as being manipulations of levels of possibilities between the physical and the actual. Thus we can view this either as a transformational process (when we make a transformation in search space we are moving “upwards” in this hierarchy) or explorational (whatever we are doing, we are just messing around in the space of physically realizable possibilities).

7.2 Transformation within exploration via substructure-building

A second challenge to the distinction between exploratory and transformational creativity comes from search spaces with the capacity for building substructures, as discussed earlier. Such a space would traditionally be seen as offering scope for exploratory creativity. However, when such substructures are capable of being built, perhaps a kind of transformational creativity is possible. Is building a vocabulary and grammar of substructures within a large search space more transformational than a pedestrian transformation of a trivial search space?

8 Questions

We end with a few unresolved and open questions.

- What sorts of search processes and search spaces are amenable to creativity?
- Are there kinds of creativity that cannot be understood from within the search paradigm?
- Is the process-agnosticism of the search stance always justified?
- Is the distinction between transformational and exploratory theories of creativity always justifiable?
- Should we pay more attention to the differences between sizes of search spaces when we are explaining things using the search stance?
- Is the paradigm of moves based on evaluation of search-space members always the best way to understand the exploration of search spaces?
- What role does modularity and substructure-capacity of search spaces play in this?
- What would a search based on connotation rather than denotation look like?
- Can we formalise some of these arguments, e.g. using the framework of Wiggins (2001; 2003).
- How many of these distinctions are just limits on the ideas and analogies that we have for explaining creativity?

References

- Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. Abacus.
- Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence*, **103**(1–2), 347–356.
- Colton, S. (2002). *Automated Theory Formation in Pure Mathematics*. Springer.
- Colton, S., Bundy, A., and Walsh, T. (2000). On the notion of interestingness in automated mathematical discovery. *International Journal of Human Computer Studies*, **53**(3), 351–375.
- Davis, R. and Lenat, D. B. (1982). *Knowledge-based systems in artificial intelligence*. McGraw Hill.
- Dennett, D. (1991). *Consciousness Explained*. Penguin.
- Dennett, D. (1995). *Darwin’s Dangerous Idea: Evolution and the Meanings of Life*. Penguin.
- Freitas, A. A. (1999). On rule interestingness measures. *Knowledge-Based Systems Journal*, **12**(5–6), 309–315.
- Johnson, C. G. (2001a). Finding diverse examples with genetic algorithms. In R. John and R. Birkenhead, editors, *Developments in Soft Computing*, Advances in Soft Computing, pages 92–99. Physica/Springer-Verlag.
- Johnson, C. G. (2001b). Understanding complex systems through examples: a framework for qualitative example-finding. *Systems Research and Information Systems*, **10**, 239–267.
- Miller, J. F., Kalganova, T., Job, D., and Lipnitskya, N. (2002). The genetic algorithm as discovery engine: Strange circuits and new principles. In P. J. Bentley and D. W. Corne, editors, *Creative Evolutionary Systems*, pages 443–466. Morgan Kaufmann.
- Wiggins, G. A. (2001). Towards a more precise characterisation of creativity in AI. In *Proceedings of the ICCBR 2001 Workshop on Creative Systems, Vancouver, British Columbia*.
- Wiggins, G. A. (2003). Characterising creative systems. In *Proceedings of the IJCAI’03 Workshop on Creative Systems*.