

Kent Academic Repository

Full text document (pdf)

Citation for published version

Johnson, Colin G. (2005) Varieties of Openness in Evolutionary Computation. In: Zimmer, Robert, ed. Proceedings of the Workshop on Creative Evolutionary Computation. Goldsmiths College

DOI

Link to record in KAR

<https://kar.kent.ac.uk/14358/>

Document Version

UNSPECIFIED

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>

Varieties of Openness in Evolutionary Creativity (extended abstract)

Colin G. Johnson
Computing Laboratory
University of Kent at Canterbury
Email: C.G.Johnson@kent.ac.uk

January 2005

Introduction: what do we mean by *openness*?

One characteristic of systems in which creativity can occur is that they are *open*. That is, the space being explored appears to be (theoretically or pragmatically) unbounded, and there is no easy way in which the structure of the space can be simply summarized.

In this talk I would like to consider a number of questions that relate this characteristic of openness to the ability of evolutionary processes to explore spaces in a creative way.

Openness, creativity and search

The predominant view of evolutionary processes and similar is that they are *search* processes. That is, they move from point-to-point (or perhaps from populations-of-points...), using some information about the points that have previously been visited. Typically this search process is used (in computational operations research and similar areas) to find “good enough” solutions to complex optimization problems where an exhaustive search of the space would be computationally intractable.

Typically such applications would not be seen as embodying a creative process. In such applications the search process is typically seen merely as a pragmatic shortcut to finding a solution to a problem for which an exhaustive process would eventually find the perfect solution.

Nonetheless search processes can be applied to problems that are traditionally seen as creative. Typically such processes are seen as creative for one of three reasons. Firstly because the criteria for evaluation are not easy to capture in a rulebound fashion. An example of this is searching a space of melodies for “interesting” or “tuneful” melodies.

Secondly because the search space is seen as having some complexity which belies “easy” search. Examples of this ideas include the use of search to explore 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 (e.g. it is not possible for a “naive” thinker to conceive of how to specify and order the “all possible” designs).

Thirdly, because the search space is seen as being extensible. 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 always be extended.

It is these latter two characteristics which seem particularly to capture the idea of “openness” in creativity.

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 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.

Size and structure of search spaces

One way that we might begin to understand the problem described in the previous section is in terms of the structure of the search space. There is a connection here to compressibility, though the link is subtle. 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 this space is typically complex, with areas of very different sizes having a certain qualitative characters.

Searching for diversity, structure and interestingness

What implications does this have for creative search using evolutionary techniques? One idea is generating search strategies that are driven by diversity rather than by searching for an optimum. An example of this is our work on diversity-driven genetic algorithms, which combine GAs with a data structure similar to a tabu list in tabu search; however this tabu list is generated 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 (according to some user-defined measure of similarity) 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.

Another example of a search space which is open to creative exploration are the program spaces which are explored in genetic programming. They present a canonical example of an open search space: the space of all programs (provided a suitably liberal choice of operators and terminals is made!). However, despite the successes of this, it is still not easy to apply these techniques to large-scale programming. Perhaps what is needed is a search strategy that explicitly recognises the rich complexity of the “fitness space”, e.g. by taking inspiration from biological development or analysing them as programs rather than just distilling fitnesses down to simple evaluation of the program on test data.

Another concern which is related to these problems is whether it is possible to encode notions of “interestingness” in the fitness component of evolutionary search. We will describe a number of attempts that have been made to do this and examine how they can be fitted into an evolutionary framework.

These examples all illustrate that a creative, open search process is about both the search algorithm and the search space needing to have aspects of openness.