

No Free Lunch in the Search for Creativity

Dan Ventura

Computer Science Department
Brigham Young University
ventura@cs.byu.edu

Abstract

We consider computational creativity as a search process and give a *No Free Lunch* result for computational creativity in this context. That is, we show that there is no *a priori* “best” creative strategy. We discuss some implications of this result and suggest some additional questions to be explored.

Introduction

It seems natural to interpret the creative process, particularly in a computational context, as one of search. This has been done since the early years of thinking about computational creativity (Boden 1992; 1998), and, more recently, Wiggins has suggested a rather more concrete formalization of the idea (2006). Here we take up the idea of creativity as search and ask the question, “Is there a best creative (search) strategy?” Not surprisingly, perhaps, under some reasonable assumptions, we can show that the answer turns out to be an emphatic, “No.” To show this, we present a simple reformulation of some classical ideas from the search, optimization and machine learning literature known collectively as the *No Free Lunch* (NFL) theorems (Wolpert and Macready 1995; Wolpert 1996; Wolpert and Macready 1997).

For simplicity, we will limit our discussion to a discrete, finite domain D containing “artefacts” to be discovered. As is typical, we consider the problem of discovering novel, useful artefacts, and here we focus on the discovery process, attributing greater creativity to strategies that make quick discoveries. This is not unreasonable given the fact that with enough time, even exhaustive search can discover good artefacts, and these ideas have been formalized elsewhere (Ritchie 2007; Ventura 2008). Indeed, it is often suggested that part of creativity is an aspect of surprise (Boden 1995; Macedo, Coimbra, and Cardoso 2001), and another way to look at rapid discovery is as a surprising result (i.e., if an observer, unaware of the search strategy used, cannot produce the result nearly as quickly [or at all], they are likely to be surprised by the result).

Main Result

To begin with, we will consider the case that there is one best element $a \in D$, which we will call a^* . We are interested

in how long it will take a particular creative (search) strategy π to discover a^* (note that we include heuristics, background knowledge, etc. in the concept of search strategy). In the general case, π can be probabilistic, and so the number of steps j required to find a^* should be represented as a probability distribution. Also, since the creator employing the strategy may have experience, exposure to an inspiring set, etc., which we will represent as I as in (Ritchie 2007), this probability distribution can be conditioned on this, and we can write $P_\pi^{a^*}(j|I)$ to mean the probability, given I that strategy π will discover a^* in exactly j steps. Then, we are interested in the cumulative distribution function

$$C_\pi^{a^*}(n) = \sum_{j=0}^n P_\pi^{a^*}(j|I)$$

which gives the probability that π will discover a^* in n or fewer steps¹.

Ideally, we would like to find a strategy π^* such that

$$\forall a \in D, a = a^* \implies C_{\pi^*}^{a^*}(n) \approx 1$$

for some small, finite $n \ll |D|$. In other words, we would like to find a strategy that quickly discovers *the* artefact, no matter which artefact in the domain is *the* artefact. Some reflection should suggest that such a strategy is unlikely to exist. However, perhaps we can at least find a strategy π^+ that dominates all other strategies, so that

$$\forall \pi \forall a \in D, a = a^* \implies C_{\pi^+}^{a^*}(n) \geq C_\pi^{a^*}(n)$$

That is, perhaps there is a strategy that will at least find *the* artefact as fast or faster than any other strategy. Unfortunately, there is the further complication that, in fact, we do not know a^* , so we can not compute $P_\pi^{a^*}(j|I)$. Thus, we must sum over all possible $a \in D$, redefining our goal as finding a strategy π^+ that dominates all other strategies, independent of the artefact $a \in D$ for which we are searching, so that

$$\forall \pi, \sum_{a \in D} C_{\pi^+}^a(n)P(a) \geq \sum_{a \in D} C_\pi^a(n)P(a)$$

¹We have completely ignored here the structure of the domain D as well as the mechanism of strategy π ; both are abstracted into the probability distribution $P_\pi^{a^*}(j|I)$.

where $P(a)$ is shorthand for $P(a^* = a)$. Because we have no way of knowing *a priori* for which artefact we may be looking, we must, in essence, have a strategy that will find any artefact in the domain faster than any other strategy (at least in an expected sense, weighted by the likelihood). Of course, we do not know the likelihood distribution, $P(a)$, either, so for now we will assume $P(a)$ is uniform; that is, we will assume that all artefacts in D are equally likely to be *the* artefact we are seeking. The question now is, given this uniformity assumption, is there a “best” creative strategy? The following theorem says that no such strategy exists.

Theorem 1. *For a fixed, finite domain D , an integer $0 \leq n \leq |D|$ and any strategy π ,*

$$\sum_{a \in D} C_{\pi}^a(n)P(a) = \frac{n}{|D|}$$

Proof:

$$\begin{aligned} \sum_{a \in D} C_{\pi}^a(n)P(a) &= \sum_{a \in D} \sum_{j=0}^n P_{\pi}^a(j|I)P(a) \\ &= \sum_{a \in D} \sum_{j=0}^n P_{\pi}^a(j|I) \frac{1}{|D|} \\ &= \frac{1}{|D|} \sum_{a \in D} \sum_{j=0}^n P_{\pi}^a(j|I) \\ &= \frac{1}{|D|} \sum_{j=0}^n \sum_{a \in D} P_{\pi}^a(j|I) \\ &= \frac{1}{|D|} \sum_{j=0}^n 1 \\ &= \frac{n}{|D|} \end{aligned}$$

The first equality is by definition; the second is by assumption of uniformity; the next two are simple algebra; the fifth is because the probability that *some* $a \in D$ is found is unity; the last is obvious. \square

What the theorem says is that, in the absence of biasing information about the creative task, the probability of discovering a^* is independent of π , the search strategy². In other words, if we do not know anything about the creativity task, no creative strategy is to be preferred over any other.

Now, let us relax or remove some of our simplifying assumptions and ask if this makes a difference. First, we can consider the possibility that more than one $a \in D$ is desirable; that is, we are searching for any member of a set

²The dual version of this says that the expected number of steps required to find a^* is

$$E_{\pi}^{a^*}[n] = \frac{|D|}{2}$$

and, in particular, is independent of π .

$A^* \subseteq D$ of desirable artefacts³.

Since $P_{\pi}^A(j|I) \leq \sum_{a \in A} P_{\pi}^a(j|I)$ (with equality if the probabilities are independent), the consequent of Theorem 1 takes the form

$$\sum_{a \in D} C_{\pi}^{A^*}(n)P(a) \leq \frac{|A^*|n}{|D|}$$

and, notably, is still independent of the choice of π .

Next, we can consider the non-uniform case for $P(a)$. In this case, the consequent in the theorem statement requires an additional integral, taken over the continuous space of possible distributions, and the resulting form is

$$\int_{P(a)} \sum_{a \in D} C_{\pi}^a(n)P(a)dP(a) = \frac{n}{|D|}$$

In other words, not assuming anything about the probability distribution $P(a)$ has the same effect on our expected success as does assuming $P(a)$ is uniform.

Note that these generalizations naturally compose, so that we can make a statement about the distribution-free probability of finding one of multiple desirable artifacts⁴:

$$\int_{P(a)} \sum_{a \in D} C_{\pi}^{A^*}(n)P(a)dP(a) \leq \frac{|A^*|n}{|D|}$$

Finally, we can mention the case of non-stationary D (i.e. the case for transformational search). While we will not say much about this here, we will note that there are variations of the NFL theorems for optimization that treat the case of a changing objective function (Wolpert and Macready 1997), and similar results will likely hold for transformational search.

Discussion

On the one hand, anyone familiar with NFL-type results will not be surprised that one applies here. Indeed, even the original NFL theorems could be seen as, in some ways, “formalizing the obvious”. On the other hand, the result, whether

³This can be thought of in terms of a fitness function $f : D \rightarrow [0, 1]$ that measures the desirability of an element of the domain, and a threshold θ , such that $A^* = \{a | a \in D, f(a) > \theta\}$. Or, we can eliminate the hard constraint and compute with the fitness f more directly. In this case, rather than summing over the different elements for which we might be searching, we integrate over the different fitness values we might find, weighted by the probability of finding a domain element whose fitness is that particular value. Then, the probability of finding the $a \in D$ with the highest fitness (which we assume is 1), assuming the distribution of fitness values $P(f)$ is uniform becomes:

$$\int_{f=0}^1 C_{\pi}^f(n)P(f)df = \frac{n}{|D|}$$

⁴Our most general statement, distribution-free, directly including a fitness function, and making no assumption about the distribution of fitness values becomes:

$$\int_{P(a)} \int_{P(f)} \int_{f=0}^1 C_{\pi}^f(n)P(f)df dP(f) dP(a) = \frac{n}{|D|}$$

surprising or not, has profound implications for computational creativity and gives us a framework in which to discuss general principles.

For example, the characterization of creativity as search as probability, allows a statistical interpretation of many aspects of computational creativity, and, in particular suggests that the optimal (in the Bayesian sense) approach to any creative endeavor (that can be cast as search) is to use the following search strategy:

$$\pi(I) = \operatorname{argmax}_{a \in D} P(a)$$

However, it is not even clear what knowledge of $P(a)$ means; and, of course, even if we did somehow know $P(a)$, for any interesting domain D , explicitly implementing such a search is completely intractable. So, the obvious question is how to approximate the Bayes optimal search. Perhaps it is possible to dynamically choose the search strategy π^a for any $a \in D$ such that

$$\forall \pi \forall a \in D, C_{\pi^a}^a(n) \geq C_{\pi}^a(n)$$

In other words, we would like a method for biasing our search strategy towards *the* artefact we are looking for. Short of *a priori* knowledge of a , this at the least requires some meta-knowledge about the creative task in question that can be used to guide the choice of π^a . We note here the similarity to meta-learning in the field of machine learning and, further, suggest a close tie to the case of transformational search. If the domain D is transformed, becoming B , we must assume that a^* has likely changed as well, becoming $b^* \in B \setminus D$ (if not the case, how do we explain/justify the domain transformation?) If this is the case, we must have some mechanism of changing our search bias to match, switching from strategy π^{a^*} to π^{b^*} .

In the case of machine learning, the NFL result says, crudely, that no learning algorithm is better than any other over all possible learning problems. The standard rejoinder to this result is that, in fact, we don't (and Nature doesn't) care about all possible learning problems, many of which represent "learning" scenarios that are not interesting or do not represent "real-world" scenarios. This dogma is universally accepted in the field of machine learning, and does seem intuitively appropriate. Further, it leads to interesting questions about which problems are the "interesting" ones, and how can we tell, and, knowing this, how can we build learning algorithms that are biased toward these types of problems.

In our current discussion of computational creativity, the analogical argument would be either that we are not likely to be searching for any possible $a \in D$ (and thus universal quantification over a is too strong a constraint) or, perhaps that Nature will favor certain members of D (and thus universal quantification over $P(a)$ is too strong a constraint). This sort of argument, of course, leads to inquiries regarding which members of D might be interesting or which distributions $P(a)$ might represent Nature; however, it is, at least at this point, much less clear that, in fact, we can make such a claim for computational creativity.

Since our result states, essentially, that all search strategies are equally effective over all possible search problems, we are asking whether all search problems are in some way "interesting" and, as a result, whether all search strategies are valuable. If we content ourselves (for the moment) with equating creativity with search strategy, then, in turn, we are asking whether all creative approaches are valuable, or whether some can be shown to be inherently better than others.

One might be tempted to claim that for a specific D , that a^* is fixed and thus, that a domain fully specifies a "creativity scenario". This would, indeed, make further analysis somewhat more tractable; however, it is unlikely that such a strong assumption is reasonable (e.g., if D is the set of all possible paintings, creating the "ultimate" painting is not likely to be either temporally or spatially consistent; indeed, the very definition of D , of what constitutes a painting, is likely to change over time and very possibly across locales as well). Thus, it is not clear that this question can be answered even for a specific domain D , let alone in the more general case, but it is, certainly, an interesting question to consider.

Acknowledgments

This material is based upon work that is partially supported by the National Science Foundation under Grant No. IIS-0856089.

References

- Boden, M. 1992. *The Creative Mind*. London: Abacus.
- Boden, M. 1995. Creativity and unpredictability. *Stanford Humanities Review* 4:123–139.
- Boden, M. 1998. Creativity and artificial intelligence. *Artificial Intelligence* 103:347–356.
- Macedo, L.; Coimbra, E.; and Cardoso, A. 2001. Creativity and surprise. In *Proceedings of the AISB01 Symposium on Creativity in Arts and Science*, 84–92.
- Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:76–99.
- Ventura, D. 2008. A reductio ad absurdum experiment in sufficiency for evaluating (computational) creative systems. In *Proceedings of the Fifth Joint Workshop on Computational Creativity*, 11–19.
- Wiggins, G. 2006. A preliminary framework for description analysis and comparison of creative systems. *Knowledge-Based Systems* 19:449–458.
- Wolpert, D. H., and Macready, W. G. 1995. No free lunch theorems for search. In *Technical Report SFI-TR-95-02-010*. Santa Fe Institute.
- Wolpert, D. H., and Macready, W. G. 1997. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1(1):67–82.
- Wolpert, D. H. 1996. The lack of a priori distinctions between learning algorithms. *Neural Computation* 8(3):1341–1390.