Creative Search Trajectories and their Implications

Kyle E. Jennings
Department of Psychology
University of California, Davis
Davis, CA 95616 USA
kejennings@ucdavis.edu

Abstract

Creative search trajectories are chronologically organized intermediate products (such as sketches and drafts) from the creative process. We discuss what sorts of conclusions can be made when these trajectories show non-monotonic progress toward the final creation. We introduce several key distinctions that are often overlooked, and argue that two null hypothesis processes must be rejected before non-monotonicity can be claimed to support more complex processes. We show that these null hypotheses are in fact difficult to rule out definitively using the sorts of evidence that past research has offered.

The sketches, drafts, revisions, and rejected ideas that creators leave in their wake on the way toward great masterpieces offer glimpses of the mental processes that are responsible for their achievements. Ordered in time, these artifacts trace a trajectory through a mental space, and may be the signatures of the specific exploration strategies that differentiate great thinking from the mundane. Even the differences in creativity that can be observed among study participants may be explainable by tracing and analyzing the detailed steps taken with simple creative problems.

This approach—which we call trajectory analysis—borrows from the problem solving research of Newell and Simon (1972) and others. However, whereas most problem solving research uses problems that are concrete, objective, and formalizable in terms of specific states, operators, and goals, the creative problems faced by artists, writers, and scientists are less easily reduced to symbols, rules, and computational steps. Though there are examples where researchers have meticulously examined creative search trajectories using objective features (Weisberg 2004), most research has let subjective judgments stand in for complete formal details (Kozbelt 2006; Simonton 2007; Kozbelt and Serafin 2009; Damian and Simonton 2011). Despite taking different approaches, many of these results point to the conclusion that creative outcomes are not arrived at directly, but rather through the twists and turns of false starts and retraced steps. These observations have been taken as evidence that the creative process is tentative and experimental rather than deliberate and informed.

Given how complex it would be to devise comprehensive and objective descriptors of intermediate states in creative problems, it may not be practical to overcome all of the imperfections of subjective judgments. However, this does not mean that conclusions drawn about the creative process on the basis of these judgments do not need to be rigorously justified. While it seems intuitively clear that a process that is tentative, uncertain, or experimental would produce trajectories that do not move monotonically closer to the final solution (Simonton 2007; Damian and Simonton 2011) or that do not show monotonic improvement over time (Kozbelt 2006; Kozbelt and Serafin 2009), caution must be exercised before concluding that trajectories with these features could *only* have been produced by such processes.

This paper aims to clarify what can and cannot be concluded about the creative process on the basis of creative search trajectories. Though we are ultimately optimistic about the potential of this approach and advocate the view that creativity requires uncertainty and experimentation, we will show that existing evidence that creative search trajectories are non-monotonic is in fact compatible with very straightforward search processes, and that more care and precision must be used when analyzing search trajectories. We begin with a more detailed discussion of existing trajectory analysis approaches, and then describe in overview the distinctions that past work has overlooked. We then illustrate these distinctions by presenting simple but formally complete examples that demonstrate the need for caution when drawing the conclusion that non-monotonic search trajectories necessarily reflect something other than a straightforward process. We conclude by discussing the burden of proof placed on researchers who wish to infer the causes of non-monotonicity and by offering suggestions for future work.

Background

While there are many approaches that people have taken to using intermediate products to understand the creative process (Getzels and Csikszentmihalyi 1976; Finke, Ward, and Smith 1992; Hennessey 1994; Ruscio, Whitney, and Amabile 1998; Rostan 2010), this paper focuses specifically on techniques that characterize the nature of the changes between successive revisions of the work (Kozbelt 2006; Simonton 2007; Kozbelt and Serafin 2009; Damian and Simonton 2011). Common to these analyses is the prediction that creativity should be associated with non-monotonicity,

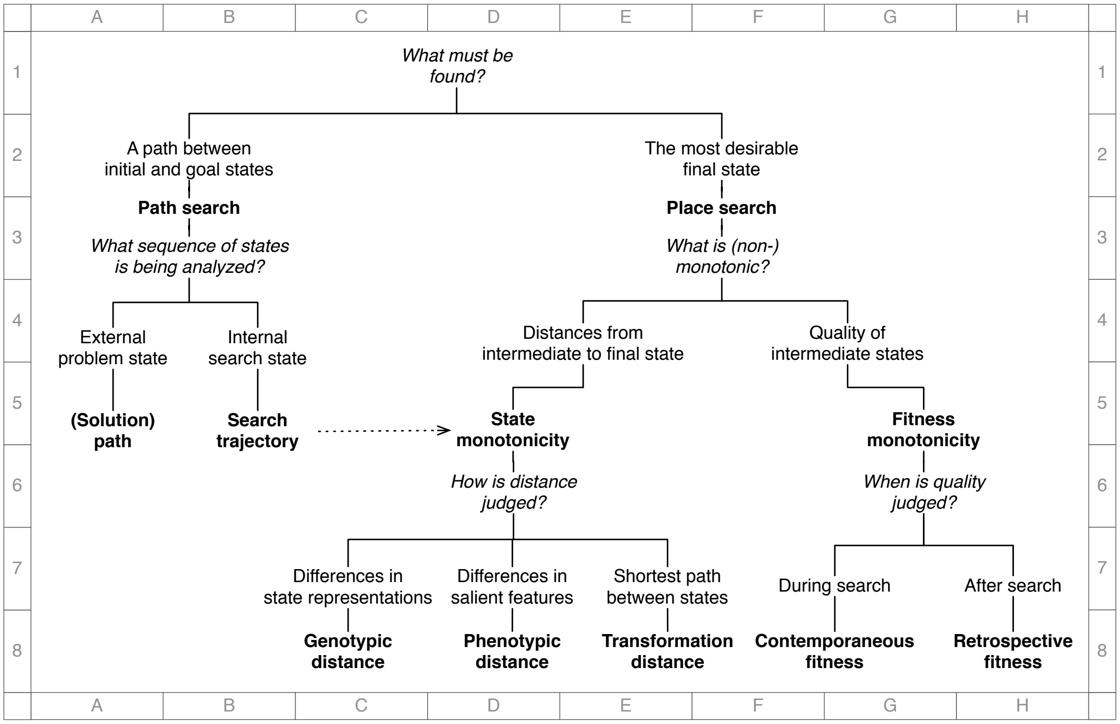

Figure 1: Depiction of the key questions, answers (italicized), and terms (bolded) that impact how search trajectories are analyzed.

which is perhaps best understood as the opposite of direct, incremental progress toward the final creation (other theories predict this, too, e.g., Perkins 2000). Beyond this commonality, the theories differ in their reasons that non-monotonicity should be expected, and indeed in how they operationalize monotonicity.

One approach to characterizing monotonicity is that taken by Kozbelt (2006) in his analysis of the sketches preceding Matisse's *Large Reclining Nude*. His approach focuses on the evaluation of each sketch, and in particular on whether the sketches become monotonically better with time. Given that Matisse's own evaluations are not available, Kozbelt has both artists and non-artists rate each sketch (presented in a random order) on 26 items, which are then analyzed in order to extract a latent quality dimension. The artists' ratings are found to be maximal at the final image, but beyond this they show non-monotonic variations over time (the contour of which is replicated in the non-artist sample). Similarly, Kozbelt and Serafin (2009) analyze intermediate sketches from drawings by non-eminent artists, and find that artwork that had been previously rated as more creative resulted from less monotonic trajectories. This is suggested to be due to an "interactive, hypothesis-testing dynamic" (Kozbelt and Serafin 2009, p. 358), though the specifics of this process are not articulated.

The other approach taken to characterizing monotonicity involves looking at the structure of the sketches themselves. The preponderance of this evidence stems from sketches Picasso left of *Guernica*. After an informal suggestion by Simonton (1999) that these sketches showed "false starts and wild experiments" (p. 197), Weisberg (2004) undertook a detailed analysis of the features shared between sketches and concluded that they were elaborations on a basic idea that itself had precedent in other work, including Picasso's own *Minotauromachy*. In response, Simonton (2007) undertook an alternative analysis in which various raters arranged the sketches in the order they would most logically have been generated, which reliably resulted in an order that did not match the actual temporal order. Later, Damian and Simonton (2011) had raters judge the similarity of the components from *Minotauromachy* to their counterparts in the *Guernica* sketches, and again found that the *Guernica* work did not get monotonically closer (or further) from the prototypes. Simonton claims this as evidence in support of the Blind Variation and Selective Retention (BVSR) theory of creativity (Simonton 2003; 2010), which holds that both desirable and undesirable variations will be generated during the creative process.

The Simonton work, in particular, has generated a great deal of recent controversy (Dasgupta 2011; Gabora 2011), much of it focused on the computational and algorithmic specifics of BVSR theory. In fact, neither Kozbelt's nor Simonton's analyses offered precise and detailed accounts of the processes that would lead to non-monotonicity, with Gabora (2011) suggesting that BVSR wouldn't even predict non-monotonicity. While Simonton has begun better formalizing BVSR theory (Simonton 2011; 2012), the fact is that there are basic problems for the trajectory analysis approach that any theory of the creative process must overcome. Therefore, rather than specifically addressing the claims by Kozbelt and Simonton, this paper will instead try to address some basic inconsistencies regarding how monotonicity is conceptualized and operationalized, and will demonstrate that non-monotonicity can result from processes that are more straightforward than most theories suggest.

Overview of Distinctions

In this paper we introduce several distinctions that are essential when analyzing search trajectories. As depicted in Figure 1, these distinctions revolve around questions about what the search must find, what aspect of trajectory monotonicity is of interest, and how monotonicity is measured. Throughout the paper we will define and illustrate the terms that these questions delineate, using the coordinates in the margins to reference to the relevant portions of the figure.

The first and most important question is what the search must find (Figure 1, D1). Problem solving research describes problems by the set of states and the operators that move between them. The problem solver's goal is to find a sequence of operator applications that transforms the initial state into (one of) the goal state(s). Because the goal state is already known, the thing being searched for is the path itself, which is why we refer to these as *path searches* (Figure 1, AB23). (See also Jennings, Simonton, and Palmer 2011.)

Because the goal states are well known in path search, solution quality depends more on the quality of the path than on the specific goal state reached, with shorter (or less costly) paths being better. Though this situation aptly describes many problems (e.g., proving a theorem, inventing a process for synthesizing a given protein) there are other problems where the end points are not known in advance. For instance, in painting the artist seeks to depict a certain scene, theme, or emotion using brush and paint. In most cases we compare paintings not by the set of brushstrokes that led from an empty canvas to the completed image, but rather by that image itself. Thinking of these final images as places in a solution space, we refer to this as a *place* search (Figure 1, EF23).

Creativity is possible with both path and place search, and most real problems involve some element of each (e.g., choosing a place and then finding the path to it). Though we'll speak of these as distinct kinds of searches in this paper, we recognize that understanding joint path-place search is an essential task for future work.

Monotonicity in Path Searches

Our discussion begins with path search. For simplicity we'll consider the classical Towers of Hanoi problem. (Though this problem leaves little room for creativity, it nicely illustrate our key points.) As depicted in Figure 2, there are three disks of decreasing size that are initially stacked on the leftmost of three pegs. The problem is to move the disks to the rightmost peg by moving one disk at a time without placing larger disks on top of smaller disks. The figure shows the shortest sequence of states that solves the problem, which together constitute the path found in this path search.

The Towers of Hanoi is often used to illustrate the failure of an heuristic called difference reduction, which entails iteratively applying the operator that most reduces the discrepancy between the current state and the goal state (Anderson 1993). In fact, solving this problem requires selecting operators that temporarily make the current state less similar to the goal state, and for this reason it could be argued that even a simple problem like the Towers of Hanoi has a non-monotonic solution. By this logic, there is no controversy behind claiming that *creative* problems exhibit non-monotonicity. However, we will see that the Towers of Hanoi isn't inherently non-monotonic, at least in sense that matters when making inferences about the search process.

Let's begin by looking at the monotonicity of the sequence of states forming the shortest path in Figure 2. Though we'll ultimately conclude that these are not necessarily the states that we should be analyzing when making inferences about the search process, they conveniently illustrate the different ways that monotonicity can be judged. The difference reduction heuristic works by economically comparing the current and goal states. For example, we could compare states by looking directly at their representations. Each state in Figure 2 can be described as an ordered triple indicating which pegs the smallest, middle, and largest disks are on. (This is sufficiently descriptive since larger disks cannot be on top of smaller disks.) Thus, the starting state is $(1, 1, 1)$, the final state is $(3, 3, 3)$, and the intermediate states are $(3, 1, 1)$, $(3, 2, 1)$, etc. By analogy to genetics, we can think of the state encoding as being a genotype. We'll define *genotypic distance* to be the dissimilarity between the encodings of two states, which here we can define as the sum of absolute differences in state representations (Figure 1, C78). For instance, the third state, $(3, 2, 1)$ differs from the final state, $(3, 3, 3)$, by $|3-3|+|3-2|+|3-1| = 3$.

We could also compare states by looking at their salient features, which may or may not be directly related to the state encoding. Again by analogy to genetics, we'll call this *phenotypic distance* (Figure 1, D78). For the Towers problem, let's calculate phenotypic distance by counting the number of disks that are on different pegs. Thus, the first and final states differ by three, the second and final by two, and so on.

As Figure 2 shows, the best solution to the Towers problem is indeed non-monotonic in genotypic distance (D_G) and phenotypic distance (D_P), with the solution becoming less similar to the goal at various points, and difference reduction would indeed fail with either of these methods for choosing operators. However, the fact remains that the sequence of moves shown is minimal (the path is the shortest path possible). Defining the *transformation distance* (D_T) as the length of the shortest path between two states (Figure 1, E78), then we can see in Figure 2 that the path does get monotonically closer to the final state. On the one hand this is a useless insight since performing difference reduction with D_T just pushes the work of finding the solution into calculating D_T. On the other hand this reveals that the solution really isn't non-monotonic, in the sense that there

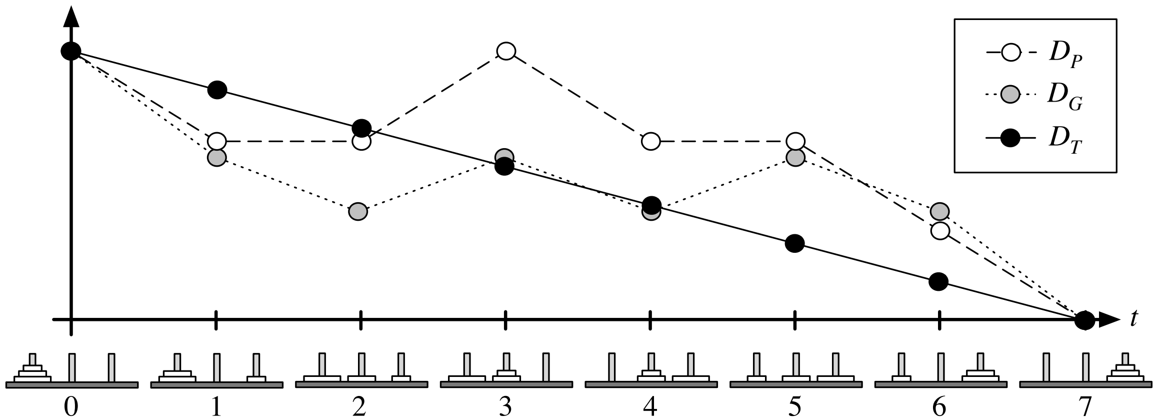

Figure 2: Illustration of non-monotonicity in genotypic distance, D_G, and phenotypic distance, D_P, but monotonicity in transformation distance, D_T, with the Towers of Hanoi problem. Distances are between the current state and the goal state, and have been normalized. The sequences of moves shown is optimal.

is no point when the path takes steps leading away from the goal state.

Having defined the various distance metrics that can apply to states, we need to ask whether we're in fact looking at the states that are relevant when making process inferences. Recall that in path search the solution is itself a sequence of states connected by operators. Thus, the states at the bottom of Figure 2 jointly constitute the solution path (Figure 1, A45). If the problem solver only represents the current state and the goal state, essentially treating each step in the solution as a new problem, then each of the individual states in Figure 2 may indeed describe the problem solver's internal state in each step, and we would conclude that the path is in fact monotonic since no state is ever revisited.

Suppose instead that the problem solver uses a technique like means-ends analysis (Newell and Simon 1961). In this case the internal search state would need to represent subgoals and paths with gaps. For example, the first subgoal in means-ends analysis would be to move the largest disk to the rightmost peg. This might be represented as:

$$(1,1,1) \rightarrow ??? \rightarrow (?,?,3) \rightarrow ??? \rightarrow (3,3,3)$$

The path would then be built recursively by filling in the steps before $(?,?,3)$ and then the steps after $(?,?,3)$. Whether this process is monotonic depends on, for example, whether the problem solver ever fails to achieve one subgoal and tries another one. Monotonicity would have to be evaluated according to the sequence of internal states (the search trajectory; Figure 1, B45), not the states of the problem itself (Figure 1, A45).

The importance of looking at internal states is clear when we realize that any search process that successfully finds a monotonic path may have explored longer paths that were ultimately edited before emitting the solution. For instance, a mathematician may prove several different lemmas as part of the proof of a larger theorem before realizing that they are all part of one general lemma and collapsing them accordingly. The final published proof would reflect this realization, but that does not imply that the mathematician's own thought processes followed the minimal path presented in the publication.

To summarize, when analyzing path searches, the relevant states to consider are the internal states, which will contain but may not be identical with the problem states. Non-monotonicity should be judged using transformation distance, as this is the most direct measure of whether the trajectory includes apparently wasted effort.

Monotonicity in Place Searches

We're now ready to shift our emphasis to place searches, where the major aim of the search is to find the most desirable final state (Figure 1, F23). Here we'll adopt a landscape metaphor, wherein states, $\vec{x}$, are thought of as being topographically organized according to the available operators. Each state's desirability is denoted $f(\vec{x})$, which we'll call it's fitness in keeping with genetics-inspired language used for genotypic and phenotypic distance.

Our discussion in this section considers search processes that maintain a single current state that is iteratively improved. Though these processes may consider several alternate states in each iteration, only one survives into the next iteration. (As with path-place searches, we leave place searches where multiple states are under simultaneous refinement for future work.) We're not going to attempt to show that any particular process fitting these parameters is most plausible. Instead, we discuss two processes that are relatively implausible and yet not straightforward to rule out with trajectory analysis. The first is a computationally implausible process that always finds the most efficient path between the initial and final state, which we call the *direct* process.[1] The second, which we call the *hill climbing* process, takes the psychologically implausible steps of evaluat-

[1]Practically speaking, finding evidence for the direct process suggests that the entire search occurred mentally, making trajectory analysis the wrong analytical approach for that problem.

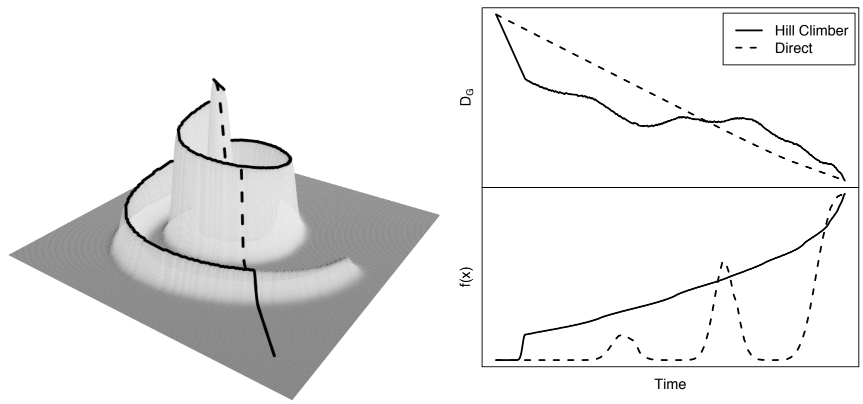

Figure 3: Illustration of (a) the hill-climbing process, which shows non-monotonicity in genotypic distance but monotonicity in fitness (solid lines), and (b) the direct process, which shows monotonicity in genotypic distance but non-monotonicity in fitness (dashed lines). Times and distances have been normalized.

ing *every* available move and *always* choosing the one that most improves fitness, stopping when no improvement is possible (regardless of how low the current state's fitness).

Whereas monotonicity in path search always referred to the states in the search trajectory, place search lets us look at the monotonicity of both the intermediate states and their fitnesses (Figure 1, F3). We'll call the monotonicity of the intermediate states *state monotonicity* (Figure 1, DE45). The same distinctions between genotypic, phenotypic, and transformation distance apply, and as before transformation distance is the most relevant form of monotonicity (though usually not the most convenient to calculate). We'll call the monotonicity of the fitness function over time *fitness monotonicity* (Figure 1, G45). As we'll see, different conclusions can result from considering fitness monotonicity as assessed during search (Figure 1, FG78) or after search (Figure 1, H78).

In the following we'll present visual examples of searches over two-dimensional grids, with the states in the x-z plane and fitness on the y-axis. In this way, the ideal endpoint for a place search is the highest point on the landscape. We'll allow single-unit moves in the up-down, left-right, and diagonal directions. Note that in this case D_G, defined as the Euclidean distance, is a good proxy for D_T.

Insufficiency of Either Monotonicity Alone

Now we can evaluate whether fitness or state trajectories are individually sufficient to rule out either the direct or hill climbing processes. For the landscape shown in Figure 3, a hill-climber that follows the fitness function, f, will trace a path like the solid line shown in the figure. This trajectory is non-monotonic in genotypic distance but monotonic in fitness. The direct process would form a trajectory that is monotonic in genotypic distance but non-monotonic in fitness. Therefore, finding one but not both of state or fitness non-monotonicity is not sufficient to rule out both null hypotheses processes.

Partially Observable States

The trajectory traced in Figure 3 is a spiral. While this is literally a roundabout path to take, it is at least free of cycles where the trajectory revisits previously-encountered states. A trajectory with cycles seems to clearly contradict both null hypotheses. Indeed, the presence of a cycle can rule out the direct process, as editing out the cycle will not prevent the path from reaching the same final state. A cycle is likewise impossible with hill climbing, *unless* the internal state is only partially observable (cf. Gabora 2011).[2]

Consider the landscape in Figure 4, where the complete state $\vec{x} = (x, \hat{x})$ has an observable part (x) and an unobservable part ($\hat{x}$). Projecting the full trajectory into x shows a cycle, but the cycle disappears with $(x, \hat{x})$.

Criteria Change

In a case almost identical to the previous example, suppose that people's criteria change during search. This is akin to having an unobserved portion of the state that affects the evaluation function. Here again, cycles could emerge in the observable state that are in fact monotonic if the evaluation function's hidden parameter were observable. (The cases are not quite identical since changes to this hidden parameter are dependent on some higher-level process, and hence aren't fully explainable by reference to f.) If one considers changes to evaluation criteria as separate from the main search process then it's still possible for this apparent non-monotonicity to occur with a hill climber.

A special problem can occur when criteria change occurs since the same state will be evaluated differently before and after the change. This could lead to retrospective evaluations of fitness over time showing non-monotonicity while in fact fitness was experienced as increasing monotonically during the search (see Figure 1, FH78). Suppose that the

[2]Cycles could also occur in the special case where all of the states in the cycle are indistinguishable in terms of fitness.

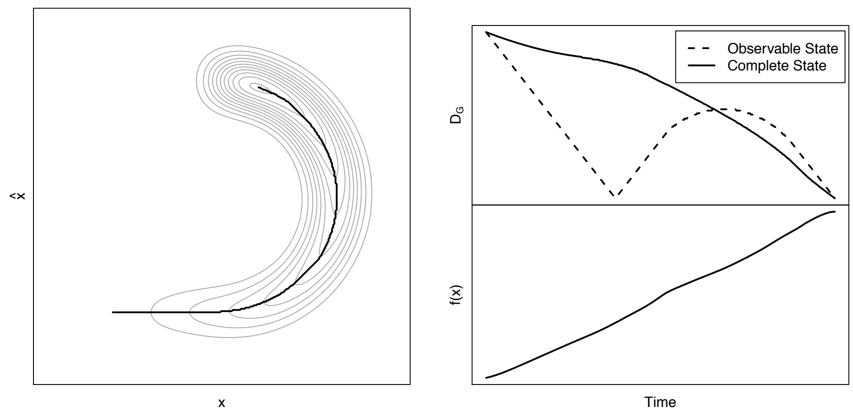

Figure 4: Illustration of an apparent cycle using a hill climbing process where the state is only partially observable. Projecting the trajectory into observable space (horizontal axis) shows a cycle, while the complete trajectory is non-cyclic.

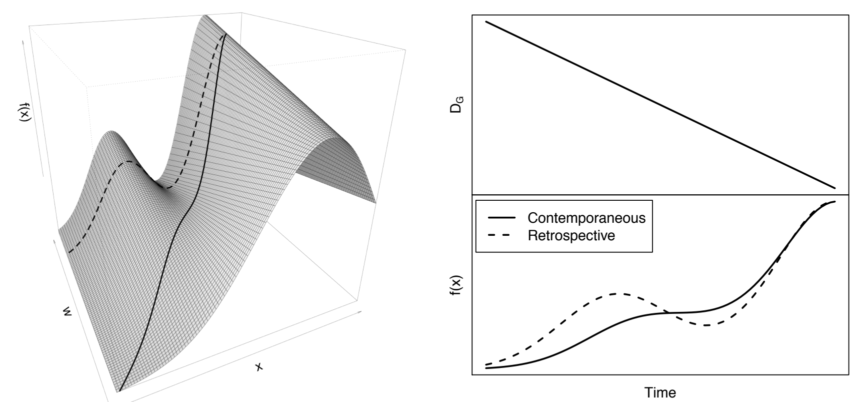

Figure 5: Illustration of a search over x while a criteria weight w is simultaneously changing. The path (solid line) starts in the foreground and proceeds to the background. As shown in the right graph, contemporaneous fitness is monotonic. However, retrospective evaluation of the traversed points would be non-monotonic, as shown by the dashed lines in both graphs.

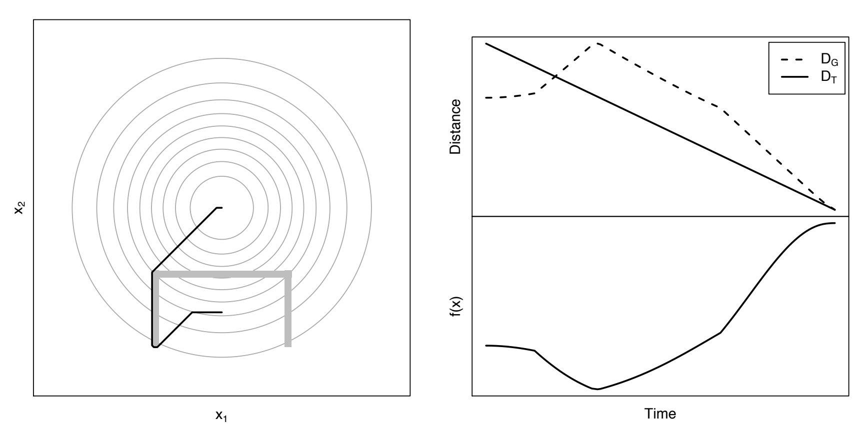

Figure 6: Illustration of a path that is non-monotonic with D_G but monotonic with D_T. The gray region in the left plot shows infeasible configurations that can be represented with the genotype but that cannot be realized.

state is a scalar x and that $f(x)$ depends on a weight parameter w such that $f(x) = w \cdot f_1(x) + (1 - w) \cdot f_2(x)$. Figure 5 demonstrates how retrospective fitness evaluation could show non-monotonicity even though contemporaneous fitness was monotonic.

Infeasible Regions

So far we have considered problems where D_G is a good proxy for D_T. However, there are problems where the state representation does a poor job reflecting the actual difficulty of transforming one state into another. Consider the landscape in Figure 6, where the gray regions are infeasible states. Though a hill climber would fail on this particular landscape, the direct process would find the path shown. As can be seen, the state trajectory appears non-monotonic with D_G but isn't when judged with D_T. The fitness trajectory is also non-monotonic, though that could occur with any direct process.

Phenotypic Distance

In real-world situations, both transformation distance and genotypic distance can be impractical to compute. In contrast, phenotypic distance, which is based on comparing the state's salient features, can be assessed fairly directly, such as by having several human raters make intuitive similarity judgments for pairs of intermediate products. However, there are two issues with this approach. From a psychological standpoint, it has long been known that human similarity judgments are not well-behaved metrics, meaning that they can be context-dependent, asymmetric, and intransitive (Tversky 1977). These properties could introduce non-monotonicity that was is not in the original stimulus (Gabora 2011).

Another problem with similarity judgments lies in the fact that genotype and phenotype may be non-monotonically related. Let the state be a scalar $x \in [0, 1]$ and suppose that this maps to a single salient property, $p(x) = \sin(2\pi x)$ (see top half of Figure 7). If a search proceeds linearly from $x = 0$ to $x = 1$, D_G (the absolute difference in state representations) will be monotonically decreasing but D_P (the absolute difference in the level of the property) will be non-monotonic (see bottom half of Figure 7). This is true regardless of what sort of landscape or search process is used.

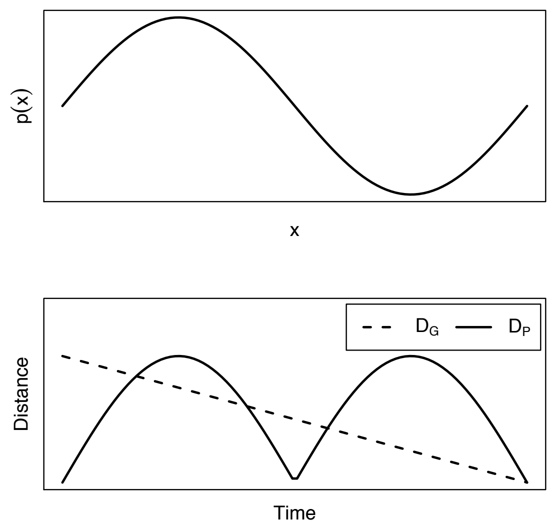

Figure 7: The top graph shows the relationship between the state x and its salient property, $p(x)$. The left graph shows monotonicity in D_G (and presumably D_T) as x goes from zero to one, but non-monotonicity in D_P.

Discussion

This paper has introduced a set of important distinctions that must be made when analyzing creative search trajectories, as summarized in Figure 1. We have argued that path searches and place searches must be understood differently. With path searches, we have claimed that the important measure of the efficiency of the search process is the transformation monotonicity of the problem solver's internal states, which may or may not be the same as the states of the problem itself. With place searches, we have argued that researchers claiming that people use complex search processes must first reject two null hypothesis processes, the direct and hill-climbing processes. We have shown how rejecting these processes entails demonstrating not just that there is non-monotonicity in both states and fitness, but also that:

- State non-monotonicity occurs with transformation distance, not just with the more convenient indices of genotypic and phenotypic distance
- The observed state non-monotonicity does not reflect movement on unobservable dimensions or the effects of criteria changes
- Fitness non-monotonicity assessed retrospectively would also have been non-monotonic when assessed contemporaneously

We have also illustrated how non-monotonicity may result from processes that are at either end of the spectrum from intelligent and informed (the direct process) and mechanical and uninformed (the hill climbing process). Regarding existing empirical work, we have shown that authors have differed on whether they've analyzed state (Simonton 2007; Damian and Simonton 2011) or fitness (Kozbelt 2006; Kozbelt and Serafin 2009) monotonicity. Though we don't claim that the counterexamples provided here disprove the claims made by these authors, we have raised important issues that future work must address.

We also acknowledge that our work has several significant limitations. First, for simplicity we have falsely dichotomized path and place search, even though we are quite sure that real creativity involves elements of each. Indeed, the joint operation of these searches may map nicely onto the problem finding/problem solving distinction that Getzels and Csikszentmihalyi (1976) introduced. Second, though we have alluded to criteria change as an important consideration when analyzing place searches, we have avoided discussing how and when this would occur and what could drive it. This decision was ultimately practical, since any discussion of changes to criteria leads to the question of what guides criteria selection, and whether this itself can change—leading to an infinite regresss that may not be resolvable at the level of an individual creator (cf. Jennings 2010a). We have also treated criteria change as compatible with the hill climbing process, which we recognize may not be without controversy. Third, we recognize that we have not carefully considered the role of operators, and in particular how the discovery of new operators mid-search may affect conclusions about state monotonicity. Finally, we are fully aware that all of our counterexamples involve highly abstracted problems, and so ultimately they serve more to illustrate our points rather than provide an existence proof. We look forward to addressing these and other limitations in future work.

Our purpose here is not to cast aspersions on the trajectory analysis approach. Indeed, we are actively pursuing empirical techniques that rely upon trajectory analysis (Jennings 2010b; Jennings, Simonton, and Palmer 2011), although these techniques promise to reveal more about the underlying problem than can be obtained from a search trajectory alone. Beyond this, we continue to believe that trajectory monotonicity is an eminently practical way to study the creative process, particularly in cases where the available data are slim. The objections that we've raised in this paper do not seal the fate of this approach, but rather offer constructive critiques that should be addressed in future research.

References

Anderson, J. R. 1993. Problem solving and learning. *American Psychologist* 48(1):35–44.

Damian, R. I., and Simonton, D. K. 2011. From past to future art: The creative impact of Picasso's 1935 *Minotauromachy* on his 1937 *Guernica*. *Psychology of Aesthetics, Creativity, and the Arts*. Advance online publication. doi: 10.1037/a0023017.

Dasgupta, S. 2011. Contesting (Simonton's) blind variation, selective retention theory of creativity. *Creativity Research Journal* 23(2):166–182.

Finke, R. A.; Ward, T. B.; and Smith, S. M. 1992. *Creative Cognition: Theory, Research, and Applications*. Cambridge, MA: MIT Press.

Gabora, L. 2011. An analysis of the blind variation and selective retention (BVSR) theory of creativity. *Creativity Research Journal* 23(2):155–165.

Getzels, J. W., and Csikszentmihalyi, M. 1976. *The Creative Vision: A Longitudinal Study of Problem Finding in Art*. New York: John Wiley & Sons.

Hennessey, B. A. 1994. The consensual assessment technique: An examination of the relationship between ratings of product and process creativity. *Creativity Research Journal* 7(2):193–208.

Jennings, K. E.; Simonton, D. K.; and Palmer, S. E. 2011. Understanding exploratory creativity in a visual domain. In *Proceedings of the Eighth ACM Conference on Creativity and Cognition*.

Jennings, K. E. 2010a. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines* 20:489–501.

Jennings, K. E. 2010b. Search strategies and the creative process. In *First International Conference on Computational Creativity*, 130–139.

Kozbelt, A., and Serafin, J. 2009. Dynamic evaluation of high- and low-creativity drawings by artist and nonartist raters. *Creativity Research Journal* 21(4):349–360.

Kozbelt, A. 2006. Dynamic evaluation of Matisse's 1935 *Large Reclining Nude*. *Empirical Studies of the Arts* 24(2).

Newell, A., and Simon, H. A. 1961. GPS, a program that simulates human thought. In Billing, H., ed., *Lernende automaten*. Munchen: R. Oldengourg. 109–124.

Newell, A., and Simon, H. A. 1972. *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Perkins, D. 2000. *The Eureka Effect: The Art and Logic of Breakthrough Thinking*. New York: W. W. Norton.

Rostan, S. M. 2010. Studio learning: Motivation, competence, and the development of young art students' talent and creativity. *Creativity Research Journal* 22(3):261–271.

Ruscio, J.; Whitney, D. M.; and Amabile, T. M. 1998. Looking inside the fishbowl of creativity: Verbal and behavioral predictors of creative performance. *Creativity Research Journal* 11(3):243–263.

Simonton, D. K. 1999. *Origins of genius*. New York: Oxford.

Simonton, D. K. 2003. Scientific creativity as constrained stochastic behavior: The integration of the product, person, and process perspectives. *Psychological Bulletin* 129(4):475–494.

Simonton, D. K. 2007. The creative imagination in Picasso's *Guernica*: Monotonic improvements or nonmonotonic variants? *Creativity Research Journal* 19:329–344.

Simonton, D. K. 2010. Creativity as blind-variation and selective-retention: Constrained combinatorial models of exceptional creativity. *Physics of Life Reviews* 7:156–179.

Simonton, D. K. 2011. Creativity and discovery as blind variation and selective retention: Multiple-variant definition and blind-sighted integration. *Psychology of Aesthetics, Creativity, and the Arts* 5(3):222–228.

Simonton, D. K. 2012. Creativity, problem solving, and solution set sightedness: Radically reformulating BVSR. *Journal of Creative Behavior* 46(1):48–65.

Tversky, A. 1977. Features of similarity. *Psychological Review* 84(4):327–352.

Weisberg, R. W. 2004. On structure in the creative process: A quantitative case-study of the creation of Picasso's *Guernica*. *Empirical Studies of the Arts* 22(1):23–54.