Published in: Leonardo Volume 51, Issue 2, p. 165-172, April 2018, MIT Press

Art in the Sciences of the Artificial

Kenneth O. Stanley (kstanley@cs.ucf.edu)
Department of Computer Science
University of Central Florida, Orlando, FL 32816

Abstract

This position paper argues that the sciences of the artificial (artificial intelligence and artificial life) have a special relationship to art that is absent from much of science. Just as art is often a depiction or interpretation of nature, so are the algorithms in the sciences of the artificial. This observation is important because the discourse in these fields largely ignores the relevance of subjective resonance with nature to their own scientific progress. Yet progress is potentially stifled if we cannot discuss such resonance openly. To support this view, examples in this paper illustrate how the subjective impression of such resonance led to novel encodings and algorithms in the author's own career. An important conclusion is that there may be more to gain than to lose by allowing some level of subjectivity to enter the discourse in the sciences of the artificial.

1 Introduction

While the argument that artistic sensibility is important to science is not new [3, 16, 20], this position paper argues that the relationship of the broad fields of artificial intelligence (AI) and artificial life (alife) to art is unique within science and deserving of more open discussion. Both AI and alife share the word "artificial," and anything that is artificial relates at some level to a natural counterpart. For AI and alife, the natural counterparts are unambiguous: Both *intelligence* and *life* are highly salient natural phenomena. In this way, AI and alife diverge from areas of science that are not predicated on creating artificial versions of phenomena seen in nature. The consequent connection to art is that much of art is *also* concerned with reproducing phenomena seen or experienced in nature. Of course, it is important to acknowledge that not all art is concerned with the reproduction of nature, especially in the modern era [2], and that art often speaks beyond only its representational aspects to the cultural and historical context of the artist. However, depictions and interpretations of nature and civilization were key historical drivers of the development of art as a cultural phenomenon and remain central to much of art today [4, 15, 18].

Yet while art often depicts or interprets the outside world, that does not mean in such cases that it is a perfect reproduction of its subject matter. For example, a painting may aim to capture the essence of an event, a mood, or a person [15] without reproducing every detail. Furthermore, of course there is no expectation that the painting *is* that event or that person; rather it is an *artificial* reproduction of it. As Gombrich [18, page 49] puts it, "What a painter inquires into is not the nature of the physical world but the nature of our reactions to it." Similarly, music may capture a feeling or literature a situation. None of these modes of expression require that the reproduction is perfect, complete, or "photographic." In fact, one of the powers of art is to emphasize (or even exaggerate) a particular quality that may come at the expense of total accuracy (e.g. in caricature [26] or impressionist art). On the other hand the artist may select particular resonant features to approach an idealization of reality [4].

Another way to articulate this idea of acceptable distortion is to say that when art serves as a window to our experience it is not inherently about accuracy but rather about *resonance*. That is, art is often most moving when it resonates deeply with some aspect of our experience of the real world, such as our experience of nature [15, 18]. (Merriam-Webster defines this sense of *resonate* as *to have particular meaning or importance for someone: to affect or appeal to someone in a personal or emotional way* and *to relate harmoniously: strike a chord* [25].) To resonate in this way, art must capture a fundamental essence of its subject filtered through the artist's intent and present it to us such that we feel its resonance with its natural counterpart. Leonardo da Vinci [10] describes the challenge of achieving such resonance: "Painting ... compels the mind of the painter to transform itself into the very mind of nature, to become an interpreter between nature and art" (quoted in Bialostocki [4]).

For example, van Gogh's famous painting *Starry Night* [43] showcases the value of resonance without the notion of accurate reproduction. It is unlikely that any night on Earth has looked *literally* like the starry night depicted by Van Gogh, yet at the same time, by exaggerating some details and stylizing others, the painting resonates with a deeper aspect of the splendor and mystery inherent in such a night. In effect, if the viewer contemplates the scene in the painting, she may come to appreciate some deeper truths about *real* starry nights that she would never have realized when experiencing the real thing. In some cases, a recognition of the artist's intent can further enrich such an experience. In this way, the art becomes a kind of lens that magnifies some qualities at the expense of others, ultimately enriching our awareness of reality by subtly distorting it [1, 15].

Interestingly, the relationship between a painting and its natural counterpart is not unlike the relationship between any artificial reproduction of nature and its inspirational counterpart. In this way, an *algorithm* in the spirit of AI or alife is in effect a kind of "painting" of nature, only viewed through the lens of *process* rather than through image. Yet the reproduction of nature through a description of process is in principle no less interpretive or revelatory than through any other artistic medium. The main difference is only that process is often described through a code that only those experienced in the art of programming or computer science can easily appreciate.

The idea that an algorithm's value can be traced to the resonance of its process with a counterpart process in nature opens up an important discussion on our conventions as a scientific community for evaluating the impact of algorithms in the "sciences of the artificial" (to borrow the term from Simon [36]). As with any other art, a feeling of resonance with nature is possible to achieve without conventional measurable accuracy, and in fact exaggeration or emphasis on some aspects of natural processes over others may even be intended [18], thereby opening the mind of the "viewer" to a hidden aspect of nature just like in *Starry Night*.

This perspective raises a genuine practical question for the fields of AI and alife because their artifacts are so commonly evaluated from an antithetical objective perspective. This sharply objective inclination is exemplified by a preoccupation with *performance* in AI and machine learning [31], and *accuracy* in alife worlds [8, 27]. The problem is that these conventional concerns stifle the discussion and promotion of the more subjective notion of resonance with nature. Perhaps an artificial intelligence can be a kind of *Starry Night*, prompting us to view the mind anew. In this way, a new AI algorithm could serve as a different perspective on intelligence rather than a better performer. Few would demand a "comparison" between *Starry Night* and a real starry night before admitting it into the world of bona fide art.

It is important to note of course that resonance is not entirely excluded from consideration in the sciences of the artificial. A number of such exceptions from AI and alife are listed in online Appendix A. Scientific communities are broad and diverse, so it makes sense that some smaller subcommunities will sometimes defy the trends of the whole. Thus the argument here is not against the straw man that *all* decisions of import in AI and alife are entirely objective and performance-driven. Rather, the problem is that while such subjective considerations do occasionally influence thinking, they are often relegated either to entirely private discourse where resonance cannot be contemplated in the literature of record, or they seep through only in highly specialized venues such as EvoMUSART calibrated for such discussion, yet outside the mainstream (e.g. AAAI, NIPS, or ICML).

While the benefits of debating subjective resonance may still seem abstract, the fundamental problem with ignoring this artistic side of exploring the artificial is that the chain of innovation that leads from one idea to another in these fields may ultimately require traversing *some* ideas that stand out not for their accuracy or performance, but for their resonance. After all, even in mainstream AI ideas that present to us a familiar landscape in an entirely new light are exactly the kind of stimulus we will need in the long run to be able to provoke imaginative new approaches. While it goes without saying that of course not *all* progress in these fields requires artistic inclination, to avoid handicapping our capacity for progress, we may need to acknowledge more broadly the special nature of artificiality within some sciences, which connects more strongly to art than in other scientific disciplines.

This paper is an attempt to argue through evidence that the subjective interpretation and experience of artifacts from alife and AI should be admissible in the professional discourse of these disciplines (e.g. in publications and reviews), not just in the service of generating aesthetic results, but in achieving the ultimate ambitious aims of the fields. Because a thorough examination of art within science necessitates at least some level of subjectivity, I have consciously chosen to present much of this position paper from my own first-person perspective. Yet despite taking this unusual first-person perspective for an article in a scientific journal, it is important to emphasize that the intent is *not* primarily to be autobiographical. Rather, my hope is that anecdotes from my own career can effectively support the general position that art is essential to AI and alife. My broader hope (perhaps overly ambitious) is that through initiating a dialogue on this important subject, it will become easier to discuss subjective and aesthetic aspects of our algorithms, both in published work and in the reviews of article submissions, without the present stigma of being called *unscientific*.

2 Past Reflections at the Intersection of Science and Art

Because science and art are both creative endeavors, it is not surprising that their relationship would intrigue many thinkers. As Albert Einstein said, “The greatest scientists are artists as well” [6]. This perspective that the practice of science is at least in part a kind of art pervades much of the commentary on the subject, often focusing on the role of intuition in scientific discovery.

William Beveridge’s classic book, *The Art of Scientific Investigation* [3], exemplifies this approach to connecting science and art, including an entire chapter on the role of intuition in science. While conventional scientific tools such as objective observations and hypotheses are discussed as well, one important message is that a sense for aesthetics and elegance also plays a role in effective scientific intuition. Online Appendix B reviews several other commentaries that strike similar themes. Of course, while an artistic inclination can benefit scientific discovery, science can also serve art, both as a method of explaining its mechanics [19, 28, 32, 35] and as a vehicle even for artificially generating it [30, 37]. In the latter case, when machines generate art, many interesting questions arise concerning the role of artistic intent and authorship [5].

This paper diverges both from conventional arguments for the role of artistic inclination within science and from explorations of science to explain or support art. Instead, the suggestion in this paper is that branches of science concerned with the artificial are *themselves* intrinsically artistic (and not only in a metaphorical sense), an idea that is significantly less discussed or debated. Because the sciences of the artificial are especially concerned with reproducing phenomena observed in nature artificially, this relationship is unique even compared to the relationship of art to other branches of science. In this way, this paper aligns more strongly with the literature on art-within-science and is thereby also not a conventional exposition of the virtues of evolutionary or artificially-generated art. The next section begins to explain the unique relationship between art and the sciences of the artificial by examining art as a tool for research in these areas.

3 Art as a Tool for Science

This section recalls the origins of two ideas from my own career. In both cases, the initial spark that led to the idea is a subjective observation about pictures that resonated for me with something deep in nature. While of course not all scientific discoveries in AI or alife will originate with subjective observations, the question is whether discouraging such observations in public discourse might cut off some future discoverers from what could have been their inspirations.

3.1 Compositional Pattern-Producing Networks

Art is not new to artificial intelligence or especially evolutionary computation (online Appendix C reviews in particular work in the field of *evolutionary art*). In 2004 I was exploring two new evolutionary art programs called DelphiNEAT-based Genetic Art (DNGA [12]) and SharpNEAT-based Genetic Art (SNGA [13]). Both of these programs interested me in part because they are based on the NeuroEvolution of Augmenting Topologies (NEAT; [40, 42]) algorithm, which I had recently introduced with Risto Miikkulainen. Because NEAT is a method for evolving increasingly complex artificial neural networks (ANNs) that was originally designed for control problems, it had not occurred to me that the ANNs evolved by NEAT might be asked to output pictures. Particularly intriguing to me was that because NEAT gradually adds complexity to the networks it evolves, the images in an evolutionary art system based on NEAT could in principle become more complex over time. However, evolutionary art was familiar to me at the time primarily as a toy.

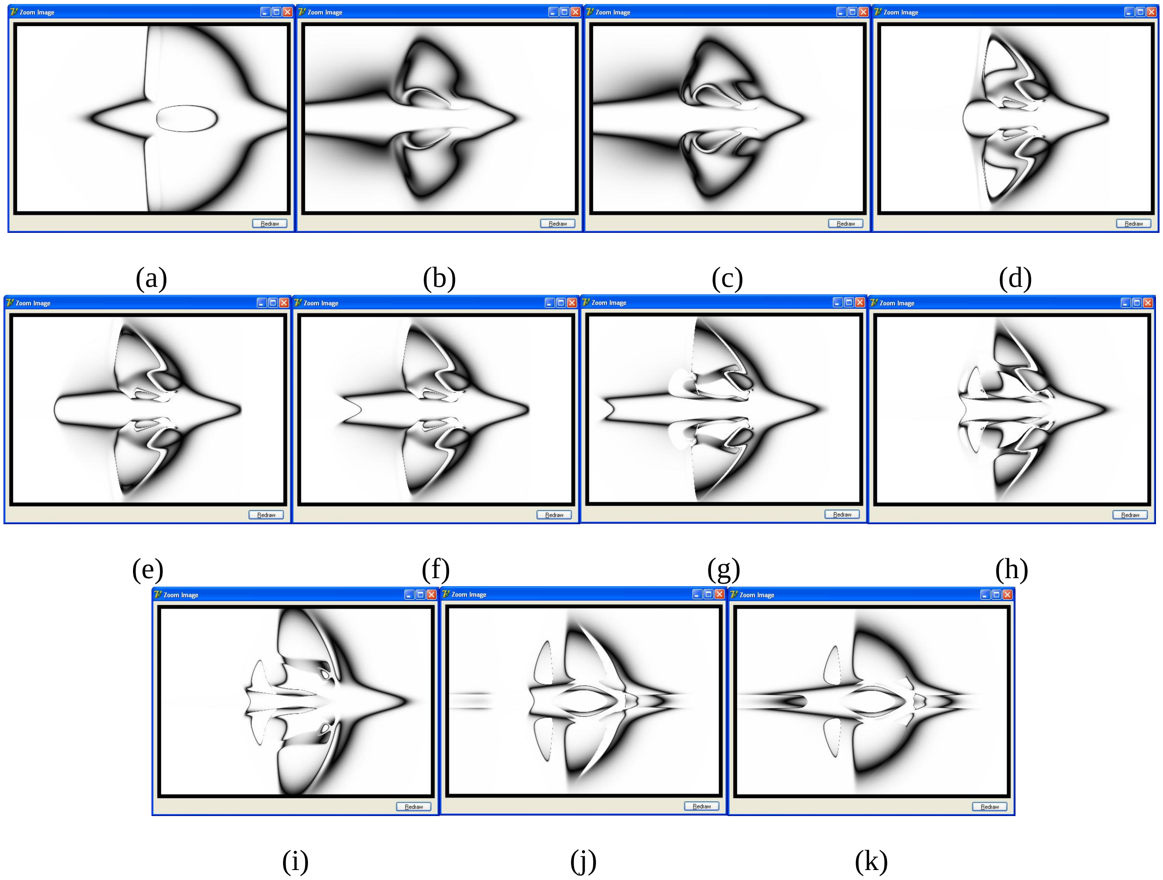

Figure 1: **Sequence of Interactively Evolved Spaceships that Inspired CPPNs.** The spaceships in this sequence are the original inspiration for CPPNs.

That changed one day in 2004 when I interactively evolved a sequence of spaceships with the DNGA tool (a selection of this sequence is shown in figure 1). This single experience of evolving a sequence of spaceships struck a deep chord because the sequence *resonated* with something subtle in nature. However, instead of reminding me of spaceships (which to me is their superficial likeness), these images reminded me of the way evolution progresses in nature, which is a feature of *process*. In effect they appeared a kind of allegory for natural descent and elaboration. For example, the sequence begins with a discovery of bilateral symmetry and elaborates from there, adding increasingly intricate elaborations that then become fodder for further elaboration. I was particularly astonished when the tail fins appeared (figure 1h). After all, DNGA knows nothing of rockets, planes, or tail fins. While the spaceships themselves may not be great art with respect to spaceship depictions, in the resonance of the whole sequence with an abstract conception of nature there *is* an artistic connection. In this sense, the sequence to me is like a stylized interpretation of the way natural evolution unfolds.

Over time, as I studied the images and their underlying representations in NEAT I began to understand that a key property enabling this kind of elaborated regularity is *function composition*, and I started to relate function composition to more traditional abstractions of biological development (which often appear in alife research) like cell chemistry and grammatical encodings [41] (online Appendix D gives several examples of such encodings). It turns out that function compositions of increasing complexity make an analogy with the way genetic representation in nature elaborates form over generations [7]. This realization led to the new encoding called *compositional pattern-producing networks* (CPPNs), which are in effect networks of functions evolved by NEAT [38]. The idea is that these functions can be conceived as representing steps in pattern-formation during a developmental process, thereby introducing a new developmental abstraction.

CPPNs went on to become the basis of numerous applications and experiments in AI and alife (several dozen examples ranging from generating music to robot morphologies are documented in online Appendix E). They also became the basis of the new kinds of evolutionary ANNs in HyperNEAT [17, 39], which itself became the basis of a new research area with practical applications in autonomous control and decision making (online Appendix E). These active areas of research would not currently exist if not for that first impression of a few images of spaceships on my computer screen.

One other application of CPPNs that deserves special mention is Picbreeder (http://picbreeder.org) [33, 34]. The potential for evolving CPPNs to lead to such interesting progressions as in figure 1 suggested to my research group that more such discoveries might emerge if a large-scale, crowd-sourced effort to evolve CPPN-generated images was made available on a public website[1], which became Picbreeder (figure 2). However, while the hope for better evolved art is a plausible motivation, after the experience of the spaceship, that was *not* my personal motivation. Rather, I had begun to believe that the products of evolutionary art are so rich with scientific implications that *something* else (though I did not know what) would likely emerge from such a large-scale artistic collaboration on the Internet. The resultant discovery is described next.

3.2 Novelty Search

Unlike most evolutionary algorithms that search for a particular objective or set of objectives, novelty search [22, 23], first introduced by Joel Lehman and myself in 2008, searches only for novelty. This kind of search is interesting in part because by ignoring objectives it can avoid the deception inherent in many objective functions [22, 23, 29]. That is, sometimes it can evolve a solution to a problem more reliably than an evolutionary algorithm that is actually rewarded for approaching the objective. However, novelty search is

[1] This idea was first proposed in a lab meeting by Jimmy Secretan.

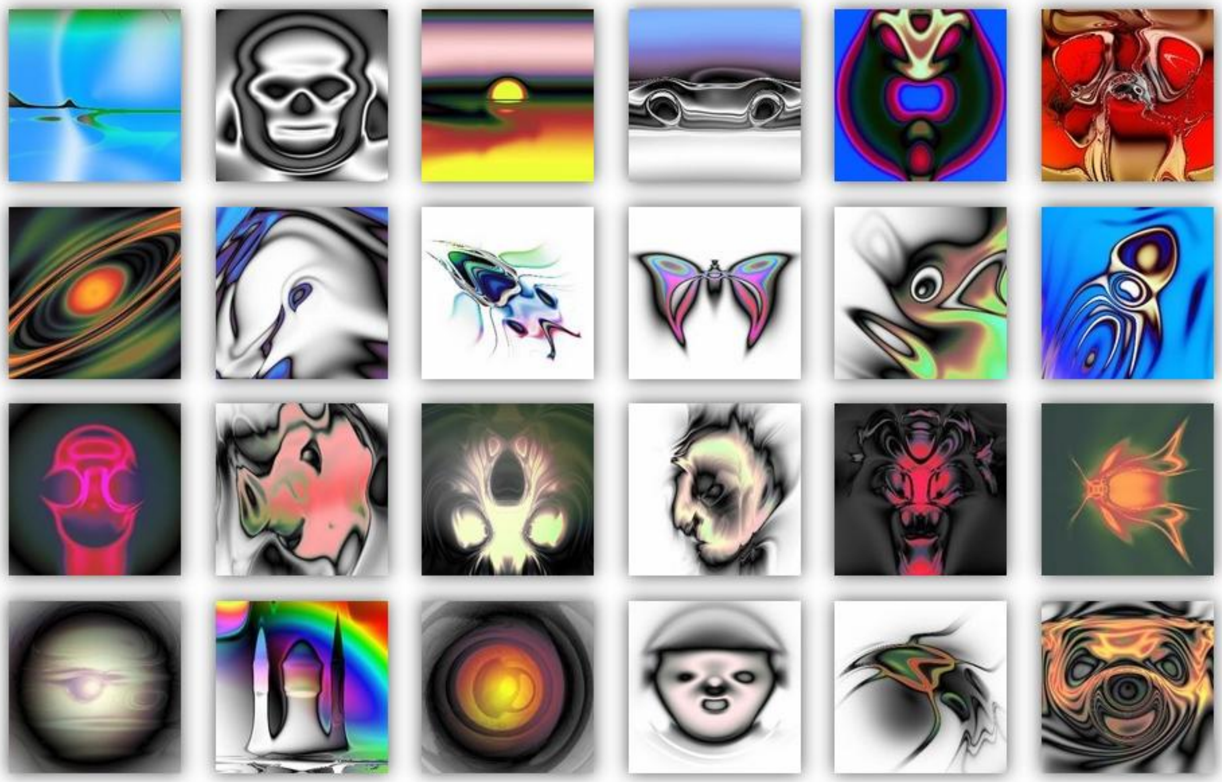

Figure 2: **CPPN-Encoded Images Interactively Evolved by Picbreeder Users.** A sample of the over-9,000 collaboratively evolved images is shown.

interesting more deeply because it can potentially *collect* interesting items in the search space, whether they are preconceived or not [9, 24]. In this way, unlike most evolutionary and optimization algorithms that are *convergent*, novelty search is a *divergent* search and thereby lends itself to new kinds of applications. Like CPPNs, novelty search has provoked a significant body of work extending and analyzing it (several dozen examples, from practical applications like robot ambulation to tools for discovery and invention, are given in online Appendix F).

While much has been written on novelty search and its variants, its original inspiration has not previously been discussed in publication. As images like those in figure 2 accumulated on Picbreeder, it became apparent that humans could discover interesting images fairly consistently. Of course, the *interestingness* of images is subjective, but to learn something from the results of a system that produces images requires at least initially some openness to subjective judgment. Serendipity ultimately revealed to me how such images are discovered. My intent initially was not to make a scientific discovery, but rather just to play with Picbreeder to see what I could evolve. An important option on Picbreeder is the ability to *branch* an interactive session from an image previously evolved by another user. The site provides a searchable

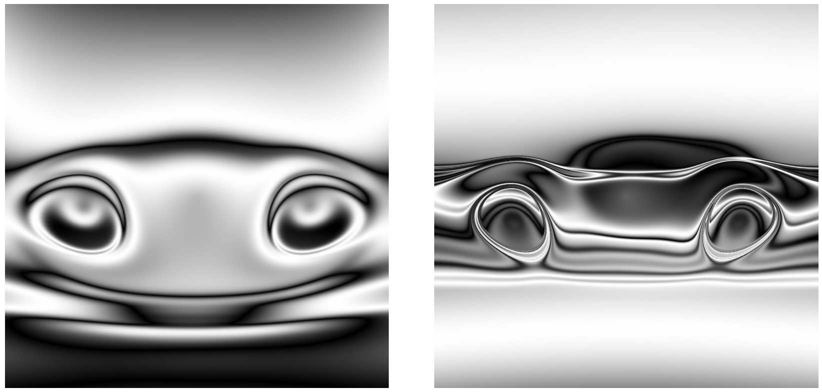

(a) Alien Face (b) Car

Figure 3: **Images that Inspired Novelty Search.** The image of the alien face (a) previously published on Picbreeder turned out to be an unexpected stepping stone to evolving the car (b).

catalog of images previously evolved by users to facilitate such branching. In this particular case, I decided to branch from a previously-evolved image that looked to me like an alien face (figure 3a). I thought that I might evolve more interesting aliens from this original alien.

However, that is not what happened. Instead, the eyes of the alien began to descend relative to the head, and I realized that it was beginning to look like a car. From there I was able to push it further in the direction of a car until I had evolved to my surprise a plausible image of a car (figure 3b). This experience was strangely reminiscent of the spaceship. At first I was amazed that I had been able so easily to evolve a car. However, the more I thought about it, the more something seemed wrong: I had not been *trying* to evolve a car. How was it that I evolved something I did not set out to evolve?

Even more strange, I realized that not only had I not set out to look for it, but that the only reason I found it was *because* I was not looking for it. After all, I began with an alien face because I wanted to find more alien faces. Therefore, if I had actually been trying to evolve a car, I would not have started with an alien face, and as a result I would have pruned away a promising path to the car. In short, it seemed that the only way for the discovery to happen was by *not* trying to look for it. This realization contradicted a lot of what I had previously believed about search, which is that the way to find something is to set it as an objective and then optimize towards it. I began to wonder, what if that is not how interesting discoveries are really made?

Of course, it is possible that the story of the car is somehow unique and that other discoveries are not made in that way, but when I looked back at the lineages of other images on Picbreeder (such as those in figure 4), the conclusion was unavoidable – almost *every* interesting discovery in Picbreeder works that way. Users almost always first branch from something so radically different from what they ultimately discover

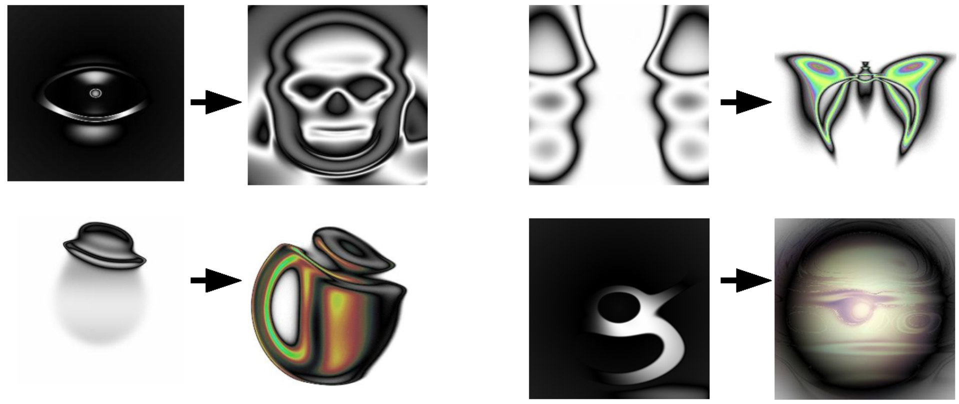

Figure 4: **Deceptive Stepping Stones.** In each pair of evolved images, the image on the left was initially evolved by a different user from the one (or more) who ultimately *branched* from it to create to image on the right.

that there is no plausible argument that they had it in mind from the start. Instead, surprisingly, users on Picbreeder discover interesting images only when they are not looking for them.

In this way, the observation that users only find what they are not seeking, provoked initially by evolving the car, was the initial impetus behind novelty search, an algorithm that searches without an objective. The idea is that maybe a good way to find something interesting is to explore possible stepping stones (like the alien face) without trying to control where they might lead. I discussed this observation with my then-student Joel Lehman, and together we developed it into the novelty search algorithm.

The common thread from CPPNs to novelty search is that *pictures* and their subjective resonance with nature triggered the key insight, in part through their own subjective resemblance to natural artifacts, but also in part through their chronological depiction of *processes* reminiscent of those in nature. Importantly, that does not mean CPPNs are *equivalent* to natural DNA or that novelty search is the *true* process of search in nature. Rather, they are resonant with such natural counterparts without being equivalent. In a sense they can be regarded as *interpretations* of nature that focus on a particular important facet, just as art often interprets the world without reproducing it precisely.

Thus it is important that we are open to the subjective nature of art-like or aesthetic output from systems and models in the sciences of the artificial. After all, the initial inspiration for both AI and alife is nature, and the root of this inspiration is not in any quantifiable fact but rather in a *feeling* about the profound significance of nature's achievements. Should we be afraid to acknowledge that as scientists?

4 Algorithms as Artistic Artifacts

The idea that subjective impressions should sometimes serve as admissible inspiration in AI and alife suggests more generally that the algorithms and encodings we create themselves can be regarded as art, and perhaps sometimes should be judged in that spirit. Just like any other art, algorithms in the sciences of the artificial are interpretations of phenomena observed in the real world, and thus stand alongside paintings, literature, sculpture, and other art forms as *statements* on the artist's view of nature. Yet outside the field of generative art (which focuses explicitly on algorithms in the context of art production [5, 11]), algorithms in alife and AI are rarely judged publicly in this way. If they were, then sometimes the statement made by the artist through the artifact would be more important than its results because what would matter is how it opens our mind to a new perspective rather than how it performs or how accurate it is.

Nevertheless, a likely criticism is that such conversation belongs outside the *sciences* of the artificial because it does not in itself advance the science. However, if indeed our fields are about ultimately replicating some of the most impressive aspects of natural phenomena, then surely the resonance of our artifacts with such phenomena is relevant not only to philosophical discourse but to concrete progress as well. That is, if we view these fields as largely creative endeavors whose practitioners are searching for ever more powerful or accurate algorithms, then each published algorithm becomes a potential stepping stone to new ideas and yet more innovative algorithms. There is no a priori reason to believe that the path through the space of algorithms is lit exclusively by performance and accuracy instead of sometimes by resonance and elegance. The result is that *if* algorithms whose primary highlight is resonance with nature might be stepping stones eventually to algorithms with greater performance (such as CPPNs leading to HyperNEAT), then we are effectively *pruning* many such chains of innovation out of the search space of the sciences of the artificial.

Whether intended or not, many algorithms in the sciences of the artificial are thereby inevitably statements about nature as well. They reveal a personal perspective on behalf of their authors about what resonates in nature. In this way, they might even be regarded as mediums of self-expression, as with any other art. In short, as a field we might sometimes benefit from following the gradient of resonance over the gradient of performance in the search for new ideas. Interestingly, even beyond the benefit to the science (which is the focus of this article), the novel mode of understanding opened by an algorithmic interpretation of nature may in its own right offer value, even outside its utility as a tool for science.

Perhaps the easiest way to show what I mean by algorithms as artistic artifacts is to paint some of my own work in this unusual light: NEAT is a meditation on the elegance of increasing complexity in nature.

HyperNEAT is a contemplation on the relationship between the geometry of the outside world, which is often regular, with the geometry of the inner world of the mind. Novelty search is a depiction of nature as a free spirit and a rebel – a process that becomes endlessly playful and innovative, but only when it is not told what to do or where it should go. These could be ultimate truths, or as art simply my personal truths, feelings about the world expressed by algorithms that thereby become their own stepping stones to new ideas. Sometimes algorithms can be odes to certain aspects of nature rather than faithful reproductions of them. And so it goes with art.

5 Synthesis

The traditions and culture of each scientific discipline in part help to protect its practitioners from a flood of baseless speculation (in what Kuhn [21] calls "normal science"). Each of us can only read so much, and the expectation that claims are supported by objective evidence helps to filter the discourse of the field to only the most scientifically rigorous.

Yet the danger of a flood of mediocrity trades off with the danger of pruning too much [14]. The risk of lost opportunity arises in particular if the point of the idea is subjective. In such cases, no amount of objective evidence can flip the subjective into the objective. Especially in the sciences of the artificial, where inexact reproduction of nature itself is on offer, there may *never* be such an objective resolution for some questions. Are the spaceships in figure 1 *objectively* reminiscent of some deeper aspect of nature?

The question then is whether we would be better off not sharing such subjective impressions (or only sharing them within the narrow context of art-based subdisciplines like evolutionary art). My suggestion is that the response to this question should be carefully balanced. It is clear that traditional objective evidence is essential to scientific progress, but some room for the discussion of subjective impression would be healthy for the fields as well. That way, a future researcher can share his or her "spaceships" with the scientific community without the need for objective pretense and maybe the next CPPNs will result after some discussion.

More broadly, an author should have the space to point to resonance with nature, both in the results of experiments and in the motivation for algorithms themselves. In some cases, such resonance should even trump performance. While a good algorithm is sometimes one that performs well, sometimes a good algorithm is instead *one that leads to other algorithms and new frontiers*.

The inevitable protest is *how can we know* whether such slippery intuitive ideas are genuinely good.

Might some clever author pull the wool over our eyes and fool us with pretty pictures? But that is the same problem faced by art. It is a messy problem to address, but that does not imply it should be avoided.

The answer is simply that we are qualified to do this. Professional scientists in the sciences of the artificial have spent years observing nature and thinking about its connection to computation. While our opinions may differ (as humans will inevitably differ on art as well), that does not mean our opinions are baseless or without substance. Rather, as experts in these fields, we are the most qualified to judge such resonance and discuss it openly. Feeling is behind all inspiration, and an iron rule that it must be hidden implies that the spark of inspiration can never light from one scientist to another unless through entirely objective means. How can that really make sense in fields predicated on the artificial interpretation of reality?

References

[1] Ackerman, J. S. (2002). *Origins, imitation, conventions: Representation in the visual arts*. Cambridge, MA: The MIT Press.

[2] Arnason, H., and Wheeler, D. (1986). *History of Modern Art: Painting, Sculpture, Architecture, Photography*. Prentice-Hall.

[3] Beveridge, W. I. B. (1950). *The art of scientific investigation.* Melbourne, London, Toronto: William Heinemann Ltd.

[4] Bialostocki, J. (1988). *The Renaissance Concept of Nature and Antiquity*. Austria: IRSA.

[5] Boden, M. A., and Edmonds, E. A. (2009). What is generative art? *Digital Creativity*, 20(1-2):21–46.

[6] Calaprice, A. (2000). *The Expanded Quotable Einstein*. Princeton, NJ: Princeton University Press.

[7] Carroll, S. B. (2005). *Endless forms most beautiful: the new science of evo devo and the making of the animal kingdom*. WW Norton & Company.

[8] Channon, A. (2003). Improving and still passing the alife test: Component-normalised activity statistics classify evolution in geb as unbounded. *Artificial Life*, 8:173–181.

[9] Cully, A., and Mouret, J.-B. (2013). Behavioral repertoire learning in robotics. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, GECCO ’13, 175–182. New York, NY, USA: ACM.

[10] da Vinci, L. (1958). Treatise on painting. 41,55.

[11] Dorin, A. (2001). Generative processes and the electronic arts. *Organised Sound*, 6(01):47–53.

[12] Fagerlund, M. (2005). DelphiNEAT-based genetic art homepage. `http://www.cambrianlabs.com/mattias/GeneticArt/`.

[13] Ferstl, H. (2006). SharpNEAT-based genetic art homepage. `http://www.cs.ucf.edu/~{}kstanley/GenArt.zip`.

[14] Feyerabend, P. (1975). *Against method*. Verso Books.

[15] Funch, B. S. (1997). *The psychology of art appreciation*. Museum Tusculanum Press.

[16] Gabriel, R. P., and Sullivan, K. J. (2010). Better science through art. In *Proceedings of the ACM international conference on Object oriented programming systems languages and applications (OOPSLA 2010)*, 885–900. New York, NY, USA: ACM.

[17] Gauci, J., and Stanley, K. O. (2010). Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation*, 22(7):1860–1898.

[18] Gombrich, E. (1960). *Art and Illusion: A Study in the Psychology of Pictorial Representation*. London: Phaidon Press Limited, Regent's Wharf.

[19] Kemp, M. (1990). *The Science of Art: Optical themes in western art from Brunelleschi to Seurat*. New Haven, CT: Yale University Press.

[20] Kemp, M. (2006). *Seen Unseen: Art, Science, and Intuition from Leonardo to the Hubble Telescope*. Oxford University Press.

[21] Kuhn, T. S. (1970). *The Structure of Scientific Revolutions*. International Encyclopedia of Unified Science, vol. 2, no. 2. Chicago: The University of Chicago Press. Second, enlarged edition.

[22] Lehman, J., and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*. Cambridge, MA: MIT Press.

[23] Lehman, J., and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.

[24] Lehman, J., and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *GECCO '11: Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 211–218. Dublin, Ireland: ACM.

[25] Merriam-Webster, I. (2014). *The Merriam-Webster Dictionary Online*.

[26] Navasky, V. S. (2013). *The Art of Controversy: Political Cartoons and Their Enduring Power*. New York, NY, USA: Alfred A. Knopf Incorporated.

[27] Ofria, C., and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial life*, 10(2):191–229.

[28] Ramachandran, V. S., and Hirstein, W. (1999). The science of art: A neurological theory of aesthetic experience. *Journal of Consciousness Studies*, 6(6–7):15–51.

[29] Risi, S., Hughes, C., and Stanley, K. (2011). Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 18(6):470–491.

[30] Romero, J. J., and Machado, P. (2008). *The art of artificial evolution: A handbook on evolutionary art and music*. Springer.

[31] Russell, S. J., and Norvig, P. (2009). *Artificial intelligence: A modern approach (3rd edition)*. Englewood Cliffs, NJ: Prentice Hall.

[32] Schmidhuber, J. (2009). Simple algorithmic theory of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *Journal of SICE*, 48(1):21–32.

[33] Secretan, J., Beato, N., D.Ambrosio, D. B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J. T., and Stanley, K. O. (2011). Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation*, 19(3):345–371.

[34] Secretan, J., Beato, N., D'Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008). Picbreeder: Evolving pictures collaboratively online. In *CHI '08: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 1759–1768. New York, NY, USA: ACM.

[35] Shubnikov, A. V., Koptsik, V. A., and Harker, D. (1974). *Symmetry in science and art*. New York,NY: Plenum Press.

[36] Simon, H. A. (1981). *The Sciences of the Artificial*. Cambridge, MA: MIT Press. Second edition.

[37] Sims, K. (1991). Artificial evolution for computer graphics. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '91)*, 319–328. New York, NY: ACM Press.

[38] Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162.

[39] Stanley, K. O., D'Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.

[40] Stanley, K. O., and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127.

[41] Stanley, K. O., and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.

[42] Stanley, K. O., and Miikkulainen, R. (2004). Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100.

[43] van Gogh, V. (1889). *Starry Night*. New York, NY: The Museum of Modern Art. Oil on canvas.

Online Appendix to Art in the Sciences of the Artificial

Kenneth O. Stanley (kstanley@cs.ucf.edu)
Department of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816

Appendix A: Exceptions: Systems Recognized for their Resonance

Although it is uncommon that achievements in AI or alife are explicitly recognized for their resonance with nature, there are some exceptions. Examples include John Conway's Game of Life [20] (reminiscent of elementary particle interactions in nature leading to organized structures), Craig Reynold's boids algorithm [83] (reminiscent of animals flocking), Aristid Lindenmayer's L-systems [68, 69] (reminiscent of the growth of plants), Thomas Ray's Tierra [81, 82] (reminiscent of evolution within natural ecologies), Karl Sims' evolved three-dimensional creatures [93] (reminiscent of the ambulation of natural organisms), etc.

Appendix B: Commentaries on Art within Science

The view expressed by Beveridge [7] that that science can benefit from artistic intuition is echoed in many places [34, 53, 73, 89, 102], often focusing on particular angles like imagery [73] (which is also an important focus in this paper) or creativity [89].

Appendix C: Evolutionary Art

Evolutionary art [5, 29, 43, 70, 72, 92, 94] is a branch of interactive evolutionary computation (IEC; [98]) going back decades to Dawkins' early work with Biomorphs reported in his book, *The Blind Watchmaker* [28], and the work of Todd and Latham [99].

Appendix D: Abstractions of Biological Development in Evolutionary Computation

A number of artificial encodings meant to abstract essential properties of biological development were conceived prior to CPPNs [1, 6, 8, 50].

Appendix E: Applications and Extensions of CPPNs and HyperNEAT

The introduction of CPPNs spawned a new research area with numerous applications and extensions [2, 3, 4, 9, 10, 11, 12, 13, 14, 15, 17, 18, 19, 24, 25, 26, 27, 31, 32, 33, 35, 35, 36, 36, 44, 45, 46, 47, 48, 49, 51, 57, 74, 80, 84, 86, 90, 91, 96, 97, 101]. In fact, because CPPNs are a kind of generic pattern-generator, their applications are diverse, including music [45, 46, 47, 48, 49], three-dimensional objects [15], dance [33], more evolutionary art [90, 91] (even winning awards at an art competition [78]), robot morphologies [3, 11, 86], and the design of robust mechanical structures [12, 84].

Perhaps the most significant new direction based on CPPNs is the HyperNEAT method for evolving large-scale artificial neural networks, which itself opened up a new research area [19, 24, 25, 26, 27, 32, 35, 36, 36, 44, 57, 74, 85, 97, 101]).

Appendix F: Applications and Extensions of Novelty Search

The introduction of novelty search led to numerous applications and extensions [16, 21, 22, 23, 30, 37, 38, 39, 40, 41, 42, 52, 54, 55, 56, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 71, 74, 75, 76, 77, 79, 85, 87, 88, 95, 100, 103, 104]. Some of these works apply it to practical problems like biped or quadruped locomotion [62, 74, 85] while others explore its potential as an automated tool for discovery and invention [23, 63].

References

[1] Angeline, P. J. (1995). Morphogenic evolutionary computations: Introduction, issues and examples. In McDonnell, J. R., Reynolds, R. G., and Fogel, D. B., editors, *Evolutionary Programming IV: The Fourth Annual Conference on Evolutionary Programming*, 387–401. MIT Press.

[2] Auerbach, J., and Bongard, J. (2010). Evolving CPPNs to grow three dimensional structures. In

Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2010). New York, NY: ACM Press.

[3] Auerbach, J. E., and Bongard, J. C. (2012). On the relationship between environmental and mechanical complexity in evolved robots. In *Proceedings of 13th International Conference on the Synthesis and Simulation of Living Systems (ALife XIII)*.

[4] Bahceci, E., and Miikkulainen, R. (2008). Transfer of evolved pattern-based heuristics in games. In *Proceedings of the IEEE Symposium on Computational Intelligence and Games (CIG-2008)*. Piscataway, NJ: IEEE Press.

[5] Baluja, S., Pomerleau, D., and Jochem, T. (1994). Towards automated artificial evolution for computer-generated images. *Connection Science*, 6(2 and 3):325–354.

[6] Bentley, P. J., and Kumar, S. (1999). Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-1999)*, 35–43. San Francisco: Kaufmann.

[7] Beveridge, W. I. B. (1950). *The art of scientific investigation.* Melbourne, London, Toronto: William Heinemann Ltd.

[8] Bongard, J. C., and Pfeifer, R. (2001). Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In Spector, L., Goodman, E. D., Wu, A., Langdon, W. B., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., and Burke, E., editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, 829–836. San Francisco: Kaufmann.

[9] Buk, Z., Koutnk, J., and Snorek, M. (2009). NEAT in HyperNEAT substituted with genetic programming. In *Inernational Conference on Adaptive and Natural Computing Algorithms (ICANNGA-2009)*, 243–252. Berlin: Springer.

[10] Cheney, N., Clune, J., and Lipson, H. (2014). Evolved electrophysiological soft robots. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, vol. 14, 222–229.

[11] Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013)*. New York, NY: ACM Press.

[12] Cheney, N., Ritz, E., and Lipson, H. (2014). Automated vibrational design and natural frequency tuning of multi-material structures. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO '14, 1079–1086. New York, NY, USA: ACM.

[13] Clune, J., Beckmann, B., McKinley, P., and Ofria, C. (2010). Investigating whether HyperNEAT produces modular neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. New York, NY: ACM Press.

[14] Clune, J., Beckmann, B., Pennock, R., and Ofria, C. (2009). HybrID: A hybridization of indirect and direct encodings for evolutionary computation. In *Proceedings of the European Conference on Artificial Life (ECAL-2009),*.

[15] Clune, J., and Lipson, H. (2011). Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In *Proceedings of the European Conference on Artificial Life (ECAL-2011)*, 141–148.

[16] Clune, J., Mouret, J.-B., and Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings of the Royal Society of London B: Biological Sciences*, 280:20122863.

[17] Clune, J., Ofria, C., and Pennock, R. (2008). How a generative encoding fares as problem-regularity decreases. In *Proceedings of the 10th International Conference on Parallel Problem Solving From Nature (PPSN 2008)*, 258–367. Berlin: Springer.

[18] Clune, J., Pennock, R. T., and Ofria, C. (2009). The sensitivity of HyperNEAT to different geometric representations of a problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2009)*. New York, NY, USA: ACM Press.

[19] Clune, J., Stanley, K. O., Pennock, R. T., and Ofria, C. (2011). On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*.

[20] Conway, J. (1970). The game of life. *Scientific American*, 223(4):4.

[21] Cuccu, G., Gomez, F., and Glasmachers, T. (2011). Novelty-based restarts for evolution strategies. In *2011 IEEE Congress on Evolutionary Computation (CEC)*, 158–163. IEEE.

[22] Cully, A., Clune, J., and Mouret, J. (2014). Robots that can adapt like natural animals. *CoRR*, abs/1407.3501.

[23] Cully, A., and Mouret, J.-B. (2013). Behavioral repertoire learning in robotics. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, GECCO '13, 175–182. New York, NY, USA: ACM.

[24] D'Ambrosio, D., Lehman, J., Risi, S., and Stanley, K. O. (2010). Evolving policy geometry for scalable multiagent learning. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2010)*, 731–738. International Foundation for Autonomous Agents and Multiagent Systems.

[25] D'Ambrosio, D., Lehman, J., Risi, S., and Stanley, K. O. (2011). Task switching in multiagent learning through indirect encoding. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 2011)*. Piscataway, NJ: IEEE.

[26] D'Ambrosio, D. B., Goodell, S., Lehman, J., Risi, S., and Stanley, K. O. (2012). Multirobot behavior synchronization through direct neural network communication. In *Proceedings of the 5th International Conference on Intelligent Robotics and Applications (ICIRA-2012)*. New York, NY: Springer-Verlag.

[27] D'Ambrosio, D. B., and Stanley, K. O. (2013). Scalable multiagent learning through indirect encoding of policy geometry. *Evolutionary Intelligence*, 6(1).

[28] Dawkins, R. (1986). *The Blind Watchmaker.* Essex, U.K.: Longman.

[29] Dawkins, R. (1989). The evolution of evolvability. In Langton, C. G., editor, *Artificial Life*, 201–220. Reading, MA: Addison-Wesley.

[30] Doucette, J., and Heywood, M. I. (2010). Novelty-based fitness: An evaluation under the santa fe trail. In *Proceedings of the European Conference on Genetic Programming (EuroGP-2010*, 50–61. Springer.

[31] Drchal, J., Kapra, O., Koutnik, J., and Snorek, M. (2009). Combining multiple inputs in HyperNEAT mobile agent controller. In *19th International Conference on Artificial Neural Networks (ICANN 2009)*, 775–783. Berlin: Springer.

[32] Drchal, J., Koutnk, J., and Snorek, M. (2009). HyperNEAT controlled robots learn to drive on roads in simulated environment. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC-2009)*. Piscataway, NJ, USA: IEEE Press.

[33] Dubbin, G., and Stanley, K. O. (2010). Learning to dance through interactive evolution. In *Proceedings of the Eighth European Event on Evolutionary and Biologically Inspired Music, Sound, Art and Design (EvoMUSART 2010)*. New York, NY: Springer.

[34] Gabriel, R. P., and Sullivan, K. J. (2010). Better science through art. In *Proceedings of the ACM international conference on Object oriented programming systems languages and applications (OOPSLA 2010)*, 885–900. New York, NY, USA: ACM.

[35] Gauci, J., and Stanley, K. O. (2008). A case study on the critical role of geometric regularity in machine learning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-2008)*. Menlo Park, CA: AAAI Press.

[36] Gauci, J., and Stanley, K. O. (2010). Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation*, 22(7):1860–1898.

[37] Goldsby, H. J., and Cheng, B. H. (2010). Automatically discovering properties that specify the latent behavior of uml models. In *Model Driven Engineering Languages and Systems*, 316–330. Springer.

[38] Gomes, J., and Christensen, A. L. (2013). Generic behaviour similarity measures for evolutionary swarm robotics. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, GECCO '13, 199–206. New York, NY, USA: ACM.

[39] Gomes, J., Mariano, P., and Christensen, A. L. (2014). Avoiding convergence in cooperative coevolution with novelty search. In *Proceedings of the 2014 international conference on autonomous agents and multi-agent systems (AAMAS '14)*, 1149–1156. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

[40] Gomes, J., Mariano, P., and Christensen, A. L. (2014). Novelty search in competitive coevolution. In *Parallel Problem Solving from Nature PPSN XIII*, 233–242. Springer.

[41] Gomes, J., Urbano, P., and Christensen, A. L. (2013). Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, 1–30.

[42] Graening, L., Aulig, N., and Olhofer, M. (2010). Towards directed open-ended search by a novelty guided evolution strategy. In Schaefer, R., Cotta, C., Kołodziej, J., and Rudolph, G., editors, *Parallel Problem Solving from Nature – PPSN XI*, vol. 6239 of *Lecture Notes in Computer Science*, 71–80. Springer.

[43] Greenfield, G. R. (2000). Evolving expressions and art by choice. *Leonardo*, 33(2):93–99.

[44] Haasdijk, E., Rusu, A. A., and Eiben, A. (2010). HyperNEAT for locomotion control in modular robots. In *Proceedings of the 9th International Conference on Evolvable Systems (ICES 2010),*.

[45] Hoover, A. K., Rosario, M. P., and Stanley, K. O. (2008). Scaffolding for interactively evolving novel drum tracks for existing songs. In Giacobini, M., editor, *Proceedings of the Sixth European Workshop on Evolutionary and Biologically Inspired Music, Sound, Art and Design (EvoMUSART 2008)*, 412–422. Springer.

[46] Hoover, A. K., and Stanley, K. O. (2009). Exploiting functional relationships in musical composition. *Connection Science Special Issue on Music, Brain, and Cognition*, 21(2 and 3):227–251.

[47] Hoover, A. K., Szerlip, P. A., , and Stanley, K. O. (2011). Generating musical accompaniment through functional scaffolding. In *Proceedings of the 8th Sound and Music Computing Conference (SMC-2011)*.

[48] Hoover, A. K., Szerlip, P. A., and Stanley, K. O. (2011). Interactively evolving harmonies through functional scaffolding. In *GECCO '11: Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 387–394. Dublin, Ireland: ACM.

[49] Hoover, A. K., Szerlip, P. A., and Stanley, K. O. (2012). Generating a complete multipart musical composition from a single monophonic melody with functional scaffolding. In Maher, M. L., Hammond, K., Pease, A., Perez, R. P. Y., Ventura, D., and Wiggins, G., editors, *Proceedings of the 3rd International Conference on Computational Creativity (ICCC-2012)*.

[50] Hornby, G. S., and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, 8(3).

[51] Jaskowski, W., Krawiec, K., and Wieloch, B. (2008). Neurohunter - an entry for the balanced diet contest. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2008) Contest Program")*. New York, NY: ACM Press.

[52] Joachimczak, M., and Wrobel, B. (2012). Open ended evolution of 3D multicellular development controlled by gene regulatory networks. In *Proceedings of the Thirteenth International Conference on Artificial Life (ALIFE XIII)*, 67–74. Cambridge, MA: MIT Press.

[53] Kemp, M. (2006). *Seen Unseen: Art, Science, and Intuition from Leonardo to the Hubble Telescope*. Oxford University Press.

[54] Kistemaker, S., and Whiteson, S. (2011). Critical factors in the performance of novelty search. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, GECCO ’11, 965–972.

[55] Krcah, P. (2010). Solving deceptive tasks in robot body-brain co-evolution by searching for behavioral novelty. In *10th International Conference on Intelligent Systems Design and Applications (ISDA)*, 284–289. IEEE.

[56] Krcah, P., and Toropila, D. (2010). Combination of novelty search and fitness-based search applied to robot body-brain co-evolution. In *Proceedings of the 13th Czech-Japan Seminar on Data Analysis and Decision Making in Service Science*.

[57] Lee, S., Yosinski, J., Glette, K., H, H. L., and Clune, J. (2013). Evolving gaits for physical robots with the hyperNEAT generative encoding: the benefits of simulation. In *Applications of Evolutionary Computing*, 540–549. New York, NY: Springer.

[58] Lehman, J., and Miikkulainen, R. (2014). Overcoming deception in evolution of cognitive behaviors. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO ’14, 185–192. New York, NY, USA: ACM.

[59] Lehman, J., and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In Bullock, S., Noble, J., Watson, R., and Bedau, M., editors, *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*. Cambridge, MA: MIT Press.

[60] Lehman, J., and Stanley, K. O. (2010). Efficiently evolving programs through the search for novelty. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, GECCO ’10, 837–844. New York, NY, USA: ACM.

[61] Lehman, J., and Stanley, K. O. (2010). Revising the evolutionary computation abstraction: minimal criteria novelty search. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, GECCO ’10, 103–110. New York, NY, USA: ACM.

[62] Lehman, J., and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.

[63] Lehman, J., and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *GECCO '11: Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 211–218. Dublin, Ireland: ACM.

[64] Lehman, J., and Stanley, K. O. (2011). Novelty search and the problem with objectives. In *Genetic Programming Theory and Practice IX (GPTP 2011)*. New York, NY: Springer.

[65] Li, J., Storie, J., and Clune, J. (2014). Encouraging creative thinking in robots improves their ability to solve challenging problems. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO '14, 193–200. New York, NY, USA: ACM.

[66] Liapis, A., Martínez, H. P., Togelius, J., and Yannakakis, G. N. (2013). Transforming exploratory creativity with delenox. In *Proceedings of the Fourth International Conference on Computational Creativity*.

[67] Liapis, A., Yannakakis, G. N., and Togelius, J. (2013). Enhancements to constrained novelty search: two-population novelty search for generating game content. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, GECCO '13, 343–350. New York, NY, USA: ACM.

[68] Lindenmayer, A. (1968). Mathematical models for cellular interaction in development parts I and II. *Journal of Theoretical Biology*, 18:280–299 and 300–315.

[69] Lindenmayer, A. (1974). Adding continuous components to L-systems. In Rozenberg, G., and Salomaa, A., editors, *L Systems, Lecture Notes in Computer Science 15*, 53–68. Heidelberg, Germany: Springer-Verlag.

[70] Lund, H. H., Pagliarini, L., and Miglino, P. (1995). Artistic design with GA and NN. In *Proceedings of the 1st Nordic Workshop on Genetic Algorithms and Their Applications (1NWGA)*, 97–105.

[71] Martinez, Y., Naredo, E., Trujillo, L., and Galvan-Lopez, E. (2013). Searching for novel regression functions. In *2013 IEEE Congress on Evolutionary Computation (CEC)*, 16–23. IEEE.

[72] McCormack, J. P. (1993). Interactive evolution of L-system grammars for computer graphics modelling. In Green, D. G., and Bossomaier, T., editors, *Complex Systems: From Biology to Computation*, 118–130. Amsterdam: IOS Press.

[73] Miller, A. I. (2000). *Insights of genius: Imagery and creativity in science and art*. Cambridge, MA: MIT Press.

[74] Morse, G., Risi, S., Snyder, C. R., and Stanley, K. O. (2013). Single-unit pattern generators for quadruped locomotion. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013)*. New York, NY, USA: ACM.

[75] Mouret, J.-B. (2011). Novelty-based multiobjectivization. In *New Horizons in Evolutionary Robotics*, 139–154. Berlin Heidelberg: Springer.

[76] Mouret, J.-B., and Doncieux, S. (2012). Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evolutionary computation*, 20(1):91–133.

[77] Naredo, E., and Trujillo, L. (2013). Searching for novel clustering programs. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, GECCO '13, 1093–1100.

[78] Nguyen, A., Yosinski, J., and Clune, J. (2015). Deep neural networks are easily fooled: High confidence predictions for unrecognizable images.

[79] Ollion, C., and Doncieux, S. (2011). Why and how to measure exploration in behavioral space. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, GECCO '11, 267–274. New York, NY, USA: ACM.

[80] Pugh, J. K., and Stanley, K. O. (2013). Evolving multimodal controllers with HyperNEAT. In *Proceeding of the fifteenth annual conference on genetic and evolutionary computation*, GECCO '13, 735–742. New York, NY, USA: ACM.

[81] Ray, T. S. (1992). An approach to the synthesis of life. In *Proceedings of Artificial Life II*, 371–408.

[82] Ray, T. S. (1993). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1(1_2):179–209.

[83] Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.*, 21(4):25–34.

[84] Richards, D., and Amos, M. (2014). Evolving morphologies with CPPN-NEAT and a dynamic substrate. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, vol. 14, 255–262.

[85] Risi, S., , and Stanley, K. O. (2013). Confronting the challenge of learning a flexible neural controller for a diversity of morphologies. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013)*. New York, NY, USA: ACM.

[86] Risi, S., Cellucci, D., and Lipson, H. (2013). Ribosomal robots: Evolved designs inspired by protein folding. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013)*. New York, NY: ACM.

[87] Risi, S., Hughes, C., and Stanley, K. (2011). Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 18(6):470–491.

[88] Risi, S., Vanderbleek, S. D., Hughes, C. E., and Stanley, K. O. (2009). How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2009)*. New York, NY, USA: ACM Press.

[89] Root-Bernstein, R. S. (1989). *Discovering*. Cambridge, MA: Harvard University Press.

[90] Secretan, J., Beato, N., D.Ambrosio, D. B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J. T., and Stanley, K. O. (2011). Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation*, 19(3):345–371.

[91] Secretan, J., Beato, N., D'Ambrosio, D. B., Rodriguez, A., Campbell, A., and Stanley, K. O. (2008). Picbreeder: Evolving pictures collaboratively online. In *CHI '08: Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, 1759–1768. New York, NY, USA: ACM.

[92] Sims, K. (1991). Artificial evolution for computer graphics. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '91)*, 319–328. New York, NY: ACM Press.

[93] Sims, K. (1994). Evolving 3D morphology and behavior by competition. In Brooks, R. A., and Maes, P., editors, *Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems (Artificial Life IV)*, 28–39. Cambridge, MA: MIT Press.

[94] Smith, J. R. (1991). Designing biomorphs with an interactive genetic algorithm. In Belew, R. K., and Booker, L. B., editors, *Proceedings of the 4th International Conference on Genetic Algorithms (ICGA-91)*, 535–538. San Mateo, CA: Morgan Kaufmann.

[95] Soltoggio, A., and Jones, B. (2009). Novelty of behaviour as a basis for the neuro-evolution of operant reward learning. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, GECCO '09, 169–176. New York, NY, USA: ACM.

[96] Soros, L., and Stanley, K. O. (2014). Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria. In *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 793–800.

[97] Stanley, K. O., D'Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212.

[98] Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capacities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296.

[99] Todd, S., and Latham, W. (1992). *Evolutionary Art and Computers*. London: Academic Press.

[100] Velez, R., and Clune, J. (2014). Novelty search creates robots with general skills for exploration. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO '14, 737–744. New York, NY, USA: ACM.

[101] Verbancsics, P., and Stanley, K. O. (2010). Evolving static representations for task transfer. *Journal of Machine Learning Research (JMLR)*, 11:1737–1769.

[102] Wolpert, L. (1994). *The unnatural nature of science*. Cambridge, MA: Harvard University Press.

[103] Woolley, B. G., and Stanley, K. O. (2011). On the deleterious effects of a priori objectives on evolution and representation. In *GECCO '11: Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 957–964. Dublin, Ireland: ACM.

[104] Woolley, B. G., and Stanley, K. O. (2012). Exploring Promising Stepping Stones by Combining Novelty Search with Interactive Evolution. *ArXiv e-prints*.