MINIMAL CREATIVITY, EVALUATION AND FRACTAL PATTERN DISCRIMINATION

Jon Bird
Creative Systems Lab
University of Sussex
Brighton, BN1 9QG, UK
jonba@sussex.ac.uk

Dustin Stokes
Centre for Research in Cognitive Science
University of Sussex
Brighton, BN1 9QG, UK
d.stokes@sussex.ac.uk

Abstract

Evaluation is considered to be an important component of any creative process. This paper explores how evaluation can be incorporated into our minimal model of creativity, which we have been developing using a combination of conceptual analysis and evolutionary robotics. Specifically, we consider how to extend our approach so that the robots themselves can evaluate mark patterns that they, or other robots, have made on the floor of their environment.

Evaluation can, we suggest, be distinguished into a descriptive and a merit assignment component. To evaluate an object or event F is, (a) to discriminate some feature of F and (b) to assign some merit or de-merit to that feature of F. Component (a) is descriptive and (b) is what would more traditionally be called 'evaluative.'

In simulation, our robots discriminate fractal from random patterns and demonstrate this by stopping on target regions of the arena floor that are covered with a fractal texture. We argue that in so doing they perform the descriptive component of evaluation. However, it is debatable whether the robots are performing the second, merit-assignment component of evaluation. Currently, the discrimination mechanism is hard-wired and does not develop during an agent's lifetime. In future experiments we will investigate how to artificially evolve agents that perform what might be described as 'minimal evaluation' by attempting to incorporate a preference element.

Keywords: Minimal creativity, evolutionary robotics, evaluation, fractal pattern discrimination

1 Introduction

A creative process, one might argue, involves evaluation. Without the second, you don't get the first. Consider Harold Cohen's AARON drawing system, which has now been generating images for more than 30 years (McCorduck, 1991). The computational structure of AARON has changed over the years. These details are theoretically interesting, but set them to one side for the moment. No matter the computation involved, and no matter the status of the artificial agency of AARON, the following is clear. Cohen, in selecting the images to display and eventually sell, must evaluate those images. He selects the ones that he prefers, and these images are the ones that you and I will see in a gallery. This is evaluation if anything is. And without it, one might say, you don't have a creative process. In fact, this may be the reason that one hesitates to call AARON creative: without Cohen, there is no evaluation. Indeed, Cohen makes the same concession (Cohen, 1999). This may motivate one to take evaluation to be a necessary condition for creative processes.

We are sympathetic to this general suggestion. And we are willing to take it seriously enough to incorporate it into our minimal approach to modelling creativity and to see what new results it may engender. Our ultimate goal is to evolve artificial agents which behave in, at least minimally, creative ways. Our methodology is also minimal, both in terms of conceptual assumptions, and the constraints we place on the robot controllers.

2 Minimal Modelling and Creativity

To this point, our conceptual assumptions about creativity have been sparse. First, we propose that creative processes require some degree of agency. And agency involves autonomy. We understand autonomy very broadly: it includes any behaviour not strictly imposed by a programmer or designer. This 'no strings attached' agency is weaker than a rich philosophical notion; it does not require deliberation or cognition. Thus a remote controlled robot would not be an agent in our sense, while an individual agent in an artificial life simulation would be. Second, we suggest that creative processes involve novelty. We follow Boden (2004) in rejecting the assumption that only historical novelty is theoretically interesting. Instead, one can acknowledge relative novelty by specifying different reference points. We are interested in two types of relative novelty. Some behaviour is *population-relative novel* for some agent A just in case that behaviour is novel relative to all of the agents in a population of which A is a member. And a behaviour is *individual-relative novel* for some agent A just in case it is novel relative to the behavioural history of A. This provides a great deal of flexibility in how we understand novelty. And of course, as the richness of the creativity in question goes up, so does the novelty, perhaps moving towards something like historical novelty. We thus propose an agency and a novelty condition for minimally creative processes (Bird and Stokes, 2006a,b; Bird et al., 2007).

To meet the 'no strings attached' agency condition our initial model consists of a simulated robot, situated in a walled arena environment, whose behaviour is solely determined by its sensory motor activity. The simulation is based on a Khepera robot with 6 IR sensors, a floor sensor that can detect lines directly below the robot and a pen that can be raised and lowered. Given the difficulty of hand designing robot-environment interactions we employ an evolutionary robotics approach to designing the artificial neural network that controls the behaviour of the robot. The fitness function explicitly rewards (a) correlated changes in the the state of the line sensor ('on to off' and 'off to on') and the pen ('up to down' and 'down to up'), and (b) marks made over a large area of the arena floor. It also implicitly penalises robots that crash into walls. This methodology also enables us to minimise any implicit or explicit biases we have about how creativity should be modelled as well as having the potential to generate models that exhibit unpredictable, and potentially novel behaviour. Our initial results demonstrate that we can evolve simulated robots that meet our minimal agency and population-relative novelty conditions (Bird et al., 2007).

3 Missing Evaluation

We've encountered many sceptics in presenting this work over the past year or two. The sceptic will often say something like the following. "You've got agents that perform reactive behaviours, some of them novel in your specified sense, but there is no space for these agents to make choices, to judge what they are doing, to evaluate. And this is what's missing: no evaluation, no creative process. So call the behaviours 'minimally creative' if you like, but they are too minimal to connect with rich creativity in any interesting way." As we said at the outset, we are sympathetic to this general line of criticism. It needs, however, to be separated into distinct criticisms, some of which we will address here, some of which we hope to address later in our research.

1. *Non-evaluative processing worry*
   This is the issue we are most concerned with at present. For agents to act creatively, they must be performing some kind of evaluation. The agent must be choosing among features or stimuli in its environment, where these choices inform (and indeed may be derived from) their behavioural activity. The mark-making behaviour of our agents lacks this feature. Their behaviours are mere reactions to the arena boundaries and previous mark-making in the arena, constrained by their own sensory motor morphologies and, from an evolutionary perspective, the performance of previous agents in that environment. There is, it seems, no evaluation by the individual robots of their mark-making activity.

2. *Myopic worry*
   A related worry concerns a feature of the sensory motor morphology of the agents we are using. As a symptom of their physical structure, our simulated mark-making agents can only see marks on the floor below them in a 2mm by 2mm area. There is thus no sense in which the agent can achieve a global perspective on the marks it is making. There is not, as there is usually for a human artist, an opportunity to 'stand back' and consider the overall pattern. This underwrites the non-evaluative processing worry, since this myopic perspective seriously limits prospects for evaluation.

3. *No-stopping worry*
   Our agents have no stopping mechanism. That is, there is no point where the agent will complete the marks, no analogue to the artist who happily steps back and says to herself 'It's finished!'. Our agents may stop, but any termination of mark-making is independent of the patterns made by the marks. This, in the theoretical context of creativity, is a problem.

4. *Aesthetic value worry*
   A final worry concerns the results of the agents' behaviours rather than the behaviours themselves. Some members of our research team would ultimately like to see some robot drawings which can be exhibited in a gallery. Controversies about aesthetic value and definitions of art to one side, displaying works in a gallery generally requires that the works possess some aesthetic merit by one criterion or other. Our results may be interesting, provided an observer has enough information about the artificial agency that generated them. But on their own, it is debatable whether they have any aesthetic interest. The challenge, then, is to achieve some aesthetically interesting results, while maintaining our theoretical stance towards creativity. In other words, we want our robots to create something aesthetically valuable, while minimising our influence over how they do it. So far, we fall short of this goal.

4 A Proposed Solution: The Fractal Framework

We think one solution to these various worries lies in the use of fractals. In simple terms, we intend to endow our agents with a capacity for discriminating fractal patterns, and a preference of sorts for completing such patterns. This is a broad enough constraint on our agents to allow, as it were, creative freedom. Fractal patterns – understood broadly as patterns which display self-similarity across a range of scales – can vary greatly in their appearance (for example, texture), and so our agents will only have a general pattern 'preference'. The resulting images thus stand to be surprising or unexpected, but the behaviours of the agents are nonetheless constrained: some images are fractal, others are not [1].

Before discussing the details of our modelling strategies to these ends, we should ask, from a purely conceptual point of view, how a fractal framework helps with the four evaluation worries described in Section 3. We consider each in turn.

1. *R1: The non-evaluative processing worry*
   One simple reason our agents do not evaluate is that they don't have anything to *look for*. They are not, at present, pattern discriminators and so *a fortiori* are not pattern evaluators. Incorporating fractal pattern discrimination and, eventually, a preference for making fractal mark patterns dissolves this particular aspect of the problem. The agents are thereby endowed with a minimal evaluation technique, preferring certain marks over others and acting in ways that display that basic choice [2].

2. *R2: The myopic worry*
   In order that our simulated agents can discriminate fractal patterns, rather than just detect marks, we are replacing the 2mm x 2mm mark detector with a 50mm x 50mm camera that points at the floor. However, even with the addition of this camera, the agent's viewpoint is still extremely limited and does not counter the myopic worry. This therefore remains a challenging problem and may have to be resolved by either using a bird's eye view camera or some type of memory in the agent's controller so that it can evaluate a larger part of the arena at any one time. This issue is independent of the fractal framework that is the focus of this paper and we do not address it further here.

3. *R3: The no-stopping worry*
   Incorporating fractal patterns provides a natural stopping point. An agent making a fractal pattern will stop once it has generated a pattern that is self-similar across that range of scales that it can discriminate. The agent thus finishes the mark-making *and* in a way that is dependent upon the patterns made.

4. *R4: Aesthetic value worry*
   This particular worry is perhaps too far downstream from our current state of research, so a suggestion for now will suffice. People like fractals. This seems true on an intuitive level and has in fact been demonstrated in the psychological literature (Spehar et al., 2003). What's more, if an audience member is armed with the knowledge that the agents she is watching are attempting to complete a fractal pattern, a very basic identification is achieved [3]. That is, the agents' behaviours are no longer unintelligibly random, since one understands an agent's general goal and may watch as the agent attempts to achieve that goal. A fractal framework thus affords potentially aesthetically valuable finished products, as well as a sharpened insight into the behaviours – and, perhaps, the creative processes – of the artificial agents. This, we suspect, moves us closer to the aesthetic merit needed for a gallery display.

5 Theoretical Analysis of Evaluation and Pattern Discrimination

What do we mean by 'evaluation'? Evaluation is a topic of rich debate in the philosophical literature ranging from aesthetics to moral theory to action theory, among others. This diversity provides a hint: evaluation is a context-bound enterprise and so theorizing it should be shaped by the relevant context. We are primarily concerned with art and art-making, so our analysis will err towards an analysis of aesthetic evaluation.

Most intuitively, to aesthetically evaluate some thing is to assign value to that thing. Less circularly, to aesthetically evaluate some thing is to assess the merit of that thing. We say that this work is 'good', that one 'bad', this one 'beautiful', that one 'poor.' We can, and often do, offer such assessments in general and unqualified ways. However, these assessments are typically made for particular reasons; and we provide these reasons if our assessment is called into question. I assign merit to some work because it has this or that property, and I appeal to the latter in justifying the former. This reveals something important about aesthetic evaluation: aesthetic evaluation, at least typically, involves a descriptive element.

There is an historical dichotomy between the evaluative and the descriptive. The distinction is – as is often the case with supposed dichotomies – a fuzzy one. The 20th century aesthetician Frank Sibley offers an insightful analysis of this distinction, with special emphasis on its place in philosophical aesthetics. Sibley's emphasis is on evaluative versus descriptive terms, but so long as we take a use of such terms to be indicative of the corresponding judgment, we can generalize from his analysis to evaluative and descriptive acts (Sibley, 2001).

First, Sibley suggests, some terms are used to indicate that an F has value (or does not), without indicating why or how F has this value. One may, for example, call a thing 'good', 'bad', 'nice', 'nasty', 'worthless' and so on, without attributing any particular properties to that thing. Sibley calls such terms solely evaluative. These terms and their correlative use are better understood when contrasted with a second class of terms. Some terms indicate a property the possession of which is a merit (or de-merit) relative to some category. 'Sharp' is such a term relative to the category of knives (and, oppositely, so is 'dull'), 'level' for billiard tables, 'round' for basketballs, and so on. Such terms, Sibley suggests, are often taken for evaluative ones. However, we might instead think of them as straightforward property terms, since to use them is to as-

[1] This, in fact, is a point of controversy in the literature on fractals (Halley et al., 2004). Given the methods of fractal measurement we employ, however, some patterns will meet the standard while others will not. Whether this satisfies parties on both sides of the debate over what we might call 'fractal realism' is an open question.

[2] Whether this amounts to anything we reasonably call 'evaluation' is debatable. The non-evaluative processing worry is central to this paper, so more on this point in sections 5, 6, and 7.

[3] If and when we should be in a position to exhibit in a gallery, the display would be not merely of resulting images, but of the embodied robots making marks on the floor of a walled arena.

cribe a property to an object and indeed can be done without an additional evaluative assessment of that object. One can describe a knife as sharp without knowing that that sharpness enables proper performance of a knife's function. In Sibley's words, "[i]n general, one does not need to know, with such a term, 'P', though one often will, that the property counts as a merit in something in order to be able to ascertain that the thing may correctly be called 'P"' (Sibley, 2001, p.92). Indeed use of these terms is common in many spheres, including aesthetics and criticism. Consider critical and appreciative practice: we do offer descriptions of artworks (for example, of their formal, art-historical, generative, or socio-political properties) which may imply an assessment of merit, but which could be offered without such an assessment. The important point of contrast is that to use a term of this type, unlike solely evaluative terms, is to identify a particular feature of that object, and this is a descriptive rather than an evaluative act.

As Sibley admits, use of a descriptive merit term often, even if not by conceptual necessity, implies an assignment of merit. To this end, we might, again following Sibley, introduce a third category of terms: evaluation-added terms. Such terms involve both a descriptive and an evaluative element. In using a term of this type, one describes an F by indicating that F has some property G, and then adds to this description an assessment of merit, where this assessment takes place on the basis of the description of F's possession of G. Sibley offers 'tasty' as one example. If I call a meal 'tasty', I am no doubt giving a positive evaluation of the meal and this evaluation implies a certain description, namely, that the meal has a lot of flavour.

Sibley goes on to question whether aesthetic evaluation carves up so neatly, and indeed questions how much of so-called aesthetic evaluation is even evaluative rather than descriptive. No matter. We can safely glean the following lesson from Sibley's analysis. Aesthetic evaluation generally involves both a descriptive and a merit-assigning element; it roughly resembles the use of what Sibley calls 'evaluation-added terms.' Consider a familiar scene from a gallery or exhibit. My friend Jon says to me, "This sculpture is lovely." An eyebrow raised, I respond, "Really, how so?" Jon's response might go something as follows. "Well, it possesses a certain balance. Notice how the curve of her hip echoes the positioning of her opposite elbow. And the face is just expressive enough: the eyes are blank but wide open; the lips are slightly curled at the corners, not quite smiling and not quite frowning." Jon may go on and on until I have had enough. "You're right. It's lovely. Time for a drink". This is a simple example of appreciative practice. One may initially offer an unqualified evaluation of a work (using what Sibley calls a 'solely evaluative term'). However, as is often the case, one has reasons for making that evaluation and will offer them when asked. These reasons, as with Jon's reasons, often, perhaps always, involve descriptions of the thing evaluated. They involve an indication of the properties which underwrite one's assessment of merit. This is enough to motivate the following understanding of evaluation.

Evaluating an F involves:

1. d indicating that F possesses property G; and
2. m indicating that one finds merit/de-merit in G as possessed by F.

d is thus the descriptive element, and m the (traditionally) evaluative element. One can think about d and m as necessary and conjointly sufficient conditions for evaluation if one prefers, but we see no need to make this commitment. For our purposes, it suffices to say that evaluation, especially in contexts of aesthetic appreciation and criticism, typically involves both a property description d and an assignment of merit m. Indeed, this seems to capture a fundamental schema for evaluation, in whatever realm it should be. The difference, for example, between moral evaluation and aesthetic evaluation lies in the kind of merit that is assigned and, perhaps, in which properties are relevantly discriminated. Common to both kinds is the presence of a descriptive and an evaluative element.

We are taking a property description to be no more than the indication that an F has property G or, if one prefers, the discrimination of G as possessed by F. This kind of discrimination is clearly descriptive (as contrasted with evaluative), but it need not be any kind of rich description. Indeed it need not be linguistic: a picture, pointing at something (given the right context), the firing of a feature detecting neuron, or pushing one button rather than another could each just as well serve this indicative role. One may have reservations about calling these acts and events 'descriptions' given the heavy philosophical baggage that comes with this term. We are sensitive to these worries, and indeed our analysis of evaluation does not require us to think of descriptions in any special way. To be clear: the only point that need be granted is that evaluation involves (partly) an indication that the object under evaluation possesses some property or properties. And this activity, following Sibley, can be performed in purely descriptive, non-evaluative ways.

This kind of evaluation is no less a feature of art-making than it is of art appreciating. When making an artwork, an artist constantly makes choices which are informed by evaluation of the work up to its present state. These evaluations involve an assessment of merit – where this assessment is informed by property descriptions – of the properties possessed by the work in progress. There may be a difference in the degree to which or frequency with which an artist justifies her ongoing evaluations by appeal to the underwriting descriptions, but this is merely a contingent social fact. If we forced artists to work under the sociological microscope they would, like Jon, offer descriptions of the work in progress which justified their assignment of merit or de-merit and the corresponding decision that came with that assignment. The fact that they are more often pressed, *post facto*, to explain their evaluations does not imply that they made them in any other way.

6 Fractal Pattern Discrimination

In this section we describe our approach to implementing real-time fractal pattern discrimination on a simulated robot. The key property of a fractal object is that it is self-similar over a range of spatial scales. Three types of self-similarity are found in fractals. An object can be *exactly* self-similar at different scales, for example, Cantor dust (Figure 1), the Sierpinski carpet, Koch snowflake and other fractals which are generated by an iterated function system (which uses a geometric replacement rule). Objects can also display *approximate* or *quasi-*, rather than exact, self-similarity at different scales. These fractals contain distorted copies of the entire fractal at different scales. For example, fractals generated using an escape-time technique, such as the Mandelbrot and Julia sets, are quasi-self-similar. In the weakest form of self-similarity, statistical measures (such as 'fractal dimension') are preserved across scales. For example, fractals generated by processes such as diffusion-limited aggregation are *statistically* self-similar.

Only mathematical fractals display self-similarity across an infinite number of scales. Natural fractal objects display quasi- or statistical-self-similarity over a limited range of scales.

In contrast to Euclidean objects, fractals usually have non-integer dimensions. The fractal dimension measures the extent to which an object fills the Euclidean space in which it is embedded (Mandelbrot, 1982). A set of points along a line will have a fractal dimension between 0 and 1; a set of points on a plane have a fractal dimension between 1 and 2.

6.1 The Box-Counting Approach to Measuring Fractal Dimension

Box-counting is the simplest and most widely used technique for measuring fractal dimension and involves superimposing a series of regular grids over the data set. A regular grid consists of square boxes with a side length s. The measurement process is carried out using grids with a range of different side lengths. The first grid is layed over the set of data points and the number of occupied boxes, $N(s)$, counted. A box is occupied if it contains at least one data point. $N(s)$ is then plotted against $1/s$ for all box sizes. On a log-log graph, the slope of the graph is an estimate of fractal dimension. A fast $O(n \log n)$ algorithm (where n is the number of data points) was proposed by Liebovitch and Toth (1989) and implemented in C by Sarraille and DiFalco (Sarraille and Myers, 1994). This FD3 code is open source and available from: ftp://www.cs.csustan.edu/pub/fd3/. We have adapted FD3 to enable our robots to perform real time fractal pattern discrimination.

In order to confirm that a structure is fractal, it is necessary to show that it is self-similar over a reasonable number of scales. What constitutes a 'reasonable number' is a matter of some controversy. The range of scales is defined as: $log_{10}(Lmax/Lmin)$, where $Lmax$ is the largest or coarsest scale, and $Lmin$ the smallest or finest. In the physical sciences, the scale ranges tend to be small (Mandelbrot, 1998) and this can lead to incorrect estimations of fractal dimension or erroneously describing non-fractal structures as fractal ('apparent fractality') (Hamburger et al., 1996). Halley et al. (2004) recommend a scale range of greater than two orders of magnitude to avoid these problems, but this is not always possible.

It is important to note that although the box counting technique, described above, can employ a very large range of box sizes, the usable range is generally a lot smaller. For example, the largest box size used in FD3 is 2^{32} larger than the smallest box size (giving over 9 orders of magnitude scale range) but estimates from the smallest and the largest boxes have to be discarded, often resulting in a usable scale range of less than one order of magnitude. At very fine scales, none of the boxes contain more than one data point (depletion); at coarser scales all of the data points can be contained in one box (saturation). At these limits, the box counting algorithm will incorrectly estimate the fractal dimension. Consequently, fractal analysis is generally limited to a range of box sizes. The two largest box sizes are ignored. The smallest box size s used to estimate the fractal dimension meets the condition $N_B(s) \ll N/5$, where N is the number of data points and $N_B(s)$ the minimal number of boxes required to cover the data set at scale s (Liebovitch and Toth, 1989). FD3, which we use to measure fractal dimension, follows this convention. Our robots process small pixel arrays where the usable scale range is typically between 1.2 and 1.8.

Hamburger et al. (1996) investigated the fractal dimension of a number of small discs randomly scattered on the plane and demonstrated that this intrinsically *non* self-similar pattern can exhibit an almost linear relationship between $1/s$ and $N(s)$ over two orders of magnitude [4]. "Whether an apparent straight line on logarithmic axes really suggests a fractal or not is obviously a difficult and fundamental question" (Halley et al., 2004, p.259). Knowledge of the process that generated a pattern can sometimes help determine whether it is legitimate to describe it as fractal or not. However, it is important to note that using a fractal dimension measurement tool, such as Fd3, 'off the shelf' and without any consideration of the pattern that is being measured can lead to erroneously labelling a non-fractal object as fractal.

6.2 Lacunarity Analysis

A further issue, relevant to our project, is that two genuinely fractal objects can have the same fractal dimension and yet be very different in appearance, for example, Cantor dusts (Figure 1). One way that such patterns can be discriminated is in terms of their texture. Lacunarity is a useful measure of texture, introduced by Mandelbrot (1982), that quantifies the heterogeneity of the gaps in a pattern. Patterns that have gap sizes that are distributed over a greater range have a higher lacunarity index than patterns where the gap sizes are more similar. Objects that have a low lacunarity index are translationally-invariant because of their uniform gap sizes (Plotnick et al., 1993). Intuitively, one could shift sections of a pattern without altering its overall appearance. For a high lacunarity pattern,

[4] It is a matter of debate in the fractal literature whether patterns generated by random processes, for example, Brownian motion and self-avoiding random walks, should be considered apparent or actual fractals. We are clear that we want our robots to generate non-random self-similar mark patterns and we therefore want them to discriminate these patterns from those generated by random processes.

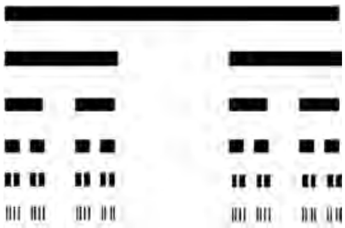

Figure 1: Cantor dust - an exactly self-similar fractal. Each of the lines has the same fractal dimension (0.6309) but a different texture, which is dependent on how many times the replacement rule (remove the central third of each line segment) has been applied. For illustrative purposes, the lines have been thickened.

shifting sections would become very apparent because the pattern is *not* translationally-invariant. It is important to note that lacunarity is a scale-dependent index - an object can have a homogenous texture at one scale but a heterogeneous texture at another scale. Further, one can measure the lacunarity of objects which are not self-similar.

There are a number of different algorithms for calculating the lacunarity of a pattern and there is not always agreement between the values that they generate (Halley et al., 2004). The most widely used technique is the gliding box algorithm (Allain and Cloitre, 1991). In this method a series of square boxes with varying side lengths s are placed over the data set and the number of points in each box (the box mass) is counted. The first box is placed in the top left hand corner and the box mass measured and then the box is systematically moved over the data set so that the position varies by one column or row. Unlike in the box counting method for measuring fractal dimension, in the gliding box algorithm the boxes overlap. For a square pattern with side M, there are $(M - s + 1)^2$ positions for a box of side length s where the box mass is measured. For each box size s, the mean and variance of the box mass is calculated. Lacunarity (Λ) is calculated as:

$$\Lambda = variance(s)/mean(s)^2 + 1.$$

When s is equal to 1, that is, a single pixel, or the grain of the data set, Λ is a function of the number of points in the data set and is independent of their spatial distribution. In this case, $\Lambda = 1/\%\ of\ on\ pixels$, where, in the case of a black and white image, an $on\ pixel$ is black. When a box is the size of the data set, the variance is 0 and so Λ is 1. In between the lower and upper bounds, lacunarity varies according to the range of gap sizes (or alternatively, clump sizes) at a given scale.

If an object is a *homogenous* fractal then the same scaling law applies at all positions of the object. Further, a log/log plot of lacunarity versus box side length generates a straight line where the slope is equal to the fractal dimension (D) minus the Euclidean embedding dimension (E); that is, D - 2 for all patterns on the plane (Allain and Cloitre, 1991). We use this relationship between fractal dimension and lacunarity to hard-wire a fractal discrimination mechanism in our robots.

6.3 Discriminating Random from Non-random Spatial Patterns

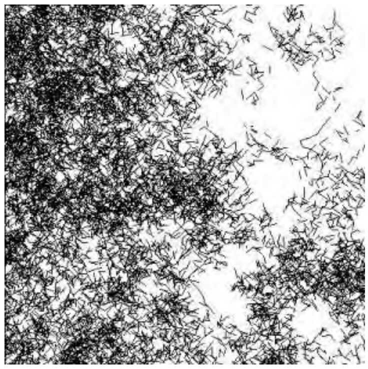

Figure 2: A randomly generated pattern of line segments which has the same % of black pixels (36%) as the two images in Figure 3. This pattern is discriminated from a self-similar fractal pattern, such as the left hand side image in Figure 3 using the method outlined in Section 6.3.

In this section we describe how we are using a box counting approach to get real-time estimates of fractal dimension and lacunarity of the spatial pattern in a simulated robot's visual field. Given the small number of points in the data set (limited by the small visual field of the robot) and the generally limited scale range, the technique we are employing could erroneously estimate self-similarity. At this preliminary stage, our goal is to develop a real-time method that enables our agents to discriminate random from non-random self-similar mark patterns. The class of non-random self-similar patterns that the robot can discriminate should have a structure such that:

1. the robot can produce members of this class with its pen;
2. and be sufficiently 'interesting' such that some members of this class are suitable for display (see 'aesthetic value worry' in Section 3).

At each sensory-motor update, the robot controller processes its 50mm x 50mm visual array in the following way:

1. using a box counting algorithm it estimates the fractal dimension (D) of the pixels over a range of scales where the ratio of the largest to the smallest box length side (s) is 2^{32}: 1 (the data points are re-scaled in order to achieve this range of box sizes);
2. using a box counting algorithm it estimates the lacunarity (Λ) at the same range of scales [5];
3. using regression analysis it measures the goodness of fit ($R2$) of the lacunarity curve over the range of scales which form the basis of the fractal dimension estimate;

[5]By using a box counting approach we estimate Λ on the basis of far fewer samples than is used in a gliding box algorithm. However, real time processing constraints forced this compromise in these preliminary experiments.

4. if $R2$ is greater than 0.95, it calculates the % error between the slope of the lacunarity curve and the estimate (D - E) of the lacunarity curve;

5. if $R2 > 0.95$ AND % error $< 20\%$, then the fractal pre-processing of the pixel array returns 1, otherwise 0.

This processing technique successfully discriminates a wide range of self-similar images from randomly generated images – not only images consisting of randomly distributed points but also randomly distributed lines segments, such as Figure 2, which are more likely to be generated by our robots than points.

6.4 The Fractal Discrimination Task

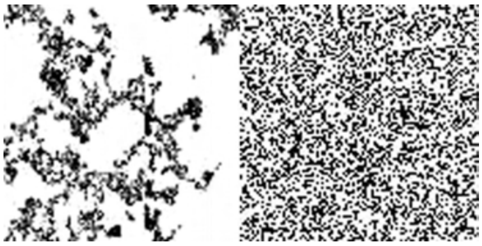

Figure 3: The left image is the texture on the target floor region in the evolutionary robotics experiment; the right image is the random point texture on the rest of the arena floor. Both patterns have the same percentage of black pixels (36%). The method outlined in Section 6.3 enables a simulated robot to discriminate between these two patterns.

The experiment described in this section was performed in a modified version of the Evorobot simulator (Nolfi, 2000). This software simulates a Khepera robot (Mondada et al., 1993) acting in user specified environments that can comprise of walls, large and small round objects and lights. Sensor readings taken from a physical Khepera robot are used to model the environment/sensor interactions.

The robots are controlled with neural networks based on Nolfi's (1997) emergent modularity architecture which has been successfully used to control complex robot behaviours, such as garbage collection. In the experiment reported here, each controller consists of 7 sensors (6 IR and 1 floor camera that is directly under the centre of the robot and has a visual field of 50mm x 50mm) and two pairs of motor units, controlling the right and left motor respectively. Each sensor connects to each motor neuron, giving 28 connections in the network. For details of the neural network update algorithm see Bird et al. (2007).

We used a genetic algorithm (GA) to evolve the biases of the 4 motor units and the connection weights between the sensor and motor neurons. The population size was 100 and the experiments were run for 600 generations. The initial population was randomly generated, each genotype consisting of the 32 neural network parameters encoded as an 8 bit integer-valued vector (range [0,255]). The mutation rate was 0.01 per allele and we did not use crossover. For more details see Bird et al. (2007) where the same GA was used to evolve robot mark-making behaviour.

The fitness function rewards proximity to the target area in the arena, with extra fitness if a robot is positioned on the target area at the end of a trial. The target area is 100mm x 100m and placed in one of two positions adjacent to the wall of a 400mm x 400mm arena. The target region consists of a fractal texture (Figure 3- left image); the rest of the floor has a random point pattern with the same overall ratio of black to white pixels as the target texture (36%)(Figure 3 - right image). If the robot crashes into an arena wall the trial is stopped and the fitness accumulated up to that point is averaged over the total number of time steps, thereby implicitly penalising robots that do not avoid obstacles.

Each genotype is instantiated as a robot controller and the robot is placed in a random position and orientation in the central area of the arena. Each individual is tested over 10 trials and the position of the target region is placed in two different positions, both adjacent to the wall of the arena. Every robot in the population was tested on the same series of initial positions and orientations each generation, and these changed every generation.

6.5 Preliminary Results

In early generations the robots do not move very far from their initial position or if they do they crash into walls. However, within 100 generations the majority of the population avoid obstacles and perform wall following. After 500 generations, the fittest individuals move in a straight line until they come close to a wall then follow the wall until they are over the target area and then stop. This is not a particularly surprising result as the patterns covering the target region and floor were chosen so that the fractal discrimination mechanism could clearly discriminate between them. However, it does demonstrate that this mechanism can be used to control the real-time behaviour of a robot and enable it to discriminate between random and fractal patterns on the floor. As in our previous experiments (Bird and Stokes, 2006b) the robots have evolved to use the arena walls (a constant and reliable feature of the environment): in the current experiment they provide a means of finding the target region which is always positioned adjacent to the edge of the arena.

7 Discussion

What does the preliminary theoretical analysis of evaluation tell us about our artificial agents and their pattern discrimination behaviour? By discriminating a fractal pattern from a random pattern are our agents evaluating their environments? Not obviously. But they are on their way. By discriminating fractal patterns, our agents are providing property descriptions, sparse though such descriptions may be, of their environment. When an agent stops on a fractal pattern and not on random patterns, it indicates the presence of a property – namely homogenous self-similarity – in the target location. This is just to say that this behaviour 'reports' that some part of the environment is a certain way, and other parts are not. The report is no richer than this, but neither is my report that this object is a square, while that one is not. Both kinds of reports are minimally descriptive of the world.

As we outlined in Section 5, evaluation comprises a descriptive element. Our agents are thus performing one element of evaluation: their pattern discrimination is a simple property description of their environment. However, what about the second, assessment of merit, component of evaluation: do our agents do this? This is a trickier issue. But here is a speculative suggestion. If our conception of evaluation is accurate, then it requires assigning merit to the properties of the object that have been discriminated. One assigns merit to the things that one likes or prefers, and de-merit to the things that one does not like or prefer. Preferences drive an agent, motivating it to act in whatever ways it does. Assignment of merit is thus to indicate, at base, what philosophers and psychologists call a 'conative attitude' [6]. This is true of human beings as well as laboratory rats; evaluation, no matter how rich, involves conation. Our agents, to evaluate in any interesting sense, thus need a preference or some degree of conation.

It would be misleading to describe the agents in our simulation experiments as possessing *individual* preferences. That is, our agents have not developed in their 'life spans' a preference or conative attitude of any kind. However, in this respect, do the agents differ from a laboratory rat, whose conative attitude of hunger motivates it to push the lever on the right and not the one on the left? The rat's preference is no more individually developed than is the fractal preference of our simulated agents. In the case of both the rat and the artificial agent, the preference has, in some sense, evolved in the population to which each belongs[7].

We might therefore say that fit robots, by successfully distinguishing fractals from non-fractals (and later in our research, by completing fractal patterns by mark-making) are thereby performing a number of simple evaluations which inform their respective behaviours. This is not rich evaluative behaviour, but if the rat is evaluating, then so is our robot. This, we suggest, is the basic route from mark-making and detection via pattern discrimination to pattern evaluation.

Acknowledgements

The *Computational Intelligence, Creativity and Cognition* project is funded by the AHRC and led by Paul Brown in collaboration with Phil Husbands, Margaret Boden and Charlie Gere.

[6]Conative attitudes are motivational states and are traditionally contrasted with both cognitive, or knowledge acquiring, states and states of affect. They comprise desires, values, preferences, likings, and so on. They are proactive and without them, agents do not act.

[7]In the case of the rat both the preference and the mechanism that underpins it evolved whereas in our robots the discrimination mechanism was hard-wired by us and the robot evolved the ability to use it appropriately. By prescribing what patterns it can discriminate, it might be argued that we have compromised the autonomy of the agent and consequently future fractal mark-making behaviour cannot be described as 'creative' (see Section 2), even though the fractal pattern discrimination might be minimally evaluative.

References

Allain, C. and Cloitre, M. (1991). Characterising the lacunarity of random and deterministic fractal sets. *Physical Review A*, 44:3552–3558.

Bird, J. and Stokes, D. (2006a). Evolving fractal drawings. In Soddu, C., editor, *Proceedings of the 9th Generative Art Conference*, pages 317–327.

Bird, J. and Stokes, D. (2006b). Evolving minimally creative robots. In Colton, S. and Pease, A., editors, *Proceedings of the Third Joint Workshop on Computational Creativity*, pages 1–5.

Bird, J., Stokes, D., Husbands, P., Brown, P., and Bigge, B. (2007). Towards autonomous artworks. *Leonardo Electronic Almanac*. In press.

Boden, M. A. (2004). *The Creative Mind*. Routledge, London, second edition.

Cohen, H. (1999). Colouring without seeing: A problem in machine creativity. *AISB Quarterly*, 102:26–35.

Halley, J. M., Hartley, S., Kallimanis, A. S., Kunin, W. E., Lennon, J. J., and Sgardelis, S. P. (2004). Uses and abuses of fractal methodology in ecology. *Ecology Letters*, 7:254–271.

Hamburger, D., Biham, O., and Avnir, D. (1996). Apparent fractality emerging from models of random distributions. *Physical Review E*, 53:3342–3358.

Liebovitch, L. S. and Toth, T. (1989). A fast algorithm to determine fractal dimensions by box counting. *Physics Letters A*, 141(8-9):386–390.

Mandelbrot, B. B. (1982). *The Fractal Geometry of Nature*. Freeman, New York.

Mandelbrot, B. B. (1998). Is nature fractal? *Science*, 279:783–784.

McCorduck, P. (1991). *AARON's Code: Meta-Art, Artificial Intelligence and the Work of Harold Cohen*. Freeman, New York.

Mondada, F., Franzi, E., and Ienne, P. (1993). Mobile robot miniaturisation: A tool for investigation in control algorithms. In *Proceedings of the Third International Symposium on Experimental Robotics*, pages 501 – 503, Berlin. Springer Verlag.

Nolfi, S. (2000). *Evorobot 1.1*. Institute of Cognitive Sciences and Technologies. http://gral.ip.rm.cnr.it/evorobot/simulator.html.

Plotnick, R. E., Gardner, R. H., and O'Neill, V. O. (1993). Lacunarity indices as measures of landscape texture. *Landscape Ecology*, 8(3):201–211.

Sarraille, J. J. and Myers, L. S. (1994). FD3: A program for measuring fractal dimension. *Educational and Psychological Measurement*, 54(1):94–97.

Sibley, F. (2001). Particularity, art and evaluation. In Benson, J., Redfern, B., and Cox, R., editors, *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics, Frank Sibley*.

Spehar, B., Clifford, C. W. G., Newell, B. R., and Taylor, R. P. (2003). Universal aesthetic of fractals. *Computers and Graphics*, 27:813–820.