# Implicit Fitness Functions for Evolving a Drawing Robot

Jon Bird, Phil Husbands, Martin Perris, Bill Bigge and Paul Brown

Centre for Computational Neuroscience and Robotics
University of Sussex, Brighton, UK

**Abstract.** We describe an approach to artificially evolving a drawing robot using *implicit* fitness functions, which are designed to minimise any direct reference to the line patterns made by the robot. We employ this approach to reduce the constraints we place on the robot's autonomy and increase its utility as a test bed for synthetically investigating creativity. We demonstrate the critical role of neural network architecture in the line patterns generated by the robot.

## 1 Introduction

The Drawbots project is a multidisciplinary investigation into creativity involving philosophers, adaptive systems researchers and an artist. A theoretical goal is to investigate the question: what is the simplest mechanism that can be described as creative? To this end we artificially evolve wheeled robots that move around an arena making pen marks on the floor. These 'embodied thought experiments' help clarify some of the necessary conditions for 'minimal creativity' (autonomy, novelty and evaluation) and how they can be embodied in a robot (see [1] for a detailed consideration of these issues). An artistic goal of the project is to generate aesthetically interesting line drawings that are suitable for exhibition. This is distinct from, and potentially at odds with, the theoretical goal. For example, by incorporating artistic knowledge into fitness functions we might enhance the aesthetic appeal of the resulting line markings, but at the expense of compromising the autonomy (and therefore 'minimal creativity') of the robots. This paper describes our use of implicit fitness functions to evolve a drawing robot where we minimise any direct reference to the line patterns and our focus is on investigating minimal creativity.

## 2 Evolutionary Robotics

The main synthetic, bottom-up methods used in the project are those of evolutionary robotics (ER). ER is a biologically inspired approach to the automatic design of autonomous robots. The field encompasses a wide range of work where one or more (sometimes all) of the following aspects of robot design are in the hands of an evolutionary search algorithm: the control system; the overall body morphology; and sensor and actuator properties. The evolutionary process uses a fitness measure based on how good a robot's behaviour is according to some evaluation criteria: a key distinction here is between *implicit* and *explicit* fitness functions [2]. An explicit fitness function rewards specific behavioural elements - such as travelling in a straight line or maximum velocity achieved - and hence shapes the overall behaviour from a set of predefined primitives. Implicit fitness functions operate at a more indirect, abstract level - reward is given for completing some task but the robot is free to achieve it in any possible way. The number of variables and constraints defined in a fitness function determine where it falls on the implicit-explicit dimension: implicit fitness functions have no or very few such components. Fitness is tested either in simulation, in the real world or using a combination of the two. Typically some form of artificial neural network acts as the nervous system of the robot; properties of the network will invariably be evolved even if other aspects of the robot design are not. By artificially evolving control architectures from suitably low level primitives, the final controller "need not be tightly restricted by human designers' prejudices" [3, p.83]: ER can therefore potentially generate novel models of creativity and art-making machines that are not necessarily constrained by the artist's (systems designer's) stylistic 'signature'.

## 3 Implicit Fitness Function Experiments

In this section we describe two sets of ER experiments that aimed to minimise our influence on the resulting robot behaviour by using *implicit* fitness functions that did not specify the types of marks that a robot should make. The first 'sensory-motor correlation' fitness function was tested in simulation; the second 'ecological' fitness function was initially tested in simulation but some of the resulting controllers have also been successfully transferred and tested on the Drawbot (Figure 3).

### 3.1 Sensory-Motor Correlation Fitness Function

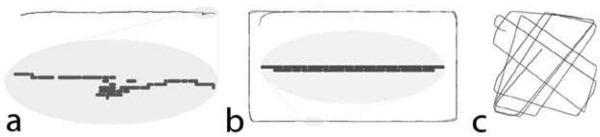


**Fig. 1.** Results from the implicit fitness function that rewarded correlated activity between the pen movement (up/down) and line detector (on/off). a) mid-fitness individual; b) high fitness individual; c) the patterns that result from adding a selection pressure to mark the entire arena.

Initial experiments were carried out in simulation using an accurate model of a Khepera robot, a standard ER platform, augmented with a drawing pen placed between its drive wheels. In the simulation, each robot controller was a neural network consisting of six motor neurons: two for each of the left wheel, right wheel and pen position (up or down) motors. At each time step in the simulation, the most strongly activated neuron of each motor pair controlled its associated actuator. The robot has seven sensors (six frontal IR sensors and one line detector positioned under the robot that can detect marks made by the pen). Each of the seven sensors was connected to each of the six motor neurons. A genetic algorithm was used to determine the strength of each of these connections and the bias of each of the motor neurons. The fitness function rewarded controllers that correlated the changes in state of their line detector and pen position. For example, if a line was detected and the robot's pen was then raised or lowered within a short time window, the robot accumulated fitness. This fitness function resulted in robots that used the arena walls (a constant feature of the environment) to guide their drawing behaviour. Mid-fitness individuals follow a wall to a corner and then gain fitness by repetitively moving forwards and backwards over a mark and appropriately co-ordinating the movement of their pen (Figure 1a). High fitness individuals initially follow the arena walls for one circuit making a continuous line and on their second circuit raise and lower their pen making marks adjacent to the initial line (Figure 1b). Different behaviours evolve when the fitness function also rewards robots for the extent to which they mark the whole area of the arena: the robots turn away from the walls at angles and mark the central parts of the arena as well (Figure 1c). In all these experiments crashing into walls is penalised by stopping the evaluation and thereby giving the robots less time to accumulate fitness.

### 3.2 Ecological Fitness Function

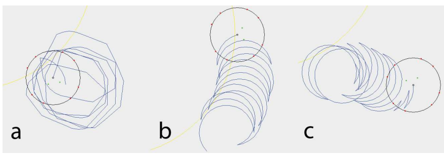


**Fig. 2.** Line patterns generated in simulation using an ecological fitness function and a simple motor model. a) is the typical pattern generated by a 20 neuron network; b) is an 'orange segment' pattern occasionally (approximately 30%) generated by a 20 neuron network; c) is the typical pattern generated by a 40 neuron network (which after further evolution looks like Figure 3). The robot is the circle with the dot at its centre.

The controllers evolved in the experiments briefly described in this section were Continuous Time Recurrent Neural Networks (CTRNNs), a rather more complex network than those used in the earlier experiments described above.

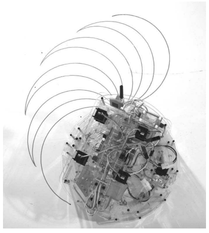

**Fig. 3.** Top-down view of an 'orange segment' line pattern generated by a Drawbot in the real world which was evolved in simulation using the implicit ecological fitness function, a 40 neuron network and a simple motor model.

The networks consisted of either 40 or 20 fully connected nodes. The connection weights, time constants, biases and gains were encoded as a string of real numbers in the range [0,1] and linearly scaled to values in the range [-5,5], [0.04,4], [-10, 10] and [0.01, 10.01] respectively. The state of each neuron was initially set to 0 plus a small random value. 6 of the neurons had external inputs from the sensors and 3 neurons acted as motor outputs: one for each wheel and one to lower and raise the pen. For full details see [4].

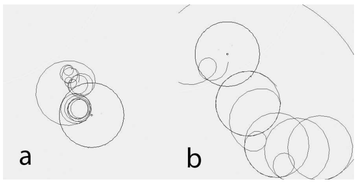


**Fig. 4.** Circle patterns generated in simulation using a 40 neuron network and a more sophisticated motor model with inertia and momentum. The robot is the circle with the dot at its centre.

Robot controllers were initially evolved in simulation using an 'ecological' fitness function. Small circular pieces of 'food' were randomly scattered in a target area of the arena (either a central rectangle or a semi-circle adjacent to a wall). Fitness was gained when a line drawn by the pen intersected one of the food particles. Each robot started with a fixed amount of energy which was used up at a constant rate while the pen was down but not while it was up; the robot could move and 'draw' freely for a fixed time period (1 minute) or until its energy ran out, whichever was sooner. The robot started in a random position and fitness was the lowest score achieved in a set of test trials.

In the initial experiment the most fit robots all displayed similar behaviour: they made sweeping curves ('orange segments') which alternated in direction and fanned out over a reasonable area of the target area (Figures 2c and 3). In the patterns generated by the fittest individuals, the separation of the segments

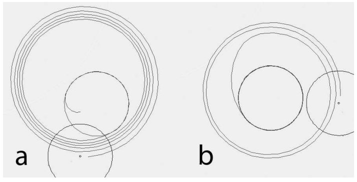


**Fig. 5.** Line patterns generated in simulation with a 20 neuron network and a more sophisticated motor model with inertia and momentum. The robot is the circle with the dot at its centre.

is just larger than the diameter of a food particle. This is a good strategy for systematic coverage of an area without crossing a food particle more than once (an individual can only score one point per food particle, regardless of the number of times its lines intersect it). The image produced by the real robot (Figure 3) is qualitatively very similar to those found in the simulation but the semi-circular curves are closer together and the robot tends to draw a full circle at the start. We halved the number of neurons in the CTRNNs from 40 to 20. When driven by the simple motor model the 20 neuron controllers tended to produce looping patterns (Figure 2a) and occasionally overlapping 'orange segments' (2b) - the pattern always generated by a 40 neuron network (2c - with further evolution the segments stop overlapping and look like 3). Although the 40 neuron controllers transferred well, the simulation did not model inertia or momentum and the robots were restricted to high speeds. In order to overcome these limitations a further series of experiments were carried out with a more sophisticated motor model. The change in motor model resulted in 40 neuron controllers generating circular patterns of varying diameter (Figure 4) and 20 neuron controllers generating spirals (Figure 5) - an effective solution for covering an area and minimising multiple intersections of the same food particle if the gap between the spirals is larger than the food particle diameter.

In all the above experiments the target area was located in the centre of the arena and although the controllers use the light as an energy source (they stop working if the light is switched off) they did not use if for directing their movement. We therefore conducted an experiment where the location of the target area varied in each trial and was always placed adjacent to a wall so that robots had to actively use their IR sensors to avoid crashing. A light was placed above the wall to indicate the centre of the semi-circular target region. The fitness of a robot was determined by its ability to perform phototaxis as well as the number of food particles it drew over. Crashing was again penalised. We found that a more distributed architecture facilitated the evolution of successful controllers in this task. The pen neuron was only connected to the light sensors and two other neurons and the threshold above which the pen was lowered was also evolved. Successful individuals make the majority of their marks in the target region, regardless of their starting position and orientation in the arena. The looping line patterns are less structured than the circles, spirals and orange segments prodcued in previous experiments (Figure 6), again illustrating that the robot's embodiment (change in network architecture), as well as the environment, influence the line patterns generated.

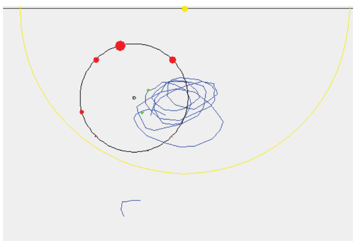

**Fig. 6.** Line patterns generated by a robot that had to perform obstacle avoidance and phototaxis in order to find the target regions where its lines would gain the maximum fitness. The top of the image is the arena wall, the dot on this wall represents the light and the semicircular area is the target region. Note that the robot marks a small curved line on the way to the light. The robot is the circle with the dot at its centre.

## 4 Conclusion

When investigating minimal creativity, our working hypothesis is that is advantageous to use an implicit fitness function in order to maximise the robot's autonomy. If we want to exhibit work produced by the robots, then a more explicit fitness function that embodies artistic knowledge about 'aesthetically pleasing' line patterns seems worth exploring and this is the focus of current research. However, even the patterns generated by implicit fitness functions can have an artistic impact, especially if the drawing process underlying the drawings is made evident, for example, by exhibiting the robots behaving in an arena rather than displaying the resulting drawings on a wall.

This paper forms part of the research supported by AHRC grant number B/RG/AN8285/APN19307.

## References

1. Bird, J. and Stokes, D.: Minimal Creativity, Evaluation and Pattern Discrimination. In A. Cardoso and G. Wiggins (Editors). Proceedings of the 4th International Joint Conference on Computational Creativity. (2007) 121–128.
2. Nolfi, S. and Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press/Bradford Books. Cambridge, MA. (2000)
3. Cliff, D., Harvey, I. and Husbands, P.: Explorations in Evolutionary Robotics. Adaptive Behavior, 2 (1993) 73–110.
4. Perris, M. Evolving Ecologically Inspired Drawing Behaviours. MSc dissertation, Department of Informatics, University of Sussex. (2007)