

Surviving on a sand-pile: *an investigation as to the type of behaviour that evolves in the presence of crises*

Bruce Edmonds

Centre for Policy Modelling

Manchester Metropolitan University

Aytoun Building, Aytoun Street, Manchester, M1 3GH, UK.

<http://www.bruce.edmonds.name>

Abstract

A simulation of artificial creatures is used to investigate the kinds of behaviour that might evolve under different environmental conditions. In particular, whether one might expect to find that evolution has resulted in behaviours that seem “well honed” with respect to the niche. The hypothesis is that this will only occur when the environment is relatively stable, that when the environment is swept with unpredictable crises there will be a greater variety of genes and behaviours. To investigate this a 3D artificial world with creatures evolving on a sand pile was created. The creatures move, eat, turn, mate and propagate depending upon tree-structured program genes which are interpreted in each situation to produce the behaviour. These genes are type-sensitively tree-crossed when mating occurs to make the genome of the offspring. The environment is tuneable so that different sizes and frequencies of unpredictable avalanches occur. Initial results seem to support the hypothesis. If it is true then this has implications for evolutionary explanations of behaviours, since it may not be the particular behaviour that is important but the range of behaviours that is significant. This has implications for the designers of animats and robots.

1. Introduction

(der Boer 1999) points out that biological evolution is characterised by the demise of organisms rather than by the success of the fittest. He calls this the *non-survival of the non-fit* in contrast to the common characterisation of evolution as the *survival of the fittest*. This is more than merely a play on words because of the grounding of “fitness” in the context of biological evolution. If ideas

such as the *survival of the fittest* are to be meaningful the fitness of an genome can not be simply defined as a label for those genes that survive, because this would reduce the idea to a trivial analytic truth. Rather it has to be given a meaning in terms of observations of the organism, for example its success in mating, avoiding predators, or finding food. To avoid confusion with other interpretations of “fitness” I will call this “behavioural fitness”. The theory of evolution can then be summarised by the two principles: firstly that the *behavioural fitness of the phenotype is a substantial cause of the survival and propagation of the genotype* and secondly that the *genotype substantially causes the phenotype*. The argument is to the extent to which, in the long term, the existence of a particular genotype is better predicted by the extent to which organisms *optimise* their behaviour in normal times (in terms of food collection, attractiveness etc.) or by its ability to *survive* periods of extreme crisis.

In evolutionary computation it is common to provide a single, fixed fitness function which determines the survival and propagation of the genetic information into future generations. Thus in a classic instrumental Genetic Algorithm we have Holland's *schema theorem* (Holland 1975) which states that a gene is spread into the next population in proportional to its fitness (a defined by the given fitness function). Here we *do* have the survival of the fittest but this in the protected world of an algorithm where the only possible threat is the competition of out-performing peers. This is very different from the circumstances that most organisms find themselves in, namely an environment riddled with: fundamentally unpredictability; catastrophic events; and complexity beyond their ability to comprehend. The feedback of the environment to organisms is complex, multidimensional and context-dependent.

This difference is nicely illustrated by contrasting two approaches to providing feedback to firms making tractors.

The story¹ is that in the Soviet Union firms making tractors were given a fixed fitness function: during a certain five-year plan they were rewarded according to the total weight of their output in tractors. In the west tractor firms were only rewarded when a customer was willing to buy a tractor, which might be due to any one of a thousand reasons. The result was that whereas the tractor firms in the west generally learnt to produce tractors that met their customer's needs (even if this was merely prestige), the Soviet firms produced immensely heavy tractors that quickly broke down. Regardless of whether this story is true, it vividly illustrates the limitations of a fixed fitness function.

It is the thesis of this paper that, in realistic circumstances, optimisation is only successful in the short-term. Where "in the short term" means *in the absence of major catastrophes*. The reason for this is due to the fact that optimisation (even where it is feasible to attempt) is inimical to the maintenance of variety. In other words, in the long run survival rather than optimisation dominates and in an unpredictable world swept with critical changes it is the heterogeneity of phenotype (and hence niche) that is the best way of ensuring survival. (der Boer 1999) hypothesises that it is the variety of habitat that is the main factor correlated with surviving crises. This is a special case of variety in general, since crises may not be geographically local but local to a niche across locations (e.g. a virus spread due to a particular behavioural trait).

A classic illustration of the efficacy of variety over optimisation is the pervasiveness of sexual reproduction (e.g. Jaenike 1978, Getz 2001). Sexual reproduction does not result in near clones of a successful organism but is a mechanism for propagating genes whilst creating and preserving the maximum variety of combinations of genes. Only for the simplest organisms, those exploiting immediate surroundings (as in plant propagation) or the most stable of environments (as with certain subterranean organisms) is sex not present.

In this paper I exhibit a model of simple creatures that move, eat, reproduce and die. They do all these things in an environment which has a (parameterised) tendency to avalanches, namely a sand-pile. The sand-pile model follows (Bak 1997). Grains of sand and food are randomly added to the pile causing unpredictable avalanches of different sizes. The avalanches can kill creatures in their path (by burying them until they starve or are dissipated) but also may uncover new food and living buried creatures. The question of interest here is the extent to which the frequency of catastrophes influence the behaviour that is evolved by the creatures. The hypothesis is that if there during long periods of stability (absence of big avalanches)

optimising strategies might do all right, but in the long-term (or equivalently during frequent and severe avalanches) the best strategies are not optimising ones but ones which promote variety.

2. Model Set-up

The central idea of the model is to capture the evolution of (genetically determined) behavioural strategies in an environment which has a tuneable level of crises. The environment I have chosen is a version of a sand-pile which is prone to unpredictable avalanches. The creatures exist on (and under) this. Their behaviour is controlled by programs which are crossed in the style of genetic programming when they mate (but only then). The kinds of behaviours that evolve are then compared to the number and severity of avalanches.

The creatures live on a torus, i.e. a square 2D grid with opposite edges joined. More than one creature can inhabit the same cell (by sitting one on top of the other). Each cell contains a column of particles upon which the creature in the cell stand upon (or be buried within). Only the top-most creature can move to another cell. Particles are continually added to the pile. Some of the particles are food particles which the creatures can eat while others are inert.

Energy in the form of food particles are being constantly added to the system: 70% of 200 particles per time period each with a food value of 15, making a total of 2100 units per unit of time (of course some of this is either dissipated or buried before it can be eaten). New creatures have 20 units of energy (which are obtained at the cost of the parent(s) except the initial population). When a creature eats its energy level is increased by 15. Each time period the creatures energy level is decreased by 0.2. Trying to eat, move, mate and propagate costs it: 0.1, 0.5, 0.5 and 1 units of energy respectively. If a creature's energy reaches zero it dies and becomes an inert particle. Thus in terms of total energy mating and propagating costs the same as in mating the costs are shared by the parents. If creatures happen to be good at eating and moving (so they are not buried for too long) and they do not mate or propagate they can survive indefinitely.

If the column of particles becomes taller than a certain critical level the particles topple over into adjacent cells. If these toppling particles fall upon a creature they can bury and hence trap it (in which case it might starve and die).

When columns topple some of it is dissipated but the rest randomly falls into adjacent cells. These might extend these columns beyond the critical level causing them to topple in subsequent time periods. These sequences of topplings form avalanches. It is known that the size of these avalanches form a power law distribution, that is all sizes of avalanches occur, but the larger the avalanche the more rarely it occurs. However the occurrence and size of

¹ The story is told to me by a colleague who referred me to (Zaslow 1989), however I have not yet been able to check its veracity.

avalanches are essentially unpredictable. To stop the particles piling every higher, a certain proportion of the particles are dissipated when the columns topple. The rate at which particles are dropped into the system, the critical toppling level and the dissipation rate determine the distribution of avalanches. The dynamics of sand piles and other systems with similar properties are extensively discussed in (Bak 1997). By changing the dissipation rate and the rate of new particles we can make the system more or less catastrophe ridden.

The creatures' behaviour is determined by a series of six programs, one for each of the potential actions: eating, mating, propagating, moving forward, turning right and turning left (turning right and left simultaneously is equivalent to turning around to face in the opposite direction). The form of these programs is a strongly typed-tree which, when interpreted, is a function that outputs a Boolean, following (Montana 1995). These trees are interpreted each time for each creature and each kind of action. If the result is "true" the creature attempts to perform that action. These attempted actions are only brought about if they are possible, e.g. one can only actually eat if the creature tries to eat and there is food immediately below, one can only mate if the creature tries to mate, it has enough energy and there is another creature immediately below it who also tries to mate and has enough energy.

The leaves of these trees are either constants or inputs whose value, when interpreted, depends upon the state of the creature and its immediate environment. These inputs represent what the creature can perceive. The creatures can perceive: the relative height of adjacent cells (higher, lower or level) in the different directions (relative to the way it is facing), as well as the presence of food, other creatures and recently fallen objects there. They perceive the time since they had last: eaten, mated, propagated or moved. They also perceive their own energy level.

The operators of these programs include Boolean operators (AND, OR, NOT and IF THEN ELSE), arithmetic operators (+, *, -, /), comparisons (>, <, =), directional operators (to the right of, to the left of, behind, ahead), and a simple memory operator (LAST). Any combination of leaves and operators that preserve the appropriate types in the trees can occur. This allows for a huge variety of possible behavioural strategies to be encoded in the creature's genome. Some examples of trees are shown in figure 1 and the full syntax of these trees in listed in the appendix.

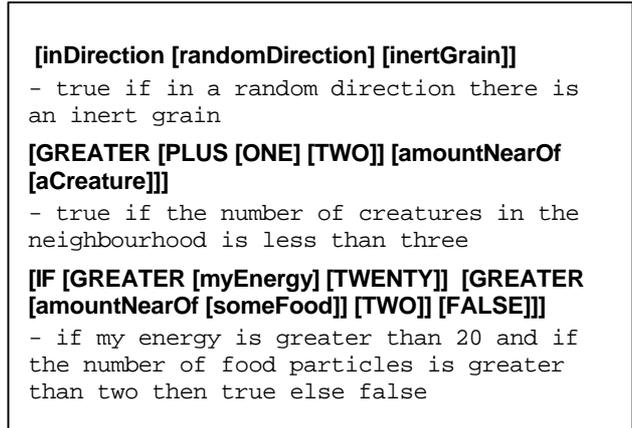


Figure 1. Some examples of program trees and some interpretations of them in terms of behaviour

A warning: although I have used comprehensible mnemonics to specify the structure of their behavioural genes this is not accessible to the creatures. I *only* did this to make it easier for me, the modeller, to understand the behaviours they have. Thus the creatures have no explicit knowledge about their environment apart from the fact of their perceptions and what happens to them. They do not have any hard-wired indicators (such as pleasure, pain hunger, lust etc.) to aid their choice of action. In particular they do not know *a priori* that, for example, that if there is food here then it is a good idea to eat or even that it is a good idea (from the point of view of survival) to eat at all. The creatures merely "blindly" execute their program genes and attempt to perform the action these entail, regardless of whether they are sensible or even possible. Thus a creature might continually try to propagate when it has insufficient energy to do so and so greatly hasten its own demise since even attempting to propagate expends a lot of energy- of course, over the course of the simulation one would expect that such creatures would quickly die out leaving only those that happen to be more suited to their environment.

Channon and Damper (Channon and Damper 2000) argue that using a GP like structure will prevent evolution occurring since this makes the fitness landscape too rugged to allow for continual and open-ended evolution. He uses a variable length GA which is mapped into neural nets using a system of gene-expression. However this is dependent upon the language of the trees (Albuquerque et al. 2000). In this case the use of strongly typed trees means that crosses are much more likely to be correlated in terms of fitness that with an untyped GP. Further the nodes and terminals were chosen so that crosses (that occur mostly near the leaves, Angeline 1996) will be related to what was there before. Thus, although I find the arguments of Channon and Damper are in general persuasive, the extent of the ruggedness in this case is unclear. It is also not entirely certain that the current model allows for truly open-ended evolution, but does *seem* to be the case. A

further study involving many long simulations would be necessary to determine this for sure. As it was the simulations took about a day for 2000 creatures over 800 time cycles with an average tree depth of 4, so considerably greater computational resources would be required for this.

If the creatures have enough energy they may propagate. In this case the genome of the offspring is a copy of the parent. However there is a small probability of a mutation (0.01 per tree copy). The mutation is realised by crossing the original tree with a new randomly generated tree. If they have a high enough energy level and meet another creature also with a enough energy they may mate with that creature. Thus mating is inherently much more difficult to achieve that propagating, even though *in total for the system as a whole* the energy requirements are the same.

At the beginning the initial population is initialised with random programs. When the creatures clone themselves their offspring are identical to themselves, when the creatures mate a type-sensitive tree-crossover occurs as in strongly typed Genetic Programming (Montana 1995) pairwise for each of the six trees and the offspring is created as a result.

The model was implemented in SDML (Moss et al. 1996). More details about the model, including how to obtain the code, are in the Appendix of this paper.

3. Results

At the present only preliminary results are available. Here I contrast three runs where the dissipation parameter is: 0.4, 0.2 and 0.05. The higher the dissipation rate the fewer severe avalanches there are because one column toppling is less likely to cause other to topple. All three were run for 800 time periods and had the following parameter settings:

- Initial population: 2000
- Energy of a food particle: 15
- Energy a creature is born with: 20
- Initial depth of genome trees: 4
- Size of grid: 40 by 40
- Proportion of new particles that were food: 70%
- Critical height of a column: 10
- Number of new grains each time period: 200
- Mutation rate: 1%
- Energy required just to exist: 0.2
- Energy required to attempt to eat: 0.1
- Energy required to attempt to move: 0.2
- Energy required to attempt to mate: 0.5 (each parent)
- Energy required to attempt to propagate: 1

In a typical run, the population soon crashes to about 30% of the initial population as the totally dysfunctional creatures (e.g. those that don't ever try to eat) die. This leaves about 200 who happen to have a rudimentary ability to survive. Out of these some will simply be good at surviving. Others will happen to propagate or even (if they

are very lucky) mate, others. This last category of creature forms the basis for the evolutionary process, which (typically) allows the population level to climb to a level of 800-1200 where it levels off.

The higher the dissipation rate, the fewer the avalanches and the easier it is to survive. Figure 2 shows the number of living creatures for each run and figure 3 shows the number of falling items (using the same key as figure 2). In each case the population evolves to survive better over time, but the creatures in the most avalanche prone have a more difficult task to evolve solutions to.

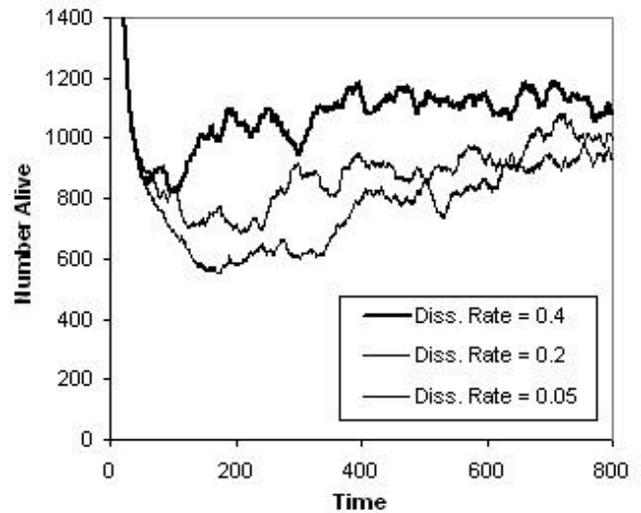


Figure 2. The population over time in the three runs.

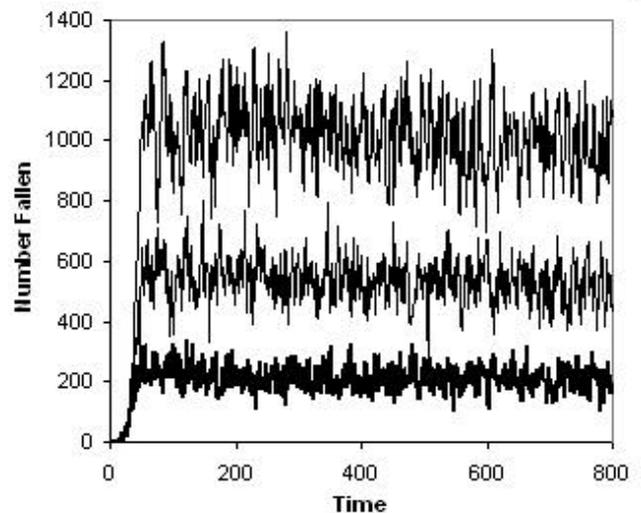


Figure 3. The number of falling items (inert grains, food and creatures) over time in the three runs.

In each case there was vastly more propagation occurring than mating. This is not surprising as for mating to occur two creatures have to be on top of each other, both with enough energy and both wishing to mate, while for propagation only one creature with enough energy and

wishing to propagate is required. Contrary to expectations the more stable runs had a higher number of matings that the unstable ones (figure 4). However this is due to the much greater difficulty of mating in the unstable runs, since many of the creatures will be buried for periods of time before another avalanche brings them back to the top.

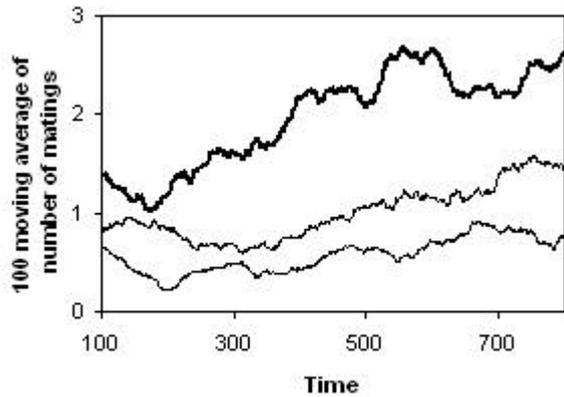


Figure 4. Number of matings in three runs (key as before)

In each of the three runs considerable variety was found to exist, even by time period 800. This may be an indication that the runs need to be a *lot* longer, it may also simply be a reflection of its environment. Table 1 shows the distribution of occurrences of frequencies of genes at the end of each run. Thus in the run with dissipation rate 0.4 there were an average of 1121 unique genes over the six gene types (one for each action), an average of 157.17 of genes that occurred twice in the population, 62.17 that occurred thrice etc.

Frequency of gene	Diss. Rate 0.4	Diss. Rate 0.2	Diss. Rate 0.05
1	1121.00	957.00	1014.00
2	157.17	116.83	111.50
3	62.17	46.17	41.50
4	42.17	32.17	33.00
5	36.83	29.50	30.33
6	31.00	27.00	26.17
7	29.00	26.00	24.17
8	25.50	26.00	19.50
9	25.50	26.00	19.50
10	22.50	21.50	19.50
11	19.17	19.83	17.83
12	17.33	19.83	17.83
13	15.33	19.83	13.83
14	13.17	13.33	9.50
15	10.83	8.67	9.50
16	8.33	8.67	9.50
17	5.67	3.33	6.83
18	0.00	3.33	6.83
19	0.00	3.33	3.83
20	0.00	3.33	3.83
21	0.00	0.00	3.83

22	0.00	0.00	3.83
23	0.00	0.00	3.83

Table 1. the occurrence of different frequencies of genes at the end of the runs, averaged over the six genes.

Some evidence for the effect of the level of catastrophe was found in the kurtosis of the (average) distribution of genes at the ends of the three runs (table 2). The Kurtosis is the fourth moment (mean is the first, variance is the second, and skewness is the third). The index of kurtosis is the fourth root of the kurtosis (just as the standard deviation is the second root of the variance). This is a measure of how the “fat-tailed” the distribution is. A number greater than one indicates that the distribution has higher tails and a sharper “point” than the normal distribution. All three distributions had very high levels of kurtosis, but the lower the dissipation rate the higher the distribution, indicating that in the crisis prone run there was more of both a few frequent genes and a lot of unique genes, while the most “stable” run had more genes with an intermediate frequency (it was closer to a normal distribution).

Dissipation Rate	Kurtosis index of gene distribution
0.40	18.57
0.20	20.03
0.05	22.91

Table 2. The Kurtosis index for the spread of genes in time period 800

Thus some evidence was found for the hypothesis but it is not very strong. Clearly many more runs of this model need to be made before its properties can be definitively characterised and then great caution must be exercised in projecting any conclusions upon biological evolution. However it may well be useful as a guide for animat developers and designers.

4. Discussion

It is almost inevitable that the characteristics of strategies that are successful in the short and long term will be different, but this is particularly true in the presence of unpredictable crises. Optimisation may or may not be possible between unpredictable crises but is utterly impossible *over* them. In the second case any mechanism that encourages variety (in habitat, genome, phenotype etc.) might ensure that a gene survives a crisis and thus helps ensure its longer-term existence.

One consequence of this is in the kind of behaviour one might expect to be evolved in different circumstances. In a relatively stable niche one might expect well-honed behaviour that “fits” its environment, but in situations of rapid and catastrophic change one might expect a variety of

behaviours or behaviours that allow a variety of niches or locations to be inhabited so that at least some of the organisms may survive new and unpredictable events.

This limits when an evolutionary explanation of a particular behaviour is feasible. For in situations prone to crises it may not be the detail of the behaviour that is significant but the range of behaviours of which the particular example may be but one.

This also provides a cautionary tale for those who would *design* or *develop* animats or robots. Endowing such a creation with a particular well-honed behaviour (whether obtained by design, from observing animals or from the results of simulations), may not provide it with the best chance of survival/success (unless the environment it is to inhabit is very limited and/or stable). Rather, it may sometimes be better to provide mechanisms to produce and maintain a great variety of behaviours, the better to ensure that some survive/succeed during an unpredictable crisis. Sexual reproduction is one such mechanism, mechanisms that produce essentially random elements to behaviours is another.

Another aspect of the thesis of this paper is that it is impossible for an organism (or species) to learn from its own extinction. Successful organisms often develop “proxies” for extinction that allow it to adapt to avoid situations where extinction is likely. Thus we have pleasure, pain, lust, hunger etc. in animals and profit, cash-flow, dividends, etc. in firms. These can indicate danger (or opportunity) ahead of (or in the absence of) a critically dangerous event. However these are only proxies for (or models of) the real thing - survival. In general only a complex collection of these indicators make up a sufficiently good model of survival to be useful in real-life situations. If an organism concentrates over-much on optimising a single indicator (for example just eliminating pain in animals or only going after short-term profit in firms) they are less likely to survive in the long term. Thus over an evolutionary process one might expect that: *firstly*, such indicators would evolve and *secondly*, that over time they would increase in number and sophistication.

Appendix – Model Specification

Design Sources

There have been many models populated with artificial creatures roaming a 2D world. Perhaps the closest to this one was made by Channon (Channon and Damper 2000). The strongly-typed Genetic Programming (Koza 1992, 1994) structure of the genome is taken from (Montanan 1995) and is similar to those in other of my models (e.g. Edmonds 1999). The structure of the sand-pile is taken loosely from (Bak 1997).

Static Structure

There is a square 2D grid of cells joined into a torus. Each cell has a list of items representing the pile of objects at that cell. Each creature has 6 program trees which, when interpreted outputs a Boolean for each behaviour.

Dynamic Structure

There are three things that change as the simulation progresses:

- The contents of each cell representing the pile of grains, food and creatures are updated to reflect new grains, fallen items, creatures moving, being born and dying.
- The energy level of each creature varies according to whether it has eaten, mated or propagated.
- The trees which determine the creatures' behaviour when interpreted in each circumstance are crossed in the manner of Genetic Programming when the creatures mate and produce a child.

Important Parameters and Settings

The model had the following parameters and settings (defaults in brackets):

- Grid Size (40)
- Initial number of creatures (4000)
- The energy creatures are created with (20)
- The amount of energy gained by a creature when eating a food particle (15)
- The number of particles that fall into the environment each time period (200)
- The proportion of new particles that are food (rather than inert) (70%)
- The depth of the initial random program trees of the creatures (4)
- The critical depth for a pile of objects at a location above which the pile topples (10)
- The amount of energy lost in a time period if the creature does not eat (0.2)
- The energy cost of trying to move (0.2)
- The energy cost of trying to mate (0.5)
- The energy cost of trying to propagate (1)
- The energy cost of trying to eat (0.1)
- The energy reserves necessary for a creature to propagate itself (> energy creatures are created with)
- The amount of energy lost in the process of propagation (= 1 x energy creatures are created with)
- The energy reserves necessary for creatures to mate (> 1 x energy creatures are created with for each mate)

- The amount of energy lost in the process of mating (= 0.5 x energy creatures are created with for each mate)
- The mutation rate (1%)

Structure of the Internal Genome

Each creature has six trees that are interpreted in each situation for each creature and each possible type of behaviour to determine its intentions. The behaviours are whether to: eat, move forward, turn right, turn left, mate and propagate. The trees are strongly-typed trees following (Montana 1995), that is each node and terminal has a specified output type and type of its inputs (if any). During propagation the genome is copied to the offspring with a 5% chance of a mutation occurring to each tree (the mutation is a random tree-crossing with a new randomly generated tree). During mating the parents trees are pairwise crossed to form the offspring's trees.

The syntax of the trees are the same, as follows:

Types:

- numeric
- Boolean
- direction
- action
- property

Terminals – type: list of labels

- **property:** aCreature, someFood, inertGrain, aFallenObject, downHill, upHill, level;
- **action:** eaten, moved, propagated, mated;
- **Boolean:** TRUE, FALSE, randomBoolean;
- **numeric:** ZERO, ONE, TWO, FIVE, TEN, TWENTY, FIFTY, myEnergy;
- **direction:** ahead, behind, left, right, here, randomDirection.

Nodes – type: the a list of: label [list of argument types]

- **Boolean:** AND [Boolean Boolean], OR [Boolean Boolean], NOT [Boolean], IFBoolean [Boolean Boolean Boolean], LASTBoolean [Boolean], LESS [numeric numeric], GREATER [numeric numeric], EQUAL [numeric numeric], inDirection [direction property];
- **numeric:** PLUS [numeric numeric], MINUS [numeric numeric], TIMES [numeric numeric], DIVIDE [numeric numeric], IFNumeric [Boolean numeric numeric], LASTNumeric [numeric], randomNumeric [numeric], amountNearOf [property], timeSince [action];
- **direction:** rotateRight [direction], rotateLeft [direction], reflected [direction], IFDirection [Boolean direction direction], LASTDirection [direction].

Initialisation

At the start of the simulation an initial population of creatures is randomly distributed about the cells, cells with no creatures at the start have three items of food. All the initial creatures are provided with random behaviour trees with a depth of 4.

Dynamics and algorithms

Each time period the following occurs:

1. The states of the cells are updated to reflect actions at the end of the last time period: new creatures (from matings or propagation), creature movements, the falling of grains and creatures, dissipation and grains input into the system;
2. The position of creatures and their energy levels are updated according to the actions last time period;
3. The creatures behavioural models are interpreted in their current situations which result in the attempted actions of the creatures;
4. The actual actions of the creatures are worked out where intended actions are possible.

Results claimed as significant

That the sort of behaviours evolved in the model depends upon the level of number and severity of avalanches.

Intended interpretation

The result suggest that optimisation may be only approximated to during periods of stability.

Other details considered unimportant for the results but which were necessary for the implementation

It is not thought that the exact topology of the world is critical to the results but it may be that having edges might add to the variety of the world and promote diversity of genotype. The toppling of the columns of objects is somewhat abstract. In particular, the avalanches are too slow, progressing only at the rate a creature can move. More realistic rules to reflect this are possible and this might lead to different specific behaviours being evolved but it is not thought that this would lead to different conclusions as far as this paper is concerned. The tree-structured genotype is not as realistic as a GA combined with a gene expression mechanism (as in Channon), it may be that the current genotype makes the fitness landscape more rugged and hence evolution more difficult.

Language and System Environment

The simulation was implemented in SDML 4.1. A description of SDML can be found in (Moss et al. 1996).

Information about SDML can be found, including downloading is at <http://www.cpm.mmu.ac.uk/sdml>.

Source code

The source code (as an SDML module) can be downloaded from: www.bruce.edmonds.name/soasp

Example Output

Some parameter settings with example output can be found at: www.bruce.edmonds.name/soasp

References

- Albuquerque, P., Chopard, B., Mazza, C. and Tomassini, M., (2000) On the Impact of the Representation on Fitness Landscapes. in *Genetic Programming, Proceedings of EuroGP'2000*, (Edinburgh, 15-16 April 2000), Springer-Verlag. EvoNet, LNCS, **1802**. 1-15.
- Angeline, P.J., (1996) An Investigation into the Sensitivity of Genetic Programming to the Frequency of Leaf Selection During Subtree Crossover. in *Genetic Programming 1996: Proceedings of the First Annual Conference*, (Stanford University, CA, USA, 28-31 July 1996), MIT Press, 21-29.
- Bak, P. (1997) *How Nature Works: The Science of Self Organized Criticality*. Oxford University Press, Oxford.
- Channon, A.D. and Damper, R.I. (2000) Towards the evolutionary emergence of increasingly complex advantageous behaviours. *International Journal of Systems Science*, **31** (7). 843-860.
- den Boer, P.J. (1999) Natural selection or the non-survival of the non-fit. *Acta Biotheoretica*, **47**. 83-97.
- Edmonds, B. (1999) Gossip, Sexual Recombination and the El Farol Bar: modelling the emergence of heterogeneity. *Journal of Artificial Societies and Social Simulation*, **2** (3). <http://www.soc.surrey.ac.uk/JASSS/2/3/2.html>
- Getz, W. M. 2001: Competition, extinction, and the sexuality of species. — *Annales Zoologici Fennici* **38**: 315–330.
- Holland John, H. (1975) *Adaptation in natural and artificial systems : an introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan Press, Ann Arbor.
- Jaenike, J. (1978) An hypothesis to account for the maintenance of sex within populations. *Evolutionary Theory* **3** 191–194
- Koza, J.R. (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Koza, J.R. (1994) *Genetic Programming II: Automatic Discovery of Reusable Subprograms*. MIT Press, Cambridge: MA.
- Montana, D.J. (1995) Strongly Typed Genetic Programming. *Evolutionary Computation*, **3**. 199-230.
- Moss, S., Gaylard, H., Wallis, S. and Edmonds, B. (1996) SDML: A Multi-Agent Language for Organizational Modelling. *Computational and Mathematical Organization Theory*, **4** (1). 43-69.
- Zaslow, D. (1989) Modernization Of Soviet Agricultural Machine Building - The Example Of Tractor Production. *Soviet Geography*, **30** (7): 559-575.