Overview of Artificial Life Research

Richard Gaywood

December 2000

Abstract

Artificial life is a challenging and intrinsically cross-discipline subject. It covers many areas, fields and skills in an attempt to study the roots of complexity.

This document attempts to give an overview of the current state of Artificial Life research and outline the system I will be developing in my project.

Contents

1 Introduction

Artificial life is a fascinating subject, combining attributes of many different areas into a distinct discipline.

The core concept at the heart of artificial life is to construct, from simple, microscopic components, a system that can exhibit macroscopic, complex behaviour. In doing so, it takes in disciplines from physics, computer science, biology, and many other areas.

However, the main subject being studied in Artificial life experiments is complexity. It allows the construction of systems where the experimenter can observe complex behaviours arise from simple components, and it is this which is principally of interest. It offers insight into how simple rules defined by the equations of modern physics can lead to the complex world we see around us. It also provides a bridge between the bottom-up approach of the theoretical physicist and the top-down approach of an observational biologist, leading ultimately to a better overall understanding of the world around us.

2 Overview of A-life

Fittingly, a-life systems are as diverse and as hard to generalise as life itself. In order to give an idea of how the systems actually work, I will discuss a simple example - that of single-celled bacteria living on an Agar plate.

The model has two key elements; the creatures in the world, and the world itself. The world is a two dimensional grid with an "energy" distribution, so that each cell of the grid has an associated energy store. This energy is distributed according to diffusion rules.

Each creature has an internal energy store, which increases when it eats energy from the world and falls when it moves. When it increases beyond a particular threshold, the animal reproduces asexually, splitting into two or more offspring. Each offspring inherits a proportion of the parent's energy. Such a system is detailed in [1].

Thus, the model can be iteratively looped; each timestep, each creature can make a choice between moving, eating and reproducing based on the environment around it and execute one such move.

Now, one final addition is made - mutation. One of more of the creature's properties is placed under genetic control; for example, the amount of energy needed before reproductive splitting takes place. This is then varied by say, $\pm 10\%$ in a small proportion of offspring. By a long period of many generations, Darwinian evolution can then lead to a fitter type of creature than was originally created.

This is one of the simplest possible a-life systems, but it forms a useful base for more complicated systems to be developed.

2.1 Properties of a generic A-life system

As previously mentioned, there is very little that is consistent from one a-life system to the next. However, there are several elements which, whilst not totally ubiquitus, are present more often than not. This can therefore provide a framework for the exploration of the different kinds of systems that can be built.

At it's core, an a-life system must, of course, have one or more types of organism. These organisms will typically have a number of attributes controlling how they move around, sense the environment, reproduce, and take other actions. One or more of these attributes will be evolutionarily determined, with it varying by random mutation from one generation to the next.

With the creatures defined, it is necessary to build an environment for them. This is usually a two dimensional grid, which sometimes maps onto a torus (ie moving off the bottom of the grid places the creature at the top, and similarly for side to side, giving an infinite plane). The environment will contain features of one or more types; including

- food
- uncrossable obstacles
- hostile areas that damage the creatures

Finally, in order to direct the actions of evolution, it is necessary to introduce some element of competition between the organisms. This generally takes the form of limited food supply, but can conceivably include location within the environment, carnivore creatures feeding from herbivore prey, or any number of other possibilities.

The sticking point in many of these experiments is the degree of competition used; too little, and survival of the fittest will have too weak an effect for observable results; too much, and the population will be driven to extinction. Clearly, this factor must be monitored carefully if an experiment is succeed.

3 The development of an A-life system

The bulk of the work to be undertaken as part of this project is the development of an a-life system. The goal is to develop a toolkit that allows the performance of a-life experiments, so a balance must be struck between ease of use and flexibility.

3.1 The system

The basic framework of the system is a series of C++ classes. There are two base classes, Resource and Critter; the first represents a food type or other resource, and the second a generic creature. Each class is designed with several generic methods; for example, Resource has methods to add food to the world, whilst Critter has methods to decide on an action, move around, reproduce, and eat.

However, the methods are designed to allow easy overriding in subclasses. Thus, two subclasses of Critter might be constructed, Herbivore and Carnivore, with appropriate changes to the decision and eating methods.

3.2 Experiments to be performed

The system described above has been defined to be as generic as possible whilst still being useful. However, there are several specific experiments I hope to perform with the system.

In each case, actually performing the experiment will involve executing many runs of the system and logging the results. Then, a statistical test will be carried out to attempt to demonstrate that one or more of the behavious below has indeed been observed.

3.2.1 Darwinian vs Lamarckian evolution

Darwin's theory of evolution states that over time, the gene pool of a species acquires variation through mutation, and that through a process of natural selection the fittest genes are passed from generation to generation [2]. However, this has not been the only attempt made to describe the mechanism that drives evolution.

An earlier theory was advanced in 1809 by the French scientist, John Baptist Lamarck [3]. This theory claimed that acquired characteristics - in other words, the things an animal trains itself to do during it's lifetime - can also be inherited by offspring. For example, if an adult giraffe stretches it's neck to reach the topmost leaves, then that elongated neck is also passed on to it's young.

Although this has been discredited as a theory of biological evolution, it still has applications in the field of sociobiology; for example, it can be used to model how a parent teaches it's children.

A suitable stastical test here might involve running two populations - one with Darwinian evolution code and one with Lamarckian code - and seeing which is more able to cope with harsh conditions or adapt more quickly to rapidly changing conditions. This allows the use of quite simple metrics such as population size.

3.2.2 Predator-prey co-evolution

This experiment involves setting up a predator/prey couple and observing how the pressures they exert on each other effect their evolution.

An example of this is given in [4].

This is a relatively easy experiment to perform, which can produce easy-to-quantify results. The critical problem will be in keeping the populations alive and viable long enough to observe some co-evolution at work. In theory, it should be able to produce co-cyclic population growth; ie., the population of the prey will follow a sine-like pattern and that of the predator a cosine-like pattern.

3.2.3 Speciesisation

Given a sufficiently large and complex environment, it should be capable to observe a genetically homogeneous population split into two or more distinct sub-species. It is proposed this can be tackled by using a "nutrition" gene to control the herbivore/carnivore tendencies of the population. Hopefully, over time this will polarise into two extremes, giving rise to a population of predators and prey.

Depending on the type of speciesation, this could be relatively straightforward to analyse for; it is relatively easy to detect for initially omnivorious organisms becoming herbivorus and carnivorus, for example.

3.2.4 Emergent behaviour in an ecosystem

Again, given sufficiently complex conditions and sufficient run time, it should be possible to look for emergent behaviour.

This will likely involve attempting to construct a stable network of a dozen or so species; it will be extremely challenging to avoid extinction under these circumstances, and the amount of computer run time necessary might become prohibitively large. If these issues can be addresses, then statistical techniques will have to be employed to sift the data and look for emergent complex behaviour. For example, a possible behaviour might be co-operation amongst animals.

A major difficulty in this type of experiment is how to determine when behaviour is truly emergent. Clearly, the initial set-up must allow the possibility of emergent behaviour, without pre-biasing the experiment. It is likely that performing this experiment will involve a very large number of runs and some statistical analysis to search for emergent behaviour; this then leads to the problem of defining a suitable metric. Even for simple behaviours, this is a non-trivial task and so this is the most complex experiment (and thus the most difficult) presented here.

4 Progress

A prototype of the final system has been written in the scripting language Perl; this shows the class structure which will be adopted in the final project. This prototype has been used to complete a prelimary version of the bugs-on-agar system, and allows the bugs to move around the world, gather food, and reproduce. The final details of the mutatation system are not in the prototype because of a technical restriction of Perl. However, the design of this mechanism is complete.

5 Further work

The next stage is the completion of the final version of the system, which will be written in C++ for performance. Hopefully, with some careful tuning, this will allow some reasonably complex calculations to be undertaken in reasonable computation time.

Work on this is already at an advanced stage, and will be completed by the beginning of next semester. Once the C++ system is in place, it is simply a matter of implemented subclasses as appropriate to perform one or more of the experiments detailed earlier. The emphasis here will be on the scientific method; one advantage of computer-based experiments is that it allows absolute, rigourous control over the experimental conditions. It is likely that many of the final results will take the form of statistical analysis of a large number of runs of the system to attempt to demonstrate that evolution has taken place - ie., that the fitness of the population has increased over time or that some emergent behaviour that was not initially present has been introduced.

References

[1] N. H. Packard: *Intrinsic Adaptation in a Simple Model for Evolution*, Artificial Life: Proceedings of the Sante Fe Institute Vol I, pg. 141

[2] C. Darwin: *The Origin Of Species*, 1859. Electronic copy available at http://www.mala.bc.ca/ johnstoi/darwin/darwin.htm.

[3] J. B. Lamarck: *Zoological Philosophy*, 1809. Partial electronic copy available at http://www.mala.bc.ca/ johnstoi/LAMARCK/tofc.htm

[4] C. E. Taylor, D. R. Jefferson,S. R. Turner & S. R. Goldman: *RAM: Artifical Life for the Exploration of Complex Biological Systems*, Artificial Life: Proceedings of the Sante Fe Institute Vol I, pg. 275