

4th November 1994

Ian,

This is probably a load of rubbish. I know it's very long but for a while I've wanted to write a single document containing a bit of information about everything to do with the project, and this seemed to be a good time.

Feel free to suggest any changes you want, but I need to know by Tuesday lunchtime, as it takes a long time to print half a dozen copies!

Steve.

Small Furry Creatures

The Definitive Article

This document is an attempt to combine in one place all the information that might be required by commercial partners, team members, journalists, technologists and anybody else who is interested in the "Small Furry Creatures" project.

Each chapter will interest different people to different degrees. I have therefore included a brief overview at the head of each section, for those who have more sense than to read the whole text.

Steve Grand (designer & programmer).

Copyright Millennium Interactive Ltd. 1994

NOT FOR DISTRIBUTION OR REPRODUCTION

1. The Concept

Overview

Owning a pet is a joy to people of all ages, and from all walks of life. “Small Furry Creatures” is the working title of a computer entertainment product (not a game) that enables people to keep a pet inside their PC, either at home or in the office. These ‘virtual pets’ are intelligent, lifelike and somewhat bizarre creatures, who live inside a virtual world filled with objects and locations for them and their owners to explore. Virtual pets can be taught a rudimentary language, need care, attention and an education if they are to thrive, and can even, in expert hands, breed and produce unique offspring.

Some of the key features are as follows:-

- A believable, **living world** inside the computer. Creatures to interact with; places and objects to explore.
- New Artificial Intelligence technologies provide the **first ever commercially available artificial lifeforms**.
- **No plot** or imposed sequence. **Something to ‘play with’, rather than something to ‘play’.**
- Capitalises on the user's **imagination, curiosity** and **sympathy**. Largely **non-aggressive** and also **not male-biased**.
- **Easy to use, but not limited in scope**. Opportunities for everyone from young children, through pet lovers and the generally curious computer user, right up to the specialist A.I. and A-Life enthusiast.

Topics

‡ Storyline - What the user thinks he’s getting

‡ Origins and Precedents - Where we are coming from

1.1 Storyline

Perhaps the best way to explain the scenario is to show how it might be described to the end-user. The following is an extract from the current on-line help file. It is merely work-in-progress, and will be completely re-thought for the final product. It also uses terminology and ideas which are less than satisfactory. However, it will serve to demonstrate the basic scenario:



to the fascinating world of vanimal-keeping. If you have never kept vanimals before, please read the following two sections to help you get the most out of this new and absorbing hobby.

What is a vanimal?

Vanimals are a product of the most recent research into Artificial Intelligence. They are virtual animals; that is to say they are real, living creatures who exist in a virtual world generated and maintained by your computer.



Unlike normal life-forms, such as you and I, vanimals are not constructed of Carbon compounds but of *information*. With help from the robotics industry, these creatures could be given hard, physical bodies and could live in our world. However, there is really no need: given a sophisticated enough virtual world to live in, these fascinating creatures can survive inside soft, 'informational' bodies, where they can be made accessible to anyone with a sufficiently powerful computer.

Keeping a vanimal is much the same as keeping a domestic pet, and they can be a source of amusement, delight and companionship. However, vanimals are more intelligent than many Carbon-based domestic creatures, and also somewhat more delicate, requiring careful rearing. They are also a completely new phenomenon, and very little is currently known about them. Thanks to the ability of information to be duplicated rather more readily than skin and bone, vanimals have quickly been made available to the general public, and active research into Artificial Life has moved from the laboratory into the home.

Where do they live?

Vanimals can't be left to fend for themselves in the world outside, like a dog or cat can. Nor can they be left to roam freely inside your computer, where they could cause untold havoc. Vanimals need a 'cage' to live in. However, they are much more intelligent and sensitive than, say, a hamster or rabbit, and their cage must contain a much richer environment than a food bowl and a few basic toys.



Another issue is the difficulty of passing objects through from our physical world into their virtual world. We can't simply feed them through the bars; we must give them their own supply of virtual food, and perhaps even the ability to grow more for themselves when that runs out. A vanimal's environment must therefore be far larger, richer and more self-contained than that for a conventional captive animal, and the 'cage' we have provided in this package amounts to an entire world, complete with weather, plant life,

technology and many other resources.

Programming such a complex world by hand would be very difficult, so we have modified the *Biomimetic* process that we use to generate vanimal life-forms, and have used that technology to create the world for us (see "More about Biomimetics" in a later topic). The Biomimetic system uses its Biomorphing facilities to create your vanimal or vanimals, and then its Terramorphing functions to create a world in which they may live. The disadvantage of this method is that we at Millennium have little control over exactly what sort of world the system produces. We can guarantee that it contains at least enough food and resources to support your vanimal for a reasonable length of time. However, beyond that we can say very little, and the geography and technology of the vanimal world are as unexplored and unknown as the vanimals themselves. See the disclaimer below for more information.


Disclaimer: The Biomimetic process used to generate the virtual world and its life-forms is not 100% reliable in its actions. We guarantee that there will be at least enough resources to keep your new-born vanimal alive without major human interference for a moderate time. However, we do not accept responsibility for any loss caused by negligence, mis-conduct or insufficient care on your part, and we emphasise that maintaining a suitable environment for these creatures requires effort and thought. Also, during the research and development phase for this product, it is possible that certain test species of vanimal became lost somewhere inside the virtual world. We accept no responsibility for any damage to livestock or property that such life-forms may cause.

More about Biomimetics

Located at Millennium's research facility in Cambridge, England, is the huge Biomimetics system, known as Darwin.

Darwin is a massively parallel computing device developed for the creation of virtual life-forms and environments. It is roughly divided into two sections: the biomorphing layer, where synthetic life-forms (vanimals) are created, and the terramorphing layer that has been used to generate the environment in which these vanimals live.

The terramorphing equipment includes a huge CD-ROM-based database, comprising information on physical laws, chemistry, geology, botany, human cultures, architecture and many other subjects. From this mass of data, Darwin has created a life-support system for vanimals; a large virtual world which Millennium's research technicians have named *Albia*. The Albian landscape is fairly strange, and its largely unexplored regions contain a selection of living and non-living objects which Darwin has judged necessary for sustainable life and intellectual and cultural development. We're not entirely sure that Darwin has got it right, and its interpretation of suitable architectural, technological and biological styles seems to us a little bizarre, but who are we to argue? The important point is that intelligent life needs noise, complexity and subtlety around it, if it is to thrive, and Darwin has done its best to provide us with it.



Unlike terramorphing, the biomorphing aspects of the Darwin system are much more within our direct control, but that's not to say we've found it easy to develop suitable life-forms. In fact, there have been a number of earlier prototype vanimal 'species', most of which were unsuccessful for one reason or another and have been abandoned (although there are rumours that one or two examples slipped the net, and even now are roaming around the more unexplored regions of Albia).

These various vanimal breeds have all been given names by their creators, taken in traditional fashion from mythology. The first creatures we developed were known as *Ettins*, and were just too stupid to justify releasing to the vanimal-keeping hobbyists. One or two of them are probably still hanging around somewhere in Albia, but they won't give you any bother. This isn't true, however, of the Mark II creatures, which we called *Grendels*, after a vicious monster in the Beowulf saga. Grendels turned out to be true to their name and had an excess of bellicosity which made it very difficult for research to continue for some time, as the aggressive and unfortunately rather mindless creatures would attack, kill and usually eat anything that got in their way. Any raggedness and dilapidation you notice around Albia are probably down to the ravages of Grendels, but we are absolutely certain, well, fairly sure anyway, that all Grendels have now been destroyed.

A much more pleasant personality was characteristic of the Mark III vanimals, which we call *Norns*, and it is these creatures that we have decided to release to the general public, and with which, we hope, you will enjoy interacting. Norns are fairly gentle, unaggressive creatures, whose only vice seems to be a penchant for impishness. In fact they behave remarkably like mischievous human infants. They are not very bright, it has to be said, and they need a good deal of running around after. However, we believe that they are intelligent enough to learn the rudiments of a spoken language, and should therefore be fairly easy to control and interact with, given a little patience.

The only other prototype species yet researched is a much more sophisticated race, known officially as Mark IV, but unofficially named *Side* (pronounced “shee”), after an ancient mythical Irish race, famous for its magic, intellect and wit. Unfortunately, the Side seem to consider themselves intellectually more capable than our research team and tend to drift off into the more sophisticated regions of Albia, where we can’t find them and they feel free to meditate in peace. This has made it rather difficult for us to make much progress with the Mark IV types, and so at the moment we are restricting ourselves to working with the Mark III Norns, who aren’t so snobbish. You probably won’t come across a Side creature in Albia, but if you do, can we have it back, please?



1.2 Origins and Precedents

A product like this has, as far as we know, never been attempted before. However, there are some precedents, and these show us that the idea is innately interesting to people and very likely to ‘touch a nerve’ with a significant proportion of our huge potential audience.

The first, and perhaps only, *game* in this style was Activision’s Little Computer People (early 1980s). This started from the premise that everyone’s computer contained a little house, occupied by a solitary human being who contentedly lived out his daily routine, modified only slightly by the intervention of the user. The level of intelligence was exceedingly low, and there was only very limited scope for user interaction, but the joke was a good one, and within the limits of eight bits it worked fine. It certainly sold in large numbers and on many platforms.

Several of us at Millennium have long thought that the idea was nowhere near exhausted, and that new technology and software techniques would allow for a much bigger and more interactive environment of this kind. Given 32-bit power, a decade’s active interest in Artificial Life and Neural Net technologies and the experience gained from more conventional games with similar qualities, we decided the time was right for ‘Little Computer People II’.

The design of this game also owes something to a book by A.K.Dewdney, called “The Planiverse: Computer Contact with a Two-Dimensional Universe”, which unfortunately is out of print, at least on this side of the Atlantic. The story dealt with an accidental fusion between a simulation written by a group of computer science students and a real, two-dimensional world occupied by strange and fascinating lifeforms, whose nature the book went on to study in detail. Unfortunately, a direct computer rendition of the book would be a bit too intellectually taxing, both for the programmer and the user. For example, tubes can’t exist in a 2D world, and so a creature with a digestive tract would

simply fall apart, if it weren't for the evolution of the zip-fastener. What's more, in a 2D world, creatures cannot pass each other: instead, one has to climb over the other, leading to all sorts of social consequences!. The book is a good read if you can get hold of it (I've lost my copy).

Other archetypes for this product include the huge number of fictional accounts of 'virtual' worlds, such as Tolkein's "Lord of the Rings", and other common obsessions with 'little creatures' such as dolls, robots and, of course, domestic pets.

2. The Product

Overview

The product is CD-based and runs under Windows 3.1 or Windows'95.

It can be used by more than one person sharing a machine (each has his or her own pet, but those pets co-exist in the same world).

The world is large enough to require extended exploration, which can only be achieved by encouraging your creature to travel. You and he thus encounter new landscapes, opportunities and problems as you both attempt to circumnavigate Albia.

Some of the objects you encounter are useable by the creatures themselves; others can be dragged out of the world, and become small Windows applications that allow the user to control or examine Albia's landscape and lifeforms in various interesting ways.

The whole system has a component-based, programmable architecture. This means that, for example, new facilities (objects, tools and even landscapes) can be added by the user after the initial purchase.

Topics

- ‡ System requirements
- ‡ The Pet Shop - choosing your norn
- ‡ Albia - the main virtual world
- ‡ Tools and Toys - applets to control and monitor Alban events
- ‡ Summary of features

2.1 Screenshot

This picture shows the Albion landscape, seen through the ‘camera’ window, surrounded by some of the possible tools and toys (see below).



2.2 System requirements

The initial target platform is an MPC-compliant IBM PC. The software will come on CD-ROM and will install/deinstall portions of itself onto the hard drive as required.

The software runs under Windows (including Windows'95), for several reasons:-

- Windows has an attractive and consistent user interface, making the game visually appealing and easy to learn.
- The game is something to ‘dip into’, rather than something to ‘play’ over an extended session. It is also likely to appeal to business PC users as much as home users. Writing the software for Windows therefore allows the user to ‘keep an eye on his pet norn’ while he works, and flick quickly and easily between the game and, say, a spreadsheet or wordprocessor. Owning a “small furry creature” is an extension of the desire to customise and personalise our desktops.
- Windows is multitasking, and offers a wide range of features that allow applications to interact with each other. This allows the product to be developed in a modular, extensible way. It also opens the possibility of interactions between the game world and other Windows apps, such as Notepad, Paint and Excel.
- Windows’ device independence makes it much easier to support a huge number of PC configurations. The Windows multimedia standards also make video, sound input and other

technologies more easily accessible.

2.3 The Pet Shop

The first thing the user sees when he installs the game will be a pet shop, containing a variety of creatures (each being a 'breed' of the species we call Norns). He will be able to choose the creature that most appeals to him, and will be asked to 'register it' by naming his creature and allocating a password. In future, he will be asked log-in with that password on start-up. This system allows several users to share the same world (on a single PC), but prevents unauthorised direct control of someone else's pet (of course, you can still try to get your own norn to attack or otherwise interact with another's pet, if you wish). Logging-in is also a way of telling the system which creature the 'camera' is to track, and which creature is the target for any medical or other tools the user may operate.

Allowing the user to choose his own pet is an embarrassingly simple solution to a problem that has taxed us for some time, namely, what should these creatures look like? Each of us had our own ideas, and we could never agree. All we have to do now, however, is to provide a good range of types, and let the user make his own choice.

Each creature is unique. If you 'buy' a particular creature from the pet shop, then no-one else can have one like it. If your pet dies, then you will have to choose another. All this is to develop a sense of ownership, and to make people care what happens to their norn. This is crucial if people are to become deeply involved in the product.

Each of the creatures in the pet shop is a 'thoroughbred' example of its type. Its appearance and its temperament are determined by its 'genetics'. If a user is able to breed two creatures then the offspring will be genetically derived from its parents, and will show a combination of their physical features and personalities. This opens the possibilities for such esoteric activities as breeding norns for sale or to give to others, or attempting to breed a 'master race' of norns! On the other hand, we will deliberately make norns very delicate and difficult to breed, so that a successful birth will be the pinnacle of a norn-owner's achievements.

2.4 Albia



The main application is a ‘camera’ view onto the planet Albia. Albia is a ‘wrap-around’ world approximately 32 screens wide and 4.5 screens high. The camera automatically tracks your norn, but the user can also pan it a short distance in any direction. The only way to see the whole world is to get your norn to travel from place to place. This involves tempting, cajoling and commanding him to walk or to ride in vehicles, such as the lifts that take him from burrow to burrow, a ship and submarine that allow him to cross the ‘Meniscus Sea’ or a hot-air balloon that can cross over the ‘Mountains of the Ragnarok’.

Your norn arrives on Albia in the Jungle region, where a small tree house contains the bare essentials to get you started. Your norn can attempt to pick up, drop, activate or deactivate the various objects scattered around the landscape, and will soon discover what things can be eaten, how to operate lifts, why walking into walls is not a good trick, and so on. Many of these objects can also be carried or activated by you, by clicking on them.

One of the first things you may try to do is to teach your norn some simple words. You can teach him names for things by drawing his attention to them and ‘saying’ the word. More abstract qualities such as verbs need to be taught using one of the blackboards you find in Albia. Each blackboard has a set of selectable ideograms - pictures representing ideas such as “make it do something”, or “approach it”. One useful ideogram represents “myself”, and allows you to teach your norn its own name. By writing words on the blackboard, the norn will learn to associate those words with their ideograms. Not only will that allow you to understand him as he tells you what he is doing, but it will allow you to give him commands. Whether he will obey those commands, however, depends on a number of factors. Once your norn knows a reasonable vocabulary you can communicate with him via simple ‘pidgin’ sentences, such as “ron start lift”, “come me”, “get ball ron”. Ron will likewise describe his thoughts and intentions at times, in increasingly understandable ways, as his flair with words increases. Note that you may teach your norn any language you please. Surprising things can happen with language; for example, your norn may pick up new words that he has heard other creatures use; likewise, children will learn from their parents.

You and your norn will probably progress around the planet Albia, through the jungle, across the Meniscus Sea and into Outgard (a pleasant, rural area well-stocked with interesting objects), thence across the Desert of Volund and over the Mountains of the Ragnarok. As you travel, the climate, the availability of food, resident creatures and other hazards form potential sources of hardship, damage or disease for your norn. On the other hand, you will come across new objects to examine; some of which may turn out to be useful. Some examples of Albian objects are: vehicles, foodstuffs, seeds (for growing food), toys (balls, musical instruments, video cassettes, etc.), medicinal herbs, books (for you to read - a herbal, perhaps?), medical equipment, research tools (for studying norns’ brains), useful machinery (irrigation pumps, etc.)...

2.5 Tools and Toys

Some of the objects you discover in Albia are more use to you than to your norn. These tools and toys are small applications that you can ‘remove’ from Albia and place onto your desktop. Some of these applets are just fun, for example a musical instrument to play to your norn —will he dance or will he cringe?. Others are useful to most people, for example a ‘lifeform detector’, so that you can track down other creatures. Still other tools will only interest certain people, particularly those who have a degree of technical interest in how it all works. These are the ‘research tools’, which allow interested users to monitor their creatures’ brain patterns, or to stimulate individual neurones to try and find out how a norn’s brain works.

One important technical point about these tools is that they are separate Windows applications, which ‘register’ themselves with the main Albia window and thus can be launched from there. These applets then communicate with Albia and its creatures using the Windows Direct Data Exchange mechanism. One of the reasons for all the technical complexity involved in this approach

is that it allows new tools to be provided as add-ons, after publication. This may not be required, but it is as well to provide the hooks, just in case!

A side effect of this design, which would perhaps only interest the absolute enthusiast, is that other Windows applications could be programmed to interact with Albia. For example, a really keen norn-owner could write a program in Visual Basic to alter the Albian climate, or an A.I. enthusiast could hook up an Excel spreadsheet to his norn's brain & record interesting data. In the other direction, objects inside Albia are potentially capable of controlling external Windows applications. Imagine if your norn sent you a letter in Word, or got bored and started up your Media Player so he could listen to some music. All of this may be way beyond our plans for a first-generation project, but it *is* possible.

It is not only the tools themselves that are object-oriented and modular. Objects and creatures in the world are also independent entities, capable of being extracted from and added to the world. One consequence of this is that we can include a 'comms' tool somewhere in the Albian landscape. This would allow a user to send his norn to a friend, to look after while he goes on holiday! Such a mechanism also allows norns to be traded and distributed down the Internet. This may not be of any great practical value, but the mere fact that it can be done helps to build the product's appeal.

If we find ourselves in a position to publish add-on products, then as well as objects, creatures and toys, whole new worlds could be supplied. Drop a rocket ship into Albia, get your norn to climb aboard, pull the lever and go off to explore another planet together!

2.6 Summary of features

- ☺ Own your own pet. Let it share the world with other creatures belonging to friends or family.
- ☺ Explore the landscape. Travel in strange vehicles over varied terrain.
- ☺ Discover and play with objects. Find out what they do. See how they affect your creature.
- ☺ Teach your pets to communicate. Give them orders. Listen to their thoughts.
- ☺ Watch your norns play whilst you work.
- ☺ Help your norn through difficulties. Nurse him if he falls sick.
- ☺ Attempt to breed new creatures.
- ☺ Send your norn to a friend to look after while you go away.
- ☺ Grow food for them to eat. Study the effects of medicinal herbs.
- ☺ Read Albian books. Look at magnified images of Albian objects and listen to a commentary.
- ☺ Watch a video documentary on Artificial Life.
- ☺ Learn about the system through comprehensive on-line help.
- ☺ Study your creature's brain activity. Find out how it works. Log long-term mental and physiological changes to really get to know your norn. Participate in a world-wide study of norn behaviour.
- ☺ Connect to the world and its inhabitants through other Windows programs.

☺ Buy new objects, creatures and worlds.

3. The Technology

Overview

The main innovation in this product is the use of neural net technology to confer lifelike behaviour on the creatures. The neural structure is my own design, developed over a number of years.

Each creature in the system contains a brain, composed of roughly 1,500 'neurones', arranged in a heterogeneous array of 'sub-systems'. This technology makes the creatures lifelike, sensitive to their environment, able to learn from mistakes and capable of generalising from past experiences.

The idea of producing believable life inside a computer has fascinated me for a long time, and can be seen as part of a growing trend towards research into 'Complexity' - a synthetic, rather than reductionist approach to scientific issues. The science of Artificial Life (A-Life) is currently being studied in various places, notably the Santa Fe Institute. It has little in common with the older discipline of Artificial Intelligence.

Topics

- ‡ Brain Questions and Answers - information about the technology
- ‡ Artificial Life Research - what other people are doing

3.1 Brain Questions and Answers

It's very difficult to cover such a vast, complex and technical topic in a relatively short document, particularly where trade secrets may be involved. However, some of you will want to know something about the technology that makes the product tick, so I've decided to try to explain some of it in the form of a mock interview. If only I could be this ready with my answers in real interviews!

Q. How do the norms' brains work, exactly?

A. I'm not telling you!

Q. Alright, just give me a rough idea. What sort of mechanism is involved?

A. The basic approach is called a Neural Network. In other words, a large array of relatively simple processing units, representing neurones, is wired up to 'detection' functions, representing the creature's senses, and 'action' functions, representing the creature's muscles, voice, etc. Sensory

information enters the network at certain points and ripples through it, eventually to arrive at one or more of the output regions, where it results in an action from the creature.

Q. I've heard of neural networks. Aren't they used for character recognition and things like that?

A. Well, yes they are. However, I'm glad you asked me that, because it allows me to explain that the neural network used in norns' brains is very different from that which has become the 'standard' neural network design. Most neural nets are called Hidden Layer nets, and usually consist of three layers of very simple 'neurones'. One of the trickiest bits about neural nets is how you make them learn from their mistakes. The mechanism used in conventional nets is called Back Propagation. This is a fairly mathematical way of sharing out responsibility for 'errors' in a net's response to a given set of inputs, so that the next time those inputs occur, the response will be less 'wrong'. My network is not a connectionist-style, hidden layer network, nor does it use back propagation for learning.

Q. So what's wrong with back propagation?

A. Nothing, in its place. However, B.P. nets are usually used for pattern recognition tasks in circumstances where a) you can determine how far from the 'correct' answer the actual output was (quantifiable errors) and b) you have the opportunity to 'train' the net on a set of known questions and answers, before exposing it to 'real' problems. Neither of those features is true for living creatures like norns: they have no opportunity to train - life itself is their training ground as well as the source of 'real' problems; also, in real life there are no 'correct' answers, only good ones, and there is no way to quantify the difference between good and bad solutions.

Q. If it's not a simple, three-layer structure like you say most nets are, then what structure does it have?

A. My net is not derived from the mathematical theories and approaches that so-called connectionists use (as in the three-layer, B.P. net). My mentor is not Maths but Biology. I have developed a system based on what I know of the structures and functions of parts of real brains, combined with a few hunches, the odd botch and a lot of experience of complex dynamical systems in general. My networks, although essentially simple, are *heterogeneous*, in that they are made of several different varieties of neurones, wired up into a collection of discrete and functionally different circuits. Different regions of a norn's brain deal with different tasks, or different aspects of the same task. For example, one region deals with directing the creature's attention to features in its environment, another pair of regions co-operate to handle the main activity of making decisions and learning from mistakes. Still other regions are involved with needs and drives, and aspects of language.

Q. So you are saying that the system is 'wired up' in various ways. Does that mean that the creatures' actions are simply programmed, albeit in an indirect way?

A. No, absolutely not. Just like the real brain, my brains are wired up at birth to form a basic *framework*, that determines how one group of neurones interacts with another. During life, however, the creatures' brains *alter* their structure; not drastically, but in myriad small ways. Some of the changes are to the internal 'chemical' and 'electrical' properties of neurones; others involve changes to their interconnections. All of those changes are produced by the environment, not by me. Therefore, I don't determine a creature's actions, only the mechanism by which a creature decides for himself which actions to take.

Q. Real animals like dogs and cats have huge brains. Surely you can't emulate something that large in a computer? And if you can't, surely your creatures are going to be very stupid indeed?

A. The Human brain contains roughly 10,000,000,000 neurones. My creatures' brains contain 1,500. This is still a lot when you consider that they need to be processed one after another, ten times per second, and the world may contain a dozen or more active brains. However, I'm not too worried about the size discrepancy for several reasons: Firstly, natural brains seem to have a great deal of redundancy - neurones are dying all the time, never to be replaced, yet we need to retain memories and skills for decades. Artificial neurones don't die, so you don't need massive redundancy to cover up for their loss. Secondly, for a norn to lift his arm requires a single pulse from a single neurone (in fact a whole co-ordinated action such as walking only requires one pulse); a human has to control dozens of muscles and counter gravity, friction and other forces in order to do so effectively. This obviously requires a huge number of neurones. Lastly, for a carbon-based creature to recognise his mother, a massive amount of pattern recognition must be performed, requiring a large percentage of the total brain mass. In a virtual world, I can cheat tremendously - my creatures can recognise an object by 'sight' simply by detecting its ID number.

Q. But isn't that cheating? If complex actions such as walking, and complex sensory tasks such as visual recognition are handled without needing neurones, what's left?

A. It is cheating a bit, but I can't crack all the problems in one go. In any case, serial processors could never handle the volume of information required. It doesn't matter, though, for the purposes of this application. It would be nice if it took a fledgling norn a long time to learn to recognise his mother, and it would be even nicer if he had to *learn* to walk, rather than knowing how to from birth (maybe in version 2?). However, the most important task I wanted to handle using neural techniques was the act of *decision making*. Having determined it can see some food and it is hungry, and knowing already how to pick things up, walk towards them, jump on them and so on, what does the creature actually *decide to do*? And if he makes a poor decision, how does he learn from his mistake?

Q. But what about this learning from mistakes business, though? If a pattern-recognition network fails to recognise the letter 'A', it just sighs and readjusts itself, ready for next time. If a real creature thinks that a good thing to do to a lion, however, is to tickle it under the chin, then that creature won't get a second chance to try a new approach!

A. Another good observation, Watson. Glad you're paying attention. This is (to my mind) a key element in the thought processes of living things, and involves the one small scientific contribution I've perhaps been able to make. The point is that creatures *generalise*. Whenever they meet a novel situation, they relate it to all their past experiences of *similar* situations, and use the net result of that knowledge to determine how to deal with it. The tricky thing is: how do you decide what 'similar' means, how do you 'remember' those past experiences, how do you use that information to guide you towards one or more sensible options, and how do you learn from that process? The answer to this is a key element in my system and is one of the bits I'm not going to tell you about!

Q. So, these creatures think on their own, generalise and learn. What can't they do?

A. Um, well, quite a lot, really. As I've said, they only have neural mechanisms for dealing with what are called Declarative processes, such as recalling past experiences. They currently don't have intelligent control over Procedural processes: i.e. they can't learn skills. What's more, their ability to *reason* is very limited. They cannot take abstract concepts and combine them in order to make judgements based on chains of logic. Norns are never going to make good chess players, I'm afraid. However, I've achieved most of what I set out to achieve, that is to say, I've created creatures with subtle, sensitive, believable, *lifelike behaviour*.

Q. You say lifelike, so are your creatures actually alive?

A. Tricky. It depends on what you mean by 'life'. Most attempts at a definition of life start by detailing features *characteristic* to living things, and then declare those features to be *necessary*

qualities for life. It's like saying that all coins are made of metal, and therefore, if it's not made of metal, it can't be a coin! This is a bit unfair. It can also get silly: a universally quoted characteristic of living things is their ability to reproduce; but does that mean that a man who has had a vasectomy is no longer alive? I think the easiest way out is to declare that Living and Non-living are not clear-cut distinctions but simply the two extremes of a continuum. Some things are more alive than others. A rock is not alive at all (though some were once alive); a fish definitely is. A thunderstorm, on the other hand, is more alive than a rock, but less so than a fish. However, there are some critical features that *arise* out of living things but are not seen in non-living things - so-called emergent phenomena. I think my norns show enough emergent features in their behaviour to count them as *not* non-living. Enough of the philosophy - yes, I believe my norns *are* alive. Or at least *virtually* alive.

Q. If they are alive, would it be immoral to kill them?

A. Look, this is getting too hard - I think this schizophrenic interview thing has gone on long enough. I kill off norns all the time. I have to. I do sometimes, however, feel bad about it. If I can make other people feel bad about it too, then I think I will have done a good job.

Q. OK, one last question. If you get an opportunity to do a second product involving this technology, what will you improve?

A. This project is really only a 'proof of concept'. There are many things I'd like to improve, and many more that will become possible as computer power increases. One of the next steps is to take some time out to research a new brain model, without the pressures that come from having to develop it at the same time as a saleable product. I already have the basic theory for 'brain II', but it needs some further work, which will have to wait until this project is published. Besides the brain, there are plenty of things to improve in the virtual world: I chose a simple, flat viewpoint, with limited physics, to make sure I didn't bite off more than I could chew. A more complex, modifiable, 3D world would be much more fun but would require a lot more intelligence. Given enough outside interest, time and commitment, I think I could really do something exciting.

Q. Thanks.

A. That's OK. Very smart questions, I thought.

3.2 Artificial Life Research

We are not the only people mad enough to try to create artificial life. Dr. Frankenstein aside, there has been a significant amount of work in recent years in the area of A-Life, although most of it has been fairly abstract or concerned more with reproduction and genetics than lifelike behaviour and intelligence. This section is a short background on the subject.

3.2.1 What is A-Life?

A-Life is a part of a new branch of science. It maintains that 'life' is merely a property of certain kinds of *complex systems*, rather than a unique result of the chemical properties of Carbon or a metaphysical 'substance' bestowed on objects by, for example, a deity. Moreover, A-Lifers believe that, since computers are capable of modelling an indefinite range of complex systems, then they are also capable of modelling systems that embody life. Note that the implication is not that computers can *model* life, but that these computer-modelled systems *are* alive — this is a fine, but important distinction!

A-Life is a small part of a new movement in Science that in total amounts to a completely new paradigm. Complexity Theory, Chaos Theory and the mathematics of Fractal Geometry are all associated fields, and the resulting 'New Physics' involves a shift away from the Mechanistic, Reductionist approach that has characterised Science since the 17th Century. The prime mover for this change was the invention of the computer and the new viewpoints that this device has given us.

Artificial Life as a broad concept is as old as computing itself (even older, in some ways), and John von Neumann (a 'father' of computer theory) was one of the first people to think in depth about self-replication and other A-Life issues. However, the 'birth' of A-Life as a distinct field probably dates from a conference held at Los Alamos in 1987, organised by Chris Langton. He and John Holland (a neural network pioneer) both work at the Santa Fe Institute in New Mexico, where a number of A-Life projects have been attempted and most Complexity research is currently concentrated.

The motives of A-Life researchers range from those of the biologist or psychologist, who is interested in modelling lifelike systems to see what they tell him about 'real' life, through to the roboticist, who is interested in developing intelligent and adaptable machines, regardless of how dissimilar his design is from that of Carbon-based life.

3.2.2 What's the relationship between A-Life and A.I.?

Artificial Life research is fundamentally a 'bottom-up' approach to creating intelligent or adaptable systems, and has surprisingly little in common with the older discipline of Artificial Intelligence. A.I. approaches the problem from a 'top-down' viewpoint, and deals with such concepts as 'symbol processing' and 'decision trees'. Where A-Life uses computers to create *systems* which are intelligent, A.I. started with the more naive approach that a computer could itself *be* intelligent. A.I. as a philosophy has been seriously discredited in recent years (largely due to its over-optimistic predictions), although it has found a valuable niche in such un-lifelike but practical fields as Expert Systems and Computer Translation. A-Life is also a much broader field than A.I., covering artificial genetics, self-replication, self-organisation and other lifelike characteristics, as well as intelligence.

3.2.3 What has been achieved so far?

'Not a lot' is perhaps the honest answer, but then the field is very young. In terms of developing a paradigm and inventing new approaches (such as the genetic Algorithm and Cellular Automata), A-Life has already come a long way. However, in terms of results that would impress a layman, not much has been achieved (partly because many non-computer people still labour under the misapprehension that computers are *innately* 'clever', and thus miss the point). More significant though, is the fact that very few A-Life techniques have filtered through into real products.

A few examples of the State of the Art in research-oriented A-Life systems are:

- PolyWorld — an 'ecosystem' of simple predatorial lifeforms with a genetic structure and very basic behaviour (eat, mate, move). These coloured blobs evolve different survival strategies, for example cannibalism, without guidance from the programmer.
- Boids — a group of blobs that execute a very simple set of movement rules, yet show sophisticated flocking behaviour. An example of emergent group behaviour that would be difficult to predict from the individuals' own characteristics.
- Loops — a cellular automaton (an array of basic 'computing devices', like the Game of Life) that shows the ability to self-replicate.

- Beer's Cockroach — a 'virtual insect' (on a screen, not a robot), whose ability to walk, navigate around the landscape and seek food is controlled by a semi-hard-wired neural network.

3.2.4 Neural networks

A-Life encompasses many things: genetics, evolution, group behaviour, etc. It also concerns itself with *intelligence* — not the chess-playing, illness-diagnosing, A.I. sort of intelligence, but the “where will I find food?”, “how can I stop him hitting me?” kind, that characterises all living things. One of the most useful approaches in this field is the neural network. Neural networks were developed from the neurological work of Hebb and others in the late 1940s. They flourished briefly in the new era of computing, but then fell out of disfavour, due to the theoretical difficulties involved in making them learn. A very rarefied notion of the neural net returned to popularity in the late 1970s, and these highly mathematised and idealised nets are currently in vogue for a number of practical pattern-recognition tasks.

The basic premise behind neu-nets is this: brains are made of neurones; neurones are (relatively) simple, predictable devices that could be modelled on a computer; perhaps if we lumped enough neurones together, then intelligence would *emerge* as a result (this follows the A-life, ‘emergent behaviour’ thinking, rather than the top-down, work-out-the-rules-and-code-them methodology of A.I.). The only snag with this argument is that there are an enormous number of ways that we could connect up these neurones, and we don't know which ones lead to intelligence. Given ten billion neurones, each with tens of thousands of connections, there are more ways to wire up a human-sized brain than there are particles in the universe — and for all we know, only a handful of them actually work!

3.2.5 What's the point?

At such an early stage in the field's history it's not surprising that A-Life hasn't yet reached public awareness. So far, there's not much you can do with it. However, there is one field that is practically crying out for the use of A-Life techniques, and that is Computer Entertainment.

Pure games, with their explicit rules amenable to algorithmic solutions, have no need of the unpredictable qualities of living things, which would only get in the way. On the other hand, strategy games need one or more computer players for the user to pit his wits against — creating such ‘intelligent’ systems using conventional techniques (hard-coded rules) rarely works well and is often a disaster. Playing chess against a computer is one thing (chess algorithms being the A.I. community's Service to Mankind), but playing a war game against a computer opponent is rarely satisfying, since no programmer can figure out an effective algorithm for it. Even computer football players leave a lot to be desired. Give a computer player a neural network, however, and you potentially have a system that's capable of *learning* how to play the game. Not only that, but the system *continues* to learn and adapt to the player after he's taken it out of the box. This is not an easy thing to achieve, but techniques such as neural nets have a much greater chance of achieving it than hard-coded rules do.

While strategy games could benefit from a hidden ‘mastermind’ to control them, simulation games have a much more explicit role for computer lifeforms. Until recently, simulations were limited by slow, clumsy graphics. However, the next generation of machines will be able to create beautiful, 3D, texture-mapped, virtual worlds, full of believable, freely manipulable objects. The potential of simulation games is enormous, and in the future people will spend an increasing amount of time treading virtual landscapes. However, the more believable these worlds look, the more a lack of intelligent inhabitants becomes a glaring omission. In the long term, VR games will undoubtedly be multi-user, but even then, who's going to want to play the bit-parts? Who wants to be your virtual dog, or the waiter from whom you order your virtual meal? Without *artificial* life-forms, such virtual worlds are going to appear sterile and lifeless. And in the shorter term, before networking

becomes universal, *all* the inhabitants of virtual worlds will have to be artificial. The techniques for rendering 3D graphics are now well established, and some work has been done on making the physics of virtual objects convincing, but almost nobody is working on how to give those objects *minds*.

4. Where might we go from here?

Overview

This project has been a long time in the making, and still has some way to go. However, we have cut our teeth on a number of new things here - principally, a new technology that had never been tried before and which we were not sure was going to work, plus a new product concept that falls well outside of existing pigeonholes. On the assumption that we were right to trust our judgement in these areas, we can start to consider how we might build on the hopeful success of Small Furrries in the future.

Products involving virtual animals and eventually, perhaps, virtual humans are one strand of possibility. Another is the application of neural nets of a broadly similar type to other product areas, in conventional games and elsewhere.

Whatever happens, I am sure that this is a line of thought that has tremendous value in the computer entertainment business, and in fact our business is a surprisingly good environment in which developments of this kind can be nurtured and utilised.

Topics

- ‡ Add-ons to Furrries I - ways to make the product last
- ‡ Small Furry Creatures II - a sequel on the same or a new theme
- ‡ Virtual Reality - brains and future software technologies
- ‡ Other applications for 'sentient' technology - computer players and pie in the sky

4.1 Add-ons to Furrries I

As I've said above, if Small Furry Creatures is a success, we may be in a position to support the product with add-ons. These should be fairly quick and easy to produce, and can include: new tools and toys, new creatures, new objects for norms to play with and new landscapes for them to explore. Some of those landscapes could incorporate 'game' elements, such as a 'battle world', stuffed with weapons and warlike creatures.

4.2 Small Furry Creatures II

The current project has been a learning experience in several ways: a pilot for a developing technology, a scratchpad as we thought our way through this novel style of product and a volatile piece of code that has had to adapt to a rapidly changing PC environment. I think it is coming together nicely, but it shows some of the scars of that development process: The graphics and terminology retain echoes of long lost storylines, plots and game-designs, while the code has been battered and bruised in its passage from 16-bit DOS, through 32-bit extended DOS into Windows and is now about to undergo another extensive overhaul to make it more extensible and modular. I can't go on rewriting it for ever - we need it to be a real product. However, if it is the success that some people tell us it's going to be, then it may justify a sequel, to *really* get it right. Given all this experience, we would be able to make much clearer decisions about the design. But what should this sequel be like?

One of the most conservative aspects of Furies I is the graphical viewpoint. Furies II could be much bolder and use a full 3D or semi-3D (isometric) view. This places greater demands on processing power, intelligence and 'virtual physics' but would result in a much more believable environment. Small Furry Creatures inside a Doom-style environment would be something to behold!

As we've discovered during the early stages of this project, A-Life systems are not very applicable to conventional game formats like adventures, where their unpredictability makes it virtually impossible to lay down a controlled gameplay. For Furies I, we have developed a fully free-form, 'computer pet' concept, with no attempt at a game. Furies II could develop this theme, but it is also possible that we could introduce a goal, or at least a strong sense of progression, if we chose the right storyline. One possibility is the evolution of a humanoid race. Imagine a virgin landscape into which a group of none-too-smart, human-like 'cavemen' migrate. On their own, they would settle into communities, build shelters and maybe discover farming. But what would happen if you interfered and gave them the Wheel? What if you gave them the Atom Bomb first? What happens if you place a bridge over a wide river, and thus allow two hitherto isolated tribes to interact? What happens if, when they've become sophisticated and discovered cell-phones and fax machines, you introduce severe global warming and make their comfortable lives untenable? This seems like something which we are now well-placed to make possible, and which would be very difficult to do convincingly without the Furies technology.

Obviously Furies II is also an opportunity to develop new brain techniques. I am already part way towards a 'Brain II', and would hope to raise the level of intelligence by at least a factor of two.

4.3 Virtual Reality

Thinking further into the future, VR is going to become a much more common feature of computer-based entertainment. The Information Superhighways and advanced computers of five years hence are going to support startlingly real, walk-in environments. Enormous efforts have gone into developing 3D surface rendering techniques, and some work has been done on 'virtual physics', to develop creatures which move realistically and objects which interact in complex, almost unlimited ways. However, no-one else seems to be working on ways of incorporating intelligent living things into those VR environments. This is something we may be able to capitalise on.

4.4 Other applications for ‘sentient’ technology

I have already mentioned the possibilities for using neural nets to provide the intelligence behind ‘computer players’ in sports and strategy games. One of the things we already plan to do is to develop a ‘black box’ brain system, that other programmers in the company can drop into their own game software and modify to suit their own needs.

Such applications are immediately practical, but, to finish, here’s some pie-in-the-sky:

In my spare moments, I’m starting to build a small robot, to use as a testbed for future brain designs. This is just for fun. However, despite being well outside of our commercial territory, I will end by mentioning one more application, far-fetched as it may seem. What about sentient toys? Children are used to dolls that wet themselves, move their arms and cry, but imagine a doll that stops crying when you soothe it to sleep with a lullaby, or that learns to copy your words. What about ‘real’, tactile pets that respond to you and move realistically, yet don’t need to be cleaned out and can just be switched off when they cease to amuse? I know that these ideas would involve all sorts of technical and cost problems, but one day such things will be possible, and somebody will make a lot of money out of them.

