# The Meme Machine

Susan Blackmore

1999
Oxford University Press

## Contents

## Chapter 16

# Into the Internet

In our house we have four telephone lines, two fax machines, three television sets, four hi-fi systems, seven or eight radios, five computers and two modems. And there are only four of us. We also have many thousands of books and a few compact disks, audiotapes and videotapes. How did all this stuff come to exist and why?

If you have never asked yourself the question you might think the answer is obvious. All these things are great inventions, created by other people to make our lives better or more fun. But is this the right answer? Memetics provides an entirely different answer, one that is somewhat counter-intuitive.
I suggest that memetic selection created them. As soon as memes appeared they started evolving towards greater fidelity, fecundity, and longevity; in the process, they brought about the design of better and better meme-copying machinery. So the books, telephones, and fax machines were created by the memes for their own replication.

This may sound odd when we know that memes are just information being copied from one person to another. How can bits of information create radios and computers? But the same question could be asked of genes – how can bits of information stored in DNA create gnats and elephants? The answer is the same in both cases – because the information is a replicator that undergoes selection. This means the evolutionary algorithm runs, and the evolutionary algorithm produces design. The design of computers by memetic selection is, in this sense, no more mysterious than the design of forests by genetic selection. The consciousness of a designer is not the causal factor in either process. Design comes about entirely from the playing out of the evolutionary algorithm.

We are used to the idea of animals and plants being designed by natural selection, but we must also think about the evolution of the replication machinery that makes natural selection possible – for both have evolved together. This is the analogy I shall to draw here. Memes do not yet have

precise copying machinery as DNA has. They are still evolving their copying machines and this is what all the technology is for.

It is helpful to look back at what must have happened in the case of genes – the only other replicators we know much about (Maynard Smith and Szathmáry 1995). When the first ever replicator arose on this planet it was presumably not DNA but some simpler precursor, or even some completely different replicating chemical. Whatever it was, we can be sure that the cellular machinery for copying it did not exist. Natural selection in the very early days of life was not selecting between complex organisms like cats and dogs, or even different kinds of simple cell, but between little bits of protein or other chemicals. Any of these proteins that got copied more often or more accurately, or that lasted longer, would have survived at the expense of the rest. Gradually, from these beginnings, natural selection would produce not only more proteins but proteins that took part in the copying of other proteins. Eventually, there evolved the system of groups of replicators, replicating machinery, and vehicles that we see today. The system settled down so that all creatures on the planet use the same (or a very similar) replication system which produces extremely high-fidelity copying of long-lasting replicators.

I suggest that the same process is now going on with memes, except that it is still in its infancy. As Dawkins put it, the new replicator is ‘still drifting clumsily about in its primeval soup’ (Dawkins 1976, p. 192). That soup is the soup of human culture, human artefacts, and human-made copying systems. You and I are living during the stage at which the replication machinery for the new replicator is still evolving, and has not yet settled down to anything like a stable form. The replication machinery includes all those meme-copying devices that fill my home, from pens and books to computers and hi-fi. Looking at it this way we can see all sorts of critical inventions of human culture as phases in the evolution of meme replication. I have already explained how treating language this way provides a new theory of the origins of language. I want now to go on from spoken language itself to the invention of writing, and then to modern information-processing technology. As before, we should expect the evolutionary process to involve increases in the fidelity, fecundity, and longevity of the replicators.

## Writing

Writing is obviously a useful step for memes because it increases the longevity of language. We have already seen how language itself increases the fecundity and fidelity of copyable sounds; the problem was longevity. Stories told using language can be remembered in human brains but, that aside, the sounds of language are necessarily ephemeral. Writing is the first step towards creating long-lived language.

No one knows how many times writing was independently invented from scratch, but the task is formidable. To start from scratch means making a large number of decisions about how to divide up speech and how to organise the marks that are going to stand for that speech. The Sumerians of Mesopotamia invented writing about five thousand years ago; the Mexican Indians some time before 600 BC; and Egyptian and Chinese systems may also have arisen independently. Sumerian cuneiform began, like many writing systems, as an accounting system representing sheep and grain. It started with clay tokens and gradually evolved into a system of marks on clay tablets, with conventions about making the marks in order from left to right and top to bottom. Other systems, naturally enough, use different conventions. From a memetic point of view we can imagine lots of people trying out different ways of using the marks and some ways being copied more than others. This selective copying is memetic evolution at work, and the result is better and better writing systems.
Many writing systems have taken a starting point from other systems, or even just borrowed the idea of writing itself. In 1820, a Cherokee Indian called Sequoyah observed that Europeans made marks on paper and went on to devise a system for writing down the Cherokee language. Although he was illiterate and knew no English, his observations were enough for him to devise a writing system so successful that Cherokees were soon writing, reading, and printing their own books and newspapers (Diamond 1997).

I have suggested that human consciousness is not the driving force behind the creation of language (or anything else for that matter) and Sequoyah looks like the ideal case to prove me wrong. In fact, I chose him as a perfect opportunity to explain what I mean. Sequoyah was presumably as conscious as any human being. In discussions about creativity people often assume that consciousness is somehow responsible for creativity, but this view meets with serious problems

as soon as you try to imagine exactly what it means. You are almost forced into adopting a dualist position, with consciousness as something separate from the brain, that magically leaps in and invents things. A more common view in science is to ignore consciousness and treat creativity as a product of the intelligence and ability of the individual concerned – ultimately taking the process back to brain mechanisms. This escapes from the dualist trap but leaves out the importance of all the ideas already available in the creator's environment. The memetic view includes all this. What I am proposing is this.

Human brains and minds are a combined product of genes and memes. As Dennett (1991, p. 207) puts it 'a human mind is itself an artefact created when memes restructure a human brain in order to make it a better habitat for memes'. In Sequoyah's case he must have had an exceptional brain, with exceptional determination and motivation, and he happened to come across a writing system that was already available at a time when his own people were in a position to take up his ideas and use them. Sequoyah's thinking was an essential part of the process, but was itself created out of the interplay between memes and genes. All this is a wonderful example of replicators creating design out of nowhere. As ever, there is really no designer other than the evolutionary process.
There are basically three strategies for writing systems. Signs can be used to stand for whole words, for syllables, or for just single sounds. The differences are important for the memes each syllable will be able to transmit. A system based on whole words is clumsy because there are so many words. Every time a new word is invented a new sign has to be created as well. At the other extreme systems using signs for single sounds can use few signs and combine them in many different ways, such as the alphabet of twenty-six letters in which this book is written. The cognitive load placed on the brains of people using the system varies in the same way. It is relatively easy to learn twenty-six letters and their sounds, although even this typically takes schoolchildren many months or even years of work. But learning Japanese kanji takes much longer, and unless you know two or three thousand of them you cannot read a Japanese newspaper.

For many reasons, writing systems based on sounds can convey more memes for less effort, and therefore are likely to win out in competition with other systems. Of course, the competition is not straightforward. The historical process by which writing systems are created means that there are all sorts of

quirks, oddities, and arbitrary conventions which, once learned by sufficient numbers of people, attain some kind of stability. In biological evolution an important principle is that evolution always builds on what it has available at the time. There is no evolutionary God who can look at the design of the eye and say 'it would be better if we got rid of this bit and started again'. There is no starting again. The same applies to the design of writing systems. They evolve gradually from wherever they have got to at any point. So, the alphabet of twenty-six letters is far from the ideal that a memetic God would create, but it is better than many other systems, and therefore, when direct competition arises, tends to win. Many languages, like Turkish for example, have changed over from more cumbersome systems to the Roman one. Many languages use variations on the system, adding umlauts or circumflexes, diphthongs, or even new letters. We have yet to see whether the economic and cultural power of Japan is enough to ensure the survival of its complicated writing system in a world in which the transmission of memes is everything, and English written with the Roman alphabet is dominant. For memetic reasons I suspect it will not be.

A similar argument applies to numerical systems. Arithmetic is formidably hard using Roman numerals but easy with any system that relies on the position of a numeral, like the Arabic system that we, and most of the world, now use. This drive towards uniformity is interesting, and is stronger than was the case for the evolution of language. In the case of writing, the invention of a new system is so difficult that borrowing one from elsewhere is more common, and novel systems are at a disadvantage. Once an adequate system has begun to evolve it has a natural advantage, in spite of any shortcomings due to historical accident and arbitrary conventions. When just a few systems exist, the one that produces slightly more, slightly better or slightly longer-lived copies begins to fill the world with its products, and the products take the idea of that copying system 'with them. The result is pressure towards one copying system taking over entirely from all the others.

We are all too familiar with this process. The standard QWERTY keyboard was devised to prevent the letters sticking together in the earliest manual typewriters; it is far from the optimum for modern keyboards and yet is almost universally used. Once music could be recorded and stored, vinyl disks of just two sides and three rotation speeds captured the market, but have now mostly disappeared. Standard reel-to-reel tapes hung on for a while after the invention

of the much smaller cassette tape but then cassettes persisted in a single format until the compact disk appeared, and may or may not continue to survive alongside it. Whether they do or not should be predictable from memetic principles. The number of memes that can be crammed on to a CD is dramatically greater than the number on tape, and CD technology allows rapid random access. Therefore, once cheap CD copying devices become available, CDs will surely outnumber cassettes, carrying with them the memes for that copying mechanism. The number of compact disks in the world is now so huge, not to mention the number of factories legitimately making them, and the even larger number illegally copying them, that an enormous step forward in fidelity or fecundity of copying would be needed to oust the system for a new one. The same has happened with the format of computer disks.

Bearing in mind the dangers of comparing memes and genes, we can speculate that the same process works in both cases, producing a uniform high fidelity copying system capable of creating a potentially infinite number of products. The genes have settled down, for the most part, to an exquisitely high fidelity digital copying system based on DNA. The memes have not yet reached such a high-quality system and will probably not settle on one for a long time yet.

Returning to writing, I have described its evolution as a step towards greater longevity for memes based on language. That step opens the way for further steps in increasing fidelity and fecundity. Spelling can vary greatly, leading to ambiguity and low fidelity. Many languages began with optional spelling that gradually gave way to ‘correct’ ways of spelling every word, dictionaries that specifies the correct spelling and, more recently, spell checkers that enforce the rules in electronically stored text.

Fecundity is obviously limited when writing is slow and difficult, as it was for marking clay or making clay tokens to stand for words. For most of its history, writing was a skill confined to a few specially trained scribes. This made political sense because of the power it gave to rulers. They alone could command scribes to keep records of barter, financial transactions, and taxes, or to maintain holy texts for the justification of oppression and war. In any case, the early writing systems were only capable of recording limited kinds of information. It took political and economic changes, as well as changes in

writing itself, before writing could be used for poetry, novels, personal letters and recording history. Widespread literacy came later with its dramatic increase in the number of memes stored and passed on as marks on paper.

The printing press was a critical step for both fecundity and fidelity. Up until the fifteenth century, all copying of texts in Europe was done by scribes, often monks who spent a large proportion of their time copying and illuminating religious works. The work was slow and they made many errors. These errors are now of great interest to historians tracking the history of texts, but they certainly did not help fidelity. The time taken meant that few copies could be made, and books were an expensive commodity for only the richest and most powerful people. This restricted the ideas in books to those for which there was financial backing – that is, ideas maintaining political, economic, and religious power. Once books were cheaply available the kinds of memes contained in them could proliferate and change. Written material is no longer confined to lists of taxes and religious tracts, but is constrained by quite different market forces.

The memes took a great step forward when they got into books.
Memes in books provide a good example of a selection system at work. In this system, the replicators are the memes: the ideas, stories, theories or instructions conveyed in the printed words. These either get copied or not, and their content affects the likelihood of their being copied. The copying machinery is the publishing houses, printing presses, and factories in which the books are made. The selective environment is the minds of authors in which memes compete to get into the final text, a world full of bookshops that stock the books or not, the book reviewers and magazines that publicise the books or not, and the people who buy and read them and recommend them to their friends – or not. We humans are, obviously, critical to the whole process. However, our creative role is not that of an independent designer conjuring ideas from nowhere. Rather we are the copying machines, and parts of the selective environment, in a vast evolutionary process driven by the competition between memes.

As I write this book I think of my mind as a battleground of ideas. There are far more of them than can possibly find their way on to the final printed pages. ‘I’ am not an independent conscious entity creating the ideas out of nowhere.

Rather, this brain has picked up millions of memes from all its education, reading, and long hours of thinking, and they are all fermenting in there as the fingers type. After this internal selective process is over and the manuscript is sent off there will be more selection, by the readers chosen by the publisher, and ultimately by the reviewers, bookshops and readers out there in the world. Whether the book sells a few hundred copies or a few hundred thousand copies will depend entirely on that selective process,

## Communications

Railways, roads and ships may not seem to be directly concerned sprite memetic copying, but they play a role in speeding up the process of memetic competition. They carry to distant places the letters in which memes are written and the goods and people who convey ideas. They also increase the number of people who are in contact with each other which provides a larger and more varied meme pool. Just as biological evolution produces more species on large landmasses than on small islands, so memetic evolution produces more developments when more people are joined together into a memetic system, Roads, railways, and airlines connect larger and larger numbers of people together, just as common languages and writing systems do.

In a 1901 classic, *Cosmic Consciousness*, the mystic Richard Bucke predicted that with the invention of 'aerial navigation', cities would no longer be needed and rich people would live in beautiful places, evenly spread out across the globe. In fact, cities have increased dramatically in population and rural depopulation is the norm. Why is this? A memetic answer, though a slight digression from copying technology, takes a familiar form. People who live in cities meet more people and therefore pick up and pass on more memes than people who live in isolated places. Among these memes are behaviours that are only possible (or are much easier) in cities – eating out and going to pubs, going to cinemas, theatres, museums and art galleries, visiting friends at a moment's notice, or having a high-powered job at the centre of the action. The citydweller not only picks up these memes but meets other people who also have them. Once these habits are picked up they are hard to drop.

Meanwhile, the people who live in the country meet fewer people, and do not have the opportunities to pick up the habits of exciting city life – unless they go to the city, in which case they may be lured by all the memes they find there. There is a critical imbalance operating here. When city-dwellers go to the country they meet few rural dwellers because they are widely spread out, and pick up few rural memes because few exist; but when country folk go to the city they meet lots and lots of city people and lots of new ideas. The consequence is memetic pressure for city-dwelling.

You may object that people make their choices about where to live either out of economic necessity or by freely choosing the life they know will make them happier. But is this really so? Economic necessity is often not a question of food and clothes for the family, but of buying televisions and cars and all the other trappings of a meme-rich life. The more we are exposed to memes the more we seem to acquire a hunger for them that is rarely satisfied. And happiness is very hard to judge. We may think that having a more exciting life, closer to the centre of the action, will make us happier, but we may be wrong. I suggest that we are, to a far greater extent than we would like to believe, driven to our choices by the pressure of memes.

This memetic argument suggests that there will be pressure for people to live in vast cities whenever the following conditions obtain: first, that there is enough communication between the country and the cities to set up the imbalance, and second, that people's main form of communication is still face to face, or via cheap local phone calls. If memetic transfer were truly independent of distance then the demographic pressures would change.

The telegraph and telephone, radio and television, are all steps towards spreading memes more effectively. They increase the fecundity of the copying process, and the distance over which it operates. People have often been unable to predict how such inventions would actually be used and which would last and which not, but from a memetic point of view prediction should be relatively easy. Anything with higher fidelity, fecundity, and longevity than its rivals should be successful. From the first electric telegraph, in 1838, to the telex machine and fax, fidelity and fecundity have gone on increasing – opening up new niches for further development along the way.

The telephone was bound to be a success. People are genetically evolved to chat and gossip (Dunbar 1996), and want to exchange news and views, creating lots of memes in the process. They can spread the memes by letters which take minutes or hours to write, and days to arrive, or they can ring each other up. People who use the phone will get more ideas spread simply because it is quicker, and those ideas include the idea of using the phone. Mobile phones have progressed very rapidly from being an executive luxury to being indispensable to every doctor, plumber, and aspiring teenager.

Letters will win out only when there is a need for longevity over fecundity. Fax machines combine the fidelity and longevity of writing with the speed (and hence fecundity) of the telephone. Photocopiers were a fantastic step for fecundity. Interestingly, people keep predicting the end of books. When radio came along predictions were made that no one would read any more. The same was proclaimed with the advent of television and then personal computers. In fact, books of a TV series can sell millions, and bookshops are selling more, not fewer, books than ever. Perhaps this is because memes can take different routes to success, just as genes do with their alternative strategies under r-selection and K-selection (p. 100). Electronic-mail messages go for high fecundity, low fidelity, and low longevity (people send out lots, do not bother to write carefully or correct the mistakes, and throw them away). Letters go for low fecundity, high fidelity, and high longevity (people write fewer letters, construct them carefully and politely, and often keep them). Books are high on all three.

All this makes a lot more sense if you look at the process as memetic competition. Any copying process that produces a successful combination of high-fidelity, long-lasting copies of memes will spread more memes and, in the process, spread itself. As this process continues more memes spread faster and faster. Note that the consequence of this is a headache for humans. Competition in business, publishing, the arts and science all depends on the transfer of memes. As memetic transfer speeds up so the competition speeds up, and people without the latest technology fail in that competition. We are driven by the latest technology to have to read all those books today, send that fax now, or be on the end of a phone line to Japan at three in the morning. We may think all this progress is designed for our own happiness, and indeed we may sometimes very much enjoy our meme-rich lives, but the real driving force behind it all is the interest of the memes.

## From copy-the-product to copy-the-instruction

So far, I have talked about increasing fidelity in rather general terms. I want now to be more specific and apply two further principles to how copying systems increase their fidelity. The first is the switch from analogue to digital systems and the second the switch from copy-the-product to copy-the-instruction. Digitising information is a good way to increase fidelity because it reduces errors in storage and transmission (p. 58). Language includes discrete words and is therefore more digital than other communications such as cries, howls, and calls. Writing extends the digitisation by committing certain sounds to certain letters, enforcing standard spellings and, above all, by allowing the vagaries of handwriting to be ignored by anyone who has learnt an alphabet. The ability of humans to read scrawly idiosyncratic handwriting is amazing, and computers are still bad at it. We are essentially able to interpret a wide variety of scribbles as being the letter 'p' or the letter 'a', thus creating a digital signal out of an analogue one. The same has been true of sound-receiving technology as it switched from grooves in disks or analogue magnetic signals stored on tape, to digital recording and storage. Indeed, it was the advent of digital sound recording that made it obvious that digital is better than analogue. Many radio stations have already changed over to entirely digital systems with a significant improvement in qualify. The copying of DNA has built-in error-correction mechanisms that far exceed anything the memes have yet created.

The second step is to copy the instructions rather than the product. I previously gave the example of a recipe for soup. It may be possible for a cook to taste the soup and copy it, but the copy is likely to be better if he works from a recipe. Why? The general principle is that following recipes is not a reversible process, whether we are talking about the genetic instructions for making a body, or the recipe for a cake (Dawkins 1982). Follow the genetic instructions in the right way and under the right conditions and you get a body, but you cannot take the body and follow the instructions back to arrive at a person's genome. The same is true of the soup. Of course you can try, but errors are bound to creep into the reverse engineering required to copy the product. You have to work out how it was done, and then do it yourself. If copies of copies are made the errors are compounded, and any good tricks invested in the original product are soon lost. It is far better to have clear instructions to follow.

The invention of writing makes possible all sorts of steps in this direction. Recipes for food are only one example, others are car-maintenance manuals, instructions on how to get to the party, user manuals for hi-fi systems or gas ovens, instructions for building model airplanes or decorating your house in the latest fashionable styles. In these, and many other cases, you may see a product or action and guess at how it was made, but verbal or rewritten instructions are a great help.

Copying written instructions is also far more secure. Writing is digital and highly redundant so that errors in spelling or syntax, or degradation by photocopying, are routinely ignored in passing on recipes or instructions. The same instruction can be copied to millions of people, as many computer manuals have been, and each person receives the same information. The booklet can be passed on to reader after reader without losing any detail.

I am returning to this principle because it has been so important in the computer revolution. Computer programs are instructions. They work on the basis of copy-the-instruction not copy-the-product. Take a familiar suite of programs such as the word processor I am using to write this book – Word 6.0. Word has evolved gradually through several stages and there are now millions of copies of its various versions in the world living inside millions of PCs in offices and houses. Some people buy them on disks or CDs, others copy them (legally or not) from each other. When installed the programs all do the same things. They put up letters on a screen, move text around when commanded by the user, send data to printers, and so on. No human being, from watching the word processor at work or seeing the documents it creates, could reconstruct the machine code on which it is based. The fantastic success of the memes inside Word 6.0 is due not only to its usefulness to the humans who use it, but to the digital copying machinery on which it operates and the fact that it is instructions and not products that are copied. These memes, or some of them, will outlive Word 6.0. If Word 8 or 9 is made it will doubtless reuse much of the code that formed earlier versions.

Note that the billions of products created by these word processors are not copied in the same way as the memes inside the word processor itself. But nor are they irrelevant to the copying process. If people were not happy with the

program, and could not easily write all their letters, articles and books with it, then Word 6.0 would not be copied at all. It is the quality and quantity of the documents created that determines the success of the word processor they were created with. We can now see that these documents play for the memes, a similar role to that played by organisms for the genes. In this sense they are a vehicle, except that they do not carry the replicators around inside them. The documents themselves may disappear, but their existence determines which of the instructions for making them are copied and which are not. And potentially these instructions can be copied on forever, just as genes can.

Many meme-copying steps have gone into the creation of the computers on which all this depends. They include the invention of language, its increased longevity by writing, increased communication between people by the building of roads and railways, the invention of telephones and televisions, the invention of digital computers, programming languages, digital storage devices, and finally the creation of user packages such as word processors, statistical packages, spreadsheets, and databases, which consist of memeplexes whose vehicles are the documents they make possible. We may expect this process to continue with the creation of more and more computer-based instructions whose operations are inscrutable to their users but whose products determine whether they are replicated or not.

Note that this evolutionary process has made memetic-copying mechanisms more similar to genetic ones. One of the great worries for memetics was the accusation that memes are passed on by Lamarckian 'inheritance of acquired characteristics'. We can now see that with further developments of meme copying technology the tendency is, just as it presumably was for genes, towards a non-Lamarckian mechanism – that is, copy-the-instruction not copy-the-product. The precise way it is done will always be different for memes and genes but the basic evolutionary principles are the same. The competition between replicators forces the invention of better and better systems for copying those replicators. The best systems are digital, have effective error-correction mechanisms, and copy the instructions for making the products, rather than the products themselves.

## Caught in the web

In 1989, the World Wide Web was invented. The Internet had already been expanding for many years, and what had begun as a small scheme linking a few government scientists, rapidly became a worldwide system through which anyone with a computer and modem could fetch stored information from all around the world. This was a great step for the memes. Memes can now be stored on the hard disk of a computer in, say, Melbourne, and at any time of day or night be copied almost without errors over phone lines or satellite connections to another computer in London, Florence, Chicago, or Tokyo, using the energy resources of countless human beings along the way.

These memes can be used to create other products (say school projects or business plans). They can be saved on disks at the new location, or to save space just the link be saved and the information called up again whenever needed. This last fact mirrors an interesting trick used by the human visual system. The visual world is so complex that storing even tiny fractions of the changing image would overwhelm even the vast storage system of the human brain. Instead, the brain throws away most of the information and relies on our ability simply to look again. We may have the impression that when we look out of the window we have a beautifully rich visual image, but in fact all our brains are holding is a little piece of the central image, a very rough sketch of the rest, and the ability to respond quickly to change and look again when necessary (Blackmore et al. 1995). In the same way, when using the Net we can mark information we might want again without actually keeping it on our own computers. The memes stay where they always were, in Sydney or Rome; and we just have a quick route for getting them again.

Use of the World Wide Web is free. This may change, but at the end of the twentieth century you pay only for the computer and phone lines that connect you to the system. Out there in cyberspace lie all the stories, and pictures, and programs, and games, that millions of people have lovingly put onto their Web sites, creating a virtual world of digital information. There are multi-user domains, or MUDS, that are imaginary places that people have constructed for others to come and play in. For some people these virtual worlds are more real than ordinary life (Turkle 1995). There are controls on who can enter a MUD but they are not financial controls. This is odd if you think of the Internet as

something that humans created for their own benefit, because you might expect them to pay for it. It makes more sense when you think of the memes as having created the Web to aid their own replication, and competing with each other to get your attention. If memes can get copied they will, and the Internet copies a lot of memes.

Does the Net need us? Yes, at the moment it does, though not necessarily for ever. We made the hardware and software on which it depends and we need to keep maintaining it, or the copying system will collapse. More importantly, our biologically evolved nature still drives, to a very large extent, which memes are successful. They are, naturally enough, those to do with sex, food, and fighting. The most common topic for searches on the World Wide Web is sex. MUDS allow people to take on invented identities and engage in meeting, chatting up, and having virtual sex with people whose location and even biological sex they may not know. The vast majority of computer games are based on killing and warfare. Any memes which can get into or tag along with such memeplexes are more likely to succeed. In this sense, the Internet still needs us, and is driven by human genes as well as memes.

However, many changes lie ahead. Already there are free-floating programs which move around in cyberspace, called bots (short for robotic programs). The way forward in artificial intelligence seems to be to build small and stupid units that together do clever things. We can imagine the Net becoming full of such autonomous stupid creatures that run about doing useful jobs. For example, as the Net increases in size and complexity, which memetic principles dictate that it must, there will be increasing problems of traffic flow and control. One idea is to create little programs based on insects laying chemical trails, that move around providing information about traffic flow on different routes. Others might perform error-correction tasks or censorship duties. At the moment, the only viruses or parasites are ones deliberately created by malicious (or just mischievous) human beings, but could bots mutate into viruses and start clogging up the system? Certainly, copying errors happen in any system and occasionally they lead to a product that proliferates. General evolutionary principles suggest that this may occur if the fantastic copying and storage system of the Net is maintained for long enough.

Other programs simulate people; they can carry on conversations and do things like psychic readings, or take part in games. There are 'chatterbots' with whom you can converse when you get lonely. In multi-user games people have been fooled by bots claiming to be real people. In a large system over a long time such bots could presumably mutate into more and more efficient 'people'. Many people seem to assume that because we built the machinery on which the net runs, we are in control of it. This is clearly not so. British Telecom can no longer understand its own telephone network, and the whole worldwide system looks set to become bigger and more complex still. Indeed, if the memetic analysis I have given here is correct, then so long as human beings maintain the infrastructure, the system will proliferate out of anyone or anything's control – like a vast natural ecosystem.

The same applies to robots. At the moment, they mostly carry out simple tasks under human control, but memetics raises the following interesting possibility. For robots to become like humans – in other words, to have humanlike artificial intelligence and artificial consciousness – they would need to have memes. Rather than being programmed to do specific tasks or even to learn from their environment as some already can, they would have to be given the ability to imitate. If they could imitate the actions of people or other robots, then robot memes would begin to spread from one to another, and a new kind of memetic evolution take off, perhaps inventing new kinds of language and communication. The robot memes would drive the robots to new activities, giving rise to motivations that we could only guess at. We humans might not be capable of imitating everything the new robots did and so we might be quite excluded from their kind of cultural evolution. We would certainly not be in control of it.

All this raises interesting, and perhaps frightening, questions about the nature of human control and human identity. In any case, memetics raises those questions from its very foundations. I have carefully avoided them so far but the time has come to ask the difficult ones. Who am I, and what am I here for?