Interview

Harry Barrow

Interviewed for Expert Systems by Kenneth Owen

Harry Barrow

Introduction

Harry Barrow returned to Britain from the United States in 1988 to be Professor of Artificial Intelligence in the School of Cognitive and Computing Sciences at the University of Sussex. After six years as a research fellow at Edinburgh University, working on computer vision and robotics, he moved in 1975 to Menlo Park, California to pursue these topics at Stanford Research Institute (now SRI International). In 1980 he helped to set up the Fairchild Laboratory for Artificial Intelligence Research (FLAIR), later renamed Schlumberger Palo Alto Laboratory, and remained there until his return to the UK. At Sussex, his main research interest is in neural networks, with particular emphasis on vision and on integration with conventional-style AI.

Kenneth Owen: *Neural networks have attracted a lot of hype, just as expert systems did a few years ago. Do you see a danger there?*

Harry Barrow: There's always a danger, in fact there are always two dangers. One is that the hype will be believed, and that much more credence will be given to the technology than is appropriate. The other danger is that it will be completely disbelieved, because the hype is recognised, and in that case you throw the baby out with the bathwater.

Certainly it's true that expert systems, despite their initial hype, have a great deal of value, particularly commercial value. They form the basis for doing a lot of interesting things. They also form the basis for exploring reasoning processes and so forth. The same is true of neural networks. Expert systems as we currently understand them cannot do everything easily and well — and the same is true of neural networks, in the present state of understanding. We should recognise that the neural networks that are currently being experimented with are extremely simple when you compare them with real neural networks. I don't mean just that they don't have many neurons, though we have millions and billions of neurons in our heads. I mean that the architecture of the networks is much, much simpler than you find in real networks. Real neural networks are very rich and very complicated.

KO: *In your invited talk at the 1989 AISB conference you said there were two opposing camps: neural network proponents who believed that AI was moribund; and AI people who believed that neural networks were a snare and a delusion. Are neural networks basically different from expert systems?*

HB: At the moment, apparently, yes, but I think in the long term the distinction will become less and less. The emphasis in much of AI has been on symbols — manipulating them, processing them, reasoning about them. However, intelligence is not simply reasoning of the logical variety. Intelligence involves a strong perceptual component. For example, I've worked for many years in computer vision and it's an interesting field because you get to work in just about everything in AI. At the front end you work on signal processing, image processing and calculations based on the physics of the imaging situation to try and recover information about the scene. Then you get to make the transition from that sort of numerical representation to a symbolic representation and you start talking about features in the image. Gradually you make the progression up towards thinking about objects, not only their geometry but also their more general attributes and the ways they behave. So vision spans this tremendous range of representations and styles of processing.

If you think of an activity such as how to prove theorems, very often there is a perceptual component to part of that process. When you're sitting down staring at a problem, scratching your head and trying to find a solution, very often you perceive the structure and something about how to solve the problem — not in a step-by-step logical way, as the proof is written, but more in an analogue way, perhaps. Once you've perceived how to go about finding the solution, you start filling in the details, and eventually you finish up with a proof which is a logical sequence of inferences, step by step, and each step is completely valid, although they were not generated in that order.

But there's a very big difference between this watertight proof and the process by which you arrived at it. This latter process can be very much seat-of-the-pants, empirical, heuristic. The sort of processes that underlie how you come up with the idea for a proof are very similar to the sort of processes you find in perception. So there's a lot of untapped scope for perceptual components in what appears to be high-level logical reasoning.

This is apparently drifting away from neural networks, but not really. The high-level reasoning processes that go on in our heads are actually implemented by low-level neural network hardware — real live neurons. We have some limited understanding of the sorts of computation these things can and do perform, in real animals and in hypothetical models. We have a reasonable understanding of how they implement some fast perceptual processing. We don't yet have a very good understanding of how high-level reasoning can be implemented on such a neural network system, for a variety of reasons.

So I don't believe there is really a clearcut difference between neural networks and symbolic reasoning in AI. Any long-term understanding of natural intelligence is going to involve an understanding of the neural substrate, how it's implemented in animals and people, and we're going to understand the relationship between those two levels of description of our mental processes.

KO: *You're saying AI and neural networks are coming together?*

HB: They will come together eventually. There's always been an overlap. There's always been a spectrum of interests of people in AI. On the one side there have always been people who are fundamentally interested in how people behave intelligently. They're interested in unravelling that problem and building working models of the way that people do it. On the other hand there have always been people who are interested in practical engineering, to make machines that can perform the sort of intellectual functions that humans perform — automating intellectual activity as distinct from muscle activity.

Both of these types of research are going on in parallel, and they overlap in the middle, because if you're working on the engineering side then you are interested in things like how you might implement certain processes on computers at all, what sort of fundamental properties and limitations these processes have, and so on. And on the other hand, if you're interested in the psychological aspects, you're trying to build models that capture human abilities, warts and all, so you can understand them and their limitations.

KO: *For the engineering people, are you saying that knowing how the human mind does it is a useful input?*

HB: Yes. How the human mind does it might give one some inspiration, but maybe sometimes an engineer will think he can do better than that. And when you have a particular application, if you can really do better that's a good thing to do.

KO: *So the psychology and cognitive science should be of interest to expert system professionals?*

HB: Certainly the work on psychological cognitive modelling may well be of interest to expert system professionals, because we're developing a better understanding of human reasoning processes, how people have insights into problems and can see solutions. When an expert tells you how he performs some task, to some extent there's a rational reconstruction process going on. He can tell you about some of the standard methods and algorithms and formulae that he uses, but exactly how he decides which ones are appropriate and when and how to use them — there's a great deal of feel and intuition that goes into that process, a great deal of perception of the nature of the problem.

KO: *Is that a kind of knowledge that you can't acquire and express in a set of rules?*

HB: No, I wouldn't say that. But sometimes it is quite difficult for the expert to introspect about that process. When one has been performing a task for a long time and has become very skilled at it, then a lot of the activity becomes internalised, and built-in, and compiled-in, so it becomes very difficult to give a running commentary describing what you're doing. You have to use a variety of techniques to help to extract that information.

KO: *So, where do neural networks fit into the scheme of things?*

HB: They fit in in a couple of ways. The obvious way that everybody seems to think of when you talk about neural networks and expert systems is the notion of learning — the fact that you can present the network with some data, and you can tell it what conclusion you want it to come to, and you give it enough examples and you train the network and it eventually discovers for itself some way of getting the conclusions from the original data. That's obviously quite a handy tool for building systems.

It may not be the only tool. As I said, we don't understand yet how to implement these very high-level, sequential logical processes on neural networks. So it might indeed turn out that for a substantial part of the system you want to have some techniques other than neural networks. But for certain areas, for certain activities, you may well find that a neural network is an ideal tool. For example, when it's very difficult for an expert to articulate how he looks at a set of data and comes up with some inferences, it may be extremely difficult for him to articulate a set of rules that hang together coherently and consistently, and cover all the examples. But a neural network might be able to see those rule sets, to generate internally what corresponds to an appropriate set of rules that cover the examples that you've been given. That's one reason why they might be useful. Another is that neural networks are good at interpolating between data samples.

KO: *Are there cases where, to do a given task, you would choose either an expert system or a neural network?*

HB: Neural networks at the moment tend to be quite good at perceptual style tasks, which are precisely the ones that it is difficult to build expert systems for. What neural networks are not good at at the moment is reasoning that involves an iterated process. I'm not talking about simply stacking up some logical decisions, I'm talking about a computation that is inherently iterative and may involve an indeterminate number of steps. To make that a little bit clearer, SRI's Prospector expert system internally is based on an inference network, and the inference network is a collection of nodes each representing some particular rule. Each node receives inputs from some preceding nodes, and the inputs represent probabilities or likelihoods, and what each node tries to do is to compute from its inputs the likelihood that the hypothesis which that node represents is true.

And the Prospector system internally has very much the flavour of a neural network. This cascaded logical computation takes some weighted sum of the inputs and uses a Bayesian inference rule to decide what the output is going to be. But Prospector does not go through an iterative application of rules to the data. You present it with some input, and the data propagates through the structure (which you can imagine as being a static structure) to the answers.

On the other hand, there's a hoary old example that was studied by Minsky and Papert: you're looking at a black and white image, and you want to decide whether or not all the white points in the image are connected. It turns out that that's very simple to do with an iterative calculation, which just steps through the image array, but it's much harder to perform with a parallel computation than you might think. It does demonstrate that there are some computations which are very difficult to perform economically in a highly parallel way.

KO: *So the two are suitable for different things. Do you see different aspects of a task in the future being tackled by different tools, including both expert systems and neural networks?*

HB: Yes. I believe very firmly that each type of tool, the neural network and the rule-based system, has a wide range of uses and applicabilities. But the combination of the two, certainly in the immediate future, will be more powerful than either of them individually, and will probably be quite useful in many practical applications.

KO: *What are the main application areas for neural networks at present — speech and vision?*

HB: Certainly they include speech and vision. There are networks which can recognise isolated words; for example, Teuvo Kohonen in Finland has developed a hardware implementation of a particular style of neural network for this. There are some commercial systems for visually recognising objects or faults in objects on production lines. There are companies that are marketing such hardware. Neural networks have been used for various types of control, such as the control of robot arms.

There has been some work on using neural networks in process control. I have heard of a large chemicals company in the United States that has been able to apply neural networks to something like 20 different internal applications, and they're so pleased with the results that they've set themselves up quite a sizeable group to work on neural networks and continue that development.

KO: *Is the increasing general interest in neural networks perceived as a threat by some AI traditionalists?*

HB: Yes. In some parts of the AI community there are people who are very concerned about neural networks. They see them as a threat for a variety of reasons, the main reason being as a competitor for resources. And they're very concerned about the hype, and that resources may be taken away from productive work in classical AI and put into this possibly charlatan field of neural networks. Equally, there's a strong feeling in some quarters of the neural network community that the shoe is on the other foot. They blame the AI community (and in particular the work by Minsky and Papert) for the cutting-off of funding in the early days of neural network research.

Well, it's not quite as simple as that. If you look at what went on in the late 1950s and early 1960s you find a great deal of enthusiasm and a good many overinflated claims. For example, Rosenblatt, talking of his Perceptron, said at one point 'For the first time we have a machine which is capable of having original ideas', which was going a little bit off the deep end.

There was a lot of excitement, a lot of promises were made, but it became quite difficult to follow this through. After about ten years of work on early Perceptron-style approaches it was proving very difficult to scale up the applications and to apply them to more activities. So there was a feeling already that things weren't coming through in quite the way they were expected. However, AI did seem to be making progress, and perhaps Minsky and Papert's analysis of the limitations of Perceptrons was a significant contribution to reducing the funding for that kind of research.

KO: *In an interview last year (*Expert Systems, *April 1989), Ed Feigenbaum of Stanford contrasted the symbolic level of logic-based traditional AI with the sub-symbolic level of neural networks, which he described as grounded in the statistics of ensembles rather than logic, and went on to suggest that neural network research 'will have little impact on models of the central cognitive functions involving logical problem-solving, and in my opinion is not of much current interest to those who are working on expert systems'. Do you agree?*

HB: I can agree with much of what Ed says in his full comment, but I would not make too firm a distinction between the sub-symbolic and symbolic levels of computation. You can view any computation at a very abstract level as a highly symbolic logical process. You could view what goes on in a theorem-proving program that way, for example.

But you can equally well view the theorem-proving program as operations on individual bits with logic gates inside the computer. That may not be a terribly good way to view computing if you are only interested in the theorem-proving aspects of what is going on. But if you're interested in implementing theorem provers then you may well be interested in manipulating bits. So I think there really is a continuum; there are many levels of viewing the computation that goes on.

One thing that I find exciting about neural networks is that we're looking at a different style of computation than we've looked at in the past. In computer science, the parallel processing that we're used to is concerned with processes, particularly a few large-scale processes, communication between processes, passing complex messages, deadlock, and things like that. In electronic engineering, the view we have of parallelism is large numbers of logic gates operating synchronously, each computing a simple logical function of a few inputs.

Neural networks are neither of these things. They are performing parallel computations on a large scale. Each element is producing particularly simple results and is computing a fairly simple function of its inputs. But its behaviour can be quite interesting by virtue of the large number of inputs that it's receiving, which is something we haven't really explored much before.

We still don't understand all the ramifications of that, but I think it's going to be very worth while exploring this new style of computation. When we have a better feel for it, we will understand better how to build a higher level of computation.

KO: *So you wouldn't agree with Ed that neural networks are not of interest to expert systems people?*

HB: No. I disagree with that because I believe they have been demonstrated already to be of interest. I believe they have already been demonstrated as being useful, certainly in some areas of expert system applications. There has been some exploration of applications to financial problems, for example, assessing creditworthiness, where the inputs to the network are responses to questions on a questionnaire, and the output is some indication of creditworthiness or various other attributes. And that's very much an expert system sort of application.

KO: *How would you sum up your assessment of these two styles of computing?*

HB: In the very short term we do have some applications that can be completely performed by neural networks. We have applications where parts of the computation can be usefully performed by neural networks, and the ability to train them is actually quite valuable. In the longer term we shall see combinations of the two styles of approach being used more frequently, and in the very long term we shall understand how they really fit together, in an explanation of how the human brain works.