Artificial Intelligence Review (1990) **4**, 147–155

VIEWPOINT

'The strange story of the Perceptron'

Richard S. Forsyth
Department of Psychology, University of Nottingham, UK

Abstract. A number of bald assertions about the history and future of neural computing are made. An alternative research agenda is proposed which is a plea for pragmatism (in two parts).

'In my opinion this problem of making a large memory available at reasonably short notice is much more important than that of doing operations such as multiplication at high speed. Speed is necessary if the machine is to work fast enough for [it] to be commercially valuable, but a large storage is necessary if it is to be capable of anything more than rather trivial operations. The storage capacity is therefore the more fundamental requirement.'

Alan M. Turing (1947)

The Renaissance of Connectionism

Connectionism (also known as neuro-computing) is not new. As soon as the digital computer was invented, the leading figures in computing research— including Von Neumann, Shannon and McCulloch— began looking at the nervous system as a source of inspiration for the design of intelligent computer systems (Von Neumann, 1958). However, the idea that machine intelligence could be achieved by mimicking the brain became discredited in the mid-1960s and has only recently been revived.

The story of the burial and rebirth of neuro-computing is, in broad outline, well known to most people in the field of Artificial Intelligence (AI). There is not room here to recount it in the detail which it deserves, except to point out what is possibly the most curious aspect of this whole curious affair; namely the fact that a practical solution to the problem that effectively killed off the first phase of neural-net research was known at the time of its fall from grace (Widrow & Hoff, 1960; Widrow, 1987) and others were discovered soon afterwards (Albus, 1971, 1981; Aleksander & Stonham, 1979).

By the time that Minksy & Papert (1969) wrote the epitaph for the first phase of neuro-computing, researchers were already turning away from the connectionist paradigm. Minksy & Papert's book confirmed this trend by showing that a particular device, the Mark I Perceptron (Rosenblatt, 1958, 1962), could not learn to solve the exclusive-OR problem or other parity problems. Thus they erected intellectual 'KEEP OFF' signs around an entire field of computing research.

Their primary concern was whether any training rule for simulated neural networks could be proved to converge on an optimum set of weights in a finite number of trials. The dilemma they posed was essentially this:

> Convergence can be proved for single-layer neural nets, but single-layer networks cannot solve parity problems; multi-layered networks can solve parity problems, but there is no proof of convergence for them.

Not until the mid-1980s, when various workers independently discovered the *back-propagation* algorithm, and proved that it would converge with a multi-layered feedforward network, was this dilemma resolved. Yet in the meantime the few workers who had persevered with neuro-computing had found several different ways of training multi-layered networks, some more effective in practice than back-propagation.

Indeed one of the very earliest workers in this field, Oliver Selfridge, whose Pandemonium model (1955, 1959) antedated the Perceptron, had worked successfully with multi-layered networks from the outset, without worrying over much about convergence proofs.

Perhaps the best description of the situation during the 'wilderness years' (1965–1985) was given by Bernard Widrow in his invited speech to the IEEE in June 1987:

> The adapation rule, the golden rule, is to assign responsibility to the neuron or neurons that can most easily assume it. It turns out it works; we never published it. Why? We couldn't prove it; we were too embarrassed. We couldn't prove that it converged. It was an embarrassment. (Widrow, 1987)

In short, by 1963, Widrow & Hoff had extended their *delta rule* to what they called the MADALINE (Multiple ADALINEs), but they could only demonstrate empirically that it seemed to work. Once the research community had turned against neuro-computing, however, empirical demonstrations were no longer enough; and even when Albus (1971) proved that a simple Mark I Perceptron, modified to incorporate an expansion recoder, could be taught to solve the exclusive-OR problem using either Rosenblatt's or Widrow & Hoff's training procedure, it had no effect on the widespread belief among computer scientists that neuro-computing was something that had been tried and had failed.

It was not until

(a) a proof of convergence (for back-propagation) had been found; and
(b) a 'changing of the guard' had taken place at Defense Advance Research Projects Agency (DARPA).

that the resurrection of neuro-computing took place. Point (a) made neural-net computing respectable again; while point (b) made it fundable.

The result is that we are now awash with hype. Neural-net computing is back in fashion and the number of conferences, books, papers, products, programs and workshops bearing the terms 'neuro' or 'neural' in their title is already uncountable.

That, then, is a brief history of this subject. The question is: what policy should a nation like the UK adopt towards a technology with such a chequered past? The rest of this document attempts to provide answers to that question.

1 Substantive issues

1.1 *Don't forget about memory*

As the neuro-computing bandwagon is well under way, and has been gathering momentum for over two years, there is little to be gained by merely stimulating awareness, either in academia or in commerce. That will happen anyway, and to the extent that it is successful will suck in imports, both hardware and software, which have been developed in non-EC countries.

In fact there is a case to be made— though it should not be pressed too strongly— for 'counter-cyclical' funding. This would entail sticking with the IKBS paradigm, which over the last 18 months has lost its attractiveness (just at the point when serious large-scale knowledge-based systems are beginning to prove their worth). However, while IKBS research should not be summarily dropped as 'yesterday's fashion', it is unrealistic for a country which wishes to remain at or near the forefront of technological advance to ignore the neuro-computing boom altogether.

So my own proposal would be to concentrate on areas which are not being over-exploited already in the USA and Japan. This implies a preference for the *memory* rather than the *processing* function of neural networks.

It is a mistake to see neural networks (natural or simulated) only as processing devices: they are not just parallel distributed **processors**, they are also parallel distributed **memories** at the same time. To call the human brain a database engine is more accurate than calling it a supercomputer. Yet the majority of existing work on neuro-computing concentrates on the processing aspect. (See, for instance, excellent surveys by Rumelhart & McClelland, 1986; DARPA, 1988.)

Two workers who have stayed in this field through boom years and bust and whose work seems to me to have been systematically undervalued are James Albus and Igor Aleksander. Both in their different ways have explored the idea of the brain as a giant fuzzy look-up table. Table look-up (even the sophisticated kind exhibited by CMAC and WISARD) is not very glamorous, but it works; and it works by finding parallelism where it is already latent (in RAM) rather than trying to impose it somewhere else (in the CPU).

A general emphasis on distributed memory models (as opposed to distributed processing models) would therefore suit the UK as a central plank of any national research programme. There are two additional reasons for advocating such an orientation:

(1) Igor Aleksander, one of the world leaders in the field, is based at Imperial College, London, and has the nucleus of a 'centre of excellence' already in place;

(2) VLSI memories exist in vast quantities exhibiting massive parallelism which is currently being exploited in the most unimaginative possible way.

The second point may need further elaboration.

It is an informal axiom of computer science that one can trade speed for space. That is to say, a program can be made to run faster if it is allowed to use more memory. Alternatively, memory can in general be saved by using a more powerful processor. This trade-off is a general principle, and applies even with the most advanced computer architectures. Yet the vast bulk of research into novel computing devices has been concerned with parallelizing the processing component (which has turned out to be difficult) not with parallelizing memory (which is easier).

CMAC (Albus, 1981) and WISARD (Aleksander & Burnett, 1984) point towards ways of exploiting the hidden parallelism of distributed memory systems, while most of the rest of the world struggles to build parallelism into multi-processor systems—without conspicuous success. Here then is an opportunity for world-class research which, if handled carefully, could lead to the capture of an important market niche. It is my contention that this area should be singled out as a major and distinctive constituent of a national research programme in neuro-computing.

1.2 *Don't propagate back-propagation*

The positive emphasis outlined above on neural networks as memory devices will need to be balanced by a negative attitude towards an alternative class of neural-net systems.

It is all too predictable that the announcement of a national research programme in this field will bring forth dozens of proposals that amount to little more than the application of a 'connectionist cliché' to yet another problem domain. This connectionist cliché can be described as a network which

(1) has 2 or 3 internal layers;
(2) is fully interconnected between layers;
(3) operates in feedforward mode; and
(4) is trained using back-propagation.

Back-propagation systems have achieved some notable results (e.g. Sejnowski & Rosenberg, 1987); but we can do without endless variations on a single theme. There is no need to waste taxpayers' money on such things, as they will spring up of their own accord in plague proportions. On the contrary, to find projects that have a chance of advancing human knowledge in this area, it will be necessary to look for work that breaks away from this cliché in at least one respect. Specifically, research worth funding is likely to investigate neuro-computing architectures that

(1) have many layers or do away with rigid layering altogether;
(2) have neurologically plausible degrees of fan-out between neurons, probably randomly assigned and possibly alterable;
(3) allow feedback during operation as well as during error correction;
(4) are trained by a method attempts to improve on back-propagation (which is merely a gradient descent

technique and not a very efficient one).

In a nutshell, a good rule of thumb for picking projects that are likely to uncover new knowledge would be: Don't back back-propagation (except for teaching purposes).

2 Towards a research strategy

2.1 The Bletchley Park spirit

My personal preference is for a Churchillian style of research management:

> Bring together a dozen or so bright young graduates, mix in a few older and wiser heads, dump them in a group of Nissen huts out in the English countryside, give them a long list of insoluble problems requiring urgent attention, starve them of adequate funding, and stand well back to await an explosion of creative energy.

This method worked extremely well at Bletchley Park from 1939 to 1945, but was abandoned at the height of its success, and is no longer favoured by funding agencies.

Nevertheless the 'Bletchley Park spirit' can be fostered even in peacetime, provided that certain pitfalls are avoided. The following suggestions are intended to help in the avoidance of some of the more serious pitfalls.

2.2 Fuzzy objectives, and how to clarify them

Any research programme should begin with a clear set of objectives. In the present context it will be more important than usual to set clear goals, as neuro-computing systems can be used for such a variety of purposes.

In theory, artificial neural nets can compute any computable function, being formally equivalent to Turing machines (McCulloch, 1965). Knowing this is not very helpful. Even being more precise, and saying that neural-net systems can be taught input–output mappings of arbitrary complexity, does not get us much further. What we really want to know is what can be done better with neural-net systems than with alternative techniques, and why. If neuro-computing research is to be publicly funded, it should be with a view to plugging holes that exist in conventional computing i.e. to doing things with neural-net systems that cannot be done, or are done badly, with traditional techniques.

This gives us a shopping list of desirable end-products:

(1) neural-net systems doing tasks that cannot be done any other way;
(2) neural-net systems doing things more efficiently than alternative computing systems;
(3) neural-net systems integrating with more conventional systems to achieve results that would not otherwise be feasible;
(4) theoretical work giving principled reasons when and why neuro-computing methods should be chosen.

The objectives of the research programme should be based on this list.

An additional objective should be the promotion of syntheses between competing approaches. An example is the learning system of Whitley (1989) which uses a genetic algorithm to optimize the weightings in a neural network. This sort of cross-fertilization between methods that were previously seen as competitors is valuable in itself.

2.3 *Antidotes to arm waving*

Already there is plenty of excited talk about neural nets and the wide blue yonder. This needs no further encouragement. To keep researchers on the straight and narrow, there must be an emphasis on practical results.

What we need is a suite of benchmark problems, rather like the test suites that are used to validate compliance of PASCAL compilers with the international standard. It would be well worthwhile to set up a clearing-house for neuro-computing applications to maintain a list of test problems and datasets. The items on such a list could be culled from various sources including, ideally, a trawl round British industry seeking to identify computing bottlenecks that might be susceptible to a neuro-computing solution. Jobs that are already done well should not be entirely ignored either. If a company says 'we already catch 96% of component defects using statistical methods' that does not exclude the possibility that they could catch 99% with a better system.

Sets of recognized benchmarks tend to arise as part of the folklore of a discipline. An example is Lin & Kernigan's (1973) 318-node tour for testing the merits of travelling-salesman heuristics: it has become a *de facto* standard. But there is much to be said for formalizing the gathering of test problems. Once an initial set has been agreed, researchers can be invited to submit proposals for attacking one or more of these problems with the aim of finding solutions that are faster, more comprehensive or cheaper than competing methods.

Here a sporting metaphor may be helpful. The clearing-house could be given the responsibility of conducting a contest (the Euro-Neuro Computing Championships, or some such) and judging the performance of the contestants. Some serious and explicit thinking about

—allowances for different classes of machine;
—standardization of benchmark problems;
—scoring of speed versus accuracy;

and other difficulties brought into the open by such an enterprise would, in itself, justify the work involved. It might even be possible to draw up a 'tariff'— as in Olympic diving— based on the practical pay-off expected from solving each particular problem. This would aid in the evaluation of all projects, not just the 'competitors'.

2.4 *Evaluation of individual projects*

This brings us to the question of evaluating projects after their completion. Optimization requires an objective function; learning requires knowledge of results. Yet the evaluation of research projects tends to be done in a thoroughly unscientific manner.

The number of papers published in refereed journals by the project team is not a good measure, but it is frequently used for want of anything better. The definition of success for individual projects needs to be tightened up (along lines outlined in Section 2.3) and promulgated in advance to would-be participants. This implies facing up to the possibility of failure, which is often more instructive than success.

During the writing of this piece, a review of one of the deliverables of the Alvey program, GLIMPSE, happens to have crossed my desk. GLIMPSE is a knowledge-based front-end for the GLIM software package, which is a statistical system for generalized linear modelling. GLIM is not easy to use by statistically untrained people, so it was hoped that an expert front-end could make it more accessible to inexperienced users. This is what the reviewer, a professional statistician, had to say.

> I needed to refer to [the manual] constantly. This seems to defeat the main aim and characteristic of an Expert System, that it should be user friendly and hence aid the user through an analysis from beginning to end. . . . the terminology was intimidating and the explanatory messages from the system were unintelligible. (Al-Doori, 1989)

What went wrong?

We do not know, but it should be somebody's job to find out. Post-mortems are not for the squeamish, but they serve an important purpose.

2.5 Learning from past mistakes

What applies to individual projects also applies to the research programme in which they are embedded. The Alvey programme achieved much, but I am not aware that any systematic appraisal has ever been made of it's failure. Part of the reason is almost certainly a feeling that it would be tasteless to enquire too closely into the things that did not go according to plan.

But if that attitude prevails, there is little hope of improvement in future. Of course there will be failures in a programme embracing hundreds of different projects; but it is vital to face up to them. Unless both the success stories and the failures are analysed with the benefit of hindsight, and the results of such analysis published, it is inevitable that many mistakes will be repeated and many successes will not be copied.

2.6 Fat-cat funding, and how to avoid it

There was a certain amount of hand-wringing after the Alvey programme ended concerning the lack of participation by small firms; but no administrative machinery has been put in place to prevent a recurrence of the same problem.

What is the point of pumping public money into a company which could fund the entire programme ten times over but chooses not to risk investing its 'cash mountain' in leading-edge research? We might just as well pay Richard Branson to organize a mission to Mars in a hot-air balloon— at least then the taxpayer would get entertainment value.

Excessive take-up of funding by large companies is especially galling in view of the fact that most true scientific and technical innovations arise from the efforts of small teams or single individuals, usually working outside the corporate ethos. Yet it could be discouraged in a simple way: just earmark a percentage of available budget (20–25%) for groups composed of fewer than 8 people, working for organizations with annual turnovers below a specified limit, asking for less than £100,000, and proposing to use off-the-shelf equipment in innovative ways. (It is also most important that the paperwork needed for small grant applications be short and simple, and that decisions should be swift.)

A good long-range prediction is that even a small fraction of the total resources distributed under such conditions would produce more useful results than the rest of the programme put together.

2.7 Competitive research

Academic rivalry is a powerful motivating force which is seldom harnessed to constructive ends. The proposed neuro-computing initiative should be designed to promote it.

That is to say, teams from different places should be allowed, and even encouraged, to work on the same problem, using alternative paradigms. This would enable comparative quality standards to be established on the clearing-house problems. It would also gather badly needed data on how well different neuro-computing architectures tackle different kinds of task. More important, it would fill gaps in our knowledge about how neural-net techniques compare against

—statistical methods
—other kinds of machine learning
—knowledge-based systems
—conventional optimization techniques
—genetic algorithms
—simulated annealing

and other approaches to the same sort of problems.

This implies that some of the groups working under the auspices of the neuro-computing initiative would not be using neural-net methods at all, but would be part of a 'control group' for purposes of comparison.

2.8 Conclusions

A research programme into clearly defined areas within the field of neuro-computing is likely to be beneficial to the nation, but steps must be taken to ensure (1) that clear objectives are stated at the outset and (2) that unusually thorough evaluation is carried out (and published) afterwards; otherwise much of the potential benefit of such a research effort will be lost.

Finally, it goes almost without saying that such a programme should have a European dimension, even though it is primarily British. There already is an ESPRIT project under way in this area centred at AERE, Harwell (ANNIE: Applications of Neural Networks in Industry in Europe). It would be logical to treat ANNIE as a springboard for launching the new initiative.

References

Albus, J. (1971) A Theory of Cerebellar Function. *Mathematical Biosciences*, **10**, pp. 25–61.

Albus, J. (1981) *Brains, Behavior & Robotics*. Byte/McGraw-Hill, Peterborough, NH.

Al-Doori, M. (1989) A Review of GLIMPSE. *The Professional Statistician*, **8**(10), 21–22.

Aleksander, I. & Stonham, T. J. (1979) A Guide to Pattern-Recognition using Random-Access Memories. *IEEE Journal of Computers & Digital Technology*, **2**(1), 29–40.

Aleksander, I. & Burnett, P. (1984) *Reinventing Man*. Pelican Books, Middlesex.

DARPA/AFCEA (1988) *DARPA Neural Network Study*. AFCEA International Press, Fairfax, VA.

Lin, S. & Kernighan, B. (1973) An Effective Heuristic for the Travelling-Salesman Problem. *Operational Research*, **21**, 498–516.

McCulloch, W. (1965) *Embodiments of Mind*. MIT Press, Cambridge, MA.

Minsky, M. & Papert, S. (1969) *Perceptrons: an Introduction to Computational Geometry*. MIT Press, Boston.

Rosenblatt, F. (1958) The Perceptron, a Probabilistic Model for Information Storage and Organization in the Brain. *Psychological Review*, **65**, 386–404.

Rosenblatt, F. (1962) *Principles of Neurodynamics*. Spartan Books, NY.

Rumelhart, D. & McClelland, J. (eds) (1986) *Parallel Distributed Processing*, Vols 1 & 2. MIT Press, Cambridge, MA.

Sejnowksi, T. & Rosenberg, C. R. (1987) Parallel Networks that Learn to Pronounce English Text. *Complex Systems*, **1**, 145–168.

Selfridge, O. (1955) Pattern Recognition and Modern Computers. *Proceedings of Western Joint Computer Conference*, Los Angeles, March 1955.

Selfridge, O. (1959) *Pandemonium: a Paradigm for Learning: in Mechanization of Thought Processes*. HMSO, London.

Turing, A. (1947) Lecture to the London Mathematical Society, February 1947: in *Alan Turing, the Enigma of Intelligence*. (ed. A. Hodges) 1985, Unwin Hyman Ltd, London.

Von Neumann, J. (1958) *The Computer and the Brain*: Yale University Press, New Haven, CT.

Whitley, D. (1989) Applying Genetic Algorithms to Neural Network Learning. *Proceedings of the 7th SSAISB Conference*. (ed. A. Cohn) Pitman, London.

Widrow, B. & Hoff, M. (1960) Adaptive Switching Circuits. *Inst. Radio Engineers, Western Electronics Convention Record*, Part 4, pp. 96–104.

Widrow, B. (1987) Transcript of invited talk to IEEE Conference on Neural Networks, June 1987.