


Computer vision**Easy on the eyes**Apr 4th 2007

From The Economist

A computer can now recognise classes of things as accurately as a person can

NEVER underestimate a computer. Never overestimate one either. For many years Garry Kasparov, a world chess champion, said that a computer would never beat him (or, indeed, any other human in his position). In May 1997 he had to eat his words. Deep Blue, an invention of IBM, did just that.

This was impressive, but it demonstrated processing power rather than intelligence. Computers are generally good at solving specific problems, not specifically good at solving general ones. Deep Blue did not learn to play chess from experience. It was painstakingly programmed with thousands of “tactical weighting errors” devised by human experts. So whenever it selected a move, it used these to work through multitudes of possible options and their possible responses. No one is quite sure how Mr Kasparov’s processor operates but it certainly does not do that. One theory goes that the human brain recognises strategic positions in a general way, and that this helps to reduce the problem to a manageable size.

Thomas Serre and his colleagues at the Massachusetts Institute of Technology have built a computer processing system that tries to work in this general way. Among the tasks that computers are bad at is recognising broad categories of images. Tell one to search for something specific, such as a rectangle or even a human face, and it can make a reasonable fist of the task. Ask it to find “animals” among photographs of dragonflies, trees, sharks, cars and monkeys, and it falls over. Indeed a monkey—or even a human baby—would leave it in the dust.

That, at least, was how it used to be. But as Dr Serre describes in this week’s Proceedings of the National Academy of Sciences, his computer handles this problem rather well. In a recent test it even did a little better than humans.

Picture perfect

Given the briefest of glances at a picture, most people believe they have not had time to recognise anything in it at all. Ask them whether they saw an animal and they consider themselves to be making a futile guess. Yet those guesses are right much more often than they are wrong. That is because the brain can carry out immediate visual processing even when it does not have time for any cognitive back-chatter. A neuroscientist trying to understand how people recognise objects would thus start with this simplest of systems.

That is the purpose of Dr Serre's computer. His project is nothing less than an attempt to reverse-engineer the relevant part of the brain. That part is the ventral visual pathway. Anatomy shows that it is organised into numerous areas. Experiments on monkeys, in which researchers have recorded what excites individual nerve cells in each of these areas, give strong hints about how it works.

The pathway is hierarchical. Signals from the retina flow to the most basic processing area first; the cells in that area fire up others in the next area; and so on. Those in the first area are fussy. They react to edges or bars in particular orientations. By combining their signals, however, cells in the second area can respond to corners or bars in any orientation. And so the system builds up. Cells in the final area can recognise general things, animals included.

Dr Serre considered his computer's processing units analogous to nerve cells, and he organised them into areas, just as they are in real brains. Then he let the machine learn in much the same way that babies do. First he mimicked early development when nerve cells are plastic. At this stage babies' brains tune their nerve cells to visual features according to how common those features are in the world around them. That is why kittens raised so that they see only vertical lines have brains that look different from those raised in an environment with purely horizontal ones. Dr Serre's processor developed sensitivities in a similar fashion when he showed it lots of photographs. That stage complete, he then told the computer when what it "saw" contained an animal, and when it did not.

The result was a model that closely imitates the ventral visual pathway. Processing units in each area are sensitive to the same set of features as nerve cells in the brain's analogous areas, and they are linked together as they are in the brain. This artificial recognition system correctly distinguishes photographs containing animals from those without creatures 82% of the time; Dr Serre's students get it right 80% of the time. Moreover, his computer and his volunteers tend to slip up on the same images—and turning photographs on their sides makes poorer animal-recognisers out of both, by roughly the same amount.

A system like this has obvious applications (it may, for instance, soon be put to use searching for child-pornography sites on the internet). But it also brings more subtle benefits. Based as it is on how brains work, it may give insights into what happens when they go wrong. Real neuroscientists rely on lesions (that is, damaged areas of a brain) to help them understand what is going on in brains by seeing what happens in response to particular sorts of damage. Dr Serre has therefore "lesioned" his computer system in similar ways. So far, this has demonstrated the importance in visual recognition of the rare connections that bypass a unit or two in the hierarchy. A computer chess-player could not have told you that.