have become extensions of our drive to compete with each other), and he speculates on the forces that lead human societies to subdivide labour. Although some biologists may find his ideas irritating (can complex behaviours really become innate all of a sudden?), everyone, biologists or not, should find them provocative.

Perhaps these three books — showing in their several ways the importance of studying whole organisms and ecosystems — may help to draw in more research money for ecologists and naturalists. As budgets shrink, these areas are likely — but wrongly — to be among the first to be squeezed. The head of the US National Science Foundation fears that its budget will shrink by 30 per cent in the next five years.

Long gone are the days when, as Bonner reminisces, a letter to the foundation that said "I tried this, that and the other, and nothing really worked" was greeted with the reply: "Don't worry about it — that is the way research goes sometimes. Maybe next year you will have better luck". □

Olivia Judson is a science writer on The Economist, *25 St James's Street, London SW1 1HG, UK.*

Object lessons in vision

Anya Hurlbert

High-level Vision: Object Recognition and Visual Cognition. By Shimon Ullman. *MIT Press: 1996. Pp. 412. £33.95, $40.*

LEGEND has it that about 30 years ago, at the Massachusetts Institute of Technology, the artificial-intelligence pioneer Marvin Minsky instructed a graduate student to connect a camera to a computer and make the computer describe what it saw. What was meant to be a summer project exploded into a huge continuing endeavour, and great advances have since been made in computer vision systems.

But there is still no computer — however well endowed with visual sensors, moving parts and clever algorithms — that can enter an unfamiliar room, immediately process the scene and respond: "Hmm, I'd like to join that little boy playing with the plastic yellow truck, but I'd better not stumble over this pile of puzzles and blocks… hey, is that a red squirrel cavorting on the lawn?"

As Shimon Ullman emphasizes, humans could easily perform the tasks of visual recognition required, even with no previous experience of the scene. Having seen only large grey squirrels before, they could apprehend the small furry red animal as a member of the same family and know that the squirrel was outside, viewed through a transparent surface. Never having met the child, they could identify its gender and approximate age. And, in the general clutter, they could distinguish individual toys, even where parts of one were obscured by another. Our visual ability to recognize objects is astonishingly flexible and robust.

To recognize an object, the human brain must in some way compare the eye's image of it with a stored neural representation — to "retrieve information associated with an object, or a class of objects, that is not apparent in the image itself". But one object may give rise to a multitude of different images, depending on the viewing angle, light-source position, and scene content — and there is not enough room in the brain for all of them.

The problem that the brain has solved is how to match one of the multitudinous incoming images with one of relatively few stored models. Ullman's book graphically demonstrates just how difficult the problem is and explores possible solutions the brain might have adopted.

The book is not a review or textbook; it is more an idiosyncratic, personal history of Ullman's own thoughts and contributions to the field. But these contributions are so varied and important, and conveyed with such crystalline logic and precision, that the global overview is inescapable. Half of the book is devoted to the problem of how to recognize an unvarying object from changing viewpoints. In championing the 'image alignment' solution, Ullman brings high-level vision down to where it can be explained by straightforward image manipulations that the brain could implement quickly.

One of Ullman's seminal contributions has been to prove that a three-dimensional model of an object and its viewpoint-varying image may be aligned by matching just three labelled points on each. Beyond that, Ullman shows that it is enough to store a small collection of two-dimensional views of the object, and then to recognize a new incoming view by aligning it with a weighted sum of the stored views (the 'multiple views' solution).

These are solutions in the machine-vision tradition, mathematical in conception, computer-oriented in implementation. But brains are biological organs, not machines. Ullman's computational approach to object recognition was nurtured at the Massachusetts Institute of Technology but has remained distinctly individual because he steers so close to the behavioural and physiological evidence on how the brain works.

He argues convincingly from experimental results that the brain probably uses the multiple-views solution. He presents a biologically plausible model of how the brain might streamline recognition by operating top-down and bottom-up matching sequences in parallel. He fixes on the most remarkable features of human vision and seeks to explain them.

Why is it that infants learn first how to recognize objects at the class level (horse versus dog, man versus tree) and that in brain-damaged adults the ability to identify individual faces is lost more easily than the ability to recognize faces in general? Fascinatingly, in machine vision it is the other way around: basic-level classification is a more difficult problem than individual identification.

In one of Ullman's own experiments, he shows that our brains must employ special class-based recognition routines for human faces: we are intrinsically poor at recognizing upside-down faces, although we can improve with training. But we remain very bad at recognizing upside-down faces under varying illumination, whereas we can easily recognize upright faces under any shadow. So there must be automatic mechanisms for discounting shading variations, specific to upright faces.

The book is illustrated throughout with black-and-white line drawings of objects ranging from Saabs to roosters. The primal impact of these simple contours emphasizes how powerful our vision is, and how much machines still have to learn. □

Anya Hurlbert is in the Department of Physiological Sciences, Medical School, University of Newcastle upon Tyne, Framlington Place, Newcastle upon Tyne NE2 4HH, UK.

Mum's a test-tube

Benno Müller-Hill

Quest for Perfection: The Drive to Breed Better Human Beings. By Gina Maranto. *Scribner: 1996. Pp. 335. $25.*

THE prize-winning science journalist Gina Maranto's book about the history, present state and possible future of test-tube fertilization begins light-heartedly. Jamie Grifo, a reproductive endocrinologist at New York University, is "the kind of doctor nurses adore". Grifo has a "repertoire of quips and jokes (some of them slightly ribald) which he deploys throughout the day (What does President Clinton say to his wife after sex?… Hi, honey, I'll be home in 20 minutes)".

The reader who expects more amusement of this sort will be disappointed. In the first half of the book, Maranto writes about the history of her subject. What did people do with their babies during the Palaeolithic period? What did the