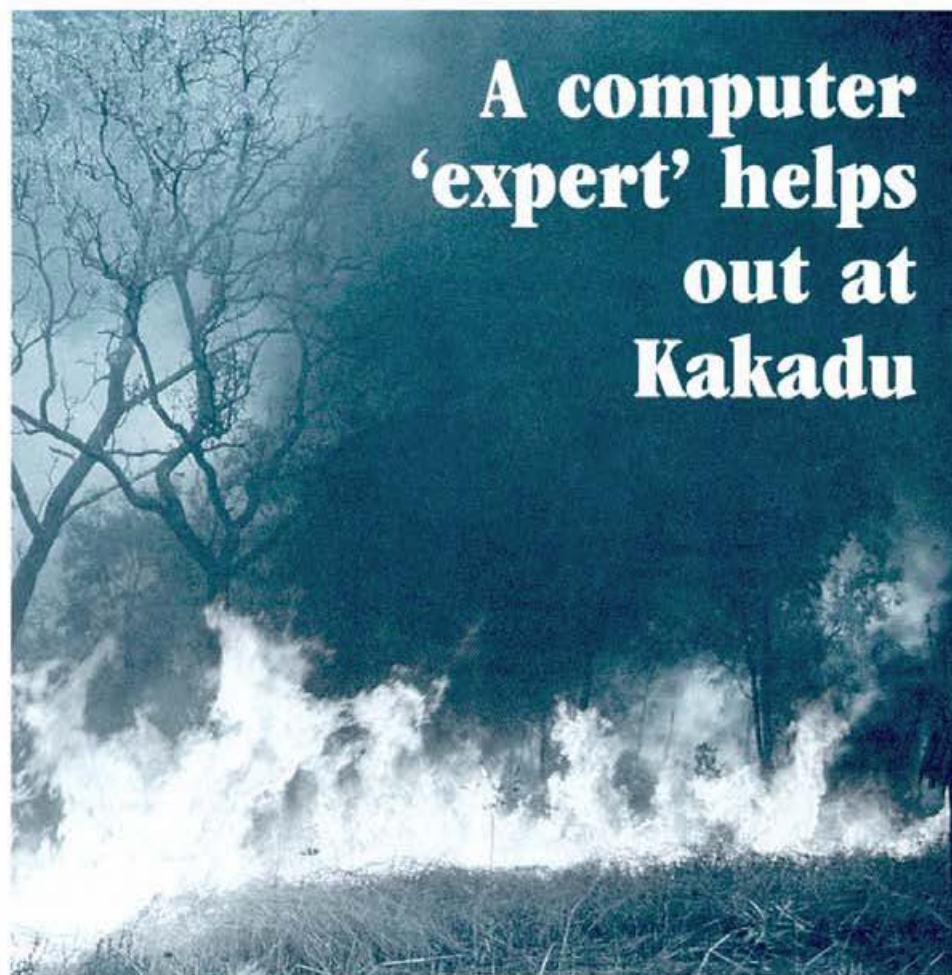


A computer 'expert' helps out at Kakadu



Park managers in the magnificent Kakadu National Park can now consult a computer to foresee the consequences of setting free, or constraining, that potent agent of change, fire. They are using FIRES, perhaps the first-ever 'expert system' installed in an environmental management agency.

An expert system is a computerised attempt to mimic the way a human expert deals with a problem. Instead of crunching numbers, however, as most computers are programmed to do, expert systems juggle symbols that represent factual knowledge.

The computer's machinations are based on rules of logical inference. Feed it a few pieces of evidence and this modern Sherlock will, in a twinkling, flip millions of logic switches and deduce the necessary conclusion.

A manager will key in the prevailing conditions of weather and fuel, and the disembodied expert will flash on the screen an estimate of flame height in any region of the Park, and (for three — so far — management regions) its effect on vegetation there.

The FIRES (fire expert system) attempts to reflect the knowledge about fire and its effects acquired by Mr James Hoare of the

CSIRO Division of Forestry and Forest Products from 10 years' experiments in Kakadu (see *Ecos* 30). His accumulated knowledge also derived from experienced Park managers and documented evidence of Aboriginal burning practices. The people who transplanted this knowledge into the computer's circuits were Dr Richard Davis, a 'knowledge engineer' at the CSIRO Division of Water Resources, and his colleague Mr Paul Nanninga, an artificial-intelligence expert.

Translating Mr Hoare's knowledge into formal statements of logic took dozens of lengthy discussion sessions extending over several months. When it was done, what was known on the subject had been distilled into 105 rules of fire behaviour and 74 relating to its effects.

All this knowledge, programmed into the Park's microcomputer, is constantly available for interrogation by Park staff.

The computer can respond (in stilted English) to the operator, justifying its conclusions by citing each step of its deductions. It can make do with incomplete information, and can attach a degree of certainty to the correctness of its answers.

Furthermore, anybody — managers, rangers, or others — can contribute their own experience by adding appropriate rules to FIRES' knowledge base. In this way, the system can make use of years of careful observation by many people without requiring it to be in numerical form.

The Australian National Parks and Wildlife Service (ANPWS) helped fund the CSIRO research that created FIRES. The problem the ANPWS faced at Kakadu was that it had expertise on fire management in people's heads and research findings in the scientific literature, but no ready way of making them both accessible to all those who were called on to set or control fires within the 13 000 sq. km of the Park.

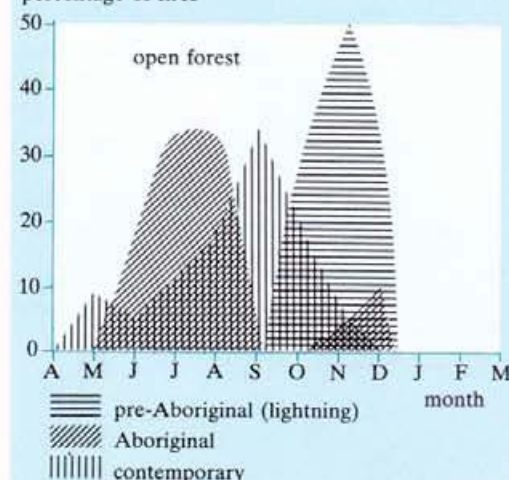
Back to Aboriginal fire

Fire is a common visitor to the Top End. Nearly all the region's forests and woodlands (which carry a tall-grass understorey) are burnt every 1-2 years, the result of lightning strikes or, more commonly, accidental or deliberate lighting. Although these uncontrolled fires are of generally low intensity compared with bushfires in southern forests, and seldom threaten human life, in the long term they may have a substantial ecological impact.

According to the results of Mr Hoare's decade of experiments, fires in Kakadu have been too frequent for the general good of the Park's vegetation. They also burn at times of the year when they are most intense, and burn larger areas than is desirable for the long-term persistence of vegetation, soils, and animals. In particu-

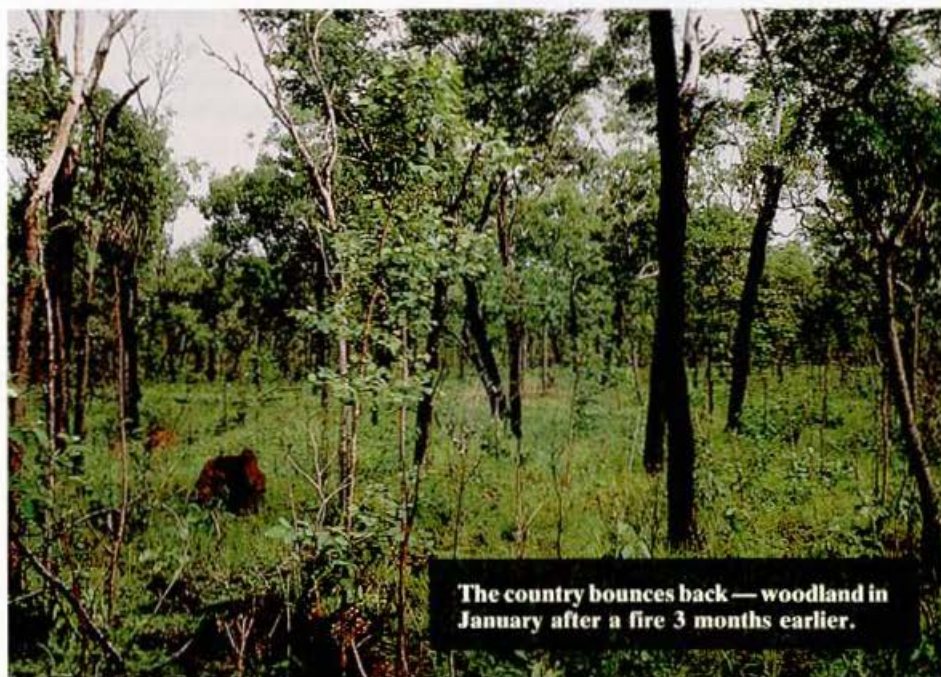
The present fire pattern has prevailed only over the last 100 years, but the inference from Mr Hoare's experiments is that it has considerably changed the structure and condition of vegetation, particularly mid-storey species like eucalypts.

Fires then and now
percentage of fires





The immediate result of fire. But what are the long-term effects?



The country bounces back — woodland in January after a fire 3 months earlier.

lar, fiercely burning tall grasses of the sorghum family have caused serious disruption to the recruitment of taller species, notably eucalypts, into the overstorey. And frequent hot fires eat into pockets of monsoon forest.

Mr Hoare suggests that, to maintain the full diversity of plants and animals, fire patterns should return to those prevailing before Europeans settled there about 100 years ago. Judicious patch burning early in the dry season should decrease the incidence of wildfires and help recreate the original rich mosaic of life forms, he says.

The ANPWS has accepted, in its 1986 Plan of Management for Kakadu, the desirability of implementing such a burn

The 179 rules in FIRES' knowledge base concerned with the physical behaviour and effects of fire are connected in this way. A consultation with FIRES involves tracing a path upwards to reach a conclusion about flame height and its consequent effects.

strategy. But the Service has lacked a system for predicting the severity of any proposed controlled burn and its effects. Formulae for fire-danger ratings routinely applied in southern climes — based on weather and fuel indexes — have yet to be developed for the tropical north.

Presently, managers use their experience to judge suitable times for controlled burning. They assess an area's fuel load, its previous fire history, and the present and recent weather conditions. They have learnt that a fire lit too early in the Dry won't carry, whereas lighting one a little too late can produce an unstoppable fire that may continue for many weeks.

In effect, they learn, over several years, some rules of thumb concerning the propagation of fire under diverse conditions. But managers come and go, and when this happens the replacement will normally need a number of years to acquire competence in fire management — unless, of course, an expert system gives immediate access to previously acquired knowledge. It is hoped that, once adept, trainees will, in turn, tell the expert system what they have learnt.

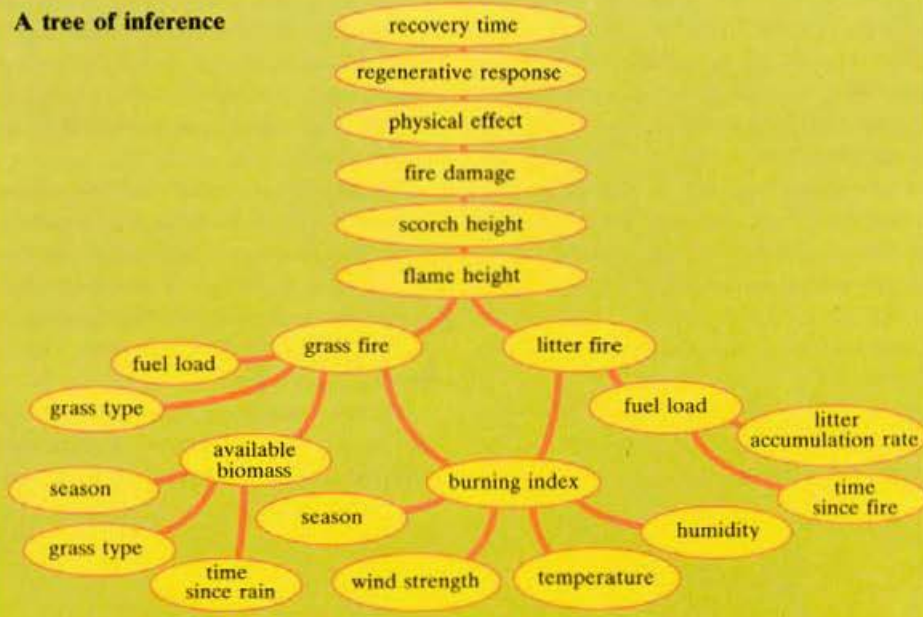
Accumulating knowledge this way appears an attractive proposition for the ANPWS, which is keen to see the system work. It was installed in September 1987, and so far reaction from staff is favourable.

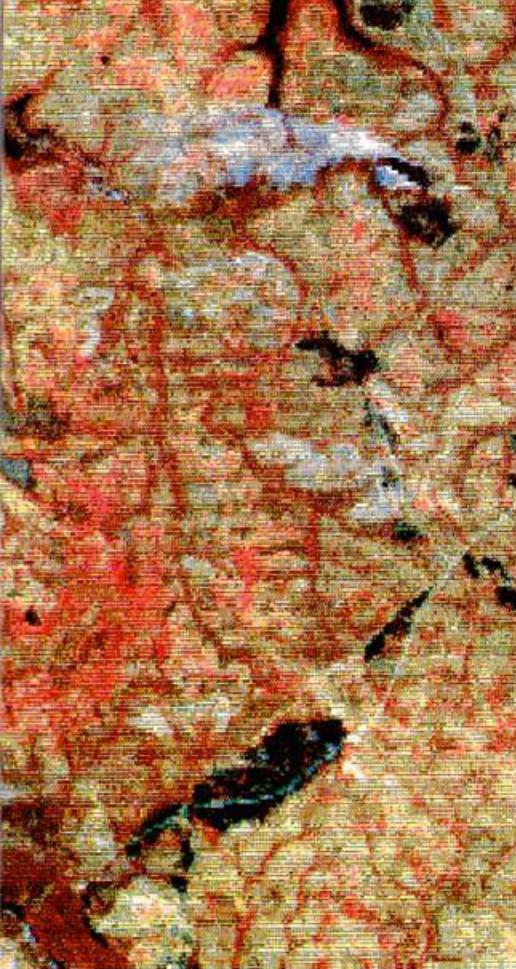
What's an expert system?

Expert systems are difficult to define precisely in a few words, but let's say that each is intended to mimic the role of an expert in situations where decision-making is difficult and a human expert is unavailable.

Expert systems emerged about a decade ago as an outgrowth of artificial-intelligence research. One of the first examples was MYCIN, which tackled the problem of diagnosing infectious blood diseases. Its core of knowledge comprised more than 400 diagnostic rules. Although MYCIN has never been put to practical use, when a panel of human experts evaluated the comparative performances of MYCIN, interns, and experienced medicos, the computer's performance was generally judged to be as good as the others'.

A tree of inference





Smoke from a large fire occupies the centre of this Landsat image of Kakadu. A large fire scar is also visible along the road at the bottom of the picture. (The image was taken on 14 May 1984, and processed on CSIRO's microBRIAN image analysis system.)

Other expert systems have found application in chemical analysis, air-traffic control, nuclear power production, and computer systems configuration. There's even one in England that helps counsel AIDS victims, and another in the United States that knows how to authorise American Express card purchases.

The two essential components of an expert system are its 'knowledge base' — a set of rules reflecting the accumulated understanding of experts — and an 'inference engine' or 'shell' — a mechanism that can apply the rules to solve a given problem.

One section of FIRES' knowledge base predicts the physical behaviour of fires (particularly flame height) over a wide range of fuel loads and weather and vegetation conditions. The other principal section, so far limited to three management regions, predicts the consequences of fire on 91 species of Park flora, and includes information on the damage sustained by particular plant species, the biological effects of the damage, the regenerative response of the species, and the estimated time for it to recover to its pre-burn condition.

To cover the other 81 regions not documented in detail, the knowledge base includes data simply listing the dominant vegetation type in each management region, whether the fuel type is annual sorghum, perennial grass, or litter, and its expected fuel load under average conditions.

Each of the system's 196 rules consists of an if-then pair. For example, rule 32 states that:

IF the season is cool, and
the 3 p.m. humidity is less than 40%, and
the wind strength is high, and
the fuel type is one of
[annual sorghum, perennial], and
the degree of curing is 100%
THEN there is reasonably strong evidence that the fire danger is high.

Each rule is labelled with an author, a list of regions where it is valid, a level of certainty (reflecting the expert's estimation of the reliability of the rule), and information on its origins.

The knowledge base takes, in symbolic form, the shape of a tree (see the diagram), and a run of the system — called a 'consultation' — can be pictured as tracing a path through the tree. The inference engine controls the progression through the tree, driven by the sequential 'firing' of the rules. A rule is said to 'fire' when its preconditions (the 'if' component) are fulfilled: the system then moves on to the next branch of the tree, and so on until it has answered the question posed by the operator.

A typical consultation with FIRES might go:

FIRES: Please type in your name
RANGER: Boss
FIRES: What is the goal parameter?
RANGER: Flame height
FIRES: What region are we to consider?
RANGER: Kay 1
FIRES: What do you reckon the season is Boss?
RANGER: Cold
FIRES: Gimme a value for the 3 p.m. humidity
RANGER: 30%
FIRES: What is the wind strength?
RANGER: Moderate
FIRES: (slight pause) It has been determined that there is strong evidence that the flame height is 2.0–2.5 m

The system has to search through the complete knowledge base, trying potential paths, back-tracking if these reach a dead end, and trying other paths until it succeeds in establishing a conclusion. Inferencing in this manner requires a computer to do a

considerable amount of searching and matching, and this means in practice a computer fast enough to make quick responses and with enough memory to keep track of where it is in its search strategy.

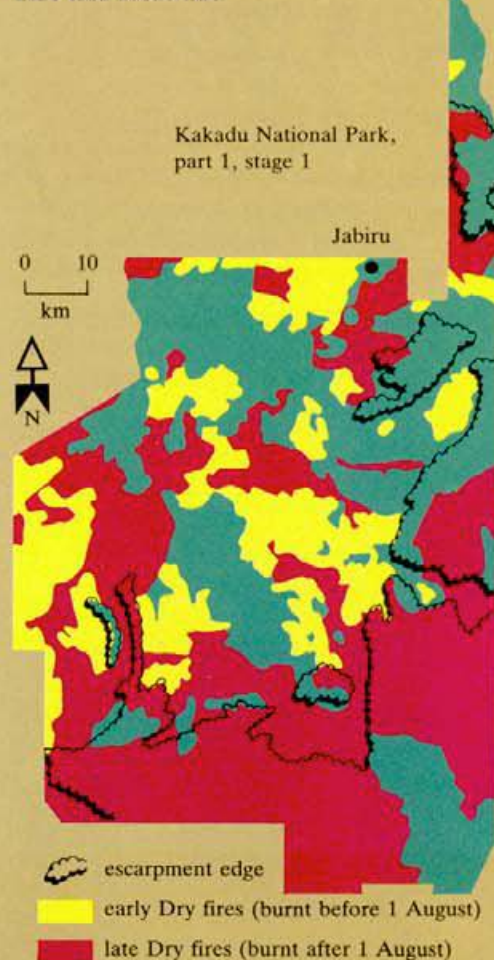
The FIRES system operates on a standard 16-bit microcomputer, a Hewlett-Packard Vectra with 640 K of live memory and a 20-MB hard disk. Its efficiency has to do with its shell, the first expert system shell specifically designed for land management. These days, a common approach to building an expert system is to take a readily available shell 'off the shelf' and load it with your own knowledge base.

But nothing available could fit the bill for the Kakadu problem, so the researchers developed their own shell: it can reason with spatial information, include mathematical models, and deal with the taxonomy and height range of the vegetation. It also incorporates an editor for easily expanding the rule base. And all this must run on a microcomputer in a way friendly to the user (menus, prompts, help commands, and default options make it all run smoothly).

The final product, written in the computer-logic language PROLOG, displays the handiwork of Dr Davis, Mr Nanninga, and Mr Grahame Williams, then a doctoral

Landsat shows how much of Kakadu was burnt in 1982.

Fire and more fire



Fire behaviour in Kakadu

The key factor in predicting how fire will behave in the wet-dry tropics is the season. Each of the fire seasons recognised by the resident Aborigines — cool, cold, hot, early storms, monsoon, and 'knock-em-down' storms — has particular fuel and meteorological conditions.

In summary, fires don't burn during the monsoon or knock-em-down seasons (which extend over about 5 months from mid November to early May) because the grass is uncured or, where cured fuel has accumulated from previous years, because of its high moisture content.

The cool season (early May to early June) represents a transition period of about 6-8 weeks between the Wet and the Dry. Early in the cool, the fire's intensity and its pattern of burning are restricted by patchy curing of grass; later on, when fuel is almost fully cured, high humidity during the day and dew at night also keep fire at bay.

During the cold and hot seasons, which take in about 5 months of the Dry from early June to early October, the vegetation progressively dries out, and at some point

fire will begin to burn uniformly and through the night as well as during the day.

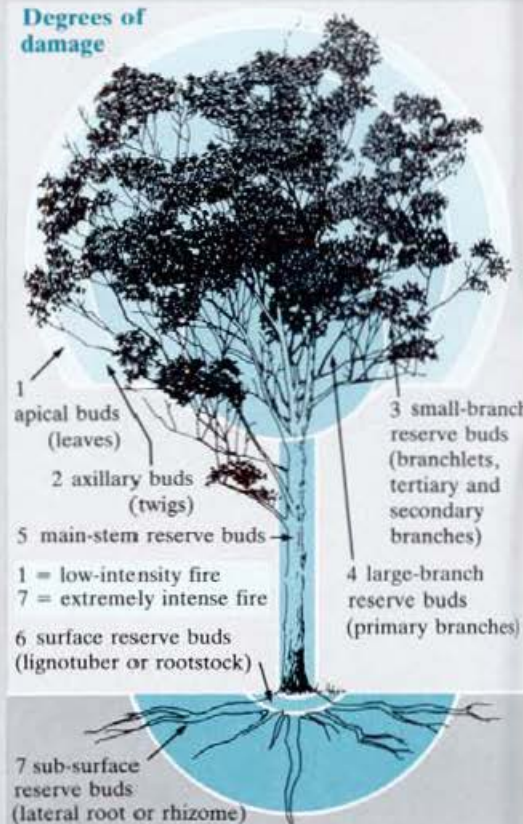
In the early storms season (early October to mid November), fires can behave very differently depending on fuel quantity and dryness, wind, temperature, and humidity.

Mr Hoare's long-running experiments, together with information provided by other researchers and Park managers, have enabled him to relate these factors to how intensely a fire will burn (measured as flame height).

A model of fire behaviour and resulting damage to vegetation for predicting the biological effects of fire on plant communities at Kakadu National Park. J. Hoare. *CSIRO Division of Water and Land Resources Technical Memorandum No. 85/2, 138-52, 1985.*

As fire increases in intensity, from 1 to 7, it will kill more of a plant's reserve buds: FIRES uses flame height as a measure of intensity and, calling on data on each species' growth form and position of reserve buds, works out how the vegetation will regenerate.

Degrees of damage



student. They started in 1985 with the shell of the MYCIN expert system and, after gaining experience with it, developed their own shell specially for FIRES.

Words into symbols

One very clear lesson the researchers learnt was that constructing an expert system is not for the faint-hearted. Although it may seem a conceptually simple task to construct a knowledge base, it is difficult to do so consistently and comprehensively.

An expert system needs a large number of rules so that it can respond to a wide

Some fires need to be controlled; others can be left to burn themselves out. FIRES can help in the choice between management options.

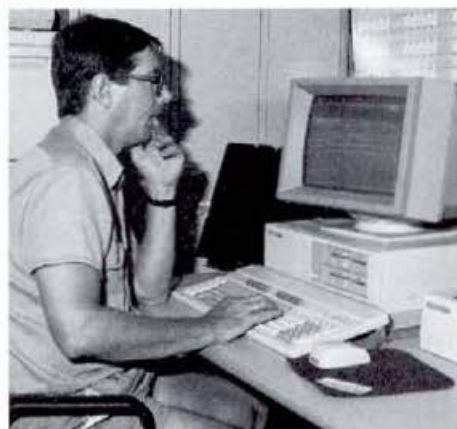
range of situations, yet the rules cannot be randomly supplied and unconnected. Not only must the expert's knowledge be elicited in a coherent form, but the 'knowledge engineer' must translate it accurately into symbolic language. Then the rules must be checked for consistency and overlap with each other. It's a long slow task.

For example, the level of understanding about Kakadu fire behaviour is still rudimentary and can be stated in just a few paragraphs (see the box on this page), but it took several months to reformulate it into 105 rules. We won't list those 105 here, but it's worth looking at the English language version to get an idea of the scope of a practical expert system.

Dr Davis recognises that a constant problem with expert systems is to ensure that they don't encompass just the trivial. (After all, an expert is commonly somebody who knows everything about virtually nothing.) On the other hand, if the task is too complex, it becomes fiendishly difficult to get the rules complete and consistent.

Once the two collaborators had completed the fire-behaviour rules, and had taken care of omissions and redundancies, it was time to do the same with the rules for the effect of fires on the vegetation. These two parts of the knowledge base are

A Park manager at Kakadu, Mr Tony Press, consulting with the FIRES expert system.



Expert systems — what they can (and can't) do

Only six or so expert systems are in daily practical use in this country. With all the benefits they offer (a captive expert in your own computer), and when the know-how to create them has been around for at least a decade, why are there so few?

One answer would be to point to the time and expense involved in locking up an expert and a knowledge engineer in a single room for long enough to elucidate the many rules that constitute the system's knowledge base.

But a second answer, more philosophical and critical, calls into question the whole enterprise of trying to encapsulate human knowledge into computer circuits. Some critics say that the difficulties are symptomatic of a fundamental dichotomy between human intelligence and the machine variety.



The issue revolves around the notion of common sense: it's a characteristically human quality that expert systems obviously lack. And its absence means that by no stretch of the imagination could you call expert systems 'intelligent'. They are tools that aid thinking, but they fall short of a human expert because general knowledge is conspicuously absent.

While an expert system's rules do give it problem-solving ability, this is largely at a fairly simple level. The explicit rules are only condensed fragments of the real expert's knowledge. And the obstacles along the road to deepening the computer's understanding appear formidable.

Perhaps with a vast amount of computer memory, a system making use of a greater breadth of knowledge would emerge. But more fundamentally, the human reasoning

process makes use of a whole host of knowledge that is impossible to express as logical rules — humans often do not know what rules they use, and the real world is not neatly ordered and logical anyway.

Nevertheless, expert systems like FIRES can be very useful. Dr Davis views them as tools that assist our own thinking process, and he thinks that the potential for Australian development of such systems looks good.

A number of institutions and companies are working in this fledgling field, and many industries recognise the potential advantages that expert systems offer. The chemical, petroleum, and coal industries are some easily identified, and so are those making iron and steel, machine tools, textiles, and electrical machinery.

Many examples were presented at the National Workshop on Expert Systems held in 1986 by the Department of Industry, Technology and Commerce.

Within CSIRO, a number of expert systems have been developed or are being worked on.

At the DIVISION OF WATER RESOURCES, Dr Davis and his colleagues have been concerned with, in addition to the FIRES system:

- ▷ a system to assess the 'trafficability' of Cape York to army vehicles under various conditions
- ▷ a program to assess tick risk, written in collaboration with the DIVISION OF ENTOMOLOGY in Brisbane, which estimates the risk of tick infestation and tick-borne diseases
- ▷ a shrub-control model, SHRUBKIL, developed with the DIVISION OF WILDLIFE AND ECOLOGY at Deniliquin, N.S.W., to help extension officers advise pastoralists if and when they should burn shrub regrowth
- ▷ a program to predict water and fertiliser needs for wheat crops in the Griffith area (with the former CSIRO CENTRE FOR IRRIGATION AND FRESH-WATER RESEARCH)

high-intensity burns. The damage wrought by various scorch heights will, in turn, depend on the growth form, height, and morphology of the species in each region.

The FIRES system recognises three types of fuel — annual grasses (sorghum), perennial grasses, and litter — and the

The Garvan Institute for Medical Research in Sydney uses an expert system to interpret some 10 000 thyroid hormone assays annually. Through continued modification of the rule base, it's estimated that the system now achieves accuracy of 99.7% in clinical diagnosis. But the modification process has gone so far that further refinement has become difficult. The DIVISION OF INFORMATION TECHNOLOGY is collaborating with the Garvan Institute to redesign the knowledge base so that changes to it can be made more easily and reliably.

The Divisions of PLANT INDUSTRY and INFORMATION TECHNOLOGY, together with the Digital Equipment Corporation, are transforming the cotton-management tool SIRATAC from its original Fortran format into a comprehensive expert system that advises cotton-growers on when and what to spray for insect-pest control. (SIRATAC was developed by the Division of Plant Industry and the New South Wales Department of Agriculture as a means of reducing pest resistance problems; one-quarter of our cotton is now grown under its management.)

The DIVISION OF CONSTRUCTION AND ENGINEERING has created the Knowledge-Based Systems Research Centre to work with industry in the development of expert systems. So far, these scientists have written:

- ▷ WINDLOADER, which allows designers and engineers to apply wind-loading Codes of the Standards Association of Australia to a proposed building
- ▷ WATERPEN, a program that allows construction designers to assess if windows can resist water penetration during storms
- ▷ TIMBER EXPERT, a prototype system that reduces the complexity of timber-beam designs
- ▷ a program for the Melbourne and Metropolitan Board of Works, which advises operators of its sewerage and water-supply network on how to respond to sudden faults detected in the network by automatic monitors

geographic data base details which fuel type predominates in each region of the Park. Flame height is worked out from fuel type, a weather- and season-dependent burning index, and fuel load.

Fuel-load figures (in tonnes per ha) can be supplied by the user. If they are

linked through the flame height, which the fire-behaviour rules establish.

Mr Hoare's experiments had shown that the flame height is a good indicator of scorch height, which gradually increases from twice the flame height for low-intensity fires to four times flame height for



Kakadu burning.

unknown, FIRES can use the standing average figures for grasses in its data base to make estimates based on the season and the number of days since the last fall of rain.

Advantages

A manager can use FIRES to predict the immediate damage to vegetation, and assess the long-term biological effects, resulting from any fire that may start today. In this way, the manager can decide whether to make strenuous efforts to control a fire, or to let it burn.

Alternatively, with specific objects in mind — such as habitat management or species conservation — he or she can

A late dry-season fire stops where an earlier fire reduced fuel loads.



A wet-season fire has removed *Sorghum intrans* at the left. That remaining on the right will promote a dry-season fire.

determine when, and under what conditions, a fire should be deliberately set.

At the end of a consultation, the manager can ask the system why it offered the advice it did. This justification process provides the user with values for all the factors upon which the result was based, how these values were obtained, and the reliability attached to each. For example, on a day when the system suggests that the vegetation damage will be severe, it does so because under rule 45 it concluded that 'there is strong evidence that the flame height will be greater than 3 m', and based its conclusion, in turn, on the 'high fuel load and hot windy conditions'.

Dr Davis and Mr Hoare believe that FIRES offers advice on the effects of fire in Kakadu that, for management purposes, is the best available. They doubt whether any conventional computer program using mathematical models could (at this stage) do as well.

Initially, the scientists saw expert systems as a convenient way of helping non-experts. But Mr Hoare, in particular, found that the demand to articulate his knowledge so

clearly helped him to see gaps and logical problems in it, and to devise better models of what was going on.

The ANPWS envisages further improvement to FIRES coming from incorporation of the fire knowledge of the several hundred Aborigines living in the Park. For example, they know that when a certain plant begins to flower it's time to put the fire stick to work.

Andrew Bell

More about the topic

Developing a fire management expert system for Kakadu National Park, Australia. J.R. Davis, J.R.L. Hoare, and P.M. Nanninga. *Journal of Environmental Management*, 1986, **22**, 215–27.

Transferring scientific knowledge to natural resource managers using artificial intelligence concepts. J.R. Davis, P.M. Nanninga, J.R.L. Hoare, and A.J. Press. *Ecological Modelling*, 1988, **40**, (in press).

Towards an expert system for fire management at Kakadu National Park. J. Walker, J.R. Davis, and A.M. Gill (eds.). *CSIRO Division of Water and Land Resources Technical Memorandum No. 85/2*, 1985.

A knowledge-based approach to land management. J.R. Davis, P.M. Nanninga, J.R.L. Hoare, and J.L. Clark. In 'Desk-Top Planning', ed. P.W. Newton, R. Sharpe, and M.A.P. Taylor. (Hargreen: Melbourne 1988, in press.)

The FIRES expert system knowledge base. J.R.L. Hoare, J.R. Davis, and P.M. Nanninga. *CSIRO Division of Water and Land Resources Technical Memorandum No. 86/19*, 1986.

Fire damages top-end forests. A. Bell. *Ecos* No. 30, 1981, 18–20.