

Expert Systems: Where Are We? And Where Do We Go From Here?

Randall Davis

*MIT AI Laboratory
545 Technology Square
Cambridge, Massachusetts 02139*

WORK ON EXPERT SYSTEMS has received extensive attention recently, prompting growing interest in a range of environments. Much has been made of the basic concept and of the rule-based system approach typically used to construct the programs. Perhaps this is a good time then to review what we know, assess the current prospects, and suggest directions appropriate for the next steps of basic research.

I'd like to do that today, and propose to do it by taking you on a journey of sorts, a metaphorical trip through the State of the Art of Expert Systems. We'll wander about the landscape, ranging from the familiar territory of the Land of Accepted Wisdom, to the vast unknowns at the Frontiers of Knowledge. I guarantee we'll all return safely, so come along...

The itinerary Our trip has three basic purposes in mind. I want first to assess and calibrate the Accepted Wisdom, second to understand its limitations, and third to characterize the nature of the research that will produce the next generation of systems. The talk is correspondingly divided into three parts. We begin with a tour through the Land of Accepted Wisdom, where we survey the field, trying to assess what's known about building and using these systems. What do we know and what can we do with that

knowledge? Where has the Accepted Wisdom succeeded?

One important goal of this part of the journey will be to calibrate the current state of the art—Given what we know, what can we do and how quickly can we do it? Does building a system typically require several months or several years? Can we turn them out routinely or are we still hand-crafting them? Is the technology a laboratory curiosity or is it ready for mass production?

As the name suggests, any journey through the Land of Accepted Wisdom invariably means encountering familiar territory. But a review of fundamentals is important if we are to accomplish our second goal—the careful determination of the nature and location of the Frontiers. We need to examine the Accepted Wisdom closely, to see where the ragged edges are and to see where experience may have worn holes in some of the cherished clichés. That careful examination forms the second part of our journey.

The shortcomings we find will lead us to our final goal, an understanding of which way to go next, as we head toward the Frontiers of Knowledge. Our vehicle will be a problem in equipment diagnosis that nicely characterizes the weaknesses of the existing technology and brings sharply into focus the challenges we face in building the next generation of systems.

A glimpse ahead. It's standard practice in a talk to give away the punchlines early on, so here's a preview:

This article has benefited from careful readings by and suggestions from Mike Brady, Bruce Buchanan, Ed Feigenbaum, José González, Pete Szolovits, and Pat Winston

Editor's note: This article was the author's "Invited Lecture" at IJCAI 7

The Central Issues

The Accepted Wisdom has provided an important foundation and a useful set of tools

Those tools can be particularly powerful where performance is based largely on "compiled experience"

Additional varieties of knowledge are necessary beyond empirical associations

Expertise is characterized by a range of behaviors

New problem domains require additional mechanisms.

We need to reconcile the new directions with the advantages of the Accepted Wisdom

In examining the intellectual foundation and the tools provided by the Accepted Wisdom, we'll see that they can be particularly powerful in domains where expertise is well described as "compiled experience." In clinical medicine, for example, a considerable amount of diagnostic expertise arises from exposure to numerous examples. Typically, relatively little of a physician's diagnostic capability is based on reasoning from detailed anatomical and physiological considerations. Instead, what makes a clinician an expert is the large collection of empirical associations he accumulates by dint of experience in the field. In domains of this sort, the existing set of tools can be useful and powerful.

But we have to press on beyond the conventional wisdom, because additional varieties of knowledge are going to be necessary in the next generation of systems. There are several reasons for this. First, as we'll see, expertise is characterized by a wide range of behaviors, of which problem solving performance—the focus of current work on expert systems—is really only the most evident. Experts do much more than simply solve the problem and it's not clear that empirical associations are capable of supporting those other forms of behavior.

Second, there are domains where problem solving clearly relies on more than compiled experience. Other varieties of knowledge are involved, knowledge of structure and causal models. To take advantage of the power they can provide, we're going to have to press beyond the existing technology.

We'll see that the desire to use those additional kinds of knowledge leads us off in directions that conflict with some of the advantages that the Accepted Wisdom has provided. Hence one of the long term challenges we face is determining how to reconcile conflicts between the new directions and the advantages of the Accepted Wisdom.

A word about terminology Let me focus for a moment on the phrases "empirical associations" and "causal models." They are, I think, more appropriate than "shallow" and "deep," words that have been used quite a bit recently in reference to much the same thing. The problem with "shallow" and "deep" is that they are simply labels whose technical content is at best obscure (i.e., they are shallow), while their overtones from nontechnical use make them far too evocative.

THE NATURE OF EXPERTISE

A Range of Behaviors: solve the problem
explain the result
learn
restructure knowledge
break rules
determine relevance
degrade gracefully

Figure 1

The technical issue is, what grounds exist for believing an inference like *If A and B then C*? If the strongest argument we can make is of the sort, *Previously, when A and B held, C was also found to be true*, then the inference is justifiably characterizable as an *association* that grew from accumulated *empirical* observations.

If, on the other hand, we can argue that *If A and B are true, C follows because of the way this device (or system, organ, etc.) works*, then we are basing our belief on an understanding of structure, function, or *causality*.

Such terms focus on the relevant issue and have the virtue of avoiding the unfortunate connotations of "shallow" (i.e., requiring little thought) and "deep" (i.e., profound) that carry over from everyday use. Clearly, for example, anyone who has a *shallow* system is missing something (and probably doing shoddy work). Anyone who's got a system based on *deep models*, on the other hand, is clearly doing something impressive.¹

We should be wary too of the images suggested by the terms Accepted Wisdom and Frontiers. They are convenient abstractions and are rarely as clearcut and static as may be implied here. Nor is the population of each Land easily determined—those who create the Accepted Wisdom at one moment are often found across the borders at the next. The metaphor is, nevertheless, a useful abstraction that allows us to categorize what we know and measure what we have yet to do.

The Nature of Expertise

Enough preliminaries. Let us begin by considering the nature of expertise. It includes a whole constellation of behaviors (Fig. 1). Problem solving is only the most obvious and while necessary, it is clearly insufficient. Would we be willing to call someone an expert if he could solve a problem, but was unable to explain the result, unable to learn anything new about the domain, unable to determine whether his expertise was relevant, etc.? I think not.

¹ *Deep* is a word like *theory* or *semantic*—it implies all sorts of marvelous things. It's one thing to be able to say "I've got a theory," quite another to say, "I've got a semantic theory" but ah, those who can claim "I've got a deep semantic theory," they are truly blessed

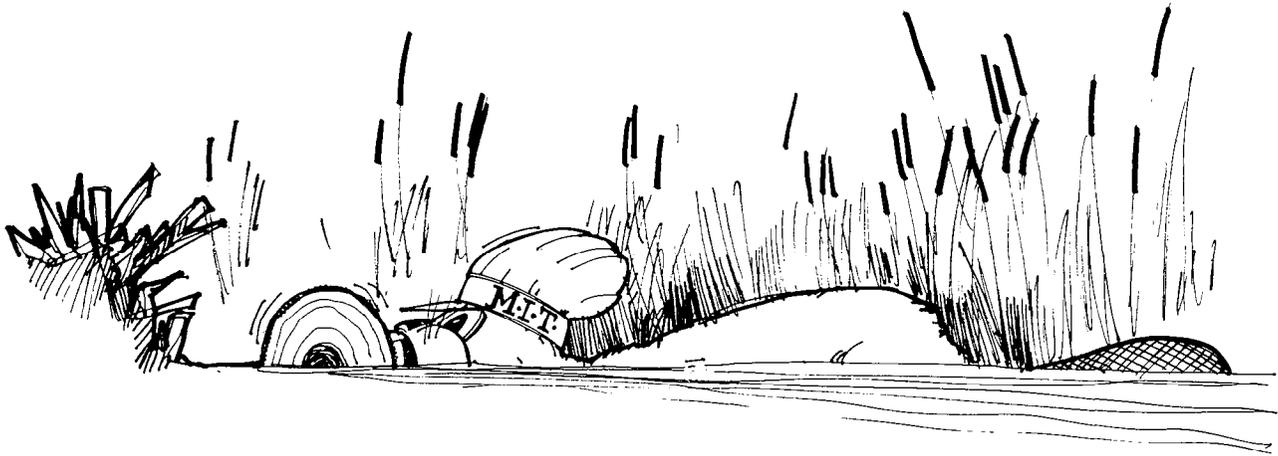


Figure 2. An early expert system.

Work in expert systems has to date explored only the first three of these behaviors in any depth. First generation systems like DENDRAL (Lindsay et al., 1981) and MACSYMA (Macsyma group, 1974) focused solely on performance, while second generation systems began to explore behaviors like explanation (e.g., MYCIN, Shortliffe, 1976; the Digitalis Advisor, Swartout, 1981) and knowledge acquisition (e.g., TEIRESIAS, Davis, 1979; Version Spaces, Mitchell, 1978). Of these, performance is still the best understood; our efforts at explanation and knowledge acquisition have only scratched the surface.

By and large the other topics have been almost totally unexplored. What would it mean for example to restructure knowledge? One example particularly relevant for this audience came out of the so-called procedural vs. declarative controversy. We, as experts on knowledge representation, went to work structuring and restructuring knowledge, so that procedural became declarative, got turned back into procedural again, and so forth. I think it became clear after we went around that loop a couple of times that the problem is at least in part in the eye of the interpreter, but nevertheless, there we were, restructuring and reorganizing knowledge.

What about breaking rules? One of the most frustrating things to an apprentice is that he no sooner learns a rule from an expert than he discovers that there are almost as many exceptions as there are rules. Experts clearly understand not merely the letter of the rule but the spirit as well.

What about determining relevance? Experts also know when they're beyond their depth. They know when a problem is outside their sphere of expertise and they know when the best answer they can give is to suggest asking someone else. Clearly none of our programs can do that as yet.

Experts also degrade gracefully. That is, as they get close to the boundaries of their expertise, they get less and less proficient at solving problems. But their skill decreases smoothly, rather than precipitously as do most of our programs.

To illustrate this I'd like to introduce a colleague of mine

(Fig. 2) who'll be accompanying us on this journey. It's been suggested that the beaver, engineer of the animal kingdom and (not entirely by coincidence) mascot of MIT, would make a perfectly good expert system if performance alone were the issue. After all, he does a marvelous job of building a dam.

But performance isn't all there is. As I was saying a moment ago, experts can explain their behavior, know when to break rules, and can learn more about their domain. On that score, of course, our friend seems to fall short (Fig. 3)

He will, nevertheless, accompany us as we head off down the road to see what lies ahead in the Land of Accepted Wisdom (Fig. 4).

In the Land of Accepted Wisdom

On the first stop in our journey, I'd like to consider where we are by characterizing the stages of development of a field and then use that to assess the state of expert systems work.

The first stage of development of any field is traditionally *case studies*. Ideally, we test one single dimension of a design—one idea on knowledge representation, or control, or system architecture. Each of these is in effect an isolated point in the design space.

As the set of such experiments grows, a collection of *architectural principles* may emerge. By examining many case studies we may begin to understand the shape and character of the design space. This allows us to make empirical observations about which parts of the space make sense for which kinds of problems. Note that these are simply empirical observations—we don't really understand why they hold, we can only say that, with our experience this design looks like the right one to use.

Eventually, and of course it's a big eventually, we get to something worthy of being labeled a science. For our purposes, that stage is characterized by an understanding that goes beyond a set of empirical observations. There is instead an understanding of *why* a particular design is appropriate for a particular task, and perhaps a much better understanding of the character and shape of the design space.

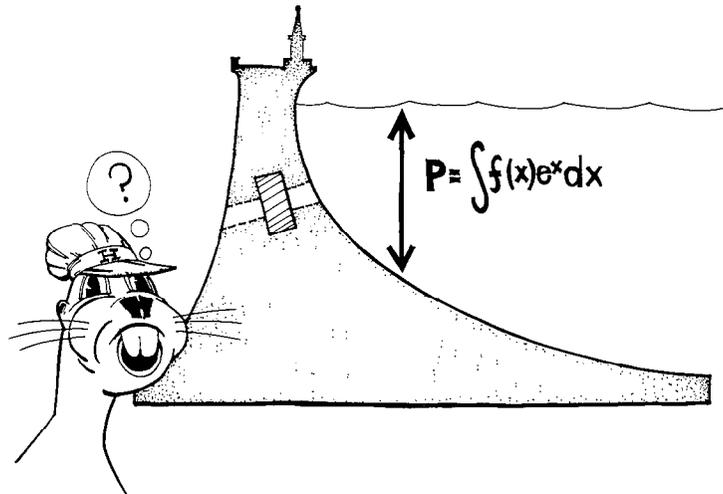


Figure 3. Outmoded by a new technology.

I think it's fair to say that work on expert systems is currently somewhere between the stage of case studies and architectural principles. As is traditional for fields at that stage of development, and as is particularly appropriate for expert systems, the existing knowledge is well captured as a collection of informal rules of thumb.

Knowledge engineering homilies. So on our next stop in the land of Accepted Wisdom, we pause for a moment to survey the sights and see what it is that we know (Fig. 5).

Figure 6 presents six of the basic commandments. Many of you will of course recognize these from a talk given four years ago by Ed Feigenbaum (Feigenbaum, 1977). I've glorified their form, but not their content. Let me review them briefly.

Perhaps the most fundamental observation—*In the knowledge lies the power*—suggests that problem solving performance often arises from extensive stores of knowledge about the task at hand, not from a large collection of domain-independent methods.

We note that the knowledge is *inexact and incomplete*, because, almost by definition, the kinds of problems attacked in AI very rarely have complete laws or theories. Rather we have to deal with an inexact and informal knowledge, the stuff of half-baked theories and guesses.

The knowledge is often *ill-specified*, because the expert himself doesn't always know exactly what it is he knows about his domain. That means in turn that one of the things we have to do is help the expert to explicate his knowledge, help him to make clear what it is that he knows and how that is applied to the problem at hand.

Amateurs, whether they're people or programs, *become expert incrementally*. Unlike Athena, they don't emerge fully formed. Acquiring expertise is instead a process of incremental approach to competence.

One impact of that is the importance of systems that are *flexible and transparent*. These systems are going to spend most of their lives being changed, updated, and improved. If it's too difficult to change them—they are inflexible—the

whole process is going to come to a grinding halt. Transparency is similarly motivated: If we're going to improve the system we have to know what it did in order to determine what it is we ought to change. If the system is a black box, it becomes impossible to make that determination and system evolution will cease.

Architectural principles. In addition to the basic credo, some architectural principles have begun to emerge (Fig. 7). One suggestion is to *separate the inference engine and the knowledge base*. By doing so, the knowledge in the knowledge base is more easily identified, more explicit, and more accessible. If the two are intermixed, domain knowledge becomes spread out through the inference engine, it becomes less clear what we ought to change to improve the system, and flexibility suffers.

A second architectural principle is *uniformity of representation*. This cuts down the number of mechanisms required, keeping system design simpler and more transparent. Each time a new representation is added to the system, something else in the system has to be able to handle it, has to know its syntax or semantics to be able to use it. Hence fewer representations means a simpler, more transparent system.

Keeping the inference engine simple helps in several ways. Explanations, for example, are easier to produce. Since they are currently generated by replaying the actions of the system, keeping those actions simple means less work is necessary to produce comprehensible explanations. Knowledge acquisition is similarly easier. When the inference engine is less complicated, less work is needed to determine exactly what knowledge to add to improve system performance.

A fourth principle—*exploiting redundancy*—is nicely illustrated by work on HEARSAY (Erman et al., 1980) that illustrated how redundancy can be a remedy for incomplete and inexact knowledge. The trick is to find multiple overlapping sources of knowledge with different areas of strength and different shortcomings. Properly used, the entire collection of knowledge sources can be a good deal more robust

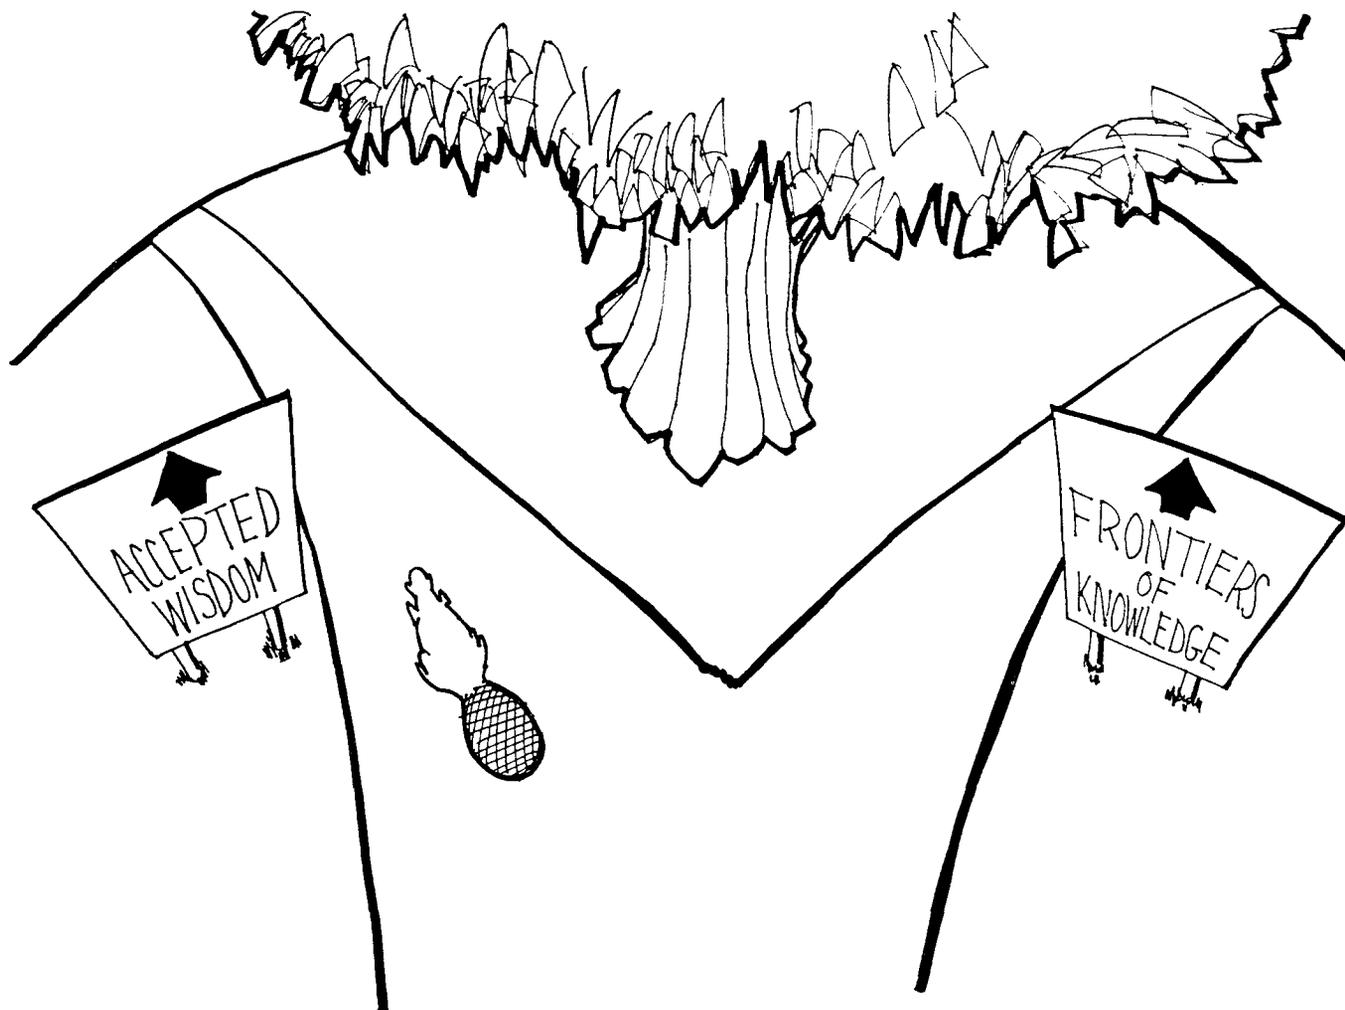


Figure 4

than any one of them taken alone

The current state of the art. What then is the current state of the art? We can characterize it roughly as suggested in Figure 8. Expert systems have to deal with very narrow domains of expertise; we have to very sharply constrain what it is we hope to achieve with them. As we get closer to the boundaries, their behavior gets rather fragile; rather than degrading gracefully, it tends to fall apart precipitously

Since the emphasis and the effort in these systems is on accumulating large knowledge bases, the knowledge representation used is typically one of the simpler ones, like attribute-object-value triples, production rules, and so forth. And since the natural language problem has yet to be solved, we're stuck with limited interaction languages, usually keyword-based parsing of input and template-generated production of text on the output.

Our model of explanation is useful but limited at the moment to recapitulation of the system's actions. We simply hope that's sufficiently informative that it will explain the behavior of the system. And finally, we don't yet know

very much about dealing with multiple experts. How do you reconcile differing and perhaps conflicting views among acknowledged experts? At the moment we don't know, so we appoint one expert as knowledge base czar and attempt to build the system in his image.

Successes of the Accepted Wisdom. So much for the current state of the art. Where does that get us? In what ways has the Accepted Wisdom led to success?

One success is the existence of the credo itself. Interestingly, not very long ago those six commandments were highly controversial. Not very long ago Joel Moses was told, when working on MACSYMA, "That isn't AI. Don't know what it is, but it isn't AI." Not very long ago the comment on DENDRAL was "That's interesting chemistry, but what has it got to do with AI?" The past five or six years have seen a significant shift in emphasis: Perhaps it *isn't* cheating to build a large store of knowledge about a domain. Perhaps part of intelligence really *does* arise from the accumulation and use of lots of information about a task.

A second contribution of the existing wisdom is its utility as a simple but effective tool for getting started in problems

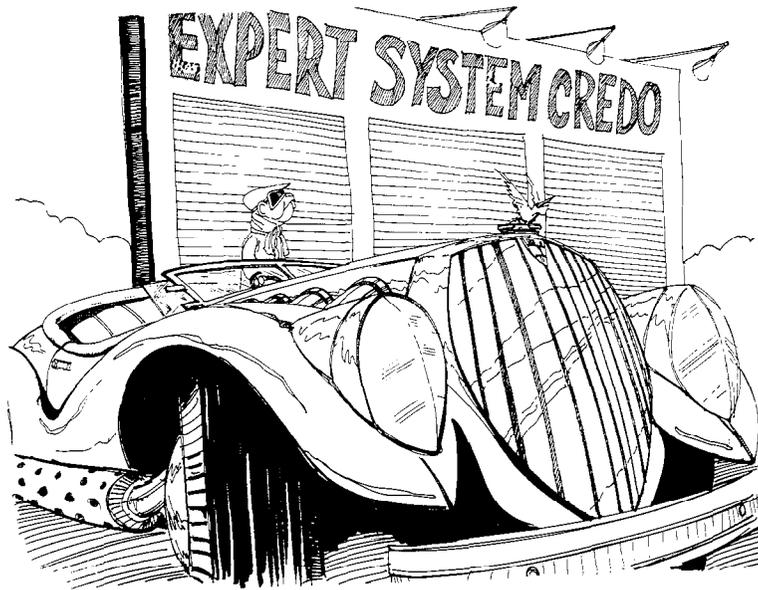


Figure 5

involving extensive amounts of informal knowledge. It suggests extracting knowledge from experts by working through sample problems with them to find out what they know. It suggests as well a methodology for knowledge acquisition: the detailed comparison of system vs. human performance on carefully selected problems.

The Accepted Wisdom also offers a way of using empirical associations to deal with inexact knowledge. The problem of inexact inference is difficult and far from solved, but the Accepted Wisdom provides a tool with a certain pragmatic utility.

The Accepted Wisdom has also given us at least four systems in current use, systems that are solving interesting and difficult problems. DENDRAL, MACSYMA, PUFF (a program for diagnosing some classes of lung disease; Kunz et al., 1978), and R1 (system configuration for VAXes; McDermott, 1980) are no longer the focus of work in research labs, they are instead being employed by the appropriate user community.

Those are the technical successes of the Accepted Wisdom. There have also been some important sociological effects. It first became evident at last year's AAAI conference and is clearly manifested here as well: The Accepted Wisdom

has succeeded in arousing considerable interest in expert systems (and AI in general) outside the academic centers. We're beginning to see more and more industrial research centers becoming interested in the technology. Why is that? I think that part of it is simply coming of age of the field in general. Certainly much the same thing is true in natural language and robotics where commercial systems are beginning to be available.

That interest has had a number of results, perhaps the most notable being the imbalance between the tremendous amount of interest in the technology and the rather sharply limited manpower supply capable of building these systems. That in turn has led to intensive recruiting of engineers of almost every sort (Fig. 9).

Calibrating the technology. For some, this growing collaboration between industry and AI is cause for concern. There are issues of manpower spread thin over a wide range of sites and concern about manpower available to train the next generation. Equally important, however, now that AI has become a focus of strong attention, are the expectations that industry has about what AI and expert systems can supply, and how quickly that promise can be fulfilled.

It is important to calibrate as closely as we can the distance between research results and commercial products. It has become important to set expectations that make clear

CREDO

- In the knowledge lies the power.
- The knowledge is often inexact, incomplete.
- The knowledge is often ill-specified.
- Amateurs become expert incrementally.
- Expert systems need to be flexible.
- Expert systems need to be transparent

ARCHITECTURAL PRINCIPLES

- Separate the inference engine and knowledge base
- Use as uniform a representation as possible
- Keep the inference engine simple
- Exploit redundancy

Figure 6

Figure 7

STATE OF THE ART

Narrow domain of expertise
Fragile behavior at the boundaries
Limited knowledge representation language
Limited I/O
Limited explanation
One expert as knowledge base "czar"

Figure 8

what can in fact be expected from expert systems technology as it currently exists.

One way to calibrate is that collection of statements about the state of the art. All those catch phrases about limited domains of expertise, fragile behavior of the boundaries, etc., help to characterize the state of the art and set the appropriate expectations about what can be done with today's technology.

To push beyond the level of catch phrases, I conducted a small survey. I called a number of people who have helped build expert systems and asked them to estimate the number of man years that were required. The results are shown in Figure 10.

There are a number of interesting things here, but let me start with a few caveats about what's *not* here. For a number

of reasons, the curve can't be taken too literally. First, the systems mentioned here vary drastically with respect to the size of the problem they are trying to solve. MACSYMA is probably two orders of magnitude larger than PUFF, for example. Others differ radically in the amount of knowledge they need and in the amount of effort required to build the knowledge base. So there's at least one dimension not shown—the scale of the problem being solved.

Another dimension is the level of performance that was achieved. Some of these systems are out in the world solving problems routinely, while others are research vehicles that never made it past the laboratory stage. That too helps to account for some of the difference in the amount of time involved.

But with those caveats in mind let's plunge ahead and see what this graph tells us. What's going on? Clearly there is a decline in the time required to build a system. We start off with numbers like thirty or forty man-years and come eventually to numbers only a quarter that size. I would claim in fact that the curve seems to be approaching an asymptote, somewhere in the range of five man-years.

What accounts for the reduction in system construction time? One answer is history: The older programs started first. All other things being equal, the younger programs would have accumulated fewer man-years. But this would predict only a linear decrease.



Figure 9

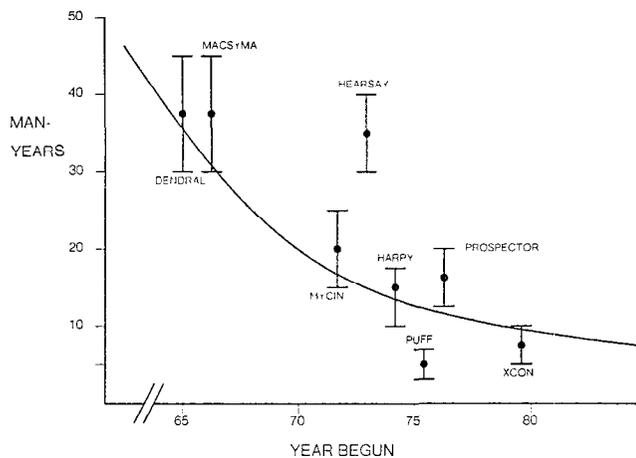


Figure 10

A second, more interesting answer is accumulation—the accumulation of ideas, of code, and experience over the past fifteen years. These aren't unrelated programs—there's a genealogy here: DENDRAL begat MYCIN, which begat PUFF; HEARSAY provided an important foundation for HARPY and then profited from its experience. We get as a result a form of the iceberg phenomenon: a great deal of work goes into the first of these in each line of succession, then the next can be built as a tip on the iceberg/foundation.

A third phenomenon that helps to account for the decline is a matching of tools to tasks. Early on—for DENDRAL and MACSYMA—it was the attempt to solve problems in chemistry and mathematics that led to the development of certain kinds of technology. Later on, as the technology matured, a feeling developed for appropriate applications. In a sense the positions reversed: now the tools are helping us to choose appropriate tasks. The two become as a result more closely matched, hence the shorter development time for more recent systems.

The creation of knowledge base development tools (e.g., TEIRESIAS; EMYCIN, VanMelle, 1980; or KAS, Reboh, 1980) has also accelerated the process of system construction.

But beyond all the uncertainties of the graph, one observation appears to stand out clearly. In all the experience available, the curve never drops below five man-years. I think that sets an important expectation, a very important lower bound. It implies that:

Even for the best-understood problems, experienced researchers using the best-understood technologies still require at least five man-years to develop a system that begins to be robust.

That's important to keep in mind, with all its qualifiers: *the best-understood problems*—the problem must be clear and the knowledge required to solve it well documented (rarely true); *experienced researchers*—like most arts, this one is learned by doing, and the doing provides an important base of concepts and software that speeds subsequent efforts; *the*

STAGES OF DEVELOPMENT

1. System design
2. System development
3. Formal evaluation of performance
4. Formal evaluation of acceptance
5. Extended use in prototype environment
6. Development of maintenance plans
7. System release

Figure 11

best-understood technologies—the Accepted Wisdom must be well-suited to the problem. Fail to satisfy any one of these and the lower bound will climb yet higher.

One final observation about the graph. Everyone I called used the same approach to estimating the amount of time spent: They multiplied the total lifetime of the project by an estimate of the average number of full-time people involved. The interesting thing was that everyone used a number between two and five for the full-time manpower level.

Now that's not terribly profound, but it is a useful order of magnitude indicator. It means that widespread experience agrees that just a few people—between two and five, not 20 to 50—function as the prime movers in building these systems. That's probably not limited to expert systems; it's likely true of large software projects in general. But what does this mean? It suggests that you can't cheat the five man-year boundary. You can't say "It's going to take five man-years, so let's get ten people and we'll be done in six months." That won't work. There are limits to what can be done with division of labor and we encounter one of those limits here.

Stages of development. There's another way to calibrate the effort involved. Several years ago Ted Shortliffe and I speculated on the stages of development in the lifetime of an expert system and came up with the informal list of Figure 11. While these may not be exactly the right stages for every project, they do offer another yardstick useful for calibrating effort.

The important point is that, with few exceptions, no expert system to date has made it beyond Stage 3. Fewer than a half-dozen of the many efforts have made it even that far. Two of them—MACSYMA and DENDRAL—sneaked their way through as the pioneering efforts; they were the "initialization conditions" for the field, and as such, intermixed all the phases noted here.

HEARSAY and HARPY had somewhat clearer development histories and were tested sufficiently to provide a rough calibration of accuracy on 1000-word vocabularies.

Of the systems built to date, R1 had by far the most clearly defined development process, evolving through a sequence of stages similar to those listed here. In its first formal evaluation, the system was tested on approximately 20 cases. The results suggested that R1 would soon solve prob-

lems correctly 90% of the time. This was very encouraging and indicated that the program was ready to be distributed to its user community for more extended testing. But when placed in that setting, users criticized system performance 40% of the time.

Performance had suddenly plummeted to the 60% level. What happened? To be fair, some of those were mistakes on the part of the users resulting from incorrect data or a misunderstanding of program operation. But there was a more basic lesson: research environments, no matter how carefully tailored, are simply not identical to user environments. They differ with respect to the problem mix, how familiar users are with the program, and a range of other factors. As a result, evaluations in the research environment can be at best only rough approximations to the results expected when the program is placed in the user community.

Yet many expert systems developed in the laboratories don't even reach the stage of formal evaluation in the research environment. There's a long road between systems as they are described in conferences like this one and systems ready to be employed by a user community.

In yet another interesting observation about R1's development, McDermott notes that "R1's knowledge grew at least as much during the last stage of its development as in any of the previous four stages." The program appeared to be in its final stage of development, yet the knowledge base was still growing rapidly. What does that mean? It means that a system like this is not built and then polished up and distributed. The process is instead one of constant, incremental growth and improvement of the knowledge base. *Fundamental and important growth of the system is a process that will continue all of its useful life.* The knowledge engineering task is a continual one. The road between systems described here and systems ready for a user community is not traversed in a single journey; it's a continuing process.

One final calibration point arises from characterizing the systems described in conferences like this. Those systems are designed to test out new ideas in knowledge representation, control, acquisition, etc.; they are new sample points in the design space. They are most often Stage 2 systems; rarely have they reached Stage 3. And when the developers say the system "works," *they mean the ideas work.* Typically they have a fragile prototype which has just begun to display credible performance. That alone is generally a significant accomplishment and one with non-trivial implications: it indicates that the basic design of the system is likely to be sound. But it is a long way from being a system ready for distribution and use.

Let me summarize the calibration issue. As the curve suggests, building an expert system means a substantial investment of time and manpower. As the stages of development suggest, there is a long road between the systems described here and a system ready for a user community. As the characterization of research systems suggests, a system described here is only the beginning of a continuing development effort.

Toward the Frontiers

Thus far on our journey we've seen the successes of the Accepted Wisdom, both technical and sociological, and tried to calibrate them. The successes are important and they're numerous, but like all such they have to be understood in the appropriate context and their power carefully measured.

But my central task here is to look forward, to look for the holes in the Accepted Wisdom, the places where it starts to wear thin. To do that, let's take off toward the Frontiers (Fig. 12). Let's see what the Accepted Wisdom can't do and find out what we have to deal with as we get closer to the Frontiers.

Where the Accepted Wisdom falls short: The architectural principles again. One place where the Accepted Wisdom begins to wear thin, interestingly enough, is that set of architectural principles we passed by earlier. Before we pass over that foreboding boundary at the Frontiers of Knowledge, let's take a quick glance backward at that set of principles (Fig. 13).

They have not, alas, proved to be uniformly valid. Consider for example the suggestion that we should separate the inference engine and the knowledge base. There are two situations where that seems not to work out very well:

1. where the "chunks" of knowledge are inherently large, and
2. where we're dealing with basically sequential information.

Consider the issue of large grain size. The concept is not precisely defined, but let me characterize it for a rule-based representation as a rule whose "then" part has a large number of actions to carry out. That is, there are many things we want to do once we've satisfied the precondition part of the rule. How can we encode such information? Well, the obvious way is simply to list each of those actions in the order we want them carried out. But if there are a lot of them, we probably want some control over them. We may want to be careful or clever about the sequence in which they're executed. One option is to further decompose the set of actions into more rules, but there are situations where that would destroy the comprehensibility of the rule. It may hang together as one large chunk of knowledge but become lost in the trees if we chop it into a collection of rules. So we keep it all together by writing long right-hand-sides that begin to look like small programs.

But wait a moment, isn't that what the inference engine is all about? Isn't control of those actions just what the inference engine is supposed to take care of? If we build large action parts in a rule, have we really made that separation?

In domains where the knowledge falls naturally into large chunks, then, we simply may not be able to obey the dictum that the inference engine and the knowledge base should be kept separate.

Now consider domains where the information we want to encode has a strongly sequential character. Let me return

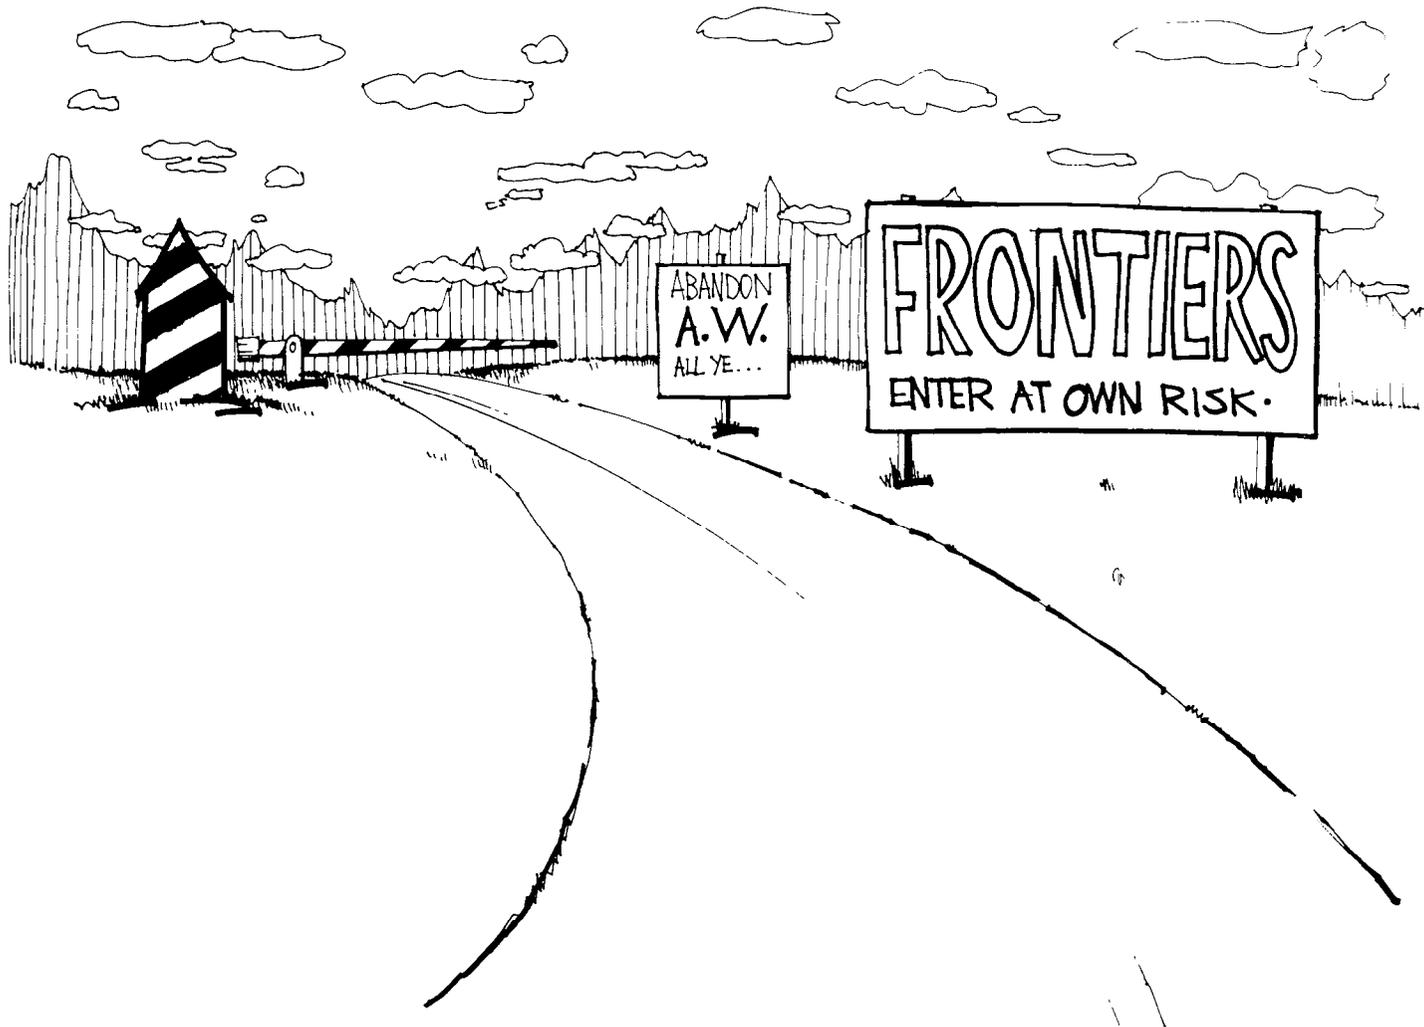


Figure 12

to R1 for an example. As McDermott notes, “R1’s rules can be viewed as state transition operators” and, indeed, that’s what they are. The whole process of configuring a system can in fact be viewed as a passage through a sequence of states.

What impact does this have on encoding knowledge? That depends on the number of such states and how recognizable they are. If there are relatively few states as compared to the number of rules, and if those states are relatively well-recognized, then it is not difficult to separate what to do (the knowledge base) from when to do it (the inference engine).

But if there are almost as many states as rules, and if the state names have no intuitive meaning and begin to look like arbitrary identifiers, well then perhaps there is a better approach. Suppose we label the states with numbers instead of names, and suppose we use a counter to keep track of which state we’re in. Then, to fire the appropriate rule, all we have to do is check the number in the register and fire the corresponding rule. We could call it a “program counter.”

Even though rules written that way look independent,

even though they appear as individual inference steps, in fact they are locked together in a very tight and important structure—the sequence of state transitions that defines what to do next.

Note that the point is not simply whether rules are an appropriate representation here. The point is that some information is inherently sequential. Sometimes what we want to know is that *after you’ve done W, do X, then do Y, then do Z*. The knowledge in such a domain is knowledge of the correct sequence of actions. In such cases the inference engine and knowledge base become nearly indistinguishable.

Our second architectural principle saw virtue in uniformity of representation. Yet there is strong evidence to the contrary: the experience of two of the larger expert systems efforts suggests exactly the opposite. The designers of MAC-SYMA and HEARSAY-II have both commented explicitly on how important it is to have available a range of specialized representations.

In the HEARSAY-II experience, for example, a uniform representation was useful for communication between knowledge sources, but the attempt to enforce the same uniformity

ARCHITECTURAL PRINCIPLES

Separate the inference engine and knowledge base
grain size
sequential information

Use as uniform a representation as possible
MACSYMA, HEARSAY-II experience
Specialization is worth the cost of translation

Keep the inference engine simple

Figure 13

on operations within a single knowledge source "... either failed completely or caused intolerable performance degradation" (Erman et al., 1980). Instead, within a single knowledge source they used representations appropriately specialized to the task at hand. The word recognizer, for example, had its own graph network representation of words; the word sequence recognizer used a bit matrix to represent word adjacencies. The common language was used for writing on the blackboard to communicate with other knowledge sources, but within the scope of a single knowledge source the appropriate specialized representation was used.

The MACSYMA experience is similar: a multiplicity of representations makes possible the efficient implementation of a wide variety of algorithms.

The lesson here is clear: It's so expensive to work with the wrong representation that you're better off working with the right one and paying the cost. Or to put it more positively:

Specialization is often worth the cost of translation.

Our third architectural principle suggested keeping the inference engine simple. Reasonable enough, but if we take heed of the warnings we've just seen and begin to allow a range of specialized representations, we may have a problem. The inference engine has to keep step, that is, it has to be able to deal with each of those representations. And it's not clear that we can maintain simplicity if we allow the representations to multiply.

Where the Accepted Wisdom falls short: Explanation and knowledge acquisition. One reason, then, for marching across the Frontiers is the mixed experience with the set of still-emerging architectural principles. There's another reason as well, and for that I want to go back to my comment about the nature of expertise. As we noted earlier, when the only concern is problem-solving performance, empirical associations can be a powerful technique. Collecting them may be a formidable task, but a large collection of them can be effective.

But what happens when we focus on some of the other behaviors, like explanation or knowledge acquisition? Let's consider explanation and examine the kinds of responses we get from MYCIN. If the system infers the presence of an organism called a bacteroides, and you ask why it believes that,

the response is of the form, "A gram negative, anaerobic, rod-shaped organism is typically a bacteroides." And if you ask why *that's* true... well, in fact you can't ask why *that's* true. I think the system ought to be able to say "Damned if I know," because in fact, damned if anybody knows. It really is simply an empirical observation that seems to hold true.

MYCIN does try to provide an answer of sorts by having justifications for rules. These are usually pointers into the literature where one can discover in medical jargon that "the ultimate etiology is yet to be established" (i.e., damned if they know either).

The point is simple: the use of empirical associations precludes any more substantive form of explanation. If all the system knows is that "A and B suggest C," then it *can't* explain why, beyond repeating the rule. For domains where such rules are all that's currently known, that is indeed the best we can do. But where more fundamental insight is available we should be able to answer more substantively, and we will need representations other than simple associations to make that possible.

Next let's consider knowledge acquisition. I think there was a very interesting lesson learned several years ago in work on Meta-DENDRAL (Buchanan and Mitchell, 1978). The task was to induce new rules about mass spectrum analysis from examples of spectra and their analyses. One of the clear lessons from that work was that some model for understanding events in the domain was crucial for guiding and grounding the induction process. The model needn't be elaborate; Meta-DENDRAL needed little more than the ball and stick model of chemistry. Yet that model offered a simple way of understanding *why* certain molecules fragmented in certain ways, and, equally important, why some other fragmentations were highly unlikely. In doing so it provided a crucial filter capable of distinguishing plausibly meaningful hypotheses from those that were likely to be mathematical coincidences in the data. Without such a filter, the number of hypotheses generated is simply much too large; the information cannot be distinguished from the coincidental noise.

Once again, pure associations fall short. If we wish to be able to learn about the domain, we need some understanding that pushes beyond that simple form of knowledge.

Across the Frontier

We now know a little about what empirical associations can do for us. In some domains they can provide an interesting and useful level of performance, offer a reasonable level of explanation, and support a primitive form of knowledge acquisition. As we have seen, they begin to display important weaknesses in dealing with the last two of these, and offer virtually no support for the remaining varieties of expert behavior noted earlier. In domains where expertise really is simply a large collection of such associations, this is likely to be the best we can do.

The question then is, what can we do in domains where

more than that is available? What can we do for problems where we don't have to be satisfied with a statement like "That's just the way it is"? In particular, what can we do in domains where there exist mechanistic or causal models?

Work on such questions is proceeding in several places. At MIT for example, a group of us is working on the problem of computer diagnosis in the literal sense. That is, how can we detect and identify malfunctions in computer hardware? A group at Stanford is working on a similar problem of hardware diagnosis, while another group at MIT is developing a similar approach to medical diagnosis. In all these cases, the effort is based in part on understanding the structure, function, and design of a device well enough to use that knowledge to guide the diagnostic process, whether the "device" is a RAM chip, a communications multiplexer, or a kidney.²

Let me illustrate this by considering our domain of digital hardware troubleshooting. It's clear there's much to know. Part of the expertise of a good troubleshooter lies simply in knowing a lot about the device at hand.

What is the character of that knowledge? One kind of information is knowledge about structure—the "anatomy" of the device. We encounter it in forms ranging from block diagrams of an entire machine down to schematics indicating the devices and wires making up a single register. Information about function—the "physiology"—resides in such things as our understanding of signals propagating along wires, our model of the operation of a RAM cell or multiplexer, or the function of an ALU. Knowledge of design comes from an understanding of the intent of the designer, the purpose and intended function of each component.

Because the machine is a human-designed artifact we have, at least in principle, relatively easy access to such information. Such is not often the case elsewhere: in medicine, for example, our knowledge of human anatomy and physiology is still growing, and Lord knows we can't get our hands on the design specs.

An Example

I'd like to use the rest of this talk to explore in some detail the kinds of technical advances that will be necessary in constructing the next generation of expert systems. To do that we'll walk through a sample problem from the domain of field engineering, paying careful attention to the kind of knowledge needed.

Let me give you a preview of the issues we're going to encounter. First, I want to use this example to indicate how the existing expert systems technology falls short of providing the machinery we will need. We will see in several

places that a large collection of rules is simply not useful or convenient enough to be a feasible approach to the problem.

Second, we'll see the need for a range of specialized representations. As we'll see, there are too many different kinds of knowledge required to make feasible a uniform representation scheme.

Third, we'll need multiple representations. We're going to want to talk about electrical configuration so that we can reason about electrical connectivity and current flow. We're going to want to talk about the physical configuration in order to reason about where something is or how it's connected physically to something else. We're going to want a representation of functional organization as well, since, as we'll see, we reason in large measure by observing behavior and making inferences about function.

Finally, we'll see a very simple but instructive example of what can be done when we have available something more fundamental than empirical associations. We'll see that even some very simple mechanistic and causal models can offer significant reasoning power.

There are two caveats to keep in mind as we proceed. First, our initial efforts are focused on non-transient hardware errors. That is, we do not attack the problem of distinguishing between hardware and software errors, and will not be dealing with transients. Eventually both of these will have to be considered, but we focus first on the more tractable cases.

Second, this example indicates our ultimate goal, not our current capabilities. We have just begun building a system capable of the kind of representation and reasoning we're about to use. The example illustrates the state of our aspirations, not the state of our art. The example is, however, quite real and the details are accurate.

The problem.

The phone rings at the local field service center. .

"Field Service Can I help you?"

"Yeah, my machine doesn't work!" says a frantic voice

"What seems to be wrong?"

"Well, when I sit at the console terminal and load the system, everything's just fine. The hardware powers up, the operating system loads, and it's clear that the operating system is running. From the console terminal everything is just fine. But the user terminals don't work. They're turned on and they look OK, but when you type something, absolutely nothing happens. So what's wrong?"

Think about that for a minute. What *is* wrong? More precisely, *What kinds of information do we need in order to determine the answer to the question?*

Clearly, in order even to get started, we need some information about the *structure of the system*. Figure 14 presents the configuration of the system we're talking about. There are two main busses with an interface between them; a disk, the cpu, and memory hang off bus 1, with the user terminals on the other.

²All of these efforts have benefited substantially by the foundation provided in work like that of deKleer (1979), Sussman (1978), and Rieger and Grinberg (1978). Their ideas about representing and using knowledge about structure, function, and design have provided an important foundation for thinking about causality, qualitative reasoning, and related issues.

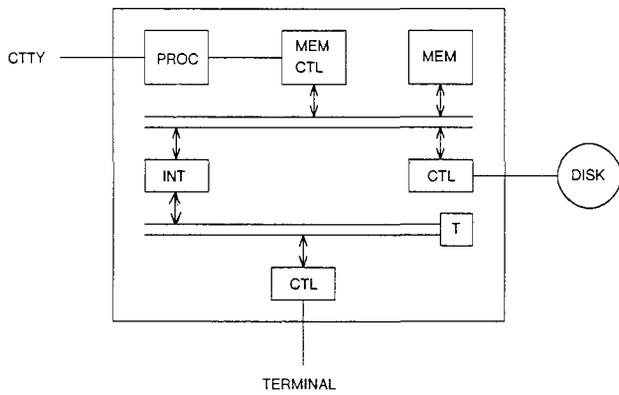


Figure 14

What does this give us? It supplies the anatomy of the system at the coarsest level of detail. And what can we do with it? Well, not too much until we supply some information about the *function* of the various components. For example, disks and memory are used to store information, busses transmit information, and cpu's execute instructions.

Now we can begin to use that information about structure and function to reason about the *behavior* of the system. Loading the system means reading the code off the disk into memory, then executing that code. Given the structure indicated by the diagram, we know that that process must involve the disk, its controller, bus 1, the MEM, memory controller, and cpu.

We can annotate our "anatomical diagram" to indicate our conclusions about how well the components are working. For example, the system load succeeded and that tells us quite a lot. It means the transmission of information from disk to memory must have succeeded, so we believe the source (the disk), the sink (memory), and the transmission medium (bus 1) are all functioning properly. The system runs, meaning that the hardware is capable of executing a program (the system code), hence the the path from memory to cpu must be functioning. Thus the memory, bus, memory controller, and memory control to cpu link are all probably functional. Finally, since the program in question—the system itself—is a large program with a lot of nontrivial code, we suspect that the cpu itself is working properly. Since the system clearly responds to the console terminal, we suspect that is working properly as well.

And what of the problem with the user terminals? Once again, *an understanding of behavior derives from reasoning about structure and function*. The terminals work in full-duplex mode, meaning that they can be viewed in the simplest terms as a transmitter and receiver, sending characters to the cpu and printing the characters echoed back. If they don't work then there must be something wrong with the source (the terminals), the sink (the cpu), or the transmission path. On those grounds alone we ought to suspect the terminals themselves, their controller, bus 2, the bus interface, bus 1, memory control, and cpu. But we already

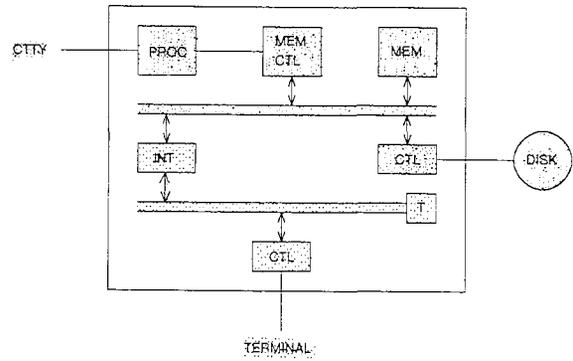


Figure 15

--- probably good
 --- suspect

have pretty good evidence that the last three of those (bus 1, memory control, and cpu) are all functional, so we contain our suspicions to the first four. The results of this reasoning process are summarized in Figure 15.

We can begin by narrowing our focus. All of the user terminals are malfunctioning identically (no response). We have either an extraordinary coincidence, or more likely, the problem lies "upstream," at a point in the data path common to all the terminals. We can thus direct our suspicions to the bus interface, the second bus, and the terminal controller.

Conveniently enough, there is a set of diagnostic programs designed to test much of this hardware. Life isn't always so easy (as we'll see below), but it is useful to be able to take advantage of diagnostics where they already exist.

So we retrieve the relevant set of diagnostics and run them. The results are shown in Figure 16 (Tests 1-5 gave no error messages, diagnostics were halted after Test 8). What do they tell us? Once again the relevant issue is really *What kinds of information do we need in order to determine the answer to the question?*

DIAGNOSTIC RESULTS				
TEST 6				
EXPECTED	AAAA	3333	F0F0	FFFF
RECEIVED	AAA8	3331	F0F0	FFFD
TEST 7				
EXPECTED	0000	0000		
RECEIVED	0031	000D		
TEST 8				
EXPECTED	0A0A	0312		
RECEIVED	0A08	0310		

Figure 16

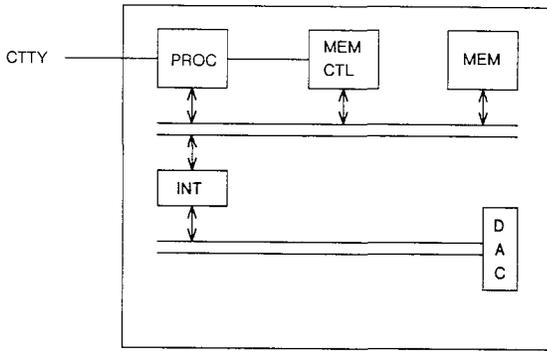


Figure 17

We need to know at least two things:

- 1 the *intent* of the test that encountered the error, and
- 2 the *behavior* of the device responsible for the error.

The first of these is important because the error messages make no sense unless we know what the diagnostic was trying to do. The second is important in allowing us to determine how the device under test may have failed.

If we're lucky, the intent of the test is made clear in the documentation; if we're not so lucky we may have to read the code. Clearly, no program we build anytime soon is likely to be able to do either of these. What reasonable substitute is there?

It appears that a reasonable substitute might be constructed from a suitably annotated, simplified view of the hardware. Since the author of a diagnostic is typically trying to test just a single part of the machine, much of the remainder becomes temporarily irrelevant. We can take advantage of this to provide a much simplified view of the sys-

tem. Consider Figure 17, for example. We've ignored the system disk, terminals, and terminal controller, and added a level of detail to the terminator, which for our current purposes can be viewed as three registers.

We then describe the diagnostics in terms of operations performed on this simplified machine: The tests in question attempt to send bit patterns back and forth between the processor and each of the terminator registers.

Now that we know what the test is trying to do, we still need to know something about the behavior of the device being tested—a register. For our purposes a register can be viewed as a single memory cell (Fig. 18). To make sense of the error messages, we need to understand the behavior of that device. In particular, we need to determine how its behavior changes as a consequence of various forms of failure.

To do that, we need a reasonably detailed model of how it works, a model capable of supporting the sort of reasoning that says "if the write-enable line were stuck low, then no new information would ever get into the cell, and every time we read it the same thing would appear." If we could do that, then we could determine that the symptoms in Figure 18 are not consistent with errors in any of the control lines, but *are* consistent with a dropped bit on the data-in or data-out path. We have in effect a simple induction problem, but we can guide the induction process by using our understanding of device function to constrain the set of plausible hypotheses.

We can thus make sense of the error messages by relying on both an understanding of what the diagnostic is trying to do and a model of the behavior of the device being tested. The result, of course, is the hypothesis that Bit #1 is being dropped somewhere on the path from the cpu to the data register (the register under test at the moment). Notice that the hypothesis is consistent with all the results of this test,

DIAGNOSTIC RESULTS				
	TEST 6			
EXPECTED	AAAA	3333	FOFO	FFFF
RECEIVED	AAAB	3331	FOFO	FFFD
	TEST 7			
EXPECTED	0000	0000		
RECEIVED	0031	0000		
	TEST 8			
EXPECTED	0A0A	0312		
RECEIVED	0A08	0310		

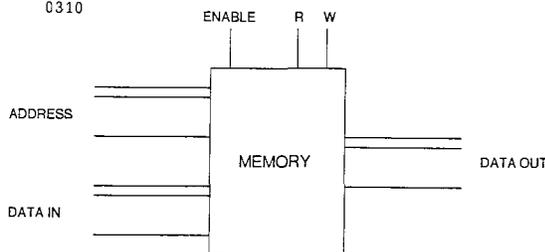


Figure 18

both the failures (where expected and received differ) and the successes (expected and received are the same).

But what can we make of Test 7? For the moment, nothing. Test 8 is once again consistent with our hypothesis. Can we come up with an hypothesis that accounts for all the data? Yes, if we invoke one reasonable heuristic: each diagnostic in a sequence is strongly dependent on the success of the previous tests. Like the syntax checks in a compiler, if one test fails, the results of successive ones can be misleading or nonsensical. So our overall guess is that we are dropping Bit #1 and that the results of Test 7 can be chalked up to the vagaries of diagnostic interdependence.³

To summarize our progress thus far: the initial set of symptoms suggested a problem in the communication path from the user terminals to the cpu. The fact that the system ran from the console teletype suggested that the cpu, memory, and the bus connecting them was functional. A prepackaged set of diagnostics was used to test part of the path from user terminals to cpu; the diagnostics indicated that the bug was a dropped bit.

So we have a hard error that's been found and characterized by the diagnostics. It's often suggested that a problem that well identified is effectively solved. But let's continue on, because I think it will become clear that there is much non-trivial reasoning left to do.

Take another look at Figure 15. At the moment we know only that we're dropping a bit somewhere on the path between the cpu and the data register of the terminator on bus 2. But there's a fair amount of hardware along that path, so we have to do some signal tracing to find the source of the problem.

Before we go to the trouble of probing, there are some easy inferences we can make. The diagnostic gives the same results for all three of the registers in the terminator. As with the user terminals, either all three are broken in an identical fashion, or the problem lies further back along the data path. Occam's razor once again suggests the latter.

From that we get the view of Figure 19, with the region of suspected hardware narrowed down a little. We no longer suspect the terminator and we've focused in on one bit of the data path.

In keeping with that narrower focus, we need a more precisely focused test. The diagnostics are effective as broad tests of large parts of the machine, but now we need a finer probe capable of testing just that part of the machine that we suspect is malfunctioning.

We want to test one bit of the datapath, but there are really two routes—one heading out to the terminator and one heading back. They overlap substantially, but do differ in places. We can test the outbound path by writing to the terminator, the inbound path by reading from it. Which shall we do? The choice seems clear—we ought to write from the cpu to the terminator—but the reasoning is slightly subtle. We want to put a 1 on Bit #1 of the datapath. Since

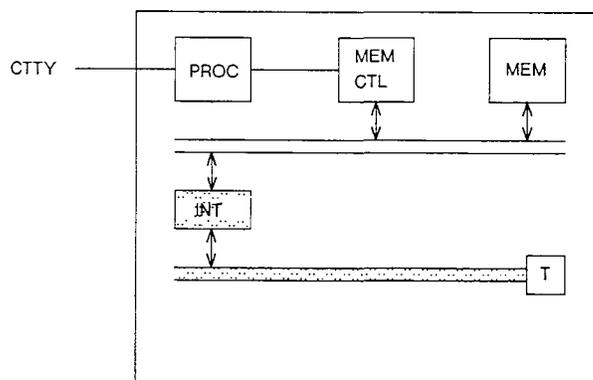


Figure 19

we never successfully wrote and read back a 1 along that bit, we don't know what's in the terminator register, hence reading from it is not a predictable test.

We might capture some of this reasoning in a general principle like "when faced with a choice use hardware known to be functional." But a more satisfying answer would tell us how to generate appropriate tests based on the sort of reasoning about structure and function at issue here.

Having decided to test the outbound path, we put the cpu in a loop that continually writes a 2 to the data register and proceed to trace the signal. Figure 20 presents an expanded view of the path, running from the point where bus 1 meets the bus interface, through the interface, along the backplane, and out to the data register. We've moved down one level of detail and have focused on one bit slice, but we're still dealing with an abstracted view of the hardware.

The hexagons are backplane pins; they're notable primarily because they're easily accessible test points. That's a standard piece of diagnostic reasoning of course—use the easy, "non-invasive" tests first, since the more complex tests often have the risk of serious side effects

Just to make sure we're on the right track, we probe first at B7, and find the 1 is indeed making it this far (more evidence that we were correct in assuming that the cpu, memory, and bus 1 are functional). We probe AF2 and find that the bit is missing; probing C12 shows the same thing. We've managed to narrow the focus still further—the signal is getting lost somewhere between bus 1 (pin B7) and bus 2 (pin C12).

Comparison with the Accepted Wisdom. One of the central points of this talk is that the Accepted Wisdom would have a difficult time supporting some of the reasoning we need to do. Let's pause for a moment and to explore that point.

Consider the signal tracing we've just done and imagine trying to do it with the Accepted Wisdom. One plausible way would be to write a set of If-Then rules, that might say

³As we'll see, this is wrong.

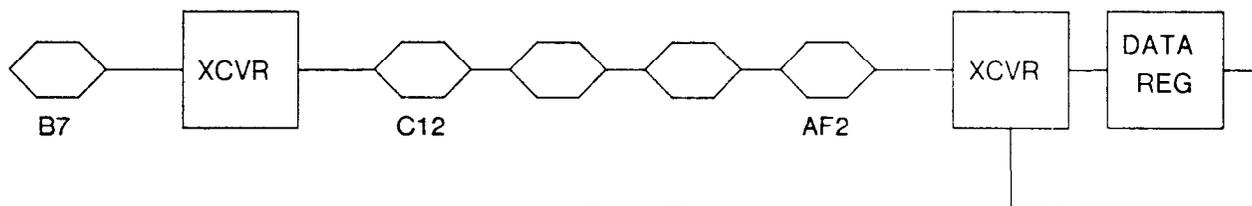


Figure 20

something like

If the signal is OK at B7 and
the signal is blocked at AF2
then the signal is being lost somewhere between B7 and AF2.

If the signal is OK at B7 and
the signal is blocked at C12
then the signal is being lost somewhere between B7 and C12.

That would work and it certainly captures the effect of the signal-tracing logic we used earlier, so we can hardly argue that the Accepted Wisdom is incapable of such reasoning. But it does have a number of problems. In order to do signal-tracing, we need to know about the structure of the hardware. The rules above contain such knowledge, but it's bound up with the art of signal-tracing and therein lies the source of many problems.

In particular, the knowledge about structure is represented only *implicitly*, it is *inaccessible* for all purposes other than signal-tracing, it is *hard-wired* into the system and inflexible as a result, and it fails to recognize signal-tracing as a distinct skill.

To see that the knowledge of structure is implicit, consider the rules above when viewed in isolation. The rationale behind them is completely mysterious until we refer to something like Figure 20.

To see that the knowledge about structure is inaccessible, consider trying to use it to do simulation. We ought to be able to predict the path of a signal inserted at B7, yet we can't do that without "unraveling" the information about structure embedded in the rules.

To see that the knowledge is inflexible, consider using exactly the same signal-tracing technique on another piece

of hardware. We would have to write out a new set of rules; none of the existing rules would be useful. Yet the underlying skill would be entirely the same.

Finally, the rules miss the sense of what signal-tracing is all about. It's a process of examining structure, but it is distinct from that structure in much the same sense that an interpreter stands apart from the program it interprets

Well, you say, no one would have suggested trying to encode that in rules anyway. What would they have done? Presumably they would have written a set of rules that tried to capture just the signal tracing skill, and had a separate description of structure. Yes, that's just the point. And if we go to the trouble of building a separate structure description, then we ought to do it right, which means a hierarchy of multiple levels of detail, information about connectivity, substructure, etc.

The point is simply that the Accepted Wisdom *focuses on the use of rules embodying empirical associations*. It does not offer us any tools for constructing structural descriptions of the sort we need, it does not offer us any techniques for using those descriptions to guide diagnosis, and perhaps even more important, *it does not even lead us to think in such terms*.

The example continued. Let's return to the example. We've determined that the problem lies somewhere between B7 and C12, so let's look at that in more detail (Fig. 21). Bit #1 comes in at B7 on the left, passes through the gate array, into the transceiver in chip E4, and out again to pin C12. We're losing it somewhere between those two points.

How can we locate the source of the problem? There's more signal tracing to be done, but once again let's try to be as clever as we can before we go probing.

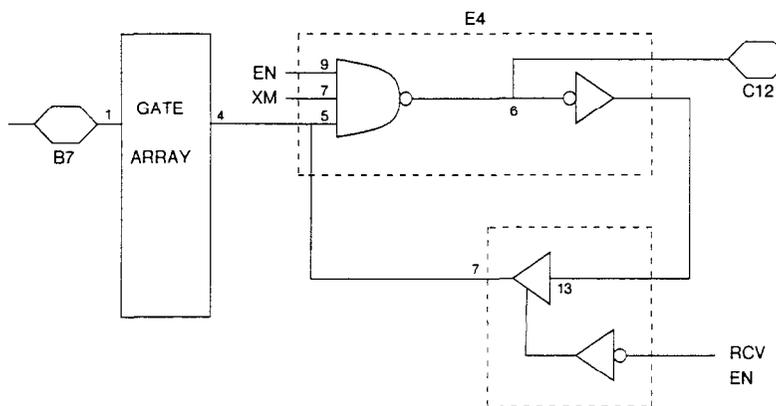


Figure 21

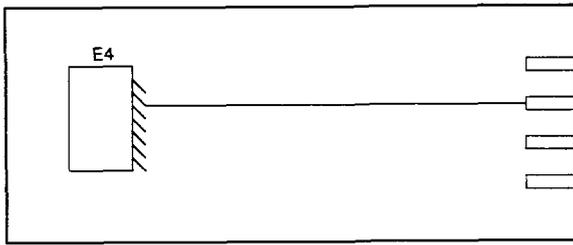


Figure 22

Experience tells us that gate arrays have a long time to failure, so we expect that the signal will make it through to the transceiver. Let's also assume that, unless we have evidence otherwise, wires are functional; if a signal appears at one end, then it will appear at the other end.

If the signal makes it to the transceiver, what might prevent it from emerging? If the ENable or XMit signals on the Nand gate were missing, then the bit would be blocked. But it turns out that both the ENable and XMit signals are shared across four bits of this data path. Hence if either one of them were bad, we would have noticed four bits dropped, not simply one.

Note the hypothetical reasoning we used. We postulated a fault and then asked whether that was consistent with the symptoms noted. That's a useful and common enough form of reasoning, but is more subtle than it might appear. We can't, for instance, simply insert a fault and then ask "How does the machine behave?" There's too much hardware there and too wide a variety of behavior to ask such a broad question. We need instead a very directed and constrained form of simulation that asks how a fault would effect the particular piece of behavior under study. In this case that tells us that the ENable and XMit lines are likely to be functional.

So it seems that the signal probably made it from B7 through the gate array and through the Nand gate of the transceiver. By checking pin 6 of chip E4 we can easily verify all of this, and indeed the signal is good there.

Well then, what's going on? We said earlier that we will assume wires are functional and propagate signals. From the diagram we have nothing more than a wire between chip E4 pin 6 and backplane pin C12, so we had better check our assumption about it being good. To do that, we need to *shift representations*, and consider the *physical* structure underlying the path between pin 6 and pin C12. That structure is suggested in Figure 22—there is an etch on the board that runs from pin 6 out to the pad at the edge. We can easily check that with an ohmmeter and it proves to be good.

Going back to Figure 21, then, we seem to have a bit of a contradiction on our hands. We have the signal getting to C12 on the diagram and yet it doesn't get to the "real" C12 pin on the backplane. To unravel the problem we once again have to shift from the predominantly electrical view of Figure 21 to a physical view (Fig. 23). That figure makes it

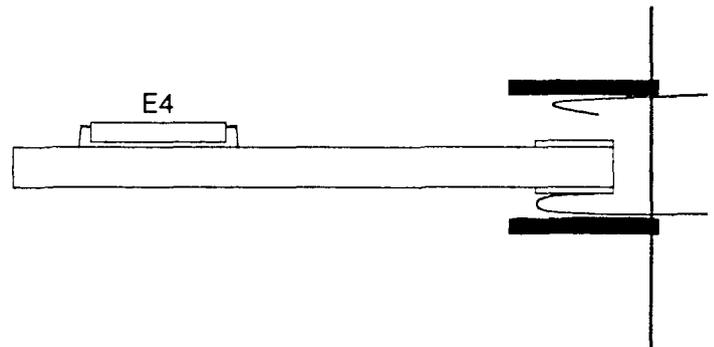


Figure 23

clear that the board fits into a slot of the backplane. The "physics" which underlies the electronic structure makes it clear that between the board and the backplane pin there is a set of spring-loaded connectors. And in this case one of them is bent and not making contact. There's our dropped bit.⁴

The example: Summary. The example, even as simple as it is, helps to illustrate a number of points about this undertaking (Fig. 24). First, the central thrust of this work is to do diagnosis based on an understanding of structure and function. From the outset, knowledge of function (loading the operating system, typing on a terminal) and structure (system configuration) interacted to allow us to infer the likely health of various components of the system. We were able to focus our efforts because we knew how devices were interconnected, knew how they were supposed to behave, and knew something about how they were likely to fail.

This understanding also led to a more precise diagnosis than might have resulted from the traditional approach of module swapping. Early in the diagnosis it became clear that the bus interface was dropping a bit. Had we replaced the existing board with a new one, at best that would have failed to solve the problem. At worst, the act of removing and replacing a board might have jiggled the connector sufficiently that we got a temporary contact. The same problem might then reoccur, while in the meantime a perfectly good board was sent to be repaired.

Second, we illustrated how even very simple causal models might be used as a way of "understanding" behavior. A model of the function of a RAM cell, for example, made it possible to reason about the likely cause of malfunction.

⁴In truth, it was actually a piece of tape that had been inserted to simulate the fault. But these connectors do, in fact, get broken or bent out of shape by people who are a little too rough in removing and reinserting boards (as we occasionally have been).

And Test 7, whose error message didn't fit into the dropped bit hypothesis? As it turns out, there is a bug in that diagnostic (someone wrote a multiply where they meant a left-shift). This is yet another reason why it's important to be able to interpret test results in context, not just accept them.

With the bug fixed, results of Test 7 are consistent with the others.

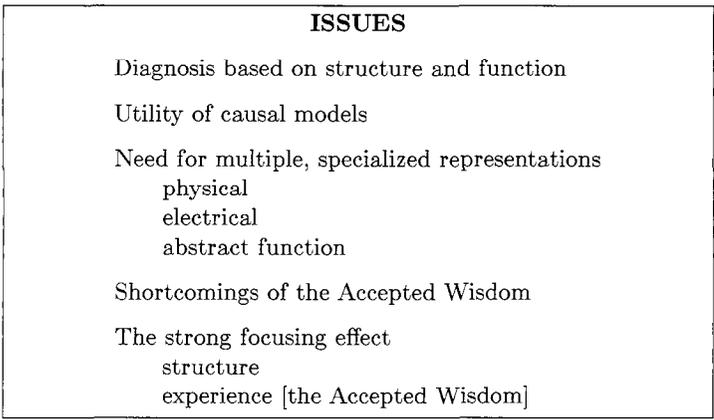


Figure 24

Third, we saw how a multiplicity of specialized representations aided the reasoning process. We used several different views of the machine, including an electrical (the schematics), physical (the board in its slot), and a more abstract functional view (boxes labeled “transceiver,” etc.). Each in turn was useful in understanding how the system worked.

Fourth, we have illustrated how the Accepted Wisdom, with its focus on collecting large numbers of empirical associations, fails to provide some of the intellectual framework and tools required for solving this problem. In particular, it fails to supply a mechanism appropriate for describing and reasoning about structure.

Fifth, the example demonstrates the interesting phenomenon illustrated in Figure 25. Consider the enormous focusing that occurred as our analysis proceeded. We went from “there’s something wrong with the machine,” to “there’s one broken connector in the backplane that’s causing one bit to be dropped.” How did we manage to focus in so sharply in relatively few steps? Clearly the structure

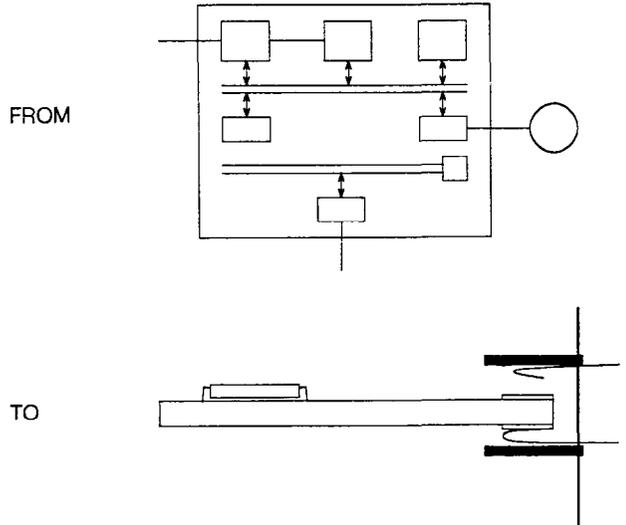


Figure 25

of the machine itself was in part responsible: Computers are traditionally conceived of in terms of successive levels of structural and functional organization, and we put that to good use.

But part of the responsibility lies in a contribution made by the Accepted Wisdom, in the form of empirical associations and “compiled experience” about where to go looking. It’s clear, for example, that a troubleshooter doesn’t respond to a problem by writing down Ohm’s law for the entire machine and trying to account for the malfunction in those terms. Instead what he appears to do is to ask “Where have I seen this before and what was the problem last time?” And when we had some guess about where the problem might be, we retrieved a set of diagnostics simply because we “knew” they were appropriate for the part of the system under examination. No deeper reasoning was involved, nor was it necessary. The Accepted Wisdom may thus have a role to play in this problem, functioning as a focusing mechanism. It helped us to zero in on components of the machine that were appropriate for more detailed examination.

Challenges. Since this work is still in its early stages, we have only begun to construct the basic foundations of the mechanisms needed to produce the required reasoning behavior. Let me comment on a few of the more important challenges we face.

Representing structure and function is an interesting task, in part because of the range of levels of detail we encounter. We would like to describe both of them at levels ranging from the basic system configuration—where “cpu” or “memory” are basic elements—on down to the most detailed levels like the logic diagram, where individual gates, wires, and pins are the basic elements.

A number of structure description languages have been written and that accumulated experience may provide a good starting place on the problem. Behavior is a bit more difficult. Several fundamentally different approaches have been tried, ranging from state transition diagrams, to petri nets, to input/output specifications, and chunks of code. No clear winner emerges and indeed the appropriate language will likely depend on the purpose at hand.

It’s easy to suggest that we’ll have to use multiple, specialized representations. It’s somewhat harder to determine the appropriate representation for each and to know which one to use when. We used a simple strategy in the problem above of starting with the functional organization, using the electrical organization next, and employing the physical representation last. This seems to be a reasonable default, since it reflects, to first order, at least, the levels of abstraction in the design of the machine. Functional organization provides the framework within which an electronic implementation is designed, and the electronics provide a framework for the physical organization. But this is at best a first order approximation. Life is rarely that well organized in design; we have no reason to expect it will be that easy in diagnosis either.

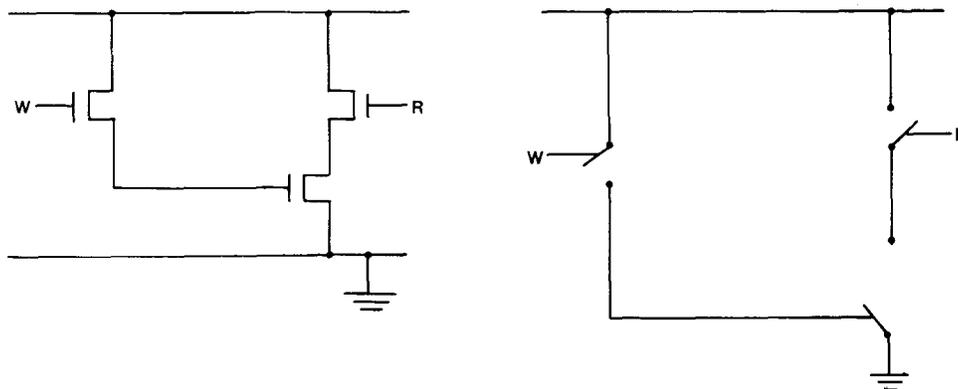


Figure 26. VLSI memory cell.

One of the longer term challenges we face is reconciling increased complexity of these systems with the desire, well-articulated by the Accepted Wisdom, of maintaining transparency and flexibility. Structural and functional models may increase the power and depth of understanding displayed by the system, but they will inevitably increase the complexity as well. How then can such systems continue to be transparent to the user and easily modified? Without transparency, they are unlikely to be widely accepted and our efforts may be in vain. Without the flexibility, they are likely to ossify long before they are fully developed, and we run the risk of never being able to finish.

One of the convenient bits of handwaving I did in the example was to dip occasionally into reasoning via empirical associations, rather than the reasoning from structure and function on which most of the example was based. As noted, both can be useful, but the obvious question is, how are we to determine which to use when? The empirical associations can supply a quick and crude guide about where to go looking, but how do we know when to shift over to more detailed mechanistic models?

Finally, there is an interesting and difficult problem of selecting the appropriate level of abstraction. In the example above I selected just the right level for each step of the reasoning, providing a powerful way of simplifying the reasoning. Consider, for example, the enormous amount of detail that is suppressed in our initial view of the system (Fig. 14). But it isn't obvious how this selection process could be automated. How do we know what level of detail is sufficiently simple that it makes reasoning easy, yet sufficiently detailed that it isn't misleading?

To illustrate just how challenging this problem can be, let me show you another example. To do that I'm going to shift domains for a moment and take a problem from design rather than diagnosis, but I think it illustrates the issue nicely.

Figure 26 shows a VLSI design for a single bit of random access memory. To see how this works, consider the simplified diagram on the right. To store a bit, put the appropriate logic level on the wire at the top and close the

switch labeled "w" ("write"). If the logic level is high, that will close the switch at the bottom labeled "b"; if the logic level is low, that will open the switch. The bit is thus stored as the state of the switch at the bottom.

To read the information out, we have to charge up the wire at the top and then close the "read" switch (labeled "r" in Fig. 26). If the charge drains away to ground then the switch at the bottom must be closed, hence we must have stored a 1 there. If the charge stays around, then the switch must be open, so we must have written a 0 there.

The design sounds perfectly plausible at that level of description, but will it work? Transistors are considerably more sophisticated than the simple switch we've been using to model them. A number of more complex models are available, so one approach would be to try simulating the design with successively more complex models. We can model it as a switch with some inherent resistance, we can add capacitances, we can view it as a voltage-controlled current source, and all of these more complex models will still suggest that the design is sound.

But the device doesn't work; the design is indeed faulty. It doesn't work because the transistor, in addition to all the other things it is, is also a leaky switch. It never quite turns off completely.

When we write a 1, we charge up the base of transistor b, then let the write line go low, turning off the write transistor. But it doesn't turn off completely, and eventually enough charge leaks off the base of b that it is no longer "on." The bit has decayed away.

The leak of charge through the write transistor is, for almost all purposes, too small ever to worry about. Typically we deal with time scales on the order of nano- and microseconds, and on that time scale the leak doesn't amount to much. But remember, this device is a memory. And we'd prefer it retain what we told it, not just over the course of microseconds but over seconds, hours, and days. And on that time scale, the leak is big enough that it can result in loss of information.

What does all that illustrate? It shows that even in a very simple design, it may not be easy to determine what

level of model to use. It isn't simply complexity that makes necessary the more elaborate models; it is instead something more like sensitivity. We have to ask what the device is most sensitive to. And in this case, it turns out to be enormously sensitive to an effect which is, in almost all other situations, so small as to be negligible.

Summary

And it is here, at the Frontiers of Knowledge, that we come at last to the end of our journey. What have we seen? We started out in the Land of Accepted Wisdom, where we explored the nature of expertise, finding it to be a whole collection of interesting behaviors, of which only a very few have begun to be automated. We surveyed the extent of the Accepted Wisdom, finding it effective where expertise consists of large numbers of empirical associations, "compiled experience."

We then stepped back to evaluate the state of the art. Two important calibration points were

1. the magnitude of manpower investment necessary to build a robust expert system, and
2. the stage of development typically reached by most current expert systems.

Data from existing efforts, though meager, seems to suggest that, even in the best of cases, at least five man years' worth of effort is necessary before an expert system even begins to perform reliably. It is also revealing to note that most expert systems to date have been developed only through the stage of construction of the basic knowledge base. Few have proceeded to the extended testing, further development, documentation, etc., that must occur before these systems are ready for distribution to a user community.

In approaching the Frontiers, we uncovered some of the shortcomings of the Accepted Wisdom, finding that some of the accepted architectural principles may not in fact be as feasible or desirable as first expected. It proves difficult, for example, to separate the inference engine and knowledge base when knowledge in the domain is inherently sequential or where the chunks of knowledge are large. Uniformity of representation can exact a considerable price, while the efficiencies of specialization can often more than offset the cost of translating between representations.

In journeying across the Frontiers, we explored a problem from the domain of hardware troubleshooting to illustrate the kinds of challenges that we face in constructing the next generation of expert systems. We found that the Accepted Wisdom was capable of encoding some of the knowledge we needed, but discovered that using simple rules can render the underlying knowledge inaccessible, hard-wired, and inflexible. This illustrated the utility of being able to represent and reason about structure and function. We saw that even simple causal models can provide a mechanism for "understanding" behavior of devices. We used a number of different, specialized representations that included physi-

cal, electrical, and abstract functional views of the device. And finally, we saw that while all these ideas can be an effective foundation for diagnostic reasoning, there is a whole range of difficult and interesting problems to be solved before we travel very far across the Frontier.

And with that, my young friend and I will remove our official Tourguide hats, take our bows, and depart.

References

- Buchanan, B. G., & Mitchell, T. M. (1978) Model-directed learning of production rules. In D. Waterman & F. Hayes-Roth (Eds.), *Pattern-directed inference systems*. New York: Academic Press.
- Davis, R. (1979) Interactive transfer of expertise: Acquisition of new inference rules. *Artificial Intelligence* 12:121-157
- deKleer, J. (1979) Causal and teleological reasoning in circuit recognition. Tech. Rep. MIT AI-TR-529, Computer Science Dept., Massachusetts Institute of Technology
- Erman, L., Hayes-Roth, F., Lesser, V. R., & Reddy, R. D. (1980) The HEARSAY-II speech understanding system: Integrating knowledge to resolve uncertainty. *ACM Computing Surveys* 12:213-253.
- Feigenbaum, E. A. (1977) The art of artificial intelligence: I. Themes and case studies in knowledge engineering. *IJCAI* 5, 1014-1029.
- Kunz, J. C., Fallat, R. J., McClung, D. H., Osborn, J. J., Votteri, B. A., Nii, H. P., Aikins, J. S., Fagan, L. M., & Feigenbaum, E. A. (1978) A physiological rule-based system for interpreting pulmonary function test rules. Rep. No. HPP-78-19, Computer Science Dept., Stanford University.
- Lindsay, R., Buchanan, B. G., Feigenbaum, E. A., & Lederberg, J. (1981) *Applications of AI For Organic Chemistry: The DENDRAL Project*. New York: McGraw-Hill.
- McDermott, J. (1980) R1: The formative years. *The AI Magazine* 2(2):21-29.
- MACSYMA group (1974) The MACSYMA reference manual. Computer Science Dept.: Massachusetts Institute of Technology.
- Mitchell, T. M. (1978) Version spaces: An approach to concept learning. Rep. No. CSD-78-711, Computer Science Dept., Stanford University. (Doctoral dissertation.)
- Reboh, R. (1980) A knowledge acquisition environment for expert consultation systems. Ph.D. dissertation, Dept. of Mathematics, Linköping University, Sweden.
- Rieger, C. J., & Grinberg, M. (1978) A system of cause-effect representation and simulation for computer-aided design. In J. C. Latombe (Ed.), *Artificial intelligence and pattern recognition in computer-aided design*. New York: North-Holland
- Shortliffe, E. H., (1976) *Computer-based medical consultations MYCIN*. New York: American Elsevier.
- Sussman, G. J. (1978) Slices: At the boundary between analysis and synthesis. In J. C. Latombe (Ed.), *Artificial intelligence and pattern recognition in computer-aided design*. New York: North-Holland.
- Swartout, W. (1981) Explaining and justifying expert consulting programs. *IJCAI* 7, 815-822
- VanMelle, W. J. (1980) A domain-independent production-rule system for consultation programs. Ph.D. dissertation, Dept. of Computer Science, Stanford University.