

EXPERT SYSTEMS AND KNOWLEDGE BASE - PART 2

TONY STARFIELD: So now you have seen a small expert system shell, and you've seen how a knowledge base can be fed into it, and what the interaction is between a user and a knowledge base. And you will recall that the knowledge base had only seven rules. So nobody can argue about the fact that it was simple. The question we have to ask is, "Was it trivial?" And the answer, I think, is, "No." And I can give you two reasons why I don't think it's trivial.

The first is that if you take somebody who doesn't know much about burning, and you get them to sit and run through that particular knowledge base, have three or four different interactions, answering the questions in different ways, they're going to get a pretty good understanding of how one goes about making these decisions to burn or when to burn parts of the Kruger Park.

My second evidence is this. At the time that that knowledge base was developed, there happened to be two schools of burning in South Africa. And whenever there was a conference on fire or grassland management or something like that, people who adhered to the two schools would argue past each other. They would reach

dissimilar conclusions, but they would not be able to understand why they disagreed. It so turned out that the person whose expertise went into the knowledge base that you saw had studied at a university where one school of burning had been developed. It also so happened that one of the gurus of the other school of burning was at the university that I was teaching at.

And so one day, purely out of mischief, I pulled the guru into my office and said, "Why don't you run this expert system? Tell me what you think of it." And the person ran the expert system and then got very thoughtful and said, "You know, I think I'm beginning to understand where I disagree with them." And it turned out there was a kind of subtle interpretation of priorities whereby, under some circumstances, he would have come up with a different answer--or a different decision, given the same set of answers. And so that very, very simple knowledge base, because it explicitly spelt out how that person was thinking, enabled somebody else to follow it and figure out where they disagreed.

So one good reason why simple expert systems are not trivial is because they are explicit. They spell out how somebody is thinking about a problem.

There's another more fundamental reason as to why small expert systems can be useful. If you look at recent models of what goes on in the human brain, and you try and model the human brain in the same way as you think about a computer, it turns out that within our brains we have fantastic computing power, but we don't have direct access to it. It's kind of going on in the back of our brains. The part that we are aware of occurs in the forefront of our brains, and it turns out that that particular part of the brain is very limited in its capacity. A person can only keep four or five or maybe six ideas in consciousness at the same time. We're kind of like jugglers that can't handle more than four or five balls.

If you ask why phone numbers have hyphens between three or four numbers, that is because the brain then can take each little group of numbers, and it can remember three things. Same with your credit card. Give somebody a credit card where you don't put hyphens in between, and you ask them to remember it, and they will forget it very, very quickly because their brains are overloaded.

Now, if you've been to a meeting, or you've interacted with somebody who strikes you as super intelligent, the kind of person that can keep their fingers on a whole number of components

in a discussion, what you are probably dealing with is somebody who is capable of handling six or seven ideas instead of four or five. And that makes a difference between somebody who's really impressive at a meeting and somebody who is just run-of-the-mill at a meeting.

So what happens if you have a problem that has six or seven components to it that you need to juggle, but you have a brain that can only juggle four or five ideas? Well, it turns out that your brain doesn't give you an error message. The back part of the brain doesn't say, "Overload, overload. We can't handle this." What the back part of the brain does is it sends the difficulty to the emotional side of the brain. And the emotional side of the brain decides which four or five things you will concentrate on and drops the rest, but it doesn't tell you that it's done that. So you've got a problem with seven components, and you think you are thinking it through, but, in fact, you are only keeping track and thinking consciously of four or five of those components.

And that is why you can get a group of people sitting around a table with the best will in the world trying to solve a problem, and it seems as though each one of them is dealing with a different problem. And the answer is, "Yes, each one of them is

dealing with a different problem because they're dealing with different subsets."

Now, take that analogy and think about that same group of people sitting down and trying very explicitly to develop a knowledge base. Now when you sit down and try to develop the knowledge base, you are breaking those seven components down into little pieces, and so your brain can handle them, two or three at a time. And if different people are thinking about different things, they can share that information.

So you can now produce a knowledge base that handles seven or eight or nine components and enables you to think through those in a way that no single human being can do. And that is why I believe small expert systems can be useful.

Okay. "But," you might say, "if you're going to keep it that simple, why bother to go to the computer? Why not just get the group of people to sit down and develop a decision tree? A decision tree would be 'Ask this question first, and then if the answer is this, go in this direction. But if the answer is that, go in that direction.' And you can put it on one or two pieces of paper, and anybody could follow the logic or come to a conclusion."

And I think there are two answers to that. The first is that the decision tree doesn't contain all that explanation, and the explanation is really a valuable part of the interaction. It's hard enough sometimes to develop a decision tree that doesn't look as though it needs to be pruned. It will be even harder if you tried to tag on explanations. So the explanation feature of the expert system shell becomes something that's really important.

The second reason for using a computer rather than a decision tree is, I think, a subtle point in the psychology of a human being. If somebody tells you, "Here's a decision tree. This is how you solve the problem," your reaction is kind of, "You're telling me what to do. You're cutting out my ability to think. You are undervaluing me. You're giving me a cut-and-dried answer to a problem." If you take that same person with exactly the same information and have them interact with an expert system, the psychology of it is, "Gee, the computer is responding to my problem. It's helping me think through this example." And that sounds kind of kooky, but I think it's important.

Okay. So where, in conservation and resource management, do I think one can use small expert systems? The first point I would make is that having a set of small expert systems will help an

organization be more professional.

Remember when we were talking about structured decision making. I made the point that for an organization to be professional, if they're making decisions, they need to be able to make them in an organized fashion so that they can talk about what worked, and they can improve, and they can deliver best practices to people.

The same goes for expert systems. If you are dealing with certain types of situations within an organization, coming up with an expert system on how you do that, developing a knowledge base and using expert systems, enables you to spell out that this is the way we, as an organization, deal with this kind of problem. Once you do that, you first of all ensure that everybody in the organization is giving similar answers, but secondly, you're creating a way to improve what you're doing.

Remember when we looked at the little example about burning, I showed you how coming out of the discussion you've had with the computer, you can get a transcript of how you answered the questions and what decisions were made.

So imagine you are a user making a real decision, and you get a

transcript, and you say, "You know, that's what the computer told me to do, but I disagree, and I'm going to write on that transcript why I disagree." And then depending on the latitude you have, you might or might not go along with a decision. Can you see if you had a number of users using the same knowledge base and thinking about what it told them and how they might change it, and then if you had an update meeting once a year and reviewed the knowledge base, you are going to have an evolving expertise within the organization? And that is valuable, okay?

At the same time, you're going to be able to deliver this expertise to new people who are coming in. So if you've got somebody who has just taken a job that has been vacated by somebody who has been doing this over and over again, they can use the expert system to get up to speed. So it becomes a learning tool and provides for continuity within an organization.

It also becomes a tool for communication between scientists and practitioners. It's a way for delivering science to somebody who has to use it. Scientists usually say, "This is the science," and they never think about the nuances and the difficulties and the real-world situations where that science might be applied. Getting scientists to work with practitioners to develop expert systems is going to lead to a better transfer between science

and practice, and that's important.

Expert systems can also be used to communicate to stakeholders and to the public. Going through an expert system of how a decision is made tells somebody--and that decision might be controversial--will tell people why the decision was reached, particularly if there's a good explanation feature. So one could have little expert systems available at websites so that people who are interested in certain things can follow through and see what the thinking is behind the decisions that are made on a routine basis.

And, finally, developing the knowledge base. Having a group of stakeholders with perhaps different objectives and different views sitting around and coming up with a knowledge base that is a compromise on how they're actually going to deal with a difficult situation can be a form of mediated modeling, just the tasks of having people who might argue with each other, who might have large divisions in the way in which they think about a problem. Having them sit down and do something cooperative so that the computer becomes the enemy, not the other person that you're working with and so that you can start breaking down your differences whereby you might say, "I would never do that," and somebody else on the other side would say, "I always do that."

Well, you can then go to the person who would never do it and say, "Maybe under these circumstances you would do it." And so you chip away at their position. And then you go to the other side and say, "Maybe under those circumstances you wouldn't do it." And so you can chip away at the differences and try and come up with useful compromises. Maybe you would never use the expert system, but the process of developing it could be really useful.

So there are some of the ideas of things that one can do within a professional organization if one develops a culture of building, thinking about, improving, using and communicating small expert systems.

Well, to wrap up, let me give you some examples. The first example I'm going to give you is an example of an expert system that I have never developed but wanted to. Remember I talked about how I got interested in expert systems. Well, one of the papers I read really sparked my attention. It was a paper on an expert system that was meant to give career guidance to a kid in high school. Now, I've no idea how good or bad career guidance is in the high school you went to, but I will tell you in a moment how bad it was in the high school that I went to when I was a kid. So the idea of having a computer help you might be a big im-

provement on having a professional career guidance person help you.

But what was interesting about this particular expert system was that it was in three parts. In part one, it asked you something about your value system. In part two, it questioned you about your abilities. And then in part three, it asked you what career you wanted to follow. And what the expert system did was to take the answers from these three points, from these three different areas and put them together and either say, "Yes, that makes sense," or else point out inconsistencies in your thinking.

So, for example, if you said in the first part, in terms of your value system, that you were not interested in money, that you did not like working with people, that you would like to live in a remote part of the world where you can be close to nature. In the second part, you might have said that you are dyslexic, that you don't have very good motor skills, and that you are a D student. And then in the third part, you say you want to be a brain surgeon. And the computer would then point out the inconsistencies in your thinking.

This excited me because what I wanted to do was try and develop a modeling expert system, one that said, "What is the problem

that you're dealing with? Give me the characteristics of the problem that you're dealing with. Tell me something about your objectives in dealing with it. What are you trying to achieve with your model? And then tell me something about a standard modeling technique that you think you might use." And then the expert system would try and put those three parts together and come back at you and say, "This makes sense," or maybe, "It doesn't." In fact, I've never got around to that because I've never been able to figure out how I would think about it.

By the way, going back to my experiences with vocational guidance, when I was about 15, we all had to do an IQ test, and then we went to see the teacher who was going to give us career guidance. And, as a timid little guy, I walked into the teacher's office, and the teacher looked up at me and said, "Starfield, you should go and work in the post office." And this kind of surprised me. I never thought of a career in the post office. And I had the temerity to ask, "Yes, sir. Why, sir?" And the teacher said, "You're smart. You'll get to the top very quickly."

Now, notice the importance of the explanation feature in that discussion. If I just walked into the teacher and been told, as a result of my IQ test, that I should go and work in the post

office, I might've lost a lot of confidence in terms of my abilities. His explanation told me how the teacher was thinking.

Okay. So much for an example that wasn't developed. Two more quick stories. I used to teach a class on modeling, which included expert systems and structured decision making as part of the Conservation Biology Program at the University of Minnesota. And a colleague at the university once asked me whether she could sit in on the class, and she was going to do all the assignments. And she came in, she did all the assignments, and one of the assignments was that, having seen how an expert system worked, students had to choose a topic and, working in small groups, develop a small knowledge base over a two-week period.

And she came to me and said, "I'd really like to develop a small knowledge base, and I've got an example of something I'd like to try, and please tell me whether it will be useful or not." She said, "I have just published a paper. It's a paper on wetland restoration, and it basically guides people through the things you need to think about and how to do wetland restoration." And she said, "Would it be a waste of time for me to take the ideas in that paper and put it into an expert system?" And I said, "No, go for it."

And she came back in two weeks' time, and she said, "That was a really interesting experience. It was interesting because by the end of it, I realized that I thought I had written a really good paper, and I had written a bad paper, what wasn't as good as I thought. Why? Because it is one thing to just talk about or write about what to do. It is another thing to go through the discipline of thinking exactly how you would do it."

And so she found that going through the discipline of the expert system shell did two things for her. One is it showed gaps in her paper, and the other thing is, because she would now have to justify the logic of what she was doing, it made her re-think some of the advice that she was putting in the paper. Hence, my point that a knowledge base and expert system is a much better way of delivering science to a practitioner than writing an article. It forces you to think hard about the user, and it forces you to explain to the user why you're giving them a particular type of advice.

That explanation feature relates to the second example that I want to talk about. Way back when I was first developing small expert systems, and I'd been talking about them to conservation organizations in South Africa, there was the equivalent of a state department of natural resources that had been managing a

small game park. And they had a set of management rules that, it had turned out, were not working well in practice.

And what they decided to do was call a meeting of scientists, managers, rangers, and pull them together to try and come up with a new way of thinking through their management strategy. And having heard about expert systems, they said to me, "Tony, we're going to have a five-day meeting. If we try and thrash out the problem in the first four days, do you think you could come along on the fifth day and help us put it into an expert system shell?" And I said, "Great, I'd love to do that."

And so I got there on the end of the fourth day, and they were really excited. They said, "We've had the most wonderful meeting. We really think we've figured out what we want to do, and it's going to be super. It's going to go like clockwork putting it into the expert system shell tomorrow."

And so we sat down, and we had the expert system shell, and we put in the decisions, and it went smoothly. We put in the questions. It went smoothly. We put in the rules, and everybody had agreed on them.

The one thing they hadn't done during the first four days was

figure out the explanations. And so the last thing we had to do was put in the explanations on the questions and the rules. And as we started doing that, the whole agreement unraveled in front of their eyes. They thought they had reached consensus. They thought they had come up with a good management plan, but when they started thinking about why they were doing it, they started disagreeing, and they started doubting what they had come up with. And so, as a consequence of the exercise of trying to put it into an expert system shell, what happened was that they threw out four days of work and said, "We're going to have to think of a way of re-doing this exercise."

And you might say, "Well, that's not a great example of the use of an expert system shell," but I actually think it's a very pertinent example, because so often in modeling, the value in the modeling exercise is doing it and how it helps you to pull your thinking together, even if you never actually use the model. It's the process that can be very, very important. So there's some real-life examples.

I'm going to end this little segment with an entertaining, apocryphal story about expert systems. Remember when we began talking about expert systems, I mentioned that the whole subject grew out of this idea in artificial intelligence of trying to

produce a computer that could fool you into thinking you were interacting with a human being. And some of the people working on this--and I think this was in England--tried to develop a computer program that would keep a conversation going with a user. So the idea was to keep a conversation moving. And the computer program had to be able to understand English grammar to a certain extent.

And they decided that the best area in which to test this was to try and develop what they thought of as a computer psychiatrist. And they had the view that might upset a psychiatrist, that the purpose of a psychiatrist was to keep you talking. And the way the system worked was that if you typed into the system, from some remote computer or terminal, something, it would take what you had written, look for some phrases, try and twist it and come back at you with a question to keep the conversation going.

And the story is that one Saturday morning, a professor who was sort of peripheral to this project was sitting at a computer in some remote area and wanted to try and have a sort of computer chat with one of the professors working on the project. And so he accessed, or dialed into the project, and thought he was dealing, interacting, chatting to maybe a grad student or somebody who was working there. In fact, he was interacting with the

computer program.

So he dialed in, and the computer said, "Hello? Can I help you?" It didn't say; typed it to him. And he typed back, "Yes. Is Professor so-and-so there?" And the computer types back, "Are you looking for Professor so-and-so?" And he says, "Yes." And the computer says, "Why are you looking for Professor so-and-so?" And he says, "I want to talk to him." And the computer says, "Why do you want to talk to him," keeping the conversation going. And the person who's trying to get through is getting uptight now, and said, "It's really none of your business. Is Professor so-and-so there?" And the computer says, "You seem very anxious to reach Professor so-and-so. Why?" And the conversation starts going downhill from there, and eventually, the person types in something fairly rude or aggressive.

Now, the expert system on the computer, when it reached an impasse, when it didn't know how to keep the conversation going, was programmed to start up a new conversation. And it had a list of maybe 40 or 50 new ways to start a conversation, and it would pick one at random and throw it in to try and re-start the conversation. So, at this point, the computer types back to the real person, "Tell me about your mother."

[END OF VIDEO]

Document: blm0211.doc

Transcribed by: <http://www.hiredhand.com>