DEVELOPING AN EXPERT SYSTEM FOR HAZARDOUS WASTE REMEDIATION

John D. Keenan

Department of Systems
School of Engineering and Applied Science
University of Pennsylvania
Philadelphia, Pennsylvania 19104

Timothy Finin

Unisys Corporation
Paoli Research Center
Paoli, Pennsylvania

Expert system methodology has been receiving considerable attention as a useful tool in the solution of those complex problems requiring the application of expertise. Familiar examples include medical diagnostics. Another area in which the role of expertise is crucial is in the remediation of sites which have become contaminated as a result of the disposal of hazardous wastes. In the following paragraphs, the development of such an expert system is detailed.

EXPERT SYSTEMS METHODOLOGY

An expert system may be defined very generally as a computer program which solves problems in some particular domain. These problems are such that they normally require some specialized knowledge (i.e., expertise) on the part of humans who successfully solve them. The operant terms in this definition are solves problems, particular domain, and specialized knowledge. This definition would exclude programs which are not characteristically problem solvers (e.g., a text editor), attempt to be very general (e.g., a theorem prover), or solve problems which are considered straightforward implementations of well-defined theories (e.g., computing FFT's) or just mundane (e.g., simple record keeping). Typical domains for expert systems are medical diagnosis, signal interpretation, fault diagnosis, and computer configuration.

Unfortunately, the definition is still much too broad. This definition encompasses most of what used to be called "application programs." In order to capture what is new and valuable in recent work on expert systems, it is useful to enumerate some of the major features which are found in "good" examples of expert systems. In addition to the partial definition offered above, a "good" example of an expert system will be likely to have the following characteristics: separate knowledge base, multiple use of knowledge, special knowledge representation languages, in the knowledge lies the power, explanations, and shell architecture.

Separate Knowledge Base. The special domain knowledge that the expert system uses is explicitly represented in a module (the knowledge base or KB) that is separate from the components which use it.

Multiple Use of Knowledge. Since the KB is explicitly represented as a separate module, it can be used in several different ways. For example, the KB may be used to make decisions, to construct explanations, to construct tutorials, etc. This requires that the knowledge be represented in a more general way that does not favor one use at the expense of another. An additional goal is to express the knowledge in a general way that allows it to be reasoned about as well as reasoned with.

Special Knowledge Representation Languages. The knowledge in an expert system is usually encoded in a special representation language. Most current artificial intelligence (AI) systems use representation languages which rely on one or more of three general techniques: rules, frames, and logical relations.

In the Knowledge Lies the Power. This often quoted slogan underscores the fact that the real intelligence in the expert system lies in the domain knowledge represented in the KB, not in any of the more general components. As a corollary, any strategies for searching for a problem solution should be expressed in the general KB rather than in the code which uses the KB.

Explanations. When appropriate, an expert system should be able to provide justification for its conclusions. This can be done by offering explanations which describe the systems reasoning which led to the conclusions. It may also be appropriate for some systems to be able to explain why other (plausible) conclusions were not reached.

Shell Architecture. A typical architecture for an expert system is to have a shell for the expert system which includes a number of general, problem-independent components which is combined with one or more specific knowledge bases which encode the problem specific information. The general components can include an interaction manager, a general "inference engine," knowledge acquisition system, KB debugger, KB editor, explanation generator, etc.

None of these features is required of a system in order to characterize it as an expert system. Many AI problem solving systems, for example, do not use a rule-based approach. Explanations may not be relevant in some problems. However, all of these features have been found to be quite useful in building powerful problem solving systems.

A very important technique for building expert systems is to organize the knowledge in the system in the form of rules. In simple terms, a rule based system has two major components: a knowledge base and an inference engine. The knowledge base contains a set of facts and a set of rules which, together, represent the system's general domain knowledge and specific knowledge about the current problem. A fact is generally considered to be an atomic proposition which is true in the world and a rule to be a conditional with an antecedent (if part) and a consequent (then part). Thus a rule is a piece of knowledge of the form:

IF antecedent THEN consequent

Depending on the nature of the antecedent and consequent parts, a rule can be one of several different varieties. If both parts are viewed as logical propositions then the rule can be viewed as a rule of inference. If the antecedent is interpreted as a partial description of some state of the factual knowledge base and the consequent can be any arbitrary executable expression then the rule is similar to the notion of a production rule as used by cognitive scientists. Many rule-based systems have a need to deal with uncertain data and rules which do not always hold. A common variation, then, is to associate with each fact or rule a degree of uncertainty which represents the system's confidence that it is true.

The inference engine is a kind of rule interpreter which can take the set of rules and facts and compute additional facts (or rules) which hold. An inference engine can use rules in one of two important ways. Forward chaining is a process of reasoning from the initial set of facts to derive additional facts which must hold. The inference engine identifies a rule whose antecedent part is satisfied by the facts currently in the KB. Backward chaining is a process of reasoning from a given goal backward to a set of facts which, if in the current KB, would support the goal. Given a goal to satisfy, the inference engine looks for an appropriate fact or rule in the KB. If a matching fact is found then the goal has been satisfied. If a rule is found whose consequent matches the goal, then the engine attempts to recursively satisfy the antecedent of the rules by setting it up as a sub-goal.

POTENTIAL ROLE IN HAZARDOUS WASTE HANDLING

Expert systems are applicable to a range of problems involving the handling and disposal of hazardous wastes at any one of several levels. These include licensing of disposal operations, siting new disposal operations (activities in conjunction with RCRA, Resource Conservation and Recovery Act), and remediation of existing sites (activities in conjunction with CERCLA, Comprehensive Environmental Response, Compensation, and Liability Act). The expert system described in this paper deals specifically with CERCLA activities.

The Comprehensive Environmental Response, Compensation, and Liability Act of 1980 provides the U.S. Environmental Protection Agency with the authority and responsibility to develop response actions at specific uncontrolled waste sites. The plan for enacting CERCLA provisions has been published as the National Contingency Plan. The National Contingency Plan requires the preparation of a detailed Remedial Investigation (RI) and Feasibility Study (FS) in order to ensure the selection of cost-effective remedial alternatives. Detailed procedures for the RI/FS process have been developed (U.S. EPA, 1985).

The purpose of this paper is to present a progress report detailing the development of an expert system dedicated to remediation activities at hazardous waste disposal sites. This expert system, termed Toxic Waste Advisor, deals specifically with CERCLA activities. The function of the system is to define the technical alternatives appropriate for a site specific remediation program.

A long term goal is the inclusion of health/safety and public health effects into the system. This effort will be delayed until a later point in the process of developing the expert system. The starting point has been taken to be subsurface contamination by organic liquids which are either volatile and/or nonvolatile hydrocarbons, e.g., solvents and fuels. The knowledge base will be expanded at a later date to include other classes of hazardous materials.

PROTOTYPE EXPERT SYSTEM

Toxic Waste Advisor has been developed to define the feasible technical alternatives appropriate for the cleanup of a site at which the soil and/or the groundwater has been contaminated. As noted above, for the purpose of the initially developing expert system, the contamination is assumed to be either volatile organic solvents and/or hydrocarbons. The contaminated aquifer is assumed to be a single, homogeneous, isotropic formation.

The prototype expert system was constructed by developing a series of engineering decision rules. These were expressed as IF . . . THEN rules and implemented as an expert system using the Texas Instruments Personal Consultant expert system shell.

Subdivision of the Problem. The problem of remediating a contaminated site involves the consideration of four sub-problems. These are contamination of the unsaturated zone, contamination of the saturated zone, treatment of contaminated groundwater, and disposal of treated groundwater. The engineering details associated with each of these aspects of the overall problem are discussed below.

Unsaturated Zone Contamination. The first consideration is whether the unsaturated zone is contaminated. In the event that

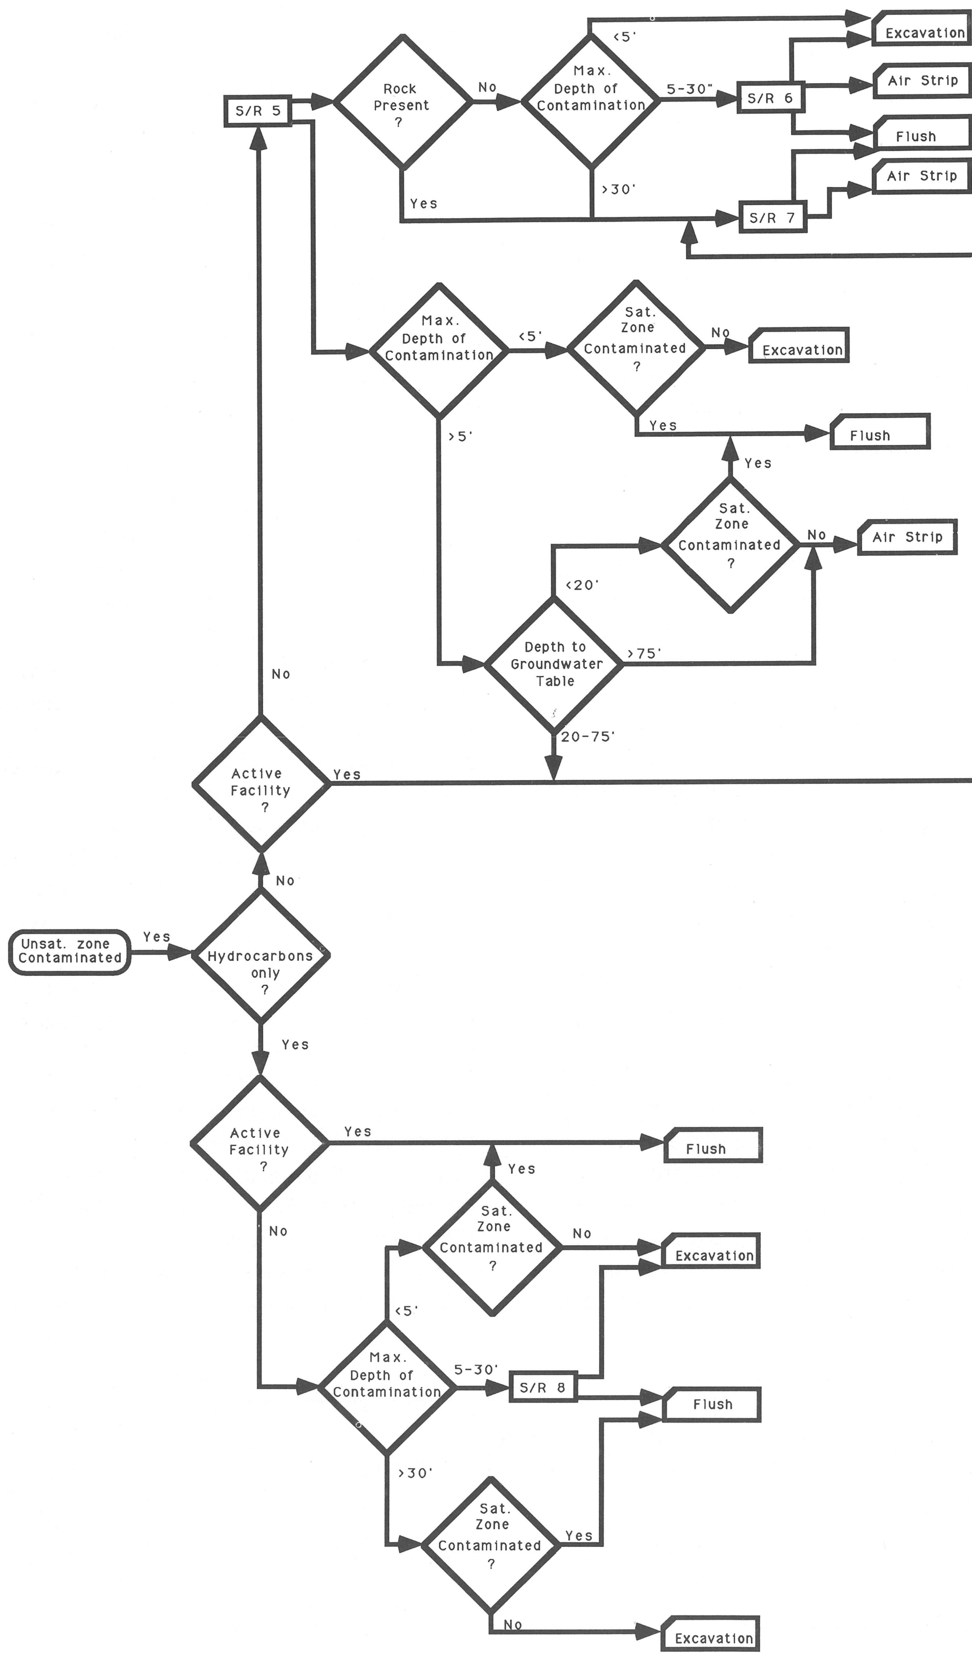

FIGURE 1
Schematic of Engineering Decision Rules for the Case of a Contaminated Unsaturated Zone

it is, the appropriate technical alternatives are excavation of the contaminated soil material, flushing via circulating water, and air stripping via the movement of air through the contaminated material. The decision making flow chart for the case in which the unsaturated zone is uncontaminated is shown in Figure 1. The selection of a given alternative is determined by a number of variables. The nature of the hazardous material itself is a factor in that its volatility and its solubility will determine the efficacy of air stripping and flushing, respectively. The status of the site, i.e., whether it is active or not, will influence the excavation option. The permeability of the subsurface medium may mitigate against the selection of flushing as a feasible option. The status of the saturated zone is another key factor. If the saturated zone is clean, then flushing is not an alternative. And, finally, an economic evaluation is may be needed to select from otherwise equivalent alternatives. These are shown in Figure 1 as sub-routines (S/R). For example, sub-routine 8 (S/R 8) will be used to evaluate the costs of excavation versus flushing for cases in which the unsaturated zone at an inactive site is contaminated with hydrocarbons at moderate depths.

Saturated Zone Contamination. In the instance where the saturated zone is contaminated, the engineering options include extraction wells, drain tiles, and barrier walls. The flow chart for this case is shown in Figure 2. The principle variables include the depth to the water table, the thickness of the saturated zone, the thickness of the contaminated region, the nature of the subsurface medium, and economics. Again, the role of the sub-routines is to perform an economic evaluation of competing alternatives. In some cases, the sub-routine is linked to a groundwater flow model. An example is S/R 2, in which the effect of barrier walls on the drain tile option is evaluated.

When the saturated zone is contaminated, the first decision variable is the depth to the water table. When both the depth to the water table (<5 ft) and the contaminated thickness (<10 ft) are small, then drain tiles or extraction wells are indicated, depending on the presence of quick conditions. As the depths increase, the situation becomes more complex. As noted in Figure 2, when the water table is found at an intermediate depth (5–40 ft), but the saturated thickness is small (<5 ft), then drain tiles are selected. Otherwise, the decisions depend upon the presence of rock, and the depth and the permeability of the formation. The purpose of subroutines 2, 3, and 4 is to determine if barrier walls are needed in conjunction with drain tiles and extraction wells.

Groundwater Treatment. The options which are available for groundwater treatment include phase separation, air stripping, gaseous phase treatment and disposal, filtration, and carbon treatment and disposal. The principal variables are the chemical nature of the contaminated groundwater and process economics. The flow chart is shown in Figure 3. The presence of a second phase in the contaminated groundwater will require phase separation to remove the floating layer. The second phase will then require disposal.

The aqueous phase is subjected to further treatment. If the aqueous phase is contaminated by hydrocarbons only then carbon treatment and disposal of the carbon is indicated. Air stripping is needed to remove volatiles from the aqueous phase. In turn, the gaseous phase from the air stripping operation may require carbon treatment depending upon the mass flow rate of the volatiles. Following air stripping, the aqueous phase may require carbon treatment if initial concentrations were sufficiently high. Furthermore, filtration for removal of iron and manganese is indicated in certain instances.

Disposal of Treated Groundwater. Subsequent to treatment, recovered groundwater may be disposed via recharge, remedial recharge, sewer, or surface water. The main variables are the overall remediation plan as well as process economics. Figure 4 includes a flow chart for these options.

Rules. After the flow charts representing the engineering decision-making process (Figures 1 through 3) were completed, the information was reexpressed as a series of IF . . . THEN rules. As an example, one of the rules is shown in Figure 4. This rule refers to contamination in the saturated zone, and it states that if the depth to the water table is less than or equal to 5 ft, and if the thickness of the contaminated soil is not more than 10 ft, then there is suggestive evidence (75%) that the treatment method is drain tiles and weakly suggestive evidence (25%) for extraction wells.

Another example of a rule is presented in Figure 5. This rule states that if the unsaturated zone was treated by excavation

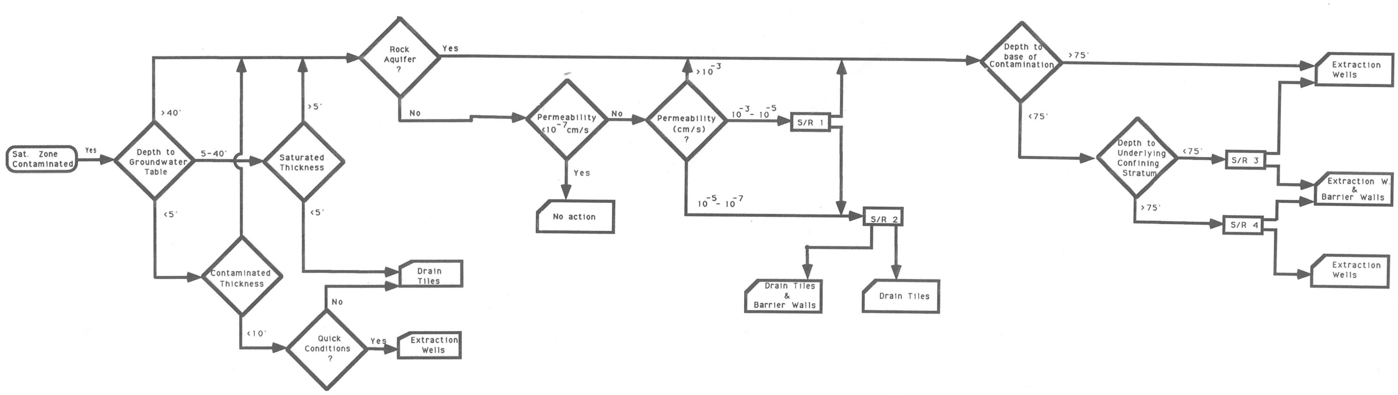

FIGURE 2
Schematic of Engineering Decision Rules for the Case of a Contaminated Saturated Zone

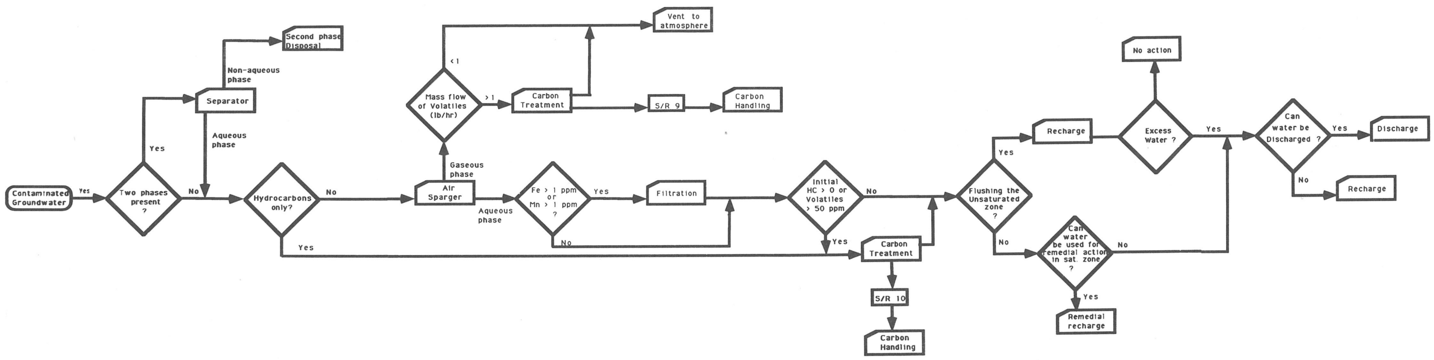

FIGURE 3
Schematic of Engineering Decision Rules for the Treatment of Contaminated Groundwater

and if the contaminant was volatile organics, then the treatment for the solids is on-site thermal treatment.

Implementation. The rules were implemented as an expert system on the TI Personal Consultant expert system shell.

```
RULE020 [SOLID-POST-TREATMENTRULES]
-------

If 1) the treatment method(s) for collecting the toxic waste
      in the unsaturated zone of SOLID-POST-TREATMENT is
      EXCAVATION, and

   2) the major type of contaminant in the UNSATURATED ZONE
      is Volatile,

Then it is definite (100%) that post-treatment for solid
     waste collected from SOLID-POST-TREATMENT is
     ONSITE-THERMAL.
```

FIGURE 4
A Typical Rule (Referring to Contamination in the Saturated Zone)

```
RULE001 [SAT-ZONERULES]
-------

If 1) the DEPTH in feet to the watertable at SAT-ZONE is
      less than or equal to 5, and

   2) the thickness of contaminated volume of soil in the
      saturated zone is less than or equal to 10,

Then 1) there is suggestive evidence (75%) that the
        treatment method(s) for the saturated zone is TILES,
        and

     2) there is weakly suggestive evidence (25%) that the
        treatment method(s) for the saturated zone is WELLS.
```

FIGURE 5
A Typical Rule (Concerning Treatment of Excavated Solids from Remediation Site)

The Toxic Waste Advisor has been designed to provide clear menus, help on demand for the user, and an explanation of the chain of reasoning. The explanation facility may be used both to indicate why the system is asking a particular question and how the final recommendation is reached. The Toxic Waste Advisor represents a working expert system devoted to CERCLA activities.

During a session with the Toxic Waste Advisor, the user is prompted for information defining the nature and scope of the problem to be considered. Space does not permit the reproduction of each screen which would be presented to the user during such a session. However, a few of the screens will be shown and discussed. A typical request from the system is shown in Figure 6. Here, the depth to the water table is requested. In Figure 7, the request is for the status of the facility. Note that one of the function keys may be pressed to provide help with the selection. At any time, the user may ask why a particular question is needed. As shown in Figure 8, the system responds with an answer which explains the need for asking the question. The system provides its recommendation concerning the treatment options as shown in Figure 9. If the user

```
session 2  page 4
================================================================================
  Knowledge Base :: Toxic Waste Advisor(TWA)
--------------------------------------------------------------------------------
What is the depth (in feet) to the water table in SITE-1?
--------------------------------------------------------------------------------
Enter a positive number.
--------------------------------------------------------------------------------

        30

------------ ------------  ----------- -----------  ----------- ----------
 Up   Down    CF  Unknwn    Done  ---   Why   How     Help Undo      Stop
------------ ------------  ----------- -----------  ----------- ----------
 F1     F2    F3      F4    F5     F6   F7     F8     F9   F10        ESC
================================================================================
```

FIGURE 6
User Interface. The Toxic Waste Advisor Has Requested the Depth to the Water Table

```
session 2  page 18
==========================================================================
   Knowledge Base :: Toxic Waste Advisor(TWA)
--------------------------------------------------------------------------
 Is there an on-going facility at UNSAT-ZONE-1?
--------------------------------------------------------------------------
 Select one of the following:  (Press 'F3' for selection help.)
--------------------------------------------------------------------------
        YES
       *NO




 ------------ ------------  ----------- -----------  ----------- ----------
 Up    Down   CF  Unknwn    Done  ---   Why   How     Help Undo       Stop
 ------------ ------------  ----------- -----------  ----------- ----------
 F1      F2   F3      F4    F5      F6  F7      F8    F9    F10        ESC
==========================================================================
```

FIGURE 7
User Interface. The Expert System Has Requested the Status of the Facility.

```
session 2  page 21 (WHY)
==========================================================================
   Knowledge Base :: Toxic Waste Advisor(TWA)
--------------------------------------------------------------------------
 Is there a hot spot in UNSAT-ZONE-1?
--------------------------------------------------------------------------
 Why this question is needed:
--------------------------------------------------------------------------
Whether hot spot in UNSAT-ZONE-1 is needed to determine
 method of treatment for UNSAT-ZONE-1

 RULE 008
 If 1) hot spot in UNSAT-ZONE-1, and
    2) maximum depth (in feet) of contaminant in UNSAT-ZONE-1
       is less than or equal to 5,
 Then There is evidence that method of treatment for
      UNSAT-ZONE-1 is EXCAVATION (750)  FLUSHING (250)
      AIR-STRIPPING (250)

      ... end   --  press RETURN

 ------------ ------------  ----------- -----------  ----------- ----------
 Up    Down   CF  Unknwn    Done  ---   Why   How     Help Undo       Stop
 ------------ ------------  ----------- -----------  ----------- ----------
 F1      F2   F3      F4    F5      F6  F7      F8    F9    F10        ESC
==========================================================================
```

FIGURE 8
Expert Explanation. The System Has Responded with an Explanation of Why a Particular Question Has Been Asked

asks how was this conclusion reached, then the system responds as shown in Figure 10.

SUMMARY

A prototype expert system, Toxic Waste Advisor, has been developed. The objective of the expert system is to assist the user in the selection of one or more techniques for remediation of hazardous waste sites involving soil and/or groundwater contaminated with solvents and/or hydrocarbons. The decisions

```
session 2  page 23
==========================================================================
   Knowledge Base :: Toxic Waste Advisor(TWA)
--------------------------------------------------------------------------
 Method of treatment for UNSAT-ZONE-1 is as follows:
 AIR-STRIPPING (75%) FLUSHING (50%) EXCAVATION(25%)

        ... end   --  press RETURN




 ------------ ------------  ----------- -----------  ----------- ----------
 Up    Down   CF  Unknwn    Done  ---   Why   How     Help Undo       Stop
 ------------ ------------  ----------- -----------  ----------- ----------
 F1      F2   F3      F4    F5      F6  F7      F8    F9    F10        ESC
==========================================================================
```

FIGURE 9
Recommendations. Based on User Input, the Toxic Waste Advisor Recommends Treatment Options.

```
session 2  page 26
==========================================================================
   Knowledge Base :: Toxic Waste Advisor(TWA)
--------------------------------------------------------------------------


 UNSAT-TREAT was determined to be: AIR-STRIPPING with certainty
 75% by using RULE 013

 If 1) hot spot in UNSAT-ZONE-1 is not true, and
    2) maximum depth (in feet) of contaminant in UNSAT-ZONE-1
       is greater than 5, and
    3) depth (in feet) to water table in UNSAT-ZONE-1 is
       greater than 20, and
    4) depth (in feet) to water table in UNSAT-ZONE-1 is less
       than or equal to 75,
 Then There is evidence that the method of treatment for
        UNSAT-ZONE-1 is EXCAVATION (250)  FLUSHING (500)
        AIR-STRIPPING (750)

 UNSAT-TREAT was determined to be: FLUSHING with certainty 50%

        more ...   --  press RETURN
 ------------ ------------  ----------- -----------  ----------- ----------
 Up    Down   CF  Unknwn    Done  ---   Why   How     Help Undo       Stop
 ------------ ------------  ----------- -----------  ----------- ----------
 F1      F2   F3      F4    F5      F6  F7      F8    F9    F10        ESC
==========================================================================
```

FIGURE 10
Expert Explanation. When Prompted, the Toxic Waste Advisor Explains Why a Particular Recommendation Was Made

reached by the expert system are affected by factors such as site characteristics (permeability, depth to groundwater, depth to bedrock, etc.), the type of hazardous material involved (volatility, extent of contamination, concentration), and an economic analysis of the alternatives. The Toxic Waste Advisor has been implemented on the Texas Instruments Personal Consultant expert system shell.

LITERATURE CITED

Handbook. Remedial Action at Waste Disposal Sites (Revised). U.S. EPA. Publication EPA/625/6-85/006 (October, 1985).