Classifying Industrial Sludge Using a Knowledge-Based Expert System

Allen C. Chao[1], William J. Rasdorf[2] & William E. Paige[3]

INTRODUCTION

Sanitary landfills have been used as the most common method for disposal of most industrial sludge materials (1). However, since hazardous constituents may leach out of the landfill, migrate into ground water, and cause potentially serious problems, regular sanitary landfill facilities designed for municipal solid wastes are not allowed to accept hazardous industrial sludge. Industries in the State of North Carolina are required by law to submit to the state information concerning the sludge they produce. Environmental experts in the State Department of Human Resources and Community Development then determine the sludge classification based on the information submitted. The criteria for identifying the characteristics of hazardous wastes listed in the N. C. Hazardous Waste Management Rules & Solid Waste Management Law (6) are used as bases by the North Carolina regulatory agencies to determine whether an industrial sludge is non-hazardous and thus can be disposed of in a municipal sanitary landfill.

To implement the regulations, human experts need to possess a detailed understanding of the procedures for implementing these criteria and regulations. An inexperienced engineer may need an extended period of training before he can be entrusted to successfully carry out sludge classification activities. Even for experienced experts, inconsistent conclusions may occasionally be drawn by different individuals due to varying interpretations of the criteria and regulations (2,7). Additionally, a serious drawback of the current method of sludge classification is that the criteria and regulations frequently change. Thus, a person who has not followed those changes for a period of time may become unfamiliar with the current criteria and regulations.

A prototype knowledge-based expert system has been developed to overcome these problems while assisting environmental experts in classifying industrial sludge. This prototype program is written in M.1, a microcomputer-based knowledge engineering tool for building small knowledge systems. The main objective of this program is to test the feasibility of utilizing knowledge-based expert systems to assist an inexperienced environmental expert in classifying industrial sludge. Successful test runs of the prototype expert system have led to a plan for developing a more versatile system using a more powerful knowledge engineering tool running on a minicomputer.

CONSIDERATIONS AND CRITERIA IN CLASSIFYING INDUSTRIAL SLUDGE

From among the numerous items listed in the hazardous waste management criteria and solid waste management law for identifying hazardous sludge, there are two criteria considered to be the most pertinent for classifying industrial sludge. The first criterion is that if a sludge exhibits certain hazardous characteristics, it is classified as hazardous. Four categories of characteristics and their test methods have been specified: ignitability, corrosivity, reactivity and EP toxicity. In each of these categories, the waste characteristics and their test methods are described in detail. For example, a solid waste or sludge is considered to exhibit the characteristic of ignitability and thus is classified as hazardous if a representative sample of the waste or sludge exhibit one of the following properties:

1. It is a liquid, other than an aqueous solution containing less than 24 percent alcohol by volume, and has a flash point less than 60°C (140°F), as determined by a Pensky-Martens Closed Cup Tester, using the test method specified in ASTM Standard D-93-79 or D-93-80 (incorporated by reference, see Section 260.11), or a Setaflash Closed Cup Tester, using the test method specified in ASTM Standard D-3278-78 (incorporated by reference, see Section 260.11), or as determined by an equivalent test method approved by the Administrator under procedures set forth in Sections 260.20 and 260.21.

2. It is not a liquid and is capable, under standard temperature and pressure, of causing fire through friction, absorption of moisture, or spontaneous chemical changes and, when ignited, burns so vigorously and persistently that it creates a hazard.

3. It is an ignitable compressed gas as defined in 49 CFR 173.300 and as determined by the test methods described in that regulation or equivalent test methods approved by the Administrator under Section 260.20 and 260.21.

4. It is an oxidizer as defined in 49 CFR 173.151.

An industrial sludge which shows a positive result to one of the above tests should be considered hazardous.

The second criterion concerns the source or the constituents of a waste or sludge. If the source has an assigned EPA hazardous waste number or if the constituents are included in

EPA's list of hazardous constituents, the sludge or waste is classified as hazardous. For certain industrial operations, materials of known hazardous characteristics are routinely used. As a result, the waste discharges from these operations are generally considered to contain these hazardous substances; hence the source is assigned a specific hazardous waste number. There are 15 nonspecific sources and 80 specific sources listed in the N. C. Hazardous Waste Management Rules & Solid Waste Management Law. Additionally, there are 657 substances defined as hazardous constituents that have been assigned an EPA hazardous waste number. The hazardous waste number consists of a letter followed by a three digit number. For example, F019 refers to a hazardous waste source containing treatment sludges from the chemical conversion coating of aluminum while P023 indicates a hazardous substance known as chloroacetaldehyde.

In addition to the two categories of criteria listed above, several other sludge characteristics are considered when classifying industrial sludge. Among them are:

1. pH: The ideal pH range for non-hazardous and non-toxic waste sludge is 5.5 - 9.

2. Moisture Content: Moisture content of the sludge is determined by the use of the Filter Test. If the moisture content is greater than 70% or if the solid content is less than 30%, the sludge may not be suitable for disposal in a landfill site designed for non-hazardous and non-toxic municipal waste and sludge.

3. Impurities: The sludge may not contain more than ten times the quantities of impurity constituents as listed in the drinking water quality standards. If it does, chemical fixation is recommended to fix these soluble impurities.

4. Chemical Fixation: If a sludge is chemically fixed to stabilize some of its constituent impurities then it may be considered favorably for disposal in a landfill for non-

hazardous and non-toxic waste and sludge. When the sludge contains EP (Extraction Procedure) toxicity and impurities in quantities exceeding ten times those specified in drinking water standards, the industry is encouraged to conduct chemical fixation of the sludge in order to change its classification.

Another classification consideration is that some industrial sludge originates from a point source whose effluent quality is limited by a National Pollutants Discharge Elimination System (NPDES) discharge permit. The sludge from such a source is therefore regulated by other State regulatory agencies and hence, must be referred to these agencies for a permit to be disposed of in a municipal landfill.

A KNOWLEDGE-BASED SYSTEM FOR CLASSIFYING INDUSTRIAL SLUDGE:

Expert System Architectures:

Expert systems consist of an inference engine, a knowledge base, and a cashe. The inference engine is the processor of the system. It uses the knowledge stored in the knowledge base as well as the facts provided by the user to reason and make heuristic decisions. The knowledge which is used by the inference engine to derive conclusions consists of facts and rules and is contained in the knowledge base (3).

To build a knowledge base, the knowledge engineer excerpts domain-specific knowledge from an expert and converts the knowledge into coded rules and facts. In M.1 , these rules and facts are stored on a disk and are loaded into the M.1 run time system prior to consultation. During a consultation, the user interacts with the system to provide it with background information that describes the problem. In the M.1 system, intermediate results and final conclusions are temporarily stored in the M.1's memory which is called a cache. The contents of cache can be saved on a disk for future continuation of the consultation. Unlike programming languages, an expert system tool is modular to the extent that execution of its rules are independent of their order in the knowledge base and that each rule or fact can be entered, altered, or removed in a manner that is largely independent of the others (4, 5).

M.1 Knowledge Engineering Tool and its Main Features:

M.1 is a microcomputer-based knowledge engineering tool developed by TEKNOWLEDGE, Inc. for building small knowledge systems with limited computational requirements. An M.1 knowledge base may contain up to 1000 facts and rules that enable it to solve structured selection problems such as those encountered in diagnosis-prescription situations (8). It is a goal-driven system such that once a goal is defined, M.1 asks for input information only when it is needed. Additionally, it can reach tentative conclusions using incomplete or inexact information. The system can justify conclusions and offer explanations.

Knowledge Representation:

Like most production systems, the knowledge base of M.1 consists of a collection of rules and facts. A fact is a description of the value of a variable or the relationship between variables while a rule derives new facts based on given facts. Each rule consists of an **IF** and a **THEN** part. The **IF** part specifies a set of conditions and the **THEN** part specifies a set of actions. When the conditions are satisfied, the actions specified in the THEN part are executed. As an example consider a rule taken from the sludge classification prototype:

```
Rule 1:     IF
            npdes_test = passed and
            hazardous_test =passed and
            ph_test = passed and
            moisture_content_test = passed and
            impurities_test = passed and
            solubility_test = passed and
            adsorption_test = passed and
            biodegradation_test = passed and
            volatilization_test = passed
            THEN
            sludge_classification = non_hazardous
```

In this rule, if all of the conditions are satisfied, i. e., if all of the specified tests are passed, then the sludge is classified as non-hazardous. However, to evaluate each condition requires the execution of another rule. For example, to determine a value for the NPDES_test requires the execution of the following rules:

```
Rule-2:     IF
            npdes_sludge = no
            THEN
            npdes_test = passed.
```

This rule states that if the answer to the question "Is the sludge an NPDES sludge?" is "no", then the npdes test is passed. Otherwise the sludge is classified as an NPDES sludge by the following rule:

```
Rule-3:     IF
            npdes_sludge = yes
            THEN
            sludge_classification = npdes_sludge.
```

To determine a value for the conditions of Rules 2 and 3, the M.1 inference engine must ask the the user whether or not the sludge is an NPDES sludge. Thus a question is posed to the user and its M.1 form is as follows

```
kb-3:      question(npdes_sludge) =
           [Is the sludge an NPDES sludge?]
```

Once rule kb-3 is activated, M.1 displays on the monitor the question included between the two square brackets. The user then types in his answer and the variable "npdes_sludge" is bound to that answer. Since in Rules 2 or 3 the value of "npdes_sludge" must be either "yes" or "no", it is necessary to limit the answer that a user may provide to "yes" or "no" by specifying the allowable legal values that "npdes_sludge" can posses as follows:

```
kb-4:      legalvals(npdes_sludge) = [yes,no]
```

Variables can also be used used in rules as is shown in the sludge pH rules:

```
rule-44:   IF
           sludge_ph = X1 and
           X1>=6 and
           X1<=9
           THEN
           ph_test = passed.

rule-44:   IF
           sludge_ph = X1 and
           X1<6 and
           X1>9
           THEN
           sludge_classification = hazardous_sludge.

kb-64:     question(sludge_ph) =
           [What is the pH of the sludge ?].

kb-65:     legalvals(sludge_ph) = real.
```

In this example, sludge pH is a variable whose value is provided by the user and is limited to real numbers. When the sludge pH specified by the user falls between 6 and 9, the pH test is passed, otherwise the sludge is classified as hazardous.

Inference AND/OR Tree

The order of rules and facts in the M.1 knowledge base is indifferent to the execution of an M.1 program. When a rule is activated depends on its goal and the knowledge needed to achieve that goal. For example, the goal of "sludge classification" is to classify the sludge which is represented by a rule called kb-1 as follows:

```
kb-1:     goal = [sludge_classification]
```

Once the main goal is established, the M.1 inference engine will first test whether the solution is already stored as a fact in the cache. If it is not and "sludge_classification" is unknown the inference engine will then determine what conditional rules need to be activated in order to achieve the goal. In this example, the rules which lead to achieving the goal of classifying sludge as non-hazardous will be activated before other rules:

```
Rule 1:   IF
          npdes_test = passed and
          hazardous_test =passed and
          ph_test = passed and
          moisture_content_test = passed and
          impurities_test = passed and
          solubility_test = passed and
          adsorption_test = passed and
          biodegradation_test = passed and
          volatilization_test = passed
          THEN
          sludge_classification = non_hazardous
```

This rule is numbered "Rule 1" in the knowledge base. However, it is activated only when the goal of classifying sludge has been established.

The order of activating rules can be better illustrated by an AND/OR tree. Each node represents the goal, or subsequently, the sub-goals to be sought. The goal of classifying sludge is represented by the root goal. The root goal for a non-hazardous sludge is further divided into sub-goals for the various tests. All of these tests must pass in order to achieve a goal evaluation of non-hazardous sludge.

The M.1 inference engine uses a left to right evaluation of the nodes of the inference tree. The subnodes correspond to the conditions of the rule and are in the same order as its nodes. The conditions of Rule 1 correspond to the nodes of Figure 2. Each condition may itself be the root of another set of sub-goals. The M.1 inference engine will traverse through each sub-goal and activate the appropriate rules to establish the conditions for satisfying the rule. Thus, the order of rules that are actually activated to solve a particular problem depends on whether the knowledge it embodies is needed and whether it is available to achieve a goal.

Building the Prototype

A prototype expert system has been built to test the feasibility of classifying industrial sludge based on the aforementioned criteria. The program consists of 55 rules and 69 knowledge base entries.

Since M.1 has a limited amount of available memory space for its cache, data entered by the user in response to questions has been generally limited to yes/no answers which currently serve our needs. For example, when it is necessary to find out whether the sludge is from a specific or non-specific source or the sludge contains hazardous components, a preferred approach is to first instruct the user to enter the source identification number as follows:

```
kb-x:      question(sludge_id_number) =
           [Enter the Sludge ID number].
```

The data item "sludge_id_number" can then be used in subsequent M.1 rules as follows:

```
rule-1x:   IF
           sludge_id_number = N and
           sludge (N) = Hazardous_description
           THEN
           sludge_type = Hazardous_description.

rule-2x:   IF
           sludge_type = Hazardous_description
           THEN
           sludge_classification = hazardous.
```

In the above example, the user types in the sludge identification number N, and the sludge type and sludge classification are then determined by M.1 if facts in the form of "sludge(N) = Hazardous_description" are provided in the fact portion of the knowledge base. Examples of such facts are:

```
sludge(K021) = "Aqueous spent antimony catalyst waste from
                fluoromethanes production".

sludge(K022) = "Distillation bottom tars from the production
                of phenol/acetone.".

sludge(K023) = "Distillation light ends from the production
                of phthalic anhydride from naphthalene".
```

This list can be continued to cover all types of specific and nonspecific hazardous sources and constituents. Providing such rules and facts is convenient for the user and is essential for proper expert system operation. However, the list of hazardous sources and constituents is lengthy and thus cannot be fully accommodated by the available M.1 memory space. An alternative approach is to instruct the user to look up in the tables of hazardous sources and constituents and provide a yes/no answer to the follow question:

```
kb-54:    question(hazardous_source) =
          [Is the sludge from a hazardous source?].

kb-55:    legalvals(hazardous_source) =
          [yes,no,unknown].
```

The answer to question kb-54 could be "yes", "no", or "unknown". A response of "unknown" indicates that the user does not know the answer and a follow-up question instructing him(her) to determine the value from a table is posed:

```
kb-58:    question(nonspecific_source_test) =
          [Please refer to Section 261.31, North Carolina
           Hazardous Waste Management Rules and Solid Waste
           Management Law,  Is the sludge from any of the
           nonspecific source?]

kb-59:    legalvals(specific_source_test) = [yes,no].
```

After checking the manual, the user must enter either "yes" or "no" to the question. This approach saves memory space but requires that the user perform manual operations prior to answering the questions. A similar approach is also necessary for entering other information such as ignitability, corrosivity, reactivity, and EP hazardousness.

Testing the Prototype

The prototype M.1 knowledge based expert system for classifying industrial sludge has been operated and tested. hazardous sludge classification. Other trial tests have been carried out using conditions for hazardous sludge with different characteristics.

EVALUATING THE USE OF M.1 FOR CLASSIFYING INDUSTRIAL SLUDGE

Advantages and Drawbacks

The prototype knowledge-based system has demonstrated its advantages as well as drawbacks for classifying industrial sludge. Clearly a tool like M.1, when used in an application like that described here, is easier to program and debug than conventional programming languages. It is particularly advantageous that the order of the rules and facts in the knowledge base does not affect the execution of the program. The result is that the program can be easily modified to accommodate changes in governing regulations without much effort.

The program is considered to be user-friendly enabling one to review all considerations for classifying industrial sludge. It is especially helpful for assisting an inexperienced environmental engineer in carrying out the task of classifying sludge. Additionally, the recommended classification by the program is rather uniform regardless of the user's familiarity with the governing regulations. Hence, this knowledge-based system in its current form is regarded as a reasonably valuable