To Mary Kay O'Connor Process Safety Center Home Page
To Program details for Day 1
To Program details for Day 2

Expert Systems in Emergency Response

Albert J. Slap, Daniel Hillman, and David A. Moore, PE, CSP
AcuTech Consulting Inc.
100 Pine Street, Suite 2240
SanFrancisco, California 94111

Paper presented at the 1st Annual Symposium of the Mary Kay O'Connor Process Safety Center,
"Beyond Regulatory Compliance, Making Safety Second Nature",
George Bush Presidential Conference Center, College Station, Texas
March 30-31, 1998.

Abstract

Ensuring the proper response to emergencies is an ideal application for expert systems. Rules can be developed that comprise the appropriate response to either preplanned, site-specific incident scenarios or general incidents at the facility. This allows responders with less experience to make better decisions and fewer errors than otherwise possible without the involvement of a "human expert." An expert system for emergency incident management can allow for better pre-planning, training, and emergency response.

An expert system can assist with making critical emergency response and incident management decisions, coordinate and automate emergency notifications, and track and log actions that were taken. As a decision-support system, it can integrate facility drawings and other reference data. Having an interactive expert system as a resource allows emergency information and procedures, and the knowledge and experience of emergency response experts to be captured for easy access and use by all

personnel at a facility. This paper discusses the use of expert systems for emergency response, and provides an example of a working system specifically addressing emergency management and incident response.

Expert Systems in Emergency Response

By Albert J. Slap, Daniel Hillman, and David Moore

Introduction

This paper will discuss the use of expert systems in the field of emergency response. Emergency response experts can make more informed decisions during the complex moments of a disaster. And, yet, there are far fewer experts than problems for them to address. Even so, the pressures of handling the emergency and the likelihood and consequences of mishandling the event call for expert computer systems to assist the decision-makers in their myriad tasks.

The proper response to emergencies is an ideal application for expert systems. Rules can be developed that comprise the appropriate response to either preplanned, site-specific incident scenarios or general incidents at the facility. This allows responders with less experience to make better decisions than otherwise possible without the involvement of a "human expert."

An expert system can assist with making critical emergency response and incident management decisions, coordinate and automate emergency notifications, and track and log actions that were taken. As a decision-support system, it can integrate facility drawings and other reference data. Having an interactive expert system as a resource allows emergency information and procedures, and the knowledge and experience of emergency response experts to be captured for easy access and use by all personnel at a facility.

Expert Systems

Expert systems are a type of artificial intelligence. Artificial intelligence (AI) is an attempt to use machines to mimic or enhance human intelligence. Expert systems are used today more than any other type of applied AI technology. See, Decision Support Systems and Intelligent Systems, Turban and Aronson, Prentice Hall, 1998, Ibid.

The basic idea of an expert system is simple. Expertise is transferred from an expert to a computer. The expert’s knowledge is stored in the computer’s memory and, users can access the knowledge as it is needed. Like a human expert, the expert system can make inferences and arrive at specific conclusions. The expert system is a decision-making or problem-solving computer program that can achieve levels of performance comparable to or even exceeding that of a human expert. Id.

Role of Experts in the Decision Process

Let's look at how individuals or organizations make decisions in the first place. In the case of a simple decision, the information and factors that go into the decision may be easily obtained and understood by an individual or organization. In the case of a complex decision, however, individuals and organizations often turn to experts for advice. "… [E]xperts have specific knowledge and experience in the problem area. They are aware of the alternatives, the chances of success, and the benefits and costs the business may incur." Id. at 17. The more undefined the situation, the more specialized and, perhaps, expensive the expertise becomes. Expert systems seek to mimic the way human experts solve problems.

If an expert system mirrors a human expert, how, then, does the human expert do his/her expert work? What are the attributes of "expertise" that we want the machine to mimic? According to Turban and Aronson:

> Human experts tend to specialize in narrow problem-solving areas or tasks. Typically, human experts possess these characteristics: They solve problems quickly and fairly accurately, explain what (and sometimes how) they do, judge the reliability of their own conclusions, know when they are stumped, and communicate with other experts. They can also learn from experience, change their points of view to suit a problem, and transfer knowledge from one domain to another. Finally, they use tools, such as rules of thumb, mathematical models, and detailed simulations to support their decisions.

These are the attributes that we also expect to see in an AI expert system. The human expert generally solves problems using a two-step process. First, the expert will obtain the information he/she needs from the external world. This information may be received from people, computers, the expert's memory or other media. Secondly, the expert uses inductive, deductive or other problem-solving approaches on the information that has been collected. The result is a recommendation on how to solve the problem. The expert system (ES), therefore, attempts to imitate the reasoning process and knowledge of experts. It, too, will obtain data from external sources; it can consult with mathematical models, and then uses its inferencing capability to determine the best course of action—that is, make recommendations.

Organizational Use of Expert Knowledge

Organizations, like individuals, look for successful solutions to problems that require expertise. An expert's recommendation may be *part* of a solution that involves the broader task of "managing knowledge." The expertise of an expert may need to be coupled with "experience", that is, a company's knowledge of what has or has not worked before in a given situation and data resources, including on-line information. These additional factors provide valuable information to the expert and non-expert alike to enhance decision-making performance. When an expert system is coupled with resources such as databases, models, and communications tools, it is transformed into a powerful decision-support system (DSS). Other terms for these configurations are: management support systems, group support systems, and executive information systems.

Uses of Expert Systems

The purpose of these systems is the same: higher decision quality, improved communication, cost reduction, increased productivity, time savings, and improved employee satisfaction. The key to any AI application is not necessarily *what* the system is called, but whether it can solve managerial or organizational problems faster and better than without the application.

Expert systems were developed by the AI community as early as the mid-1960's. Special purpose expert systems were first developed at Stanford University. One of the early Stanford systems was used for diagnosis and cause of hospital-borne infections. The first expert systems ran on special computers. Since the late 1980's, however, expert systems are run on most standard computers, including PCs. Id.

Virtually thousands of expert systems are in use today in almost every industry and functional area. For example, COMPAQ, the world's largest manufacturer of PCs, uses an expert system to allow its customers to solve problems with network printers on their own. American Express now makes charge approval decisions in less than 5 seconds, compared to about 3 minutes before implementation of its on-line expert system. Id. Other expert systems operate in such diverse areas as: stock portfolio management, libraries, retail chains, and farm management. Expert systems are in use in military and space programs and throughout industry in a wide variety of applications.

The Rules-Based Expert System

What are the reasons for an organization to use an expert system and particularly a rules-based expert system? According to Turban and Aronson:

> Knowledge is a major resource, and it often lies with only a few experts. It is important to capture that knowledge so others can use it. Experts get sick or become unavailable, so knowledge is not always available when needed. Books and manuals can capture some knowledge, but they leave the problem of a particular application up to the reader. Expert systems can provide a direct means of applying expertise. The purpose of an expert systems is not to replace the experts, but to make their knowledge and expertise more widely available. An expert system permits nonexperts to increase their productivity, improve the quality of their decisions, and solve problems when an expert is not available. Id.

Typically, an organization will consider utilizing an expert system in the following situations:

> when the solution to the problem has a high payoff;
>
> when the expertise is needed in many different locations;
>
> when expertise is needed in hostile or hazardous environments;
>
> and when the ES is needed for training as well as decision-making.

Other factors favoring the use of an ES include:

> the cost of maintaining expertise within the organization is high;
>
> large amounts of data must be sifted through in the decision-making process;
>
> an error in the decision-making process could lead to disastrous results;
>
> there is a shortage of experts available to the organization; and,
>
> expertise is needed to augment the knowledge of junior personnel.

Most commercial ES are rules-based systems; that is, the knowledge is stored mainly in the form of rules. These rules are a type of knowledge known as heuristics or "rules of thumb." Rules are typically expressed in the manner of IF/THEN statements. These expert rules are then arranged in a fashion similar to a decision- tree matrix. A typical rule may look like this**: IF** the release is **CHLORINE** and is **AT THE RAILCAR OFFLOADING STATION**, **THEN**, the **RESPONSE ACTION** is **CLOSE LOADING SYSTEM ISOLATION VALVE**. A rule is said to "fire" when all of the rules hypotheses or antecedents are satisfied. Rules can also be used to activate other areas of the computer system such as mathematical models or databases.

So, for example, a rule may also be constructed as follows: **IF** the spill is **CHLORINE**, **THEN** obtain **CHEMICAL PROPERTIES** from **TABLE W** and **METEROLOGICAL DATA** from **WEATHER TOWER X** and activate **PLUME DISPERSION MODEL Y** and **DISPLAY** the results graphically over **GIS MAP Z**. Another example of operation of a rules-based ES system would be as follows: **IF** the spill is **CHLORINE**, **THEN** activate **AUTOMATED EMERGENCY NOTIFICATION SYSTEM** (autodialer) and **SELECT CHLORINE** pre-recorded message and **NOTIFY** chlorine response list. **IF** the Chlorine release is **DIRECTED TOWARDS THE SCHOOL**, **THEN**, the **EVACUATION ACTION** is **CALL SCHOOL AND SHELTER IN PLACE**.

An Emergency Response Expert System

The idea of using an expert system in real-time for emergency response has been a goal of AI professionals for over a decade. An early concept for an emergency response expert system included rules sets and recommendations and activation of mathematical models for oil spills. In the case of the theoretical oil spill expert system, the computer ran a mathematical model of the movement of oil on a body of water and displayed recommendations to the user. In this way, both the expert and nonexpert could see and utilize the results of the model in their decision-making processes much faster than would be humanly possible utilizing manual methods.

In the early 1990's, an emergency response expert system was prototyped for hazardous material releases in the petrochemical industry. This system was conceptualized as a complete incident command decision support tool. The system, called PlantSafe^{TM}, contained heuristic rules, a toxic gas dispersion model, access to maps and diagrams, and contained other databases. These other databases included: material safety data sheets, building and tank information, personnel records, notification lists, various checklists, and other information valuable to emergency responders.

The PlantSafe system was beta tested in 1993-1994 at a large chemical plant in Pennsylvania. During the Beta test phase, the developer decided to couple the expert system with an automatic telephone messaging system, called TeleSafe^{TM}. The purpose of this addition was to increase the speed at which the response personnel on-site at the beginning of an incident could summon help or notify the community of the situation.

Since 1994, the PlantSafe^{TM} and TeleSafe^{TM} systems have gained commercial acceptance at some of the largest chemical plants and refineries in the world.

The PlantSafe system was written in "C" running under DOS, with stand-alone capabilities, only. With the advent of object-oriented, rule-development software methods and client-server architectures, the system is presently being migrated to the Windows NT operating system. This will allow multiple users throughout a plant or corporation to respond to emergencies simultaneously using PlantSafe on their individual workstations. The ability to use the ES cooperatively more accurately reflects actual needs and usage patterns of an incident command team during an emergency or training.

The NT system will also utilize databases that are resident on other servers, mainframes, or even the Internet. These may include: process control data, personnel or maintenance records, etc. The PlantSafe NT rules will operate various customer-designated mathematical models, from plume dispersion, to blast effects, to oil spills. A simulator program will allow customers to assign a performance time factors to each rule or recommendation. This feature will allow incident command members to check their performance during real-life emergencies or during training exercises. The expert rules will also manage the TeleSafe automated emergency notification system. TeleSafe is now expanded to communicate with the latest alphanumeric paging equipment. The PlantSafe NT system will also display the latest GIS maps and CAD overlays. New software technology will enable "on-line chat" among incident command members over the local area network, while working inside of the PlantSafe program. This feature adds an important internal communications tool to this decision support system. Finally, the ES will accept digital or analog signals from various parts of the distributed control system (process controls, alarm management), that will then trigger rules firing and display specific recommendations.

Implementation of the Expert System

Purchase and implementation of an emergency response expert system at a major petrochemical plant is a process that involves a number of people, both inside and outside the organization. The purchasers of these systems typically tend to be the emergency response managers and fire chiefs, with approval of by the plant manager. The users of the system are the members of the incident command teams, shift supervisors, dispatchers, and hazardous materials technicians. The corporate management information systems (MIS) department is usually involved in the implementation of the system and the integration of the PlantSafe software with the existing software and hardware at the plant. At the corporate level, the environmental, safety and health department, corporate crisis management, and information systems may also evaluate the system for company-wide applicability.

In order for the company’s corporate emergency response knowledge to be incorporated into the rules-base within the ES, a knowledge engineer (part of the developer’s staff) works with domain experts, within the company or outside consultants, in a collaborative approach to transfer the company’sexpert’s knowledge to the machine. This is the process of rules creation.

The system users are also very important in the process of configuring the ES. User input will help organize the maps and CAD overlays, specify the order of cues, and identify plant-specific recommendations, and other similar elements required to implement an effective site-specific decision support system.

Because an expert system evolves, it is never really complete. The user will also be The ES must be revised on a regular basis with regard to the applicability of the rules, the integrity and quality of the data inputsfeeds, the use of the interlinked databases, and so on. New "rules of thumb" become available as the state-of-the-art in each domain area improves. Experts are constantly training themselves on new situations, or reorganizing their knowledge to account for unencountered situations. Thus, an expert system must be adjusted for these cases. In addition, software and hardware bugs must be fixed as found and the system must be upgraded to run on new software and hardware platforms. Id.

The Future

The "first generation" of AI has seen the development and commercial acceptance of expert systems, and rules-based systems, in particular. The second generation of AI is focusing on a concept called "machine learning". That is, programs which allow computers to "learn" from their surroundings. This is also called neural computing. Neural computing is a technology that endeavors to use pattern recognition as a means of computer learning. Thus, a computer may look for patterns in data and, using mathematical based on algorithms, and develop rules or recommendations based on theose patterns in the data. In this way, raw information can be is "automatically" converted into "knowledge" by the computer that can be acted upon by other programs, such as expert systems. Another way of looking at this is that neural computing is a way for a computer to obtain heuristic "rules of thumb", without having that knowledge directly transferred to it by human experts.

According to Turban and Aronson: "An artificial neural network can be useful for fast identification of implicit knowledge by automatically analyzing cases of historical data. The ANN analyzes the data sets to identify patterns and relationships that may subsequently lead to rules for expert systems." Id. at 674. Artificial NN's canN's learn from historical cases. The learning (training) produces the required values or weighting of information, which is then used in the inferencing process., which make the computed outputs equal (or close) to the desired outputs For example, neural networks are used by John Deere & Co. to invest part of its pension fund (about $1billion).

Another branch of AI is called genetic algorithms. This area of AI has been defined as an "iterative procedure maintaining a population of structures that are candidate solutions to specific domain challenges." A genetic algorithm system is in use by United Distillers, to control blending of whiskey and minimize the movements of whiskey casks. Id. at 702.

One of the latest areas of AI is "fuzzy" logic. This area of AI deals with the problem of uncertainty. Fuzzy logic uses the mathematical theory of "fuzzy sets" to simulate the process of normal human reasoning by allowing the computer to behave less precisely and logically than conventional computers do. According to Turban and Aronson:

> In a standard rule-based system, a production rule has no concrete effect at all, unless the data completely satisfy the antecedent of the rule. The operation of the system proceeds sequentially, with one rule firing at a time; if two rules are simultaneously satisfied, a conflict resolution policy is needed to determine which one takes precedence. In a fuzzy rule-based system, in contrast, *all* rules are executed during each pass through the system, but with strengths ranging from "not at all" to "completely," depending on the relative degree to which their fuzzy antecedent propositions are satisfied by the data. Id. at 706

Conclusions

In conclusion, the use of artificial intelligence (AI) via a rules-based expert system in the area of real-time emergency response has been described. As the area of AI develops, new techniques for improving emergency response decision support systems (DSS) will be incorporated into existing and new computer software. In the not too distant future, such decision support systems may operate over the Internet and may keep themselves automatically refreshed with new rules and new relevant data by sifting through data on the Internet.

Ensuring the proper response to emergencies is an ideal application for expert systems. An expert system for emergency incident management can allow for better pre-planning, training, and emergency response.

To Mary Kay O'Connor Process Safety Center Home Page
To Program details for Day 1
To Program details for Day 2