This Is a Publication of
The American Association for Artificial Intelligence

This electronic document has been retrieved from the
American Association for Artificial Intelligence
445 Burgess Drive
Menlo Park, California 94025
(415) 328-3123
(415) 321-4457
info@aaai.org
http://www.aaai.org

(For membership information,
consult our web page)

CHEMREG

Using Case-Based Reasoning to Support Health and Safety Compliance in the Chemical Industry

Kirk D. Wilson

■ CHEMREG is a large knowledge-based system used by Air Products and Chemicals, Inc., to support compliance with regulatory requirements for communicating health and safety information in the shipping and handling of chemical products. This article concentrates on one of the knowledge bases in this system: the case-based reasoner. The case-based reasoner addresses the issue of how proper communication of public health and safety information can be ensured while rapid and cost-effective product evaluation is allowed in the absence of actual hazard testing of the product. CHEMREG generates estimates of hazard data for new products from similar products using an existing relational database as a case library. Implementation of the case-based reasoner in rules and objects using a commercial knowledge-based system shell is described. Although some refinements remain, the performance of the case-based reasoner has met its design goals.

C HEMREG is a knowledge-based system (KBS) begun in 1993 by Air Products and Chemicals, Inc., of Allentown, Pennsylvania. Today, CHEMREG consists of three knowledge bases, a large relational database of product information, and an online system that is used at our facilities throughout the United States and Europe. The purpose of the system is to support compliance with regulatory requirements for communicating health and safety information on chemical products by generating shipping descriptions and material safety data sheets.

Background: Purpose and Justification of CHEMREG

The chemical industry is heavily regulated. Every hazardous chemical product must have a set of shipping descriptions that conform to strict regulations issued by the Department of Transportation (DOT), the International Maritime Organization (IMO), and the International Air Transport Association (IATA). Shipping descriptions provide a concise characterization of the hazards a chemical can present during transportation (figure 1). Failing to comply with transport regulations can result in penalties ranging from delayed shipments to heavy fines or even incarceration of corporate officials. In addition, each chemical product has a material safety data sheet (MSDS) that conforms to Occupational Safety and Health (OSHA) and American National Standards Institute (ANSI) standards. Unlike shipping descriptions, *MSDSs* are lengthy documents of 8 to 10 pages that provide a detailed description of the health hazards a product can pose in the workplace (figure 2). They also contain information on procedures for storing, handling, and disposing of a chemical. Inadequately prepared MSDSs can lead to substantial product-liability lawsuits against the company if the product is involved in an industrial accident.

The ultimate goal of these regulations is to ensure proper communication of health and safety information for the protection of the public. Air Products is committed to the initiative of the Chemical Manufacturers Association (CMA) known as Responsible Care. This initiative focuses on the safe manufacturing, distribution, use, recycling, and disposal of chemicals. Proper communication through accurate shipping descriptions and full disclosure of hazard information in the MSDS plays a key role in fulfilling obligations under Responsible Care.

Maintaining shipping descriptions and MSDSs requires a major effort. Most corporate systems are intensely manual. Shipping regulations change periodically, requiring transporta-

Specifically regulated chemical:

Methanol // 3 // UN1230 // PG II // RQ (Methyl alcohol)

Chemical regulated by family:

Alcohols, n.o.s. (METHYL-3-BUTYN-2-OL,2-) // 3 // UN1987 // PG II

Hazardous Chemical "Not Otherwise Specified":

DOT *(road)*: Corrosive liquids, flammable, n.o.s. (Dimethylcyclohexylamine) // 8 // UN2920 // PG II // (3)

DOT *(rail)*: Corrosive liquids, flammable, n.o.s. (Dimethylcyclohexylamine) // 8 // UN2920 // PG II // HazMat STCC=4936601 // (3)

Product with different shipping descriptions across agencies:

DOT: Xylenes solution // 3 // UN1307 // PG III

IMO: Flammable liquid, n.o.s. (Xylenes) // 3.3 // UN1993 // III // F.P. 37.8 C // HazMat STCC=4910185 // EmS No: 3-07 // MFAG No: Refer to subsection 4.2

Products with different shipping descriptions for the same agency:

Product 1: DOT *(packaged)*: Chemicals, N.O.I. - Not DOT Regulated

DOT *(bulk)*: Combustible liquid, n.o.s. (Silicone) // NA1993 // PG III

Product 2: DOT *(packaged)*: Environmentally hazardous substances, solid, n.o.s. (Toluenediamine) // 9 // UN3077 // PG III // RQ Toluenediamine)

DOT *(bulk)*: Elevated temperature liquid, flammable, n.o.s. // 3 // UN3256 // PG III // RQ (Toluenediamine)

Figure 1. Sample Shipping Descriptions.

Shipping descriptions vary depending on the regulatory agency, the characteristics of the product, and the containers in which the product is shipped.

Air Products and Chemicals, Inc.
7201 Hamilton Boulevard
Allentown, PA 18195-1501

AIR PRODUCTS

Telephone (610) 481-4911

TARGET ORGANS

Eye
Skin
Respiratory system

SIGNS AND SYMPTOMS OF EXPOSURE (Acute effects)

Product vapor in low concentrations can cause lacrimation, conjunctivitis and corneal edema when absorbed into the tissue of the eye from the atmosphere. Corneal edema may give rise to a perception of 'blue haze'' or 'fog'' around lights. The effect is transient and has no known residual effect. Burns of the eye may cause blindness. Inhalation of vapors may cause irritation in the respiratory tract. Contact of undiluted product with the eyes or skin quickly causes severe irritation and pain and may cause burns, necrosis and permanent injury. Ingestion may cause death unless treated promptly.
Inhalation of aerosols and mists may severely damage contacted tissue and produce scarring.
Risk of exposure to hazardous concentrations of vapor under normal working conditions in a well ventilated space is minimal. However, conditions such as spraying, or sudden release of hot liquid, which generate an aerosol, mists or fog should be avoided.
Product is absorbed through the skin and may cause malaise, discomfort, injury and death unless treated promptly.

SIGNS AND SYMPTOMS OF EXPOSURE (Possible Longer Term Effects)

Repeated and/or prolonged exposure may cause allergic reaction/sensitization.
Repeated and/or prolonged exposures may result in: adverse respiratory effects (such as cough, tightness of chest or shortness of breath), adverse eye effects (such as conjunctivitis or corneal damage), adverse skin effects (such as rash, irritation or corrosion)

MEDICAL CONDITIONS GENERALLY AGGRAVATED BY EXPOSURE

Asthma
Chronic respiratory disease (e.g. Bronchitis, Emphysema)
Eye disease
Skin disorders and Allergies

CARCINOGENS UNDER OSHA, ACGIH, NTP, IARC, OTHER

This product contains no carcinogens in concentrations of 0.1 percent or greater.

SECTION 4 - FIRST AID

EYE CONTACT

Hold eyelids apart and immediately flush eyes with plenty of water for at least 15 minutes. Seek medical advice.

SKIN CONTACT

Remove product and immediately flush affected area with water for at least 15 minutes. Remove contaminated clothing and shoes. Destroy contaminated leather apparel. Cover the affected area with a sterile dressing or clean sheeting and transport for medical care. DO NOT APPLY GREASES OR OINTMENTS. Control shock, if present. Launder contaminated clothing prior to reuse.

POLYCAT• 8 CATALYST   10/14/1997 - Page 3 of 10

*Figure 2. A Typical Page Extracted
from a Material Safety Data Sheet.*

tion groups to scramble to reevaluate thousands of shipping descriptions within a limited time. Because corporations usually maintain MSDSs as separate documents in word-processing files that have been written over a long period of time, keeping their contents up to date and consistent constitutes a significant challenge.

These problems were well known to Air Products and Chemicals. We are a $4 billion, international industrial gas and chemicals company. Approximately one-third of our revenues are generated from the sale of chemicals. Our product lines consist of a substantial number of multicomponent, formulated products, including polymer emulsions, catalysts, specialty performance additives, and amine products. In 1993, the chemical industry confronted a major change in DOT shipping regulations. At the time, we deemed that manually reevaluating all products for new shipping descriptions was impractical. We had already investigated automating the process of generating shipping descriptions and had identified the following goals for an automated system:

First was to significantly reduce the 10 to 30 days it was taking for new products to go through the regulatory compliance process. This goal would have to be reached while it retained 100-percent compliance with the regulations.

Second was to provide a consistent level of service without increasing staff despite rapid new product introductions and acquisitions of new product lines. In addition, the system had to allow efficient response to changes in the regulations when significant numbers of products need to be reevaluated.

The decision was made to develop a knowledge-based system, using AION, to fulfill these goals and meet the new regulatory requirements. Programming on the transportation knowledge base began in early 1994—just nine months before the regulatory changes were to take effect.

The goals of CHEMREG dovetailed with another long-standing goal of Air Products: to have a central database of product information. Thus, the basic architecture of CHEMREG was born (figure 3). An online system was distributed to Air Products business managers and laboratory technologists throughout the United States and to several of its facilities in Europe for them to add technical data to a central database. The design specifications of CHEMREG do not require that shipping descriptions be generated in real time. Therefore, the knowledge base runs twice daily as an offline process that generates shipping descriptions
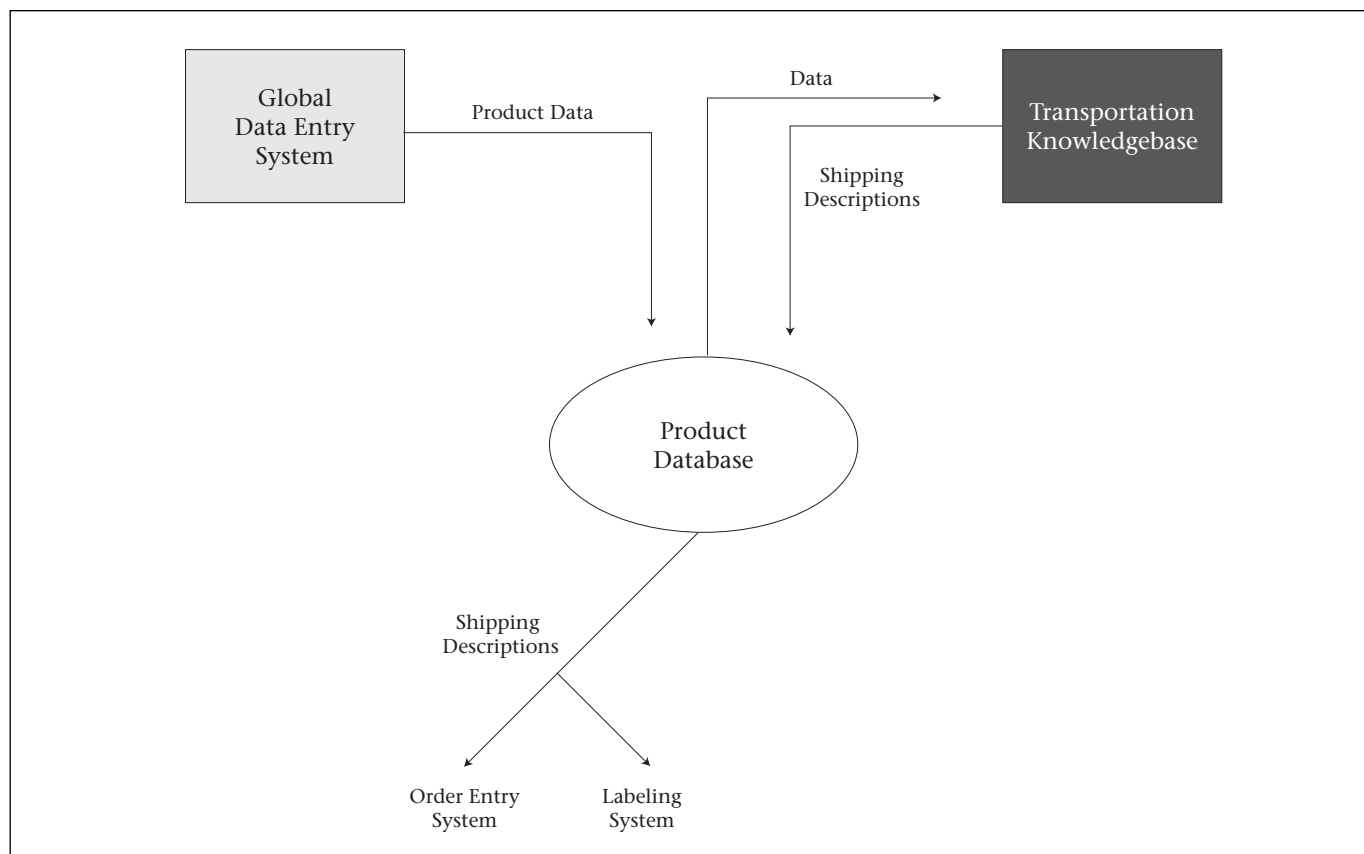
Figure 3. Initial Configuration of the CHEMREG System.

for products on which maintenance has been performed since the last cycle. The knowledge base writes shipping descriptions back to the database where they are available to the order-entry system, which produces bills of lading. The descriptions are also fed to the container labeling system.

The transportation knowledge base was completed in August 1994 in sufficient time to meet the new DOT requirements. In addition, the knowledge base generated shipping descriptions under IMO (sea) and IATA (air) regulations.

The knowledge base has accomplished its operational goals. It has reduced the time required to process regulatory information for new products to two days, the minimal time to accommodate plant labeling operations. It has also allowed the transportation professionals to meet business demands for rapidly introducing new products and respond quickly to regulatory changes without increasing staff.

A second knowledge base was added to the system in 1995 for generating (authoring) the contents of the MSDS. MSDSs cover a myriad of interrelated topics pertaining to health and safety issues for a chemical product. The MSDS

knowledge base represents the contents of the MSDS as a semantic network of the various topics covered in the MSDS. Today, the semantic network consists of approximately 200 nodes. Rules at each node generate what is said on a particular topic for a product.

A Common Industry Issue

Both knowledge bases share a common design philosophy of mapping classificatory rules directly into knowledge base code. These rules determine the hazard classification of a finished product whether this product is a chemical substance (reacted product) or a mixture. For example, a product might be classified on the basis of a physical property such as its flash point; for example, it is a flammable material. Regulatory rules are also based on the product's harmful effects on living organisms. Such properties include oral, dermal, and inhalation acute toxicity as well as corrosivity and irritation.

It is both time consuming and expensive to perform a complete set of tests on every product. Air Products estimates that to conduct tests on all its products would cost $1 million to $4 million. However, not testing for all haz-
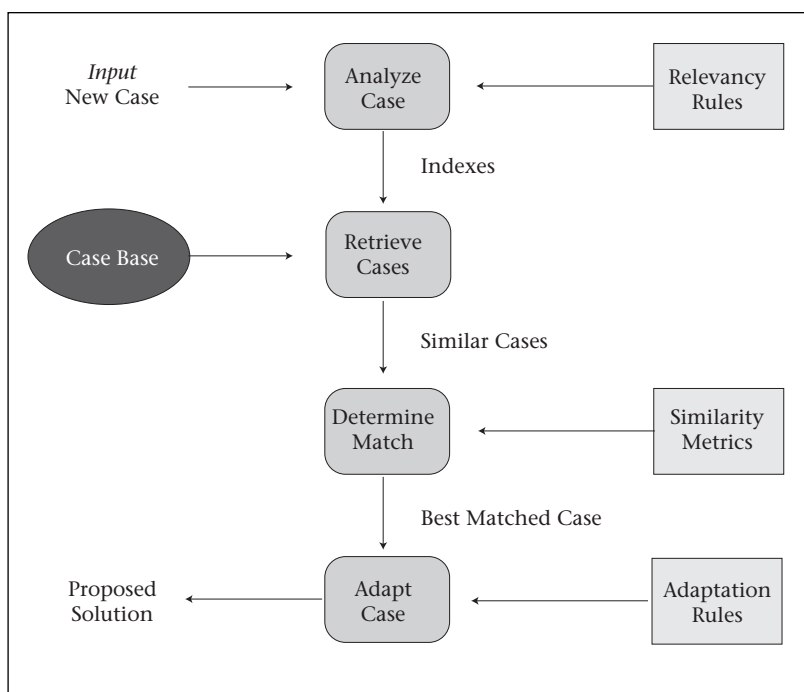
Figure 4. Problem-Solving Strategy of Case-Based Reasoning in CHEMREG (adapted and abridged from Riesbeck [1989]).

ards could allow possibly hazardous classifications to be missed, thereby contravening the goal of 100-percent compliance. Thus, both knowledge bases incur the same issue: How can proper communication of public health and safety information be ensured while the company allows for rapid and cost-effective product evaluation?

All chemical companies face this issue. We initially addressed the problem in CHEMREG by requiring estimates to be entered when actual test data were unavailable; however, this approach proved unsatisfactory. The challenge became one of finding a means to alleviate the manual entry of estimates while ensuring that reliable and auditable data would be available to the transportation and MSDS knowledge bases.

The CHEMREG Solution

To address this challenge, CHEMREG uses case-based reasoning (CBR). This section describes how (1) CHEMREG implements CBR; (2) the users of CHEMREG, the Air Products product safety experts, work with the system; and (3) the case-based reasoner is physically implemented in a knowledge base.

Automated Case-Based Reasoning

If products are allowed to assume the role of cases, this challenge becomes a classic CBR problem. CBR is a problem-solving strategy based on finding a known, similar case on which to base the solution of a new problem (Leake 1996; Kolodner 1993; Reisbeck and Schank 1989). The basic process behind CBR was not new to the product safety experts on the project. When MSDSs were manually written, a similar product's MSDS was often selected as the model on which to base a new MSDS. Indeed, the Environmental Protection Agency (EPA) endorses the use of similar products in the absence of actual test data. The EPA (1986) has stated that "the basic assumption in the recommended approach is that risk assessments on chemical mixtures are best conducted using toxicologic data on the mixture of concern or a reasonably similar mixture" (p. 34,020).

However, the manual selection of a similar MSDS cannot be exercised rigorously in a reasonably acceptable amount of time. Moreover, regulatory agencies provide no guidelines for determining similarity, preferring to leave such issues to "case-by-case" judgment (EPA 1986, p. 34,023). CHEMREG automates the similar-product approach for the first time in the chemical industry.

Figure 4 presents the flow of case-based reasoning tasks used in CHEMREG. CHEMREG must perform these steps without human intervention. The following tasks are executed:

The *analyze case task,* or situation assessment (Kolodner 1993, pp. 285–286), reads physical and hazard data along with compositional information for a new product from the database and determines what missing hazard data are needed (that is, are relevant) for classifying the product. Rules determining relevancy identify the indexes for retrieving products from the database. Indexes consist of the type of hazard data being sought and, if appropriate for the hazard, the principally hazardous component in the current product that causes the hazard to be relevant.

The *retrieve cases task* creates the appropriate SQL statements for retrieving candidate similar products and their components from a relational database. Indexes identified in step 1 are translated into SQL where clauses. Retrieval is based on whether the candidate product has actual data for the relevant hazard (all hazard data are flagged in CHEMREG about whether they are actual or estimated) and whether the product contains the principally hazardous component in the current product. Retrieved cases are parsed into appropriate knowledge base classes for similar-product data and compositional information.

The *determine match task* applies similarity metrics based on the way Air Products' safety
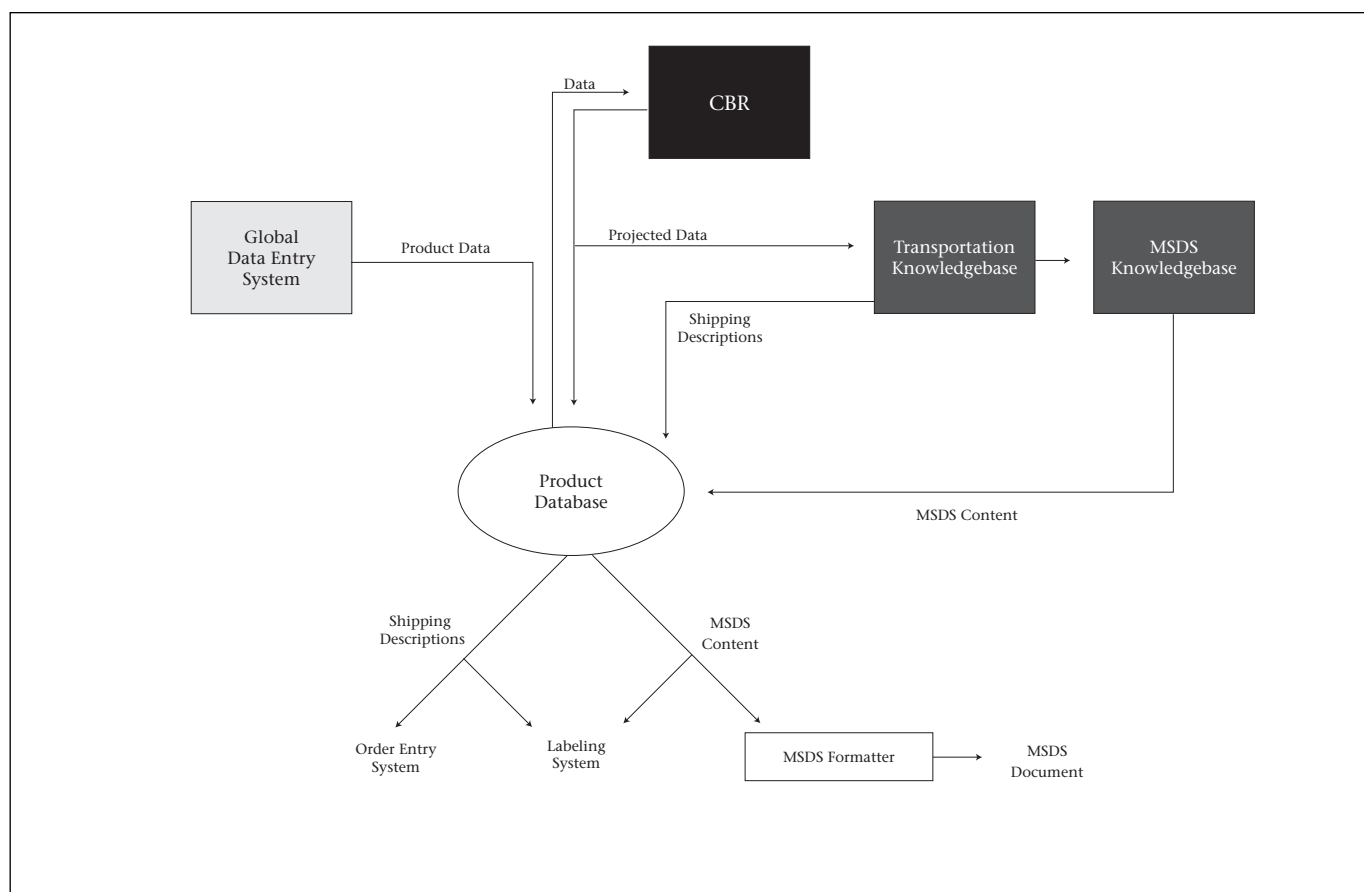
Figure 5. Configuration of CHEMREG *with the Material Safety Data Sheet Knowledge Base and the Case-Based Reasoner.*

experts would choose the most closely matching similar product. This task is the most complex of the case-based reasoner's tasks. First, rather than matching on specific product properties and aggregating measurements of the distances between their values, CHEMREG must measure structural and performance similarity. Products are primarily defined by their *structure,* that is, by their components. Although the compositional level of a particular component can be considered a specific feature, each component might or might not be matched in another product. Other metrics measure performance similarity by quantifying the aggregate effect and interaction of the components in a mixture. Thus, performance must be inferred from compositional structure. Second, the metric for measuring similarity under one type of hazard could differ from the metric for measuring similarity with respect to other hazards. Therefore, the CHEMREG case-based reasoner needs to accommodate context-sensitive matching (Kolodner 1993, pp. 328–329). If a product has multiple relevant hazards, the case-based reasoner will make sev-

eral passes over the product, each time considering the product in a different context, perhaps selecting different similar products to meet the needs of the different hazards. Finally, it was necessary to distinguish two types of similarity metric: (1) *objective metrics,* which make a quantitative measurement of similarity, and (2) *selection heuristics.* These types of metric are discussed later.

The *adapt case task* adapts the hazard data found on the selected similar product to the current product when a significant difference in the similarity measurement between the current product and the selected similar product still exists. Adaptation rules always occur in pairs—one rule defines the procedure for increasing the severity of the hazard data on the selected similar product, the other for diminishing the severity of the data under the reverse direction in the difference between the two products. The CBR technique used is parameter adjustment (Kolodner 1993, pp. 404–407).

Case-based reasoning integrated easily into the CHEMREG system. Figure 5 shows the current
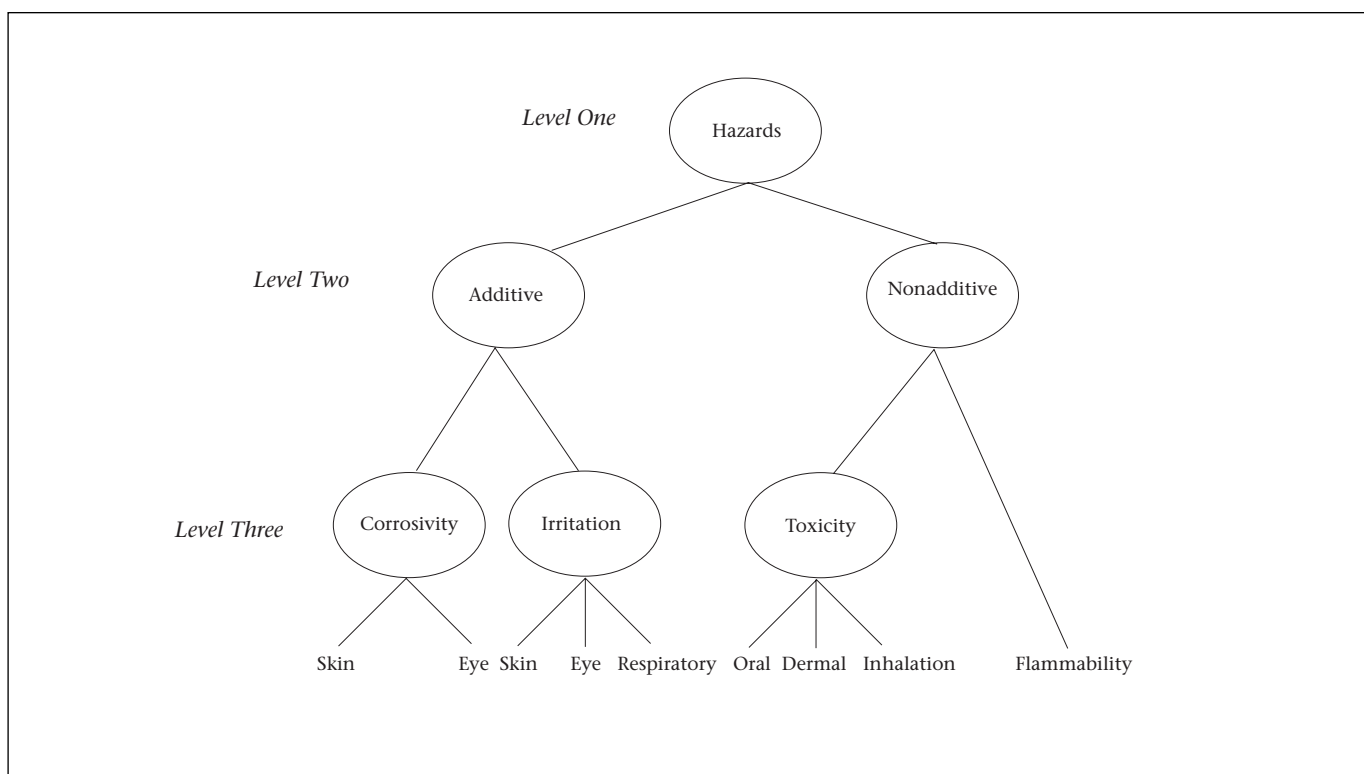
Figure 6. Class Hierarchy of Hazards.
Leaf nodes are classes that contain slots for holding test results from a similar product.

configuration of CHEMREG. The case-based reasoner acts as a front end to the other knowledge bases and generates the data that the transportation and MSDS knowledge bases use to make classifications (called *projected data*).

Operation of the Case-Based Reasoner

The CBR process just described can apply only to products that are mixtures. Product structure must be understood as compositional structure rather than chemical structure. It is a problem far beyond the scope of CHEMREG to project hazard characteristics based on similarity in the chemical structure of two molecules. Often, the addition of a single functional group can substantially alter the characteristics of a chemical substance in unpredictable ways. Our product line consists of families of mixtures whose individual products are sufficiently similar in compositional structure to make this reasoning process highly beneficial. It became clear early during testing that the case-based reasoner was able, in a significantly reduced time, to do a more thorough and accurate job of selecting similar products than could the experts themselves.

The case-based reasoner makes estimates consistently, and it tracks what product and

reasoning were used as the basis of these estimates. If the case-based reasoner fails to find a similar product for the new product, the product is kicked out to the product safety experts with a recommendation that it be tested. This strategy stands in contrast to prior knowledge-based approaches for authoring MSDSs that have used the technique of projecting product hazard classifications from component hazard classifications (for example, Sorani [1991]). This latter approach assumes that any synergistic interactions between hazardous components either can be ignored or will always be known ahead of time. CHEMREG aids product safety experts in identifying those products on which testing should be performed when there is no current model on which to base an estimation. Thus, the CHEMREG strategy represents a more conservative approach toward Responsible Care.

Nevertheless, it is possible for the product safety experts to override the case-based reasoner by specifying an estimate for a test value on a new product as input to the CBR process. When the experts enter an estimate, the case-based reasoner will still perform the complete similar-product logic and report any discrepancy between its result and the experts' estimate. However, it will project the entered esti-

mate for the new product. The need to enter estimates counts as a failure of the case-based reasoner; see the Implementation, Performance, and Maintenance section.

Architecture of the Case-Based Reasoner

In this section, I describe the architecture for implementing the case-based reasoner using traditional knowledge-based system technology (rules and objects).

The CHEMREG case-based reasoner involves three layers of inferencing: (1) the inference engine itself, which is provided by the development software and allows for backward-chaining and pattern-matching forward-chaining rules; (2) a task-control layer representing the problem-solving strategy of CBR (figure 4), which sequences the execution of tasks; and (3) the task-specific rules. *Task-specific rules* determine the relevancy of hazard data for classifying a product (relevancy rules), measure the similarity of products (similarity metrics), and adapt the selected similar-product data to the current product (adaptation rules).

The following points summarize the physical design of the knowledge base:

First, the task-control layer consists of a functional decomposition of the steps involved in CBR. Execution is controlled by procedural code, which invokes control rules that send messages to knowledge base objects to initiate specific problem-solving tasks.

Second, hazards are represented as classes in a hierarchy (figure 6).

Third, relevancy rules are implemented as backward-chaining rules whose root goals are Boolean slots of hazard classes indicating whether the hazard is relevant.

Fourth, cases, which consist of products and their components, are represented as instances of classes. Candidate similar products have a slot for accumulating their similarity scores.

Fifth, all similarity metric rules are expressed by powerful forward-chaining rules (figure 7). These rules join the knowledge base classes constituting a case. Thus, a single rule can process a particular metric involving all candidate similar products.

Sixth, similarity scores are calculated by a simplified form of nearest-neighbor weighting:

$$score = weight \times sim(f_I, f_R) \text{ ,}$$

where *weight* is the importance factor of feature *f* for the hazard under consideration, and *sim* is the objective measurement of the difference between this feature in the input case and the retrieved case, respectively. Scores for individual features are summed. (Compare Kolodner [1993, p. 355].)

```
IFMATCH Current_Comp, Similar_Comp,
        Similar_Products
   with Current_Comp.comp_id =
                Similar_Comp.comp_id
   and Similar_Comp.prodcode =
                Similar_Products.prodcode
THEN
   send (CalcCompDifference to Similar_Products
        with Current_Comp.comp_percentage,
                Similar_Comp.comp_percentage)
END
```

Figure 7. General Similarity Metric.

CalcCompDifference calculates a similarity score based on the compositional level of a particular component in the current product versus its level in a similar product.

Seventh, context-sensitive matching uses metric rules that reference level 2 of the hazard-class hierarchy (figure 8). Dynamically created knowledge base objects track similarity scores of the same product within different contexts.

Eighth, additional selection heuristics refine the selection of the highest-scoring similar products. Selection heuristics are constructed on the same pattern as similarity metric rules (figure 7) but might require matching on levels in the hazard-class hierarchy below level 2 and might reject a candidate similar product outright.

Ninth, adaptation rules are implemented in functions that apply the relevant data from the selected similar product to the current product. As more adaptation rules are identified, it becomes increasingly desirable to give these rules explicit representation as backward-chaining rules rather than embed them in procedural functions.

Figure 8 shows an example of the reasoning that is followed by the case-based reasoner when considering a case of oral toxicity (the lethal effects of ingestion).

Development Issues

This section deals with two issues encountered in developing the case-based reasoner: (1) using a relational database as the case library and (2) eliciting task-specific rules.

```
        The Knowledge Base executed the following activities when
        processing Product NEWPROD  Air Products' New Product

        Oral Tox     Oral toxicity data is relevant for
                     product NEWPROD because there is a
                     component that is classified as an oral
                     toxin.  The most severely hazardous
                     component for oral toxicity is
                     Pentamethyldiethylenetriamine at 95.00%.

        Oral Tox     For oral toxicity the following products
                     were evaluated: S_PROD01, S_PROD02,
                     S_PROD03, S_PROD04.

        Gen Comp     Product S_PROD01 is being evaluated for
                     oral toxicity from a base score of
                     216.12.

        Gen Comp     Product S_PROD02 is being evaluated for oral
                     toxicity from a base score of 195.00.

        Gen Comp     Product S_PROD03 is being evaluated for oral
                     toxicity from a base score of 163.50.

        Gen Comp     Product S_PROD04 is being evaluated for oral
                     toxicity from a base score of 120.80.

        Max Comp     Product S_PROD02 was incremented to 585.00 for
                     oral toxicity because it contains
                     Pentamethyldiethylenetriamine at 100.00%
                     vs current prod concentration of 95.00.

        Max Comp     Product S_PROD03 was incremented to 490.50
                     for oral toxicity because it contains
                     Pentamethyldiethylenetriamine at 58.50%
                     vs current prod concentration of 95.00.

        Rej NonAdd   Product S_PROD03 was rejected for oral
                     toxicity because it contains
                     N,N-DIMETHYLCYCLOHEXYLAMINE at 41.50%,
                     which is a more severely hazardous
                     component.

        Rej NonAdd   Product S_PROD04 was rejected for oral
                     toxicity because it contains
                     N,N-DIMETHYLCYCLOHEXYLAMINE at 84.20%,
                     which is a more severely hazardous
                     component.

        Comp Rej     Product S_PROD01 was rejected because it
                     contains Pentamethyldiethylenetriamine at
                     11.25% which is below the minimum
                     concentration for this component of
                     48.33%.

        Oral Tox     Product S_PROD02 was selected as the similar
                     product for oral toxicity for product
                     NEWPROD.

        Oral Tox     Using similar product data from product
                     S_PROD02 for product NEWPROD.
```

Figure 8. Case-Based Reasoning Report.

Gen Comp messages result from the general similarity metric (figure 7). Max Comp messages represent the application of the specific similarity metric rule for toxicity. Rej NonAdd and Comp Rej messages result from selection heuristics.

Relational Case Library

Technical challenges arose with the use of SQL to retrieve cases (products) from a relational database. Previous uses of relational database technology to store cases have used a simple flat record structure in which cases are represented as *n*-ary relations (Kitano and Shimazu 1996, p. 254). However, this simple case structure presents a problem because of the normalized structure of most corporate databases. Although existing databases provide ready-made case libraries, normalization typically leads to a complex case structure. In CHEMREG, for example, products and their components exist in an *m:n* relationship. Case retrieval must span all the tables in which data relevant to the case are found. To address case complexity, CHEMREG requires three passes over the database: (1) to set up (in the database) a list of products to be retrieved based on applying the relevant indexes (a step involving a three-table join), (2) to retrieve products and their components (a step involving a six-table join), and (3) to retrieve pure products of the principally hazardous component regardless of chemical family.

Moreover, CHEMREG does not use similarity-based retrieval mechanisms (for example, fuzzy matching or nearest-neighbor retrieval) when retrieving product information from the database. Retrieval is based on the candidate product being explicitly flagged as having actual data for the relevant hazard and containing exactly the same component as the principally hazardous component in the new product. All similarity measurement takes place within the knowledge base itself. This strategy allows for easily generating the reasoning report (figure 6) from a single source.

In this situation, however, there is a performance trade-off to keep to a minimum the number of case records coming into the knowledge base across the network but not exclude potentially relevant cases. In CHEMREG, this trade-off is resolved by retrieving only those products within the same chemical family as the new product. A specific chemical family is used for those product lines in which we have a significant number of products; a general family can be used to expand the search in more specialized product lines. It is the experts' responsibility to establish the appropriate chemical family on which to define the scope of the search given the chemical family of the new product.

On average, eight similar products are retrieved for each product; these might involve 60 or more database records being retrieved depending on the number of components making up these products. The case-based reasoner takes about 15 seconds to process a product, with approximately one-third this time devoted to performing physical input-output functions.

Knowledge Engineering

Four types of task-specific rule had to be elicited from the experts:

First are relevancy rules. *Relevancy rules* identify when a hazard is relevant for classifying a product according to the regulations. These

were the easiest of the task-specific rules for the experts to articulate. Each type of hazard has 5 to 10 relevancy rules. The most common form of a relevancy rule is to identify the concentration level at which a hazardous component causes a hazard to be relevant for classifying the product.

Second are objective metrics. *Objective metrics* quantitatively measure the difference between structural and performance features of the new product and a similar product. These rules, which calculate a numeric score for the similarity between structural and performance features of two products, are either *general* and apply the same metric to all cases (figure 7) or *specific* and apply different metrics depending on the hazard under consideration (context-sensitive matching). As a rule, specific objective metrics for additive hazards make an assessment of a product's performance characteristics, but metrics for nonadditive hazards must be more sensitive to compositional structure.

Third are selection heuristics. *Selection heuristics* ensure that the highest-scoring similar product meets conditions of acceptability. They express the experts' judgment for selecting similar products in anomalous situations. The need for this layer arose during system testing when the experts would reject the results of the objective metric calculations based on their interpretation of, or knowledge about, the data. The experts could articulate these heuristics only when confronted with an anomalous situation. At present, seven selection heuristics have been identified.

Fourth are adaptation rules. *Adaptation rules,* which are also context sensitive, define a reasonable way of estimating scores based on similar-product data when there is still a significant difference between the current product and the selected similar product. Adaptation rules, therefore, have to recognize when a significant gap occurs in the measurements of a structural or performance feature and define a procedure by which the severity of the hazard data from the selected similar product is to be increased or diminished. Both these points are matters of professional judgment; thus, the adaptation rules in CHEMREG are heuristic. Adaptation rules proved to be the most difficult for the experts to articulate. Only one adaptation rule for toxicity mapped to previous manual procedures for handling regulatory compliance for new products. The hazards that are amenable to adaptation (corrosivity, irritation, and toxicity) have an average of four (two pairs) adaptation rules.

The CBR process in CHEMREG does not use a feedback, or learning, cycle. Projected estimates are not themselves used to project further estimates because the need to adapt data from the product on which the original, actual test was performed might be missed if a more closely matching similar product on which these data have not been adapted is identified. The gap between the new product and the similar product with the projected data might be insufficient for the system to recognize the need for adaptation. Based on the originally tested product, however, the data might need to be adapted to the new product.

Implementation, Performance, and Maintenance

Excluding development of the relevancy rules, which were implemented independently of the similar-product logic in an initial version of the knowledge base, eliciting the similarity metrics, selection heuristics, and adaptation rules took approximately 40 hours of sessions with the experts over a period of 5 months. This time covered a substantial period of reviewing tests and refining the selection heuristics and adaptation rules.

Air Products fully implemented the case-based reasoner in June 1996. We have tracked the performance of the knowledge base for acute toxicity (dermal, oral, or inhalation), skin corrosivity, and skin and eye irritation. Table 1 summarizes the performance criteria and actual results obtained during testing and the first four months of production. We also tracked the times when the case-based reasoner proposed estimates that, on review, were preferred to the experts' own original estimates (column 5). Actual results equal the times the experts agreed with the system plus the times that CHEMREG was judged to perform better than the experts. The chart shows that in cases involving toxicity, the case-based reasoner has exceeded expectations for its performance. The areas of skin corrosivity and eye irritation require further refinement.

Maintenance requirements differ for the three knowledge bases. Maintenance on the transportation knowledge base must be performed periodically as the transportation regulatory agencies promulgate revised regulations, which occurs about every two years. The MSDS knowledge base, however, was designed as a formal system for authoring the MSDS. Consequently, the product safety experts can play a more direct role in controlling the output of the knowledge base to meet their particular needs for changing the contents of an MSDS. Currently, we must still monitor the

Type of Estimates	Performance Goal (%)	Actual Results (%)	Experts Agreed (%)	CHEMREG Judged Better (%)	Experts Disagreed
Estimates for acute toxicity	90	96	34	62	4
Estimates for skin corrosivity	95	93	50	43	7
Estimates for skin irritation	95	97	59	38	3
Estimates for eye irritation	95	82	46	36	18

Table 1. Performance Results of the Case-Based Reasoner.

performance of the case-based reasoner and consider refinements to its reasoning, particularly to its selection heuristics and adaptation rules.

Benefits and Future Enhancements

Regulatory compliance is a public and legal obligation; it is a cost of doing business. The business strategy must be to find the most cost-effective means for ensuring full regulatory compliance and communication of the hazards of the product. CHEMREG guarantees that shipping descriptions and MSDSs will comply with corporate policy and interpretation of the regulations. Moreover, periodic regulatory changes can quickly and consistently be implemented across the entire product line. The result is an assurance of superior hazard communication to the public and a reduction for the company in the potential for significant loss as a result of noncompliance.

Air Products would have incurred significant costs and effort under the manual system to bring shipping descriptions and MSDSs into compliance with the current regulatory environment. Developing the transportation knowledge base avoided an estimated $250,000 expenditure involving a substantial effort on the part of the product business managers to review all their products (a highly time-consuming activity to impose on product management) to comply with the original DOT regulatory revisions. Developing the MSDS knowledge base and case-based reasoner avoided a 1.5 to 2 person-year effort to implement the ANSI standard format for all MSDSs and ensure consistency throughout all documents. In addition, the case-based reasoner has relieved

an estimated $1 million one-time cost of selectively testing all products to develop complete health and safety information on each product in the whole product line.

Although development expenditures have nearly offset these initial savings, ongoing benefits for CHEMREG can be estimated in terms of costs that would be incurred under the manual system to achieve a comparable level of quality and consistency. In other words, what would be the cost of returning to the manual system and expecting the same results? Air Products estimates that it would take at least one to two person-days of additional work on each new product to achieve results comparable to CHEMREG's. In addition, there is an expected savings of $50,000 to $70,000 in annual expenditures for testing new products. Thus, CHEMREG will accomplish $1 million worth of otherwise additional work and avoided expenditures over 5 years.

Among planned future enhancements to the case-based reasoner is the introduction of variable relevancy rules for determining when a hazardous component in a mixture is relevant. Currently, relevancy rules stipulate a minimum concentration level as the point at which, in the experts' judgment, a hazard of a particular type and severity becomes relevant. Variable relevancy rules can depend either on particular substances (for example, phenol can be assigned a lower relevancy point than other toxic substances) or on results in the case library that indicate when a hazard actually causes a hazard to occur in a mixture.[1] The project team favors the approach of basing variability on particular substances because it is simpler to implement, and we can use internationally recognized minimum concentrations established by European regulations for specific substances.

An additional enhancement could be to add the ability to handle complex compositional structures. Currently, the case-based reasoner can compare the structure of products only on the basis of the fundamental chemical substances they contain. However, products often consist of mixtures of chemicals; for example, Air Products might purchase another chemical company's specialty mixture and add it as one component in one of its mixtures. The goal would be for the case-based reasoner to consider multilevel compositional structures in both the relevancy rules and similarity metrics. The need for this enhancement and the effect on the similarity metrics are under discussion.

The long-term project plan calls for CHEM-REG to become a global hazard communication system. In 1996, European road transport regulations (ADR) were added to the transportation knowledge base, and during 1997, the MSDS knowledge base was extended to classify products according to the European Economic Union classificatory scheme for including risk and safety phrases on an MSDS. Expansion to Latin America and the Pacific Rim is anticipated.

Summary

CHEMREG is a large knowledge-based system used by Air Products and Chemicals, Inc., to support compliance with regulatory requirements for communicating health and safety information in the shipping and handling of chemical products. This article concentrated on one of the knowledge bases in this system: the case-based reasoner. Previously written knowledge bases for generating shipping descriptions and material safety data sheets mapped classificatory rules directly into knowledge base code. To support regulatory compliance, however, this strategy required estimates to be entered for all hazard-related properties on every product. To alleviate this burden, an automated case-based reasoner was written that generates estimates of hazard data from similar products using an existing product database as its case library. Although some refinements remain, the performance of the case-based reasoner has met its design goals.

Acknowledgments

Note

1. The first alternative has been adopted by the European Union in its regulations for applying risk phrases to preparations. The second alternative was suggested by an unnamed questioner at the Innovative Applications of Artificial Intelligence conference, to whom I apologize for not giving his question its due consideration at the time.

References

EPA. 1986. Environmental Protection Agency: Guidelines for the Health Risk Assessment of Chemical Mixtures 1986. *Federal Register* 51(185): 34014–34025.

Kitano, H., and Shimazu, H. 1996. The Experience-Sharing Architecture: A Case Study in Corporate-Wide Case-Based Software Quality Control. In *Cased-Based Reasoning: Experience, Lessons, and Future Directions,* ed. D. B. Leake, 235–268. Menlo Park, Calif.: AAAI Press.

Kolodner, J. 1993. *Case-Based Reasoning.* San Francisco, Calif.: Morgan Kaufmann.

Leake, D. B. 1996. CBR in Context: The Present and Future. In *Cased-Based Reasoning: Experience, Lessons, and Future Directions,* ed. D. B. Leake, 3–30. Menlo Park, Calif.: AAAI Press.

Riesbeck, C. K., and Shank, R. C. 1989. *Inside Case-Based Reasoning.* Hillsdale, N.J.: Lawrence Erlbaum.

Sorani, G., and Lauer, R. A. 1991. The Application of Artificial Intelligence in the Field of Chemical Regulation. In *Innovative Applications of Artificial Intelligence 3,* eds. R. Smith and C. Scott, 55–64. Menlo Park, Calif.: AAAI Press.
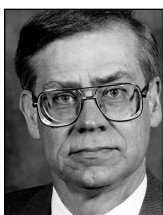
Kirk D. Wilson served as a knowledge systems analyst at Air Products and Chemicals for 12 years. He is currently employed by Platinum Technology, Inc., where he is a manager of advanced technology on the development staff of the next version of AION. He holds a Ph.D. in philosophy from the University of Massachusetts. His interests include knowledge modeling, the philosophical foundations of AI, and the synergy between object-oriented analysis and knowledge engineering. His e-mail address is kirk.wilson@platinum.com.