Learning Lessons Intelligently in the Electric Power Industry

A. Martin Wildberger and Madeleine M. Gross

EPRI (Electric Power Research Institute)
3412 Hillview Avenue
Palo Alto, CA 94304-1395
mwildber@epri.com, mgross@epri.com

Thomas J. Ayres

Exponent Failure Analysis Associates
149 Commonwealth Drive
Menlo Park, CA 94025
tayres@exponent.com

Abstract

EPRI (formerly the Electric Power Research Institute), on behalf of its member companies worldwide, has initiated a strategic research program in Human Performance (HP), designed to enhance HP management by development of approaches and tools for anticipating, monitoring, and addressing factors with adverse impact on human productivity and reliability in workplace settings. This program emphasizes proactive intervention strategies, featuring intelligent, real-time semi-automated and automated tools to identify early indicator patterns of worker, workplace, and organizational factors associated with HP problems. One major programmatic thrust is directed toward developing approaches for making optimal use of the vast amount of text and quantitative data, both event and operational, routinely recorded by EPRI's member companies (electric utilities and other energy industry settings), as well as the informal lessons learned (undocumented knowledge) of experienced workers. . A combination of artificial intelligence (AI) and computer simulation may be used to implement a system that would automate the assembly, organization, and evaluation of such data, answer user queries, and even suggest unanticipated relationships. This suggested integrated approach could use query-driven simulation, combining two common approaches to artificial intelligence: model-based and case-based reasoning.

Introduction

Traditional lessons-learned efforts are directed towards keeping track of findings or conclusions to which people have arrived based on operating experience, usually from one incident at a time. A more proactive or *intelligent* approach to learning lessons from past experience includes detecting, organizing, and evaluating both recorded and undocumented data, information, and knowledge that have reasonable implications for improving future performance of humans, machines, and human-machine systems. Then those lessons can be put into practice, and feedback established, to ensure continuous future learning by the capture and analysis of what past and present experience shows to be the most useful new information.

One major thrust of EPRI's strategic (relatively longer-term exploratory) research in human performance is directed toward the design and implementation of what constitutes an intelligent lessons learned system that would, in part, make optimal use of the vast amount of data routinely recorded by EPRI member companies (largely electric power industry settings). Such data are stored in event logs, incident and accident reports, sensor data files, maintenance requests and other task-specific documents, as well as in background operational information such as personnel and payroll records, plant operation statistics, safety programs, and managerial directives. Unfortunately, some of it only resides in the memories of experienced workers. Techniques of intelligent automated search, summarization, and information extraction (as well as storage and presentation) are relevant to exploring lessons learned from all these sources. Through the identification and use of such methods, knowledge that is otherwise undocumented (in the sense that implications cannot otherwise readily be drawn, or in that the information is maintained informally) can conceptually be transformed into documented knowledge and lessons learned.

This paper describes how a combination of artificial intelligence (AI) and computer simulation could be used not only to automate the assembly, organization, and evaluation of these data, but also to implement a system to answer user queries and even suggest unanticipated relationships. The proposed conceptual integrated approach discussed here uses query-driven simulation, combining two common approaches to artificial intelligence: model-based and case-based reasoning. Continuous learning by the system itself may be possible using evolutionary algorithms operating on user feedback (Wildberger 1999).

Necessary preliminary work is now underway to explore human performance models, data-gathering approaches, database structures, and analytic techniques that may be applicable to the overall EPRI HP program (Gross and Ayres 1998). Existing and in-development analogous sets of data are being examined, as well as techniques for identifying, classifying, and coding event (and other) data. Current studies involve looking for the best ways to handle the capture and validation of undocumented knowledge, and also investigate the use of automated text-analysis software to facilitate the recovery of data not originally stored for such information analysis and knowledge extraction purposes as those explored here.

Current Work and Progress

Largely because of the history of the perceived significance of safety-related and other human factors issues associated with nuclear power generation, most of the past HP research and development by EPRI (which emphasizes energy industry settings) has been associated with that domain. However, widespread blackouts, also an electric power industry-related concern, can be at least as dangerous as minor radiation leaks, and can have a far greater economic impact. Deregulation and competition in the electric power industry have increased the demand for productivity and efficiency, and have resulted in a narrowing of the security margin in power systems. Safety, public benefit, and economic considerations are all driving forces for the more recent broader scope of EPRI HP research and development.

Current EPRI work in improving human performance, including both base program and strategic components, comprises relatively near-term objectives, largely for nuclear power plants, as well as some of the essential preliminary analysis and development needed to support the longer-term strategic research which is emphasized in the present paper. For example, one relatively short-term effort in the nuclear generation area is directed toward development of a set of Leading Indicators – measures that provide advance information about the overall system. That project (EPRI 1999) involves developing and evaluating indicators based on a review of organizational models of human performance. This approach, using an aggregate model to suggest measures of interest, is related to work that has been undertaken in other industries (Reason 1997). Another near-term nuclear generation effort deals with the development of a software tool to facilitate linking the cause of a problem to an appropriate corrective action (see Gross and Ayres 1998).

There are three major areas in which research and development are taking place in support of longer-term, strategic goals. They are described in the next three subsections. Although work in these areas is specifically directed toward supporting strategic plans, some of the early and mid-term analyses and results are expected to have immediate or near-term application to EPRI base program-specific (e.g., nuclear) programs.

Capturing Undocumented Worker-Job-Knowledge. Any organization is subject to the loss of valuable undocumented knowledge and expertise as senior workers and managers become unavailable, and this loss can have significant negative productivity, economic, and safety consequences. The project objective is to select or build the most effective methods and tools to:

- Determine whether valuable undocumented knowledge will actually be lost when senior, experienced personnel become unavailable (e.g., by retirement or during vacation);
- Decide whether such worker knowledge is worth capturing;
- Capture and store such knowledge;
- Retrieve and present this knowledge when needed.

These considerations, of course, are also relevant to the general issue of Knowledge Management in an organization.

Autotools and Autotext. The Automated Human Performance Analysis Tools project involves exploring concepts for intelligent semi-automated and automated tools to support human performance management. The Automated Text Analysis sub-project focuses on evaluating the use of software tools for extracting information from industry documents, such as text reports of incidents (Murray, Gross and Ayres 1999), in order to make use of the large quantity of narratives and other non-quantitative data collected at utilities.

Human Performance Management - Database and Analysis. This is the central project in the current EPRI strategic HP program. This project has included a review of research on the inferred causes and correlates of human performance problems at electric utilities and in analogous settings, and of research on analytic techniques currently being applied (see Gross, Ayres and Murray 2000). It is conceptually related to, and integrated with, the nuclear power plant Leading Indicators project. A database structure is being designed that will support the combination of event, mishap, and site performance data with organizational and other information, in a way that makes it easier to identify (through innovative statistical analytic and data mining techniques) the precursors (especially organizational factors) associated with HP problems.

One long-range objective of all aspects of the program is to explore potentially effective ways to link the current data gathering and analytic efforts with innovative intelligent processing. The remainder of this paper discusses possible directions that are under consideration.

Selection of Artificial Intelligence Methods to Assist the Lesson Learning Process

In general, human workers who have discovered an improved way to accomplish a task or a potential problem to be avoided have to record, store, and disseminate that

information before the lessons they have learned from personal experience can become lessons learned by others from vicarious experience. Presently, computer assistance in recording, storing, and disseminating this information is rudimentary. It consists mainly of word processing, database management, and publishing tools. Little assistance is provided for recognizing a potentially useful lesson at either end of the process. Both the input of personal experience and the application of the vicarious experience depend heavily on the ability of individual workers to recognize its importance, generality, and relevance, as well as their initiative in reporting new lessons or searching for recorded ones (SELLS 1999).

On the other hand, many records are routinely generated concerning work in progress or completed in process industries, and most tasks have written procedures describing how to accomplish them. These records and procedures are frequently computerized to some degree, and interfaced with automatic inventory control, equipment tag-out, and other ancillary tasks. It appears to be both possible and useful to employ some AI methods in the context of a total integrated system design that would make learning lessons less dependent on individual human initiative and prior knowledge.

Although quantitative data (for instance, instrument readings and inventory counts) are important to operational and maintenance activities, most of the knowledge about how to perform these tasks is generally contained in qualitative, natural language documents and notes. Therefore, the first place to look for applicable AI techniques is in the family of symbolic AI methods, rather than the methods of computational intelligence, such as neural networks and evolutionary algorithms, although these may also have some use in a total system design (Wildberger 1991).

The methods of symbolic artificial intelligence distinguish between shallow (procedural, historical) knowledge and deep (causal, structural) knowledge (Lenat and Guha 1990). The former is typical of the formal operations and maintenance procedures used to ensure safety, consistency, and completeness in most industrial settings. The latter requires modeling the task rather than modeling how the person doing that task actually accomplishes it. Such models are characteristic of those used for optimization and control of a power system or in a simulator for training power plant operators.

Conventional Expert Systems. The most common approach to building an "Expert System" is to elicit procedural steps from an expert human and represent those steps as a set of logically interrelated "if ... then" statements (Nilsson 1998). This kind of "rule-based" expert system would seem a natural one for lessons learned, since most lessons are couched in terms of "if the following situation pertains, then perform the following actions and/or avoid performing these other actions." However, this approach does nothing to help workers specify new rules (i.e., record a lesson they have learned) or to help them find an appropriate rule to use (i.e.: learn from the recorded lesson).

Model-Based Reasoning (MBR). Model-Based Reasoning requires knowledge engineers to build a causal model representative of the underlying knowledge domain. This is essentially the same modeling task required prior to designing a computer simulation. However, rather than using precise quantitative formulations, MBR attempts to characterize the problem domain by an underlying physical model expressible in a relatively small number of general rules (Dvorak and Kuipers 1989, Bau and Brézillon 1992). The derivation of the rules is guided by the goals of the users and the known facts about the current situation being investigated. Because of the uncertainties inherent in the process of developing such an approximate physical model, as well as the need for aggregation and categorization required to reduce it to a practical size, fuzzy sets and fuzzy logic have been found to be particularly useful for formulating MBR systems (Singer, Gross and Humenik 1991). A typical MBR expert system generates guidance for its users by simulating the effects of actions the user might take in the current situation, and recommending those that are most effective and efficient within all known constraints.

Case-Based Reasoning (CBR). One convenient approach to building a procedural expert system is to use Case-Based Reasoning (CBR) (Slade 1991, Watson 1997, Aha, Breslow and Muñoz-Avila 2000). The knowledge engineer begins by collecting a library of historical and/or idealized cases that exemplify typical situations and their solutions within the problem domain. Then, each new situation is matched, with the user's help, against this library in order to find the best example for the user to follow. The success of such an expert system depends on ensuring that the cases represent useful stereotypes, cover the entire problem domain, and can be retrieved easily from sensor data and/or user input that only partially defines a complete case. The more cases stored, the greater is the confidence that an approximate match-up is meaningful. The theory of rough sets has been found to be very useful in formalizing the knowledge discovery required for successful CBR systems (Yurtsever, 1994, Yokomori and Kobayashi 1995).

Query Driven Simulation (QDS): Combining MBR and CBR. At a conceptual level, the distinction between MBR and CBR in expert systems is much the same as that between causal and statistical mathematical models in simulations. Model-based expert systems are essentially qualitative simulations of the problem domain combined with an interface that both responds to user questions and makes unsolicited recommendations. Case-based expert systems resemble databases equipped with inductive retrieval tools based on combinatorics and non-parametric statistics.

Query driven simulation (QDS) is an example of a simulation method that combines the essential features of case-based or statistical reasoning with causal or model-based reasoning. QDS, in general, is intended to provide a

user with a knowledge/data base that is capable of answering questions about data not stored explicitly but derivable by running a simulation (possibly real-time constrained) using some of the knowledge/data that is stored (Nair, Miller and Zhang 1996). One example is Active KDL, a comprehensive knowledge/data language and object-oriented knowledge/data base system in which active objects respond to queries by retrieving stored data, by inferring additional conclusions from the stored data using heuristic rules, or by deriving new information from simulations using both stored models and various logical compositions of those models (Kochut, Miller and Potter 1991). In this type of simulation, the stored information may originally be real data obtained as individual cases and a qualitatively formulated causal model may be used both to rationalize the data and to produce simulated cases for situations in which no real cases were recorded.

Using QDS for Lessons Learned

A typical lessons learned system includes the following functions:

- Gather pertinent documents and reports generated locally and elsewhere, with special emphasis on incident reports and beneficial suggestions submitted by local workers.
- Review for applicability and importance.
- Publish formal action documents with action level prescribed. Actions may range from "read and file," through "modify standard procedures," to "perform one time equipment inspection or modification."
- Follow-up to ensure action was taken.

Although emphasis and detailed implementation vary considerably, these functions seem to be essential (SELLS 1999). Within this context, a QDS system is intended to assist learning lessons in two primary areas: reviewing and extracting information from documents and diverse data sources, and incorporating the lessons into routine operations. Adequate review of potential lessons learned is highly dependent on the availability of knowledgeable personnel, who frequently have more pressing demands on their time. In many lessons learned systems, using non-mandatory lessons has been difficult to achieve because such use depends on workers taking the initiative to search a lessons learned file to find lessons pertinent to the operation or maintenance task they are about to perform (SELLS 1999).

Starting with a database of cases, each of which contains a specific lessons learned, QDS has the potential to supplement those cases and to find or construct one directly applicable to the task at hand. The QDS system conceivably could be run in the background while a maintenance request is being filled out on the computer screen or an operational procedure is being assembled for printing. If some stored case (lesson) were already associated with the task, its advice could be included in the maintenance request or operational procedure. If no such case were available, the simulation could construct one, based on cases associated with similar tasks, modified to be consistent with the simulation's internal model of the physical equipment/components/system being maintained or operated at the local facility.

Mathematical models and even actual simulations are available for most industrial process plants, discrete manufacturing facilities, vehicles, transportation systems, telecommunication networks, and military weapon systems. Even when no formal model has already been developed, each such system has its own operational logic, from which a rough model and qualitative simulation can be derived. At the very least, Computer-Aided Design (CAD) drawings could be used to constrain and guide the interpolation of cases (Wildberger and Hickok 1992).

Integrating the lessons learned with all the other computer-based systems that are used to operate and maintain the facility provides an automatic connection between the task to be performed and any lessons learned from its previous performance or from the performance of similar tasks at that facility or elsewhere. The simulation, in effect, produces a generalization of the lessons in the database constrained by the logic of the particular system on which work is being done.

Building the original database or knowledge base of lesson-cases, however, is inherently different from the integration of a lessons-learned system with other computer systems. This building process is a primary purpose of the current HP work described earlier in this paper. Industrial lessons learned programs depend fundamentally on the initiative of staff, both to submit potential lessons and to perform knowledgeable review of these and other documents and reports generated locally and elsewhere. There is no expectation of eliminating human intervention or judgment, and human interest in the system is essential to the success of any lessons learned system. The results of current EPRI HP programs, however, are expected to facilitate the work of the human analyst, as well as to bring together information from many sources to highlight potential lessons. The QDS simulation would then put these in the context of the industrial system and equipment of the local facility. This approach would be anticipated to reduce dependence on worker initiative, and make staff review less tedious and demanding.

Summary

One major thrust of EPRI's strategic research in human performance is directed toward the conceptual design and implementation of what constitutes an intelligent lessons learned system that would, in part, make optimal use of the vast amount of data routinely recorded by its member companies. The system would use a combination of artificial intelligence (AI) and computer simulation not only to automate the assembly, organization, and evaluation of this data, but also to implement a system to answer user queries and even suggest unanticipated

relationships. This Query-Driven Simulation (QDS) would, in effect, produce a generalization of the lessons in the database constrained by the logic of the particular system on which work is being done. Although this is a conceptually ambitious proposal, given the current capabilities of the relevant technologies, consideration of such an advanced system can provide the substrate for development of an intelligent, integrated system of software tools to support HP management in energy industry settings.

References

Aha, D. W.; Breslow, L. A.; and Muñoz-Avila, H. 2000. Conversational case-based reasoning. To appear in *Applied Intelligence*.

Bau, D-Y. and Brézillon, P. J. 1992. Model-based diagnosis of power-station control systems. *IEEE Expert* Vol. 7, No. 1 (Feb.): 36-44.

Dvorak, D. L. and Kuipers, B. J. 1989. Model-based monitoring of dynamic systems. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI '89)*, 1238-1243. San Mateo, CA: Morgan Kaufmann.

EPRI (1999). Guidelines for Leading Indicators of Human Performance: Preliminary Guidance for Use of Workplace and Analytic Indicators of Human Performance. EPRI Report TR – 107315, October, 1999. Palo Alto, CA.

Gross, M. M. and Ayres, T. J. 1998. Research initiative for human performance management of nuclear power. Presented to the IERE Workshop: R&D for Cost Reduction, Fukuoka, Japan.

Gross, M. M.; Ayres, T. J.; and Murray, J. 2000. Analysis of human error at electric utilities. To be presented at the 44th Annual Meeting of the Human Factors and Ergonomics Society, San Diego, CA.

Kochut, K. J.; Miller, J. A.; and Potter, W. D. 1991. Design of a CLOS version of Active KDL: a knowledge/data base system capable of query driven simulation. In *Artificial Intelligence and Simulation*, 139-145. Uttamsingh, R. J. and Wildberger, A. M. Eds. San Diego, CA: Society for Computer Simulation.

Lenat, D. and Guha, R. V. 1990. *Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project*. Reading, MA: Addison Wesley.

Murray, J.; Gross, M.M.; and Ayres, T.J. 1999. Human error in power plants: a search for pattern and context. *Proceedings of the Silicon Valley Ergonomics Conference & Exposition*, 187-191. San Jose, CA: Silicon Valley Ergonomics Institute.

Nair, R. S.; Miller, J. A.; and Zhang, Z. 1996. A Java-based query driven simulation environment, *Proceedings of the 1996 Winter Simulation Conference (WSC'96)*, 786-793. Coronado, California (December).

Nilsson, N. J. 1998. *Artificial Intelligence: A New Synthesis*. San Mateo, CA: Morgan Kaufmann.

Reason, J. 1997. *Managing the Risks of Organizational Accidents*. Brookfield, VT, USA: Ashgate Publishing Company.

SELLS (The Society for Effective Lessons Learned Sharing) 1999. *Proceedings of the SELLS Spring Meeting* (March 15-17, Las Vegas, NV), available at web site: http://www.tis.eh.doe.gov/ll/sells/index.html

Slade, S. 1991. Case-based reasoning: a research paradigm. In *AI Magazine* Vol. 12, No. 1 (spring): 42-55.

Singer, R. M.; Gross, K.C.; and Humenik, K.E. 1991. Pumping systems detection and diagnosis utilizing pattern recognition and fuzzy inference techniques. In *Proceedings of the Fourth International Conference on Industrial and Engineering Applications of Artificial Intelligence & Expert Systems* (Kauai, HI, June 2-5): 51-61. Tullahoma, TN: UTSI.

Watson, I. 1997. *Applying Case-Based Reasoning: Techniques for Enterprise Systems*. San Francisco: Morgan Kaufmann.

Wildberger, A. M. 1991. Comparison of neural network and expert system technology. In *Artificial Intelligence and Simulation*, 152-159. Uttamsingh, R. J. and Wildberger, A. M. eds., Society for Computer Simulation, San Diego, CA.

Wildberger, A. M. 1999. Survey of projects involving evolutionary algorithms sponsored by the Electric Power Research Institute. In *Evolutionary Algorithms*, 277-293. Davis, L. D.; DeJong, K.; Vose, M. D.; and Whitley, L. D. eds. New York, NY: Springer-Verlag.

Wildberger, A. M. and Hickok, K.A. 1992. Power plant modeling and simulation using artificial intelligence and neural networks. In *Progress in Simulation*, Vol. 1, Chap. 4, 100-125. Zobrist, G.W. and Leonard, J. V. eds. Norwood, NJ: Ablex Publishing.

Yokomori, T. and Kobayashi, S. 1995. Inductive learning of regular sets from examples: A rough set approach. In *Soft Computing: Rough Sets, Fuzzy Logic, Neural Networks, Uncertainty Management, Knowledge Discovery*. 69-73. Lin, T.Y. and Wildberger, A.M. eds. San Diego, CA: Society for Computer Simulation.

Yurtsever, T. 1994. Factory floor simulation controlled by fuzzy expert system. In *Proceedings of the European Simulation Symposium, ESS94*. (October 9-12, Istanbul, Turkey), 277-281. Kaylan, A. R.; Lehmann, A.; and Ören, T.I. eds. Ghent, Belgium: Society for Computer Simulation International.