Intelligent Executive Information Systems

David King, Comshare
Daniel O'Leary, University of Southern California

DEVELOPMENTS IN ARTIFICIAL INTELLIGENCE AND THE INTERNET ARE CHANGING THE BASIC NATURE OF EXECUTIVE INFORMATION SYSTEMS. THE AUTHORS OFFER A WAY TO EXTEND THE STRUCTURE OF SUCH SYSTEMS ON THE BASIS OF THESE DEVELOPMENTS.

RECENTLY, THERE HAVE BEEN substantial changes in the technology available to support executive decision making. Perhaps the most visible change has been the rapid development of the Internet and the World Wide Web. In addition, this network infrastructure has facilitated and promoted the growth of enabling technologies, such as databases and artificial intelligence. These changes in technology have led to a change in the basic nature of executive information systems. Accordingly, the purpose of this article is to extend the structure of EISs to account for this wide range of recent technological changes.

As the technologies available for EISs have changed, the location of available resources regarding EIS has also changed. Substantial information related to the discussions generated in this article is available on the World Wide Web. Hence, we searched the Web, and the "Some relevant World Wide Web URLs" sidebar summarizes the findings of that search, listing several addresses of pages related to the topics discussed in this article.

Classic executive information systems

An EIS has been defined as "a computer-based system that enables senior managers to access a common shared source of internal and external information that has been summarized in easy-to-access graphical displays. These displays can be developed and customized to meet the changing information needs of individual executives, but, since the displays are created using information from a common information store, the system helps promote a shared view of the organization."[1]

Traditional EISs initially provided the executive with data stemming almost solely from transaction-based accounting systems. Typical EISs offered *drill-down* capability to facilitate analysis of those databases. In particular, clicking on a number would give the executive more and more detail. For example, clicking on "Sales" might bring up the regional sales numbers that add up to sales. Clicking on a regional sales number could bring up detail by states, cities, or stores in that region.

Generally, in classic EISs, the databases were primarily *hard data*, such as numbers. There was limited access to *qualitative* databases—in large part, because there were few or none to access. Traditional EISs generally let the user set *alarms*. If actual data exceeded planned data by a certain range, then the alarm would go off, indicating the variance condition to the executive.

Later versions of traditional EISs also had some built-in communications capabilities. As a result, some systems included features such as e-mail, to facilitate asynchronous communication with others.

Generally, a traditional EIS could not provide much competitive information. As a result, some firms had departments whose sole job was to try to capture information about the competition and prepare that information for executive consumption. The same information was fed to all on the EIS, to generate a shared view of the organization.

Furthermore, EISs featured information for executives. Of course, those working for an executive who could access such a system would want to be able access it themselves.

Nevertheless, these systems were usually structured so that only relatively few executives could access the information.

Criticisms of this traditional approach are many. Most data was from accounting transactions, which limited the real scope of decisions that reliably could be made using that data. Most data was internal, which made competitive analysis using an EIS virtually impossible. Although there was data, there was difficulty teasing meaning out of that data. The system was aimed at executives; few others, inside or outside the firm, could access the system. Executives took, but did not contribute, knowledge from the system. This was unfortunate, because they often had important knowledge and information. EISs did nothing to facilitate information sharing. EISs were designed to be uniform, so that data was common and available to all users. This did not allow executives to exploit unique search capabilities to find useful, new information.

Changes in technology have changed our view of EISs

Since the development of that classic view, there have been many technological developments that have changed our view of executive information systems. These changes include

- dominant use of Microsoft Windows as an operating system;
- a shift to data warehouses, away from accounting databases;
- development of tools designed to facilitate capturing and sharing of data and knowledge, such as Lotus Notes;
- development and growth of the Internet and the World Wide Web; and
- development, growth, and availability of intelligent and autonomous computational agents.

Because of all these changes, there has been a movement, away from a single, comprehensive database-analysis tool, to a suite of tools that exploits qualitative and quantitative information available on the Internet and elsewhere, such as Lotus Notes.

Windows seems to have influenced the generation of solutions that provide information for keeping executives up to date. Many of the information technology products designed to support decision making are embedded in a Windows framework. This goes beyond classic spreadsheets and word processing and includes agent-based software.

Firms are using strategies that depend on increasingly detailed data available in data warehouses, and not in their accounting databases. As noted recently by a Wal*Mart executive, "Every cost, every line item, is carefully analyzed, enabling better merchandising decisions to be made on a daily basis."[2] As part of that strategy, Wal*Mart will end up with 7.5 terabytes of data, a quantity of data that is impossible for a human to peruse and make sense of without some sort of information technology-based intervention.

In addition, to facilitate capturing and sharing of important qualitative data, several firms have adopted tools labeled under the heading of an *intranet.* The resulting qualitative databases are too large for complete analysis by human analysts.

Furthermore, although the Internet offers potential opportunities for generating a wide range of information about competitors, ranging from product information to changes in competition management, it is also huge and growing rapidly (some estimates place that growth at 10% per month, regardless of what metric is used). As a result, this information makes it impossible for a human to make sense of more than only a small subset of the available data.

The availability of an infrastructure that permits communication and knowledge sharing, such as the Internet, facilitates cooperative knowledge sharing. In addition, knowledge sharing is facilitated by technology (for example, knowledge repositories) aimed at sharing knowledge in virtual groups. Such sharing allows codevelopment of proposals, reports, and so on, without members of the group ever meeting face-to-face.

Many of those changes concern basic infrastructure issues. However, the way AI is embedded in these information systems has also changed, to mitigate the amount of information that is now available to executives. In particular, perhaps the most important development is the use of agents to help find important, relevant information, and the ontologies on which they are built. The impact of AI on the next generation of EISs is the primary focus of this article.

Executive time continuity

Frequent interruptions beset most executives, with some recent research indicating that they receive a communication (electronic, paper, or oral) every five minutes. As a result, an important development is the ability to filter those communications to determine if a message is important or not. Intelligent agents can help busy executives decide which electronic mail communications are important, cutting down on interruption frequency.

E-mail filtering agents generally require that users specify keywords for various fields in an e-mail message, such as in the Subject field.[3] For example, the following rule might be used:

If subject is "Business Process Reengineering" then store in Folder "BPR".

Agents then use such generated rules to send, forward, or file the mail according to user specifications. For example, both Lotus cc:Mail and Lotus Notes Mail have agent functions where end users can specify rules and personalize their mail handling.

From databases to data warehouses

Data warehouses include substantial data not normally considered part of accounting-transaction databases. In some cases, they capture all of a company's sales transactions. As a result, data warehouses generate unprecedented amounts of information. Consider Frito-Lay, where 10,000 salespersons feed information on 100 product lines from 400,000 stores, generating around 400 billion numbers. When actuals are compared to planned numbers, the figure increases to around 800 billion numbers. Thus, it is easy to see how the amount of numbers in a data warehouse can exceed a trillion.

Data such as this substantiates the contention of Paul Saffo, of the Institute of the Future, that bringing data to the executive has succeeded. Executives and managers have been brought point-and-click tools that let them sort, drill up and down, and navigate through gigabytes of data in search of problems and opportunities. However, there is so much data that the Gartner Group estimates that as little as 5% may actually be accessed. This group also indicates that the availability of such large quantities of data actually reduces, rather than enhances, decision-making capabilities.

Recent AI developments can, as Saffo notes, help "reduce the flood of data to a meaningful trickle." Intelligent, autonomous software agents can be used to search for

changes in data and to identify patterns, all of which can be brought to the attention of the executive.

Recently, agents have become a critical component of many on-line analytic processing, relational on-line analytic processing, and EIS products. OLAP is distinguished from classic on-line transaction processing systems by its generation of multidimensional-views data, which is critical to users being able to perform ad hoc data access and analysis. Comshare's *Robot for OLAP* searches for trends and patterns within an OLAP database. Robot for OLAP finds matches and delivers its findings to the executive.

The Hertz car rental company uses Robot for OLAP to investigate competitors' pricing moves. Hertz's goal is to offer the lowest price of any competitor in a given market segment, car class, and location. However, Hertz has many competitors, several market segments, a wide range of car classes, and thousands of locations. It is impossible for a human to manually search through so much data. As a result, pricing managers define a series of problem-and-opportunity detection rules for the Robot for OLAP. A simple problem (because a car can be rented at a lower rate than at Hertz) detection rule might be,

For all airport locations in the Western Region, find those competitors whose prices today are at least $6 lower than ours.

Whereas a sample opportunity (because Hertz can still raise the rate but keep it lower than competitor rates) detection rule could be,

For all car classes at a given airport location, find those classes where our prices are at least $4 below all the competitors' prices.

After the rules have been specified, they are given to Robot for OLAP, which runs them on a daily basis against the OLAP database. Those situations that satisfy the rules are then used to alert executives to potential problems and opportunities.

Such knowledge discovery from data (data mining) is a critical function for executives, and it has numerous other applications. For example, Wal*Mart, having generated the huge data warehouse mentioned earlier, is in the process of using knowledge discovery to establish trends and relationships in data. Executives have questions such as "What products sell with each other?" Just having the data is not enough: now, firms want to extract information and meaning from it.

Some relevant World Wide Web URLs

EIS companies

List of EIS vendors	http://www.interpublish.com/notes/noteseis.html
List of vendors	http://clever.net/dbd/olap/corpinfo.html
List of companies	http://www.avontiop.co.uk/cqm/analysis.html
Comshare	http://www.comshare.com/
Forest and Trees	http://www.platinum.com/products/businfl/fortreps.htm
MicroStrategy	http://www.strategy.com/press/prs_exec.htm
Cognos business intelligence tools	http://www.cognos.com/businteli/3-2-1.html
Dunn and Bradstreet	http://www.dbsoftware.com/
SAS	http://www.sas.com/

Database information from relational and multidimensional on-line analytical processing

Case for relational OLAP	http://www.strategy.com/dwf/wp_6_x1.htm
Comshare's Robot for OLAP	http://www.comshare.com/olapalrt.htm
An introduction to OLAP	http://www.pilotsw.com/f_and_f/whitpaper/olap/olap.htm#a

Business intelligence: agents for intelligent browsing on the Web

AAAI Fall Symposium on Browsing	http://www.aaai.org/Publications/TechReports/Contents/Symposia/Fall/FS-95-03.html
	http://cs-www.uchicago.edu/~burke/index.html

Business intelligence: general and specific agents for searching the Web

InterAp	http://www.stac.com/Homepages/Software/intcont.html
Quarter Deck	http://www.qdeck.com/
Detect and Alert Robot	http://www.comshare.com/olapalrt.htm

General robot search engines

Lycos	http://www.lycos.com/
InfoSeek	http://www.infoseek.com/
World Wide Web Worm	http://www.cs.colorado.edu/wwww/
Web Crawler	http://www.Webcrawler.com/

Business intelligence and the Internet

Executives require business-intelligence information. However, historically, they have not received it from their information system, except as part of a small set of shared information about, for example, competitors. Nevertheless, recent developments on the World Wide Web greatly facilitate the ability to bring business intelligence to the executive's desk. Business intelligence on the World Wide Web can have three different primary functions: browsing the Web; doing searches on the Web; and having intelligent, autonomous, special-purpose computational agents browse the Web for executives.

Browsers such as Netscape are particularly easy to use. However, because of the Internet's size, AI might be needed to facilitate Web browsing for executives. Some browsers offer the user a tour guide for the Internet. Two of the better known systems are WebWatcher and Letizia.

WebWatcher offers several capabilities beyond browsing, including finding pages related to the current page, adding hyperlinks to the starting page to meet user search goals, and giving advice based on user preferences. With Letzia, the user browses—via Netscape, for example—while the agent collects information about the executive's preferences and tries to anticipate additional items that might interest the user. By incorporating this data and various heuristics, the system infers user interests and then feeds that information back to them, in the form of recommended pages on the Web.

An alternative approach is to have intelli-

Metasearch engines

Find-it	http://www.itools.com/find-it/find-it.html
All-in-One	http://www.albany.net/allinone/
Savvy Search	http://rampal.cs.colostate.edu:2000/
LinkSearch	http://where.com/ls/LinkSearch.html

Specific-purpose agents

BargainFinder	http://bf.cstar.ac.com/bf/
WiseWire	http://www.empirical.com
Webdoggie	http://webhound.www.media.mit.edu/projects/Webhound/

Browsers for touring the Internet

WebWatcher	http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-6/Web-agent/www/project-home.html
Letizia	http://lcs.www.media.mit.edu/people/lieber/Lieberary/Letizia/Letizia-Intro.html

External information gathering: news finding

PointCast Network	http://www.eds.com/pointcst/epc00000.html
NewsAlert	http://www.newsalert.com/fns/
Stanford NetNews Filtering Service	http://woodstock.stanford.edu:2000/
NetNews Wizard	http://www.stac.com/netscripts/newspub.html

Intelligent support on an intranet

Intranet Journal	http://www.brill.com/intranet/
Lotus Notes	http://www.lotus.com/inotes/
ContactFinder	http://www.ac.com/cstar/hsil/agents/art_cf.html
ASK	http://www.ils.nwu.edu/techrep.html
QPEA	http://ai.eecs.umich.edu/people/huffman/hufpubs.html

Virtual groups knowledge exchange/sharing

National Interchange for Knowledge Exchange	http://www.eit.com/creations/papers/nike/
Knowledge Sharing	http://www-ksl.stanford.edu/knowledge-sharing/papers/index.html
ARPA	http://www.arpa.mil/sisto/symp/Demos/I3/PR_GROUP.html
What is an ontology?	http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

gent agents find information or products for the executive. These agents could accomplish specific objectives, or they could be quite general. Some of the better-known specific-purpose agents include

- *BargainFinder*. This system helps the user find the lowest-priced CD.
- *WiseWire* (formerly NewsWeeder). This system learns a user's reading interests and then uses machine learning to find new Web pages and news articles.
- *Webdoggie*. Either on demand or periodically, the system recommends World Wide Web documents based on user-expressed preferences.

For those situations where the executive needs a general-purpose agent, several search systems are available. Much navigation on the World Wide Web uses general robot search engines, such as Lycos, InfoSeek, World Wide Web Worm, and Web Crawler. Metasearch engines that use more than a single search system are also available: for example, Find-it, All-in-One, Savvy Search, and LinkSearch.

However, one of the most interesting developments comes with the general availability of user-programmable search agents for different environments, ranging from the Internet to Lotus Notes, by a number of developers. Executives have the ability to readily generate agents that will search the Web for specifically requested information. Those agents can be "let go" at night, to do their search, and then the next morning the executive can decide if the information is appropriate and sufficient or if additional search efforts are necessary.

External information: news finding

One of the primary sources of qualitative information for executives comes from news. Hence, real-time news monitoring can be critical for an executive to gather information about competitors.

Several services are available to provide news to the executive. They range from passive news-service screen savers (such as PCN) to agent-based active services (such as NewsAlert[4] or NetNews).

NewsAlert and NetNews let executives use agents to

- routinely monitor news wires according to individual, personalized rules,
- automatically and immediately deliver alerts to a user's desktop,
- organize alerts by subject area, and
- provide tools for users to investigate the context of an alert and to communicate those findings to others.

Stanford University's NetNews is another service available on the Internet that could provide useful information to executives. Stanford's NetNews Filtering Service provides a personalized delivery service via e-mail. Users subscribe to the service by submitting profiles that describe their interests. NetNews articles come from newsgroups, and Stanford periodically sends them to subscribers.

Agents can help find news in other environments besides the World Wide Web. For example, David King and Kirk Jones describe an agent developed for use in Lotus Notes.[4] Their *Robot for Internet* filters e-mail and gathers news from news groups available to Lotus Notes.

Intelligent support for finding information on an intranet

Increasingly, Internet technology is being used for intraorganizational communication, work groups, and so on: the intranets. In addition, there is specific intranet software (for example, Lotus Notes) designed to facilitate communication and data sharing between work groups. As a result, executives must be able to find information on an intranet. As such resources continue to grow, intelligent support for information search becomes increasingly important.

Executives might need to browse informa-

tion available to them internally. Although most of the focus of browsing systems has been on the World Wide Web, there is also interest in being able to browse qualitative databases such as those developed in Lotus Notes. This has been of particular interest in the major consulting firms, such as Andersen Consulting and Price Waterhouse, which have been among the first users of Lotus Notes.

Andersen Consulting has developed a system that employs an approach similar to Richard Osgood's ASK system. Andersen Consulting's system facilitates search (of video or text databases) by employing two types of user interfaces: zooming interfaces that provide the content to the user, and a browsing interface that facilitates exploration. To help the user find useful content, the company contends that it needs to establish a zooming feature that anticipates *what* content will be useful, *when* it will be needed, and *how* it must be explained. To facilitate browsing, the system provides elaborations, explanations, comparisons, or other advice. Anderson Consulting also developed an indexing process that creates and maintains the zooming and browsing relationships.

ContactFinder, on the other hand, replaces the "good ole boy" network with the "good ole bot" network.[5] Although groupware such as the Internet and Lotus Notes lets users answer questions and share data, most users can only spend a limited amount of time reviewing questions. ContactFinder has been running on an Andersen Consulting worldwide bulletin board. To date, it has read more than 5,000 messages, has given referrals to over 10% of the questions it has seen, and has been accurate about 86% of the time.

Price Waterhouse is designing Query Planning Environment Assistant to help information-retrieval specialists (such as librarians) select and combine heterogeneous data sources to produce new information services.[6] In particular, QPEA provides decision support for planning a search, including the choice of databases.

Knowledge sharing through virtual groups

One of the most important activities that executives perform is interacting with other executives, gathering knowledge and information along the way. Virtual groups are distributed organizations and teams of people that meet and work together on line. Members of these groups rely on support systems to help gather, retrieve, and share relevant knowledge. Shared knowledge can preserve expertise or process knowledge. Support systems may also be in a position to facilitate rapid assembly of groups or teams to solve particular problems, such as *virtual tiger* teams or proposal-preparation teams consisting of individuals throughout the world.

Shared knowledge is at the core of organizational or group memory. Accordingly, one of the primary issues in the case of virtual groups is using technology to share knowledge. DARPA's Intelligent Information Services supports virtual groups with several new emerging technologies, such as

- institutional memory tools to help organizations capture expertise, including process knowledge and access to expert consultants;
- tools to support multiuser-authoring hypermedia Webs so that groups can build their own; and
- knowledge repositories that can, through use, self-organize and adapt to community needs.

AI also lets organizations rapidly form an executive group to solve a specific problem. Expert systems help executives choose a team from a database of participants and related activity. Such teams could be constructed to meet any of a variety of criteria, such as current availability, experience, or organization position.

Developing and using agents

Agents, whether for the Internet or for a data warehouse, should have certain technical capabilities. First, they need pattern-matching capabilities to perform the activities requested of them. As a result, they should be able to handle relatively complex logical comparisons. In addition, agents should have hierarchical and time intelligence. In particular, agents should be able to inherit rules from other structurally related agents. Moreover, the agent should be able to discriminate between the current month as time changes—for example, through the following rule:

For the current month, which of the products in the fall product line have sales below plan?

Furthermore, agents should be able to function both for individual users and for the organization as a whole. Because agents are for individuals, they must have the ability to be personalized. So they must be easy to build. For agents to be easy to build, problem and opportunity detection rules must be easy to use. In addition, agents should have the ability to interface with data in an intelligent manner. If the data requested is not directly available but is derivable, then, ideally, the agent should be able to generate that derivable data.

Moreover, when one agent comes back with an alert that information should be placed in context by other agents, those agents must provide the executive with information on (1) why the alert was generated, (2) where the executive can go for further information without having to generate a complicated query, and (3) possibly some recommended action for the problem or opportunity. Some agents can learn.[7] However, most depend on the user for definition and purpose.

Agent-based executive information systems change the way information is viewed. The perspective associated with the use of agents is that there is too much information for any one human to make much sense of it. The only real solution is one where autonomous, intelligent computational agents sort through the mounds of available data. At the heart of cybernetics is W. Ross Ashby's law of requisite variety: "Only variety can destroy variety."[8] In terms of data warehouses and the Internet, only computational agents can destroy the hugeness inherent in those databases.

Personalized vs. organizational agents

Much of the focus on agents has concentrated on personalized agents for executives. Alternatively, there are also organizational agents. When the agents are personalized, the executive knows what the agent is talking about, and virtually all agents report directly to the executive: agents do not need to talk to one another. So, generally, there isn't communication confusion over terms or over what was meant when agents were asked to search for particular concepts.

However, in the case of organizational agents, all requests might not be directly understood by the entire organization. In addition, agents might need to talk to other agents to find information, or there might be agents whose sole job is to summarize information

from other agents. Robert Neches and company developed a framework for knowledge sharing that included a shared ontology, a custom ontology, state-dependent runtime knowledge, domain knowledge, and problem-solving knowledge.[9] Perhaps the first two are the key components for this discussion, because they focus on the group and what the group needs to interact. A shared ontology defines a sublanguage for the specific topic area, thus defining the terms and relations that constitute the vocabulary. A custom ontology uses the shared ontology as the basis of the vocabulary for defining a group-specific model. Without a shared ontology, the framework would not ensure that agents could hold a meaningful dialogue. So, to generate a shared model of knowledge, an organization must generate an ontology or a set of ontologies to address those issues examined by the agents.

Although there has been personal use of agents by executives, there is little research regarding organization-based agents for use in EISs and the resulting ontological issues necessary to coordinate interaction between human executives and computational agents. However, using such agents might be necessary, particularly in the area of competitive intelligence searches that track competitive activity for various executives.

LET'S REVISIT THE INITIAL DEFINition of EIS and see how the technology has changed our view. First, classic EISs focused on accounting data, but the new paradigm includes both substantial internal and external data.

Second, whereas under the traditional focus there might have been an effort to increase executive use of e-mail and other communications capabilities, the new paradigm is also bent on using technology to filter through those communications so that only the important ones need to be addressed.

Third, intelligent and autonomous agents let executives browse and search the huge databases now available. Agents also let executives delegate searching and monitor several different settings, including data warehouses, business intelligence on the Internet, news finding, and internal qualitative database analysis. Competitive analysis is an important part of the new paradigm.

Fourth, whereas traditional EISs focused on getting information to the executive in a one-way communication process, new technology can help us establish a two-way flow of information. Under current configurations, knowledge sharing is part of the use of information technology.

Fifth, previously only a few executives had access to the EIS. Now, increasingly more competitors, as well as others in the firm, have access to the Internet, news sources, and even data warehouses. Although a firm might not be using its full potential by accessing these databases using intelligent agents, at least some competitors, and possibly nonexecutives, are likely to use agents for detection and monitoring in a wide range of databases.

Sixth, the common shared view must now be formalized and captured in the ontology used and shared by organizational agents in an EIS. The old definition spoke of a common shared source of internal and external information. However, now that individuals can readily have their own personal agents, this definition no longer fits. Although shared-view information exists under the new paradigm—for example, in the home pages of companies on the Internet, in the information on the firm's intranet, and through organizational agents—there is also room for substantial system-generated personal views because of the potential for personal browsing and personal agents. In fact, individual competitive advantage can accrue to executives who are able to exploit computational agents.

AGENTS SHOULD BE ABLE TO FUNCTION BOTH FOR INDIVIDUAL USERS AND FOR THE ORGANIZATION AS A WHOLE.

Acknowledgments

We would like to acknowledge the comments of Jim Hendler on an earlier version of this article.

References

1. L. Applegate, "Executive Information Systems: Technology Overview," in *MacMillan Encyclopedia of Computers*, MacMillan, New York, 1991.
2. B. DePomps, "Wal*Mart Orders a Super Warehouse," *Information Week*, Mar. 11, 1996, p. 22.
3. T. Malone et al., "Intelligent Information-Sharing Systems," *Comm. ACM*, Vol. 30, No. 5, May 1987, pp. 390–402.
4. D. King and K. Jones, "Competitive Intelligence, Software Robots and the Internet: The News Alert Prototype," *Proc. 28th Ann. Hawaii Conf. Systems Sciences*, IEEE Computer Society Press, Los Alamitos, Calif., 1995, pp. 624–631.
5. B. Kulwich and C. Burkey, "Contact Finder: Extracting Indications of Expertise and Answering Questions with Referrals," *Proc. AAAI Fall Symp. AI Applications in Knowledge Navigation and Retrieval*, AAAI Press, Menlo Park, Calif., 1995, http://www.ac.com/cstar/hsil/agents/art_cf.html.
6. S. Huffman and D. Steier, "A Navigation Assistant for Data Source Selection and Integration," *Proc. AAAI Fall Symp. AI Applications in Knowledge Navigation and Retrieval*, AAAI Press, 1995, http://ai.eecs.umich.edu/people/huffman/ hufpubs.html.
7. T. Selker, "Coach: A Teaching Agent That Learns," *Comm. ACM*, Vol. 37, No. 7, July 1994, pp. 92–99.
8. W. Ashby, *Introduction to Cybernetics*, Chapman and Hall, London, 1956.
9. R. Neches et al., "Enabling Technology for Knowledge Sharing," *AI Magazine*, Vol. 12, No. 3, 1991, pp. 36–56.

David R. King is the director of Comshare's Applied Research Section. He has been responsible for the design and development of several AI-related decision-support system (DSS) and EIS projects. His research interests include intelligent agents and using AI in organization design, modeling, and control. He obtained his PhD in sociology and mathematical statistics from the University of North Carolina. He has authored over 40 articles, as well as an upcoming book on decision support and knowledge-based systems.

Daniel E. O'Leary is an associate professor in the School of Business at the University of Southern California. He received his BS from Bowling Green State University, his MS from the University of Michigan, and his PhD from Case Western Reserve University. He has served as the Program and General Chair of the IEEE Conference on AI in Applications and has chaired the AAAI Workshop on Verification and Validation of KBS. He is a member of the AAAI, the ACM, and the IEEE Computer Society. Readers can contact him at the Univ. of Southern California, 3660 Trousdale Parkway, Los Angeles, CA 90089-1421; oleary@rcf.usc.edu.