McMaster
University

MICHAEL G. DEGROOTE
SCHOOL OF BUSINESS

RESEARCH AND
WORKING PAPER
SERIES

KNOWLEDGEABLE AGENTS FOR SEARCH AND CHOICE SUPPORT IN E-COMMERCE: A DECISION SUPPORT SYSTEMS APPROACH

By

Susan Sproule and **Norm Archer**

Michael G. DeGroote School of Business
McMaster University
Hamilton, Ontario

Working Paper # 446

July, 2000

ER UNIVERSITY
Street West
Ontario, Canada L8S 4M4
(905) 525-9140

Knowledgeable Agents for Search and Choice Support in E-commerce: A Decision Support Systems Approach

Susan Sproule and Norm Archer
Michael G. DeGroote School of Business,
McMaster University

ABSTRACT

Software agents are a major innovation in how people use information systems, and they have parallels with how Decision Support Systems (DSS) support human decision-making. A DSS approach to the development of software agents suggests a highly interactive and flexible interface between the agent and its user, and addresses some potential barriers to the successful adoption of agent technologies. Within a DSS model, agents can be classified as providing search, choice or interface support. Each of these classifications uses techniques originating from separate disciplines and requires different performance measures. We use a real estate agent as a metaphor to examine the descriptive, procedural and semantic knowledge bases that agents can use to support search and choice activities in an exchange of goods or services.

Key Words and Phrases: software agents, decision support systems, search behaviour, choice behaviour, knowledge bases, knowledge representation.

Acknowledgement: This research was supported by a grant from the Social Sciences and Humanities Research Council of Canada.

Knowledgeable Agents for Search and Choice Support in E-commerce: A Decision Support Systems Approach

1. Introduction

Software agents are computer programs that run in the background and perform tasks autonomously, as delegated by the user. Although there has been much research on this topic, usable software agents are at an early stage of development, and are only now starting to appear in real applications. A fruitful application area for software agents is in the area of electronic commerce where agents can help buyers and sellers deal with the flood of information that can be exchanged and processed.

A Decision Support Systems (DSS) approach to software agent development provides insight into how interactive systems can provide flexible and adaptive ways of approaching the complex decision-making processes involved in electronic commerce. This approach also suggests a functional classification system for agents based on their reference disciplines and provides effective ways to evaluate agent performance.

The knowledge that human agents use to provide their services is a useful metaphor for the knowledge bases that software agents may employ. A number of different knowledge representation techniques may be required for different parts of the business process and the choice of technique may depend on the nature of the information and the level of interactivity that is designed into the system. We show how research and development in areas such as Extensible Markup Language (XML) and Knowledge Query and Manipulation Language (KQML), are expected to enhance the knowledge-acquisition and knowledge-sharing abilities of future agent-based systems.

1.1 Approaches to Software Agents

In reviewing the current research and development activity in software agents it is helpful to acknowledge some of the different approaches to this new field of information technology. Software agents were originally conceived and developed within the Artificial Intelligence (AI) research community, and within this community agents are seen as a potential vehicle through which traditional AI techniques can be built into practical applications. The Computer Science/Systems (CS/S) research community views agents as a means to design the distributed and flexible systems that today's systems environment demands. Within the field of Management Information Systems, agents are most closely related to the study of DSS, and we believe that many principles and findings from the DSS field can be applied to software agents to manage the exchange of goods and services. Expert systems (ES) can be viewed as a hybrid of AI and DSS and may provide clues for the successful implementation of intelligent agents in decision support. Table 1 summarizes some basic characteristics of AI, ES, DSS, and software agents.

Table 1 – Comparing AI, DSS, Expert Systems, and Agents

	ARTIFICIAL INTELLIGENCE	EXPERT SYSTEMS	DSS	AGENTS
Objective	Replicate human intellectual activity	Replicate human (expert) decision-making	Assist human decision-makers	Make decisions on behalf of human users
Decision-Making Roles	System replaces human decision-maker	System replaces human decision-maker (often with human review of decision)	Human decision-maker interacts with system	User delegates tasks to agents
Application	Abstract problems	Real-life problems	Real-life problems	Real-life problems
Model for Knowledge Base	General knowledge	Expert	Data and models to supplement individual user's knowledge	Individual user or expert
Critical Success Factor(s)	Similarity to human behaviour	Trust	Usefulness	Usefulness and trust

AI attempts to model and replicate human intellectual activity. Reasoning and learning capabilities developed within AI provide the autonomous and adaptive behaviour that we want for agent applications. However, AI systems are designed as "black box" systems that focus on results. The classic "Turing test", where an AI system is expected to produce behaviour that is indistinguishable from that of a human being, is evidence of this focus. [Turing, 1950] Because of the difficulties involved in capturing and representing general knowledge, many AI research projects are designed to address simplified, abstract problems with limited problem domains. Etzioni [1997] argues that AI researchers who want to develop agents for use in the Web, must follow a "useful first" policy, with intelligence as a secondary goal. He also points out that AI research has not been concerned with issues such as robustness, usefulness, and usability that become critical considerations in commercial applications.

Expert systems (ES) apply rule-based reasoning capabilities developed by AI to real life problems, and have been the most successful commercial application of AI technology to decision problems. Many of the characteristics of DSS systems have been integrated into expert systems as they evolved from research projects into practical applications. We might expect to see agent technologies follow a similar path to adoption and examining how DSS principles have been applied to ES may help developers ensure that agents make this transition smoothly.

The CS/S approach to software agents uses the agent paradigm to provide modularity for designing flexible and distributed architectures for information systems (e.g. Joeris, 2000; Simons et al, 2000). The CS/S approach does not necessarily require agents to be intelligent or adaptable. A popular example of the CS/S approach is the use of agents as "wrappers" for legacy systems. Mobile agents, employed primarily to increase the efficiency of networks, can also be seen to have developed out of the CS/S approach, where the agent's existence is not apparent to the end user of the system. While the CS/S approach is frequently encountered in the agent literature, it is not applicable to e-commerce applications where the system and the user collaborate in making decisions.

1.1.1 The DSS Approach

The study of DSS examines how information systems can be used to help decision-makers make better decisions. Decision-making is a complex, multi-staged process. DSS research recognizes that computers can complete certain parts of this process faster and more accurately than people can. People, however, bring abilities such as creativity, intuition, and experience that enable them to complete other parts of the process more effectively than machines. The DSS approach is to structure parts of an ill-structured problem. These structured parts can then be performed by the system. Humans interact with the system, using their own knowledge to "join" the structured parts together and develop a complete solution to the problem.

This approach demands a lot of interaction between the decision-maker and the system. Some experienced agent developers propose an approach to the design of software agents which is remarkably similar to DSS design, describing a level of interactivity very different from the "black box" model of the AI approach. From their experience developing the *Information Lens* agent system, Malone, Grant and Lai [1997] propose two principles for agent design that fit well within the DSS paradigm:

1. "Don't build agents that try to solve complex problems all by themselves... Build systems where the boundary between what the agents do and what the humans do is a flexible one. We call this the principle of *semiformal systems*..."
2. "Don't build agents that try to figure out for themselves things that humans could easily tell them. Instead try to build systems that make it as easy as

> possible for humans to see and modify the information and reasoning processes their agents are using. We call this the principle of *radical tailorability*..." [pg. 110]

DSS are often built to support individual decision-makers, with one-time or ad hoc problems, and DSS developers recognize that their human users learn during the development process and while using the system. A fast and highly interactive development process is necessary and DSS tools allow changes to be made quickly and flexibly during the process. Similar problems arise in the design of software agents. A defining characteristic of software agents is the ability of an agent to be "personalized" for each user. Agents must be able to satisfy the needs of users with different levels of experience, different perceptions of risk, and different decision-making preferences. This will require tools comparable to those used in DSS development, where users can experiment with "prototype" agents and change their agent's characteristics as they gain experience and trust in the agent's abilities.

A critical success factor of a DSS is the system's usefulness, as defined by the user. While other organizational information systems, such as transaction processing systems and management information systems are usually "mandated" into use, the use of a DSS is generally considered to be optional. Similarly, we assume that people will choose to use an agent, and will do so only if its usefulness is clearly evident.

Current research in DSS is incorporating AI to add structure to larger and more complex areas of the decision-making process. [Siskos & Spyridakos, 1999] By segmenting the overall problem, smaller components can be defined to require very specific domain knowledge and reasoning capabilities, making them well-suited to the limitations of current AI technologies. Bui and Lee [1999] take a DSS approach to developing a system of collaborative agents to assist in crisis management. Their development process involves deconstructing the overall problem-solving process into primitive tasks, specifying the required functionality and behaviour of agents for these tasks, and deciding if use of an agent is justified. Coordination and collaboration mechanisms are then designed so that humans and software agents can integrate their activities into an overall workflow.

1.2 Agent Classifications from a DSS Approach

DSS are commonly considered to include a data subsystem, a model subsystem and a dialogue subsystem. Turban [1988] suggests that AI can be embedded into DSS to support the model, data or dialogue subsystems, the complete system, or the user. We propose a classification of agents according to whether they support search functions through the data subsystem, choice functions through the model subsystem, or interface functions through the dialogue subsystem. This is shown in Figure 1. Much of current software agent development is focused on supporting the dialogue subsystem, primarily through natural language processing (which can be viewed as a specific area within the AI approach). The following discussion concentrates on search and choice support functions but it must be recognized that the boundary between the dialogue subsystem and the other subsystems is often not clearly defined.

Figure 1 - Agents in a Decision Support System

[adapted from Turban, 1988]

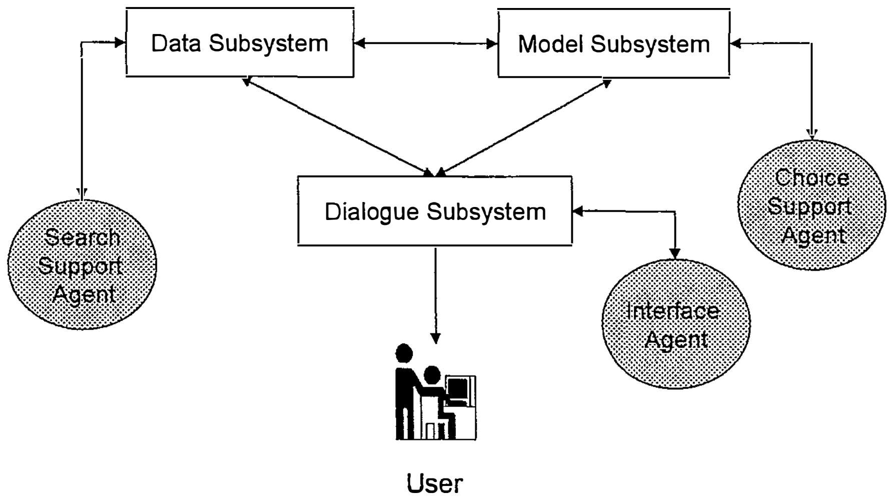

We can distinguish between search, choice and interface agents according to the disciplines from which they borrow their techniques and how their performance should be measured. Table 2 summarizes these disciplines and measures. Agents that support the search function use techniques and measures developed within the information retrieval (IR) community. Agents that support the choice function borrow their techniques from economics, psychology, management science and other disciplines that describe how people make choices between alternatives and how to improve decision quality. The different theories proposed by these disciplines result in a variety of evaluation criteria.

Interface support is based on principles developed in the study of human-computer interaction, where various measures can be used to evaluate a system's "usability". [1]

Table 2 – Agent Reference Disciplines and Measures

Type of Agent	Reference Discipline(s)	Potential Performance Measures
SEARCH SUPPORT	Information Retrieval	Precision Recall
CHOICE SUPPORT	Decision theory from: • Economics • Psychology • Management Science	Consistency of decisions Compare choice to optimal Amount of information used or processed Time to make decision
INTERFACE SUPPORT	Human Computer Interaction	Usability measures such as: • User satisfaction • Errors • Learning time

2. Knowledge Requirements

2.1 A Human Agent Metaphor

The roles of human agents can serve as useful metaphors to derive models of what software agents may do. [Jennings and Wooldridge, 1998] There are many examples of human agents that operate in traditional commerce. We will use real estate agents to

[1] Experiments with agents may use sets of measures that cross these boundaries. In one experiment, Periera [2000] evaluates three different choice mechanisms by asking users to assess their trust in the agent's abilities, confidence in the decision made, propensity to purchase, satisfaction with the decision-making process, perception of cost savings and cognitive effort involved. In another experiment, Vinaja et al [2000] compare the result to an optimal decision, measure the percent of information found and the number of sources used, and obtain user attitudes toward both the process and the outcome.

illustrate the various types of knowledge that a human agent may possess and how this knowledge is linked to the perceived value of their services in supporting search and choice activities. Real estate agents were chosen as our example because they may act for either the buyer or the seller.

Some of a real estate agent's knowledge replicates the client's knowledge. In this case the agent is valued for being able to reduce the time that the client must spend in the process. Software agents that allow the user to build and add to the knowledge base or where the agent learns from the user's actions, would be examples of systems that attempt to replicate this type of support. Real estate agents also possess knowledge that the client may not have, and in this case they are valued for their expertise. Corresponding software agents are those based on the traditional class of rule-based expert systems (ES) and collaborative agents that combine the knowledge of a number of different users to arrive at decisions or make recommendations.

2.2 Types of Knowledge

What knowledge does a real estate agent possess? The agent knows "facts" about properties, sellers, buyers and the market. This type of knowledge will be referred to as **descriptive** knowledge. The agent also knows what to do with information – how to process it to arrive at and implement decisions. We will call this **procedural** knowledge. Finally, the agent also knows what facts are important, both in general and to the client, and how various facts relate to each other. This allows them to evaluate and assimilate new information and communicate by exchanging knowledge in a meaningful way to others. We will call this **semantic** knowledge. [2]

2.3 Knowledge Representation

A fundamental concept in any AI system is knowledge representation - using symbols to build a model of the portion of the real world that is of interest. Knowledge representation techniques include logic, frames, rules and semantic nets. The choice of representation will determine how the knowledge base is processed and the type of reasoning that the system employs. [Davis et al, 1993]

AI research has explored a number of different theories of intelligent reasoning. Davis et al [1993] classify five of these theories according to the disciplines from which they originate as follows: pure logic-based systems (mathematics), probabilistic reasoning systems (statistics), frames (psychology), connectionist systems such as neural networks and genetic algorithms (biology), and utility theory and rational agents (economics). The choice of representation, and its theory of intelligent reasoning, will determine the responses that the system allows. [Davis et al, 1993]

---

[2] In the context of building knowledge-based decision support systems, Holsapple and Whinston [1996] define three primary types of knowledge (descriptive, procedural, and reasoning) and three secondary types of knowledge (assimilative, linguistic and presentation). Our definition of descriptive knowledge is consistent with their classification. We use the category of procedural knowledge to discuss knowledge that can be represented in the processing code of systems, including reasoning capabilities, or procedural knowledge that can be stored and retrieved from a knowledge base. Our classification of semantic knowledge combines the three secondary types of knowledge (assimilative, linguistic and presentation).

If we ask a real estate agent what effect a proposed price will have on mortgage payments, or if the current zoning is consistent with proposed use of the building, we want a logically sound, correct answer. However, many business decisions involve uncertainty and an answer that is "probably" true may be preferable to no answer at all. If we ask our real estate agent to identify the best neighborhoods or schools, we want an informed but necessarily subjective answer. This knowledge could be provided by an "expert" or by a poll of the community. It is important to choose a representation technique that meets the needs of the problem situation. In many e-commerce applications a combination of representation techniques, each for different parts of the overall problem, may be required.

If the user must interact with the system it is important to use models that represent the way that the user conceptualizes the system or the problem. Some models are more understandable than others, and the choice of representation may depend in large part on the degree of the user's involvement and interaction with the system. An ES "replaces" the decision-maker in the same way that we envision tasks being delegated to an agent. By studying the adoption of ES we may be able to identify important design elements that facilitate the development of trust that is necessary for such delegation. For example, while learning to use an ES, users often want to know how the system arrived at a decision. "Explanation" capabilities are now considered a necessary component of ES design. If "explanation" is a design requirement for a part of the process that we want an agent to handle, representations based on connectionist systems like neural networks should not be used because they provide "black box" solutions and no explanations.

3. Search Support

3.1 A Model of Search

Figure 2 shows a basic search model that contains an information source and its representation, an information need and its representation, and a method for comparing these representations.

Figure 2 - A Model of Search

[adapted from Belkin & Croft, 1992]

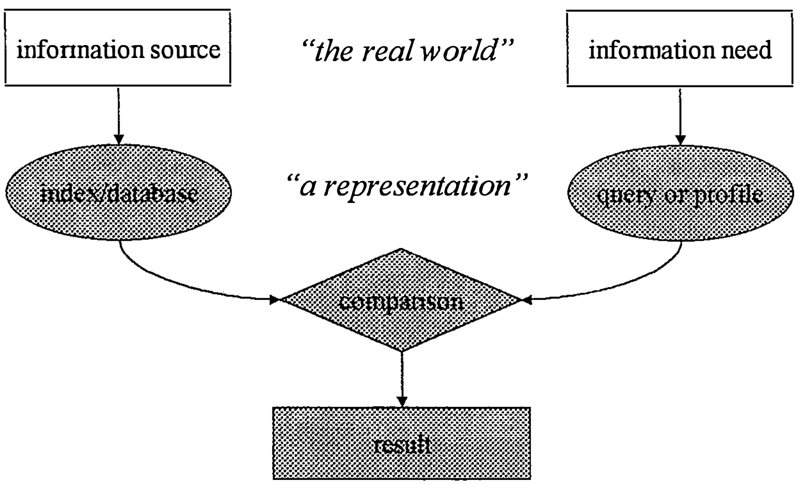

Both the information source and the information need may change over time. Information retrieval deals with a "static" set of sources and a "dynamic" set of one-time needs (queries). A real estate agent deals with this situation when a new buyer arrives. The client's needs are represented by a query and the current listings are searched to retrieve a list of potential properties. Information filtering deals with "dynamic" sources and a "static" need (a profile). [Belkin & Croft, 1992] In our example, if the initial search fails to find a satisfactory property, the buyer's query is stored as a "profile" that can be compared to any properties that are subsequently listed.

Information sources can be unstructured (e.g. full text), semi-structured (e.g. integrated catalogues) or structured (e.g. databases), and the degree of structure will affect the kind of representation used. Full text sources may be represented by sets of index terms. Catalogue items may be represented by minimal information (e.g. a product name and a supplier) with a link to the full information source. The records or objects in a database represent structured information. Similarly, queries or profiles can be unstructured (e.g. a natural language request), semi-structured (e.g. a list of key words or phrases, possibly enhanced by logical operators), or structured (e.g. an SQL command to a database).

Information retrieval often consists of finding structure in predominately free text documents, such as those that make up the Web. Structure can be inferred from features such as hyperlinks [Arocena et al, 1999], header tagging [Guan & Wong, 1999], or question-answer formats [Burke et al, 1997]. An important area of development involves Extensible Mark-up Language (XML). XML allows creators to encode additional structure into their information sources, producing more "searchable" documents by allowing more complex and complete representations to be built.

3.2 Knowledge Bases for Search Support

What knowledge does a real estate agent use in the search process and what types of knowledge bases can agents possess and use to produce the same kinds of results? Table 3 identifies some of the knowledge bases that an agent may use in the search process, classified as descriptive, procedural, and semantic knowledge. Related examples are discussed in the following sections.

Table 3 – Agent Knowledge Bases for Search Support

Search Support		
DESCRIPTIVE KNOWLEDGE	PROCEDURAL KNOWLEDGE	SEMANTIC KNOWLEDGE
Directories, matchmakers, brokers	Search strategies	Ontology
Information sources (databases, catalogues, indices and documents)	Creating representations (e.g. automated indexing, query formation)	Communications protocols
Information needs (queries/profiles)	Matching algorithms	

3.2.1. Descriptive Knowledge

A real estate agent knows where to obtain information and has access to directories and catalogues such as multiple listing services. From these and other sources (i.e. first-hand observation and secondary information from colleagues) the real estate agent can gain considerable descriptive knowledge about properties, sellers, buyers, and market conditions. For an experienced real estate agent, this information covers both the current and past states of the market. The real estate agent also determines and clarifies the client's information needs. These needs can be stated or observed.

Similarly, software agents can use directories, matchmakers, and brokers to identify potential information sources and gain access to current and historical recorded facts through databases, indices, catalogues and documents. Collaborative agents can request and obtain descriptive knowledge that has been collected by other agents [Ackerman et al, 1997]. Learning agents can derive important descriptive information by "observing" information objects [Etzioni, 1997].

Software agents can ask clients to state their information needs.[3] A form or questionnaire can be used to elicit the representation requirements where information is highly structured. However, where information needs are complex and ill-structured, more open processes of collection may be required and these processes can be time-consuming and inaccurate. Significant efforts have been made to design software agents that use proxy information to develop profiles [Rucker & Polanco, 1997] or learn their user's preferences by observing behaviour [Lieberman; Ngu & Wu, 1997]. Collaborative filtering involves comparing profiles to find users with similar information needs so that information judged relevant by one user can be shared with others. This is another way to reduce the profiling effort required by each user. [Balabanovic & Shoham, 1997] Kuflik and Shovel [2000] review a number of approaches to profile generation and updating, and suggest that further research is required to determine how these approaches perform in different application domains.

IR research has shown that efforts to improve the query will contribute more to retrieval success than similar efforts to improve indexing. [Pao, 1989] Twidale and Nichols [1998] suggest that intelligent search support should be designed to "fail gracefully". They develop a graphical representation of the search process that helps the

---

[3] While the client's information needs may indicate potential decision criteria, this is not necessary at the search stage.

user to analyze what the system has attempted to do. This same representation also provides a visual tool to clarify the information need through constructive collaboration with trained information intermediaries or other users.

3.2.2. Procedural Knowledge

We expect a real estate agent to develop an efficient search strategy that will determine the sources to be used and the order in which they are used. This strategy must consider the representation techniques in use and must then be implemented by constructing appropriate representations A buyer's real estate agent will query the MLS catalogue and may discuss the client's needs with other agents. A seller's agent will construct a listing for the MLS and perhaps develop an information sheet as a supplementary representation of further information. The agent must then be able to compare the representations to find potential matches and produce a "reasonable" number of alternatives.

Most software agents have pre-defined search strategies. Some attempts to design adaptive strategies have used query optimization [Duschka & Genesereth, 1997], the efficient use of network resources [Howe & Dreilinger, 1997], or balancing source cost against quality [Lesser et al, 2000].

Software agents are able to create representations and translate between source and need representations. Information extraction techniques such as automated indexing systems are used to create feature-based representations of web documents. Within a specified domain, systems that use more sophisticated linguistic analysis can create structured databases out of information extracted from full text sources. [Cardie, 1997]

Some meta-search agents are able to translate phrase-based requests into either keyword or phrase-based queries acceptable to popular Web search engines. [Etzioni, 1997] "Virtual service representatives" can extract key words and patterns from natural language queries [Neuromedia]. Duschka and Genesereth [1997] suggest three levels of abstraction or representation for a centralized "smart" catalogue. By using different representations at the need, the source, and at an intermediary position, new sources can be added while accommodating a dynamic set of queries. Matchmakers and brokers can provide a similar intermediary state of representation in distributed multi-agent systems.

The agent must be able to compare the source and need representations to find potential matches and produce a "reasonable" number of alternatives. Simple agents may use traditional Boolean systems of information retrieval to match queries to documents, but many agents use more advanced probabilistic systems that weight index terms or look at the statistical distribution of terms within a document. These systems also allow document to document comparisons, creating clusters of sources or user profiles that can be used in retrieval and filtering operations. [Belkin & Croft, 1992; Pao, 1989]

3.2.3 Semantic Knowledge

Computers can store vast amounts of descriptive knowledge, and process this knowledge at speeds greatly beyond human capabilities. However, it is semantic knowledge that will produce what we consider to be intelligent and adaptive systems. A real estate agent knows the relationships between objects and concepts, and can therefore determine the relevance and importance of facts. For example, knowing the age of a heating system, the real estate agent can estimate when the cost of replacement will occur. By using this knowledge of relationships, unexpected information can be assessed and the agent can broaden or narrow the search if the expected information is missing or the amount of information retrieved is overwhelming.

Ontologies are formal descriptions of the relationships between objects and concepts within a domain. These formal descriptions provide a common vocabulary, allowing agents to exchange information in a meaningful and unambiguous way. [Gruber, 1993] Frames and semantic nets are commonly used to represent these relationships. The objects and concepts in a commercial transaction or relationship can be described at many levels. A base ontology covers terms common to all transactions such as those for finance, measurement, and standard contractual conditions. Domain ontology describes objects and concepts within a product category. Individual suppliers or intermediaries can create translation ontologies that relate proprietary terms to the domain ontology. [Keller & Genesereth, 1996] An agent should be able to handle market requests in both surplus and shortage situations. Widemeyer and Lee [1986] describe the ontological requirements for an AI system that can broaden the search to include substitute products in a shortage situation. The system can also apply increasingly stringent criteria to represent the need in a surplus situation. Mukherjee et al [2000] use a similar ontology to reformulate overly specific or overly general queries and to include suggestions of alternate products.

We have indicated that the representation system should be chosen to fit the application domain and that different representations may be required within an e-commerce application. Agents must also be able to communicate with other agents and systems. The Knowledge Sharing Effort (KSE), a project of the University of Maryland

(Baltimore) has contributed in both of these areas, developing the Knowledge Query and Manipulation Language (KQML) and the Knowledge Interchange Format (KIF). [KSE, 1998] KQML provides communications protocols and has been adopted for use in many multi-agent systems, including matchmaking and brokering systems for information retrieval and filtering. KQML specifies the "intent" of the message, based on speech act theory. The message content can be written in any knowledge representation language that is understood by the recipient. [Finin et al, 1994] KIF provides a common representation language for message content. By translating their knowledge into KIF, agents can exchange knowledge even if they have different internal representations and use different programming languages to reason with their knowledge bases. [Bigus & Bigus, 1998]

4. Choice Support

4.1 A Model of Choice

In the real world, the choice problem can be described as interrelated sets of alternatives, criteria and consequences that are processed and analyzed by the decision-maker. [White, 1975] To model the problem (see Figure 3) each alternative can be represented by a set of variables. Parameters are set to represent the selected criteria and any assumptions about the problem situation. A decision model is used to process each alternative, returning a result that represents the consequences of that choice. Prescriptive models compare results and determine the best choice of alternatives, but descriptive models present the results associated with each alternative to the decision-maker.

Figure 3 - A Model of Choice

[adapted from White, 1975]

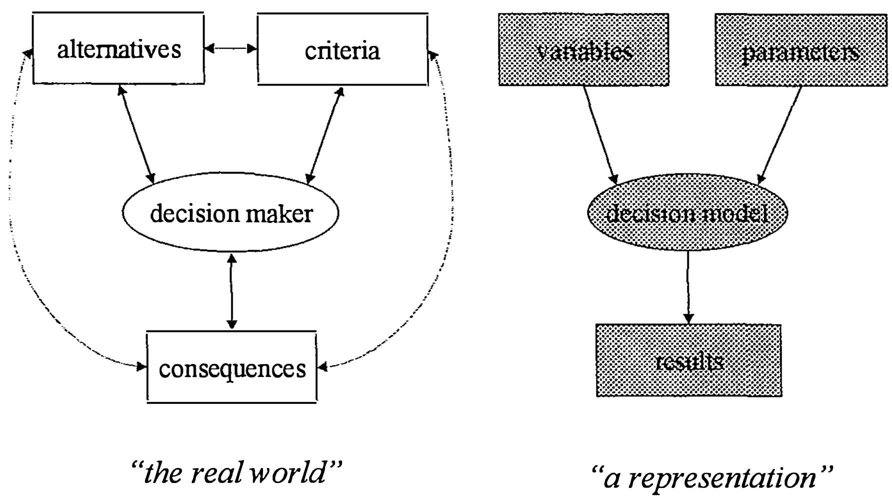

4.2 Knowledge Bases for Choice Support

What types of knowledge does the real estate agent use in the choice process and how can similar knowledge bases be used by software agents? Table 4 identifies some of the knowledge bases that agents may use to assist in the choice process. The following sections provide explanations and examples of these knowledge bases.

Table 4 – Agent Knowledge Bases for Choice Support

Choice Support		
DESCRIPTIVE KNOWLEDGE	PROCEDURAL KNOWLEDGE	SEMANTIC KNOWLEDGE
Attributes to describe alternatives	Decision models and algorithms	Restructuring
Decision criteria (weights, thresholds, trade-offs, etc.)	Process and workflow knowledge	Negotiation protocols
Cases	Negotiating strategies	Transaction protocols

4.2.1 Descriptive Knowledge

Through the search process, a real estate agent has gathered descriptive knowledge of the alternatives - a set of attributes that describe each property. There may also be information about buyers or sellers that will influence the decision process and the agent may use knowledge about market conditions to make reasonable and useful assumptions that help define the problem space. The real estate agent has also collected and refined information about the client's decision criteria including the relative importance of the various attributes, acceptable trade-offs, and threshold levels on specific attributes.[4]

Similarly, software agents have access to descriptive information about the alternatives collected through the search process and can ask the decision-maker to weight the importance of attributes or to set threshold levels for various attributes. Some theories of consumer choice argue that buyers often do not know these preferences in advance. [Bettman et al, 1998] These theories support the need for a highly interactive system, where users can see results and vary their criteria in an iterative process.

The complexity of some commercial transactions and relationships arises from the many outside factors that may or may not warrant consideration. Resource limitations, potential risk and reward, goals, time-pressure, and many other factors can change from one transaction to another. In this context, case-based systems that collect a number of features describing a situation or "case" may be the most effective way to represent the parameters involved in complex purchasing or selling situations.

4.2.2 Procedural Knowledge

A real estate agent is expected to facilitate and assist in decision-making, suggesting different ways of processing information about the alternatives. Experienced agents are likely to have a number of different decision-making techniques that they can match to the situation and the client's individual preferences. A real estate agent often handles transactions where there is more than one decision-maker (such as a husband and wife or a logistics department within a large corporation). An understanding of the information flows and decision-making processes employed within these groups can be used to ensure that the appropriate information is conveyed to each party at each stage in the process. A real estate agent also knows how and when to negotiate.

Theories of consumer choice have developed out of research in economics and psychology. Economic theories of choice assume a perfectly rational decision-maker, able to state clear preferences at the beginning of the choice process. These preferences are used to develop a utility function that can be optimized to form the decision model. Psychological theories of choice have developed out of the belief that humans have limited information-processing capabilities and often use heuristics to reduce the amount of information processing required in decision-making. Heuristic models of decision-

---

[4] In the search process the client's profile represented the information needed to identify a set of alternatives. While a search query or profile may indicate something about the way that a choice will be made, it may be important that other parties not be able to determine

making use a series of constraints to eliminate alternatives until a decision can be made with minimal effort. [Meyer & Kaun, 1991; Bettman et al, 1998]

We can find examples of agents from both of these paradigms. *Personalogic* uses a heuristic approach [Maes et al, 1999], asking the user to specify both hard and soft constraints on the attributes describing alternative brands of a product. It eliminates brands that do not meet the specified hard constraints and presents the remaining alternatives ranked in order of how they compare on the soft constraints. *Tete-a-Tete* is based on a rational model of decision-making [Maes et al, 1999], using weighted-averages and a utility function to recommend a product choice. Consumers often use a combination of decision models. [Bettman et al, 1998] In an interactive system, if the agent is to follow a process that is familiar and recognizable to the user, information may need to be passed between coordinating agents with different modeling capabilities.

The models of buyer behaviour described above are static models that assume that a buyer's choice is independent of previous purchases. Market researchers have also developed models that represent the dynamic nature of consumer decision-making, incorporating factors such as learning, loyalty, novelty seeking, or inertia. [Meyer & Kahn, 1991] Today's technologies make it possible to collect large amounts of time-series data for individual consumers. An agent that is able to predict behaviour from historical purchase information could make timely suggestions based on the loyalty, inertia or variety-seeking tendencies in that consumer's behaviour.

---

the choice criteria from the information request as this could jeopardize future negotiating strategies.

In a business to business environment, agents can use procedural knowledge to integrate activities within buying or selling organizations. While not an e-commerce application, Bui and Lee's [1999] crisis management system shows show how procedural knowledge can be used to coordinate the activities of specialized agents. Agent systems designed to assist in organizational purchasing may require similar procedural knowledge. Joeris [2000] outlines three approaches to agent-based workflow management: 1) role-based autonomous agents as cooperating actors[5], 2) reactive, task coordination agents, and 3) mobile agents that migrate a workflow instance to different "service stations".

Negotiating strategies are procedural knowledge in that they describe a plan of action that can be employed to change the set of attributes describing the alternatives. Simple, one-dimensional (price) time-dependent negotiating strategies have been used by buying and selling agents in an electronic marketplace [Chavez & Maes, 1996]. More sophisticated theories of negotiation can include cooperative and non-cooperative situations and multi-dimensional solution spaces. Negotiating agents must agree to use a common ontology and there must be a way to represent buyer and seller preferences as a utility function. Agents can be preprogrammed with negotiating strategies or equipped with ways to learn effective strategies through techniques such as probabilistic networks or genetic algorithms. [Beam & Segev, 1996]

---

[5] See Yu & Schmid [2000] for the link between role-based analysis and agent-based design. Filipe [2000] discusses policies as "normative knowledge" that can represented in role-based agents as "default rules", reducing the need to negotiate on every action while preserving the agent's autonomy.

4.2.3 Semantic Knowledge

To assist the decision-maker, a real estate agent is able to select information that the chosen model requires, transform it into the form required, and provide reasonable assumptions about missing information. A real estate agent is also expected to know the "rules" of negotiation, and how to communicate with other parties during the negotiation process in a series of offers and counter-offers. Finally, a real estate agent is expected to be able to communicate the results of a decision so that it can be implemented.

"Restructuring" refers to functions that edit, transform, and infer information so that the chosen decision model can be populated with alternatives. [Coupey, 1994] To restructure information, software agents must rely on an ontology to standardize attribute values, eliminate redundant information, and infer missing information. Restructuring can also be seen as a constructive process. Transforming attribute data into standardized values, eliminating redundant or irrelevant information, and rearranging information, may reveal patterns and regularities that suggest the use of a particular choice model. [Coupey, 1994] The constructive nature of restructuring is another indication that an interactive process may be preferred by the decision-maker. By restructuring and presenting information in different ways, the system can help decision-makers to choose models they are comfortable applying in particular situations.

A negotiation protocol defines the rules for an economic mechanism and the form of communications between parties. Negotiating agents must have knowledge of these rules in order to communicate with systems, other humans, or other agents. While a protocol is defined for a particular environment, individual agents can have different strategies as they act within the environment. Users must ensure that the chosen strategy is effective with the given protocol [Brenner et al, 1998] and that the strategy cannot be inferred by other parties. [Beam & Segev, 1997] Many electronic auctions allow participants to "instruct" agents that can monitor for certain events and act on their behalf according to the rules defined for the auction.

There are also rules that must be followed to complete a transaction. EDI messages enable systems to exchange information and create contractual agreements between parties in a transaction. Moore [1998] has shown how standard EDI messages can be interpreted in terms of speech act theory. Covington [1998] examines how KQML, based on speech act theory, can provide a way for software agents to exchange similar messages. Both Covington [1998] and Genesereth [1997] suggest improvements or modifications to simple KQML message protocols so that they can convey the level of detail necessary for commerce transactions.

5. Conclusions

We have shown how a DSS approach to software agents leads us towards flexible and interactive systems that accommodate the capabilities of AI systems and adjust to the user's individual and changing needs. Table 5 summarizes these contributions.

The DSS approach also suggests a classification system according to whether agents support search, choice or interface functions. The techniques used in these functions have different reference disciplines, and agent performance should be measured differently in each function. We have provided examples of the knowledge bases that agents may use to duplicate the services of a human agent in search and choice functions.

Software agents may be an important innovation in how people deal with distributed, complex and ubiquitous systems [Jennings & Wooldridge, 1998] such as those envisioned for e-commerce. While it is desirable to delegate certain activities to such agents, we believe that users will want to retain control of other parts of the process. The boundary between these areas may vary with the user, the task and the situation. How this boundary varies, and how to design systems that accommodate these variations, are important areas for future research.

Table 5 – Contributions from a DSS Approach

The DSS approach promotes…	In the development and design of software agents, this accommodates…
… the segmentation of a large ill-structured decision problem into smaller components	… the limited problem domains that AI applications can adequately address … the need for different representations and reasoning systems in separate parts of the problem
… flexible boundaries between the user and the system allowing for many levels of interaction	… the development of trust … user learning … dynamic situational factors … constructive search and choice behaviour
… an interactive development process with tools that allow the user to adapt and customize the system	… the need for agents to be personalized for each user
… "usefulness" as a critical characteristic of the system	… the need to consider the voluntary nature of agent use.

References:

Ackerman, M., Billsus, D., Gaffney, S., Hettich, S., Khoo, G., Kim, D., Klefstad, R., Lowe, C., Ludeman, A., Muramatsu, J., Omori, K., Pazzini, M., Semler, D., Starr, B., and Yap, P., "Learning Probabilistic User Profiles", *AI Magazine*, Summer 1997, pp. 47-56 : http://www.ics.uci.edu/~pazzani/Publications/AI-MAG.pdf (June 5, 1998)

Arocena, G. O., Mendelson, A. O., Mihaila, G. A., "Applications of a Web Query Language": http://www.cs.toronto.edu/~websql/www-conf/wsql/PAPER267.html (August 4, 1999)

Balabanovic, M. and Shoham, Y., "Fab: Content-based, collaborative recommendation", *Communications of the ACM*, 40(3), March 1997, pp. 66-72.

Beam, C. and Segev, A., "Electronic Catalogs and Negotiations", CMIT Working Paper 96-WP-1016, August 1996: http://haas.berkeley.edu/~citm/wp-1016-summary.html (July 3, 1998)

Beam, C. and Segev, A., "Automated Negotiations: A Survey of the State of the Art", 97-WP-1022, 1997: http://www.haas.berkeley.edu/~citm/auction/index.html (July 2, 1998)

Belkin, N. and Croft, W. B., "Information Filtering and Information Retrieval: Two Sides of the Same Coin?", *Communications of the ACM*, December 1992,Vol. 35, No. 12, pp. 29-38

Bettman, J. R., Luce, M. F. and Payne, J.W., "Constructive Consumer Choice Processes," *Journal of Consumer Research*, Vol. 25, December 1998, pp. 187-217

Bigus, J. P. and Bigus, J., Constructing Intelligent Agents with Java, Wiley Computer Publishing, John Wiley and Sons, Inc, New York, 1998.

Brenner, W., Zarnekow, R. and Wittig, H., Intelligent Software Agents: Foundations and Applications, Springer-Verlag, Berlin Heidelburg, Germany, 1998

Bui, T. and Lee, J., "An agent-based framework for building decision support systems", *Decision Support Systems*, No. 25, 1999, pp.225-237.

Burke, R. D., Hammond, K. J., Kulyukin, V., Lytinen, S. L., Tomuro, N. and Schoenberg, S., "Question Answering from Frequently asked Question Files", *AI Magazine*, Spring, 1997, pp. 57-66

Cardie, C., "Empirical Methods in Information Extraction", *AI Magazine*, Winter 1997, pp. 65-79

Chavez, A. and Maes P., "Kasbah: An Agent Marketplace for Buying and Selling Goods", Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, London, UK, April 1996.

Coupey, E., "Restructuring: Constructive Processing of Information Displays in Consumer Choice," *Journal of Consumer Research*, Vol. 21, June 1994, pp. 83-99.

Covington, M. A., "Speech acts, electronic commerce, and KQML", *Decision Support Systems*, Vol 22, 1998, pp. 203-211.

Davis, R., Shrobe, H. and Szolovits, P., "What is Knowledge Representation?", AI Magazine, Vol. 14(1), 1993, PP. 17-33: http://medg.lcs.mit.edu/ftp/psz/k-rep.html (August 10, 1998)

Duschka, O. M. and Genesereth, M. R., "Infomaster – An Information Integration Tool", 1997: http://infomaster.stanford.edu/infomaster-info.html (May 11, 1999)

Etzioni, O., "Moving Up the Information Food Chain", *AI Magazine*, Summer 1997, pp. 11-18.

Filipe, J., "Agent Oriented Organization Modelling using Information Field Theory", from the *Program of AOIS at AAAI-2000*: http://www.aois.org/2000/Program-AAAI.html (June 2, 2000)

Finin, T., Fritzson, R., McKay, D. and McEntire, R., "KQML – a Language and Protocol for Knowledge and Information Exchange" (Technical Report CS-94-02), Computer Science Department, University of Maryland, UMBC, Baltimore, MA: http://www.cs.umbc.edu/kqml/papers/kbkshtml/kbks.html (September 17, 1998)

Gensereth, M. J., 'An Agent-based Framework for Interoperability", from Software Agents, Bradshaw, J. M. (Ed.), AAAI Press / The MIT Press, Menlo Park, California, 1997

Gruber, T., "What is an Ontology", 1993: http://www-ksl.stanford.edu/kst/what-is-an-ontology.html (July 5, 1999)

Guan, T. and Wong, K. F., "KPS --- a Web Information Mining Algorithm": http://www8.org/w8-papers/4a-search-minimg/kps/kps.html (July 5, 1999)

Holsapple, C. W. and Whinston, A. B., Decision Support Systems – A Knowledge-based Approach, West Publishing Company, St Paul, Minnesota, 1996.

Howe, A. E., and Dreilinger, D., "SavvySearch: A Metasearch Engine That Learns Which Search Engines to Query", *AI Magazine*, Summer 1997, pp. 19-25.

Jennings, N. R., and Wooldridge, M., "Applications of Intelligent Agents", from Agent Technology Foundations, Applications, and Markets, Jennings, N. R., and Wooldridge, M. (Ed), Springer-Verlag, 1998: http://www.cs.umbc.edu/agents/introduction/jennings98.pdf (June 17, 1998)

Joeris, G., "Decentralized and Flexible Workflow Enactment Based on Task Coordination Agents", from the *Program of AOIS at CaiSE*00*: http://www.aois.org/2000/Program-CaiSE.html (June 2, 2000)

Keller, A. M. and Genesereth, M. R., "Multivendor Catalogs: Smart Catalogues and Virtual Catalogues", 1996: http://infomaster.stanford.edu/infomaster-info.html (May 11, 1999)

KSE (Knowledge Sharing Effort): http://www-ksl.stanford.edu/ (August 20, 1998)

Kuflik, T. and Shoval, P., "User Profile Generation for Intelligent Information Agents – Research in Progress", from the *Program of AOIS at CaiSE*00*: http://www.aois.org/2000/Program-CaiSE.html (June 2, 2000)

Lesser, V., Horling, B., Raja, A., Zhang, X. and Wagner, T., "Resource-Bounded Searches in an Information Marketplace", *IEEE Internet Computing*, March-April, 2000, pp. 49-58.

Lieberman, Henry, "Autonomous Interface Agents": http://lieber.www.media.mit.edu/people/lieber/Lieberary/Letizia/AIA/AIA.html (July 2, 1998)

Maes, P., Guttman R. H. and Moukas, A. G. "Agents that Buy and Sell: Transforming Commerce as we Know It," *Communications of the ACM*, March 1999, p.81

Malone, T., Grant, K. R., and Lai, K-Y, "Agents for Information Sharing and Coordination: A History and Some Reflections", from Software Agents, Bradshaw, J. M. (Ed.), AAAI Press / The MIT Press, Menlo Park, CA, 1997.

Meyer, R. J. and Kahn, B. E., "Probabilistic Models of Consumer Choice Behaviour", in Handbook of Consumer Behaviour, Robertson, T. S. and Kassarjian, H. H. (Eds), Prentice-Hall, Englewood Cliffs, NJ, 1991.

Moore, S. A., "Categorizing automated messages", Decision Support Systems, Vol. 22, 1998, pp.213-241.

Mukherjee, R., Dutta, P.S. and Sen S., "Analysis of domain specific ontologies for agent-oriented information retrieval", from the *Program of AOIS at AAAI-2000*: http://www.aois.org/2000/Program-AAAI.html (June 2, 2000)

Neuromedia, http://www.neuromedia.com (May 15, 2000)

Ngu, D. S. W., and Wu, X., "SiteHelper: a localized agent that helps incremental exploration of the World Wide Web", *Computer Networks and ISDN Systems*, Vol. 29, 1997, pp. 1249-1255.

Pao, M. L., Concepts of Information Retrieval, Libraries Unlimited, Inc., Englewood, CO, 1989.

Pereira, R.E., "Optimizing Human-Computer Interaction for the Electronic Commerce Environment", *Journal of Electronic Commerce Research*, Vol. 1, 2000: http://www.csulb.edu/journals/jecr/issues/20001/paper5.htm (May 9, 2000)

Rucker, J., and Polanco, M.J., "Personalized Navigation for the Web", Communications of the ACM, Vol. 40, No. 3, March 1997, pp. 73-75.

Simons, J., Arampatzis, A.T., Wondergem, B.C.M., Schomaker, L.R.B., van Bommel, P., van der Weide, Th.P. and Koster, C.H.A., "PROFILE – a Multi-Disciplinary Approach to Information Discovery", from the *Program of AOIS at CaiSE*00*: http://www.aois.org/2000/Program-CaiSE.html (June 2, 2000)

Siskos, Y. and Spyridakos, A., "Intelligent multicriteria decision support: Overview and perspectives", *European Journal of Operational Research*, Vol. 113, 1999, pp.236-246.

Turban, E., Decision Support and Expert Systems, MacMillan, New York, NY, 1988.

Turing, A.M., "Computing Machinery and Intelligence", *MIND – A quarterly review of Psychology and Philosophy*, Vol. LIX, No. 236, October, 1950: http://www.abelard.org/turpap/turpap.htm (August 10, 1999)

Twidale, M. and Nichols, D., "Designing interfaces to support collaboration in information retrieval", *Interacting with Computers*, Vol. 10, 1998, pp. 177-193.

Vinaja, R., Slinkman, C.W. and Mykytyn, P.P., "An Empirical Study of the Use of Internet Information Delivery and Agent Software for Business Decision Making", from the *Program of AOIS at AAAI-2000*: http://www.aois.org/2000/Program-AAAI.html (June 2, 2000)

White, D. J., Decision Methodology – A Formalization of the OR Process, John Wiley and Sons, London UK, 1975.

Widemeyer, G. R. and Lee, R. M., "Preference Elicitation in Decision Aiding: Application to Electronic Shopping", from Decision Support Systems: A Decade in Perspective, MacLean E. R., and Sol H. G. (Eds.), Elsevier Science Publishers B. V., North Holland, 1986.

Yu, L. and Schmid, B.F., "A Conceptual Framework For Agent Oriented and Role Based Workflow Modeling", from AOIS'99 Papers at http://aois.org/Papers.html : http://www.mcm.unisg.ch/people/lyu/yuAOIS99.pdf (June 8, 2000)

Michael G. DeGroote School of Business
McMaster University

WORKING PAPERS - RECENT RELEASES

420. Harish C. Jain and Simon Taggar, "The Status of Employment Equity in Canada", December, 1996.

421. Harish C. Jain and Parbudyal Singh, "Beyond The Rhetoric: An Assessment of the Political Arguments and Legal Principles on Strike Replacement Laws in North America", January, 1997.

422. Jason Schwandt, "Electronic Data Interchange: An Overview of Its Origins, Status, and Future", March, 1997.

423. Christopher K. Bart with John C. Tabone, "Mission Statement Rationales and Organizational Alignment in the Not-for-profit Healthcare Sector", November, 1997.

424. Harish C. Jain, Michael Piczak, Işik Urla Zeytinoğlu, "Workplace Substance Testing - An Exploratory Study", November, 1997.

425. S. Suarga, Yufei Yuan, Joseph B. Rose, and Norman Archer, "Web-based Collective Bargaining Support System: A Valid Process Support Tool for Remote Negotiation", January, 1998.

426. Pawan S. Budhwar and Harish C. Jain, "Evaluating Levels of Strategic Integration and Development of Human Resource Management in Britain", March, 1998.

427. Halit Üster and Robert F. Love, "Application of Weighted Sums of Order p to Distance Estimation", April, 1998.

428. Halit Üster and Robert F. Love, "On the Directional Bias of the ℓ_{bp}-norm", April, 1998.

429. Milena Head, Norm Archer, and Yufei Yuan, "MEMOS: A World Wide Web Navigation Aid", October, 1998.

430. Harish C. Jain and Parbudyal Singh, "The Effects of the Use of Strike Replacement Workers on Strike Duration in Canada", February, 1999.

431. Parbudyal Singh and Harish C. Jain, "Strike Replacements in the United States, Canada and Mexico: A Review of the Law and Empirical Research", February, 1999.

432. John W. Medcof and Jeremy Boyko, "Reinforcing, Revising and Reconciling Attributions in the Employment Interview", March, 1999.

433. Norm Archer, "World Wide Web Business Catalogs in Business-to-Business Procurement", March, 1999.

434. Diwakar Gupta and Saifallah Benjaafar, "Make-to-order, Make-to-stock, or Delay Product Differentiation? - A Common Framework for Modeling and Analysis", April, 1999.

435. Harish C. Jain, Parbudyal Singh and Carol Agocs, "Recruitment, Selection and Promotion of Visible Minorities and Aboriginals in Selected Canadian Police Services", April, 1999.

436. Harish C. Jain and Angus Bowmaker-Falconer, "Employment Equity/Affirmative Action Codes of Practice and Best Practices in USA, Britain, Canada and Other Selected Countries", May, 1999.

437. Diwakar Gupta, Yavuz Günalay, and Mandyam M. Srinivasan, "On the Relationship Between Preventive Maintenance and Manufacturing System Performance", June, 1999.

438. Jinliang Cheng, George Steiner, and Paul Stephenson, "A Fast Algorithm to Minimize Makespan for the Two-Machine Flow-Shop Problem with Release Times", June, 1999.

439. Jinliang Cheng, George Steiner, and Paul Stephenson, "A Fast Algorithm to Minimize Maximum Lateness for the Two-Machine Flow-Shop Problem with Release Times, June, 1999.

440. Norm Archer and Yufei Yuan, "Electronic Commerce and the Business-to-Business Customer Relationship Life Cycle", August, 1999.

441. Halit Üster and Robert F. Love, "Calculation of Confidence Intervals for Estimated Distances", August, 1999.

442. George O. Wesolowsky, "Detecting Excessive Similarity in Answers on Multiple Choice Exams", October, 1999.

443. Naresh C. Agarwal, "Mandatory Retirement and the Canadian Human Rights Act", November, 1999.

444. Susan Sproule and Norm Archer, "Software Agents in Electronic Commerce", March, 2000.

445. Milena Head and Yufei Yuan, "Privacy Protection in Electronic Commerce – a Theoretical Framework", June, 2000.