

Intelligent Web User Interfaces

Myron Sheu

Department of Computer Information Systems
California State University at Dominguez Hills, Carson, CA 90747
msheu@csudh.edu

Xin (James) He

Department of Information Systems & Operations Management
Dolan School of Business, Fairfield University, Fairfield, CT 06824
203-254-4000 ext. 2835, xhe@mail.fairfield.edu

ABSTRACT

This paper investigates the key components of an intelligent web user interface to facilitate online investment as a novel approach to compensating for the impersonality of e-commerce. By analyzing challenges to online brokerage services and evaluating key criteria for a viable intelligence system, we develop a decision tree based intelligent web user interfaces model. The resulting intelligent model is intended to help online shoppers avoid common mistakes by means of implicit reasoning, flexible knowledge granularity, and effective reasoning-by-exception, which is significantly different from the traditional approaches that largely rely on assistance from remote control knowledge engines. One of the key contributions of the intelligent web user interfaces model introduced in this paper is that it provides a heuristic guidance behind the scene for both online shoppers and online stores without extra structural constraints or financial burdens.

INTRODUCTION

Online shopping has been taking much longer than initially predicated to become an attractive alternative to traditional shopping channels, especially in the retailing markets where purchasing decisions are not easy to come up because a large number of related factors complicate the decision-making process (Haubl & Trifts, 2000). To make the online shopping more attractive to traditional offline consumers, we introduce in this paper an effective model that will greatly enhance the intelligence of web user interfaces to substantially improve online personalized assistances that are comparable to the traditional shopping environment. While efforts were previously made to achieve such an objective (Schafer *et al.* 2001), many of them were either structurally or financially infeasible due to extra constraints imposed by underpinning infrastructure. Zviran and Erlich (2003) found that most researchers tried to measure the success of an information system as a function of cost-benefit, information value, or organization performance. In this research, we extend the assessment function into three evaluation criteria for a viable intelligent web user interface.

First, a viable web intelligence system must incur little administrative overhead. Since online stores usually have to deal with many more potential customers than traditional stores due to near zero transportation costs and around the clock service hours, extra maintenance overhead could incur unexpected administrative costs.

Second, a viable web intelligence system should not result in any confusion to online consumers but provide personalized assistances to the online shoppers. Since there is no way to assess the intellectual competence of online shoppers before offering any kind of assistance, any significant upfront requirements, even with a brief learning curve, is deemed unrealistic (Payne *et al.* 1996).

Third, the mechanism of the intelligence must also adapt autonomously in accordance with individual needs that likely evolve over a period of time (Kobsa *et al.* 2001), while it has to keep simple and intuitive as stated in the first two criteria.

In this research, we develop an intelligent web user interface model aimed at meeting these three criteria, which departs significantly from the existing approaches in the marketplace. In addition, we define the key components suitable for the online brokerage business due to its leadership position in the e-commerce industry and its decision-making complexity. The rest of the paper is organized as follows. Section 2 provides a literature review. Section 3 elaborates the key components in the online brokerage operations. Section 4 develops the intelligence web user interfaces model. Finally, Section 5 summarizes the research findings.

LITERATURE REVIEW

While applauding the existing literature's significant contributions to online commerce, we also point out some common weaknesses against the viable criteria highlighted earlier.

Although intelligent agents that run transparently behind the web user interfaces can significantly enhance the intelligence of a web site as it interacts with online users (Wang *et al.* 2002), such an approach does not address the design issues of web user interface *per se* as such agents operate at a separate tier. As we have emphasized earlier, any efficient and effective approach to enhancing the intelligence of web user interface must incur little overhead, strive for simplicity, which is a unique advantage of online commerce. Benetti *et al.* (2002) has proposed an interesting framework for information integration online. The complexity of the proposed model nevertheless could render few benefits to both e-commerce users and owners. In addition, it also largely relies on other background processes to provide intelligence, which basically replicates the traditional expert-system approach. Recently, Fano and Kurth (2003) present a graphical web user interface model that would offer personalized advice in line with the lifestyle of an online customer. While the model seems applicable in principle, the restrictive constraints have severely weakened the effectiveness of such a web interface model for two reasons. First, its graphical presentation in fact stems from a remote process that depends on a large database, which would significantly increase the network traffic. Second, the advice is

constructed procedurally rather than declaratively so that the system may have to conduct heavy calculation since many factors have to be considered in order to render a valuable advice and many of these factors are not predictable. Therefore, an online store relying on such a model would be too complicated to pass the benefits of online shopping to end-users.

As one of the e-commerce pioneers, Merrill Lynch has recently created a web site named "holdrs.com" which provides helpful online tools for individuals to create and control their customized portal pages. While these portal pages could be generated in various formats in accordance with users' preference, these generated user interfaces possess little intelligence and don't provide any customized decision support in regard to individual financial situations and goals. Charles Schwab, another e-commerce pioneer, has recently enhanced the customization of an individual's personal web page via MySchwab, through which an individual investor is able to build his personal portfolios and to watch relevant information closely. The drawback of such a system is that the information structure can only hold plain financial data in a tabular format and provide little assistance in processing the data into useful information. This approach would generally render more benefits to those who are disciplined in conducting analysis on their own with certain prior knowledge of utilizing the tools.

In short, nearly all the existing approaches to facilitating online shopping generally render three incongruities against the three evaluation criteria stated above. First, their assistance usually better fits those who are disciplined in conducting analysis because the reasoning steps must be explicitly conducted. These firms usually recommend a product or security regardless of individual taste and, consequently, require their customers to filter and assemble the information on their own. Second, their guidance usually takes time for their customers to digest, which defeats the purpose of online shopping assistances. Third, most e-commerce web models we investigated primarily follow the traditional expert system approach to gaining intelligence. As a result, the intelligence of that kind would weaken other quality features. Thus, these approaches fail to meet the evaluation criteria for intelligent web user interface and, therefore, do not effectively and efficiently address the challenges and unique characteristics of online shopping.

DESIRED QUALITY ATTRIBUTES AND FUNCTIONALITIES

Key Quality Attributes

In developing our model, we believe that the intelligence of an e-commerce web site aimed at facilitating customer decision-making shall subsequently possess three intrinsic quality attributes: derivability, digestibility, and integrity. By derivability, we mean that the knowledge is well organized such that the whole system facilitates an informed decision making process. By digestibility, we mean that the knowledge is understandable and intuitive to the customers. By integrity, we mean the knowledge is reliable and is not excessively driven by instantaneous information. Furthermore, an intelligent web user interface shall be able to effectively assist online investors in perceiving the valuation of their portfolios in relation to the global security market and, thus, shall help them make

rational decisions. Two key functional features are proposed below to make it possible for an intelligent web user interface to possess these quality attributes: a) implicit and heuristic reasoning, and b) knowledge granulation.

Implicit and Heuristic Reasoning

An intelligent mechanism effective for delivering knowledge based on mass information is to apply constraints to an object-oriented representational model. The resulting representation is a set of constraint objects that could form an implicit reasoning bed at a chosen abstract level. It further facilitates the abstraction of knowledge and thus derives a layered structure of knowledge. The implementation of constraints initiates from defining default constraints (Reiter, 1985). On a security trading web site, a default constraint can be identified to depict the intrinsic value of a security that should be less volatile because of the arrival of news events concerning the security. To help investors stick with the intrinsic value of a security, default constraints should also be arranged into a hierarchy. Figure 1 shows that the valuation of a security, such as GMH, is assessed with a set of constraints, some of which are considered exceptions to its own parent object representing a typical valuation in the industry. In particular, DBS (direct broadcast services) is also described by a set of default constraints, some of which are overridden by its child object.

Furthermore, there could be another object above the DBS object; the grandparent object is created to describe a typical valuation in the sector, e.g., telecommunication. A set of default constraints is defined to represent the default valuation of an enterprise in the sector. Equipped with such hierarchical defaults, the web site would be able to remind online investors of the following questions, implicitly however:

1. Is GMH a typical business entity in the telecommunication sector? If yes, the valuation shall be influenced heavily by the top set of the default constraints.
2. Is GMH a typical business entity in the DBS industry? If yes, the valuation of GMH is entailed using the default constraints at the second layer to supersede the top set of the default constraints when any conflicts between two sets occur.

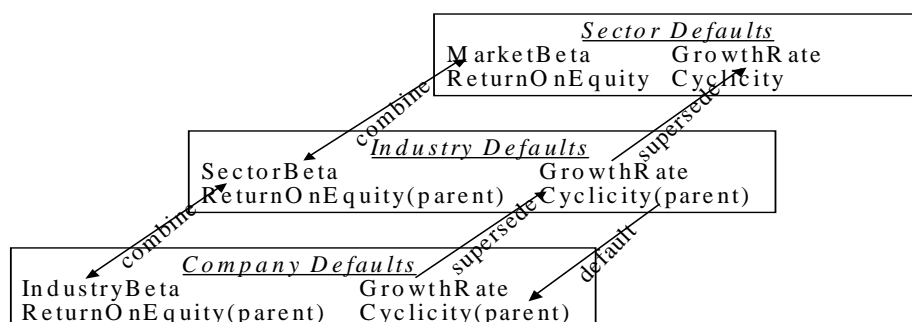


Figure 1. Defaults Relationship between Lower, Current, and Upper Layers.

Any additional constraints appearing at the third layer would imply that GMH either has additional attributes about its valuation or holds exceptional constraints superseding some constraints at the parent layer. Note that a constraint at a non-bottom layer is always considered a default constraint.

The use of constraint objects in a hierarchical structure makes complex constraints manageable, because a group of constraints may then be possibly organized to constitute a macro concept with a desirable granularity. The hierarchy in the example infers that an industry in a depressed sector should not be marked attractive unless some exceptional qualities that the industry possesses. Similarly, a business entity within a depressed industry shall likely deliver the depressing results, unless it has peculiar strength to supersede those defaults attributed by the industry to which it belongs.

Two observations are worth mentioning regarding the default hierarchy. First, the knowledge represented through a default hierarchy reminds online investors of the valuation of a typical corporation in the sector or industry, but it does not necessarily mean that the rest of corporations in the sector or industry shall be so evaluated. Instead, the force of the default hierarchy implies that, in the absence of any evidence to the contrary, only this corporation possesses such a typical valuation. Assisted with a default-reasoning framework, individual investors should not speculate without adequate counter evidence. If a P/E ratio is displayed in green but the growth rate for the same business entity is displayed in red, the contrast in color should serve as an alert to individual investors. Such assistance should be more effective than expecting them to find out the facts elsewhere. Secondly, with the default hierarchy, reasoning only proceeds down to a preferred layer. A mutual fund investor, for example, may only need to know which sector or industry holds better potentials. Likewise, an investor who is interested in value

stocks may only need to dig into one or two layers below a depressed industry or sector to find out some outstanding corporations there.

There are several advantages of adopting a hierarchy of constraint objects. The complexity of the financial market has to be highly abstracted so that an average online investor can perceive fundamentals of a security of his interest on a web user interface generated specifically for him. Since e-commerce demands simplicity, a reasoning bed with logical implications shall circumvent the needs for explicit, lengthy reasoning, in addition to any implicit reasoning processes. Therefore, default reasoning fits the scenario nicely for desired intelligence. Also, the probability of predicating true or false of the financial future mandates that the arbitrary reasoning be unsuitable. The constraint-based reasoning, through a hierarchy of objects, offers a flexible structure in which implicit reasoning would bring about a desirable intelligence.

Knowledge Granulation

Knowledge via discovering, granulating, and representing relationships among information entities is another intrinsic functional feature that makes web user interfaces more capable of supporting conceptual intelligence. A specific challenge to designing web user interfaces is to express the temporal data at various levels of abstraction and projection (Bauer & Scharl, 1999). To address this challenge, a web user interface that facilitates a variety of relationships for individuals would help online investors visualize an unbiased knowledge of both historical events and future indications (Ladkin, 1986). While a fundamental analysis focuses on the current financial situation of a company by means of a set of static objects, appropriate indices extend from the current snapshots of a company to historical data and thus facilitate a technical analysis. In addition, many distinctive perspectives of a portfolio are to be analyzed through multiple chains that help reveal causes and effects. The importance of such a mechanism can be seen through the example of the instant dissemination of available events, which makes online investors believe what they see is current with the false assumption that the influence has not yet reflected by the security price. In reality, however, the incoming information has already impacted the price of the security in question. Therefore, it is imperative for an online investor to view the related events and price changes simultaneously.

A simple, effective way to express events and prices simultaneously is the before and after chain since most events naturally fall into a sequence. In an object-oriented paradigm, events could be chained in parallel with the reactions. Furthermore, the before and after chain can be organized to deliver these events at appropriate economic occasions. In accordance with importance, the various abstractions of time interval and event granularity should also be made available for end-users to navigate efficiently. For example, at a higher-level abstraction, there could be a chain of more significant events during a time interval in parallel with sub-chains of less significant events.

The adjustable indices could be leveraged to enable an online customer to draw an idea from many snapshots of a changing target at various levels of granularity and to easily develop various knowledge granularities by changing the sampling time interval. It is

useful when the time arrival of an event is of specific interest to someone who may take similar action if the same arrival pattern occurs in the future. Most financial forecasts, however, largely depend on statistical analysis based on arrival time intervals, which is quite helpful in support technical analysis for online investors to visualize a comprehensive assessment of their holdings.

AN INTELLIGENT WEB USER INTERFACE MODEL

The intelligent web user interface model, based on implicit reasoning and knowledge granulation in section 3, applies defaults and exceptions to express major valuations of a security from various perspectives, where a hierarchical data structure is extended from the layer of a default valuation as a function of, for example, three parameters: S&P 500 beta value, S&P 500 return, and risk-free return. Figure 2 shows that the default hierarchy extends another branch at the Sector layer to adapt the Sector defaults, where the branched layers of the hierarchy displays exceptional valuations (Beck & Bommel, 1999). An example in hand is an exceptional valuation for the Sector entity in the event of a significant increase of interest rate.




Figure 2. Defaults at Child Layer Considered Exceptional to that at Parent Layer.

In support of multiple levels of confidence, exceptions could be organized into subclasses within a list of values that are separated from the ones that do not. As more attributes take exceptional values in a consistent manner, non-monotonic reasoning should allow an exceptional valuation of the stock to replace a typical valuation that is originally adopted as default.

Recognizing Exceptions

The web user interface model specified thus far would only function properly if the financial market were quite static and individuals had stable financial situations and goals. Of course, these assumptions are not true. For example, an experienced online investment environment is required to understand the individual's investment behavior and to deliver the personalized knowledge representation. In addition, the economic potential of a company or an industry is also evolving. Therefore, the ability to accommodate changes is crucial for a truly intelligent web user interface, but it needs additional reasoning mechanisms to incorporate exceptional facts.

To accurately represent the dynamic nature of investment, the first task is to recognize exceptional valuations that contradict to a default valuation derived from use of inheritance in the previously described mode. It is important however to represent exceptional observations developed through inductive reasoning even when they are yet to mature but have already possessed some significance. Although the defaults generally facilitate a financial decision by focusing on the market risk but not on unique risk, without exceptional considerations, the unique risk of an individual security is largely overlooked and subsequently is up to an individual justification. Since the unique risk could weigh significantly in comparison to the market risk, the quality of knowledge representation weakens if the unique risk is not represented qualitatively or quantitatively.

In contrast to exceptions resulting from deductive reasoning, some exceptions arise through an inductive reasoning process. This kind of exceptions originated from specific performance of a particular security or custom practice for an individual investor. After repeatedly observing the occurrences contradictory to the default results, consistent observations should be collectively identified as an exception.

Another kind of exceptions appropriate for inductive reasoning is related to specific performance or financial decisions for an individual investor. An individual investor, for example, may primarily invest in a few large companies or only in domestic companies because of limited access to the information. To help make a correction and eventually improve the long- term performance of the individual's portfolio, an exception is created to hint that the individual should restrain from further investing in these stocks. Although upon the creation of a new exception it is usually weak, its associated confidence gradually strengthens if more observations support the exception.

In addition, exceptions are classified in terms of granularity: some are applicable to individual securities while others are applicable to one's entire portfolio. For instance, the exception that could be created at the level of the entire portfolio is to suggest the avoidance of buying on Monday and selling on Friday.

Note that if no exceptions were raised, one should manage his investment portfolio by following general financial guidelines for his risk preference. That is, without particular

reasons, an investment decision is based on the principle that is configured by his financial goals.

In contrast to exceptions at the portfolio level, some exceptions are there for a single security. These exceptions usually represent unique characteristics of the security and they are useful to detect an abnormal security in the industry. In an object-oriented paradigm, by inheritance, a security may not be typically recommendable as a long-term investment, but by polymorphism, it could be considered attractive as a short-term investment opportunity when it hits a cyclic bottom. In fact, most short-term buying opportunities are represented as exceptional cases, since general representational models are designed for long-term goals.

A security-level exception is observed when a security contributes to unnecessary risk. An exception is attached to selective securities with a heavy concentration in a few industries because excessive concentration on specific industries may somewhat deteriorate the quality of an investment portfolio. Such exception is served as a reminder to rebalance the distribution of the investment baskets.


While representing exceptions is a necessity in object-oriented modeling, the dynamically nature of exception changes in light of both qualitative and quantitative significance imposes great challenges to effective knowledge representation. In the financial market, the exceptional valuations of a security often stem from a variety of reasons that are in turn caused by other reasons and tracing all causes exhaustively is neither feasible nor necessary. Since these causes are not equally influential to the long-term quality of a security, they should be evaluated accordingly. Further more, since the significance of an exception varies over time, its variation demands a dynamic representation of the exceptions so that the resulting influence to a security is favorable and appropriate. Therefore, the ability to appropriately represent these exceptions would facilitate millions of unsophisticated investors for better financial decisions.

Cultivating Exceptions

The classification of exceptions facilitates early discovery of an exception. Once an exception is detected, however, appropriate cultivation of the exception becomes crucial. Within each class of exception, one method is to have exceptions initially structured in terms of the likelihood of their occurrence and another is to raise exceptions by the order of significance. An exception often loses its initial significance and ultimately is dropped off although in an extreme situation, an exception grows so strong that it eventually replaces the default. The following equation can be used to figure out the confidence of betting on an exception by considering both occurrence and significance:

$$\text{Confidence} = \text{weight}_1 * \text{occurrence}_1 + \text{weight}_2 * \text{occurrence}_2$$

A linear structure is often applied to store the exceptions within the same class:



An exception, for instance, is determined as confidence (e_i): 0.6, derived from the equation (or $.5 * .4 + .5 *.8$). Other exceptions possessing lower confidence are also linked in the order of the confidence level. Weak confidence exceptions are not represented to the end-users due to their insignificance to the individual investors.

An exception must be associated with a reasoning mechanism with fact inference to justify the confidence of an exception, which requires minimal or no maintenance because neither an online broker nor an online investor would be able to maintain the mechanism. On the one hand, a brokerage firm does not manually keep track of thousands of stocks. As a matter of fact, even if the brokerage firm discovers many exceptions for major stocks, it is labor intensive to recommend those stocks based on individual investors' unique financial situations. On the other hand, an individual investor likely fails to grasp the kind of complexity with numerous factors in the financial market. Thus, the ability of unsupervised learning becomes a necessity to successfully cultivate exceptional valuations.

Two major reasoning mechanisms that possibly support the kind of machine learning are decision tree and neural networks. While each has advantages and disadvantages, we here only examine the former because it can $\text{\textcircled{t}}$ autonomously accommodate the growth of exceptional valuations of a security without losing simplicity.

A Simple, Feasible Schema Addressing Implementation Issues

The viability of an intelligent web user interface in this research lies in the feasibility of autonomous compilation of variations. Now, we introduce a simple, practical schema of decision trees to address specific details of the model. Since the confidence of an exceptional valuation can be derived from a decision tree that would grow by including more factors through observational learning, each factor to be included can be stored as a decision node, with each branch of a decision tree being forked out to express possible results influenced by the factor.

If several specific factors influence a stock price, in addition to factors influential to the entire industry, inheritance through default reasoning then would not reflect the impact of specific factors. Due to the potential significance of specific factors, the resulting exceptional valuations are represented in the order of the significance. However, the significance of each factor varies depending on the market and business conditions, i.e., certain factors weigh much heavier than others where the most significant ones constitute a different decision tree.

Table 2. An Action Table in Accordance With the Decision Table

Nevertheless, such a tabular structure provides only primitive data, but leaves the data digestion to the individual investors. Since this tabular structure is linear (rather than composite), fragmented (rather than granular), and segregating (rather than associative), it is often confusing to unsophisticated users and, consequently, is considered less intelligent. Figure 3, however, introduces two-dimensional a decision table with a clear preference and a coherent heuristic to average online users.




Figure 3. A Two-dimensional User Interface Presenting Decision Trees.

With an event function attached to each box in the figure, an online user can toggle horizontally to navigate and compare the results from combining different decision nodes. Dimension (a) can be used to guide a user in perceiving a better systematic assessment, whereas dimension (b) can be used to effectively show the results of aspect-focused choices. Both dimensions nonetheless are capable of representing an intelligent decision-tree structure without explicit rules.

To increase information granularity, a decision tree can be constructed at the portfolio level. Consider the following observations, for example:

1. 80% of successful professional investors would buy when the financial market is depressed.
2. 60% of professional investors would buy when a good company's stock sharply drops.
3. 65% of brokers would increase growth stocks when the financial market as a whole is depressed, or would decrease growth stocks after the financial market has advanced to a new high.

The corresponding decision tree can be constructed as follows:

1. Decision node 1: If the market condition worsens and a stock index drops by 5%, a buying decision of 80% probability would lead to a 5% gain above the average return. A selling decision would otherwise lead to a 5% loss.
2. Decision node 2: If a reputable Fortune 500 company drops by 5%, a buying decision of 60% probability would lead to a 5% gain in the long run. A selling decision would then lead to a 5% loss.
3. Decision node 3: If the market retreats by 5%, one should sell some value stocks and buy some growth stocks. As a result, the terminal wealth of his portfolio will increase by 5%.

The decision tree represents a deductive reasoning approach if the tree is developed at a brokerage firm level where a comprehensive database can be gradually cultivated through synthetic analysis of inductive data collection. The question then is when a decision tree can be embedded into an individual's web site to facilitate the decision-making process. The answer lies in what is considered statistical confidence. When a decision tree yields some statistical significance that indicates an exceptional situation with regard to a particular stock, the exception should be released along with the associated confidence. Taking a more deductive reasoning strategy, the decision-tree approach could only prevent common mistakes through the use of highly trained decision nodes. If the mistakes, such as buying as the market gone higher and selling as the market gone lower, are frequently observed, correction rules then can be first derived and then stored offline to prevent online investors from repeating the same mistakes.

To prevent uncompensated mistakes that a particular individual habitually makes, the decision nodes often must be restructured to fit the individual's objectives. One way to address the issue is to classify individual mistakes as a subset of common mistakes. Since major common mistakes are usually easier to identify because of broad recognition, personalized individual mistakes then can be singled out as a subset of the major common mistakes. Thus, this individual can be objectively advised to avoid these mistakes upon identifying the subset. However, since not all patterns of mistake are identifiable, continued pattern recognition in search of a perhaps very small subset of mistakes suitable to an individual's investment practice should heavily involve induction. By focusing on the mistakes of high risks, the resulting decision tree no longer just offers a general opinion.

To illustrate the concept, let us examine an extreme situation. Assume that a pattern of more buying trades on Mondays and more selling trades on Fridays is identified, even if the market does not necessarily reach low on Mondays and high on Fridays. If more people follow the pattern, the market tends to move in an opposite direction on these days. Therefore, an individual who trades randomly should be reminded to trade against the pattern in order to take the advantage. Accordingly, a decision node could be created to indicate the possible profit by leveraging these recognized mistakes.

5. CONCLUSIONS

Research Significance

The framework for an intelligent web user interface presented in this paper has suggested a novel approach to facilitating online shopping and to compensating for the deficiency of e-commerce. By infusing intelligence into web user interfaces, we are able to achieve personalized assistant comparable to that in the traditional offline shopping without having to constantly rely on a remote knowledge engine as often seen in most decision-support systems. The research has developed a web user interface model that integrates heuristic intelligence into an object-oriented framework in the online investment industry without losing generality, which can be generalized to facilitate online shoppers in other commercial marketplaces. This research also exhibits how the intelligence of a web user interface can evolve in order to accommodate the dynamic nature of the financial market.

Research Limitations

While the model presents a novel feasible approach to facilitating online shopping, limitations inherent to the decision-tree method do exist. First, the decision tree based web user interface model may only work effectively for a few decision factors, which in some cases may weaken the confidence of the intelligence embedded in a web user interface. As the number of decision factors increases, it may become too complex to be processed by software applets. Second, although it is reasonable to expect that the extensive intelligence of an online shop would help individuals reach a sound purchasing decision, this model nevertheless is not intended to help individuals achieve a superior performance due largely to its reliance on defaults. Instead, it is intended to provide online, heuristic intelligence for individual investors to avoid common mistakes.

REFERENCES

- Bauer, C. and Scharl, A. (1999). Acquisition and symbolic visualization of aggregated customer information for analyzing Web information systems. In R.H. Sprague, Jr. (ed.), *Proceedings of the Thirty-second Annual Hawai'i International Conference on System Sciences*. Maui, HI: IEEE Computer Society Press.
- Beck, T.J. and Bommel, M.R. (1999). Incremental Encoding of Multiple Inheritance Hierarchies. *Proceedings of the 8th International Conference on Information Knowledge Management*, November 2 - 6, 507-513, Kansas City, MO.
- Benetti, H., Beneventano, D., Bergamaschi, S., Guerra, F., & Vincini, M. (2002). An Information Integration Framework for E-commerce. *IEEE Intelligent Systems*, Jan/Feb 2002.
- Fano, A. and Kurth, S. (2003). Personal Choice Point: Helping Users Visualize What It Means to Buy a BMW. *Proceedings of the 2003 International Conference on Intelligent User Interfaces*, January, 46 – 52. Miami, Florida.
- Haubl, G. and Trifts, V. (2000). Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. *Marketing Science*, 19(1), Winter, 4-21.

- Kobsa, A., Koenemann, J., and Pohl, W. (2001). Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships. *The Knowledge Engineering Review*, 16(2), 111-155.
- Ladkin, P. (1986). Time Representation: A Taxonomy of Interval Relations. *Proceedings of AAA' 86*, 360-366. Los Altos, CA: Morgan Kaufmann.
- Payne, J.W., Bettman, J.R., and Luce, M.F. (1996). When time is money: Decision behavior under opportunity-cost time pressure. *Organizational Behavior and Human Decision Processes*, 66(2), 131-152.
- Reiter, R. (1985). On Reasoning by Default. In R. J. Brachman and H. J. Levesque (Eds.), *Readings in Knowledge Representation*, Los Altos, CA: Morgan Kaufmann.
- Schafer, J. B., Konstan, J. A., and Riedl, J. (2001). E-Commerce Recommendation Applications. *Data Mining and Knowledge Discovery*, 5(1-2), 115-153.
- Wang, H., Mylopoulos, J., and Liao, S. (2002). Intelligent Agents and Financial Risk Monitoring Systems. *Communications of the ACM*, March 2002, 83-88.
- Zviran, M. and Erlich, Z. (2003). Measuring IS User Satisfaction: Review and Implications. *Communications of the Association for Information Systems*, Vol. 12, 81-103.