Multimedia Explanations in IDEA Decision Support Systems

Giuseppe Carenini and **Johanna Moore**
Learning Research & Dev. Center and
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
[carenini, jmoore]@cs.pitt.edu

Abstract

In this paper, we present a new approach to support the decision of selecting one object out of a set of alternatives. As compared to previous approaches, the distinctive feature of our approach is that neither the user, nor the system need to build a model of user's preferences. Our proposal is to integrate a system for interactive data exploration and analysis with a multimedia explanation facility. The explanation facility supports the user in understanding unexpected aspects of the data. The explanation generation process is guided by a causal model of the domain that is automatically acquired by the system.

Introduction

With the rapid increase in the amount of on-line, up-to-date information, more and more people, ranging from professional public-policy decision makers to common people, will base their decisions on on-line sources. Thus, there is an increasing need for software systems that support interactive, information-intensive decision making for different user populations and different decision tasks. This work is a preliminary study of a possible role for automatic generation of multimedia [1] presentations in supporting interactive, information-intensive decision making.

In our investigation we focus on non-professional users making personal daily-life decision. The particular decision task we consider is the selection of an object (e.g., a house) out of a set of possible alternatives, when the selection is performed by evaluating objects with respect to their values for a set of attributes (e.g., house location, number of rooms).

This kind of decision problem has been normatively modeled in decision analysis by the multi-attribute utility theory (MAUT) (Clemen 1996). According to MAUT, when certain assumptions hold, the decision maker, to select among a set of competing alternatives, should follow a two step methodology. First, typically with the help of a decision analyst, the decision maker should build a well-defined utility model (i.e., a value function) that maps each alternative into an overall measure of its *value*. Next, the decision maker should apply the value function to each competing alternative, and select the alternative(s) that maximize this function. In the rest of the paper, we refer to this decision methodology as the MAUT-method.

Although following the MAUT-method, the decision maker is assured to find optimal solutions with respect to well-established principles of rationality, the application of MAUT may be problematic for certain users performing certain tasks. First, in some situations decision optimality is not the only outcome that matters. Especially in personal daily-life decisions, decision confidence and satisfaction may be very important factors. Second, for some decision problems a partial utility model can be sufficient, and eliciting a detailed model would be a waste of resources. Third, in some circumstances a decision analyst may not be available to help with model elicitation. Finally, in some situations what information should be used to make a decision (i.e., what are the alternatives and what are the relevant attributes) is not rigidly fixed. In these cases, a substantial part of the decision maker's task is to create new information or rearrange the information initially available to increase the quality and acceptability of the final decision.

To support the decision of selecting an object out of a set of alternatives when a strict MAUT-method could be problematic and to take advantage of the increasing amount of on-line information potentially useful for decision making, several software prototypes have been developed. (Jameson *et al.* 1995) provides a unified, abstract framework to analyze these prototypes and characterizes them as EOIPs (Evaluation-Oriented Information Provision systems). In an EOIP: "the user has the goal of making evaluative judgments about one or more objects; the system supplies the user with information to help her make these judgments". All EIOPs are mixed-initiative. As the interaction unfolds, the system tries to acquire a (partial) model of the user's value function and uses such a model to guide successive interactions with the user. For instance, in deciding when to suggest possible good alternatives to the user, or when and how to elicit new information from the user about her

This work was supported by grant number DAA-1593K0005 from the Advanced Research Projects Agency (ARPA). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of either ARPA or the U.S. Government.

[1] In this paper, by multimedia we mean a combination of text and information graphics.

preferences. For a discussion of an EIOP system publicly available see (Linden, Hanks, & Lesh 1997). This system is based on the candidate/critique framework, in which communication from the system is in the form of candidate solutions to the problem (possible alternatives), and communication from the user is in the form of critiques of those solutions.

In this paper, we discuss a different, but possibly complementary, approach to supporting object selection out of a set of alternatives. We assume that certain users, familiar with graphical displays of considerable complexity and direct manipulation techniques, may prefer to make their decision by autonomously exploring and analyzing the set of alternatives by means of visualization and interactive techniques. However, even such power-users may need system assistance when they encounter situations they do not expect or understand. The goal of our study is to investigate how effective assistance can be provided when these situations occur. We propose a multimedia explanation facility based on a causal model of the domain that is automatically acquired by the system from the same information originally available to the user. Users can invoke the explanation facility at any time during the interaction.

System Rationale and Architecture

Visualization and interactive techniques enable users to perceive relationships and manipulate datasets replacing more demanding cognitive operations with fewer and more efficient perceptual and motor operations (Casner 1991). Therefore, there is a growing body of research on developing Interactive Data Exploration and Analysis systems (IDEAs) in which visualization and interactive techniques play a prominent role (Selfridge, Srivastava, & Wilson 1996; Roth *et al.* 1997). In an IDEA system, users can typically query or visualize the information at different levels of detail, and in many different formats, to search for important relationships and patterns. Furthermore, by means of interactive techniques users can control the level of detail, eliminate part of the information, and also create new information by reorganizing, grouping and transforming the information originally available. To clarify, we briefly sketch a simplified IDEA system in the real estate domain. Such a system should at least allow the user to visualize information about houses for sale on a map. As exemplified in Figure 1(a), each house can be visualized as a graphical mark, and properties of each house can be shown by properties of the corresponding mark (e.g. in Figure 1(a) the price of a house is shown by the size of the corresponding mark). By examining such a visualization the user can easily perceive relationships in the dataset, for instance why and where some houses are more/less expensive the others. Furthermore, if the map was realistic and showed streets, parks and important buildings[2] the user could easily identify houses in desirable locations (e.g., close to a park, far from heavy traffic areas). The real estate IDEA system might also provide interactive techniques to facilitate further exploration of the dataset. For instance, a dynamic slider would allow the user to restrict the set of houses shown on the map to only houses whose price is in a certain range. As shown in Figure 2, a dynamic slider is an interactive technique that allows to specify any range of values for an attribute by sliding a drag box or its edges with the mouse.

An assumption of our study is that IDEAs represent appropriate decision making frameworks for computer- and graphics-literate users, at least when the decision is the selection of an object out of a large set of possible alternatives and the selection is performed evaluating objects with respect to their values for a set of attributes. Computer- and graphics-literate users, instead of going through the painstaking model elicitation process required by the MAUT-method, or through the system driven interaction typical of EIOPs, may prefer to autonomously explore the whole dataset of alternatives, deliberately selecting the data attributes to examine and the alternatives to rule out as the selection process unfolds. In other words, they may prefer to use the full power of visualization and interactive techniques to express their preferences and constraints and to verify how these constraints affect the space of acceptable alternatives.

However, in exploring and analyzing large datasets, users often encounter situations they do not expect or understand. There are two fundamental ways in which users can resolve these impasses. They can either produce tentative hypotheses and proceed by further exploring the data to verify such hypotheses, or, in cases where an artificial or human domain expert is available, they can ask the expert for an explanation.

The goal of this study is to design an explanation facility that can generate effective multimedia explanations for decision makers facing a situation they do not expect or understand.

In general, the design of an explanation facility requires one to address at least three issues:

- Coverage: what are the requests for explanation the intended user can pose?
- Knowledge sources: what knowledge is necessary to effectively answer these requests for explanation?
- Knowledge acquisition: how can knowledge sources be acquired?

Coverage: a typical puzzling situation for users interacting with IDEAs occurs when they find some unexpected, disproportionate, similarities or differences between some data values. Since in the daily-life decisions we are examining, price is usually the attribute users are most interested in understanding, we will use the price attribute to clarify our approach. In the real estate domain, for instance, these are sample requests for explanation that may arise: *Why do all these houses have similar prices?* or *Why are these houses much cheaper than those houses?*. Thus, the generic explanation request we assume for our explanation facility is: *Why do* $< subset - of - alternatives >^*$ *have* $< similar - or - different >$ *values for* $price$.

[2] This is not the case in Figure 1(a)

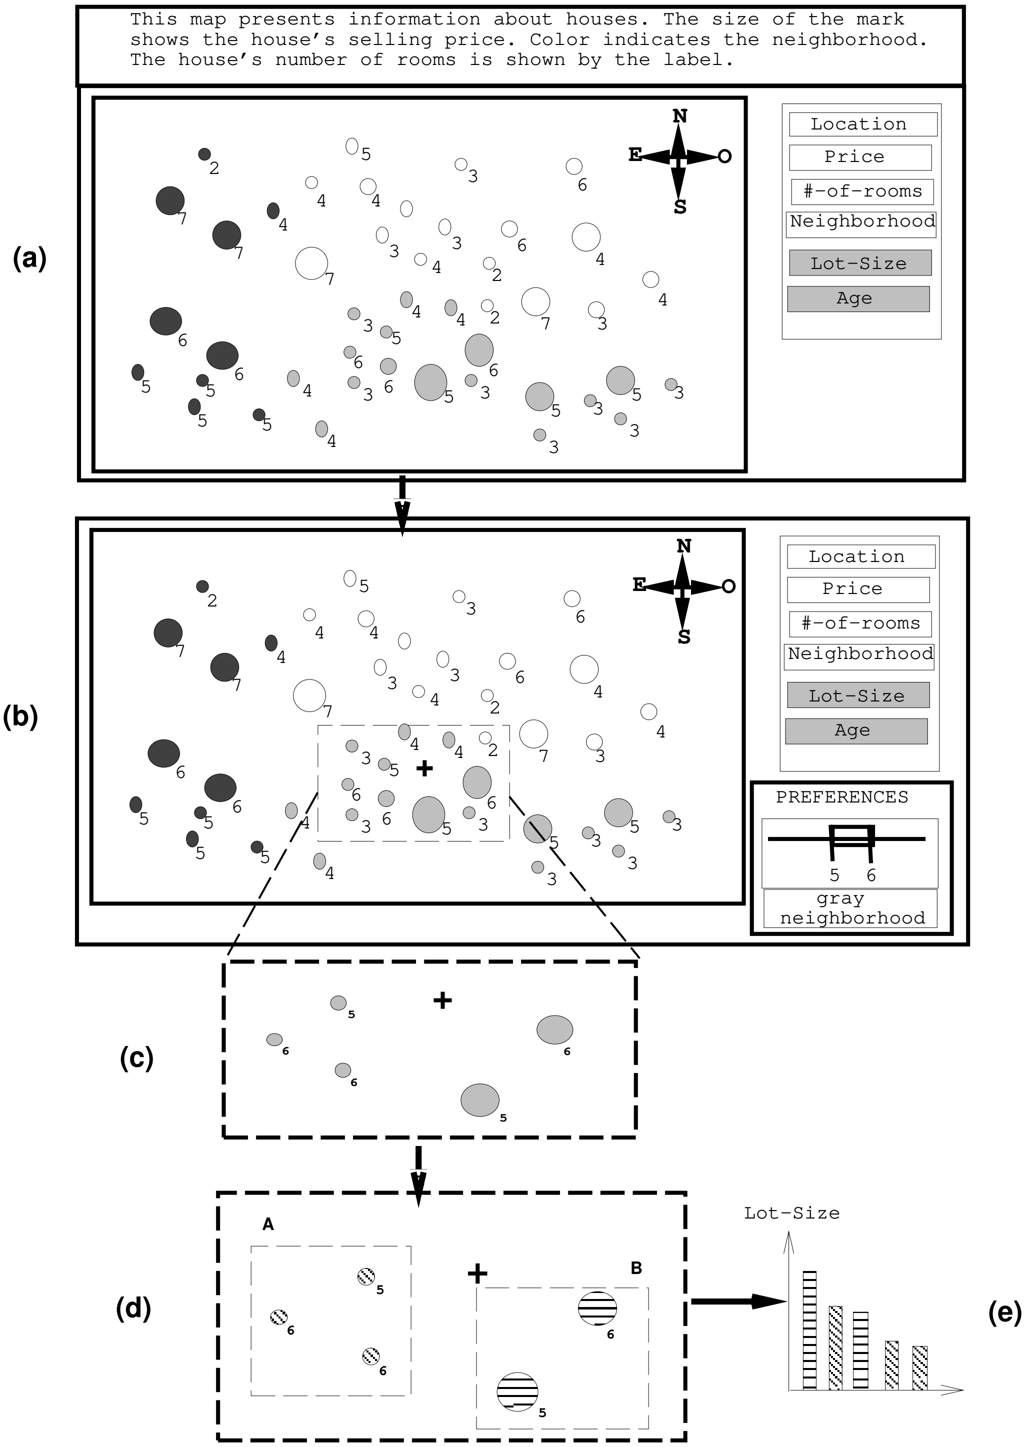

Figure 1: Sample IDEA system and sample Interaction.

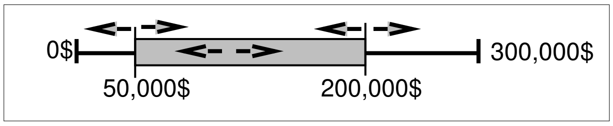

Figure 2: A dynamic slider: the user can specify any sub-range of the 0-300,000 price range by sliding the drag box or its edges in the directions indicated by the arrows.

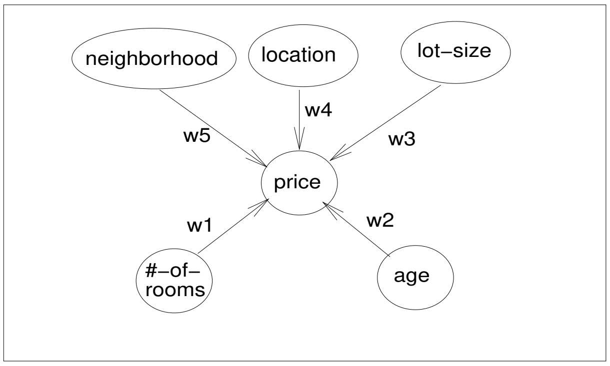

Figure 3: A minimal causal model for the price of a house.

Knowledge sources: Intuitively, a fundamental knowledge source for generating explanations about differences and similarities between attribute values of objects is a model of how some attributes values determine other attributes values, namely, a *causal model* of the domain. For example, when explanations are about the attribute price, what is needed is a causal model of how the attribute price is determined by all the other attributes. A simple possible causal model for price is a one level causal tree, in which each attribute independently contributes to price. The importance of each attribute in determining price can be specified by a weight, and price can be determined as a weighted sum of all other attributes. A graphical representation of a sample model for the price of a house is shown in Figure 3. Clearly such a model is just for illustration. A realistic causal model for the price of a house would be much more complex, including more attributes than the ones usually found in real-estate databases. For instance, it would include attributes of the neighborhood the house belongs to, such as the number of restaurants and bars, the crime-rate, and also the accessibility and quality of services, such as transportation, schools, hospitals etc.

Despite its structural similarity to a value tree, note that a causal model for the price of a house is not an utility model (a value tree) for the current user or for any users. As we said, it is a causal model of how the price of an object is determined, and consequently could be estimated and explained by other attributes of the object.

A second knowledge source that can contribute to generating fluent and coherent explanations is a *representation of the current state of the display* in terms of objects and their attributes (i.e., what and how information is currently shown).

Finally, the *history of the interaction* is also a critical knowledge source for explanation generation. However, in this study, we have not yet examined the role of this source. [3]

Knowledge source acquisition: A causal model has two main components: the qualitative causal graph that specifies causal influences between attributes, and a set of weights that quantify the importance of those influences. We assume that if the causal graph for the attribute price [4] is determined by domain experts, and the dataset is sufficiently large, the weights for the specific dataset can be inferred by using statistical methods of regression (Neter, Wasserman, & Kutner 1990), or machine learning techniques, such as neural networks (Bishop 1995).

For the acquisition of a representation of the current state of the display the explanation facility can only rely on the IDEA system. This is not a problem, because the IDEA system we intend to use, Visage (Roth *et al.* 1996), maintains by design links between the graphical objects displayed and the data objects they represent.

Let's now briefly go through the architecture of the system shown in Figure 4. The user's decision problem is to select an object from a large set of alternatives. To perform the selection, the user interacts with an IDEA system by means of information visualization and manipulation actions. Whenever the user encounters some unexpected similarities or differences between some objects for the attribute price, she can pose a request for explanation. To

[3] See (Carenini & Moore 1993) for a discussion of the role of interaction history in a system that generates textual explanations.

[4] Or for any attribute that can be reasonably expected to be determined by other attributes.

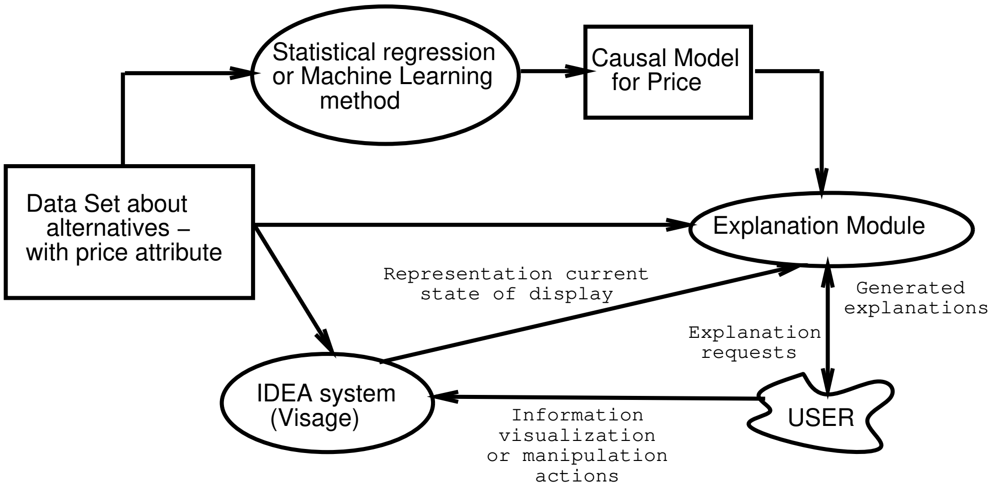

Figure 4: System architecture.

satisfy this request, the explanation module generates a multimedia explanation by consulting the original dataset, a causal model of how the price of an object is determined by other attributes of the object, and a representation of the current state of the display.

A final note on Figure 4. The distinction between the IDEA system and the Explanation Module is functional. In an actual implementation, the IDEA system and the Explanation Module would be seamlessly integrated from the user's perspective. The explanation requests would be built by direct manipulation actions and the generated explanations would be displayed and coordinated with the current state of the IDEA system display.

A Sample Interaction

The example we present is in the real estate domain. Let's assume the user is looking for a house to buy and she has particularly strong preferences for the location, the number of rooms, and the neighborhood. To get familiar with what alternatives are available, the user may ask the IDEA system to create a presentation containing all the houses for sale and showing the attributes she has strong preferences for, including price. The IDEA system generates the map shown in Figure 1(a), in which position shows the geographic location of the house, color shows the neighborhood, a numeric label indicates how many rooms it has, and price is shown by the size of the mark. Let's assume now, that the user has available a repertoire of manipulation actions to express her preferences and verify how they affect the space of acceptable alternatives. In Figure 1(b), the user expresses her preference for a location close to work (the work place is shown by the small cross) by selecting houses around the cross by means of a bounding box. She also expresses preference for houses with 5 or 6 rooms and for the neighborhood encoded by the color gray. The result of these selections is shown in Figure 1(c). We zoomed in on the result of the selections for illustrative purposes only, in interacting with an actual IDEA system the user can decide whether she does or does not want to keep the whole dataset displayed at different stages of the interaction. Going back to the example, let's assume now that examining the data in Figure 1(c) the user is puzzled by the substantial difference in price between the houses on the left of the cross (set A in Figure 1(d)) and the ones on the right (set B in Figure 1(d)). In the attempt to explain this difference, by means of manipulation actions the user creates the bar chart in Figure 1(e), which shows the lot size of the selected houses. This bar chart does not help much, differences in lot size do not convincingly explain the price differences under scrutiny. At this point the user, feeling that a more complex explanation may underlie the price differences, sends a request for explanation to the explanation module[5]. The request is: *"Why do $SetA$ $SetB$ have $different$ values for $price$?"*. Notice that this is an instantiation of the generic form presented in the previous section. Figure 5 shows the explanation generated by the explanation module to satisfy this request. The explanation combines text and graphics and is integrated with the previous state of the display. In the next section, we propose a model of how similar explanations can be generated.

Multimedia Explanation Generation

Multimedia explanation generation comprises the following steps:

- **Content Selection and Organization**: determining what information should be included in the explanations and how it should be organized.
- **Media Allocation**: deciding in what media different portions of the selected information should be realized.

[5]Notice that actual users may prefer to explore other hypotheses before resorting to the explanation module. This sample interaction is just for illustration.

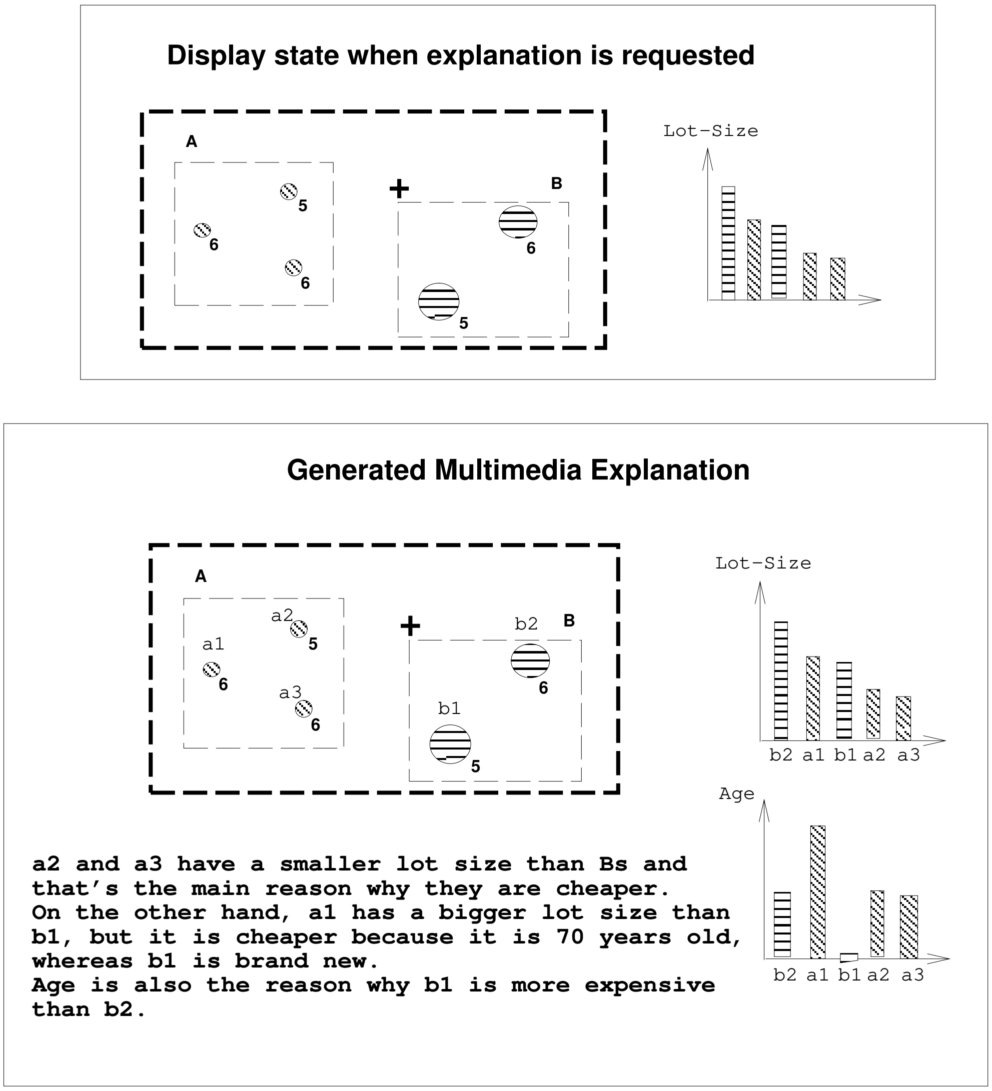

Figure 5: Sample Explanation.

- **Media Realization and Coordination**: realizing the information in the chosen media, ensuring that portions of the explanation realized in different media are coordinated.

Briefly, we argue that a causal model of the kind shown in Figure 3 can guide content selection and organization. Adapting techniques presented in (Klein & Shortliffe 1994), a causal model allows the explanation module to determine for any two objects which attribute contributes more/less to their difference or similarity with respect to a target attribute (price in our example). Depending on the strength of their contribution, attributes are, or are not, included in the explanation and different points can be made about those attributes. For instance (see Figure 5), on one hand the attribute lot size was kept in the explanation from the previous display, because it was the main reason why $a2$ and $a3$ were cheaper that the Bs; on the other hand, the attribute age was introduced, because it was the main reason why $b1$ was more expensive than $a1$. A measure of how attributes contribute to difference or similarity with respect to a target attribute can also guide how information is organized. For instance, $a2$ and $a3$ were compared as a group to the Bs, because, for both houses, lot size is the main reason they are cheaper.

Although careful media allocation and coordination are basic steps in generating multimedia explanation, in this study, we have not yet addressed these issues.

As far as media realization is concerned, for language, we are adapting techniques for generating arguments presented in (Elhadad 1995), whereas for graphics, we plan to use a tool, currently under development by our group, that designs graphics based on tasks that the users should be able to perform (e.g., search for certain objects, look-up certain values).

A final consideration. Since users, in interacting with an IDEA system, tend to produce highly populated displays, it is a very challenging task for the explanation module to maintain coordination between the new information and previously displayed information. Notice how, in the sample explanation in Figure 5, coordination is maintained by generating identifiers for the objects involved and associating them with the corresponding graphical objects. This is just a simple example, more study is necessary.

Conclusions and Future Work

In this paper, we present a new approach to support the decision of selecting one object out of a set of alternatives. With respect to previous approaches, like the MAUT-method and EIOPs, the main distinctive feature of our approach is that neither the user nor the system need to build a model of user's preferences.

Our proposal is to integrate an IDEA system with a multimedia explanation facility. On one hand, by interacting with an IDEA system, users can autonomously explore the dataset of alternatives and express their preferences and constraints. On the other hand, the explanation facility can help users understand puzzling situations they may encounter in exploring the data.

Although in the discussion we focused on the attribute price, the explanation facility we propose is quite general. No aspects of the facility are specific for price, any attribute can be explained as long as it is reasonable to assume that the attribute is causally determined by other attributes. When this is the case, the weights of the causal model can be automatically learned from the dataset, and the fully specified model can guide the explanation generation process.

Some readers might question the practical utility of the kind of explanations we presented. Although the only arbiter for utility is sound evaluation, we believe that utility of these explanations would be more evident if larger subsets were compared. In the sample explanation, we had to use two small subsets and we assumed a minimal causal model because of space limitations and to simplify the discussion of the main issues.

What information should be included in the multimedia explanations and how it should be organized is an area for further research. For instance, in certain situations, it might be beneficial to explicitly show the causal model (or a portion of it) to the user. The causal graph might provide a general framework in which more specific information about the subsets under investigation could be more effectively explained .

Finally, given the availability of different approaches to support selecting an object out of a set of alternatives (i.e MAUT-method, EOIPs and IDEA+explanations), two challenging research issues must be addressed. First, determining when one approach is more or less appropriate than another. Possible relevant factors in answering this question are, among others: the number of alternatives, the number of attributes, the ratio between these two numbers, user's computer skills, and user's familiarity with domain and decision task. The second research issue is about integration among different approaches. Either they could be applied in succession to the same problem and different outcomes compared and merged, or different approaches could be used in different phases of the same decision process. For instance, IDEA+explanations to reduce the number of alternatives, combined with the MAUT-method or EOIPs to select one object out of the reduced set.

Acknowledgments

We would like to thank N. Green, S. Kerpedjiev and S. Roth for all the useful discussions on the ideas presented in this paper. We also thank an anonymous reviewer for useful comments.

References

Bishop, C. M. 1995. *Neural Networks for Pattern Recognition*. Oxford: Oxford University Press.

Carenini, G., and Moore, J. D. 1993. Generating explanations in context. In Gray, W. D.; Hefley, W. E.; and Murray, D., eds., *Proceedings of the International Workshop on Intelligent User Interfaces*, 175–182. Orlando, Florida: ACM Press.

Casner, S. 1991. A task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics* 10(2):111–151.

Clemen, R. 1996. *Making Hard Decisions: an introduction to decision analysis*. Belmont, California: Duxbury Press. Second Edition.

Elhadad, M. 1995. Using argumentation in text generation. *Journal of Pragmatics* 24:189–220.

Jameson, A.; Schafer, R.; Simons, J.; and Weis, T. 1995. Adaptive provision of evaluation-oriented information: Tasks and techniques. In *Proceedings of 14th IJCAI*.

Klein, D. A., and Shortliffe, E. H. 1994. A framework for explaining decision-theoretic advice. *Artificial Intelligence* 67:201–243.

Linden, G.; Hanks, S.; and Lesh, N. 1997. Interactive Assessment of User Preference Models: The Automated Travel Assistant. In *Proceedings of Sixth International Conference of User Modeling*.

Neter, J.; Wasserman, W.; and Kutner, M. 1990. *Applied Linear Statistical Models*. Irwin.

Roth, S. F.; Lucas, P.; Senn, J. A.; Gomberg, C. C.; Burks, M. B.; Stroffolino, P. J.; Kolojejchick, J. A.; and Dunmire, C. 1996. Visage: A user interface environment for exploring information. In *Proceedings of Information Visualization*, 3–12. San Francisco: IEEE.

Roth, S. F.; Chuah, M. C.; S. Kerpedjiev, S.; Kolojejchick, J. A.; and Lucas, P. 1997. Towards an information visualization workspace: Combining multiple means of expression. *Human-Computer Interaction Journal*.

Selfridge, P.; Srivastava, D.; and Wilson, L. 1996. IDEA: Interactive data exploration and analysis. In *Proceedings of SIGMOD-96*, 24–34.