

Tailored Responses for Decision Support

Terrence Harvey, Sandra Carberry, and Keith Decker

University of Delaware, Newark DE 19716 USA
{harvey,carberry,decker}@cis.udel.edu

Abstract. Individuals differ in the resources that they are willing to expend on information gathering and on the importance of different kinds of information. We have developed MADSUM, a system that takes into account user constraints on resources, the significance of propositions that might be included in a response, and the user's priorities with respect to resource and content attributes; MADSUM produces a response tailored to individual users in a decision support setting.

1 Introduction

We have been investigating the design of a decision support system that can adapt to a user's resource constraints, resource priorities, and content priorities in a dynamic environment. Our approach employs a multi-attribute utility function as part of a user model. The utility function weighs the benefit of different decisions about resource usage and information selection. Our approach provides a structure in which the priorities of the user can be explicitly represented and considered in light of the environment (information currently available, the cost of getting the information, etc.) Furthermore, an agent architecture allows the system to dynamically respond to changes in the environment or user priorities.

We have applied the MADSUM architecture to decision-support in a financial investment domain, where the system must support a user in making a buy/don't-buy decision on a single investment. The MADSUM decision making algorithms and the agent hierarchy, communications, and interaction are domain-independent. Furthermore, MADSUM is easily extendible to new domains with different attributes in its utility function. However, implementation in a particular domain requires a set of domain-dependent information agents. For example, tailored decision-support in an investment domain requires domain-dependent agents that can estimate how significant a particular piece of information will be to the current user, given her current personal and financial status.

2 The User Model

The MADSUM user model has three components: User Attributes, Constraints, and Utility Function. The User Attributes component of the user model captures characteristics of the user, including appropriate domain-specific information. For the financial investment domain, this component of the user model includes the user's age, salary, expected number of years to retirement, approximate annual expenditures, current investment portfolio, and portfolio allocation goals. The User Attributes affect the significance of certain pieces of information.

The Constraints component of the user model offers the user the option of setting both soft and hard constraints for a given attribute. Soft constraints are attribute values that the user would prefer not be exceeded in constructing a response. These soft constraints affect the utility function, as described below. Hard constraints are values that an attribute *must not* exceed in a response, and are used to pare the search space before utility is calculated.

MADSUM’s utility function contains n attribute terms, each consisting of a weight w_i giving the importance of that attribute to the user, a parameter a_{value_i} that is related to the value of the attribute, and a function f_i .

$$Utility = \sum_{i=1}^n w_i f_i(a_{value_i})$$

The weights w_i , giving the importance of each attribute to the user, are extracted from the positions of sliders that are manipulated by the user in a graphical user interface. For resource attributes such as length of response or processing time, a_{value_i} is the actual value of the attribute, such as 75 words. On the other hand, information attributes capture propositions that might be presented to the user, and thus for information attributes, a_{value_i} captures the significance of a set of propositions in the environment of the user’s personal characteristics and the application domain. We call this approximation *Decision Specificity* or *DS*. Determining DS is a domain-specific task, and thus in the MADSUM architecture, the functions that compute DS are provided by the application designer as part of the domain-specific information agents that propose propositions for inclusion in the response to the user.

In the financial investment domain, we have implemented domain-specific information agents for three categories of information: Risk (the riskiness of an investment), Value (the prospects for the investment gaining in value), and Portfolio (how the investment relates to the individual’s portfolio allocation goals). For example, the significance of a proposition from the Portfolio Agent that addresses the relationship of a proposed investment to the user’s portfolio allocation goals depends on the extent that the investment would cause the user’s portfolio allocation to deviate from his goals, while a proposition that addresses the appropriateness of the investment from an age perspective may depend on how close the user is to retirement.

Each of the functions f_i that appear in the utility function map their parameter a_{value_i} into a utility value between 0 and 1. The particular function f_i that is used determines whether an increasing parameter value increases or decreases utility (and at what rate). For example, $f_{StartPlateauNorm}$ captures instances in which utility remains high over a plateau and then decreases for increasing values of its parameter, and $f_{LinearPlateauEnd}$ captures instances in which utility increases linearly for increasing values of its parameter until a plateau is reached. Our financial investment domain by default uses $f_{StartPlateauNorm}$ for resource attributes and $f_{LinearPlateauEnd}$ for information attributes, but advanced users can select from among MADSUM’s full set of predefined utility functions.

The soft constraints entered by the user adapt the utility function f_i by determining its shape. For example, the function $f_{StartPlateauNorm}$ is used by default for the resource attribute of processing time; the soft limit determines

where the plateau ends and also the rate of fall in utility after the plateau (the falling portion resembles a normal distribution whose spread is 1/2 the soft limit). This captures the notions that 1) the soft limit on processing time set by the user is the point at which the utility of the response will begin to decrease and 2) the larger the soft limit on processing time, the less severe will be the loss of utility for each second of increased processing time.

3 Agent Architecture

To address the issues of collecting and integrating information from distributed sources into a single text plan, MADSUM is implemented as a hierarchical network of independent agents. The agents bid to provide information for the response; once a highest utility set of bids is selected, the lowest level information agents pass raw information to their parent middle agents, who use the information to generate small text plan trees. As the trees are propagated further up the agent hierarchy, the middle agents assemble them using coherence rules; in doing so, the middle agents first order the text plan trees according to the utility of their highest utility proposition, and the rules for combining trees attempt to assemble larger trees with the higher ranked constituents on the left, so that the higher ranked constituents will appear earlier in the response (subject to coherence constraints). Once an assembled tree is returned to the top-level Presentation agent, it is resolved to text via templates, and the text is presented to the user. A full description of the agent architecture can be found in [1].

4 Adaptive Responses, Implementation, and Evaluation

The MADSUM architecture for adaptive response generation has been implemented and tested in a financial investment domain. The GUI interface of sliders for setting priorities is not yet part of the implemented system, so priority settings are currently entered as numbers. Formal evaluation experiments have validated MADSUM's design, such as MADSUM's strategy of balancing significance and priority in content selection and MADSUM's decisions about order of presentation of propositions.

As examples of responses produced by our system, consider a user who proposes the purchase of 100 shares of stock in IBM. The user model contains personal characteristics of the user, including her current investment portfolio and her portfolio allocation goals. Before proposing the stock purchase, the user has set soft constraints on the length of the response, the cost in dollars of any purchased information, and processing time. She has also indicated the importance she assigns to usage of different resources (length of response, cost, and processing time) and her interest in information that addresses each of the different content categories (investment risk, value, and impact on portfolio allocation).

Figure 1 displays our implemented system's response under different soft constraint and priority settings. In Figure 1a, the soft constraint on length was 75 words and the user placed a higher priority on risk information than on value

1a: Risk metrics indicate IBM has a low debt-equity ratio, suggesting the ability to weather an economic downturn; further, the company has a strong current ratio, indicating good short-term liquidity. In addition, IBM has historically maintained a moderate debt policy, and the stock has maintained a moderate risk profile. On the other hand, from a portfolio perspective you have already exceeded your allocation goal for equities. Value metrics indicate IBM has a price earnings ratio similar to the tech industry average.

1b: Risk metrics indicate IBM has a low debt-equity ratio, suggesting the ability to weather an economic downturn; further, the company has a strong current ratio, indicating good short-term liquidity. On the other hand, from a portfolio perspective you have already exceeded your allocation goal for equities.

1c: Value metrics indicate the stock has a price earnings ratio similar to the tech industry average; on the other hand, from a portfolio perspective you have already exceeded your allocation goal for equities.

Fig. 1. Three responses, derived from different soft constraints and priority settings.

and portfolio information. For the responses in Figure 1b and Figure 1c, the soft constraint on length was lowered to 35 words, resulting in the exclusion of some available propositions; the relative priorities on risk, value, and portfolio information were kept the same in Figures 1a and 1b, but were altered in Figure 1c to place a much higher priority on value information than on risk or portfolio information. Due to the 35 word soft constraint on length that was set for the response in Figure 1b, propositions had to be excluded. Since risk was given highest priority, much (but not all) of the risk information was included. However, the high significance of the proposition about the proposed investment’s impact on the user’s portfolio allocation goals (she had already exceeded her goals for equities such as IBM) caused that proposition to increase the estimated overall utility of a response containing this proposition, and thus it was included despite the length of the resulting response slightly exceeding the soft constraint on length. In Figure 1c, the user’s much higher priority for value information resulted in selection of the value proposition, even though it was of lesser significance than other available propositions. In addition, the highly significant proposition about portfolio allocation goals was included in the response. These examples illustrate the system’s ability to vary its responses depending on the user’s resource constraints, the significance of information, and the priority that the user assigns to different resources and kinds of information content.

5 Related Work

Adaptive systems have been using concepts of utility theory, either informally or formally, to tailor responses so that they take into account the user’s preferences. In previous work[2], we used a weighted additive function to reason on a model of

user preferences to detect suboptimal solutions and suggest better alternatives during a collaborative dialogue. [3] uses a similar model in ranking candidate flights in a travel domain. The MAUT Machine[4] uses a formal utility function to evaluate products in an electronic catalogue. Moore[5] and Walker[6] use a formal utility function to rank travel and restaurant options respectively, but then other mechanisms are used to identify the actual propositions that are included in the natural language response. While other systems measure the utilities of possible domain outcomes (e.g. the utilities of one purchased item vs. another) and then tailor a message accordingly, MADSUM is the first system to use a formal utility function to evaluate the utility of the message itself. This message utility includes not only a component related to the noteworthiness of a domain outcome, but also message-specific components such as length and time.

6 Conclusion

MADSUM was designed to exploit a user model in the generation of responses that are tailored to the individual user of a decision support system. MADSUM takes into account user constraints on resources, the significance of propositions that might be included in the response, and the user's priorities with respect to resource and content attributes. The output of MADSUM is a tailored response that has the highest estimated utility for the particular user. The MADSUM architecture has been implemented and tested in a financial investment domain, and the system's balancing of significance and priority in content selection and presentation order has been validated at a statistically significant level.

References

1. Harvey, T., Decker, K., Carberry, S.: Multiagent decision support via user modeling. In: Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi Agent Systems. (2005) To appear.
2. Elzer, S., Chu-Carroll, J., Carberry, S.: Recognizing and utilizing user preferences in collaborative consultation dialogues. In: Proc. of the Fourth International Conference on User Modeling. (1994) 19–24
3. Linden, G., Hanks, S., Lesh, N.: Interactive assessment of user preference models: The automated travel assistant. In: Proc. of the Sixth International Conference on User Modeling. (1997) 67–78
4. Schmitt, C., Dengler, D., Bauer, M.: Multivariate preference models and decision making with the maud machine. In: Proc. of the Ninth International Conference on User Modeling. (2003) 297–302
5. Moore, J., Foster, M.E., Lemon, O., White, M.: Generating tailored, comparative descriptions in spoken dialogue. In: Proc. of the International Florida Artificial Intelligence Research Society. (2004)
6. Walker, M., Whittaker, S., Stent, A., Maloor, P., Moore, J., Johnston, M., Vasireddy, G.: User tailored generation in the match multimodal dialogue system. *Cognitive Science* **28** (2004)