Adaptive Medical Information Delivery Combining User, Task and Situation Models

Luis Francisco-Revilla and Frank M. Shipman III
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, USA
+1 409 845 4924
{l0f0954, shipman}@cs.tamu.edu

ABSTRACT

Medical information delivery for users with different levels of expertise will be required for the manned mission to Mars due to limited potential for communication with Earth. The Mars Medical Assistant (MMA) uses a combination of user, situation, and task models to create virtual hypertext structures by piecing together medical "information components." Information components are chosen based on the semantic content and the cognitive characteristics of the component's media type. The medical assistant currently supports three tasks: 1) describing medical procedures, 2) aiding diagnosis, and 3) providing information on health concerns. Conflicting suggestions from the three models need to be resolved. Tradeoffs in the model representations and conflict resolution strategies are being explored in the context of MMA.

Keywords

Medical information systems, adaptive hypertext, user models, task models, situation models, conflict resolution

INTRODUCTION

Medical Informatics has largely focussed on providing information to medical professionals to improve efficiency and quality of care. There is also the need to develop systems supporting care providers who are not medical experts. We are designing a medical information system for one such situation – the human exploration of Mars planned for 2014. Due to the long-term self-sufficiency required of the crew and the limited potential for real-time advice from experts, the crew's medical information system must support different users with different levels of proficiency.

We are studying the use of adaptive hypermedia in order to provide access to a large amount of information facilitating the medical task without overwhelming the medic. In the following sections we present the approach used, the combination of user, task and situation models, the conflict resolution strategy, the implementation of a first prototype Mars Medical Assistant (MMA), and initial conclusions.

APPROACH

MMA is designed to assume the role of an "assistant", aiding the medic in performing the medical task by providing information and suggestions as needed. Our approach is influenced by observations about the situated nature of action limiting the usefulness of static help systems [10]. This has led to a combination of user, task, and situation models being used to determine the appropriate medical information for the current context. Based on the content of the three models, virtual structures and hypermedia composites [5] are created over a corpus of medical information. To provide hypermedia medical information in an appropriate manner, the system adapts the content of the information (level of detail, semantic type, etc.) as well as the particular presentation of it (media type).

Dynamic Hypertext

One shortcoming of traditional hypermedia is its static nature. This has led to the development of different approaches to increase the dynamic capabilities of hypermedia. For instance, PHIDIAS [8] uses a Dynamic Hypermedia approach to support design by creating virtual structures over a fine-grained hypermedia document.

Dynamic hypermedia provided a more flexible approach for presenting hypermedia documents. However, there was little discussion of alternatives for how dynamic behavior should be provided. Recently developed frameworks identify different dynamic types of hypermedia.

The adaptive hypermedia community modifies (adapts) the underlying hypermedia document based on a user model [1]. De Bra et al., in their model for adaptive hypermedia AHAM [2], distinguish content-adaptation from link-adaptation. Brusilovsky makes a similar distinction using the terms "adaptive presentation" and "adaptive navigation".

Adaptation in MMA

MMA provides an eclectic approach with roots in the dynamic hypermedia and adaptive hypermedia approaches.

Similar to PHIDIAS, it supports the creation of virtual structures out of information chunks but does so using explicit models of the user, as well as the task and situation. The result combines content-adaptation and link-adaptation.

In order to construct the aforementioned virtual structures, the medical information is divided up and categorized based on the semantic content and the cognitive characteristics of the media type. The resulting "information components" or ICs are reassembled into hypertext nodes in an approach similar to PUSH [6]. Semantic categorization allows for information selection based on the task and the medic's expertise – e.g., basic procedural information may be left out for a practicing physician but presented for other members of the crew. Media-based categorization enables selection based on the cognitive demands of the current situation and task – e.g., a diagram or video may be better when performing a task while textual descriptions may be better for learning conceptual information.

Much prior work on adaptive systems has emphasized sophisticated user or task modeling and inferences based on these models [3]. This work has focused on a particular media type, typically hypertext (text with navigational links), to allow the investigation of abstract characteristics of the adaptation mechanism. Rather than emphasizing a single model or adaptation effect, MMA has been designed to explore the complexity that naturally occurs when constraints from a variety of sources interact.

The complexity of the interactions results in a variety of types of adaptation by MMA. The current adaptability of the system determines:

- whether to initiate the dialogue or wait for the user,
- which high-level interface layout to present,
- what information to present,
- how to present the information (media and order), and
- how much information to make visible by default.

MMA uses the division of content into ICs to present the information most likely to be useful while also providing access to additional information through selection and navigation, similar to PUSH [6]. In addition, MMA's use of ICs provides support for managing and creating flexible virtual structures based on metadata particular to each IC identifying its *semantic type* and its *media type*.

Semantic types provide an indication of the questions about the topics an IC might answer. Currently we have explored the use of "explanations" and "descriptions" as semantic types. An IC is classified as an "explanation" when it addresses the question "why". In contrast, an IC is classified as a "description" when it addresses the questions "what" or "how". This initial distinction could be extended to a more complete set of semantic types based on the classes of questions answered, such as that included in MediaDoc for encoding software documentation [4].

The use of the media types allows the system to select information to be presented according to the affordances of each media type. The division of the information into ICs drives the creation and population of the database, where each IC is classified into a two dimensional matrix accordingly to its media type and to its semantic type.

USER, TASK AND SITUATION MODELS

Most adaptive hypermedia systems base their inference mechanisms on a single source of information about the context of use, typically a user model [1, 2]. This approach has been used successfully by many systems. Despite this, there is a drawback, namely, the complexity of the model tends to increase as more considerations appear to be relevant. Depending on the case, the system may require taking into account considerations that are not properly related to the person using the system, since the same user may present different requirements in different cases. We found this true as our approach consisted of designing a system to be used by different people (*users*) trying to accomplish something in particular (a *task*) within a given context (*situation*). This prompted the use of three distinct models for users, tasks and situations. All three models provide suggestions as to what and how information should be presented, but each model may present a different point of view based on the factors or actors taken into account.

The models being used are based on stereotypes [9] of users, tasks and situations. Stereotypes attempt to predict facts about the users, tasks and situations, based on observed co-occurrence of user's, task's and situation's characteristics. Currently, a small number of stereotypes have been implemented for each model. Users can be classified as novice, regular, or expert in the different topics in the medical database. Tasks supported are procedural guidance, diagnosis, and information browsing. Situations are classified as emergency, consultation, and educational.

MMA's user model includes information about the user's knowledge about each medical topic, knowledge about the system, interface/system preferences, and medical profile (because they are also the potential patients). Due to the small number of users within MMA's context of use, it is safe to assume a fine-grained model for each individual.

The characteristics of the specific task being performed that influence what information will be most useful are represented in MMA's task model. Stereotypes of tasks can help determine the semantic type of ICs selected for presentation. For instance, in the case that the user is browsing through the medical information, it makes sense to have all available information visible, providing the user more freedom for exploration. In the case of a procedural guidance, the user is expecting to find direction on how to perform some action, thus it makes sense to expand, at least initially, just the descriptive information on the topic.

While it is true that different people can react in a different manner to the same situation, it is also true that a given situation presents some constant requirements independent of the user. In the situation model, the requirements for content and presentation are clustered according to different situation stereotypes. For instance, emergency situations require the medical task to be performed promptly and precisely, without concern for the medic's long-term retention of knowledge. Thus information that expedites and enables performing the task should be presented in the media that allows rapid comprehension with little effort. This is in sharp contrast to an educational setting.

As of today, the majority of adaptive hypermedia systems have been educational applications [1, 2]. Even though other adaptive hypermedia systems have been designed for crisis situations [7], they are still designed to work within a single situation context. MMA differs in the sense that it is designed to be used in a variety of situations, and this led to the use of potentially conflicting models for deciding what information to provide. The next section describes how MMA resolves such conflicts.

CONFLICT RESOLUTION

The adaptation mechanism is divided into two levels, instantiated by the *content service* and the *presentation service*. The content service determines what information is required, specifically the medical concept and the semantic type. The presentation service determines the variables of how this information should be presented: the selection of appropriate media and composition of the resulting page.

Figure 1 shows the logical interaction between services. Each service resolves locally the possible conflicts arising from the suggestions by the different models. Then conflicts between the two services are solved by negotiation between the content and presentation service and the result is returned to the interface. The suggestion inference and the conflict resolution are rule-based mechanisms.

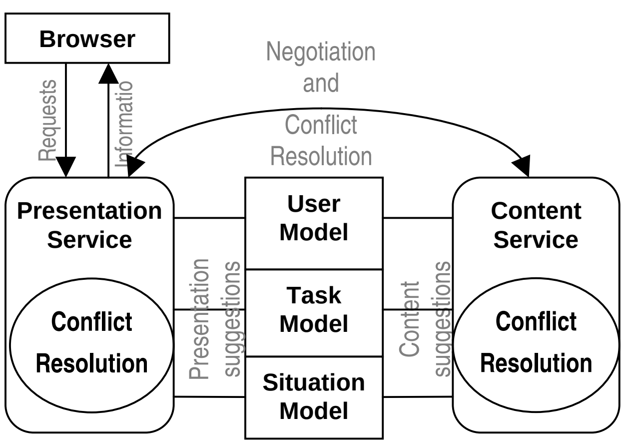

Figure 1. Interaction between models and services

In MMA both the content service and presentation service interact with the different models when making their inferences. This is indicated by the suggestion links from the models to the services in Figure 1. Within each service, conflicts can occur due to differences in suggestions from different models. For example, consider a user looking for information about decompression illness in an emergency situation. The task model suggests text as the media type and explanations as the semantic type. These suggestions attempt to maximize knowledge transfer and retention. The user model indicates pictures are the favorite media, then text, and finally audio and video. The situation model suggests video or audio as they expedite the performance of the task, since time is a limited resource and it is likely the medic has to attend to the patient at the same time. Clearly, there is the need for a conflict resolution mechanism.

Rules can be used to choose between conflicting suggestions. In the case above, since patient safety is paramount, emergency considerations will be emphasized over the medic's preferences. Once conflicts between the suggestions from the three models are resolved, ICs matching the decisions will be selected. At this point additional conflicts may occur due to the lack of ICs for a particular topic with the desired semantic and media types, resulting in negotiations between the presentation and content services to find alternatives given available ICs.

The next section describes the prototype system being used to explore the design space and identify limitations, tradeoffs, and new requirements for such a complex information delivery process.

PROTOTYPE

The MMA prototype is implemented as a Web-based medical assistant that supports for three tasks, 1) describing medical procedures, 2) aiding diagnosis, and 3) providing information on health concerns.

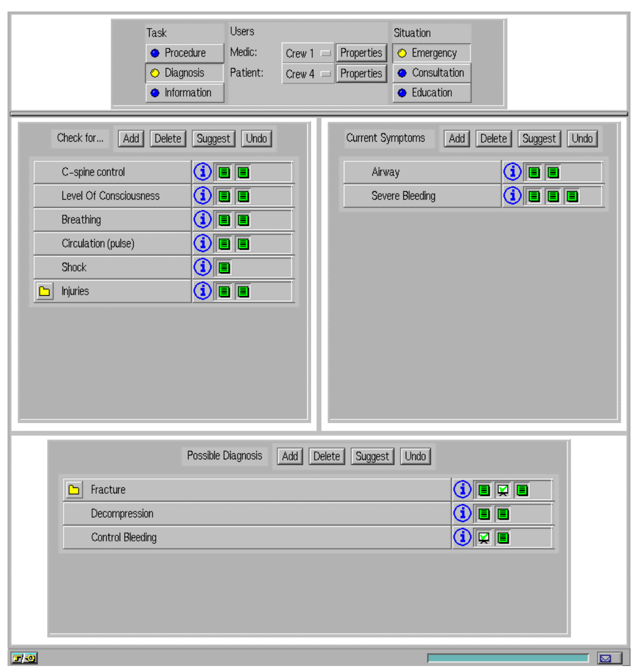

Figure 2. Diagnosis task

Figure 2 shows the system while aiding diagnosis. At the top of the screen is the control area where the user provides information about the current situation. Below this are the list of symptoms to check for (on the left), and the list of symptoms already confirmed (on the right.) At the bottom is a list of the most likely diagnoses based on the current symptoms. Each item (symptom and possible diagnoses) is represented by a concept bar as shown in Figure 3. The medic can start an action or get more information about a symptom or diagnosis by using the concept bar.

The concept bar includes a folder icon on the left enabling the expansion and contraction of hierarchically nested sub-concepts. The label allows the initialization of an action associated with the concept, such as initializing a dialog or collecting patient data. The right side of the concept bar is the information button. The number and types of ICs available for that concept and the order of their presentation in the hypertext composite node is indicated by the icons on the button. By selecting or deselecting the ICs the medic can manipulate the contents of the information page.

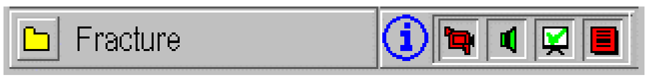

Figure 3. Concept bar

Once a diagnosis has been made, the task of the medic may switch to performing a medical procedure. In this case, a list of concept bars is used for access to information about medical procedures as shown in Figure 4. The list on the left displays a sequential view of the steps and sub-steps in the procedure. The right side of Figure 4 shows the information associated with the selected step. Similarly, all symptoms, diagnoses, and procedures have associated information that can be displayed when the user requests it.

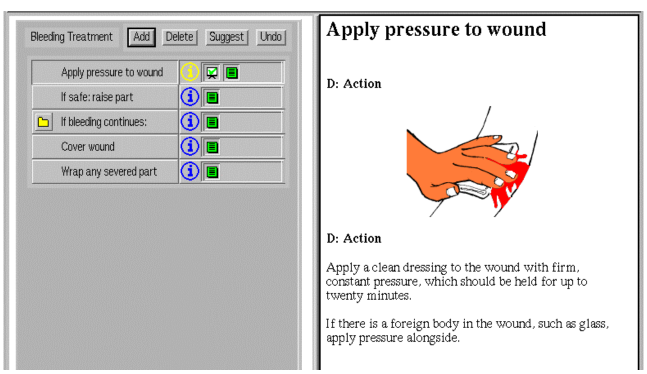

Figure 4. Procedure guidance task

The final task currently supported by MMA is browsing information within the medical database. Concept bars are used to provide a hierarchical view of the medical database. The coloring of ICs, not visible here, provides annotation for the links indicating likely relevance.

The adaptation described in the prior section involves determining which topic bars are displayed and the order and visibility of ICs on the topic bar.

DISCUSION

Applying this approach requires authoring of the hypermedia ICs. Authors of the hypermedia database must be aware of the varying nature of the final presentation. Populating the information space requires significant effort, both to collect and to categorize the ICs

An area for future research is improving the situation awareness of MMA. Since the crew will be heavily monitored, it will be possible for the system to obtain medical data directly with little overhead for the medic. Use of such information requires additional design considering the interaction style, the intrusiveness of the system, and inferred changes in the situation.

CONCLUSIONS

This project looks at the use of adaptive hypermedia to support medical information delivery to users with varying levels of expertise. The combination of user, situation, and task models are used to tailor the information display for diagnosis, procedural, and informational tasks. In order to concurrently use different models it is necessary to solve conflicts resulting from clashing suggestions.

MMA's hypertext nodes are composed by combining semantic-based and media-based information components. These composite nodes can be modified by the user. This project shows the potential for applying models characterizing users, their tasks, and their environment in domain-oriented information delivery.

ACKNOWLEDGMENTS

This work is supported in part by NSF grant IIS-9734167.

REFERENCES

1. Brusilovsky, P. Methods and Techniques of Adaptive Hypermedia, *Journal of User Modeling and User-Adapted Interaction*, 6, 2-3 (1996), 87-129.
2. De Bra, P., Houben, G.-J., Wu, H., AHAM: A Dexter-based Reference Model for Adaptive Hypermedia, *Proc. of Hypertext'99*, 147-156.
3. Encarnação, L., Multi-Level User Support through Adaptive Hypermedia: a highly application-independent help component, *Proc of Intelligent User Interfaces 1997*, 187-194.
4. Erdem, A., Johnson, W. L., Marsella, S. Task Oriented Software Understanding, *Proc. of ASE 1998*, 230-239.
5. Halasz, F. Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems, *Communications of the ACM*, 31, 7, 836-852.
6. Höök, K. Evaluating the Utility and Usability of an Adaptive Hypermedia System, *Proc. of Intelligent User Interfaces 1997*, 179-186.
7. Iba, W., Gervasio, M. Adapting to User Preferences in Crisis Response, *Proc of Intelligent User Interfaces 1999*, 87-90.
8. McCall, R., Bennet, P., d'Oronzio, P., Ostwald, J., Shipman, F., Wallace, N. PHIDIAS: Integrating CAD Graphics into Dynamic Hypertext, *Proc. of ECHT'90*, 152-165.
9. Rich, E. Stereotypes and User Modeling, in *Kobsa, A. and Wahlster, W., (Eds.) User Models in Dialog Systems*, pp. 35-51, Springer-Verlag, New York, 1989.
10. Suchman, L. A. *Plans and Situated Actions*. Cambridge University Press, New York, 1987