Empirically Based Decision-Theoretic Methods for Situated Interaction

Anthony Jameson*
Department of Computer Science, University of Saarbrücken
P.O. Box 151150, D-66041 Saarbrücken, Federal Republic of Germany
`jameson@cs.uni-sb.de`, http://www.cs.uni-sb.de/users/jameson/
Fax: +49 681 302-4136, Phone: +49 681 302-2474

1 Introduction

1.1 Main Claims

This presentation will argue for three claims about research on situated interaction in ubiquitous computing:

1. It is sometimes necessary or useful to obtain information about the situation of the user ($\mathcal{U}$) by interpreting features of $\mathcal{U}$'s behavior that indirectly reflect $\mathcal{U}$'s situation.
2. It is often useful to conduct thorough empirical research at an early stage in the design of such systems, rather than at a later stage of evaluation.
3. Decision-theoretic methods from artificial intelligence are in many ways well suited to the task of making the inferences and decisions required if a system ($\mathcal{S}$) is to adapt appropriately to the user's situation.

The arguments will be illustrated with examples from research in the project READY that has been conducted since 1996, including some previously unpublished results.

1.2 Scenarios

In the first READY scenario, $\mathcal{U}$ was an automobile driver whose car needed some minor repair while $\mathcal{U}$ was on the road. $\mathcal{U}$ obtained assistance from $\mathcal{S}$ by speech via mobile phone.

In the scenario we're currently developing, $\mathcal{S}$ is a mobile system that is lent to travelers who make use of a large airport (e.g., Frankfurt Airport). $\mathcal{S}$ multimodally answers $\mathcal{U}$'s questions about various aspects of the use of the airport (e.g., how to get to $\mathcal{U}$'s departure gate as quickly as possible).

In both scenarios, we have not hooked the system up to actual speech recognizers or synthesizers (although this extension is planned for a later phase). Instead, speech input and output have been simulated. The function of our prototypes is to make inferences and decisions.

The aspects of $\mathcal{U}$'s situation that we have concentrated on are:

1. the extent to which the situation creates *cognitive load* for $\mathcal{U}$, making it harder for $\mathcal{U}$ to perform his tasks and to interact successfully with $\mathcal{S}$; and
2. the extent to which $\mathcal{U}$ is under *time pressure*.

*This research is being supported by the German Science Foundation (DFG) in its Collaborative Research Center on Resource-Adaptive Cognitive Processes, SFB 378, Project B2, READY.

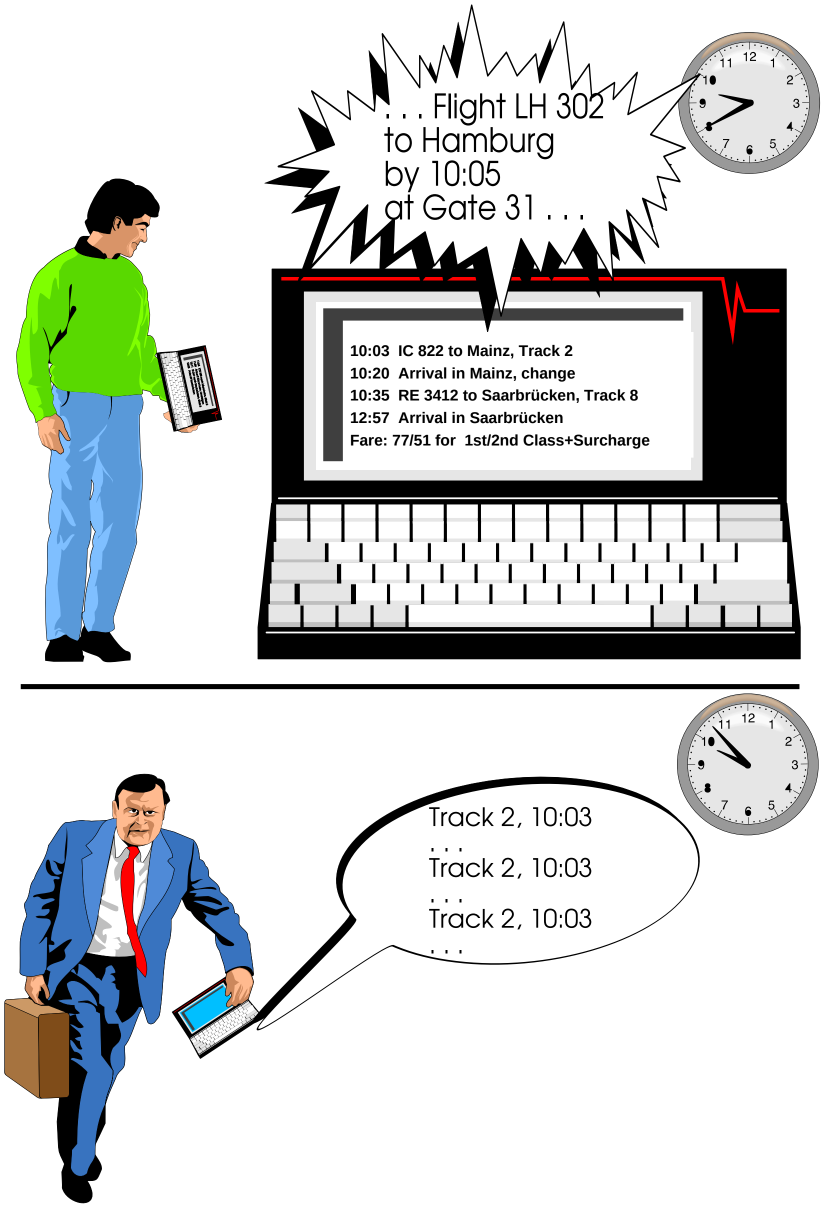

Figure 1. *Illustration of the type of situated interaction aimed at it the* READY *airport scenario.*

2 Using the User's Behavior as Evidence About the Situation

2.1 Motivation

The most straightforward way for a system to obtain information about $\mathcal{U}$'s situation is through sensors or through data that are more or less directly available to $\mathcal{S}$ (e.g., information about the estimated time of departure of $\mathcal{U}$'s plane). Although READY uses some information of these types, the focus is on analyzing $\mathcal{U}$'s behavior so as to make inferences about the situation and its effects on $\mathcal{U}$.

Up to now we've concentrated on the analysis of $\mathcal{U}$'s speech input: How can $\mathcal{S}$ infer from features of $\mathcal{U}$'s speech whether $\mathcal{U}$ is subject to unusual cognitive load or time pressure. We're now looking into similar questions about other aspects of $\mathcal{U}$'s input, such as $\mathcal{U}$'s use of a scrolling/pointing device.

There are several reasons why it may make sense to focus on $\mathcal{U}$'s behavior in this way:

1. Sometimes more direct evidence about $\mathcal{U}$'s situation is not available. For example, if $\mathcal{U}$ is communicating with speech via a mobile phone, it's unlikely that $\mathcal{S}$ will be able to get information from sensors in $\mathcal{U}$'s environment or on $\mathcal{U}$'s body.
2. Methods for interpreting $\mathcal{U}$'s behavior tend to be more generalizable than those that rely on other sources of information. For example, there are a great many factors that can cause high cognitive load in $\mathcal{U}$. The particular factors may vary from one domain or context to the next, so that it's hard to develop ways of capturing them directly. To the extent to which we're interested in the influence of the situation on $\mathcal{U}$'s psychological state, it may be best to assess it more directly on the basis of $\mathcal{U}$'s behavior.

A difficulty is that $\mathcal{U}$'s behavior often yields only unreliable evidence about $\mathcal{U}$'s current psychological state. But the same difficulty applies when more direct information about the situation is interpreted with a view to inferring $\mathcal{U}$'s psychological state.

2.2 Examples

There already exists a large body of experimental evidence concerning the question of how cognitive load is reflected in features of a speaker's speech (see Berthold & Jameson, 1999, for a brief overview). Some of the more important consequences of cognitive load are an increase in the number and duration of various types of pauses, a slight reduction in the articulation rate, and a higher frequency of sentences that are started and then broken off. We are currently conducting an experiment that yields more directly relevant data about the role of symptoms like these – and also about their dependence on other aspects of the speaker's situation, such as the relative priorities of speed and quality in the production of speech.

3 Performing Empirical Studies at an Early Stage

3.1 Motivation

When one thinks of the role of empirical studies in system design, the first thing that comes to mind is usually the evaluation of prototypes. We believe that, at least in the area of situated computing, more attention should be devoted to studies that create an empirical basis for the system design at an early stage. Without a solid empirical basis, a later evaluation study is likely to reveal that the system doesn't work very well, without giving much indication of how it could be improved.

There are a lot of questions about the consequences of features of the situation which cannot be answered reliably on the basis of a priori considerations.

3.2 Examples

Here are some examples of empirical studies that we've performed, in addition to the experiment mentioned above:

1. We did thinking-aloud studies with firemen who answer emergency calls, to see how they assess a caller's situation on the basis of the evidence that comes over the phone line and how they adapt their own behavior accordingly.
2. We analyzed the transcripts of a field study in which our first usage scenario was simulated, with the role of the system being taken by an experienced auto mechanic. This

analysis revealed, for example, which features of a user's speech occur frequently enough to be potentially useful as evidence (Berthold & Jameson, 1999).

3. We conducted a psychological experiment to see how a help system's instructions should best be adapted to the user's current cognitive load (see Jameson, Großmann-Hutter, March, & Rummer, 2000).

4 Employing Decision-Theoretic Methods From Artificial Intelligence

4.1 Motivation

As was mentioned above, much of the evidence that a system $\mathcal{S}$ can obtain about $\mathcal{U}$'s current situation and/or psychological state is unreliable: Often, it is only on the basis of multiple pieces of evidence that $\mathcal{S}$ can make a useful, though still uncertain, inference. Bayesian networks are a powerful technology for processing this type of evidence. (See Pearl, 1988, for the classic exposition and Jameson, 1996, for an introduction that includes references to many user-adaptive systems.) In particular, dynamic Bayesian networks make it possible to model properties of the situation and the user that change over time (see, e.g., Jameson, Schäfer, Weis, Berthold, & Weyrath, 1999).

Decisions that are made – implicitly or explicitly – by situation-aware systems need to take into account multiple factors and goals, as well as uncertainty about the relevant variables. Decision-making techniques such as *influence diagrams* offer ways of dealing with these complications.

Methods for *decision-theoretic planning* (see, e.g., Boutiller, Dean, & Hanks, 1999) make it possible for a system, when deciding what to do next, to consider how the next few steps in an interaction might proceed. For example, when deciding how to present a route description, $\mathcal{S}$ can consider how likely it is that $\mathcal{U}$ would fail to understand particular parts of the description—and what it would have to do to recover from such a failure.

Finally, decision-theoretic methods include ways of learning user models automatically on the basis of empirical data. It is therefore possible to base a system's inference methods more or less directly on the type of empirical data that was discussed in the previous section (see, e.g., Großmann-Hutter, Jameson, & Wittig, 1999; Jameson et al., 2000).

4.2 Examples

In the workshop presentation, concrete examples of the application of decision-theoretic techniques will be given, as time permits. The emphasis will be not on the technical aspects but on the extent to which these methods constitute useful tools for those who develop systems for situated interaction.

5 Concluding Questions

As a way of opening the discussion, the participants will be asked for their evaluation of the claims just presented. In addition, the question will be raised of how the methods presented here can best be combined with those mainly used by the other workshop participants.

6 References

Note: All of the listed papers that were cowritten by the present author are available via the READY web site: http://w5.cs.uni-sb.de/ ready/.

Berthold, A., & Jameson, A. (1999). Interpreting symptoms of cognitive load in speech input. In J. Kay (Ed.), *UM99, User modeling: Proceedings of the Seventh International Conference* (pp. 235–244). Vienna, New York: Springer Wien New York.

Boutiller, C., Dean, T., & Hanks, S. (1999). Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artifical Intelligence*, *11*, 1–94.

Großmann-Hutter, B., Jameson, A., & Wittig, F. (1999). Learning Bayesian networks with hidden variables for user modeling. In *Proceedings of the IJCAI99 Workshop "Learning About Users"*. Stockholm.

Jameson, A. (1996). Numerical uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interaction*, *5*, 193–251.

Jameson, A., Großmann-Hutter, B., March, L., & Rummer, R. (2000). Creating an empirical basis for adaptation decisions. In H. Lieberman (Ed.), *IUI2000: International Conference on Intelligent User Interfaces.* New York: ACM.

Jameson, A., Schäfer, R., Weis, T., Berthold, A., & Weyrath, T. (1999). Making systems sensitive to the user's changing resource limitations. *Knowledge-Based Systems*, *12*, 413–425.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference.* San Mateo, CA: Morgan Kaufmann.