

## What's in a user?

Karen Sparck Jones

Computer Laboratory, University of Cambridge  
New Museums Site, Pembroke Street  
Cambridge CB2 3QG, UK

I take user modelling to imply reasoning, explicitly or quasi-explicitly, about the user. Thus while systems or interfaces usually embed assumptions about their users (eg they know English or find some icon mnemonic), and defacto exploit these assumptions, this is not real user modelling.

But we cannot talk about the need for or nature of user modelling, let alone its feasibility, without reference to the critical features of the context. These are

- (1) the role of the human: agent or patient
- (2) the status of the user: central or peripheral
- (3) the orientation of the activity: outward or inward
- (4) the character of the information: deep or shallow.

In one salient class of contexts, illustrated by advice systems like UC, there is one individual who is both agent, the person operating the interface, and patient, the person about whom the system is taking decisions. The user is central rather than peripheral as the system task is to support him. The system-user activity is directed outwards towards the user rather than from the user to other system functions. Finally, the information about the user supplying the substance of the model is deep in the sense of dealing with the user's beliefs, goals and plans, as opposed to shallow in the way a simple log of mentioned topics would be. Instruction systems may also have some or all of these properties.

But users may figure in quite other ways, including those where agent and patient are distinct individuals with their own proper beliefs and goals, as with lawyer and client; those where the user is subordinate or peripheral, as in industrial plant control, so the dialogue manager with its user modeller is a secondary and not the dominant system component; those where the user is the starting rather than the end point for system action, as in document production; and those where the information about the user is extensional and not intensional, as with social security status for example.

In many contexts where it seems legitimate to refer

to user modelling it is nevertheless not obvious that this has to look below the interactive surface for underlying beliefs or goals, or even at surface beliefs and goals rather than at straight behaviour, or indeed at more than ergonomic factors: all of these can imply reasoning about the user. Thus to set more precise bounds on a discussion of the need and nature of user modelling we can ask whether we should model user(s) in cases like designing a system for

- (a) supporting a psychoanalyst and subject
- (b) providing holiday advice
- (c) managing personal bibliographic files.

The strong user modelling claim is that for even the third of these we should go beyond not just ergonomics but also recorded behaviour (allowing the system to adapt eg its classification) and even stated user requirements (as in searching) to modelling based on explanatory properties of the user. These may be objective or subjective, but subjective properties - beliefs, goals and plans - are especially important here. It is necessary, further, for the system to recognise and manipulate these properties not only where they relate to the system's primary decision making and so influence system effectiveness, as in giving appropriate advice, but also where they do not determine primary decisions but can influence system acceptability, as in presenting advice.

But it is generally so difficult to get this type of information, especially reliably and particularly for non-decision properties, that it may be safer to be more conservative and not look too deep. This implies that modelling should

- (1) consider only readily available information, if not supplied by the user then easily derived via the system's task and domain knowledge;
- (2) focus on decision rather than non-decision properties unless the latter are explicit and exploitable;
- (3) treat subjective, intensional properties only insofar as these are primary for the task;
- (4) take the user as an individual (whether or not taking a stereotype as a starting point) just as so defined by the particular user-system interaction.

These principles apply whether natural language is the dialogue medium or not, but clearly work out in detail somewhat differently according to means of interaction.

In some contexts weak modelling, through behavioural adaptation, via modifiable (though not fixed) stereotypes, or without reference to intensional attitudes, may be sufficient for effective and acceptable system performance. But stronger modelling even if not mandatory appears to be desirable, not only for obvious cases like advice systems but in any complex system context where user and system are both goal directed, but the relation between their goals is not obvious and has to be established quickly. Thus it is hard to see how a financial advisor or literature search system could not be improved by characterising the user's input in terms of knowledge and purpose, and by establishing relationships between different inputs through connected knowledge or purposes.

Even so, I believe that strong modelling can be adequately done, and in the normal absence of extensive and reliable data is best done, with only the system's task and domain knowledge, plus the general conventions of communicative dialogue, to support the given specific dialogue; and that it is neither necessary nor sensible to go looking too hard for hidden motivations. Too much refining on explanatory hypotheses is no more needed for everyday conversation with computers than it is with people.

Thus with a holiday advisor, if the user asks about travelling by train, just accept that he wants to travel by train without hypothesising why (e.g. it is cheap); or with a training system for equipment fitting, if the apprentice asks which bit is the pobjoy just say which, without considering all the bits the apprentice might think is the pobjoy and telling him they are not. Conservative strategies like this allow for some evaluation of goals and preferences in task and domain terms, as in UC; but they stop speculation getting out of hand.