Technical Report

UCAM-CL-TR-111
ISSN 1476-2986

Number 111

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Realism about user modelling

K. Spärck Jones

June 1987

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/TechReports/

ISSN 1476-2986

REALISM

ABOUT

USER MODELLING

Karen Sparck Jones

Computer Laboratory, University of Cambridge
Corn Exchange Street, Cambridge CB2 3QG, England

June 1987

Abstract

This paper reformulates the framework for user modelling presented in an earlier technical report, 'User Models and Expert Systems', and considers the implications of the real limitations on the knowledge likely to be available to a system for the value and application of user models.

This paper will appear in User Models in Dialogue Systems (ed Kobsa and Wahlster), Berlin: Springer Verlag.

25.6.87

Realism about user modelling

Karen Sparck Jones

Computer Laboratory, University of Cambridge

Introduction

The object of this paper is to explore the issues raised by user modelling, from an analytic point of view. It is generally assumed that systems cooperating with human beings, and especially systems with natural language interfaces, need, and can get, user models. There is also a widespread assumption that the more user model the better and, correspondingly that, while we may not yet see quite how to do it, there is no problem in principle about being able to construct ever richer models.

The user models of current working systems, with some specialised exceptions like teaching systems, are very simple. This paper does not describe any implementation. It considers examples that are much more complex than those that could be supported by the present state of the art, in order to lay out the elements and implications of modelling and so provide a framework for future work on actual implementations. This framework distinguishes and characterises the major factors in modelling, namely the nature of the information in a model, the function of a model, and the means by which the information for a model can be obtained. In any specific modelling enterprise it is necessary to be clear about the relevant values of these application parameters.

The direct conclusion to be drawn from the paper's analysis is a pessimistic one: even assuming very powerful system resources, far beyond the scope of today's systems, there are generally early limits to the modelling that can be achieved. Thus though the examples illustrating the argument imply modelling aims far more ambitious than those on which current system building is based, the conclusion is relevant to the long term research goals of the field.

But there is also a more positive and constructive message in the paper. This is that the useful, if limited, modelling it is reasonable to seek would be promoted by adopting the framework the paper presents. The distinctions it draws between the main components of modelling are important and should be applied.

The form of the paper

The paper attempts a comprehensive treatment of modelling considering first, what is being modelled, second, what it is being modelled for, and third, from what the modelling information may be obtained. The implications to be drawn from this discussion, and the relations between the account of user modelling given here and those presented elsewhere, are examined in the fourth, final, section on rational principles for user modelling. The points made are illustrated through extensive examples. These examples, which are designed to show what is involved in modelling taken seriously, are necessarily drawn from hypothetical systems. The paper further takes specific modelling mechanisms for granted and is not concerned with what the technical procedures involved might be. The focus is on such questions as the aims and value of modelling, taking it for granted that it can, in a strictly technical sense, be done: thus I am assuming, for example, that some piece of information can be extracted from some linguistic expression, that a suitable data structure representing this and the other information constituting a model can be created and manipulated, and that the results of this processing can in turn drive further system operations.

For the purposes of the paper I am also making some further assumptions. First, somewhat obviously, that I am modelling a human user. Second, I assume that the object is to model users as individuals, not simply to assign them to previous, explicitly defined classes. I am further assuming modelling in the context of complex computational systems of the kind represented by substantial expert systems. That is to say, it is taken for granted that the system may, for its primary purpose, exploit large bodies of knowledge and engage in extensive reasoning, with no direct connection with user modelling. I am also assuming, naturally, that there is some means of interaction between user and system: the implications of natural language interaction will be examined in the third section of the paper. In addition, I assume throughout that we are concerned with the modelling that is done in a single interactive session, taking longer-term modelling as dependent on, and not essentially different from, this.

Finally I am assuming, without attempting to define rigorously, some notion of what a model is, as constituting something more than a simple aggregate of pieces of data. Individual items of information, or a collection of these, do not constitute a model. The presumption in talking of user models is that items of information about a user may be related to one another, or to other (typically general) knowledge stored in the system, in a manner which supports predictions that can stimulate further system actions. This is not necessarily a very strong definition, since much depends on the nature and extent of the relationships implicitly or explicitly holding between pieces of information. For example in a medical system there may be a tacit assumption of causal relations between symptoms. The kind of categorisation applied in GRUNDY [RICH79] to predict book choices, on the other hand, does not make any strong assumptions about causal relations between the different user characteristics.

1. What is being modelled

This section considers aspects of users bearing on modelling, or modelling factors, drawing a number of distinctions relative to these factors. Some of the distinctions may seem over-fine. The justification for them is to be found in Section 2, where the uses to which models may be put are examined.

User roles

It is necessary, first, to distinguish the different roles users can have. In a given session, we may have different users, with distinct roles, or just one user with different roles. The two roles I distinguish are the patient role and the agent role. I define the patient as the subject of the system's decision-making processes, confining the system's decision-making processes here to those the system was primarily constructed to carry out. I assume the system was designed to reach decisions in a specific domain, in the sense in which classical expert systems have been designed to reach decisions about the diseases people have, the financial investments that might be appropriate for individuals in given circumstances, and so forth. The domain knowledge is instantiated in relation to the user, to reach a domain conclusion about him. (This therefore views personal financial decisions as ones about a human money owner and not as ones relative to abstract financial situations.) This is not of course to imply that a system reaches a decision only at the end of a session: it may take decisions continuously throughout a session, as in a teaching system.

As will be evident, a system can make decisions, in the general sense, about other matters. But I am specifically restricting "decision-making" here to activities aimed at satisfying the system's primary task goal. In the medical case "patient" has an obvious interpretation, but I am using it as a role label in a more general sense.

The agent, in contrast, is the person who conducts the interaction with the system. The agent, as such, is thus not the subject of the system's decision-making processes.

The implications of this distinction between patient and agent will become apparent as the other components of modelling are considered. Here it is necessary to point out that different person and role situations can arise as shown in Figure 1. Thus for a given system we may have a single human whose only role is as patient, as with a surveillance system where the interaction between the system and the external world, and hence with the patient, is through automated monitoring. Or we may have a situation where the single human being is only the agent, as in industrial design, where the system's decisions are about inanimate machine structures. We may have two human beings, one as patient and one as agent, so the latter is acting as the medium for information about the patient, as in the medical case with a nurse as system driver. Finally, we may have a system with a single human who is both patient and agent, as in a home computer diet advice system.

It is clear that both patients and agents, as humans, can have beliefs, goals, plans, etc. All users, whether patients or agents, are thus agents in the general sense of purposive beings. However I shall specifically reserve "agent" in this paper for the role definition I have just given, taking it for granted that any user can have beliefs, goals etc; (so if it is necessary to refer to the general meaning of "agent", this will be clearly marked). My use of "agent" is therefore more particular than that generally adopted in discussions of user modelling. My distinction between agent and patient is indeed intended to allow an analysis of aspects of modelling not covered by the normal unitary view of the user (though Morik [MORI87] makes a similar distinction between 'user' and 'protagonist').

I shall use "user" where the distinction between patient and agent is immaterial because the points made apply to either role. This makes "user model" comparable to "agent model" as the latter is generally understood, but here I am using "agent model" more restrictively to refer specifically to a model of a human in the agent as opposed to patient role.

<u>User properties</u>

It is further necessary to distinguish types of user property. Some of the distinctions to be made here are not theoretically watertight, but they are operationally important because they are reflected in material differences in the reliability and accessibility of modelling information.

The first distinction, see Figure 2a, is between <u>objective</u> and <u>subjective</u> (or mental) properties, for example between clinical thermometer readings and sensations of worry. (Note that subjective properties are defined as subjective in the possessor, not the observer.) This is an intuitive, but real enough, distinction. In general we behave, and therefore a system could similarly be expected to behave, as if it is easier to establish, for example, that a user is female than that they believe democracy is under threat. (There is no space here for the more detailed discussion this distinction requires: for example what is the status of reported properties, as in "I felt terrible pain"? A good case can nevertheless be made for the basic distinction as one with consequences that have to be taken into account in modelling. For some further points see [SPAR84].)

A second distinction, cutting across the first, and applying most obviously but not exclusively to subjective properties, is that between <u>static</u> and <u>dynamic</u> properties, for example between political beliefs that hold throughout the session and preferences for dresses that change in interaction with a mail order system. This is not to imply that political beliefs cannot change. It is rather that, in relation to a system's purposes and activities, there are properties which the system can treat as static because they are, if not necessarily static, at least relatively permanent and are unlikely to be changed by the system's activities, which are not directed towards decisions concerning them.

The distinction between static and dynamic does not refer to the effect of time on the <u>system's</u> knowledge of a user property, where a system may suppose X until the passage of time brings evidence that not-X, but the user is throughout in state not-X. It refers to changes in the <u>user</u> over time, i.e. from state X to state not-X. In practice it may be difficult for a system to know whether its knowledge or the user's state is changed, and in some cases the consequences may be the same. But I have introduced the distinction to allow for the situation where changes in the user are important because modelling is intended not merely to recognise them but (possibly) to effect them. Clearly, though, there need be no close correspondence between when a change in a user is observed by the system and when it is experienced by the user. It is important to emphasise that objective properties can change as well as subjective ones, though this may be less common: for example instructions to sell financial stocks.

There is, finally, a third property distinction which also cuts across the others. This distinction, unlike the others, is an absolute one because it is explicitly defined in the construction of the system for its primary task; and as it is fundamental to the system's purpose, it is also much more important than the other two. This distinction is one between decision and non-decision properties. As noted earlier, an expert or other complex system is intended to take decisions in a certain domain. Those user properties that pertain to this domain are a user's decision properties. According to the design purpose of the system these may be, for example, medical symptoms, or preferences for sorts of clothes. They can clearly be a mixture of objective and subjective, or of static and dynamic, as illustrated in Figure 2a, which shows a static objective property, temperature reading, and a dynamic subjective property, sartorial preference, as possible decision properties. Whether a property is a decision one or not is wholly contingent, as it depends on the specific system application. Any particular user property, e.g. political belief, may be a decision property in one system (e.g. a voting advisor) and a non-decision one for another (e.g. an arithmetic tutor).

If the system uses both non-decision properties and decision ones, the distinction between them may appear unnecessary. But it may nevertheless be important operationally. Property information about users can be exploited for different purposes, described in the next section. Clearly that for which decision property information is sought is the most important. The permanent system model of the domain phenomena to which these properties refer may therefore be much more fully developed and characterised than that of other entities or activities; and it may also be the case that the importance of decision property information may mean that it has to be handled more carefully, for example be more rigorously tested. There are therefore good practical reasons for distinguishing these two types of property, as will be more clear from the examples considered later, though there may be no intrinsic difference in the nature of the properties labelled decision and non-decision for a given system.

Though for precision it might be desirable, in relation to to the three distinctions just drawn, to refer respectively to, say, the type, mode and status of a property, I shall regard it as sufficient to refer simply to the (six) types of property, as listed in Figure 2a.

It will be evident that in general these distinctions of type are relevant to both patient and agent roles. Thus as Figure 2b shows, the same property, like salary, or age, can hold of patient or agent, indeed hold of both where there are two people. However there are some restrictions. Thus by definition decision properties can only be properties of the patient: they do not apply to the agent role.

2. What modelling is for

The previous section characterised the nature of the inputs about users with which a system could be concerned. The next question is why it should be so concerned. Specifically, if the function of a system is to take decisions, why should the agent role be considered. Why, indeed, should non-decision properties even of the patient be considered.

Both roles, and both decision and non-decision properties, are relevant because there are quite different system functions user models can serve, so non-decision properties of either patient or agent can be helpful in supporting or enhancing system performance. The possible functions models can serve are illustrated in Figure 3, for a medical system with distinct patient and agent. Thus the first, and most important function models can serve is system effectiveness. The prime object of a user model is to ensure the system, as a decision-making system, reaches the correct decision. For example, if we have a medical diagnosis system, we need a model of the patient which is adequate to support what, if independently evaluated, would be shown to be the correct decision (assuming that, given adequate evidence, the system's processing is itself sensible.) This clearly implies that the only relevant model for effectiveness is a patient model. Effectiveness is clearly also a necessary system function.

However effectiveness is not the only system function user models can serve. Models can serve system efficiency, i.e. reaching the correct decision in an economical way. A patient model is obviously relevant here too: for example the available knowledge about a patient may be used to order the tests to be made on a patient for further information so that the most potentially useful tests in the specific circumstances are carried out first. Moreover if the agent is seen as a transmitter of information about the patient, the system can also exploit an agent model to gather information in one way rather than another, as likely to be more efficient. For example, if the agent is medically experienced, this fact can be used to promote efficiency through quick answers because well-defined medical terms can be used in questions the system puts to the agent.

There is also a third function a user model can serve. This is system acceptability, i.e. supporting (or expressing) its decision making in a comprehensible, perhaps also agreeable, way. (This refers primarily to the system's operations in a session, not to the extent to which its performance over many sessions carries conviction about the validity of its decisions.) Here again both agent and patient models can be used, for instance the fact that the patient is in pain can be exploited to emphasise one aspect of the proposed regime (as likely to reduce pain) rather than another (which may be nasty), and the fact that the agent is experienced to augment the decision information with literature references to comparable case histories. Acceptability, like efficiency, is not a necessary function; but both may be very important in practice.

As the examples of Figure 3 show, the different role models may be used, in different ways, to serve distinct system functions. But this general statement disguises some relevant complexities. Thus in relating models to functions it is necessary to separate the source of the modelling information from its functional destination. In the example of Figure 3, source and destination are in fact the same, e.g. information about the agent is applied to the agent, as in the exploitation of his experience to use technical terminology. But patient data could also, for example, be applied to system functions directed at the agent (i.e. in his own right, not as an information transmitter), for example to explain the diagnosis.

There are indeed further points to notice, which also have implications for the detailed way in which models are exploited. In using models to serve functions, it is necessary to separate (logical) addressee from beneficiary (as well as actual from logical addressee: the agent may be actually addressed to transmit a question to the patient). Thus the agent may be addressed, using an agent model, to obtain information exploited for the benefit of the patient, or for the agent's own benefit; or the patient addressed to obtain information which may also, as just suggested, be exploited for the benefit of the agent. One may also have first and second order beneficiaries, so for example, benefitting the agent indirectly benefits the patient. In other words, we may have both patient and agent models serving system functions themselves directed towards both patient and agent. Either patient and agent model may naturally, moreover, be exploited differently both to gain inputs, and to direct outputs, in supporting the system's various functions.

It must further be emphasised, as has been implied and will be more fully illustrated below, that both decision and non-decision properties can be of value in relation to the system's functions. In the example of Figure 3 this is shown most particularly by the use of the agent model in connection with efficiency, since the agent's properties are necessarily non-decision ones; but if we assume, for example, that pain is not a decision property for the disease in question, a patient non-decision property is being used for acceptability (to the patient).

The set of distinctions drawn in this section seems very elaborate. But the distinctions and their implications are relevant if we want to serve system purposes by modelling, given that modelling situations range from those where it is necessary to recognise user properties to those where it is only desirable, and from those where models are used to promote positive system behaviour to those where we want to inhibit negative behaviour. This discussion of course assumes, in general, a system of sufficient complexity to suggest a need for modelling beyond that involved in constructing and manipulating the basic patient decision model. But we have still to relate the validity of the modelling that can in practice be done to its actual utility.

Example 1

The need, assuming modelling is feasible, to recognise all the distinctions discussed, relating functions and roles and also property types, can be illustrated by a further system example, summarised in Figure 4. Thus suppose we have a (notional) social security expert system with distinct patient and agent. The patient is an elderly and disabled woman who has to use a wheelchair; she is also poor-sighted, a Catholic, and honest but suspicious of officaldom, poorly informed about benefits, and believes that age is the main determiner of benefit (these are illustrative properties, and are not intended to constitute a complete patient description). The agent is an experienced clerk, but one who is a male chauvinist who believes women are unduly favoured by the benefits system. Figure 4a lists the patient's properties by type. For clarity, but in some sense artificially, user properties of different types are treated here as pertaining to separate models. Thus for the patient we have a P1 model referring to her static, objective, decision properties, in this case age etc, another model P2 referring to her static, objective, non-decision properties, poor sight etc, in fact eight models altogether, though not all are instantiated in the example. Agent models are correspondingly numbered, but as an agent cannot have decision properties, the only relevant models are the even-numbered ones.

The separation of models for a single role is somewhat artificial, but does serve to mark differences in the relative importance or stability of the information involved which, even in a single union model could be reflected in the way information in the various classes was managed and applied. Decision property information may be more carefully vetted, for example, and inconsistencies require different treatment in static and dynamic cases since static property clashes may have to be resolved where dynamic inconsistencies may legitimately reflect changes over time.

The kind of uses to which the example modelling information might be put are illustrated in Figure 4b and c. The patient's disability is directly used to reach a decision about the amount of benefit due, but may also be exploited in the interests of efficiency to organise the search for possible benefits (I assume a suitable rule structure allowing this). It may also be used for acceptability to the patient to explain the makeup of the benefit payment, for example to indicate that disability precludes some other apparently relevant factors like sex. The patient's dynamic belief that age is the determiner of benefit could be used to support data gathering efficiently through the forms of questions, or might be corrected as part of an explanation for the nature of the questions being asked.

Turning to the agent's illustrative properties, his experience could be exploited for efficient data gathering and to make the decision reached acceptable by reference to the relevant sections of the Regulations. His unfortunate chauvinism, on the other hand, might have to be counteracted, for efficiency, by careful question formulation in data seeking and, for acceptability, by an indication that sex is not the basis of the calculated benefit. (Notice the need to maintain a broad definition of acceptability, and the need for the system to discourage as well as encourage user behaviour.)

This example is developed in more detail in [SPAR85], and is compared there and in [SPAR84] with other examples illustrating systems with different purposes and other patient and agent role combinations and property descriptions. The various examples taken together are designed to provide a fuller picture of the relevance and utility of the factor (i.e. role and property type) and function distinctions I have made. The situation would be quite different, for example, in a system with the rather different factor combination represented by a learning system with patient as agent and decision properties that were all or mainly dynamic. The social security system would of course also be very different with different properties or property values.

These different examples, however, simply reinforce the important general points made here by the social security illustration. These are that when one looks at systems as wholes and not only at interactive interfaces, and at systems engaged in non-trivial tasks, it is necessary to allow on the one hand for distinct user roles, and property types, and on the other for distinct system functions. These distinctions then imply multiple relationships between pieces of information about the user and system purposes. Thus an individual user property may be exploited not only to support different functions; it may be used in different ways to serve a single function. A particular functioin may, conversely, exploit not merely different properties of the user, but different types of property. It is particularly important to recognise the need, in more complex systems and system environments, for the distinction between decision and non-decision properties and the value of non-decision information in supporting the efficiency and acceptability functions which, if they are not primary in the way effectivenes is, are nevertheless highly significant.

Thus if we consider a single user property, namely sex, this may be a decision property in some systems. But in other systems where it is not a decision property, it can still be very useful. For example (and making some arbitrary medical assumptions), if sex is not a decision property for a smallpox diagnosis system, it might still be exploited for efficiency to order diagnostic tests (supposedly justifiable by possible pregnancy implications), and/or it might be used in relation to acceptability (to the patient) to couch verdicts in different terms (assuming women are more sensitive to potential disfigurement than men); sex could indeed be used in more than one way in relation to acceptability, for instance to draw attention to the implication of verdicts for contact with children (generally rather different for men and women). Of course in choosing information to present the system is making a decision in the ordinary sense, though not in the strict sense defined earlier: properties which are not decision properties for the central task become decision properties for the subsidiary response generation task. The distinction between central decision properties and non-decision properties is nevertheless a useful one.

The rich pattern of relationships between kinds of modelling information and their functional uses that follows from going outside the decision core of a system are summarised in Figure 5. This applies the convention introduced earlier and isolates each factor combination as a submodel, so looking at the different functions models can serve gives us the distribution shown in Figure 5.

As Figure 5 indicates, the contribution of some models to some functions is indirect rather than direct. The agent or patient's non-decision properties can bear on system effectiveness only indirectly through, e.g. the transmission of incorrect decision information. It is of course the case that quite apart from the general question of whether a system can in a straightforward sense obtain reliable information (it may simply not be clear, for instance, how severe a sick person's pain is), transmitting patient information through an independent agent can lead to all kinds of distortions (the doctor may think the patient is malingering, for example).

As a further potential complexity, it should also be noted that it is possible that each submodel may really be a set of models, either because different and incompatible interpretations may be made of the available evidence, or because the different uses made of the same information impose different structural organisations on it. This sort of complexity arises, moreover, even when the factor distinctions made earlier are suppressed so all the information about one human is amalgamated in a single model. (This of course assumes the system can recognise, if it is not set up to assume, there are two people as opposed to one.)

These points are illustrated more fully in Example 2, presented in section 3 below.

3. What modelling is from

I have not so far considered the means of communication between the user (strictly agent) and the system. If we can see the point of modelling, there is still the question of how easy it is to do. This clearly depends critically on whether the user information needed for modelling can actually be got, and this in turn depends on the way user and system communicate.

The context in which this paper is being written assumes that user and system (and especially the former) communicate in natural language, and ideally in full and free natural language. Natural language is of course not the only possible means of communication. Setting aside a situation where user, and specifically patient, communication is like the autonomous and involuntary monitoring of VM expert system [BUSH84], where the patient is not a decision-making agent in the ordinary sense of "agent", there may be applications where the user may communicate, as a purposive agent, with the system through non-linguistic means. This could be the case, for instance, in a design system where drawing was the means of communication, or in a mathematics teaching system where arithmetic terms and expressions are used. Formal language, and quasi- and restricted natural language are also obvious possibilities, which may be implemented to allow more or less initiative and direction to the user: with the sort of restricted natural language interface often found with menus, for instance, the user's freedom is typically very limited.

It is important to recognise that these other communication means can perfectly well support user modelling (as further discussed in [SPAR84]): intelligent computer-aided instruction in school mathematics is an obvious example, where the instruction is specifically driven by a model derived from the user's problem-solving behaviour. Retricted communication e.g. through menus, may also be perfectly satisfactory for a particular application. I shall nevertheless concentrate on interaction through unrestricted natural language (assuming therefore a natural language processing component in the system which is far beyond out present capabilities), primarily because natural language offers the richest communicative resources and so the best opportunity to gather information indirectly if, for good reasons, it is not supplied directly. In language use information may be conveyed both directly, as in the answer to a question, and indirectly, as in the form of words chosen. As in human dialogue, it may not be feasible to gather all the modelling information directly either because the necessary explicit interchange becomes too tedious, and therefore oppressive, or because it may be too offensive to the user. Natural language concentrates many different items of information, simultaneously, in its expressions through the concurrent deployment of, for example, lexical choice, syntactic structure and form of reference, so natural language interaction offers the best chance of maximising user information. This is independent of the fact that natural language dialogue may be required or appropriate for other reasons, for example speech in telephone interaction.

It is also important to emphasise here that exploiting already-interpreted natural language communications to gain modelling information is logically distinct from having to have a model to interpret these language communications in the first place. Thus modelling may be logically parasitic on natural language communication adopted as an effective means of obtaining the system's decision-making information, as is most clearly shown where there is no human patient, but also holds for models other than decision-property ones where there is a human patient. Or it may be deemed to be a prerequisite for the utterance interpretation that supplies that decision-making (or other functionally-useful) information. There may in any case be situations, for example where the human patient is also the agent, where this distinction seems immaterial. The point of making it here is nevertheless to draw attention to the fact that in this paper the relation between language and model is being approached from one end rather than the other, so the concern is with what the linguistic expression contributes to the user model, and not with what the model tells us about the expression. This is not unreasonable, given the varied modelling purposes considered here since, though some model elements like subjective goals will clearly work both ways and play a part in both enterprises, others may have little part to play in assisting linguistic interpretation.

Example 2

The problem of extracting useful modelling information can be illustrated by an extended example. This is a hypothetical dialogue, representing the first part of an interaction, for the social security benefits system of Example 1. The users, the elderly disabled woman as patient and the experienced benefits clerk as agent, are assumed to have the properties given earlier in Figure 4, with some additional ones needed to make modelling fruitful and interesting (these will emerge in the dialogue).

The example assumes that the system has the necessary technical linguistic interpretation and generation capabilities, and also the necessary internal modelling component. Thus I am assuming that the system is in principle able to recognise instances of the kinds of properties it has available to describe patient and agent, and may indeed be capable of constructing new, particular properties out of other knowledge it already has. The example is designed to show what support natural language dialogue, and more specifically the user inputs, can offer for modelling, and hence to provide a basis for a realistic assessment of the opportunities for, and value of, user modelling in enhancing system performance.

In the dialogue the modelling conclusions drawn by the system from its inputs (by whatever means it has for doing this) are associated with the various submodels described earlier and illustrated in Figure 4, i.e. P1 the static, objective decision model for the patient, P2 and A2 the analogous non-decision patient and agent models, and so forth. Thus the presumption is that the system starts with empty models, and seeks to build up reliable and coherent submodels, with its dialogue responses seeking further modelling information as indicated. The use of the submodels is primarily, as before, to make the different modelling operations clearer. The symbol "+" below indicates the addition of information to a model, or its reinforcement if already present; "?" indicates the model proposition is hypothetical rather than certain (of course in practice there would not be a uncertain/certain dichotomy, but different degrees of certainty: however the simple dichotomy is convenient here). The comments, c, on input and output, summarise what the system is doing.

As this is the beginning of a dialogue the system's responses are primarily directed towards obtaining more modelling information, rather than applying a built model e.g. in explaining something, but as the example shows, the information seeking is frequently guided by the (tentative) models already constructed, though in some cases the system engages in quite basic information gathering. Moreover it is possible to show, even thus early in the dialogue, the system applying the information it has to different system functions. (Note that I am not concerned here with the plausibility of the assumed benefits regulations and its expert system implementation.) "O" is system output, "I" is user input. (I assume the initial logon sequence has been gone through.)

```
O: Please give me details of the applicant.

I: The applicant is female, aged 89, widowed, grade 3 disabled, and I
imagine should get a pretty lavish benefit.

c: P1 + female, 89, widow, disabled 3
   P2 + nil
   P4 + nil
   P8 + nil
   A2 + experience?
   (use of "grade 3" jargon)
   A4 + hostile to patient?
   (use of word "lavish", "and", "will" - compare "but", "should")
   A8 + nil

O: What are the applicant's sources of income?
```

c: response motivated simply by need for further decision property information, primarily to promote effectiveness (i.e. no model application yet).

I: She says she's getting widow's pension of 17pounds and 3pence and otherwise the priest sometimes gives her 2pounds out of his poor box. Her clothes seem rather worn.

c: P1 + pension 17.03, charity occasional 2.00
P2 + catholic?
(reference to "priest")
P4 + honest?
(pence details; mention of odd gift)
P4 + ignorant?
(does not realise housing status relevant)
P4 alternative + dishonest?
(concealing housing status)
P4 + suspicious (of clerk)?
(only replying to direct questions; note not necessarily an alternative)
P8 + no obvious beliefs, expectations about system?
A2 + nil
A4 + hostile to patient?
(reinforced by use of "she says", "seems")
A4 alternative + cautious observer?
(not experienced in evaluating state of old women's clothes)
A8 + exploring influence of modifications on basic income?
(apparent interest in whether she gets a clothing allowance)

O: Find out whether she buys clothes frequently. But has she been receiving anything for her disability? And ask her whether she lives with anybody.

c: response motivated by need to elaborate on P1, for effectiveness and possible efficiency; also to check on P4s of honesty and ignorance, attempting to select from alternative hypotheses, for the same reasons; by use of homely language attempting to bypass hypothesised suspicion, with additional motive of enhancing acceptability to patient; further motivation to gather information sufficient to start a useful P8. With respect to agent response motivated by attempt to distinguish alternative A4s and sidestep hostility, primarily for effectiveness but also for acceptability to patient; also using A8 suggesting focus on ancillary income to redirect attention to main sources of income. (Here the system is beginning to apply its models.)

I: She says the hospital gave her a wheelchair and that she lives by herself but her daughter sometimes visits and the District Nurse attends at basic frequency. She finds it very difficult to afford clothes.

c: P1 + no disability allowance, home owner?
P2 + nil
P4 + honest? i.e. ++
P4 + ignorant?
(apparently does not know of disability allowance, clothing allowance)
P8 + no notion of system and its capabilities?
A2 + experience?
("basic frequency" jargon)
A4 + nil
A8 + nil

O: Ask her whether she owns her home, and if she is still paying anyone for it (check for mortgage freedom). Tell her we may be able to help even if she has a house.

c: response pushing for more P1 information, for effectiveness and, we will suppose, efficiency and acceptability; also reacting to and testing P4s and P8, promoting acceptability and perhaps changing the patient attitudes embodied in the P8. The system is exploiting the agent's experience in A2 (reference to "mortgage freedom") and trying to reduce his hostility and render the system more acceptable by exhibiting care and exhaustivity in its operations.

This example illustrates many specific points about modelling, for example the possibility of agent interference in the transmission of information about the patient, and also the need to distinguish user beliefs about the nature of the domain as covered by the system from beliefs about the domain within which the system operates. For instance the patient may have beliefs about the scope of the system based on her beliefs about the scope of social security, although the scope of the system is narrower: for example in relation to housing, she may suppose the computer system covers rates rebates because social security as a whole does, though the system is in fact only concerned with housing in relation to asset ownership. The patient's mistaken beliefs could motivate replies to questions about what she is paying which were quite misconceived, but in a way very difficult to detect.

But the major, critical points made by the example are first, how heterogenous the evidence for models is, and second, how weak it can be. This is particularly in relation to properties other than decision ones, i.e. the non-decision properties of both patient and agent. Thus there are three data of different sorts, the three pence, the cash gifts, and the daughter visits, all suggesting the user is honest; however any one of these, like mentioning the three pence, is really only very weak evidence for honesty: it could rather indicate a passion for accuracy, for instance.

The essential problem about modelling is that property indicators may be unreliable or undiscriminating. Thus to take an independent example, a man being a

Catholic does not imply he is a priest, so being a Catholic is unreliable evidence for being a priest. But being a priest does not imply being a Catholic priest: as there are priests in other churches, being a priest is not discriminating. Building and applying models is clearly a very tricky business, even for decision models, but it is much more so for non-decision models. It should be noted, moreover, that these points also apply to modelling initiated by stereotypes: for example the stereotype old woman may be religious, but it may not be at all easy to verify that a given old woman is religious (in the limit even by asking) or, more importantly, to establish she is in fact not religious.

4. Rational principles for modelling

Given the manifest difficulty of obtaining reliable modelling information, what are the implications for system design? Should the system devote effort to trying to confirm the evidence? Information about the user is needed to guide action bearing on the user, but if the system does not have much, or good information, should it seek more, especially if drawing strong conclusions from weak evidence can lead to poor performance?

There is no doubt that if the decision properties the system uses are the only proper ones for the nature of the decision to be made, then every effort has to be made to get the best information about them, even if this is very difficult. The problem is much more serious with non-decision properties because it may be much harder to get useful non-decision modelling information.

This is primarily because there are potentially very many, diverse non-decision properties, not connected with one another or with decision properties. The presumption about decision properties, as mentioned earlier, is that they form a related set, which is typically also a select set, especially if the system's task is a well-defined and constrained one. But the number and range of properties a user can possess outside the decision set is large and wide, even considering only properties that might be validly exploited to serve any of the system's functions. Thus for the example system, user properties as disconnected as religion and poor sight could be usefully exploited in the interests of system acceptability, the one to encourage a visit to the priest to encourage a discussion of a low benefit, the other to direct the patient to specially-printed large type explanatory leaflets. The problem is that there may well be no motivation or leverage for a system search in one non-decision direction rather than another (unless a random pursuit is adopted on the basis that anything is better than nothing, and one thing is as good as another).

There is a particular further problem with non-decision property information, assuming it is not explicitly sought, namely that it is more likely to be conveyed indirectly than directly by the user, and thus be more likely to be subject to uncertainty. In general, the more indirect clues are the less helpful they are. Again, the deeper models are, as those concerned with subjective properties generally are, the harder the properties involved are to ascertain.

All this suggests that the realistic conclusion to draw is that rather than diverting systems into costly and probably still unprofitable searches for evidence, given that even quite powerful systems may yet lack extensive general world knowledge, we should restrict modelling to the user properties we have a chance of getting good information about, i.e. primarily, if not exclusively, decision properties, for which data, if they are not sought directly, may be more obviously supplied or, where decision properties are systematically related, inferred. The further conclusion is that modelling, especially modelling beyond decision modelling for effectiveness, should be very conservative, i.e. it should not be driven too deep for content or stretched to far in use.

It is relatively easy to imagine how this strategy could be applied to the social security example. The system would not bother, for instance, with the patient's religion as this is a high cost, low payoff property, i.e. one difficult to establish and of rather marginal utility through not having a predictable acceptability value; the system could probably do better tailoring output for acceptability by using suitable decision properties like age. Honesty and system expectations would, on the other hand, be available as non-decision patient properties, but would not be pursued unless the decision data to hand was blatantly incongruous or inadequate.

The illustration below is designed to emphasise these points, through a second and rather different application system.

Example 3

This example is summarised in Figure 6. The system here is a simple travel agent dealing with transport, so the system's decisions are recommendations on the means of getting from A to B. The decision properties are thus those setting constraints on the choice of mode of travel, e.g. date, cost etc. The user is assumed to be both patient and agent. Suppose, then, that we have the initial interchange between the travel agent T, and the client C, represented by the C1, T1, C2 in Figure 6a. The client's reference to trains can suggest three different hypotheses, (a), (b) or (c), about the user properties motivating this, namely that the user is interested in economy (trains are cheap), thinks travel by train is fun, or dislikes flying. Different choices by the system here would naturally suggest quite different responses, for example T2a for hypothesis (a) or T2b for (b). Clearly, if the system's hypothesis is wrong, its response could look very unmotivated to the user. A less risky strategy, therefore, than adopting a hypothesis in the absence of good (i.e. reliable and selective) evidence, would be to seek further evidence, allowing a motivated choice, via the response T2x which explicitly asks the user about his underlying interests. But the user may in fact have no such interests, and may find the further questioning unnecessary and hence irritating. Thus a third, even more conservative strategy would be simply to take the user's input at its face value, and to respond, as in T2y, just by listing the Istanbul train data. (Note that even here, the hypothesis that the user wants trains has properly some uncertainty attached to it, though it is reasonable to regard this hypothesis as much less uncertain than any of (a), (b) or (c).)

What this example suggests is that even where the system is dealing with decision properties, it may be better to be less active in constructing and using models, in an attempt to provide a tailored response, and simply to offer the user sufficient information for him to apply his own model of himself.

The hazards of overcommitting on evidence about non-decision properties, which may be even less adequate than that for decision properties, are illustrated in Figure 6b. Thus suppose that after C1 and T1 as before, we have C2, leading the system to hypothesise that the user has the non-decision property of being religious (being religious is clearly a non-decision property for characterising modes of travel.) This non-decision model may be exploited, with the best of intention, to produce T2p rather than T2q. But in fact, what sufficient reason is there for supposing that religion is the user's motive for not wanting to travel on a Sunday? Attempting to gather more evidence for the hypothesis, for example by questioning the user, would, moreover, be unnecessarily elaborate. The safest, and also the optimal strategy, is for the agent simply to produce the straightforward response T2r.

Conclusion

But is it really necessary to draw such a pessimistic conclusion as the paper would appear to suggest: namely that modelling is so difficult that we are unlikely to be able to provide systems with any material modelling capacity, and should therefore not even attempt to do this. After all, areas like intelligent computer-aided instruction demonstrate that successful modelling can be done [SLBR82], [CLAN86].

But it is important to recognise that modelling here is primarily decision-property modelling, made easier by the constrained nature of typical applications, namely teaching a relatively restricted skill like solving certain types of algebraic equation or basic linguistic knowledge. This is indeed not to say that building such systems is easy; but it can be claimed that user modelling here is less difficult than in many other cases because of the narrow focus such systems tend to have and the well-structured domain they normally have. It may also be possible to get leverage in systems like this by building in general characterisations of the kinds of users they may have.

Pessimism is moreover not permissible as far as decision properties are concerned. If the system has a human patient, this user has to be modelled. The correct, but of course not novel, conclusion to be drawn here from the paper is just that useful modelling critically depends having an adequate set of decision properties, and also that conservative modelling strategies are also in order.

It might further be maintained that, even allowing for other application areas more challenging than instruction (as ordinarily treated), there is really no need to take the kinds of modelling aim tacitly presupposed in the paper seriously. There is therefore no great cause for pessimism. Thus it may be possible to do very useful modelling, in a modest way, allowing for some non-decision properties, basic role distinctions and different functions, but not going beyond the more obvious sorts of information and relatively straightforward uses of them. This is in fact all that we are able to attempt now. But our performance in practice is not very encouraging here, and the examples of the paper equally suggest that it may not be so easy to identify and deploy even quite obvious information about users.

Adopting a modest approach to modelling does not, moreover, imply that there is no need for the kinds of distinctions drawn in this paper. Wahlster has argued (in his opening presentation at UM86), for example, that the distinctions are over-complex, imagining what contortions they would imply for the designer of a marriage counselling system which might have be faced with a concerned mother, son and son's wife. But as Figure 7 shows, simply considering different goals for the system implies a need to recognise the kind of distinctions I have drawn, as without them the system is unlikely to be able to satisfy its goals. Thus focusing on different parties carries different role distinctions with it. (A situation with multiple agents would require equally careful distinctions.)

As a basis for discussing and implementing user models, the general framework presented in this paper has obvious points of contact with those offered by [BELK84], [DANI86a], [KOBS87], [MORI87], [RICH83] and [WAKO86], for example. But the perspectives adopted in all these analyses vary, and there are interesting differences in granularity. Thus Morik's treatment [MORI87] of interactive settings for user and system actions in terms of the system elements open to change, the means of change, and the agencies of change, provides a rather different and in some ways more, though in other ways less, detailed analysis of modelling contexts than that provided here, suggesting that it may be useful to try to put the two together.

The general framework offered here is also neutral with respect to what may be called the orientation of a system to a particular generic task, for example explanation [BECO86], [BEEK86],[CLAN86], [SLEE86], especially explanation seen as an instructional activity. The implications of these generic tasks have obviously to be taken into account in user modelling: my point here has been rather to emphasise the need to look at user modelling outside the popular generic tasks. Thus explanation may in some cases be at most a minor system need.

Again the analysis offered here makes only weak assumptions about cooperation: cooperation with the patient may be minimal or non-existent, and sufficiently attained with the agent with very little effort. However it is clear than in individual applications it may be necessary to take into account the complexities of cooperative behaviour illustrated by [AIRE86] and [JAME87], as well as the need to apply specific response strategies in dialogue including, for example, explicit clarificatory and negotiatory dialogue, and to tailor responses to the user as illustrated by [CRBE87], [COJO87], [GOOD86] and [MCCO87], for instance. The paper's examples assumed such requirements but did not analyse them, since the paper was more concerned to emphasise the great variation of user modelling contexts and hence to allow for the fact that questions of cooperative response may not arise though user modelling itself is appropriate.

Beyond drawing a broad distinction between static and dynamic subjective properties, the paper did not attempt to give the kind of detailed analysis of the properties of beliefs, goals, plans etc carried out by [KOBS87], or to consider the detailed relations between the user's intensional states and the reference world investigated by [BECO86], [BEEK86], [CHIN87] and [MCCO87]. The broad view adopted has to be filled out in this way, as appropriate to individual system contexts. However the framework provided does draw attention to the fact that in many contexts it may be proper to talk about user modelling which is not exclusively devoted to, or dominated by, the user's intensional states, and also that users may, for example, have goals without also having plans.

The paper was also deliberately focused on the user model as opposed to any discourse or dialogue model, say of the kind illustrated by [LITM85], taking it for granted that these two are distinct but not analysing or defining the precise relationship as this is done, for example by [SCHU86] and the contributors to [UMDM87]. The analysis of user models has clearly to be supplemented for implementation purposes by a view of their relation to discourse models (and any other models like a world model). The object of the paper was to emphasise the fact that in complex systems the user model cannot be equated with the discourse or dialogue model, i.e. it is necessary to distinguish communicative properties of the user from real properties of the user, but then of course also to allow for very subtle relationships between these. It is clearly possible, in particular, to envisage further complications of the various kinds illustrated by [JAME87], [WIBI83] and the HAM-ANS hotel application [HOEP83], where dialogue participants may be deliberately deceitful, or may mix individual and institutional attitudes.

The examples in the paper illustrate the variety of possible sources of user modelling information also considered by [WAKO86], for example; and it is clear that though, as indicated, it may be very difficult to rely on structural relations between different properties or property types, especially for a wide scatter of non-decision properties, modelling has to look for relations between objective and subjective properties of the kind considered by [BROO86], [COJO87], [MORI87] and [RICH79]. How these are related is just one element of the specific modelling process; and the framework for modelling given in the paper has of course to be supplemented by an account of actual modelling procedures, i.e. of the inference mechanisms required to construct and deploy models, like those described in [WILE86] and for goals and plans in [BECO86], [CRBE87], [GOOD86] and [POLL86]. The aim of the paper has been to indicate the need to consider the various aspects of modelling, especially in relation to system functions, before detailed procedures are designed.

The analyses of the information-seeking interactions between the users and staff of specialised library services carried out by Belkin and his colleagues [BELK84], [BROO85], [BROO86], where user modelling is of the essence, serve to bring out just how hard modelling is in such complex situations. Establishing the nature of the library user's information need is user modelling in the broad sense of the present paper, though Belkin and his colleagues themselves give "user model" a much narrower meaning. Characterising someone's scientific literature needs requires a large apparatus of 'functional experts' to deal with the various facets of the modelling, and the application area is one which clearly illustrates the intractability of modelling in those cases where the user's subjective, i.e. mental properties are in question, and even more, in those cases where the user's goals, plans, beliefs etc are dynamically changing through the interaction. But the particular problem in this case is that the modelling information is necessarily inadequate, because the user is seeking information which, as he has not yet got it, he cannot properly specify. Brooks' and Daniels' detailed analyses [BROO86], [DANI86b] of human information interviews show just how difficult and how complicated capturing information needs and other relevant properties of users is: we are nowhere near being able to extract information from dialogues like this automatically.

However the interview transcripts also show quite clearly that even human beings may not find it useful or possible to seek to carry modelling very far. The overall message of my paper is that as the situations in which user modelling may be envisaged vary enormously, because they depend on combinations of user populations, system tasks, and application domains, the strategy to be adopted for modelling has to be evaluated for each individual system. But it may well then turn out that, quite apart from the fact that we cannot yet, for example, provide sufficiently powerful natural language processors, the conjunction of necessary ignorance about users and system resource limitations e.g. in relation to real time operation, may mean that if modelling is done at all, it can only be done in a very limited way. It is nevertheless still user modelling, and may be as adequate as is in fact required. This line is, on the one hand, a justification for the use of stereotypes, as in e.g. [CHIN87], and on the other for such modest applications of modelling as the individualised parsing proposed by [LECA87]. It also has the important advantage of being unlikely to mislead the user into thinking the system is smarter than it is.

Acknowledgement

This work was carried out during the tenure of a GEC Research Fellowship. I am grateful to Richard Young and to the participants in the Workshop on User Modelling held at Maria Laach in 1986 for discussion and comment.

<u>References</u>

UM86 = paper presented at the International Workshop on User Modelling, Maria Laach, August 1986.

[AIBC86]
Airenti G, Bara BG and Colombetti M (1986). Modelling the Relationship between User and System in Man-Machine Dialogue, UM86.

[BEEK86]
van Beek PG (1986). A Model for User-Specific Explanations for Expert Systems, Research Report CS-86-42, Department of Computer Science, University of Waterloo.

[BECO86]
van Beek PGF and Cohen R (1986?). Towards User Specific Explanations from Expert Systems, Department of Computer Science, University of Waterloo.

[BELK84]
Belkin NJ (1984). Cognitive Models and Information Transfer, <u>Social Science Information Studies</u>, 4, pp.111-129.

[BROO85]
Brooks HM, Daniels PJ and Belkin NJ (1985). Problem Descriptions and User Models: Developing an Intelligent Interface for Document Retrieval Systems, <u>Informatics 8: Advances in Intelligent Retrieval</u>, London: Aslib.

[BROO86]
Brooks HM (1986). <u>An Intelligent Interface for Document Retrieval Systems: Developing the Problem Description and Retrieval Strategy Components</u>, PhD Thesis, Department of Information Science, City University.

[BUSH84]
Buchanan BG and Shortliffe EH (1984). <u>Rule-Based Expert Systems</u>, Reading, MA: Addison-Wesley.

[CRBE87]
Carberry S (1987). Plan Recognition and Its Use in Understanding Dialogue, in this volume.

[CHIN87]
Chin DN (1987). KNOME: Modelling What the User Knows in UC, in this volume.

[CLAN86]
Clancey WJ (1986). Qualitative Student Models, <u>Annual Review of Computer Science</u>, 1, pp.381-450.

[COJO87]
Cohen R and Jones M (1987). Incorporating User Models into Expert Systems for Educational Diagnosis, in this volume.

[DANI86a]
Daniels PJ (1986). Cognitive Models in Information Retrieval - An Evaluative Overview, Journal of Documentation, 42, pp.217-328.

[DANI86b]
Daniels PJ (1986). The User Modelling Function of an Intelligent Interface for Document Retrieval Systems, in Brookes B (ed) Intelligent Information Systems for the Information Society, Amsterdam: North-Holland.

[GOOD86]
Goodman BA (1986). Miscommunication and Plan Recognition, UM86.

[HOEP83]
Hoeppner W et al (1983). Beyond Domain Independence: Experience with the Development of a German Language Access System to Highly Diverse Background Systems, Proceedings of the Eighth International Joint Conference on Artificial Intelligence, pp.588-594.

[JAME87]
Jameson A (1987). A Model of Impression Monitoring and Image Maintenance, in this volume.

[KOBS87]
Kobsa A (1987). A Taxonomy of Beliefs and Goals for User Models in Dialogue Systems, in this volume.

[LECA87]
Lehman, JF and Carbonell JG (1987). Learning the User's Language: A Step Towards Automated Creation of User Models, in this volume.

[LITM85]
Litman DJ (1985). Plan Recognition and Discourse Analysis: An Integrated Approach for Understanding Dialogue, PhD Thesis, Report TR70, Department of Computer Science, University of Rochester.

[MCCO87]
McCoy KF (1987). Reasoning on a User Model to Respond to Misconceptions, in this volume.

[MORI87]
Morik K (1987). Modelling the User's Wants, in this volume.

[PARI87]
Paris CL (1987). Tailoring Object Descriptions to a User's Level of Expertise, in this

volume.

[POLL86]
Pollack ME (1986). A Model of Plan Inference that Distinguishes between the Beliefs of Actors and Observers, Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics, pp.207-214.

[RICH79]
Rich E (1979). User Modelling via Stereotypes, Cognitive Science, 3, pp.329-354.

[RICH83]
Rich E (1983). Users are Individuals: Individualising User Models, International Journal of Man-Machine Studies, 18, pp.199-214.

[SCHU86]
Schuster E (1986). Establishing the Relationship between Discourse Models and User Models, Department of Computer and Information Science, University of Pennsylvania.

[SLBR82]
Sleeman D and Brown JS (eds) (1982). Intelligent Tutoring Systems, New York: Academic Press.

[SLEE85]
Sleeman D (1985). UMFE: A User Modelling Front-End System, International Journal of Man-Machine Studies 23, pp.71-88.

[SPAR84]
Sparck Jones K (1984). User Models and Expert Systems. Technical Report 61, Computer Laboratory, University of Cambridge.

[SPAR85]
Sparck Jones K (1986). Issues in User Modelling for Expert Systems, in Cohn A and Thomas J (eds) Artificial Intelligence and its Applications, Chichester: John Wiley and Sons, pp.183-195.

[UMDM87]
Papers on the relationship between user models and discourse models.

[WAKO86]
Wahlster W and Kobsa A (1986). Dialogue-Based User Models, Proceedings of the IEEE.

[WILE86]
Wilensky R et al (1986). UC - A Progress Report, Report UCB/CSD 87/303, Division of Computer Science, University of California, Berkeley.

[WIBI83]

Wilks Y and Bien J (1983). Beliefs, Points of View, and Multiple Environments, Cognitive Science, 7, pp.95-119.

1. PATIENT ONLY

e.g. automatic surveillance of a human being

2. AGENT ONLY

e.g. human operator of an industrial design system

3. PATIENT AND AGENT, DIFFERENT

e.g. medical diagnosis system with nurse agent and sick person patient

4. PATIENT AND AGENT, SAME

e.g. diet advice system with customer agent and patient

Figure 1

User role possibilities

OBJECTIVE / SUBJECTIVE

e.g. temperature reading / feeling worried

STATIC / DYNAMIC (for session)

e.g. political belief / dress purchase preference

DECISION / NON-DECISION

e.g. temperature reading / age

dress purchase preference / political belief

<u>Figure 2a</u>
<u>Property type distinctions</u>

OBJECTIVE	P / A	salary
SUBJECTIVE	P / A	anxious
STATIC	P / A	sex
DYNAMIC	P / A	belief about relative costs of brick and concrete
DECISION	P	foreign language words known
NON-DECISION	P / A	age

<u>Figure 2b</u>
<u>Property role possibilities</u>

P = patient, A = agent, assumed different people

<u>Figure 2</u>

<u>Property types</u>

EFFECTIVENESS	
e.g. disease diagnosis	P
EFFICIENCY	
e.g. diagnostic test ordering	P
terminology of system questions	A
ACCEPTABILITY	
e.g. medical regime presentation	P
case citation	A

Figure 3

System functions and relevant models

A medical diagnosis system

P = patient model, A = agent model

PATIENT

		Decision	Non-decision
	Static, Objective		
P1		aged disabled female	
P2			poor-sighted Catholic
	Static, Subjective		
P3			
P4			honest suspicious
	Dynamic, Objective		
P5			
P6			
	Dynamic, Subjective		
P7			
P8			poorly informed age matters

Figure 4a
Patient properties

P1 Static, Objective, Decision

disabled	==>	EFFECTIVENESS	e.g. determinant of benefit
	==>	EFFICIENCY	e.g. drives search establishing benefit
	==>	ACCEPTABILITY	e.g. referred to in explanation of benefit

P8 Dynamic, Subjective, Non-decision

age matters	==>	EFFICIENCY	e.g. guides data gathering
	==>	ACCEPTABILITY	e.g. referred to in explanation of benefit

Figure 4b
Exploitation of user properties for system functions

AGENT

A2 <u>Static, Objective, Non-decision</u>

experienced	==> EFFICIENCY	e.g.	data gathering
	ACCEPTABILITY	e.g.	practice manual references

male ....

A4 <u>Static, Subjective, Non-decision</u>

chauvinist	==> EFFICIENCY	e.g.	data checking for bias
	ACCEPTABILITY	e.g.	explanation of benefit basis

<u>Figure 4c</u>
<u>Agent properties and uses</u>

<u>Figure 4</u>

<u>User models and their uses</u>

<u>A social security benefits system</u>

(patient and agent distinct)

			PATIENT	AGENT
EFFECTIVENESS				
Decision	(Static /) (Dynamic)	(Objective /) (Subjective)	X	
Non-decision	(S / D)	(O / S)	(X)	(X)
EFFICIENCY				
Decision	(S / D)	(O / S)	X	
Non-decision	(S / D)	(O / S)	X	X
ACCEPTABILITY				
Decision	(S / D)	(O / S)	X	
Non-decision	(S / D)	(O / S)	X	X

Figure 5

Factor-function relation possibilities

```
C1     I want to ask about Istanbul.
T1     So you want to go to Turkey?
C2     Yes. Can I go by train?

  train ==> (a)  economy
            (b)  fun
            (c)  no flying

T2a    Even the train costs $100.
T2b    How about the Orient Express?

  train ==>  ?
             ?
             ?

T2x    Do you want economy or fun ... ?

  train ==>  wants trains

T2y    The trains are ... (times, types, costs).
```

Figure 6a

Dialogue using decision properties

```
C1     I want to ask about Istanbul.
T1     So you want to go to Turkey?
C2     Yes. I want to go by train, but not on Sunday.

  not Sunday ==>  religious

T2p    Unfortunately the only train is a so-called
         Fun Special on Sundays.

NOT

T2q    There's a super Fun Special on Sundays.

  not Sunday ==>  not Sunday

T2r    The only available train is on Sundays.
```

Figure 6b

Dialogue using non-decision properties

Figure 6

Model using strategies

A travel agent system

```
1.   GOAL :  help son

     agent     :  mother
     patient   :  son
          wife :  son's view        = P3
                  real nature       = P1

2.   GOAL :  help son and wife

     agent     :  mother
     patient1  :  son
     patient2  :  wife
          wife :  son's view     = P1.3
                  real nature    = P1.1
          son  :  wife's view    = P2.3
                  real nature    = P2.1

3.   GOAL :  help mother

     agent     :  mother
     patient   :  mother
          son  :  mother's view = P3
                  real nature   = P1
          wife :  mother's view = P3
                  real nature   = P1
```

Figure 7

Goal perspectives and modelling implications

A marriage guidance system

P1 = objective decision properties, P3 = subjective decision properties