COGNITIVE SCIENCE 3, 329-354 (1979)

User Modeling via Stereotypes *

ELAINE RICH

The University of Texas at Austin

This paper addresses the problems that must be considered if computers are going to treat their users as individuals with distinct personalities, goals, and so forth. It first outlines the issues, and then proposes stereotypes as a useful mechanism for building models of individual users on the basis of a small amount of information about them. In order to build user models quickly, a large amount of uncertain knowledge must be incorporated into the models. The issue of how to resolve the conflicts that will arise among such inferences is discussed. A system, Grundy, is described that bunds models of its users, with the aid of stereotypes, and then exploits those models to guide it in its task, suggesting novels that people may find interesting. If stereotypes are to be useful to Grundy, they must accurately characterize the users of the system. Some techniques to modify stereotypes on the basis of experience are discussed. An analysis of Grundy's performance shows that its user models are effective in guiding its performance.

1. INTRODUCTION

Scene I

Someone walks into a large library, tells the librarian that he is interested in China, and asks for some books. What sort of books does the librarian recommend? That depends. Is the person a small child who just saw a TV show about China and wants to see more pictures of such an exotic place? Is the person a high school student doing a term paper? Or maybe a prospective tourist? Or a scholar interested in Eastern thought? Can the person read Chinese? The librarian needs to know these things before he can point the reader to the right books. Some of what he needs to know he'll know before he even thinks about it, such as the approximate age of the person. Some things he'll assume until he has evidence to the contrary, such as that the person does not read Chinese. To find out other things, he'll ask a few specific questions. Only after he has a rough model of the person he's talking to can he answer the question.

Scene II

The phone rings in the information division of a large pharmaceutical firm. The caller wants information about a drug the company makes. What sort of information should be provided? That depends. Is the caller a doctor, a patient, or an FDA representative? To provide the right information, the person answering the phone needs to know some facts about the caller.

The scenes above illustrate some kinds of situations in which people need to form a *model* of the person with whom they are dealing before they can behave appropriately. They form their model by collecting a few specific pieces of information and then invoking the knowledge they have about the groups to which the current person belongs, such as scholar or medical patient.

As computers come to be used by a larger number of people to help perform a great variety of tasks, it is becomming more and more important for them to be easy for people to use. There are many factors than can contribute to the ease of use of a computer system, ranging from the good design of input devices such as terminals to the speed of the system's response, the appropriateness of its response, and the naturalness of its input and output languages. Appropriate models of the users of a system can be an important contribution because they can simultaneously affect several of those factors, such as speed and quality of response and habitability of the language interface.

Most systems that interact with human users contain, even if only implicitly, some sort of model of the creatures they will be dealing with. For example, the central assumption behind the mini-max strategy used by game playing programs is that the opponent is trying to win and will therefore make his best possible move. Although it is almost always valid to assume that the opponent wants to win, it is much less often true that he will therefore make the best move. He may, and probably does, have idiosyncracies of style or strategy that preclude that. Of course, human players know that and watch for evidence of such quirks in their opponents.
The term "user model" can be used to mean several different things. The three major dimensions along which user models can be classified are:

- Are they models of a canonical user or are they models of individual users?
- Are they constructed explicitly by the user himself or are they abstracted by the system on the basis of the user's behavior?
- Do they contain short-term, highly specific information or longer-term, more general information?

There are other significant differences between the systems using these various types of user models, but they follow from these major differences. Systems with a single model of a canonical user can have that model permanently embedded within themselves, whereas systems with models of individual users must build the model on the fly, and so must make explicit the ways in which the model influences the performance of the overall system. Systems that extract the user model from the user's behavior must grapple seriously with the issues of incorrect or conflicting information arising from the inferences that led to the model. Systems with explicitly stated user information can, on the other hand, avoid many of those issues. Systems that deal with short-term knowledge must deal successfully with the problem of detecting when things change, while longer-term systems may be able to finesse that issue. But as these differences reduce to the three outlined above, they do not need to be focused on explicitly.

The work described here has concentrated on user models in one comer of this space. Models of individual users were built because of the facility for personalization that they provide that is lacking in a canonical user model. Implicitly constructed models were used because of the inherent inaccuracy and the annoyance of requiring users to construct their own models of themselves. The choice of position along the third dimension was much less clear cut. Both short-term knowledge (such as the topic currently being discussed or the goal now being pursued) and long-term knowledge (such as the level of the user's knowledge about the problem domain or goals that are likely to be pursued) can be important to the performance of an interactive system. But it appeared likely that different techniques would be appropriate for dealing with the two types of knowledge. Models of long-term user characteristics were chosen because very little work has been done with them, while short-term user models were already being explored, particularly in the context of natural language understanding. So the rest of this work will deal explicitly with the issues involved in individual user, implicit, long-term user modeling.

2. A BRIEF DISCUSSION OF IMPLICIT, INDIVIDUAL USER, LONG-TERM USER MODELING

A user modeler is significant only insofar as it communicates with a larger performance system in which it is embedded. This suggests that the issues that must be tackled in order to implement a user model and to integrate it usefully into a system are the structure of the user modeler itself and the links between the user modeler and the rest of the system.

2.1. The User Modeler and the Links into It

The first major issue that a user modeler must confront is how to build models of users. In building a model of a person, the obvious first step is to collect some facts about the person, as, for example, his age or his experience with computers. However useful this approach is, through, it is severely limited in its effectiveness for two reasons. One is that it might take a lot of questions to accumulate all the knowledge the system needs. The other reason is that the person may not always be able to provide accurate answers, either because he doesn't know, as for example, a student talking about his incorrect knowledge, or because he doesn't want to talk about it, as for example whether someone's grandmother is interested in reading books about homosexuality. To deal with situations such as these, systems must be able to infer information about their users based on a small number of explicitly stated facts.

There is a lot of evidence in the psychological literature to support the assertion that people are not reliable sources of information about themselves. Nisbett and Wilson (1979) describe a number of experiments that suggest that people are very bad at introspecting and then reporting on their cognitive processes. A different source of unreliability in self-description is suggested by the work of McGuire and Padawer-Singer (1976). They present evidence in support of the assertion that, when asked to describe themselves, people are heavily influenced by the social groups in which they find themselves. People seem to mention those characteristics that distinguish them from the other people in the group. This clearly introduces an element of unreliability into such descriptions since it is difficult to control or measure which groups a person is identifying with at a particular instant.

All of these experiments suggest that it is important that a user modeler not rely too heavily on answers to specific questions in building models of individuals. The most obvious way to build a user model without asking many questions is to make direct inferences from a user's behavior to a model of him. For example, if the base system performs some action which the user said was not what he wanted, the system should look at the reasons that it performed that action and conclude that it is likely that at least one of these reasons is wrong, at least for this user. But even this mechanism of drawing explicit inferences from the user's behavior is not always sufficient, since it may require many interactions between the system and the user to build a user model with enough information in it to be of significant use.

2.1.1. Stereotypes. A major technique people use to build models of other people very quickly is the evocation of *stereotypes,* or clusters of characteristics. So, for example, one might know that if someone is a judge, he or she is probably - over forty, well-educated, reasonably pro-establishment, fairly affluent, honest, and well-respected in the community. Although not all of these attributes are necessarily true of any particular judge, a person would tend to assume them until shown otherwise. Besides our everyday belief that we possess such stereotypes, there is also lore, in such areas as education, that something like stereotypes is necessary. (Highet, 1950).

In addition, there has been a lot of discussion in the psychology literature about the nature of sterotypes. For a review of much of this work, see Hamilton (1979). There are many theories about why people use stereotypes, but one of the most certain explanations is that people use stereotypes as a means for dealing with the fact that the world is far more complex than they can deal with without some form of simplification and categorization. One of the ways in which stereotypes help to simplify the world is that they have a strong effect on what characteristics of a person are attended to and remembered. As a result, they will tend to be confirmed by experience since potentially disconfirmatory evidence will be ignored. (See, for example, Cantor & Mischel, 1979, and Snyder & Uranowitz, 1978.)

When computer systems use stereotypes, they may be able to avoid some of the pitfalls that plague the human use of stereotypes. Computers will not develop any emotional attachment to their stereotypes and so, to whatever extent that contributes to the inadequacies of stereotypes, computers will be immune. There are, however, very direct limitations on the cognitive ability of any individual computer system, although these limitations are not identical to those of people. So we shall have to discover to what extent a computer system's stereotypes can be altered on the basis of experience and to what extent perception is sufficiently biased that such learning is not possible. On the other hand, computer systems are definitely superior to people in that they can explicitly use stereotypes in certain situations and not in others.

Of course, a computer cannot help but be at a loss compared to a human in quickly sizing up a person on the basis of superficial characteristics, if for no other reason than that it can neither see him (to determine his age, type of clothing, or sex) nor hear him (to determine where he comes from or his self-assurance). But it does have available, because of the way communication between man and machine takes place, some information that a person would not have. For example, if a user types quickly, he is probably either a secretary, a computer scientist, a newspaper reporter, an author, etc.

In order for a system to be able to use stereotypes effectively, it must have two kinds of information. It must know the stereotypes themselves - the collections of characteristics or *facets.* A user can be characterized by a set of facets, each of which contains a value (or values). The particular set of facets used will, of course, be determined by the domain and purpose of the system as a whole, but could include such things as level of expertise with the system, specific concepts dealt with by the system that are of particular interest, or specific system tasks that are of particular concern. Stereotypes are simply collections of facet-value combinations that describe groups of system users.

A system that is going to use stereotypes must also know about a set of *triggers* - those events whose occurrence signals the appropriateness of particular stereotypes. So, for example, if a user immediately uses an advanced construct in a system, that should serve as a trigger for the "expert-user" stereotype, which should then be *activated.* When a stereotype is activated, the predictions it makes about various characteristics should be incorporated into the model of the user.

Stereotypes, as used in this context, are also very similar to scripts (Schank, 1977), frames (Minsky, 1975), and schemas (Bobrow & Norman, 1975). All these mechanisms provide a means of representing partial descriptions of frequently occurring situations. This is done by specifying a collection of the significant aspects of a situation while omitting mention of the nonsignificant ones. In the case of scripts, the aspects mentioned are events. In the case of stereotypes, the aspects are personality characteristics. Looked at in this light, stereotypes are just another instance of the increasingly well-understood phenomenon that people interpret many of their experiences in terms of stored expectations for those experiences. Another way of looking at the wide applicability of this sort of structure is that although all situations are not identical, neither are they all completely dissimilar.

As suggested by the need for the generalization relationship between stereotypes, the stereotypes in use by a given system may range from very general ones, such as man or woman, which might be appropriate in a wide variety of domains, to very specific ones, such as reader of good mysteries, which are appropriate only for specific systems and domains. All of these can contribute to the system's model of an individual user.

The use of stereotypes, then, when combined with the ability to record explicit statements by the user about himself and to make direct inferences about a user from his behavior, may provide a powerful mechanism for creating computer systems that can react differently to different users.

2.1.2. The USS. The core of the user modeler is its model of an individual user. This model, called the *User Synopsis* or USS, is built by combining the direct information provided by the user, direct inferences from the user's actions, and predictions made by stereotypes that are deemed appropriate for this user. The information in the USS can then be used to guide the rest of the system.

2.1.3. Probabilistic Inference. Very little of the information in the USS is known for sure to be true, since it arises almost entirely from probabilistic rather than certain inferences. Some behaviors suggest specific characteristics, while others suggest the appropriateness of a particular stereotype, which then suggests particular characteristics. This is very similar to the situation that arises in medical diagnosis, as described in Shortliffe, (1976), and the way these inferences should be handled by a user modeler is very similar to the way they are handled in MYCIN.
Because of the uncertain nature of behaviorally inferred knowledge, a user modeling system must associate with each piece of information that it possesses a rating representing how confident it is that that information is correct. So stereotypes actually consist of a set of (attribute, value, rating) triples. The situation in the User Synopsis is even more complicated. In order to be able to resolve the conflicts that will inevitably arise, the system should remember why it believes the things it believes. So the USS consists of a set of (attribute, value, rating, justifications) quadruples.

Notice that both stereotypes and the USS contain ratings. In both cases, the ratings indicate confidence in the associated fact. In the case of stereotypes, the ratings represent confidence that any person who fits the stereotype will exhibit a particular characteristic. In the case of the USS, the ratings indicate the confidence the system has in its belief that a particular fact is true of a particular user.

2.2. Links Out of the User Modeler

The other major aspect of user modeling is how the user models, once they are built, are used to guide the performance of the rest of the system. In other words, how are the links out of the user modeler to the base system, the language understander, and the language generator used? The solution to this problem depends heavily on the structure of the particular system in question. Some general guidelines are discussed in Rich (1979).

If both of these issues, compiling a user model and using that model effectively, can be dealt with successfully, performance systems should be able to significantly improve the level of service they provide to their users.

2.3. Choice of Task Domain

In order to have a forum in which to explore these issues and to be able to make a convincing case that some solutions have been found, it was deemed necessary to tackle the issues in the context of a specific task domain and actually to build a system into which the solutions to the issues could be incorporated. Although the actual task chosen is not important, it is important that it satisfy the two principle criteria that determine appropriateness of this sort of user model. These criteria are:

- The system should be one that will be used by a heterogenous group of users, If the set of possible users is highly homogeneous, then it is much more efficient to simply build in a model of a canonical user.

- The system should be one that has a nontrivial amount of flexibility in its operation. If it does not, then it does not need the additional information of a user model to tell it what to do.

The task that was chosen for this experiment is to suggest to users novels that they might like to read. In other words, to play librarian. This task meets both of the above criteria. Almost everyone reads novels at least occasionally. And the task is highly flexible; anyone of a large number of books could be selected. So. this looked like a good forum for the exploration of the issues raised by individual user modeling.

3. AN OVERVIEW OF GRUNDY

The main purpose of this section is to provide an overview of Grundy, a system that plays librarian. The best way to view the task of the system is that is should come as close as possible to simulating the performance of a good librarian at a public library when a patron walks in the door looking for something to read over the weekend. If the librarian already knows him, she will be able to provide some suggestions right away. This will happen almost all the time, for example, in a small town library where the librarian knows everyone. If, on the other hand, the librarian doesn't know him, she will first size him up quickly. How old is he? How well educated does he appear to be? Is he a man or a woman? Even those clues will probably not be enough, though, so she will have to ask a few questions. Maybe she'll ask him what the last book that he read and liked was. Or maybe who his favorite author is. Then she'll start recommending books. Based on his reaction to the first few suggestions, she'll modify her view of him if necessary and continue making suggestions until he is satisfied. Grundy will try to do as well as it can it performing this task. Clearly, it has several limitations. It cannot see the user. That eliminates many of the librarian's initial clues. And many questions that the librarian might ask would be very difficult for the system to use effectively. In order to get any information out of the user's response to the question about the last book he read, the system would have to know a much larger number of books than it is ever likely to know. So, although the system will try to perform the same task as the librarian, it will not try to copy exactly her performance. Rather it will try to perform the task using the techniques best suited to the system's strengths and weaknesses.

The rest of this section attempts to provide an overview of Grundy, and simultaneously an exposition of some of the issues involved in user modeling. It presents a dialogue between a user and the system, annotated to explain what the system knows and what it is trying to do at each stage. The lines printed in upper case were typed by the system; those in lower case were typed by the user.

PLEASE TYPE IN YOUR NAME

John Greene

YOU HAVEN'T USED THIS SYSTEM BEFORE, HAVE YOU?

no

> The system can now activate two stereotypes. The first, ANY-PERSON, will be activated for any new user. It predicts middle of the road values along almost all dimensions. Of course, it predicts them with very low ratings so that they can immediately be superseded by specific information. It can also activate the MAN stereotype on the basis of the name "John." It will activate both of these stereotypes and construct an image of its total current knowledge of the user before it tries to do anything else.
>
> Then it will try its first task designed to tell it as much as possible about the user.

OK

I'D LIKE TO KNOW WHAT SORT OF PERSON YOU THINK YOU ARE. PLEASE TYPE IN A FEW SINGLE WORDS THAT YOU THINK CHARACTERIZE THE MOST IMPORTANT ASPECTS OF YOURSELF

unconventional open direct honest humorous persistent adventurous

I DON'T KNOW THE WORD UNCONVENTIONAL
I DON'T KNOW THE WORD OPENI DON'T KNOW THE WORD DIRECT

WOULD YOU LIKE TO TRY SOME MORE WORDS?

> Since it didn't get the complete picture the user was trying to convey, it gives him another chance.

yes

OK GO AHEAD

relaxed jewish friendly playful thoughtful eccentric

> At this point, Grundy believes it knows quite a lot about John. Because he is male, it believes he has a fairly high tolerance for violence and suffering, as well as a preference for thrill, suspense, fast plots, and a negative interest in romance. Because of stereotypes evoked by "humorous," "friendly," and "playful" it thinks he's quite optimistic. It thinks he's interested in conflicts such as art-vs-science, exploration, God, and what it means to be human. It thinks he has particular sympathy for such strengths as kindness, courage, honesty, and perceptiveness.

HAVE YOU EVER READ EARTH ANGELS BY SUSAN CAHILL?

> To select a book to recommend, it picks one of the facets of John which has both a non-middle-of-the-road value and a fairly high rating. Then it rates all the books suggested by that facet by comparing them to John along all dimensions. It got this suggestion from the belief that he is politically liberal (which it believes from "Jewish").

no

> Grundy selects, from the things it knows about the book, those things that it thinks will make John like it.

YOU MIGHT LIKE IT.
IT IS ABOUT A GIRL WHO GOES TO CATHOLIC ELEMENTARY SCHOOL,HIGH SCHOOL, AND COLLEGE, THEN GOES ON TO BECOME A NUN, BUT FINALLY BACKS OUT A FEW DAYS BEFORE FINAL VOWS.
A MAJOR CONFLICT IN IT IS THE RELATIONSHIP OF MAN TO GOD.
ONE OF THE MAJOR STRENGTHS OF ONE OF THE MAIN CHARACTERS IS INTELLIGENCE.
ONE OF THE MAJOR STRENGTHS OF ONE OF THE MAIN CHARACTERS IS INDEPENDENCE.
IT IS POLITICALLY LIBERAL.
IT HAS A BIT OF ROMANCE.
IT HAS A MEDIUM-FAST MOVING PLOT.
DOES THAT SOUND GOOD?

no

> Grundy needs to find out where it went wrong. It knows the reasons it thought John would like the book. It asks about them to see which if any of them is incorrect. The order in which it asks about them is determined by an a *priori* rating of how likely that reason is to be the problem and the current rating of the related characteristics of John. It asks about things it is least sure of first.

WAS IT BECAUSE OF THE DEPTH OF THE CHARACTER DEVELOPMENT? (EITHER TOO MUCH OR NOT ENOUGH)

no

WAS IT BECAUSE OF THE AMOUNT OF ROMANCE? (EITHER TOO MUCH OR NOT ENOUGH)

no

WHERE YOU BOTHERED BY THE POLITICAL POSITION ESPOUSED BY THE BOOK?

no

AREN'T YOU INTERESTED IN THE RELATIONSHIP OF MAN TO GOD?

no

> It has found the problem. There may, of course, be other problems, but since a primary goal of a system like this is to ask as few questions as possible, it stops asking now. If there are other problems, it will discover them later.
>
> Although John is not interested in the relationship of man to God, it was not unreasonable for Grundy to have throught he was, since he said he was Jewish, and the mention of a religious group frequently (though not always-it may refer to an ethnic group) indicates an interest in religious ideas. But now Grundy changes its belief to reflect this more specific information.

DO YOU WANT SOME MORE SUGGESTIONS?

yes

HAVE YOU EVER READ ZEN AND THE ART OF MOTORCYCLE MAINTENANCE BY ROBERT PIRSIG?

> It selected this based on the personality characteristic "living in a separate world," which it got from "eccentic." It also matched on several other dimensions, such as an interest in the conflict of art-vs-science.

yes

DID YOU LIKE IT?

yes

> Grundy now has some additional information about John. It goes through all the characteristics of this book and makes direct inferences about John based on them.

HAVE YOU EVER READ THE CRYING OF LOT 49 BY THOMAS PYNCHON?

> It also generated this one based on "separate world."

no

YOU MIGHT LIKE IT.
IT IS A HUMOROUS WEll-WRITTEN NOVEL OF SOUTHERN CALIFORNIAN INSANITY. ONE OF THE MAJOR WEAKNESSES OF ONE OF THE MAIN CHARACTERS IS LIVING IN A DIFFERENT WORLD FROM THE REST OF MANKIND.
ONE OF THE MAJOR STRENGTHS OF ONE OF THE MAIN CHARACTERS IS INTELLIGENCE.
ONE OF THE MAJOR STRENGTHS OF ONE OF THE MAIN CHARACTERS IS ABILITY TO COPE WITH WHATEVER HAPPENS.
IT HAS ALMOST NO ROMANCE.
IT IS MILDLY PHILOSOPHICAL.
DOES THAT SOUND GOOD?

yes

GOOD
DO YOU WANT SOME MORE SUGGESTIONS?

no

OK THEN. BYE.

4. STEREOTYPES IN GRUNDY

4.1. Stereotypes and Stereotype Structures

A *stereotype,* as mentioned above, is a collection of frequently occuring characteristics of users. Figure 1 contains some sample stereotypes used by Grundy. A bit of explanation of the values contained in those stereotypes is in order. The simple linear scales (Thrill, Tolerate-sex, Tolerate-violence, Tolerate-suffering, Romance, Education, Sex--open, and Piety) have values ranging from - 5 to +5. The numbers associated with each value of the list-valued facets (e.g. Strengths, Genres) also range from -5 to 5. Ratings range from 0 to 1000, with higher numbers indicating greater degrees of certainty.

FACET		*VALUE*	*RATING*
Activated-by		*Athletic-w-trig*	
Genl		*ANY-PERSON*	
Motivations			
	Excite	3	600
Interests			
	Sports	4	800
Thrill		5	700
Tolerate-violence		4	600
Romance		-5	500
Education		-2	500
Tolerate-suffering		4	600
Strengths			
	Physical-strength	4	900
	Perseverance	3	600
		SPORTS-PERSON	

Activated-by		*Feminist-w-trig*	
Genl		*ANY-PERSON*	
Genres			
	Woman	3	700
Politics		*liberal*	700
Sex-open		5	900
Piety		-5	800
Political-causes			
	Women	5	1000
Conflicts			
	Sex-roles	4	900
	Upbringing	3	800
Tolerate-sex		5	700
Strengths			
	Perseverance	3	600
	Independence	3	600
Triggers		Fem-woman-trig	
		FEMINIST	

Figure 1. Same Sample Stereotypes.

Stereotypes can contain any number of characteristics; they will almost never contain values for all the characteristics that the system understands. For example, the FEMINIST stereotype contains values for the facets PIETY and TOLERATE-SEX, which the SPORTS-PERSON stereotype does not. But SPORTS-PERSON contains values for the facets THRILL and TOLERATE-SUFFERING, which FEMINIST does not. But since usually more than one stereotype will be active for a particular user, it will be possible to build up a fairly complete picture of a user even though the individual stereotypes contain only partial information.

The stereotypes about which the system knows are arranged in a directed acyclic graph (DAG) formed by the partial ordering relation "generalization of." This is important because it allows information not to have to be represented identically in many different stereotypes. For example, Figure 2 shows a piece of the stereotype DAG containing the stereotypes for religious people. It represents the fact that there are many characteristics shared by Christians, regardless of denomination, and others shared by all religious people regardless of religion. This generality DAG can be used to focus the attention of the system on the events that are the most likely to be of significance.

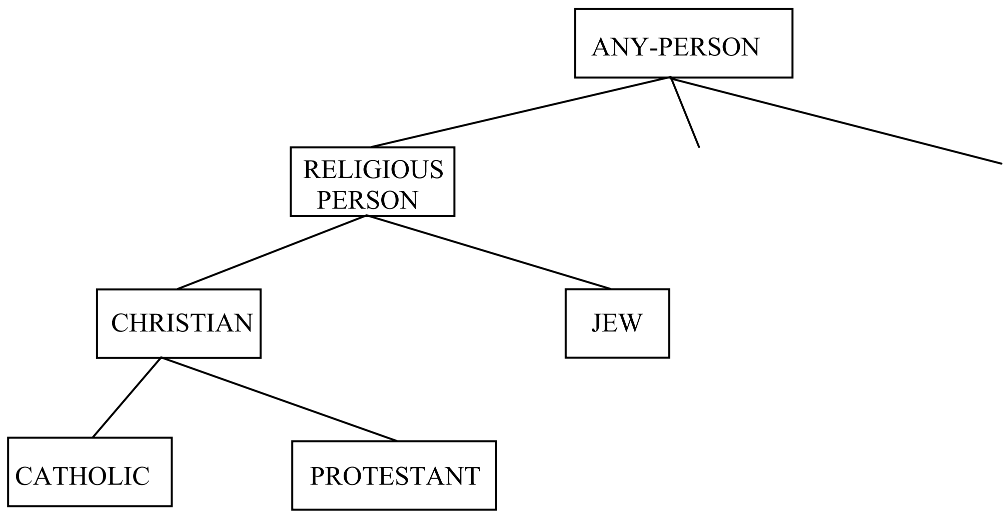

Figure 2. A Piece of the Stereotype DAG.

The most general node of any stereotype structure is the stereotype ANY-PERSON. It provides default values for all facets, but provides them with very low ratings. These values can be used to prevent the system from doing anything that might be offensive until it learns enough about the individual user to know his particular inclinations. It is a cross between a model of a canonical user and a lower bound on the user's tolerance for things.

This notion of a canonical or representative person with respect to whom individual characteristics can be compared has been useful in at least one other system. Carbonell (1979) describes a system for describing personality traits as goal structures and then using the personality traits to understand stories. He describes personality traits as changes to the goal structure one would attribute to an average person. Thus a miserly person would rate the "preserve money" goal higher than another person would. This mechanism of representing individuals as departures from a norm has two significant advantages. It provides a frame of reference and it corresponds well to the conversational notion of default. If someone's "preserve money" goal is exactly what one would normally expect, a storyteller (or a system user or a conversationalist) will not usually mention it. If it is mentioned, it will be because the value is different from what one would otherwise expect. In either case, the system's model is likely to be correct.

Although it is usually true that stereotypes will share the values of their generalizations, it is not always so. For example, although religious people can reasonably be expected to be quite pious, Jewish people are much less likely to be so than other religious groups. So the information in the JEWISH stereotype overrides some of the information in the RELIGIOUS stereotype. The algorithm for incorporating stereotype information into the USS must then handle this override. The most extreme case of overriding of general information by specific occurs with respect to the ANY-PERSON stereotype, most of whose information will eventually be overridden by other stereotypes once the system gets to know the user. This use of general information accompanied by specific overrides when necessary has also been useful in other systems (see Carbonell, 1979 or Charniak, 1977).

It is important that the structure used to represent the generality relationships among stereotypes be a DAG rather than a tree to make it possible to represent the fact that some stereotypes may have more than one immediate generalization. For example, Grundy has a stereotype of a MARXIST. That stereotype has two generalizations, each representing different aspects of a Marxist's position: LEFT-WING-PERSON and RADICAL. It is unfortunate that this structure is necessary since it complicates many of the algorithms that the user modeler uses to maintain the user synopsis.

The generality structure of stereotypes in the user modeler's data base is merely a DAG and not a lattice because, although there is a least upper bound (ANY-PERSON), there is no greatest lower bound. The stereotypes that are active at any particular time do, however, form a lattice with the user synopsis representing a single user, as the greatest lower bound. This has no effect on the algorithms used to move around in the structure, however.

4.2. Activating Stereotypes

Stereotypes can be activated as a result of almost any kind of action. A stereotype is activated by instantiating one of its triggers. A trigger is an object associated with a particular situation. In addition to its name, it contains the name of the stereotype to be activated and the rating (a number between 0 and 1000) to be assigned to the stereotype. The rating represents the probability that the stereotype is actually appropriate in the particular situation. Figure 3 shows a few representative triggers. There can be many triggers for the same stereotype. This is necessary for two reasons. The ratings to be assigned to the stereotype may differ for different situations in which it is activated and it is useful to be able to tell what situation caused a stereotype to be activated.

Each currently active stereotype has associated with it a list of the triggers that have activated it. If a trigger is instantiated, the action will depend on whether that trigger has been instantiated before. If it has, no further action will result. If it has not, then the stereotype will be activated. Here again, the action depends on what has already happened. If the stereotype has not previously been activated, it is activated now. If it has and is still active, then two things happen: its rating increases and both the current trigger and the earlier triggers get their ratings increased to reflect the confirmation they have just received. If, on the other hand, the stereotype is no longer active, the situation must be re-examined on the basis of the new information to decide whether the balance of the evidence is in favor of or opposed to the stereotype.

NO-TV-TRIG		(Besides asking for characteristic words, the other thing Grundy can do to find out about users is to ask them about TV. This trigger is activated if the user says he does not watch TV.)
FACET	*VALUE*	
Stereotype	NON- TV-PERSON	(this stereotype suggests that a person is likely to be educated and serious)
Rating	800	(maybe this person is not really a non-TV person. Maybe he just can't afford to buy one.)
SCI-ED-TRIG		(This trigger is associated with the SCIENTIST stereotype and will be activated whenever the SCIENTIST stereotype is activated.)
FACET	*VALUE*	
Stereotype	EDUCATED-PERSON	
Rating	900	
Reasons	SCIENTIST	

Figure 3. Some Sample Triggers.

The major action involved in the activation of a stereotype is updating the user synopsis to reflect the new information. This involves going through each of the facets for which information is predicted by the stereotype, seeing what information is already there and why, and deciding what the value should be, based on both the new and the old information. If the stereotype was not already active, this will always be done. If it was, then it may be necessary to propagate its change in rating to all the predictions it makes. But it is only necessary to do that if the rating change is significant. So its new rating will be compared to its rating the last time changes were propagated. If the difference is greater than an appropriate threshold the changes will be propagated.

The amount a stereotype influences the value of a facet is a function of both the rating of the stereotype and the strength with which the stereotype predicts the value of the facet, as well as the other evidence already present about the facet. There is no explicit priority given to earlier inferences over later ones, but because the earlier ones will tend to determine the course of later interactions, questions, and observations, they will end up having a more pronounced effect than will later inferences.

Stereotypes can be activated whenever their triggers are instantiated. Triggers can be instantiated almost any time. The three major types of situations in which they are instantiated in Grundy are:

- As part of a specific task designed to solicit information about the user, such as the tasks that ask about descriptive words or TV watching. For example, associated with each word that the system understands to characterize users is a list of triggers to be instantiated if that word is used.
- As a result of a facet in the USS being given a particular value. For example, if the gender facet is assigned the value male, then the MAN stereotype will be activated as a result of a trigger associated with the facet.
- As a result of another stereotype being activated. When a stereotype is activated all of its generalizations are also activated. This means that information that is common to all the specializations of a given stereotype can be represented just once in the most general stereotype to which it applies. In addition, stereotypes may contain triggers that will be instantiated whenever the stereotype is. The stereotypes indicated by the triggers will then be activated too. This is useful if the fact that someone is a member of one class suggests that he is also a member of another. For example, the scientist stereotype suggests the atheist stereotype. The ratings of the additional stereotypes will be a function of the rating of the stereotype that caused them to be invoked and the rating of the trigger that activated them. In other words, it represents the probability that the first stereotype is appropriate and, given that it is, that the other stereotype is too.

There are other situations that could arise in other contexts, for example:

- In a system with a natural language front end, stereotypes could be triggered by the use of arbitrary words, phrases, or grammatical constructions.
- In a system with a specific set of commands that the user can issue, stereotypes can be triggered by the use of particular commands.
- Stereotypes could be triggered by any other information that the system has about the user. For example, his account number might indicate his status in some way.

The thing that all of these situations have in common is that they are specific situations that the system must recognize individually and handle. So to associate stereotypes with the situations, all that is necessary is to build some structure in which the appropriate triggers are linked to the relevant situation (another piece of information for each entry in the lexicon, or command table, or user account list, or facet description, for example). Then whenever the system processes a situation, it also checks for associated triggers.

4.3. Combining Stereotypes to Form the USS

As mentioned earlier, the central structure in a user modeling system is the USS, the model of an individual user. The USS is built by combining all of the knowledge the system has or can infer about the user. Figure 4 shows an example of a USS built by Grundy. It was constructed on the basis of the user having told Grundy that she is a feminist and an intellectual. Notice that some facets contain values predicted by a single stereotype, while others have been affected by more than one. Sometimes (as in the case of SEX-OPEN) the stereotypes reinforced each other's predictions. Other times (as in the case of PIETY) the values predicted by the various stereotypes conflicted. When such conflicts occur, an appropriate value must be computed for the USS, and its associated rating must be computed to indicate the lack of confidence that the conflict suggests. In the case of linear facets such as PIETY, the value stored in the USS will be a weighted average (weighted by rating) of the values predicted by all the relevant stereotypes. For other types of facets, other conflict resolution mechanisms are appropriate. For some symbolic facets, the value selected will be the most specific object that is a generalization of all the suggested values. As the complexity of the values which facets are allowed to assume grows, the sophistication of the conflict-resolution rules will also have to grow. It is significant that even such simple rules as these enabled Grundy to perform quite well (see Section 7. 1) .

FACET	*VALUE*	*RA*TING	*JUSTIFICATIONS*
Gender	female	1000	Inference-female name WOMAN
Nationality	USA	100	ANY-PERSON
Education	5	900	INTELLECTUAL
Seriousness	5	800	INTELLECTUAL
Piety	-3	423	WOMAN FEMINIST INTELLECTUAL
Politics	Liberal	910	FEMINIST INTELLECTUAL
Tolerate-sex	5	700	FEMINIST
Tolerate-violence	-5	597	WOMAN
Tolerate-suffering	-5	597	WOMAN
Sex-open	5	960	FEMINIST INTELLECTUAL
Personalities	4	646	WOMAN
Opt-pes	0	100	ANY-PERSON
Plot-intr	0	100	ANY-PERSON
Plot-speed	-2	475	EDUCATED-PERSON
Suspense	0	100	ANY-PERSON
Thrill	-4	839	WOMAN INTELLECTUAL
Romance	3	696	WOMAN
Confusion	3	570	EDUCATED-PERSON
Real-fant	0	100	ANY-PERSON
Comedy	0	100	ANY-PERSON
Genres			
Literature	4	700	INTELLECTUAL
Women	3	700	FEMINIST
Political-causes			
Women	5	1000	FEMINIST
Strengths			
Perceptiveness	2	570	EDUCATED-PERSON
Intelligence	4	800	INTElIECTUAL
Independence	3	600	FEMINIST
Perseverance	3	600	FEMINIST
Sympathy	2	497	WOMAN
Kindness	2	497	WOMAN
Weaknesses			
Reason	3	600	INTElIECTUAL
Conflicts			
Difference	3	600	INTElIECTUAL
Upbringing	3	800	FEMINIST
Sex-roles	4	900	FEMINIST
Propriety	2	497	WOMAN
Love	2	497	WOMAN
Motivations			
Learn	4	700	INTELLECTUAL
Interests			
Ideas	4	900	INTELLECTUAL

Figure 4. A Sample User Synopsis

5. EXPLOITING THE USER MODEL

5.1. How Grundy Exploits the USS

The raison d 'etre of a user model is to provide guidance to an underlying system in the performance of its task. Grundy's task is to suggest appropriate novels and the USS guides the execution of that task in three ways. It enables the system to focus on a fairly small set of potentially appropriate books; it provides a way of evaluating each of those books in detail in order to select the best one to recommend; and it provides a way of selecting which of the selected book's many attributes would be of interest to a particular user.
For each recommendation it makes, Grundy performs the following set of actions:

1. It builds a set of all of the facets of the USS which have a fairly high rating and a non-middle-of-the-road value. This is the set of salient characteristics of the user.
2. It randomly selects one facet from the set. It then builds a set of books to consider by selecting those books with specific attributes suggested by the chosen user characteristic. For example, if the EDUCATION facet were being used and it contained a high value, then books with high PHILOSOPHY and LITERARY-MERIT values would be included in the initial set. This step constitutes the attention-focusing function of the USS.
3. Once the set of possible books is assembled, each of them is compared facet by facet to the USS in order to decide which book is the best match. This step exploits the evaluation capability of the USS.
4. Having selected a book to recommend, Grundy must decide what to say to the user about the book, since it knows much more than most people would want to hear. The USS is used again, this time to filter the most interesting of the book's features so that they can be presented

From this discussion it should be clear that a user model can serve a variety of functions even to an underlying system as simple as Grundy.

5.2. How Grundy Does Not Exploit the USS

Because of the simplicity of the overall Grundy system, there are several areas in
which it does not exploit the USS. These include:

- Deciding how to phrase things (for example, as a function of the user's educational level). Grundy's sentence generation mechanisms are so simple that they have no use for any outside information.
- Deciding what questions to ask the user about himself. This is a problem because most of the time Grundy gets enough information from its first question ("Tell me some words that describe yourself."). In a few limited situations, additional information will be sought specifically to correct what Grundy perceives to be a misconception on its part about the user. If, for example, a book is rejected by the user because of its political position and Grundy believed that the user would like the book because of its politics, it will ask the user a direct question in order to ascertain for sure his political position.
- Interpreting things the user says to the system. Because of the limited natural language facility possessed by Grundy, statements by the user are restricted to single words. There is little room for exploitation of additional knowledge of any sort in interpreting such simple statements.

These limitations in the ways Grundy exploits the user model arise from limitations within the Grundy system itself rather than from inherent limitations in the applicability of user models. See Rich (1979), for a fuller discussion of these issues, particularly the potential for the use of user models in conjunction with a natural language understander.

6. LEARNING

6.1. Hierarchical Memory

The knowledge in a system like Grundy is divided into three sections, as a function of its scope of relevance. First there is the knowledge about the domain in which the system operates. That information is true and relevant to all users and all dialogues (although of course not all of it will be relevant to anyone dialogue, it could be). At the next level is the knowledge about an individual user. That knowledge is relevant each time that user interacts with the system. At the narrowest level is the information that is relevant to this dialogue only.

In order to provide access to the appropriate parts of the system's knowledge at the appropriate times, it is necessary to arrange the system's knowledge hierarchically. Whenever the system wants to access some information, it looks first to see if that information exists in the dialogue-specific memory. If it is not there, it looks in the user-specific memory. And if it is not there, it looks in the universal memory. This hierarchical search is embedded in the data-base access routines, so it is transparent to the rest of the system. To record information, it is necessary to specify which memory the information should be stored in. This mechanism has the advantage that it is possible to override universal knowledge with more specific knowledge without permanently erasing the global knowledge.

At the end of the session, all the information in the user memory is written into a file corresponding to the user. That file will be retrieved whenever that user returns to the system. The universal memory can be written out to form a new universal file. And the dialogue memory can be forgotten.

6.2. Adapting Stereotypes

In addition to learning models of individual users, it is important for a user modeling system to be able to modify its data base of stereotypes. This importance arises from the almost total lack of real data on which to base the initial construction of the stereotypes. To bias a system forever with the prejudices of its designer would place an undue burden on it (although even those pure prejudices would probably be significantly better than nothing). In addition, computers have a significant advantage over people with respect to the user of stereotypes-they are not emotionally committed to them and thus are able to change them as warranted by experience. This advantage certainly ought to be exploited.

Almost all events that occur when a stereotype is active can shed some light on the correctness of the stereotype and its triggers. If the user behaves in a way predicted by the stereotype, it lends confirmation both to the appropriateness of that prediction and to the appropriateness of the triggers that caused the stereotype to be activated. If, on the other hand, the user exhibits a behavior that conflicts with a prediction of an active stereotype, then either the prediction is inappropriate or the triggers that caused the stereotype to be activated are inappropriate (or possibly both). It should be noted that this does not necessarily mean that the stereotype or the triggers are wrong. They may correctly predict a strong likelihood to which the current user is merely one of the exceptions. When a conflict does occur, however, the system cannot tell exactly why it arose. It can merely conclude that there is less evidence for the predictions than there was before the conflict. Because it also cannot know which of its inferences is the source of the problem (either the triggers or the stereotype itself or both), it must unreinforce all of them by lowering their ratings. If some of the inferences are in fact correct, they will be reinforced in other situations and so will be retained by the system. The inferences that are not correct will eventually be eliminated.[2] It is because of this spurious unreinforcement that may occur that it is particularly necessary that reinforcement occur whenever a prediction is borne out by experience. Otherwise it might be reasonable to simply do nothing when the system is already on the right track and to modify the data base only when it is suggesting erroneous conclusions.

A major issue that arises in adapting the stereotypes and triggers on the basis of experience is how much to change them each time a significant event occurs. If they are changed too much then it is too easy for a few atypical users to excessively contort the system. If, however, change is too slow, then it will require a great many users for whom the data base is inaccurate before significant change can occur. Ideally, the weight to be given to a new piece of information would be a function of how much information has already been accumulated. At the outset, a new piece of information would be weighted heavily compared to the old value. As more and more information accumulates to support the old value, a new piece of information would be given less and less weight. Unfortunately, to implement this approach accurately requires storing, along with each prediction in the system, a measure of that amount of data that has accumulated in support of it. That is a heavy overhead. One possible compromise is to use the rating as an estimate of the amount of supporting data. It is not always a very good estimate, though. For example, there might be a lot of evidence supporting the assertion that half the members of a particular class possess a given characteristic. And there might be much less evidence evidence in support of another prediction that eighty percent of the members of that class possess a second characteristic. But the rating of the former will be lower than the rating of the latter. Fortunately, the very nature of the adaptation process is flexibility. And in addition, all the quantities being manipulated are approximations at best. So using the rating of a prediction to serve a dual purpose as a measure of confidence of a prediction as well as an estimate of the amount of information behind the prediction is not an unreasonable compromise.

An even simpler compromise, and the one actually used in Grundy, is to weight the old value by a constant factor. This avoids the problem of having the weight depend on the rating (which may, as was just noted, be unrelated to the amount of evidence that has accumulated). It does, of course, fail to weight according to the evidence at all. If the initial values were unrelated to the appropriate values, this would not be a reasonable approach, since in order to avoid wild fluctuations, the weight assigned to the old value must be large, and so it takes a long time, even at the beginning, for change to occur. But if the initial values are reasonable, this method does provide the slow change that is appropriate.

The actual mechanisms used by Grundy to modify its stereotypes and triggers can then be described as follows. Whenever an inference or a fact is recorded in the USS, stereotypes and triggers can be modified. At present, only numerical values are altered. This includes both the values for the linear numerical scales as well as the values associated with the symbolic attributes of the list facets (such as GENRES and INTERESTS).

Each fact or inference suggests a value for a particular facet in the USS. Every stereotype that has contributed to the value already contained in the USS for that facet may be modified, as well as all triggers that have activated those stereotypes. For each of these stereotypes one of three cases arises. In the first, the stereotype and the newly acquired information neither confirm nor conflict, so nothing is done. In the second, the stereotype predicts a value that confirms the newly acquired information. If that happens, the value contained in the stereotype will be modified by the value suggested by the new knowledge according to the following formula:

$$New-Stereotype-Value = \frac{Old-Stereotype-Value * Const + New-Info-Value}{Const+1}$$

The rating to be assigned to this prediction is recomputed as follows:

$$New-Stereotype-Value = Old-Stereotype-Rating + \frac{New-Info-Rating}{Old-Stereotype-Rating}$$

Note that the rating modification formula makes it easy to increase low ratings, but more difficult to increase already high ones. Then the rating of each of the triggers that activated the stereotype is modified according to the formula:

$$New-Trigger-Rating = Old-Trigger-Rating + \frac{New-Info-Rating}{Old-Trigger-Rating}$$

The other situation that may arise is that a stereotype's prediction may conflict with the new information. In that case, the new value will be computed using the same formula given above. But the ratings will be decreased by the modification value, rather than increased as in the case of a confirming piece of new knowledge.

Thus values will tend to move in the direction suggested by the bulk of the other evidence available, and ratings will increase with confirming evidence and decrease with conflicting indications.

6.3. Creating New Stereotypes

The next step in the learning process after the modification of existing stereotypes is the creating of new ones. This could be done after a system had enough models of individual users to be able to abstract patterns from them. The construction of new stereotypes can be done using straightforward pattern classification techniques in very much the same way that TIERESIAS (Davis, 1977) builds models that describe classes of rules it has learned. Such stereotypes would have the same advantages that such automatically constructed models provide to TIERESIAS, principally lack of vulnerability to the prejudices of the system designer. This step has not actually been taken in Grundy, since it has not yet had enough experience on which to draw.

7. THE PERFORMANCE OF GRUNDY

In order to collect evidence that Grundy is successfully exploiting user models, twenty-three people were observed using the system. By the end of this experiment, Grundy knew about 153 words, 90 stereotypes, and 115 triggers. Some of these were added during the course of the experiment as they were suggested by some of the dialogues. The use Grundy made of this information was by no means uniform. Some things were very heavily used (the MAN stereotype, for example), while others were rarely if ever used.

7.1. Grundy's Success at Recommending Novels

Although building the ideal novel recommending program was not the primary goal of this research, it is important to establish that Grundy exhibits some form of intelligent action in order to show that the user modeling techniques are effective. In order to measure the system's success, a small experiment was conducted at the end of each conversation Grundy had with a user. Some books were selected randomly from the data base, and several characteristics of those books were also chosen randomly. The books were then suggested to the user. It was then possible to compare Grundy's rate of success at recommending books that looked good to the users both when it was exploiting the user model -and when it was not. Table 1 shows the results of the comparison. It shows the number of good suggestions (i.e. the user said he might like) and the number of bad ones (i.e. the user said he would not like) in both the controlled and random modes. The data clearly demonstrate that the user models do contribute significantly to Grundy's performance as a novel recommender.

TABLE 1
Grundy's Success Rate

	CONTROLLED	RANDOM
GOOD	102	54
BAD	39	60

7.2. Learning in Grundy

As has already been discussed, it is important that the stereotypes and triggers that form the bulk of th~ data base of a user modeling system be able to change to reflect the actual body of users of the overall system. On the other hand, it is important that no single user be allowed to have a momentous effect on a system's global view of the class of all users. So in running Grundy with about twenty users, only the most commonly invoked stereotypes could be expected to change significantly. The most frequently activated stereotype was the MAN stereotype, and it did show interesting and significant change over the period Grundy was run. For example, the stereotype initially predicted that men would like books that are very fast-moving and full of suspense. The values of both of these characteristics decreased considerably on the basis of Grundy's experience with its users. Both of these changes can easily be explained by the difference between the group the stereotype was originally intended to characterize (all adult male Americans) and the group it was actually applied to in practice (some adult male intellectuals). The people who used the system, although they were men and did seem to like fast-moving, exciting books, also tended to like philosphical novels and literary classics, .which tend to be much calmer. Thus by the end of the experiment, Grundy had a more accurate picture of the men it was encountering than it had had at the outset. This suggests that the learning mechanisms are capable of causing at least some sorts of improvement in the stereotype data base of a user modeling system.

7.3. The Effectiveness of Stereotypes

In looking at the results of Grundy's attempts to use stereotypes to build useful models of people, one is reminded of one important characteristic of stereotypes and thus of systems that use them. Although a stereotype may often provide highly appropriate and useful information about a person, it is, at best, an expression of a tendency and not an absolute truth. One example will serve to illustrate this point clearly.

Several stereotypes in Grundy predict that people who fit them will be interested in the conflict "art-vs-science." Those predictions caused Grundy to believe, correctly, that several users were interested in that conflict and thus to recommend books involving it. But one person, who fit three stereotypes that predict an interest in that conflict said, on the second questionnaire, that he was definitely not interested in it. Figure 5 shows all the situations in which Grundy believed the user was interested in art-vs-science.

This example just serves as a reminder of the great usefulness of stereotypes, but, at the same time, their fallibility.

User	*Stereotypes that Predicted An Interest*	*Is the User Interested?*
4	Intellectual, Contemplative, Scientist	yes
6	Intellectual	yes
7	Intellectual, Contemplative	yes
8	Intellectual	yes
9	Contemplative	yes
10	Intellectual, Artist	yes
11	Intellectual	yes
13	Intellectual	yes
20	Intellectual	yes
22	Scientist, Artist, Contemplative	NO

Figure 5. An Illustration of the Fallibility of Stereotypes.

7.4. Conclusion

Grundy is not a perfect novel recommender nor is it a perfect user model builder. But it does a respectable job at both tasks. Thus the techniques it demonstrates can be seen to be useful and to provide at least a start along the path of discovering the right way to build and exploit user models.

8. CONCLUSION

The goals of this research were to discover the usefulness of individual user models to performance systems and to explore the appropriateness of some specific techniques (principally the use of stereotypes) to the task of building such models. Although it is always difficult to generalize from a single system to the class of all systems, the experience with Grundy suggests both that user models can be of significant benefit and that the use of stereotypes to build those models can be highly productive. A more detailed study of all of these issues can be found in Rich (1979).

FOOTNOTES

[1]This is not to suggest that readers are necessarily male and that librarians are female. But in order to make this description readable, it is necessary to use different pronouns for the two people. No political statement is intended.

[2]In order to eliminate inferences from the system, it is necessary to have a threshold below which inferences are too trivial to bother with and then to compare ratings to this threshold whenever they are changed. Although this is straightforward, it has not been implemented in Grundy.

REFERENCES

Bobrow, D. G. & Norman, D. A. Some principles of memory schemata. In D. G. Bobrow & A.Collins (OOs.), *Representation and Understanding.* New York: Academic Press, 1975.

Cantor, N. & Mischel, W. Traits as prototypes: Effects on recognition memory. *Journal of Personality and Social Psychology,* 1977,35, 3S-48.

Carbonell, J. Jr. *Subjective understanding: Computer models of belief systems.* PhD Thesis, Yale Univ. Computer Science Dept., 1979.

Charniak, E. A framed painting: The representation of a common sense knowledge fragment. *Cognitive Science,* 1977, I, 355-394.

Davis, R. Interactive transfer of expertise: Acquisition of new inference rules. Proceedings of the 5th International Joint Conference on Artifical Intelligence, 1977, 321-328. Hamilton, D. L. A cognitive-attributional analysis of stereotyping. In L. Berkowitz (ed.), *Advances in ExperimentaL PsychoLogy, 1979.*

Highet, G. *The Art of Teaching.* Vintage, 1950.

McGuire, W. J. & Padawer-Singer, A. Trait salience in the spontaneous self-concept. *JournaL of Personality and*

Social Psychology, 1976,33,743-754.
Minsky, M. A framework for representing knowledge. In P. H. Winston (00.), *Psychology of Computer Vision.* New York: McGraw-Hili, 1975.
Nisbett, R. E. & Wilson, T. D. Telling more than we know: Verbal reports on mental processes. *PsychoLogical Review.* 1977,84,231-259.
Rich, E. A. *Building and exploiting user models.* PhD Thesis, Carnegie-Mellon Univ. Computer Science Dept., 1979.
Schank, R. C. & Abelson, R. P. *Goals, Plans, Scripts and Understanding: An Enquiry into Human Knowledge Structures.* Lawrence Erlbaum Associates, 1977.
Shortliffe, E. H. MYCIN: *Computer-Based Consultations in Medical Consultations.* American Elsevier, 1976.
Snyder, M. & Uranowitz, S. W. Reconstructing the past: Some cognitive consequences of person perception. *Journal of Personality and Social Psychology,* 1978,36.

*This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory Under Contract F33615-78-C-1551, while the author was at Carnegie-Mellon University. The author was supported by a grant from the Xerox Corporation.