THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Generation and evaluation of user tailored responses in multimodal dialogue

Citation for published version:

Walker, MA, Whittaker, SJ, Stent, A, Maloor, P, Moore, J, Johnston, M & Vasireddy, G 2004, 'Generation and evaluation of user tailored responses in multimodal dialogue', *Cognitive Science: A Multidisciplinary Journal*, vol. 28, no. 5, pp. 811-840. https://doi.org/10.1207/s15516709cog2805_8

Digital Object Identifier (DOI):

[10.1207/s15516709cog2805_8](https://doi.org/10.1207/s15516709cog2805_8)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Cognitive Science: A Multidisciplinary Journal

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Generation and Evaluation of User Tailored Responses in MultiModal Dialogue [★]

M.A.Walker ^a S.J.Whittaker ^a A. Stent ^b P. Maloor ^c J. Moore ^d
M. Johnston ^c G. Vasireddy ^c

^a*University of Sheffield
Sheffield, England, S1 4DP*

^b*SUNY at Stony Brook
Stony Brook, NY, USA, 11794*

^c*AT&T Labs - Research
Florham Park, NJ, USA, 07932*

^d*University of Edinburgh
Edinburgh, Scotland, EH8 9LW*

Abstract

When people engage in conversation, they tailor their utterances to their conversational partners, whether these partners are other humans or computational systems. This tailoring, or adaptation to the partner takes place in all facets of human language use, and is based on a *mental model* or a *user model* of the conversational partner. Such adaptation has been shown to improve listeners' comprehension, their satisfaction with an interactive system, the efficiency with which they execute conversational tasks, and the likelihood of achieving higher level goals such as changing the listener's beliefs and attitudes. We focus on one aspect of adaptation, namely the tailoring of the content of dialogue system utterances for the higher level processes of persuasion, argumentation and advice-giving. Our hypothesis is that algorithms that adapt content for these processes, according to a user model, will improve the usability, efficiency, and effectiveness of dialogue systems. We describe a multimodal dialogue system and algorithms for adaptive content selection based on multi-attribute decision theory. We demonstrate experimentally the improved efficacy of system responses through the use of user models to both tailor the content of system utterances and to manipulate their conciseness.

Key words: dialogue systems, user modeling, user-tailored generation

[★] This work was partially supported by DARPA grant MDA 972 99 3 0003 to AT&T Labs Research.

Email address: walker@dcs.shef.ac.uk (M.A.Walker).

1 Introduction

When people engage in conversation, they tailor their utterances to their conversational partners, whether these partners are other humans or computational systems [Brennan, 1991, Schober, 1998]. This tailoring, or adaptation to the partner, has been shown to take place in all facets of human language use, including speaking rate and response delay [Ward and Nakagawa, 2002, Darves and Oviatt, 2002], amplitude and prosodic range [McLemore, 1992, Coulston et al., 2002], lexical and syntactic choice [Levelt and Kelter, 1982, Kempen and Hoenkamp, 1987, Brennan, 1996], choice and modality of referring expressions [Joshi and Weinstein, 1981, Nadathur and Joshi, 1983, Garrod and Anderson, 1987, Brennan and Clark, 1996, Schober, 1998, Bell et al., 2000] and in higher level discourse processes such as the selection of content and form for persuasive arguments and negotiation [Joshi et al., 1984, Joshi, 1982, Webber and Joshi, 1982, Mayberry and Golden, 1996, McGuire, 1968, Walker, 1996]. This adaptive behavior is based on a *mental model* or a *user model* of the conversational partner [Brennan and Clark, 1996, Levelt, 1989, Zukerman and Litman, 2001, Wahlster and Kobsa, 1989]. Such adaptation has been shown to improve listeners' comprehension [Clark and Wilkes-Gibbs, 1986], their satisfaction with an interactive system [Nass et al., 1995], the efficiency with which they execute conversational tasks [Brennan, 1996, Clark and Wilkes-Gibbs, 1986], and the likelihood of achieving higher level goals such as changing the listener's beliefs and attitudes [Luchok and McCroskey, 1978, Carenini and Moore, 2001, 2000b, Zukerman and McConachy, 1993].

Our focus here is on one aspect of adaptation, namely the tailoring of the content of dialogue system utterances for the high level processes of persuasion, argumentation and advice-giving. Dialogue systems are one of the few examples of an intelligent artifact that can interact with humans to carry out a variety of tasks. Various hypotheses about conversational interaction can be tested in dialogue systems by implementing algorithms that control the system's conversational behavior. As such, dialogue systems provide an important experimental vehicle for cognitive science and theories of interaction. Our research also has the practical goal of improving the dialogue interaction capabilities of the MATCH (Multimodal Access to City Help) multimodal dialogue system, a system that provides information on restaurant and entertainment options in New York City [Johnston et al., 2002].

Our hypothesis is that algorithms that adapt content for higher level discourse processes, according to a user model, will improve the usability, efficiency, and effectiveness of dialogue systems. Dialogue systems have a particularly strong requirement to produce concise, informative and relevant utterances, especially during the information presentation phase of the dialogue [Walker et al., 2002a]. In this phase, the system has a number of possible options

that match a user's constraints, which need to be presented to the user. It is important for the system to present the options in a form that will help the user understand and evaluate the tradeoffs among them. Dialogue strategies for recommending particular options, or for making balanced comparisons between options, should help users make such evaluations. To be effective in spoken dialogue, these recommendations and comparisons should also be concise.

Previous work on user modeling has primarily applied models of user expertise or knowledge to the generation of user tailored texts, rather than to system utterances in a dialogue system. The first such system, developed by Rich 1979, tailored book recommendations to a user's preferences as expressed in a user model. The system first asked the user a series of (yes/no) questions in order to categorize the user into one of its known stereotypes, and adjusted this model as the (typewritten) interaction progressed. There has been considerable subsequent research on developing interactive systems that utilize models of users' capabilities, preferences or biases for recommendation, advice or explanation [Webber and Joshi, 1982, Joshi et al., 1984, Joshi, 1982, Joshi et al., 1986, Paris, 1988, Chin, 1989, Morik, 1989, Moore and Paris, 1993, Klein, 1994, Walker, 1996, Jameson et al., 1995, Thompson and Goker, 1999, Linden et al., 1997, Cawsey, 1993, Carenini and Moore, 2000b, 2001] *inter alia*, and on methods for automatically inferring such models from user actions [Goecks and Shavlik, 2000, Rafter et al., 2000, Linden et al., 1997, Rogers and Fiechter, 1999]. Generation of text recommendations based on user preferences is now being commercially deployed by CoGenTex in their Recommender system, which automatically generates natural language descriptions and comparisons of product features, using information obtained from a ranking and comparison engine [Cogentex, 2003].

Recommendations for particular options, and comparisons among options, are one form of *evaluative argument*. An evaluative argument typically consists of a main *claim*, and *evidence* relevant to the claim. Argumentation theory provides a number of guidelines for producing effective evaluative arguments [Mayberry and Golden, 1996, Miller and Levine, 1996, Corbett and Connors, 1999, McGuire, 1968, Zukerman et al., 2000], which are summarized by Carenini and Moore [Carenini and Moore, 2000a]. These guidelines require:

- (1) *Identifying supporting and opposing evidence*: Evidence must be based on a model of the user's values and preferences, e.g. superb restaurant decor can only be used to support an argument for going to a restaurant if the user is oriented to decor.
- (2) *Positioning the main claim*: placing the main claim first helps users follow the line of reasoning, but delaying the claim until the end of the argument can also be effective if the user is likely to disagree with the claim.
- (3) *Selecting supporting and opposing evidence*: An argument cannot include

all the possible evidence, so only strong evidence should be presented in detail, and weak evidence only briefly mentioned or omitted entirely.

- (4) *Arrangement of supporting evidence*: The strongest support should be presented first but, if possible, one effective piece of supporting evidence should be saved for the end to leave the user with a final impression of the strength of the argument.
- (5) *Addressing and ordering opposing evidence*: The choices are not to mention any opposing evidence, to acknowledge it without refuting it, or to acknowledge it and refute it. The opposing evidence should be presented so as to minimize its effectiveness with strong opposing evidence in the middle and weak evidence at the beginning and end.
- (6) *Ordering between supporting and opposing evidence*: If the reader is aware of the opposing evidence, then it should come before the supporting evidence, otherwise after.

These guidelines must first be formalized to be used in a computational system. The formalization requires representing the user's values and preferences (guideline 1), providing a way to measure the strength of supporting or opposing evidence (guidelines 3,4,5), representing whether the user is aware of certain facts (guideline 6), and developing strategies for ordering and structuring the selected content into coherent and persuasive arguments (guidelines 2,4,5,6).

Carenini and Moore formalized and evaluated these guidelines in the context of a system for interactive data exploration in the real estate domain [Carenini and Moore, 2000b,a, 2001, Carenini, 2000]. Their operationalization of user models is based on multi-attribute decision theory [Keeney and Raiffa, 1976, Klein, 1994]. Multiattribute decision theory provides both a way to represent the user's values and preferences and to measure the strength of supporting and opposing evidence (as we explain in more detail below). The strength of evidence measure is then the basis for strategies for selecting and structuring the content of recommendations. The user model is also used to make these recommendations concise, in a similar approach to that described here. Carenini and Moore showed experimentally that tailored recommendations were preferred over non-tailored recommendations, and that concise recommendations based on the user models were preferred over verbose recommendations.

Our research extends that carried out by Carenini and Moore in four ways. First, we test whether user models based on multi-attribute decision theory generalize across domains, by applying this approach to the problem of restaurant selection in New York City. Second, we extend this approach to multi-modal dialogue, where the requirements for interactive information presentation are different from those for text presentations. The system developed by Carenini and Moore is interactive but does not carry on a natural language dialogue with the user; instead it presents a single text recommenda-

tion in a multi-modal context. Third, we extend user-tailored generation to include comparisons as well as recommendations. We evaluate the effects of user-tailoring on these strategies. Finally, we explore the relationship between tailoring and mode of information presentation by exploring the effect of presenting these strategies using text or speech.

Section 2 describes the MATCH system and how we use it to test various cognitive hypotheses about user tailored interaction. Section 3 describes the use of multi-attribute decision theory for user modeling and provides detailed examples of user models from our user group. Section 4 describes the content selection algorithms based on the user models, and how they are utilized in dialogue strategies based on argumentation theory. Section 5 describes the design, hypotheses, and results of two evaluation experiments, which demonstrate the benefits of tailoring and the benefits of the user models in manipulating the conciseness of utterances. We sum up in Section 6.

2 The MATCH dialogue system and specific hypotheses

Fig. 1. User interacting with MATCH running on Fujitsu PDA

The MATCH system runs on a small, portable, tablet computer, providing a testbed for research on multimodal dialogue interaction in a mobile setting. Figure 2 shows the size of MATCH relative to the human hand and illustrates a user gesture. Users interact with MATCH using a multimodal user interface client. The tablet screen is divided into an area showing a street map of New York City (See Figures 3 and 4), a table showing options that match the user's current request (if any), buttons to activate the speech recognizer, and

a panel that provides feedback on recognition and the system state. The street map is dynamically rendered and can be panned and zoomed. Users may take a turn in the dialog with speech, gestures made with the pen, handwriting, or a multimodal combination of inputs [Johnston et al., 2002]. Inputs are parsed, integrated, and assigned a combined meaning representation using a multimodal language processing architecture based on finite-state techniques [Johnston and Bangalore, 2000, Bangalore and Johnston, 2000, Johnston and Bangalore, 2004]. In addition to finite-state multimodal integration, MATCH also differs from previous multimodal systems for interacting with maps, such as QuickSet [Johnston et al., 1997, Johnston, 1998, Cohen et al., 1998, Oviatt, 1999, Wu et al., 1999], in that it supports handwritten input in addition to pen gestures, provides a dynamic map display, and incorporates a multimodal dialog manager. The multimodal dialog manager, implemented in Java, is in the style of TrindiKit [Larsson et al., 1999] and features extensions to support the visual mode in addition to spoken interaction [Johnston et al., 2002, Hastie et al., 2002]. The system’s responses can be speech, changes in the map display or text window, or coordinated multimodal presentations of these different output modes. AT&T’s Watson and Natural Voices engines are used for speech recognition and text-to-speech (TTS) [Sharp et al., 1997, Beutnagel et al., 1999].

Our goal is to improve MATCH’s dialogue interaction capabilities and general utility by making it easier for users to understand the tradeoffs between different restaurant options. Our view is that this can be done in three ways by providing the system with capabilities for: (1) responding to requests for recommending one of a set of restaurants, or for comparing small sets of restaurants; (2) tailoring these recommendations and comparisons to a model of the user’s individual preferences; and (3) making the responses sufficiently concise for the user to understand and remember important information.

The role of the user model in system responses is to affect both the ranking of options returned from the database and the selection of which attributes to mention in a recommendation or comparison. Figure 2 shows a sample dialogue with MATCH exploiting the user tailored dialogue strategies described in the rest of the paper, with examples tailored for the user OR (See Figure 7 for OR’s user model). In Figure 2, in utterance U1, the user specifies the query *Show Italian Restaurants in the West Village* in speech. The system responds in S1 by presenting a map of New York, zooming to the West Village and highlighting Italian restaurants (Figure 3). At this point, the user has too many options to decide between so he decides to select a set with a pen gesture (Figure 4) and compare them (U2). S2 is that comparison. Since all the restaurants mentioned in S2 are acceptable, the user asks the system to recommend one by writing the word “recommend” (U3). The recommendation operates on the current dialogue context which is the selected set (from U2).

U1	<i>“Show me Italian restaurants in the West Village”</i>
S1	[zooms to the West Village, displays Italian restaurants. (Figure 3)]
U2	[circles a selection of restaurants and says <i>“compare”</i> (Figure 4)]
S2	<i>“Among the selected restaurants, the following offer exceptional overall value. Ugualè’s price is \$33. It has excellent food quality and good decor. Da Andrea’s price is \$28. It has very good food quality and good decor. John’s Pizzeria’s price is \$20. It has very good food quality and mediocre decor.”</i>
U3	[writes word <i>“recommend”</i>]
S3	<i>“Ugualè has the best overall value among the selected restaurants. Ugualè’s price is \$33.”</i>

Fig. 2. Example dialogue with MATCH, U = user, S = System

Fig. 3. MATCH’s graphical system response to *Show me Italian Restaurants in the West Village*.

The system utterances in Figure 2 are all generated at a fixed level of conciseness. Sample recommendations for a task of finding a Japanese restaurant in the East Village for two different users, with varying levels of conciseness as generated by our algorithms are shown in Figure 5. The user model and the conciseness parameter z lead to selection of different restaurants to recommend and to mentioning different facts to each user.

Note also that the user model reflects a users’ *dispositional* biases about restaurant selection, but these can be overridden by *situational* constraints specified

Fig. 4. User circles subset of Italian West Village restaurants for comparison.

User	Conciseness	Output
CK	Concise ($z= 0.3$)	Bond Street has the best overall value among the selected restaurants. Bond Street has excellent food quality.
BA	Concise ($z= 0.3$)	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29. It's a Japanese, Latin American restaurant.
CK	Sufficient ($z= -0.7$)	Bond Street has the best overall value among the selected restaurants. Bond Street's price is \$51 and it has excellent food quality and good service. It's a Japanese, Sushi restaurant.
BA	Sufficient ($z= -0.7$)	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29 and it has very good service and very good food quality. It's a Japanese, Latin American restaurant.
CK	Verbose ($z= -1.5$)	Bond Street has the best overall value among the selected restaurants. Bond Street's price is \$51 and it has excellent food quality, good service and very good decor. It's a Japanese, Sushi restaurant.
BA	Verbose ($z= -1.5$)	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29 and it has very good service, very good food quality and good decor. It's a Japanese, Latin American restaurant.

Fig. 5. Recommendations for Users CK and BA, for the East Village Japanese Task, of Varying Levels of Conciseness.

in a user query. For example, the user models (as described below) allow us to represent the fact that some users have strong preferences for particular food types. However, in a particular dialogue situation, these can be overridden by interactively requesting a different food type, e.g. Italian food as in Figure 2. Thus *dispositional* biases never eliminate options from the set of options returned by the database, they simply affect the *ranking* of options, and the weighting of their attributes.

The primary hypothesis that we wish to test through user interactions with the MATCH system is that *user tailored responses are more effective*. In the evaluation experiments described below, we compare the users' evaluation of dialogue responses tailored to their own model, with responses tailored to a randomly selected model of another user.

Our second hypothesis concerns conciseness. We utilize the strength of evidence defined by the user model to vary the *conciseness* of system responses. Concise utterances are defined as those mentioning just those restaurants and their attributes that are most relevant to the users preferences. We compare user's evaluation of concise, sufficient and verbose dialogue responses.

A third hypothesis concerns potential interactions between *user-tailoring and the mode in which information is presented* in a multimodal dialogue system, i.e. in speech or in text. Consistent with prior research, we expect the ephemeral nature of speech (and the resulting cognitive load) to make this a less effective output mode than text [Oviatt, 1997, McKeown et al., 1998, Mittal et al., 1995, Whittaker et al., 1991]. However we also expect that tailoring might address some of the inherent limitations of speech, having a greater effect on spoken than text presentations.

3 Multi-Attribute Decision Models in the Restaurant Domain

User models derived from multi-attribute decision theory have been shown to be effective for guiding user interaction in various types of interactive systems [Klein, 1994, Jameson et al., 1995, Linden et al., 1997, Thompson and Goker, 1999]. They have also been found to be good predictors of user's consumer behavior [Solomon, 1998]. For our current purposes, they have two other important properties, namely (a) they are quantitative, which makes them easy to operationalize (b) it is relatively easy to gather the data necessary for constructing such user models of this type.

Multi-attribute decision models are based on the claim that if anything is valued, it is valued for multiple reasons [Keeney and Raiffa, 1976]. In the restaurant domain, this implies that a user's preferred restaurants optimize tradeoffs among restaurant attributes. To define a model for the restaurant domain, we must determine these attributes and their relative importance for particular users. We use a standard procedure called SMARTER that has been shown to be a reliable and efficient way of eliciting multi-attribute decision models for particular users or user groups [Edwards and Barron, 1994].

3.1 Structure of the Model

The first step of the standard SMARTER procedure is to determine the structure of a tree model of the *objectives* in the domain. In MATCH, the top-level objective is to select a good restaurant. User interviews and data collection along with an analysis of online restaurant databases indicated that six attributes contribute to this objective: the quantitative attributes *food quality*, *cost*, *decor*, and *service*; and the categorical attributes *food type* and *neighborhood* [Whittaker et al., 2002]. These attributes are structured into the one-level tree shown in Figure 6. A more complex structure that grouped decor, neighborhood and service under a higher level objective called *ambiance* was considered, but informal questioning of users suggested this structure was less intuitive [Whittaker et al., 2002].

Fig. 6. Structure of Objectives for MATCH

The structure is user-independent with user-dependent weights on the branches as explained below. We apply this structure to a database of approximately 1000 restaurants populated with information freely available from the web. Values for each of these attributes for each restaurant are stored in the database.

3.2 Normalizing attribute values

The second step is to transform the real-domain values of attributes x into single-dimension cardinal utilities $u(x)$ such that the highest attribute value is mapped to 100, the lowest attribute value to 0, and the others to values in the interval 0 to 100. This is necessary to normalize the values of the different attributes. In the restaurant database that we accessed from the Web *food quality*, *service* and *decor* range from 0 and 30, with higher values more desirable, so 0 is mapped to 0 and 30 to 100 in our model. The *cost* attribute ranges from \$10 and \$90 and higher values are less desirable, so \$90 is mapped to 0 on the utility scale. Preferred values for categorical attributes such as *food*

type are mapped to 90, dispreferred values to 10 and others to 50.¹ Table 1 shows the attributes in the restaurant domain, with the functions mapping the values of each attribute in the web database into cardinal utilities.

Attribute	Range of values	Mapping of values to cardinal utilities
Food quality, Service, Decor	0–30	value x 3 1/3
Cost	0–90	100 - (10/9 x value)
Food type, Neighborhood	e.g. Italian, French, West Village	top values listed by user are mapped to 90, bottom ones to 10 and all others to 50

Table 1

Mapping of attribute values to utilities in the restaurant domain

The vector of $u(x)$ values are aggregated into a scalar in order to determine the overall utility U_h of each option h . The most widely used model for such aggregations is the additive model (over 95% of models used in practice are additive), and standard heuristic tests with users suggested that an additive model is a good approximation [Edwards and Barron, 1994]. Use of an additive model means that each attribute is assumed to be independent of every other one. The individual attribute utilities are combined into an overall utility using a simple additive function; the value for each attribute is multiplied by its weight and all the weighted values are summed. Thus, if h ($h = 1, 2, \dots H$) is an index identifying the restaurant options being evaluated, k ($k = 1, 2, \dots K$) is an index of the attributes, u is the function for each attribute mapping attribute values to utilities, and w_k is the weight assigned to each attribute:

$$U_h = \sum_{k=1}^K w_k u_k(x_{hk})$$

3.3 Allocating Weights to Attributes

The final step of decision model construction is the assignment of specific weights w_k to each attribute k . Attribute weights are user-specific, reflecting individual preferences about tradeoffs between options in the domain, and are based on users' subjective judgments elicited using the SMARTER elicitation procedure. SMARTER's main advantage over other elicitation procedures is that it only requires the user to specify the *ranking* of domain attributes. There is considerable experimental evidence showing that simple attribute ranking

¹ Users were allowed to select up to five preferred and five dispreferred food types. This simplification is motivated by the large number of food types available in New York City and our requirement to keep the enrollment process short and simple.

is both efficient, and nearly as accurate as more time-consuming methods, in which users allocate weights directly [Edwards and Barron, 1994, Srivastava and Connolly, 1995].

We elicit a user model when new users enroll with MATCH using the standard form of questions specified by SMARTER. The elicitation procedure is implemented as a sequence of web pages. The first web page says *Imagine that for whatever reason you’ve had the horrible luck to have to eat at the worst possible restaurant in the city. The price is \$100 per head, you don’t like the type of food they have, you don’t like the neighborhood, the food itself is terrible, the decor is ghastly, and it has terrible service. Now imagine that a good fairy comes along who will grant you one wish, and you can use that wish to improve this restaurant to the best there is, but along only one of the following dimensions. What dimension would you choose? Food Quality, Service, Decor, Cost, Neighborhood, or Food Type?* After the user chooses an attribute on this page, the scenario is repeated omitting the chosen attribute, until all attributes have been selected. Users are then asked to specify whether they have any *neighborhood* or *food type* likes or dislikes.

Given the ranking, the weights are calculated using the following equation, which guarantees that the total sum of the weights add to 1, a requirement for multi-attribute decision models:

$$w_k = \frac{1}{K} \sum_{i=k}^K \frac{1}{i}$$

3.4 Resulting User Models

To date, 29 different user models have been elicited and stored in a database that MATCH uses. Figure 7 shows attribute weightings and likes and dislikes for five of these users. What is most striking about the table are the large differences between users. When differences in categorical preferences are taken into account, no two users in our sample are alike, but even if we only consider the relative importance of various attributes, we find that only two pairs of users are identical in the ranking of attributes. For 25 of these users, we found that *cost* and *food quality* are always in the top three attributes, but user BA ranked *food type* highest, followed by *cost* and *service*. Even for users who ranked both *cost* and *food quality* in their top three attributes, the relative importance of lower ranked attributes, such as *decor*, *service*, *neighborhood* and *food type*, varies widely. For example, every user ranks *service* differently as reflected by the different weights in the Service column. User CK ranks *decor* as the least important attribute, while user OR ranks it third in importance, and users CK and SD rank *food type* as the second most important attribute while users OR and MSh rank *food type* last.

User	FQ	SVC	Dec	Cost	Nbhd	FT	Nbhd Likes	Nbhd Dislikes	FT Likes	FT Dislikes
BA	0.10	0.16	0.06	0.24	0.03	0.41	Downtown, Midtown, E. Village, TriBeCa, SoHo	The Bronx, Harlem	Cajun Creole, Greek, Italian, Japanese, Seafood	Coffeehouses, Desserts, German, Steak
CK	0.41	0.10	0.03	0.16	0.06	0.24	Midtown, Chinatown, TriBeCa	Harlem, Bronx	Indian, Mexican, Chinese, Japanese, Seafood	Vegetarian, Vietnamese, Korean, Hungarian, German
OR	0.24	0.06	0.16	0.41	0.10	0.03	W. Village, Chelsea, Chinatown, TriBeCa, E. Village	Upper E. Side, Upper W. Side, Uptown, Bronx, Lower Manhattan	French, Japanese, Portuguese, Thai, Middle Eastern	no-dislike
MSh	0.41	0.10	0.06	0.24	0.16	0.03	Flatiron, Chelsea, W. Village, Midtown East, Midtown West	Chinatown, Lower E. Side, E. Village, Upper E. Side, Upper W. Side	Indian, Mexican, Ethiopian, Thai, French	Steakhouse, Russian, Korean, Filipino, Diner
VM	0.24	0.10	0.03	0.41	0.06	0.16	Upper W. Side		Cajun Creole, Chinese, Coffeehouses, Indian, Tapas	

Fig. 7. Example User Models: FQ = Food Quality, SVC = Service, DEC = Decor, Nbhd = Neighborhood, FT = Food Type

After examining these differences qualitatively, we decided it would be useful to be able to quantify the differences among user models. We utilize a common measure of *distance*, the Manhattan or city-block distance [Mitchell, 1998], which is simply the sum of the absolute values of the differences in the weights for each attribute in the user models. That is, for users i, j and attributes k indexed from $1 \dots K$, with weights w_k ,

$$distance_{ij} = \sum_{k=1}^K (|w_{ki} - w_{kj}|)$$

For example, the distance between users CK and VM in Figure 7 is 0.84, and the distance between users CK and BA in Figure 7 is .89. The average distance between user models in our current user group is .57. The distance metric enables us to manipulate differences between models and to quantify the effect of those manipulations.

4 The SPUR Dialogue Planner

So far, we have described the nature of user models derived from multi-attribute decision theory. We now explain how these are used to generate user tailored outputs in an interactive dialogue system.

The content planning module in MATCH is called SPUR (Speech Planning with Utilities for Restaurants). The user model is used by SPUR for two aspects of content selection: (1) it ranks the options returned from a database query, and the ranking is used by SPUR to select a subset of restaurant options to recommend or compare; (2) it determines a subset of attributes that are mentioned for each option, with the size of the subset depending on the setting of a conciseness parameter.

SPUR takes as input: (1) a dialogue strategy goal; (2) a user model; (3) a conciseness parameter z , and (4) a set of restaurant options returned by the database that match situational constraints specified in the user’s query. Given the options, and the conciseness setting, SPUR constructs a content plan specific to each strategy and user model. The resulting content plan serves to filter the information presented to the user so that only options and attributes that are most relevant to the user are mentioned, contrasted and highlighted. This should make it easier for the user to evaluate the trade-offs among options in a set, reducing dialogue duration and increasing user satisfaction.

We first illustrate how the user model reranks the option set to which we then apply our system dialogue strategies, and describe how the user model affects the recommend and compare content plans.

4.1 The Effect of User Model on Option Ranking

To show the effects of user model on option ranking, we present the restaurant options that match the query *Show Japanese restaurants in the East Village*. Figure 8 shows how the user models for CK, BA and VM rank these options. The third column gives the overall utility, U_h . The subsequent columns give the attribute values and in parentheses the weighted utilities (WTD). Note that *food quality* contributes most strongly to the weighted utilities in the CK model ranking, while *cost* contributes most strongly to the ranking for both BA and VM. However, BA and VM differ in that VM’s second most important attribute is *food quality*, while for BA the second most important attribute is *food type*. This modifies the ranking of the restaurant set.

Let us consider in detail the differences in overall ranking for CK and VM resulting from different attribute weightings. Bond Street (a highly priced

User	Restaurant	U_h	FQ(wtd)	SVC(wtd)	DEC(wtd)	Cost(wtd)	Nbhd(wtd)	FT(wtd)
BA	Komodo	77	22(7)	22(10)	19(4)	29(18)	90(2)	90(36)
BA	Japonica	71	23(7)	18(7)	15(3)	37(16)	90(2)	90(36)
BA	Takahachi	71	21(6)	17(6)	14(2)	27(19)	90(2)	90(36)
BA	Shabu-Tatsu	70	20(5)	18(7)	15(3)	31(17)	90(2)	90(36)
BA	Bond Street	69	25(8)	19(8)	22(4)	51(11)	90(2)	90(36)
BA	Dojo	66	15(2)	12(2)	8(1)	14(23)	90(2)	90(36)
CK	Bond Street	63	25(34)	19(3)	22(2)	51(5)	50(7)	50(12)
CK	Japonica	59	23(29)	18(3)	15(1)	37(7)	50(7)	50(12)
CK	Komodo	59	22(26)	22(4)	19(2)	29(8)	50(7)	50(12)
CK	Takahachi	54	21(24)	17(2)	14(1)	27(8)	50(7)	50(12)
CK	Shabu-Tatsu	52	20(22)	18(3)	15(1)	31(7)	50(7)	50(12)
CK	Dojo	30	15(10)	12(1)	8(0)	14(10)	50(7)	50(12)
VM	Komodo	66	22(16)	22(7)	19(2)	29(31)	50(3)	50(7)
VM	Takahachi	61	21(14)	17(4)	14(1)	27(32)	50(3)	50(7)
VM	Japonica	58	23(17)	18(4)	15(1)	37 (26)	50(3)	50(7)
VM	Shabu-Tatsu	57	20(13)	18(4)	15(1)	31(29)	50(3)	50(7)
VM	Bond Street	56	25(20)	19(5)	22(2)	51(19)	50(3)	50(7)
VM	Dojo	56	15(6)	12(1)	8(0)	14 (39)	50(3)	50(7)

Fig. 8. Results of DB Query for East Village Japanese for users BA, CK and VM: U_h = Overall Utility. WTD = Weighted utility for each attribute. FQ = Food Quality, SVC = Service, DEC = Decor, Nbhd = Neighborhood, FT = Food Type

restaurant with excellent food quality) is fifth for VM because VM ranks *cost* first and *food quality* second. Bond Street’s 25 rating for *food quality* results in 34 utils (utils are units of weighted utility) for CK, but only 20 utils for VM. Also,

Bond Street’s price of \$51 per person results in only 19 utils for VM; all of the restaurants ranked higher by VM than Bond Street are less expensive. On the other hand, Komodo is more highly ranked for VM than CK. This is mainly because its modest price gets 31 utils for VM but only 8 for CK. Note also that Dojo, which is very inexpensive, is as good as Bond Street in overall utility for VM (both get 56) but for CK, Dojo’s lower food quality means that it has a much worse overall utility.

4.2 SPUR Dialogue Strategies

We defined two types of strategy for SPUR: (1) RECOMMEND one of a selected set of restaurants; (2) COMPARE three or more selected restaurants. For each response, SPUR outputs a content plan to the template-based surface realizer (described below), using the overall utility U_h to rank the options as described in Section 4.1. For recommendations, the algorithm selects the top-ranked

option. For comparisons, the algorithm selects a top-ranked subset of options to compare. Then the weighted attribute values are used to select the content for each option.

Conciseness is controlled with a parameter that determines whether an option or attribute is an outlier with respect to other options or attributes. Outliers are deemed worth mentioning because they deviate from the norm [Klein, 1994]. According to multi-attribute decision theory, the weighted attribute model also enables us in principle to determine the likelihood that mentioning a given attribute will change the user’s belief state. For example, compare the recommendations in Figure 5. The most concise recommendation for both CK and BA mentions one attribute. The weighted attribute values for each user in Figure 8 predict how convincing a recommendation would be that includes just that attribute. Figure 8 indicates that telling CK about Bond Street’s *food quality* should provide 34 utils (units of utility) out of a possible 63. Similarly, telling BA about Komodo’s *food type* is predicted to provide 36 utils out of a possible 77. Including more attributes makes the recommendation more convincing, e.g. adding the *food type* attribute as in CK’s Sufficient recommendation in Figure 5 should provide 46 (34 + 12) utils out of a possible 63 total utils.

In sum, we map conciseness directly onto the weighted attribute ranking of the user model: more concise descriptions select the subset of attributes that maximally affect the user’s belief state. More verbose descriptions also include lower weighted attributes. Obviously, however, there is a trade-off between maximizing expected utility, and verbosity. Mentioning more attributes increases expected utility while requiring the user to remember more information.

Below, we first describe how outliers are identified (Section 4.2.1), and then describe the algorithms for constructing each type of content plan. Section 4.3 describes the templates used to realize the content plans.

4.2.1 Defining Outliers

We define a response as *tailored* if it is based on a user’s known biases and preferences. A response is *concise* if it includes only those options with high utility, or possessing *outliers* with respect to a population of options or attribute values. We use the *z-score* (standard value) of an option’s overall utility, or of the weighted attribute value v , to define an *outlier*:

$$z(v) = \frac{v - \mu_V}{\sigma_V}$$

The *z-score* expresses how many standard deviations σ_V a value v is away from the mean μ_V of a population of values V . The population of values V

that are used to calculate μ_V and σ_V can be (a) other attributes for the same option (for RECOMMEND), or (b) the same attribute for other options (for COMPARE). Depending on a threshold for outliers, z , different numbers of options or attribute values are considered to be worth mentioning, because they stand out from other values. For example, when the threshold for z is 1.0, the weighted attribute values must be more than one standard deviation away from the mean for that attribute to be selected for expression. This threshold can obviously be modified to generate responses at different levels of conciseness. In the examples below, we illustrate responses for different settings of z , for the user models for VM, CK and BA in Figure 7.

4.2.2 Recommendation dialogue strategies

The system’s strategy for making a recommendation is to select the best option (based on overall utility) and provide convincing reasons for the user to choose it (based on weighted attribute values). Figure 9 provides the algorithm for selecting the content for the RECOMMEND dialogue strategy.

- (1) Select the restaurant option R with highest overall utility from returned options.
- (2) Using the setting for z , identify the attributes a_i whose weighted attribute values v_i for that option are outliers.
- (3) Construct a content plan with the claim that R has the best overall value, because R possesses attributes a_i with values v_i , as exemplified in Figure 16.

Fig. 9. Algorithm for recommendation generation

First, consider the effect of the user model on recommendations at a fixed level of conciseness (z -value of 0.3). Figure 10 shows sample responses, for the East Village Japanese task, for users CK, BA and VM. Because of differences in user model ranking, Bond Street is recommended to CK and Komodo to BA and VM, and different attributes are selected for the three recommendations.

User	Z value	Output
CK	0.3	Bond Street has the best overall value among the selected restaurants. Bond Street has excellent food quality.
BA	0.3	Komodo has the best overall value among the selected restaurants. Komodo’s price is \$29. It’s a Japanese, Latin American restaurant.
VM	0.3	Komodo has the best overall value among the selected restaurants. Komodo’s price is \$29 and it has very good food quality.

Fig. 10. Recommendations for Users CK, BA and VM, for the East Village Japanese Task, for $Z = 0.3$.

Now, consider the algorithm’s implementation for the BA model for varying values of z as illustrated in Figure 11. The setting for z determines the

number of attributes selected to provide evidential support for recommending Komodo. In the experiments reported here, for recommendations, z ranges from -1.5 to 1.5. Generally this means that there is at least one outlying attribute, even for the highest value of z . When there are no outlier attributes, the algorithm simply mentions the restaurants with highest overall value. Figure 8 provides the relevant utility and weighted attribute values. The weighted attribute values for Komodo are 7,10,4,18,2,36 for *food quality*, *service*, *decor*, *cost*, *neighborhood* and *food type* respectively. Outliers are calculated for recommendations with respect to the values for other attributes for the same restaurant. When z is 1.5 or 0.7, only the *food type* attribute is selected. When z is 0.3, the attributes *cost* and *food type* are selected. When z is -0.5, the attributes *cost*, *service* and *food type* are selected. When z is -0.7, the attributes *cost*, *service*, *food quality* and *food type* are selected. When z is -1.5, the attributes *cost*, *service*, *food quality*, *decor* and *food type* are selected.

User	Z-value	Output
BA	1.5	Komodo has the best overall value among the selected restaurants. Komodo's a Japanese, Latin American restaurant.
BA	0.7	Komodo has the best overall value among the selected restaurants. Komodo's a Japanese, Latin American restaurant.
BA	0.3	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29. It's a Japanese, Latin American restaurant.
BA	-0.5	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29 and it has very good service. It's a Japanese, Latin American restaurant.
BA	-0.7	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29 and it has very good service and very good food quality. It's a Japanese, Latin American restaurant.
BA	-1.5	Komodo has the best overall value among the selected restaurants. Komodo's price is \$29 and it has very good service, very good food quality and good decor. It's a Japanese, Latin American restaurant.

Fig. 11. Recommendations for User BA, for the East Village Japanese Task, of Varying Levels of Conciseness.

4.2.3 Generating Comparison Content Plans

The goal of a comparison is to mention several potential candidate options (those with highest overall utility) and provide the user with user tailored ways for choosing among them (expressed as different weighted attribute values).

SPUR's COMPARE strategy can be applied to three or more options. If there are more than five, a subset are first selected according the algorithm in Figure 12 and the content for each option is selected using the algorithm in Figure 13. Because comparisons are inherently contrastive, the algorithm in Figure 13 describes a procedure whereby if a weighted attribute value is an outlier for any option, the attribute value is realized for all options. We use this approach for two reasons: (1) it is not possible for the user to compare options without the

- (1) If the number of restaurants is greater than 5 then
 - (1a) Select the restaurant options R_i that are positive outliers for overall utility (outstanding restaurants). Add a claim C_j to the content plan that the elements of the set R_i have outstanding value.
 - (1b) If there are no outstanding restaurants, select the 5 highest ranked restaurant options R_i for overall utility U_h . Add a claim C_j to the content plan that the elements of the set R_i are the top 5 in overall value.

Fig. 12. Algorithm for selecting a subset of options to compare

- (1) For each option R_i , for each attribute a_i
 - (1a) If the weighted attribute value v_i is an outlier when compared against the weighted attribute value for other options, then add attribute to \$OUTLIER-LIST.
- (2) For each option R_i , for each attribute a_i in \$OUTLIER-LIST, add an assertion s_i to the content plan that R_i has the attribute value v_i , and a relation that s_i *elaborates* the claim C_j .
- (3) For each assertion s_i about an attribute a_i , add a *contrast* relation to the content plan with the s_i as joint nuclei.

Fig. 13. Algorithm for selecting content for subset of options to compare

same information about all of them; and (2) mentioning the same attributes about each option allows a parallel structure in the realization, which supports the user’s inference of contrast [Prevost, 1995, Prince, 1985, Meteer, 1991].

Figure 14 illustrates the effect of the user models on comparisons, for the East Village Japanese task, for a fixed level of conciseness. Each comparison selects a subset of options that are outliers for overall quality for the particular user, given the setting for z . In this case, the z -value of 0.3 selects three options for CK, two for VM and only one for BA. The selected attributes are outliers with respect to the population of attribute values under consideration. For comparisons, the population of values are those for a particular attribute **across** a set of restaurant options. The outlier attributes here are *food quality*, *service* and *decor*, which are realized as in Figure 14. The *food type* attribute is selected for user CK because of the larger set of restaurants from which outliers are calculated.

Now, consider the algorithms in Figure 12 and Figure 13 applied with z values ranging from -1.5 to 1.5. Figure 15 shows comparisons for user VM for varying levels of conciseness. Figure 8 shows the relevant values for overall utility U_h and weighted attribute values.

User	Z value	Output
CK	0.3	Among the selected restaurants, the following offer exceptional overall value. Bond Street's price is \$51. It has excellent food quality, good service and very good decor. It's a Japanese, Sushi restaurant. Japonica's price is \$37. It has excellent food quality, good service and decent decor. It's a Japanese, Sushi restaurant. Komodo's price is \$29. It has very good food quality, very good service and good decor. It's a Japanese, Latin American restaurant.
VM	0.3	Among the selected restaurants, the following offer exceptional overall value. Komodo's price is \$29. It has very good food quality, very good service and good decor. Takahachi's price is \$27. It has very good food quality, good service and decent decor.
BA	0.3	Among the selected restaurants, the following offer exceptional overall value. Komodo has very good service, very good food quality and good decor.

Fig. 14. Comparisons for Users CK, VM and BA, for the East Village Japanese Task.

User	Z-value	Output
VM	1.5	Among the selected restaurants, the following offer exceptional overall value. Komodo has very good service.
VM	0.7	Among the selected restaurants, the following offer exceptional overall value. Komodo has very good service and good decor.
VM	0.3	Among the selected restaurants, the following offer exceptional overall value. Komodo's price is \$29. It has very good food quality, very good service and good decor. Takahachi's price is \$27. It has very good food quality, good service and decent decor.
VM	-0.5	Among the selected restaurants, the following offer exceptional overall value. Komodo's price is \$29. It has very good food quality, very good service and good decor. Takahachi's price is \$27. It has very good food quality, good service and decent decor. Japonica's price is \$37. It has excellent food quality, good service and decent decor.
VM	-0.7	Among the selected restaurants, the following offer exceptional overall value. Komodo's price is \$29. It has very good food quality, very good service and good decor. Takahachi's price is \$27. It has very good food quality, good service and decent decor. Japonica's price is \$37. It has excellent food quality, good service and decent decor. Shabu-Tatsu's price is \$31. It has very good food quality, good service and decent decor.
VM	-1.5	Among the selected restaurants, the following offer exceptional overall value. Komodo's price is \$29. It has very good food quality, very good service and good decor. Takahachi's price is \$27. It has very good food quality, good service and decent decor. Japonica's price is \$37. It has excellent food quality, good service and decent decor. Shabu-Tatsu's price is \$31. It has very good food quality, good service and decent decor. Bond Street's price is \$51. It has excellent food quality, good service and very good decor. Dojo's price is \$14. It has decent food quality, mediocre service and mediocre decor.

Fig. 15. Comparisons for User VM, for the East Village Japanese Task, at Varying Levels of Conciseness.

4.3 Realization of Dialogue Strategies

We developed a template-based realizer that takes as input the content plans that SPUR produces using the algorithms described above and generates a marked-up string to be passed to the text-to-speech module.

Figure 16 illustrates a content plan for recommendations, for user BA for z of -0.7, that the template-based realizer takes as input. The representation of the plans is based on previous work [Marcu, 1997, Mellish et al., 1998], where each plan consists of a set of *assertions* that must be communicated to

the user and a set of *rhetorical relations* that hold between those assertions that may be communicated as well. Each rhetorical relation designates one or more facts as the *nuclei* of the relation, i.e. the main point, and the other facts as *satellites*, i.e. the supplementary facts [Mann and Thompson, 1987]. The content plan in Figure 16 specifies that the nucleus is the assertion that *Komodo has the best overall value*, and that the satellites are the evidential support for this assertion.

Following guidelines from argumentation theory, the strategy for realizing recommendations is to order the nucleus first followed by the satellites. The satellites are ordered to maximize the opportunity for aggregation - to produce the most concise recommendations given the content to be communicated, phrases with identical verbs and subjects are grouped, so that lists and coordination can be used to aggregate the assertions about the subject. Figures 5, 10, and 11 provide examples.

strategy:	recommend
items:	Komodo, Japonica, Takahachi, Shabu-Tatsu, Bond Street, Dojo
relations:	justify(nuc:1;sat:2); justify(nuc:1;sat:3); justify(nuc:1;sat:4); justify(nuc:1;sat:5)
content:	<ol style="list-style-type: none"> 1. assert(best(Komodo)) 2. assert(has-att(Komodo, cost(29))) 3. assert(has-att(Komodo, foodquality(verygood))) 4. assert(has-att(Komodo, service(verygood))) 5. assert(has-att(Komodo, foodtype(Japanese, Latin American)))

Fig. 16. A Content Plan representation for a recommendation for user BA for a Japanese restaurant in the East Village for z of -0.7.

The realizer also lexicalizes each attribute value of the content assertions for both recommendations and comparisons. Each attribute value except for *cost* is mapped to a predicative adjective using the following mapping: 0-13 $\rightarrow$ *mediocre*; 14-16 $\rightarrow$ *decent*; 17-19 $\rightarrow$ *good*; 20-22 $\rightarrow$ *very-good*; 23-25 $\rightarrow$ *excellent*; above 25 $\rightarrow$ *superb*. Cost is not lexicalized in this way, because user pilots showed little consensus between users about mapping absolute cost to specific lexical items, i.e. \$30 is an expensive meal for some, but cheap for others. Thus the cost attribute is referred to as *price* and its real value is given in the description.

Figure 17 illustrates a content plan for comparisons, for user VM for z of 0.3, that the template-based realizer takes as input. The option selection algorithm in Figure 12 determines that Takahachi and Komodo are outliers for overall utility, thus the nucleus is the assertion that Komodo and Takahachi are exceptional restaurants and the satellites are assertions about the selected attributes for each restaurant. Contrast relations hold between pairs of assertions about attributes. The realization template for comparisons focuses on communicating both the *elaboration* and the *contrast* relations. One way to communicate the *elaboration* relation between the nuclei and the satellites is

strategy:	compare
items:	Komodo, Takahachi, Japonica, Shabu-Tatsu, Bond Street, Dojo
relations:	elaboration(nuc:1, sat:2); elaboration(nuc:1, sat:3); elaboration(nuc:1, sat:4); elaboration(nuc:1, sat:5); elaboration(nuc:1, sat:6); elaboration(nuc:1, sat:7); elaboration(nuc:1, sat:9); elaboration(nuc:1, sat:9); contrast(nuc:2, nuc:3); contrast(nuc:4, nuc:5); contrast(nuc:6, nuc:7); contrast(nuc:8, nuc:9)
content:	1. assert(exceptional(Komodo's, Takahachi's)) 2. assert(has-att(Komodo, cost(29))) 3. assert(has-att(Takahachi's, cost(27))) 4. assert(has-att(Komodo, service(verygood))) 5. assert(has-att(Takahachi's, service(good))) 6. assert(has-att(Komodo, decor(good))) 7. assert(has-att(Takahachi's, decor(decent))) 8. assert(has-att(Komodo, foodquality(verygood))) 9. assert(has-att(Takahachi's, foodquality(good)))

Fig. 17. A Content Plan representation for a Comparison for a Japanese restaurant in the East Village for user VM for $z = 0.3$

to structure the comparison so that all the satellites are grouped together, following the nucleus. In order to communicate the *contrast* relation, these satellites are produced in a fixed order, with a parallel structure maintained **across** options [Prevost, 1995, Prince, 1985]. The satellites are initially ordered in terms of their evidential strength, but then are reordered to allow for aggregation in order to produce the most succinct descriptions. Examples are given in Figures 14 and 15.

5 Experimental evaluation

Our experiments evaluated 4 main hypotheses, concerning tailoring, conciseness, mode and the interaction between tailoring and mode.

- **Tailoring:** We expected users to prefer tailored to untailored system responses;
- **Mode:** We expected users to prefer text to speech responses;
- **Tailoring/Mode:** We expected tailoring to have a greater effect on judgments of speech as opposed to text responses because speech imposes a greater cognitive load;
- **Conciseness:** We expected users to be sensitive to the amount of information provided in system responses, and to prefer concise to verbose responses.

We test the first three hypotheses in the tailoring experiment described in Sections 5.1 and 5.2. We test the final hypothesis with a separate experiment that directly manipulates the conciseness parameter while holding the user model

constant for the particular user who is acting as subject in the experiment. This experiment is described in Sections 5.3 and 5.4.

In both experiments, the user models were collected in a separate process that took place before the experiments. We also carried out a pre-experimental survey where subjects provided information about where they live, the frequency of eating in restaurants in general, and their familiarity with Manhattan.

There were six experimental tasks altogether, each involving one or two constraints on the selection of a set of restaurant options: (a) French restaurants; (b) restaurants in Midtown West; (c) Italian restaurants in the West Village; (d) Asian restaurants in the Upper West Side; (e) cheap restaurants; (f) Japanese restaurants in the East Village. Only tasks a,b,c and d were used in the tailoring and mode experiment whereas all six tasks were used in the conciseness experiment. The tasks were chosen after extensive piloting to accommodate a variety of user models, to be fairly easy for subjects to remember, and to provide sets of restaurants large enough to be interesting. The order of the presentation of tasks is consistent across subjects.

The experimental procedure for both experiments treats the subject as an “overhearer” of a series of dialogues, each involving one restaurant-selection task [Walker et al., 2001a, Whittaker and Walker, 2002]. Each dialogue about a task consists of a sequence of dialogue exchanges between the user and the system, with each exchange presented on a separate web page. The initial web page for each task sets up the task by showing the MATCH system’s graphical response for an initial user query, e.g. *Show Italian restaurants in the West Village*. Prior research indicated that a typical dialogue structure in this domain is for users to identify promising candidate restaurants and request more information about these (compare) and then request detailed information about a single specific option (recommend) [Whittaker et al., 2002]. Therefore for all subjects and tasks, the following pages show the user first circling some subset of the restaurants and asking the system to COMPARE them, and then to RECOMMEND options from the circled subset. The sample dialogue in Figure 2 illustrates this dialogue structure. The main advantage of the “overhearer” method is that it allows users to give specific feedback about alternative system responses in the context in which they are provided.

5.1 Experimental Design for Tailoring and Mode Experiment

The first experiment tests the Tailoring, Mode and Tailoring/Mode hypotheses and consists of dialogues involving tasks a, b, c and d, as described above. To test the Tailoring hypothesis, the subject sees two types of responses on each web-page for each dialogue exchange, one tailored to her user model, and

the other tailored to the user model of another randomly selected subject. We then compare subjects' judgments of the two responses. By randomly selecting another user model (Random), and using the distance between two user models, we can both test whether having one's own model is better than someone else's, and quantify *how much* distance there has to be between two user models to make a difference in the subject's perception of system responses. The order of presentation of subject-tailored and other-tailored responses is randomized from page to page.

For each instance of a recommend, or compare strategy, the subject is asked to state her degree of agreement on a 5-point Likert scale (a standard technique for mapping subjective responses to scalar values [Likert, 1932]) with the following statement, intended to determine the *informativeness*, or *information quality*, of the response: *The system's utterance is easy to understand and it provides exactly the information I am interested in when choosing a restaurant.* The statement refers to both comprehensibility and informativeness. We asked about both these dimensions in a single compound statement because our algorithms were intended to simultaneously optimize both the exact information presented (the second part of the statement), and the format in which it was presented (the first part of the statement). Extensive piloting showed a statement about comprehensibility alone favored a short response and a statement about informativeness alone favored a long response. Responses to this compound statement are measured as InformationQuality.

Since the algorithms for recommendations and comparisons consist first of algorithms for ranking restaurant options, and then for selecting content, we ask users to provide judgements related to the ranking of options as a secondary measurement of the efficacy of system responses. For each instance of a recommendation, the subject is asked to state her degree of agreement with this statement (again on a 5 point Likert scale): *I am confident that the recommended restaurant is someplace I would like to go.* A similar statement is used to evaluate the ranking of options for comparisons between three or more restaurants: *I am confident that the restaurants being described are places I would like to go.* User responses to these questions are measured as the variable RankingConfidence in the results below.

In order to test the Mode and Tailoring/Mode hypotheses the entire sequence of web pages is presented twice. The first time through, the subject can only read (not hear) the system responses. The second time, she can only hear them. We used this read-then-hear approach (again after extensive piloting), to familiarize subjects with proper names in the restaurant domain. Prior text presentation means that proper names are primed for users in the speech condition making them less likely to be misunderstood. But the fact that text and speech presentations of the same task are 10-15 minutes apart means that users cannot remember their judgments for the previous instance of the task.

We test the Mode Hypothesis by asking users to judge the same responses in the same context first in text and later in speech. We test the Mode/Tailoring hypothesis by comparing the effects of tailoring on speech with its effects on text. We expect that providing a user model will have greater effects for speech than text because of the greater problems that users experience in processing complex speech outputs.

To summarize, each subject “overhears” a sequence of four dialogues about different restaurant-selection tasks. The entire sequence is presented twice (once for text, once for speech). The subject makes six InformationQuality judgments for each dialogue each time made up of (a) one recommendation and two comparisons tailored to the subjects user model; and (b) one recommendation and two comparisons tailored to a randomly selected user model. The total number of InformationQuality judgments per subject is 48. The subject makes four RankingConfidence judgements for each dialogue each time. The total number of confidence judgements per subject is 32. The total time required to complete the experiment is approximately half an hour per subject.

Sixteen subjects who had previously enrolled with the system took part in the experiment as volunteers. All were fluent English speakers. Most eat out moderately often (seven eat out 3-5 times per month, six 6-10 times). All sixteen currently live in northern New Jersey. Eleven described themselves as somewhat or quite familiar with Manhattan, while five thought they were not very familiar with it.

5.2 *Tailoring experimental results*

We first tested whether differences in the user model affected subjects’ rankings of the InformationQuality of the system’s responses. A paired t-test confirmed the Tailoring Hypothesis that people prefer responses generated with their own model than with a randomly assigned model ($t(383)=1.76$, $p<0.05$, for a one-tailed test).

However, although the predicted effect is significant, this is a conservative test: the Random model condition includes cases where the randomly assigned model is close to the User’s Own model. We therefore filtered the original set of judgments to exclude cases where the distance between the Random Model and the User’s Own Model was less than 0.3, to exclude these similar cases. This removed 9% of judgments from the original data set. To test our hypotheses, we conducted two analyses of variance with Model Type (Own,Random) * Mode (Speech,Text) * Strategy (Recommend, Compare) as independent variables and judgments of InformationQuality or RankingConfidence as the dependent variables. As predicted, there were main effects for Model Type,

both for InformationQuality ($F=5.6$, $df = 1,674$, $p < 0.02$) and RankingConfidence ($F=4.3$, $df = 1,674$, $p < 0.05$), showing that using the User’s Own Model significantly improved system responses, and confirming the Tailoring Hypothesis.

Our results partially confirm the Mode Hypothesis. For InformationQuality, as predicted, Mode was significant, ($F=3.8$, $df = 1,674$, $p < 0.05$), with Text responses being rated more highly than Speech. Mode has no significant effect on users’ RankingConfidence, however ($F=0.1$, $df = 1,674$, $p = 0.9$).

Finally, and contrary to our predictions, there was no interaction between Model Type and Mode ($F=0.02$, $df = 1,674$, $p > 0.05$), so the Mode/Tailoring Hypothesis was not confirmed. One possible explanation is that there were floor effects for Speech judgments and this in turn reduced the variance in these judgments. Nevertheless, there were no differences between judgments of text with the random model and tailored speech ($t(397)=1.8$, $p > 0.05$). This offers some evidence that with previous exposure to restaurant names and proper name priming it is possible to overcome limitations of speech by the use of tailoring.

5.3 Conciseness experimental design

We now describe our evaluation of the Conciseness Hypothesis. The goal of this experiment is to: (1) Test whether our manipulations of conciseness correspond to user’s perceptions of conciseness; and (2) Determine an optimal level of conciseness for recommendations and comparisons. We first addressed user’s sensitivity to conciseness and the correspondence between algorithmic conciseness and user judgments of conciseness. Our expectation was users would discriminate between different descriptions in terms of conciseness. More specifically, we expected that outputs we had operationalized as *concise* should be judged as providing too little information, outputs operationalized as *sufficient* should be judged as providing the right amount of information, and outputs operationalized as *verbose* should be judged as providing too much information.

A second focus was the relation between conciseness and information presentation strategy. Contrast the recommendations in Figure 11 with the comparisons in Figure 15 for varying values of z . Of the two strategies, comparisons inherently contain more information than recommendations, because they mention multiple options and their attributes. We should therefore expect users to judge comparisons as more verbose than recommendations.

This experiment used all six tasks described above. As before, an initial web page set up the task by showing the MATCH system’s graphical response for

a user query, and then the page showed the user circling some subset of the restaurants and asking the system to first compare and then recommend options from the circled subset. Subjects saw one page each for recommend and compare, for each task. In this case, on each page, they saw multiple system responses of differing conciseness. We operationalized *concise* responses as a z-value of 0.3, *sufficient* responses as a z-value of -0.7, and *verbose* responses as a z-value of -1.5. As before, the order of the tasks, and the order of appearance of strategies within the task was consistent across subjects. However, the order of presentation of conciseness variants was randomized from page to page. For each instance of a recommend, or compare, the subject was asked to state her degree of agreement (on a 5-point Likert scale) with the following statement, intended to determine the conciseness of the response: *When choosing a restaurant, the amount of information provided by the system utterance is; (1) far too little, (2) too little, (3) neither too little nor too much, (4) too much, (5) far too much.*

Twenty one subjects completed the experiment in approximately half an hour per subject. All were fluent English speakers. Most eat out moderately often (eleven eat out 3-5 times per month, ten 6-10 times). All subjects currently live in northern New Jersey or in Manhattan. Fourteen described themselves as somewhat or quite familiar with Manhattan, while seven were not very familiar with it.

5.4 Conciseness Experimental Results

We analyzed the user data using ANOVA. Independent measures were Algorithmic Conciseness (Verbose, Sufficient, Concise), and Strategy (Recommend, Compare). Using standard Likert scale procedures we first transformed the elicited Conciseness judgments into a linear scale, so that an output judged to provide *far too little information* was scored -2, *too little* -1, *neither too much nor too little* 0, *too much* +1, and *far too much* +2. The transformed measure of conciseness was used as the ANOVA dependent measure.

Figure 18 indicates the relationship between Algorithmic conciseness and user judgments of Conciseness. It shows both that users are sensitive to conciseness and that user judgments paralleled our algorithmic implementation. Consistent with our hypothesis, outputs generated as concise were more likely to be judged as having too little information than those generated to be sufficient, which in turn were likely to have less information than those generated to be verbose ($F(2,750) = 220.8$, $p < 0.0001$), with post hoc tests showing judged differences between Algorithmically Concise and Sufficient, and between Algorithmically Sufficient and Verbose (both $p < 0.0001$). These data clearly show that we have algorithmic control over conciseness.

Fig. 18. Relationship between Algorithmic Conciseness and User Evaluations by Strategy.

Nevertheless, Figure 18 also indicates the need for further calibration of the algorithm. If our algorithmic calibration had been correct, we would have expected verbose outputs to have been judged as providing too much information (scored as +1 or greater on the Likert scale), sufficient outputs judged as 0, and concise as -1 or less. Results show that sufficient outputs require little further calibration as they are judged at -0.2 (where “0” indicates exactly the right amount of information), but those generated to be verbose are judged as 0.5, and those generated to be concise are generated as -1.2. These observations suggest that we may be providing marginally too much information for our concise outputs and too little for our verbose outputs. This would imply a need to tune the algorithm, in particular by adding more information to the sufficient statements.

Our second hypothesis concerned the relationship between judged conciseness and strategy. Figure 18 also shows as predicted that recommendations are judged to be more concise than comparisons ($F(1,750)=19.7$, $p<0.0001$). Furthermore, there is an interaction between strategy and judgments ($F(2,750)=10.0$, $p<0.0001$), with the main difference being accounted for by users’ tendency to judge verbose comparisons as containing more information than verbose recommendations (post hoc test, $p < 0.05$). Possibly this was because verbose comparisons mention as many as 10 facts, and this is perceived to be a large additional memory burden. Finally, despite our tailoring of information content to individual users’ preferences, there are large individual differences between users in terms of their perception of conciseness, suggesting that the conciseness parameter itself should be user-tailored.

6 Conclusions and Future Work

This paper describes an approach to user tailored generation of evaluative responses for multimodal dialogue systems that is based on quantitative user models. We address a pressing problem for current systems, namely that information presentation strategies overload users, and do not effectively support them in making decisions about complex options Walker et al. [2002a]. We present new algorithms for information presentation based on multi-attribute decision theory that focus the presentation on small sets of options and attributes that are significant and salient to the user. These algorithms enable both option and attribute selection for two different dialogue strategies: recommendations and comparisons. We have implemented the algorithms for generating content plans for these strategies in SPUR, a content planner for the MATCH dialogue system. Furthermore our theoretical framework allows parameters of the content plans to be highly configurable: allowing us to generate differently concise content plans, that highlight and compare different sets of attributes and options.

Our results show that user models based on multi-attribute decision theory generalize across domains. Techniques that work for other domains are also effective in the restaurant domain, as well as being effective for multi-modal dialogue, where the requirements for interactive information presentation are different from those for text presentations. We have also extended user-tailored generation to include comparisons as well as recommendations. Our results show the effects of user-tailoring on these strategies. Users rated responses generated using their Own Model much more highly than those generated with the Random Model.

We also demonstrated effects of presentation mode. As we expected, text responses were rated more highly than speech responses. Despite our attempts to prime text-to-speech pronunciation, users complained about difficulty understanding restaurant names, which are often foreign words. Contrary to our expectations, we found that the effect of tailoring was no greater in the speech than the text condition. However this may have been due to the fact that overall ratings of responses were low in the speech condition so that effects may be observed in a domain that is less demanding for text-to-speech. In support of this, we found that users who ate out more frequently or who knew Manhattan better rated speech responses more highly. This suggests that people who are familiar with restaurant names and general restaurant information are better able to overcome perceptual limitations associated with understanding text-to-speech output.

We also successfully implemented a strategy for controlling presentation conciseness, and showed that outputs from our algorithm were consistent with

user judgments of conciseness, although some fine tuning of the algorithm will be necessary to exactly map onto absolute user judgments.

In the future we plan to conduct additional experiments in this framework. First of all, we would like to test the effect of the user model and conciseness parameters on other variables such as task completion, time to completion and user satisfaction, as in other work evaluating spoken dialogue systems [Walker et al., 2002b]. Another area of additional experimentation is in the mapping between selected content and dialogue strategy. For example, we did not vary the content plan templates for each strategy in this experiment, although in our own exploration we identified various possibilities for each strategy. Additional experiments could alter different aspects of the template, and explore subject preferences for the resulting output. In current work, we are enriching SPUR’s ability to structure the selected content, and interfacing SPUR to a sentence planner and surface realizer [Stent et al., 2004, Walker et al., 2003, 2001b, Bangalore and Rambow, 2000]. We also hope to conduct field trials of people using the system in a mobile environment.

References

- S. Bangalore and M. Johnston. Tight coupling of multimodal language processing with speech recognition. In *Proceedings of ICSLP*, Beijing, China, 2000.
- Srinivas Bangalore and Owen Rambow. Exploiting a probabilistic hierarchical model for Generation. In *COLING*, Saarbrücken, Germany, 2000.
- L. Bell, J. Boye, J. Gustafson, and M. Wirn. Modality convergence in a multimodal dialogue system. In *Proceedings of Gotalog 2000, Fourth Workshop on the Semantics and Pragmatics of Dialogue*, pages 29–34, 2000.
- M. Beutnagel, A. Conkie, J. Schroeter, Y. Stylianou, and A. Syrdal. The AT&T Next-Generation Text-to-Speech System. In *Meeting of ASA/EAA/DAGA in Berlin, Germany*, 1999.
- S. E. Brennan. Conversations with and through computers. *User Modeling and User-Adapted Interaction*, 1:67–86, 1991.
- Susan E. Brennan. Lexical entrainment in spontaneous dialog. In *1996 International Symposium on Spoken Dialogue*, pages 41–44, 1996.
- Susan E. Brennan and Herbert H. Clark. Lexical choice and conceptual pacts in conversation. *Journal of Experimental Psychology: Learning, Memory And Cognition*, 1996.
- G. Carenini. *Generating and evaluating evaluative arguments*. PhD thesis, University of Pittsburgh, 2000.
- G. Carenini and J. Moore. A strategy for generating evaluative arguments. In *Proceedings of the 1st International Conference on Natural Language Generation (INLG-00)*, Mitzpe Ramon, Israel, 2000a.

- Giuseppe Carenini and Johanna D. Moore. An empirical study of the influence of argument conciseness on argument effectiveness. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL-00)*, 2000b.
- Giuseppe Carenini and Johanna D. Moore. An empirical study of the influence of user tailoring on evaluative argument effectiveness. In *IJCAI*, pages 1307–1314, 2001.
- A. Cawsey. Planning interactive explanations. *International Journal of Man-Machine Studies*, 38:169–199, 1993.
- D. Chin. KNOVE: Modeling what the user knows in UC. In A. Kobsa and W. Wahlster, editors, *User Models in Dialog Systems*, pages 74–107. Springer-Verlag, Berlin, 1989.
- Herbert H. Clark and Deanna Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22:1–39, 1986.
- Cogentex. Cogentex recommender system, 2003.
- Philip R. Cohen, M. Johnston, D. McGee, S. L. Oviatt, J. Pittman, I. Smith, L. Chen, and J. Clow. Multimodal interaction for distributed interactive simulation. In M. Maybury and W. Wahlster, editors, *Readings in Intelligent Interfaces*. Morgan Kaufmann Publishers, 1998.
- E. P. J. Corbett and R. J. Connors. *Classical Rhetoric for the Modern Student*. Oxford University Press, 1999.
- Rachel Coulston, Sharon Oviatt, and Courtney Darves. Amplitude convergence in children’s conversational speech with animated personas. In *Proceedings of the 7th International Conference on Spoken Language Processing (ICSLP’2002)*, pages 2689–2692, 2002.
- Courtney Darves and Sharon Oviatt. Adaptation of users’ spoken dialogue patterns in a conversational interface. In *Proceedings of the International Conference on Spoken Language Processing (ICSLP’2002)*, 2002.
- Ward Edwards and F. Hutton Barron. Smarts and smarter: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes*, 60:306–325, 1994.
- S. Garrod and A. Anderson. Saying what you mean in dialogue: A study in conceptual and semantic coordination. *Cognition*, 27:181–218, 1987.
- J. Goecks and J. Shavlik. Learning users’ interests by unobtrusively observing their normal behavior. In *Proceedings of the 2000 International Conference on Intelligent User Interfaces*, pages 129–132, 2000.
- Helen Hastie, Patrick Ehlen, and Michael Johnston. Context-sensitive multimodal help. In *Proceedings of the 4th International Conference on Multimodal Interfaces*, 2002.
- Anthony Jameson, Ralph Schafer, Joep Simons, and Thomas Weis. Adaptive provision of evaluation-oriented information: Tasks and techniques. In *IJCAI*, pages 1886–1895, 1995.
- M. Johnston and S. Bangalore. Finite-state multimodal parsing and understanding. In *Proceedings of COLING 2000*, Saarbrücken, Germany, 2000.
- M. Johnston and S. Bangalore. Finite-state multimodal integration and un-

- derstanding. *Natural Language Engineering*, 2004.
- M. Johnston, P.R. Cohen, D. McGee, S.L. Oviatt, J.A. Pittman, and I. Smith. Unification-based Multimodal Integration. In *Annual Meeting of the Association for Computational Linguistics, ACL*, 1997.
- Michael Johnston. Multimodal language processing. In *Proceedings of ICSLP*, Sydney, Australia, 1998.
- Michael Johnston, Srinivas Bangalore, Gunaranjan Vasireddy, Amanda Stent, Patrick Ehlen, Marilyn Walker, Steve Whittaker, and Preetam Maloor. MATCH: An architecture for multimodal dialogue systems. In *Annual Meeting of the Association for Computational Linguistics, ACL*, 2002.
- Aravind K. Joshi. Mutual beliefs in question-answer systems. In Neil V. Smith, editor, *Mutual Knowledge*, pages 181–199. Academic Press, New York, New York, 1982.
- Aravind K. Joshi, Bonnie Webber, and Ralph M. Weischedel. Some aspects of default reasoning in interactive discourse. Technical Report MS-CIS-86-27, University of Pennsylvania, 1986.
- Aravind K. Joshi, Bonnie L. Webber, and Ralph M. Weischedel. Preventing false inferences. In *COLING84: Proc. 10th International Conference on Computational Linguistics*, pages 134–138, 1984.
- Aravind K. Joshi and Scott Weinstein. Control of inference: Role of some aspects of discourse structure - centering. In *Proc. International Joint Conference on Artificial Intelligence*, pages 385–387, 1981.
- R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, 1976.
- G. Kempen and E. Hoenkamp. An incremental procedural grammar for sentence formulation. *Cognitive Science*, 11:201–258, 1987.
- D. Klein. *Decision Analytic Intelligent Systems: Automated Explanation and Knowledge Acquisition*. Lawrence Erlbaum Associates, 1994.
- S. Larsson, P. Bohlin, J. Bos, and D. Traum. Trindikit manual. Technical report, TRINDI Deliverable D2.2, 1999.
- W. J. M. Levelt. *Speaking: From Intention to Articulation*. MIT Press, 1989.
- Willem J. Levelt and Stephanie Kelter. Surface form and memory in question answering. *Cognitive Psychology*, 14(1):78–106, 1982.
- R. Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 140, 1932.
- Greg Linden, Steve Hanks, and Neal Lesh. Interactive assessment of user preference models: The automated travel assistant. In *Proceedings of the Fifth International Conference on User Modeling*, 1997.
- Joseph A. Luchok and James C. McCroskey. The effect of quality of evidence on attitude change and source credibility. *The Southern Speech Communication Journal*, 43:371–383, 1978.
- W.C. Mann and S.A. Thompson. Rhetorical structure theory: Description and construction of text structures. In Gerard Kempen, editor, *Natural Language Generation*, pages 83–96. Martinus Nijhoff, 1987.
- Daniel Marcu. From local to global coherence: a bottom-up approach to text

- planning. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'97)*, 1997.
- K. J. Mayberry and R. E. Golden. *For Argument's Sake: A Guide to Writing Effective Arguments*. Harper Collins, 1996.
- W. J. McGuire. The nature of attitudes and attitude change. In G. Lindzey and E. Aronson, editors, *The Handbook of Social Psychology*, pages 3:136–314. Addison-Wesley, 1968.
- K. McKeown, S. Feiner, M. Dalal, and S-P Chang. Generating multimedia briefings; coordinating language and illustrations. *Artificial Intelligence Journal*, 103:95–116, 1998.
- Cynthia A. McLemore. Prosodic variation across discourse types. In *Workshop on Prosody in Natural Speech*. Institute for Research in Cognitive Science, University of Pennsylvania, 1992.
- Chris Mellish, Alistair Knott, Jon Oberlander, and Mick O'Donnell. Experiments using stochastic search for text planning. In *Proceedings of International Conference on Natural Language Generation*, pages 97–108, 1998.
- Marie Meteer. Bridging the generation gap between text planning and linguistic realization. *Computational Intelligence*, 7(4):296–304, 1991.
- M. D. Miller and T. R. Levine. Persuasion. In M. B. Salwen and D. W. Stack, editors, *An Integrated Approach to Communication Theory and Research*, pages 261–276. 1996.
- T. Mitchell. *Machine Learning*. CHECK, 1998.
- V. O. Mittal, S. Roth, J. D. Moore, J. Mattis, and G. Carenini. Generating explanatory captions for information graphics. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI '95)*, pages 1276–1283, 1995.
- J. D. Moore and C. L. Paris. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4), 1993.
- K. Morik. User models and conversational settings: Modeling the user's wants. In A. Kobsa and W. Wahlster, editors, *User Models in Dialog Systems*, pages 364–385. Springer-Verlag, Berlin, 1989.
- Gopalan Nadathur and Aravind K. Joshi. Mutual beliefs in conversational systems: Their role in referring expressions. In *Proceedings International Joint Conference on Artificial Intelligence, Austin*, pages 603–605, 1983.
- C. Nass, J. Steuer, and E. Tauber. Computers are social actors. In *Proceedings of the Conference on Computer Human Interaction, CHI-94*, pages 72–78, 1995.
- Sharon L. Oviatt. Multimodal interactive maps: Designing for human performance. In *Human-Computer Interaction*, pages 93–129, 1997.
- Sharon L. Oviatt. Mutual disambiguation of recognition errors in a multimodal architecture. In *CHI '99*, pages 576–583. ACM Press, New York, 1999.
- Cecile L. Paris. Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14(3):64–78, 1988.
- Scott Prevost. *A Semantics of Contrast and Information Structure for Speci-*

- fying Intonation in Spoken Language Generation*. PhD thesis, University of Pennsylvania, 1995.
- Ellen F. Prince. Fancy syntax and shared knowledge. *Journal of Pragmatics*, 9(1):65–81, 1985.
- R. Rafter, K. Bradley, and B. Smyth. Personalized retrieval for online recruitment services. In *Proceedings of the 22nd Annual Colloquium on Information Retrieval*, 2000.
- Elaine Rich. User modeling via stereotypes. *Cognitive Science*, 3:329–354, 1979.
- S. Rogers and C. Fiechter. An adaptive interactive agent for route advice. In *Proceedings of the Third International Conference on Autonomous Agents*, 1999.
- Michael F. Schober. Different kinds of conversational perspective-taking. In S. R. Fussell and R. J. Kreuz, editors, *Social and cognitive psychological approaches to interpersonal communication*, pages 145–174. Lawrence Erlbaum, 1998.
- R.D. Sharp, E. Bocchieri, C. Castillo, S. Parthasarathy, C. Rath, M. Riley, and J Rowland. The watson speech recognition engine. In *Proc. ICASSP97*, pages 4065–4068, 1997.
- M. R. Solomon. *Consumer Behavior: Buying, Having and Being*. Prentice Hall, 1998.
- J. Srivastava and T. Connolly. Do ranks suffice? a comparison of alternative weighting approaches in value elicitation. *Organization Behavior and Human Decision Processes*, 63(1):112–116, 1995.
- Amanda Stent, Rashmi Prasad, and Marilyn Walker. Trainable sentence planning for complex information presentation in spoken dialogue systems. In *Meeting of the Association for Computational Linguistics*, 2004.
- Cynthia A. Thompson and Mehmet H. Goker. Learning to suggest: The adaptive place advisor. In *Proceedings of the Third International Conference on Autonomous Agents*, 1999.
- W. Wahlster and A. Kobsa. User models in dialogue systems. In *User Models in Dialogue Systems*, pages 4–34. Springer Verlag, Berlin, 1989.
- M. Walker, J. Aberdeen, J. Boland, E. Bratt, J. Garofolo, L. Hirschman, A. Le, S. Lee, S. Narayanan, K. Papineni, B. Pellom, J. Polifroni, A. Potamianos, P. Prabhu, A. Rudnicki, G. Sanders, S. Seneff, D. Stallard, and S. Whittaker. DARPA communicator dialog travel planning systems: The june 2000 data collection. In *EUROSPEECH 2001*, 2001a.
- M. Walker, O. Rambow, and M. Rogati. Spot: A trainable sentence planner. In *Proceedings of the North American Meeting of the Association for Computational Linguistics*, 2001b.
- M. Walker, A. Rudnicki, R. Prasad, J. Aberdeen, E. Bratt, J. Garofolo, H. Hastie, A. Le, B. Pellom, A. Potamianos, R. Passonneau, S. Roukos, G. Sanders, S. Seneff, and D. Stallard. DARPA communicator: Cross-system results for the 2001 evaluation. In *ICSLP 2002*, 2002a.
- M. A. Walker, S. J. Whittaker, A. Stent, P. Maloor, J. D. Moore, M. Johnston,

- and G. Vasireddy. Speech-Plans: Generating evaluative responses in spoken dialogue. In *In Proceedings of INLG-02.*, 2002b.
- Marilyn Walker, Rashmi Prasad, and Amanda Stent. A trainable generator for recommendations in multimodal dialog. In *EUROSPEECH*, 2003.
- Marilyn A. Walker. The Effect of Resource Limits and Task Complexity on Collaborative Planning in Dialogue. *Artificial Intelligence Journal*, 85(1-2): 181-243, 1996.
- Nigel Ward and Satoshi Nakagawa. Automatic user-adaptive speaking rate selection for information delivery. In *Proceedings of the International Conference on Spoken Language Processing (ICSLP'2002)*, 2002.
- Bonnie Webber and Aravind Joshi. Taking the initiative in natural language database interaction: Justifying why. In *Coling 82*, pages 413-419, 1982.
- Steve Whittaker, Susan Brennan, and Herb Clark. Coordinating activity: an analysis of interaction in computer supported cooperative work. In *CHI 91*, page 1, 1991.
- Steve Whittaker and Marilyn Walker. Evaluating dialogue strategies in multimodal dialogue systems. In *IDS02*, 2002.
- Steve Whittaker, Marilyn Walker, and Johanna Moore. Fish or fowl: A Wizard of Oz evaluation of dialogue strategies in the restaurant domain. In *Language Resources and Evaluation Conference*, 2002.
- Lizhong Wu, Sharon L. Oviatt, and Philip R. Cohen. Multimodal integration – a statistical view. *IEEE Transactions on Multimedia*, 1(4):334-341, December 1999.
- I. Zukerman and D. Litman. Natural language processing and user modeling: Synergies and limitations. *User Modeling and User-Adapted Interaction*, 11(1-2):129-158, 2001.
- I. Zukerman, R. McConachy, and K. Korb. Using argumentation strategies in automated argument generation. In *Proceedings of the 1st International Natural Language Generation Conference*, pages pages 55-62, 2000.
- I. Zukerman and Richard McConachy. Generating concise discourses that addresses a user's inferences. In *IJCAI 93*, 1993.