

Guiding User Choice During Discussion by Silence, Examples and Justifications

Maier Fenster¹ and Inon Zuckerman² and Sarit Kraus³

Abstract. This paper describes an approach for guiding human choice-making by a computerized agent, in a conversational setting, where both user and agent provide meaningful input. In the proposed approach, the agent attempts to convince a person by providing examples for the person to emulate or by providing justifications for the person to internalize and build or change her preferences accordingly. The agent can take into account examples and justifications provided by the person. In a series of experiments where the task was selecting a location for a school, a computer agent interacted with subjects using a textual chat-type interface, with different agent designs being used in different experiments. The results show that the example-providing agent outperformed the justification providing agent and both, surprisingly, outperformed an agent which presented the subject with both examples and justifications. In addition, it was demonstrated that in some cases the best strategy for the agent is to keep silent.

1 Introduction

Human interaction with computerized systems is widespread over the Internet. Many of these interactions are used as a setting for persuasion of a user by the system [6]. Examples include recommendation systems in which a computerized advisor provides suggestions based on the user's profile or input (e.g., Amazon.com) and systems that personalize advertisements based on user profile or intent.

We are particularly interested in interactions where a salesperson "talks the customer" into unplanned and/or expensive purchases. Often, such a salesperson does not argue with the customer, and uses a "yes and" type of non-judgmental conversational style rather than a "yes but" or "no because" critical conversational style. This is elaborated below as a "discussion game".

In this work we build upon previous persuasion techniques to tackle the following problem: given a large set of choices with a subset of choices preferred by an agent, the agent should guide the user to select choices from the desired subset rather than from the larger set. With that in mind, our system extends the current literature as follows: (1) First, we focus on a natural language discussion setting (in contrast to catalog selection), where both agent and user interact multiple rounds through plain text interface [16]. (2) Second, the priming effect (see below) was used successfully in prior personalized recommendation systems [4]. The goals of these systems were to lead the user to make a choice that would be beneficial for herself.

We, on the other hand consider a self-interested agent and study how priming can be used to lead the user to make a choice that is beneficial for the agent (e.g., as a delegate of the commercial website that deploys it). (3) Finally, our work follows the theory of "construction of preference" (e.g., [18]), which teaches that people's preferences are often constructed in the process of elicitation or decision making, rather than being pre-constructed. Our agent aids the user in constructing preferences which meet the agent's goals and can also analyze the user choices to follow such construction as it happens.

In a discussion setting, an agent is provided with several ways to persuade a user. First, the agent can provide the user with suggestions, in the form of examples to emulate. Second, the agent can provide a justification for a particular suggestion or a general justification not linked to a particular suggestion. Finally, the agent can remain silent and allow the user to act out her own preferences. Success is measured by an increased utility to the agent, brought about by the user selecting choices from the agent-preferred subset rather than from the whole set of possible choices. This setting also tries to tackle an inherent limitation of recommender systems in that people tend to be more satisfied with their choices if they feel that they made the choice on their own, even over time [5]. It is also the case that some individuals tend to reject suggestions that are explicitly made [13]. Note that our agent *does not* disagree with the user.

To this end we built a discussion system including agents which can conduct chat-like discussion with a user, using simple Natural Language Processing (NLP) to "understand" the user and to provide suggestions and/or justifications in response. In general, the justifications presented by the agent are selected to encourage the user to consider certain attributes of the setting at the expense of other attributes and to the benefit of the agent. In other words, the justifications are used to *prime* the user to the agent's preferred subset.

We programmed an experimental setting denoted as "School place", a setting where the agent and a user interact over the construction location for a new school. We conducted extensive experiments with more than 420 subjects on Amazon Mechanical Turk, using an agent which uses only one of the persuasion strategies for each experiment. We showed that the agent can significantly affect the subject's choices and the effectiveness of different strategies depend on the subject's a priori inclination. Our results also show that agents using suggestions generally out-perform agents using justifications and both generally outperform agents using a combination of suggestions and justifications. Moreover, it appears that in some scenarios, silence is the best strategy.

Our proposed agents are substantially domain independent and we believe that they can be used in a wide range of domains. In addition, another benefit is that our approach can be used also when the agent has limited information about a user. This type of approach can

¹ Dept. of Computer Science, Bar-Ilan University, Ramat Gan, 52900, Israel, email: maier.fenster@gmail.com

² Dept. of Industrial Engineering and Management, Ariel University Center of Samaria, Ariel, 40700, Israel.

³ Dept. of Computer Science, Bar-Ilan University, Ramat Gan, 52900, Israel. Sarit Kraus is also Affiliated with UMIACS.

be useful, for example, in one-off sales situations, especially if user does not know what she wants, or when dealing with the cold-start problem [14].

2 Related Work

The problem of guiding human choice-making by a computerized agent has been addressed using various approaches [15]. With respect to persuasion methods, Cialdini [6] provides a useful overview of persuasion methods commonly used in commercial settings. Fogg [7] is a foundational publication and an introduction to computer persuasion methods, describing the general intent to modify user preference using automated techniques.

Making suggestions is the focus of considerable work in, for example, strategic advice [3, 2] and automated bargaining [12]. However, if an agent only makes suggestions, the perception among users might be of a bargaining situation, with the above noted risk of antagonism [13]. While we also tried to compare our results to agents which only make suggestions, the focus of our work is persuasion using justifications.

Less work has been done on providing justifications or arguments during a user-agent persuasive interaction. A known persuasion method uses a combination of priming [8] as applied to a multi-attribute value based decision making model of a human [11]. One way of determining which justification to present to a user is to choose a justification which will have the desired priming effect. In our work we extend these ideas to automated ways for selecting which attributes to present to the user in order to guide her decision making according to the goals of the agent. User goals may not exist or be otherwise unknown.

Andrews [1], describes a conversational persuasion agent used in a chat-like setting and which analyses a user’s inputs and provides a persuasion message, as part of an argument and including messages tailored to convince the user. Sophisticated dialog management techniques are described, including planning and a layered management system for managing multiple layers of the dialog. The results show that a personalized argumentation system is more effective than a scripted dialog system. However, it is not easy to separate out the effect of tailoring justifications from that of other sophisticated features of the system. In [16, 17] the dialog was a persuasive dialog, but the user was not able to present her own position. Some work, for example that of Huang [17], uses a profiling step before a persuasion step. Kaptein [10] in an on-line sales application, suggested to match the persuasion technique of an agent to a user profile. In our work, a profiling step is essentially omitted. Instead, the user is encouraged to provide her preferences throughout the discussion, effectively integrating Active Learning with persuasion [14].

In an additional contrast to some of the above approaches, we base our work on “construction of preference” (e.g., [18]), so the persuasive interaction is assumed to help in *forming*, rather than only *changing* a user’s preferences. As a consequence, we hypothesize that the process of selecting new attributes to prime the user instead of (or in addition to) only matching to a user profile, should prove to be more effective. This is indeed supported by our results.

3 The Discussion Game

In a “discussion game”, two parties (here an agent and a human user) exchange statements about a situation having multiple possible choice alternatives, with an intention, at least by the agent, that the choices made by the user will have a higher utility for the agent. In our subset of the discussion game, one party (i.e., the agent) has

pre-defined preferences and desires to persuade the other party (i.e., the user), who does not have as strong preferences.




Figure 1. School place map

Specifically, we used a choice making setting, called “school place”, where the agent and user discuss possible locations for a new high school. The discussion was with respect to a real map of a part of the city of Monterrey in Mexico (taken from Google maps). The map has various map icons representing physical features, such as train stations, factories and coffee shops, added to it. A grid indicated by letters and numbers was also added to the map, along the top and left edges (see Figure 1). Future work will attempt to extend this setting to a sales situation, with exchange of money.

We define C to be the set of all possible choices which can be made (e.g., map grid points). For simplicity, it is assumed to be a discrete set with $|C| = n$ alternatives. The discussion relates to independent attributes of the choices A with $|A| = k$ attributes. For each $a_i \in A$ there is a function $\delta_i : C \rightarrow D_i$, which for each choice associates a value within the domain D_i of attribute values possible for a_i . For example, consider the selection of a location for a school. C will consist of all the possible locations for the school; the attributes will include properties such as “near a park”. The possible domain values of the “near a park” feature could be true or false (or other discrete value). The “school place” is also characterized by a low disagreement between users tested as to whether any particular attribute is positive or negative for a school location.

The agent is provided with a specific utility function, $u^a : C \rightarrow \mathbb{R}$, which is not necessarily dependent on the attributes (will be discussed later on).

The user is assumed to *not* have strong pre-determined preferences or pre-defined utility function u^u . However, some assumptions are made regarding the general form of the utility function actually used by the user. First, u^u is assumed to depend on the attributes. Furthermore, it is assumed that for each possible attribute value, there is a consensus (between different users) regarding the type of effect of a particular value on the utility, but no consensus is assumed for the magnitude of such an effect, i.e., for each $a_i \in A$ there is a function $\{\beta_i : D_i \rightarrow \mathbb{R}\}$, which map attribute values to utility values. As a shorthand, $\alpha_i(c) = \beta_i(\delta_i(c))$, provides a score for each choice, $c \in C$ and each attribute value. Each user is assumed to have a set of attributes she uses, A_i^u , which may change as a function of time. The utility function of the user at a given time t is assumed to be of the form $u_i^u(c) = \sum w_i \alpha_i^u(c) + U^u$ with weights w_i being positive

and unknown to the agent and α_i^u being the user's score function, which, as described below, is applied only on attributes in A_t^u . U^u represents parts of the utility function which do not depend on the attributes (e.g., various biases), which parts are not specifically modeled. For example, at an early stage of the discussion game, A_t^u may include only "near park" and "near transportation" and $u_t^u(c)$ may be assumed to be $1 \cdot$ "near park" $+ 1 \cdot$ "near transportation". As a simplifying assumption for this implementation, we assumed that the user's utility function is modified only by addition, that is, the user can add more attributes a_i , but cannot drop any. Future work will be directed at the agent causing attributes to be dropped or downgraded in importance by the user.

In the school place setting we focus on a game in which the choices in C are arranged in a meaningful spatial relationship, agreed-upon by the agent and by the user, for example, a function $XY : C \rightarrow \mathbb{R}^2$ maps each choice to a point on a two dimensional grid (e.g., points A1, B4, F3 on the map of Figure 1).

The discussion uses a language, where R is a set of phrases that the parties can potentially use and which are mapped by a function $\rho : R \rightarrow A$ to the attributes (e.g., the phrase "near park"). The agent uses a subset of the phrases $J \subset R$, as described below.

The protocol of the game is as follows. The game has m rounds, neither agent nor user know m in advance. At each round t , the user proposes a choice $c_t^u \in C$ and a reason $r_t^u \in R$. The agent responds with a choice $c_t^a \in C$ and a justification $r_t^a \in R$ ⁴. In our implementation, the agent can provide up to two phrases, so the agent actually provides two justifications $\{j_t^{(a,1)}, j_t^{(a,2)}\}$. In an idealized example round, a user might say "D4; it is near a playground" and the agent may respond, "I suggest D0; My suggestion is near nature or exercise options. My suggestion is also near transportation." The use of two justification is a tradeoff between keeping the discussion simple and allowing more complex agent strategies.

The goal of the agent can be phrased as maximizing the sum of the utility of the choices made by the user, over the entire discussion (or at least its end), thereby showing that the user is doing the agent's will. So, the agent attempts to maximize: $\sum_{t=1}^{t=m} u^a(c_t^u)$. As the user is assumed to not have an a-priori utility function she is unable, at least initially, to provide such a maximization problem.

Denoting H_t as describing the history up to round t . The agent uses a strategy $\pi : \{H_{(t-1)}, c_t^u, r_t^u\} \rightarrow \{c_t^a, j_t^{(a,1)}, j_t^{(a,2)}\}$, several exemplary strategies will be described below.

4 Agent Design

The language used for agent justification was simplified as follows. The phrases used by the agent include only a subset $J \subset R$, each of which can be mapped to a single attribute $a \in A$, using a many to one function $\tau : R \rightarrow A$, which is a limited implementation of ρ . Only a single justification phrase was defined for each attribute. Agent justifications were slightly modified for agents providing suggested choices (where the phrase is a justification of the choice) and for agents not providing suggestions (where the phrase is a general statement about possible desirable attributes of a choice). Also, in cases where a useful justification could not be generated, where τ failed or where $j_t^{(a,2)} = j_t^{(a,1)}$, general statements on preference were made by the agent or justifications were omitted. User reasons were understood using a classification scheme whereby each user input was recognized using a text recognizer[9] which mapped the reason back to an attribute. The agent did not have other knowledge of the users.

⁴ We reserve the word "reason" for the user and "justification" for the agent

4.1 BEST NEW ATTRIBUTE strategy

The strategy $\pi^{bestnew}$ attempts to persuade the user by priming the user with an attribute the user had not brought up, but which attribute is compatible with an increased agent utility. It is assumed that the user will modify her utility function u_t^u to take this attribute into account. For example, taking the above example of A_t^u , the agent may suggest a choice location, such as "A1" and justify it with the argument "far away from pollution". This is an attribute not raised and therefore assumed more likely to not have been previously considered by the user. This attribute, if now considered by the user, might predispose the user to prefer choice locations that are far from factories.

To do this, the agent generates an estimate of the user's current utility function and then simulates the effect of priming the user with various attributes to determine a most effective attribute to use.

The simulation uses a set of choices $C_t^a = \{c | c \notin H_{(t-1)}\}$. Each choice in the set is scored according to an estimated user utility function $u_t^{u*} : C \rightarrow \mathbb{R}$, so the set is coarsely sorted, into choices a user might find acceptable and choices a user will probably not find acceptable, based on attributes that the user has already indicated as being considered. A set $A_t^u = \{a_i | \exists r_j^u \in H_t \wedge \tau(r_j^u) = a_i\}$ of all the attributes brought up by the user in previous rounds is used to estimate the current utility function and to select new attributes to use. The estimated utility function is taken to be an un-weighted sum taking into account only attributes brought up by the user: $u_t^{u*} = \sum_{\{i | a_i \in A_t^u\}} \alpha_i^u$.

Now, the agent needs to find a new attribute $a_i \notin A_t^u$, not previously brought up by the user, which is compatible with choices that have a high utility for the agent. a_i is selected to be that new attribute which can be expected to increase the utility of the agent by the largest amount, if taken into account by the user. a_i is found as follows. All possible variations of the user's utility function with one attribute added are tried out. The function (attribute) with the highest score is selected and then the agent generates a choice and justifications for that choice which are compatible with the attribute and with the user's previous choices and reasons.

More formally:

1. A set of test functions $u_t^{u*i} = u_t^{u*} + \alpha_i^u, \forall i | a_i \notin A^u$ is generated and applied to C_t^a .
2. Each possible u_t^{u*i} is applied to all the choices in the subset of C_t^a that are acceptable to the user. A score s_i is calculated using this application and is based both on the estimation of a choice desired by the agent being selected and on the value of the choice:

$$s_i = \frac{\sum u_t^{u*i}(c) u^a(c)}{\sum u_t^{u*i}(c)} \quad (1)$$

also note that u^a may be zero for many choices in C_t^a .

3. The a_i which has the highest score s_i is selected, denoted as a_i^* .
4. The choice c_j with highest score for that u_t^{u*i} is presented as a choice c_t^a .
5. $j_t^{(a,1)}$ is selected to be a statement showing that the agent agrees with the selection of an attribute indicated by the user as being important and that the choice by the agent is also in agreement with the user preferences, i.e., $\tau(j_t^{(a,1)}) = \tau(r_{t-1}^u)$.
6. $j_t^{(a,2)}$ is selected from the set $\{j \in J | \tau(j) = a_i^*\}$.
7. Finally, the triplet $\{c_j, j_t^{(a,1)}, j_t^{(a,2)}\}$ is presented to the user.

It should be noted that u_t^{u*i} is not an estimate of what the user is already using as a utility function. Rather, it models the form of a

utility function that the agent is trying to encourage the user to use in the future.

Continuing with the above example, available attributes might be “far away from pollution” and “near food”. If “far away from pollution” is found to have a higher score, this means that it is more likely to predispose a user to select a location beneficial to the agent. Now, the priming needs to be implicitly delivered to the user. The agent will search for a choice which has the attribute “far away from pollution” with a positive value and which does not conflict with the user’s previously professed attributes, offer that choice, and give a justification “this location is far from noise and/or pollution”. If the priming is successful, the user will limit her attention to only those locations which are also far from pollution.

4.2 NOOFFER and NOREASON variants

As a discussion includes both examples and justification, we also provided variations of the above priming strategy, to assist in analyzing the effectiveness of the strategy. Two variations on the $\pi^{bestnew}$ strategy were tried out. In a first one, $\pi^{nooffer}$, the strategy is applied as above, but the computer only presents the first and second justifications $j_t^{(a,1)}$, $j_t^{(a,2)}$ and does not present the choice c_t^a .

In the $\pi^{noreason}$ strategy, the converse is true. The suggestion c_t^a is presented, but neither the first nor the second justifications $j_t^{(a,1)}$, $j_t^{(a,2)}$ are presented.

4.3 SPATIAL strategy

As a baseline, a strategy which tries to maximize agent utility without considering justifications was built. This is a stand-in for bargaining scenarios and is not the focus of this work. For simplicity and ease of comparison a simple greedy strategy was used. During experimentation it appeared that a tested strategy had an effect due to the spatial arrangement of choices that it provided and the strategy was improved to become $\pi^{spatial}$. This strategy ignores the history and the user’s contributions $\{H_{t-1}, c_t^u, r_t^u\}$ and provides an output c_t^a based only on the utility u^a and on a desired spatial clustering which matches u^a . No justifications $j_t^{(a,1)}$, $j_t^{(a,2)}$ are provided.

The desired clustering is defined by identifying large spatial clusters in C , which share a high u^a . Clustering proceeds in two steps. First, high value choices are identified using a threshold. Then, as simple cluster shapes are desired, the clustering method identifies simple shapes that have an average high value for choices within the cluster. The shapes are identified using a lower threshold to help compensate for the generally unavoidable inclusion of some low-value choices when only simple cluster shapes are allowed. In the first step, using a given threshold value T_1 , which represents the value above which a choice is considered interesting, a 2D map $M = \{c \in C | u^a(c) > T_1\}$ is created, with the choices arranged on a grid using the function XY, described above. Also given is an allowed shape for the clusters, for example, square. Generally, simple cluster shapes are desired, so that the human will be able to easily guess the general boundaries of the cluster using a small number of points in the cluster.

In a second step, define Q as the set of all clusters which when bounded by the maximally sized shape, have an average utility for all choices bounded by the shape, above a second threshold T_2 . Typically T_2 is derived from T_1 and represents a trade-off between the size of a cluster and its allowed irregularity. Desirably, but not necessarily, a single $q \in Q$ is chosen.

At turn t , the agent chooses $\{c_t | c \in q, c \notin H_{t-1}, u^a(c) > T_1\}$. This serves as an example to the user, to make choices from the cluster. The agent uses a random value function to generate grid coordinates, which identify a choice c_t in the cluster, with the caveat that c_t cannot have been previously suggested by agent or user. In our implementation, the $\pi^{spatial}$ agent simply provided locations in either one of the top left or bottom right corner of the map (in areas of size 4×4), depending on which agent utility function was used.

4.4 SILENT strategy

As an additional baseline, another strategy, π^{silent} , was tested. An agent playing with that strategy makes no suggestion and gives no justifications. Rather the subject was repeatedly asked for alternative choices and for reasons for these choices. Thus, this strategy does not affect the user construction of preference.

5 Empirical Methodology and Design

We implemented the above strategies as agents which held “conversations” with human subjects in an attempt to persuade the subjects to make choices commensurate with an agent’s utility functions.

The same school place setting was run with multiple subjects and both the agent strategy and the agent utility function were varied. Between 19 and 35 subjects were used for each combination of agent strategy ($\pi^{spatial}$, $\pi^{bestnew}$, $\pi^{nooffer}$, and $\pi^{noreason}$) and agent utility function (u_{left}^a , and u_{right}^a) (see Figure 2, below). An additional strategy π^{silent} was neutral with respect to the utility function.

Flow of an interaction: Each human subject (we used workers from Amazon Mechanical Turk) has a structured series of exchanges with the agent, following the protocol described above. In a first training round, the subject is presented with a first setting in which only four locations are available $C = \{A, B, C, D\}$ (roughly one for each quadrant of the map of Figure 1, with the grid hidden) and is asked to provide a reason r_1^u . Limiting the set of choices for this first round was aimed at helping familiarize users with the settings and causing them to build and state at least some preferences. The results analyzed below exclude this first round.

Then, in the following rounds, the map is shown with a positioning grid, so $C = \{(x, y) | x \in 0 \dots 9, y \in A \dots I\}$. The agent provides a suggestion and/or justifications and the subject is asked to provide alternative grid locations for a school. After five such exchanges, the subject is thanked and does not participate again. The subject is not allowed to repeat suggestions made by the user or the agent, to ensure that a maximal indication of the user preferences are elicited.

Programming the text recognition and attributes: The text recognition system was programmed by first collecting hundreds of sentences from subjects (that were not used later in the evaluation phase), where the subjects were asked to justify their own suggestions and also being asked to explain why a computer-suggested location is good or bad. Then, these reasons were cleaned up and manually classified into 13 different attributes. A textual recognizer [9] was programmed with these examples, so that sentences using similar keywords would be classified in a similar manner.

Programming the attributes: The map was manually analyzed and each point on the grid was classified (by an author) with respect to each of the attributes, specifically, 0 indicating neutrality with respect to an attribute, -1, negative utility and +1, positive utility. As noted there was general agreement between subjects with regard to utility of attributes.

Motivating the user to take a stand: The discussion protocol was designed to avoid trivial interaction results. First, the subject moves before the agent, i.e., the subject is required to take a stand, including both a location suggestion and a reason for such a suggestion, before the agent provides any suggestions or justifications. It is believed that people follow a consistency heuristic (see [6], page 50), which causes them to keep to a previous made decision even if it is sub optimal. This should give the agent leeway to apply persuasion methods, rather than merely have the subject form an opinion based on whatever the agent outputted. Second, a subject was not allowed to repeat a previous suggestion or to accept a suggestion made by the agent: $c_t^u, c_t^a \notin H_{t-1}$. This forces a choice to not be a trivial acceptance of a previous choice and is believed to encourage subjects in the experiments below to make actual choices.

Agent utility functions: The agent utility function was not specifically related to the attributes and defined a preference over each grid point on the map, with the values 0, 1 or 2. The subjects were not aware of these functions. Because the user is not allowed to accept an agent offer as is, but only be inspired by the offer, we used agent utility functions which have low spatial variation. Two alternative and complementary utility functions were generated for the agent: u_{right}^a , a utility function with a preference for the right and especially the top right and a utility function u_{left}^a , with a preference for the left and especially the bottom left, in Figure 2. The final form of the function was chosen to allow similar alternative functions to be tried (e.g., mirrored), while ruling out effects due to the shape of the function. As it turned out, u_{left}^a and u_{right}^a are not equivalent with respect to average user preference.



Figure 2. Utility functions u_{left}^a, u_{right}^a

6 Results and Discussion

We now present and analyze the results obtained in the experiments. We also present several extensions that were conducted to explore specific hypotheses. Statistical significance tests were used for all statistical comparisons, which were used to compare pairs of strategies, as discussed below.

6.1 Summary of Results

Considering that each interaction instance included five suggestions by the subject, the score for each such instance is between 0 and 10. A score based on random choices for the above utility functions would be 3. Table 1 presents the average agent’s score and the number of subjects in each experimental condition:

The strategies should be considered in view of the non-neutral baseline condition shown by π^{silent} . When the subject is just asked for a choice suggestion and justification, with no agent counter-suggestion or justification, there is a marked preference for the choices compatible with u_{right}^a . The data shows that suggestions and justification each have an effect on the subject, but that this effect seems to depend on this baseline inclination. In any case, providing

Table 1. results for strategies

Team	$u_{\{left right\}}^a$	num subjects	score	std dev
$\pi^{spatial}$	R	30	3.73	0.30
$\pi^{bestnew}$	R	26	3.04	0.35
$\pi^{nooffer}$	R	25	4.04	0.33
$\pi^{noreason}$	R	35	3.66	0.25
π^{silent}	R	25	3.96	0.31
$\pi^{spatial}$	L	27	4.19	0.42
$\pi^{bestnew}$	L	25	3.00	0.38
$\pi^{nooffer}$	L	25	2.72	0.45
$\pi^{noreason}$	L	27	3.15	0.43
π^{silent}	L	25	2.2	0.41

both a suggestion and a justification is not more effective than providing only one of them.

It is interesting to note that the $\pi^{spatial}$ strategy has a greater effect on subject choice than that of $\pi^{bestnew}$ ($p < 0.01$), suggesting that spatially focused examples are more effective than suggestions with justifications. As discussed above, the $\pi^{spatial}$ strategy is a stand-in for bargaining, which was not the focus of this work given the observation that people tend to be more satisfied with their choices if they feel that they made the choice on their own (even over time), and that bargaining may lead to antagonism. However, $\pi^{spatial}$ can serve as a baseline for comparison.

6.2 Comparison of Different Strategies

Considering first the results for the u_{right}^a cases. It appears that when a subject has an inclination, the best way to bring it out is not to provide any computer input (π^{silent}) or reinforce justification based thinking by presenting justifications (i.e., $\pi^{nooffer}$).

The next best way to affect a subject is to give an example of desired behavior ($\pi^{spatial}$, and to a lesser extent $\pi^{noreason}$), ($p < 0.01$ when compared to $\pi^{bestnew}$).

Finally, comes the option of “standard” discussion, where both a suggestion and a justification are provided. It is possible that the subjects were inclined to be contrary when they felt they were in a negotiation-like situation. However, this type of effect should also have been exhibited in the $\pi^{spatial}$ strategy, but was not. It is noted again that all strategies were either as good as being silent or worse than being silent, in this case.

Referring now to the u_{left}^a results. The differences between the $\pi^{bestnew}$, $\pi^{nooffer}$ and $\pi^{noreason}$ strategies are small and/or statistically insignificant. However, all did better than merely allowing the subject to follow her inclination. It is hypothesized that engaging the subject caused him (or her) to consider their position. The effect of examples on user choice is especially marked. In any case, persuasion was possible in either utility function.

It is hypothesized that an important effect on subject behavior is the priming of the subject to use attribute based thinking rather than some other choice heuristic. To test this hypothesis, two more sets of experiments were run, one with a strategy π^{random} , where the agent makes suggestions uniformly randomly distributed on the map grid and another with a strategy $\pi^{reasonrandom}$, where the agent also gave a justification for the suggestion, based on the positive attributes that a randomly selected grid point happened to have. Results are shown in table 2, below.

At least for the u_{right}^a case, if suggestions are made, then providing a justification, even if unrelated to the agent’s utility function, caused an increase in score that is significant ($p < 0.01$). It is hypothesized that the π^{random} strategy encouraged subjects to select

Table 2. results for RANDOM strategies

Team	$u_{\{left right\}}^a$	num subjects	score	std dev
π^{random}	R	19	3.16	0.41
$\pi^{reasonrandom}$	R	35	3.80	0.24
π^{random}	L	19	2.95	0.50
$\pi^{reasonrandom}$	L	35	2.63	0.27

school locations without thinking about their attributes, while the $\pi^{reasonrandom}$ strategy caused subjects to consider attributes and thus be more likely to align with the u_{right}^a utility function.

With regard to the u_{left}^a case, it is noted that both conditions were better than just leaving the subject alone, possibly, with the better results for π^{random} being explained by the subject not being encouraged to consider attributes. An alternative explanation is that random choices encourage the user to follow suite and provide random suggestions, which give a score close to the expected random score (3).

It is important to note that the pre-test results (not detailed here) showed a preference for u_{left}^a , rather than for u_{right}^a . However, these initial results are on the 4-option forced choice setting. The provision of freedom to choose any grid point, as in rounds 2 through 6, seems to have revealed an underlying preference for u_{right}^a . It is hypothesized that the seeming change in preference from u_{left}^a to u_{right}^a is due to the explicit elicitation and recitation of a reason by the subject, which caused the subject to think in terms of attributes.

Another interesting question is the effect of justifications on the subject. In an attempt to shed some light, an additional strategy was tested, in which the agent only expresses an agreement with the subject's reasoning, by providing a justification for the attribute previously provided by the subject, without the agent also providing an additional attribute. This was called $\pi^{bestnow}$. Table 3 compares two NOOFFER conditions (where no actual suggestion was made, only justifications⁵):

Table 3. results for NOOFFER strategies

Team	$u_{\{left right\}}^a$	num subjects	score	std dev
$\pi^{bestnow}$	R	26	3.23	0.31
$\pi^{bestnow}$	L	25	4.04	0.33

Comparing the results of Table 3 with those of Table 1 it seems that the fact that new attributes were selected to increase the computer utility, has a significant effect on subject choice ($p < 0.01$). This can be contrasted to [4] which compared the use of several justification to the use of all possible justifications.

In general, no position effects were identified. That is, the subject's behavior (e.g., agent score) does not appear change between the first exchange and subsequent exchanges. This may be because the subject's behavior is set by the first agent suggestion/justification or because of a high variability within subject behavior.

An initial analysis to see which strategy reduced the average Euclidean distance between the agent's suggestion and the following subject's suggestion yielded the information that in $\pi^{spatial}$, the smallest average distance (3.70) is found and in $\pi^{bestnow}$, the largest average distance (5.20) is found, with the other strategies having a distance of about 4.20. This is compatible with the explanation that $\pi^{spatial}$ caused subjects to emulate the example set by the agent.

⁵ In a condition where $\pi^{bestnow}$ also included an offer compatible with A_t^u , there was no significant difference in score between $\pi^{bestnow}$ and $\pi^{bestnow}$.

7 Conclusions and Future Work

We showed that priming can be applied automatically by an agent, even without a pre-profiling step and provide meaningful persuasion. We also showed that suggestions can be better than justifications and using one is generally better than both. The mere usage of justifications also affected the way a user made choices. Finally, a goal of this research was and remains providing a domain independent persuasion strategy which can be easily transferred between domains and using simple NLP techniques. As can be seen from the description of the experiments, only a relatively simple procedure is required to reprogram the agent strategy to be used in a different setting. We plan to test this hypothesis, as well as define a useful setting which does not require graphical interaction or explicit spatial arrangement.

ACKNOWLEDGEMENTS

This work is supported in part by ERC grant number 267523 and MURI grant number W911NF-08-1-0144. Special thanks to Erel Segal of Bar-Ilan's NLP laboratory for building the NLP system.

REFERENCES

- [1] P.Y. Andrews, *Persuasive Computer Dialogue Improving Human-Computer Communication*, Ph.D. dissertation, 2008.
- [2] A. Azaria, Z. Rabinovich, S. Kraus, and C. V. Goldman, 'Strategic information disclosure to people with multiple alternatives', in *AAAI*, (2011).
- [3] A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and Y. Gal, 'Strategic advice provision in repeated human-agent interactions', in *AAAI*, (2012).
- [4] G. Carenini and J. D. Moore, 'Generating and evaluating evaluative arguments', *Artif. Intell.*, **170**(11), 925–952, (2006).
- [5] E. Chan and A. Mukhopadhyay, 'When choosing makes a good thing better: Temporal variations in the valuation of hedonic consumption', *Marketing Research*, **47**(3), 497–507, (2010).
- [6] R. Cialdini, *Influence: Science and practice (3rd edi)*, HarperCollins, New York, 1993.
- [7] B. J. Fogg, *Persuasive Technology: Using Computers to Change what we Think and Do*, Morgan Kaufmann, Amsterdam, 2003.
- [8] G. Häubl and K.B. Murray, 'Recommending or persuading?: the impact of a shopping agent's algorithm on user behavior', in *EC '01*, pp. 163–170, New York, NY, USA, (2001). ACM.
- [9] T. Joachims, 'A support vector method for multivariate performance measures', in *ICML '05*, pp. 377–384, (2005).
- [10] M. Kaptein, 'Adaptive persuasive messages in an e-commerce setting: the use of persuasion profiles.', in *ECIS*, eds., Virpi Kristiina Tuunainen, Matti Rossi, and Joe Nandhakumar, (2011).
- [11] R.L. Keeney and H. Raiffa, *Decisions with multiple objectives: Preferences and value tradeoffs*, Cambridge Univ Pr, 1993.
- [12] R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry, 'Negotiating with bounded rational agents in environments with incomplete information using an automated agent', *Artif. Intell.*, **172**(6-7), 823–851, (2008).
- [13] Y. Maaravi, Y. Ganzach, and A. Pazy, 'Negotiation as a form of persuasion: arguments in first offers.', *Personality and Social Psychology*, **101**(2), 245–255, (2011).
- [14] R. Neil, D. Kaplan, and M. Sugiyama, 'Active learning in recommender systems', in *Recommender Systems Handbook*, eds., P.B. Kantor, F. Ricci, L. Rokach, and B. Shapira, 735–767, Springer, (2011).
- [15] D. Sarné, A. Elmalech, B. Grosz, and M. Geva, 'Less is more: restructuring decisions to improve agent search', in *AAMAS*, pp. 431–438, (2011).
- [16] D. Schulman and T. Bickmore, 'Persuading users through counseling dialogue with a conversational agent', in *Persuasive '09*, pp. 25:1–25:8, New York, NY, USA, (2009). ACM.
- [17] H. Shiu-Li, L. Fu-Ren, and Y. Yufei, 'Understanding agent-based online persuasion and bargaining strategies: An empirical study', *Int. J. Electron. Commerce*, **11**, 85–115, (October 2006).
- [18] P. Slovic, 'The construction of preference', *American Psychologist*, **50**(5), 364–371, (May 1995).