A Study of User's Online Decision Making Behavior

Jiyong Zhang[1], Pearl Pu[1] and Paolo Viappianni[2]

[1] Human Computer Interaction Group,
Ecole Polytechnique Fédérale de Lausanne (EPFL),
CH-1015, Switzerland
{jiyong.zhang, pearl.pu}@epfl.ch
[2] Artificial Intelligent Laboratory,
Ecole Polytechnique Fédérale de Lausanne (EPFL),
CH-1015, Switzerland
paolo.viappianni@epfl.ch

Abstract. Understanding user's online decision making behavior is a crucial issue in developing consumer supportive e-commerce systems because if we can model user's online decision behavior precisely, the e-commerce system can be designed and developed readily satisfying end-users' various requirements and the time and the expense for developing the system can be greatly saved. The target of a consumer decision support system is to help users find high accurate choices within a low level of interaction effort and cognitive effort. In this report we study a group of subjects' decision behavior based on the FlatFinder application. We analyze user's online decision behavior in several aspects such as the preferences input, interaction cycles, and the frequency of stating preferences on each attribute. The results provide evidences to show that users have bounded rationality in stating preferences and have adaptive decision behavior during the interaction process.
keywords:online decision behavior, example-critiquing interaction, decision theory

1 Introduction

As the number of consumers who purchase online is growing sharply during the last decade, it becomes increasingly important to develop web-based product search systems to cater for online consumers' various requirements and to help them find the desired products efficiently. To do that, we must learn how consumers act in processing product information and making purchase decisions in online situations. Understanding user's online decision making behavior is a crucial issue in developing consumer supportive e-commerce systems and will be great helpful for the growing e-commerce market.

Traditionally, An individual directly tackles a decision problem without any decision aid. In earlier days people assume that an individual has comprehensive domain knowledge, a well-organized and stable set of preferences, and unlimited computational capability to solve the faced decision problem. Based on this "rational decision maker" assumption, a group of researchers had established the utility theory for making rational decisions[1] [2]. In 1950s Herbert Simon[3] questioned this assumption and argued that the decision maker only has bounded rationality. He pointed out that if the decision problem is in a high level of complexity, the decision maker is not likely to solve the decision problem by calculating the exact utility of each choice. Instead, she tends to adopt a heuristic strategy to determine the choice by defining a certain level of satisfying degree for each attribute. Only those products that each attribute has a better value than the level of the satisfying degree is to be considered as the final choice. According to Payne et al [4], individuals are often impressively adaptive in their responses to different decision situations. For example, if the problem is very simple, she may adopt an accurate decision strategy such as weighted-additive (WADD) strategy to find the desired choice, otherwise the user is likely to use some easier heuristic strategies to determine the choice to save the decision effort. According to their work, this adaptive decision behavior in offline situations can be described by an accuracy-effort framework.

While the above researches mainly from psychologies and economics have been carried out to describe the decision maker's behavior in traditional decision situations, so far it is little known about user's decision making behavior in online situations. It is obvious that user's online decision behavior is quite different to that in traditional situations because of the following observations. Firstly, in an online situation the user is no longer required to evaluate all the products one by one to make a decision. Instead, typically a pre-designed program in the computer (called decision support system) can help her to search the possible products according to her preferences. For example, if a user wants to buy a digital camera, she may input her preferred values on some attributes such as brand, picture size, and price into the system and expect the system to recommend the desired products. Secondly, in online situations the user needs to input her preferences into the computer system so to receive decision aids. Thus a certain level of effort for interacting with the computer system has to be provided. Finally, in traditional situations, the decision process is a binary relationship between the decision maker and the decision problem. But in online situations, the decision process is a tri-nary relationship among the decision maker, the decision support system and the decision problem. It's quite certain that these differences in the decision environment have important impact to users and their decision making behaviors will be altered consequently.

In some literatures a system providing decision aids to end-users is also called a decision aid agent or a recommender system. In this report we generally call it as a decision support system. We also deem the end-users of a decision support system as individuals, consumers, decision makers or simply users without difference. In this report we intend to deeply study end-users' online decision behaviors based on a set of real user studies in the apartment finder domain. We report how users act when they input their preferences, and the way they interact with the system during the interaction process. We provide evidences to show user's adaptive decision behavior in such online situation. In the rest of this paper, we first introduce some earlier work related to the study of user's decision behavior. Then the information about our experiment setup is described. Next we report the results of user's online decision behavior learned from the experiment. Finally we give conclusions and the future work.

2 Related Work

In traditional environments where no computer aid is involved, behavioral decision theory has provided adequate knowledge describing people's choice behavior and the approaches for solving decision problems. In 1990s Payne et al. [4] established a well known effort-accuracy framework describing how people adapt different decision strategies by trading off accuracy and cognitive effort to the demands of the tasks they face. Especially, they studied the following decision strategies:

- The weighted additive (WADD) Strategy. It considers the values of each outcome on all of the relevant attributes and all of the relative importance (weights or probabilities) of the different attributes to the decision maker. Each outcome is given an evaluation value by multiplying the weight and the attribute value for each attribute and summing these weighted attribute values over all attributes. The outcome with the highest overall evaluation value is chosen as the optimal solution. Actually, the WADD decision strategy is a special case of the normative approach based on Multi-Attribute Utility Theory (MAUT) [2] with the additive value function.
- Some basic heuristic strategies. They are the equal weight (EQW) strategy, the elimination-by-aspects (EBA) strategy, the majority of confirming dimensions (MCD) strategy, the satisficing (SAT) strategy, the lexicographic (LEX) strategy and the frequency of good and bad features (FRQ) strategy. Their detailed definitions can be found in [4] and [5].
- Hybrid strategies. Besides the basic heuristic strategies, people may also use a combination of several of them to make a decision to try to get a more precise decision result. These kinds of strategies are called hybrid decision strategies. For example, The elimination-by-aspects plus weighted additive (EBA+WADD) strategy uses an EBA process until the number of available outcomes remaining was three or fewer, and then used a weighted additive strategy to select among the remaining outcomes. The elimination-by-aspects plus majority of confirming dimensions (EBA+MCD) strategy first uses the elimination-by-aspects process to reduce the problem size, and then uses a majority of confirming dimensions heuristic to select the optimal outcome from the reduced set.

In online decision making environments, the way to obtain users' preferences during the interaction process is a fundamental issue for the system design. Pu. et al [6] pointed out the following principles of the preference elicitation based on the study of the decision behavior theory[4]:

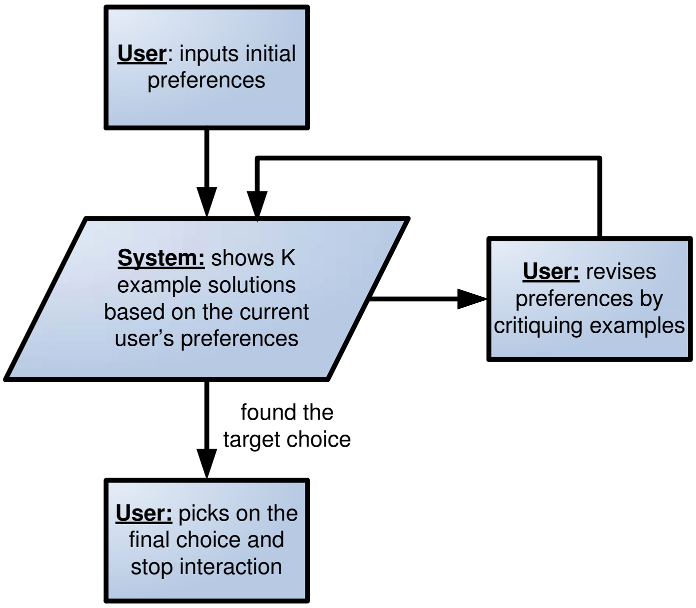

Fig. 1. The example-critiquing interaction paradigm.

– Users are not aware of all preferences until they see them violated. For example, a user does not think of stating a preference for intermediate airport until a solution includes a change of airplane in a place that he dislikes. This can not be supported by the decision tool that requires preferences to be stated in a predefined order.
– Elicitation questions that do not concern the user's true objective can force him to formulate means objectives corresponding to the question. For example, in a travel planning system suppose that the user's objective is to be at his destination at 15:00, but that the tool asks him about the desired departure time. The user might believe that the trip necessarily involves a plane change and take about 5 hours, and thus forms a means objective to depart at 10:00 to answer the question. However, the best option might be a new direct flight that leaves at 12:30 and gets there at 14:30. This solution would not be found using the elicited preference model. This phenomenon has been studied by Keeney[7] in his work on value-focused thinking.
– Preferences are often in contradiction and require users to make tradeoffs, which require users to add, remove or change preferences initiatively in any order at any time.

According to these studies, we can see that the system must provide instant feedback to users to indicate the results that they can obtain with the current preferences. A good way to implement such feedback is to implement a mixed-initiative decision support system with the *example-critiquing* interaction paradigm as shown in Figure 1. In each interaction cycle the system displays several examples of complete solutions and invites users to state their critiques of these examples. Example critiquing allows users to better understand the impact of their preferences. Moreover, it provides an easy way for the user to add or revise his or her preferences at any time in any order during the decision making process. Example-critiquing as an interface paradigm has been proposed by a variety of researchers[8] [9] [10].

Some researchers had carried out some experiments with artificial users in evaluating the performance of their systems or approaches. Payne et al.[4] introduced a simulation experiment to measure the performance of various decision strategies in offline situations. Recently, Boutilier et al. [11] carried out their experiments by simulating a number of randomly generated synthetic problems and user responses to evaluate the performance of various query strategies for eliciting bounds of the parameters of utility functions. In [12] and [13], various users' queries were generated artificially from a set of offline data to analyze the recommendation performance of the incremental critiquing approach. These related work generally assume that the user has a rather simple-minded decision behavior. For example, in [13], an artificial user was always assumed to select the one that is closest to the target product during the interaction process. In [14], artificial users were assumed to input 3 initial preferences in average into the system. It is an improvement in carrying out simulation experiments by integrating user's actual decision behavior, but it is still based on a quite simple and intuitive behavior model.

Despite of all the above work, so far it is lack of a detail study of user's online decision behavior. In this report we intend to fill this gap by studying deeply user's online decision behavior on a experimental application built on the apartment finder domain.

3 Experimental Setup

We performed our user studies using FlatFinder, a web application for students to find apartments offered from a university database [15]. Each apartment consists of 10 attributes: the type of accommodation (room in a family house, room in a shared apartment, studio apartment, apartment), the rental price, the number of rooms, furnished (yes or no), the bathroom (private or shared), the type of kitchen (shared, private), the transportation available (none, bus, subway, commuter train), the distance to the university and the distance to the city center. There are 180 different apartments for users to select during the experiment.

The system is implemented based on the example-critiquing diagram introduced earlier. The user states a set of initial preferences and then obtains a set of recommended choices by pressing the search button. Subsequently, she goes through a sequence of interaction cycles where she could refine her preferences by critiquing the displayed examples. During each interaction cycle the system maintains her current set of preferences and she could state additional prefer-

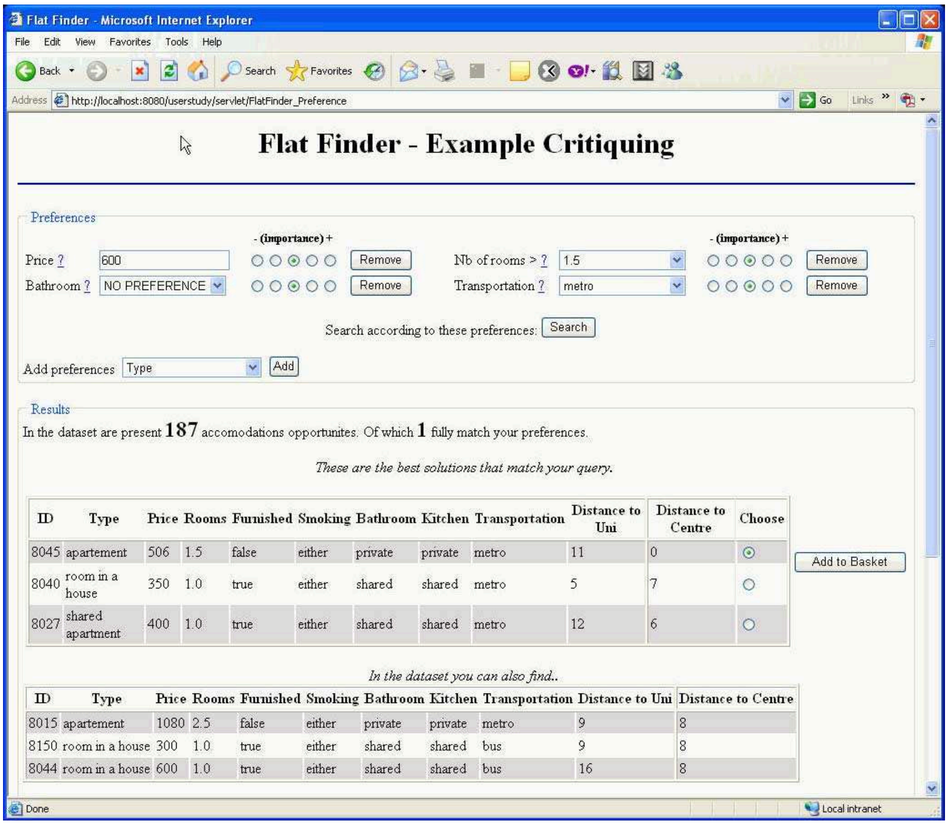

Fig. 2. User interface of the flat finder system.

ences, change the reference value of existing preferences, or even remove one or more existing preferences. Finally the process is finished with the user's final set of preferences, and a target choice chosen by the user from the displayed examples. Figure 2 shows the interface of the apartment search system developed in our study.

We recruited 40 (9 females) subjects of 9 different nationalities, mostly undergraduate students to participate the experiment. Most of them (27 out of 40) had searched for an apartment in the area before and 26 out of 40 had used online tools to look for accommodations. Importantly, all subjects were motivated by the interest of finding a better apartment for themselves. Each subject tried the search procedure twice with different strategies of showing the examples: one strategy is to show 6 similar products to the user, the other strategy is to show 3 similar products and another 3 diverse products as suggestions. The detail implementation of these two strategies can be found in [15]. In this study we don't distinguish between those two strategies when we learn user's decision behavior. Together there are 80 cases of interaction procedures in our study. After the

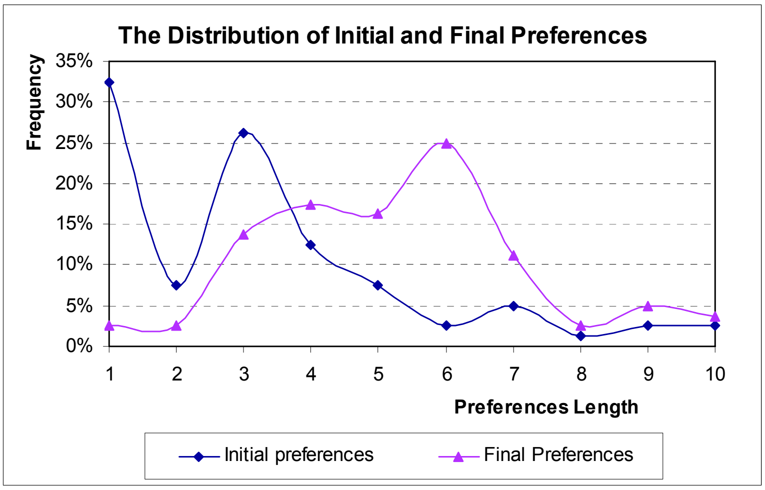

Fig. 3. Distribution of the initial and final preferences.

online experiment finished, we also asked each subject to select the most desired one by browsing the whole apartment list.

4 Results and Analysis

In this section we report our study of the user's online decision behavior. We focus on the following information: the initial and final preferences, the length of the interaction cycle and the distribution of preferences on various attributes.

4.1 Initial and final preferences

Understanding the initial and final preferences is an important part of knowing user's decision behavior. Figure 3 shows the distribution of initial and final preferences in our experiment. It shows that even though the system has no limitation to users' initial preferences, users are only willing to input around 3 (average length: 3.21)initial preferences into the system. Specifically, users would like to input either 1 or 3 initial preferences, rather than 2 initial preferences. For the final preferences, the figure shows that the distribution of the final preferences mainly lies in 4-7 preferences. The figure also shows users would like to state around 6 (average length: 5.31) final preferences when stopping the interaction process.

We also measured the number of initial preferences violated by the apartment selected by users at the end of the interaction process. In our experiment there are 40 sessions (out of the 80 sessions in total in our experiment) that at least one attribute of the initial preferences had been violated. In other word, there is only 50% of the sessions that users had stated initial preferences which are in

Table 1. The compromised and respected preferences

	Initial preferences		Final preferences	
	respected	compromised	respected	compromised
average	2.74	0.48	4.50	0.81
variance	4.25	0.38	3.04	0.89

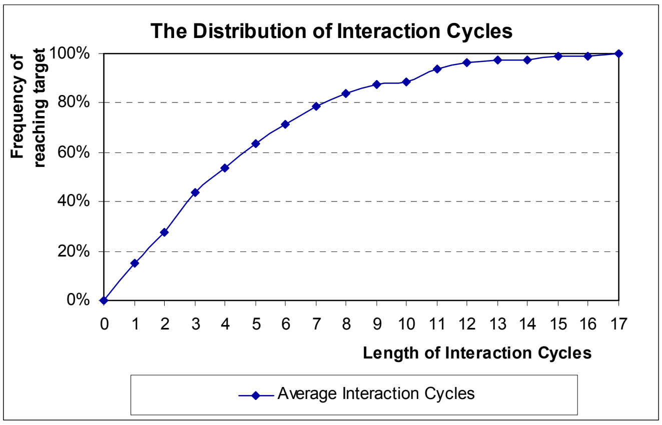

Fig. 4. Distribution of the interaction cycles.

line with the respective final choices. This shows that initial preferences are not very accurate and users have the tendency to change the preferences adaptively during the interaction process.

A more detailed investigation is to divide each set of the initial and final preferences into two parts: respected values and compromised values. If the value of an attribute in the preferences is consistent with the final choice, then we say it is *respected*; otherwise we say it is *compromised*. Table 1 shows the difference between the initial preferences and the final preferences in terms of the number of compromised and respected values. We can see that in the set of initial preferences there are 2.74 values in average which have been respected, while in the set of final preferences, there are 4.50 values in average that are respected. This result shows that users will get more respected values at the end of the interaction process than those at the beginning. For the compromised category, users have compromised 0.48 values in initial preferences in average, but still have compromised 0.81 values in final preferences in average. This shows that there are still some inconsistencies between the final preferences and the final choices.

4.2 Distribution of interaction cycles

Interaction cycle can be regarded as a unit of measuring the interaction effort. Here we try to find out in average how many cycles will be carried out during the decision process. The result is shown in figure 4.

This figure shows that users tend to finish the product search process within only a few interaction cycles. For example, when the length of interaction cycles is limited to 10, about 90% percent of interaction sessions have reached the target choices and the interaction process has been stopped. This result shows that users seldom interact with the system more than 10 cycles. The frequency of reaching the target product can be seen as a criteria for measuring the decision accuracy. This figure shows that more interaction effort indeed can help users get more accurate decisions, However, if the interaction effort exceeds a certain threshold, more interaction effort can only contribute little to the growth of decision accuracy.

We further divide the interaction process into 3 groups according to the interaction length: the group with interaction length between 1 to 5, the group with interaction length between 6 to 10, and group with interaction length longer than 10. For each group, we compare the online choice by the interaction process with the final choice selected after the experiment. Here we define the accuracy as the percentage that the online choice equals to the user's final choice. The statistic result is shown in figure 5. We can see that all groups of interactions have similar accuracy. Though the group with interaction length longer than 10 has much more interaction effort than the group with interaction length 1-5, there is no signification improvement of the choice accuracy. This result shows that longer interaction length is not necessarily for higher choice accuracy.

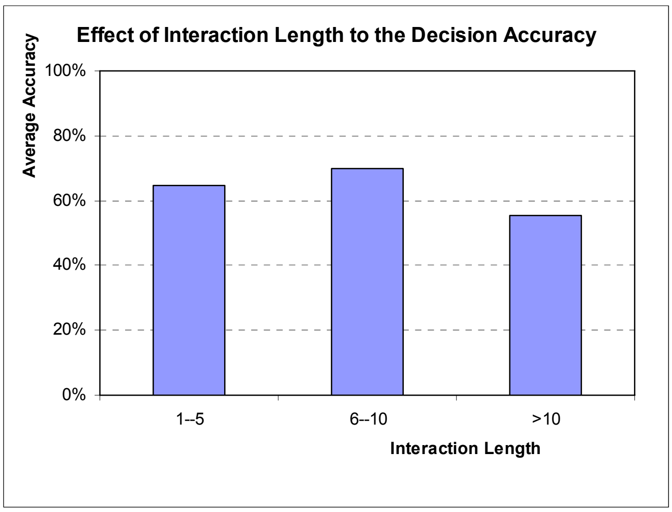

Fig. 5. Interaction length to the choice accuracy.

4.3 The frequency of users' preferences on each attribute

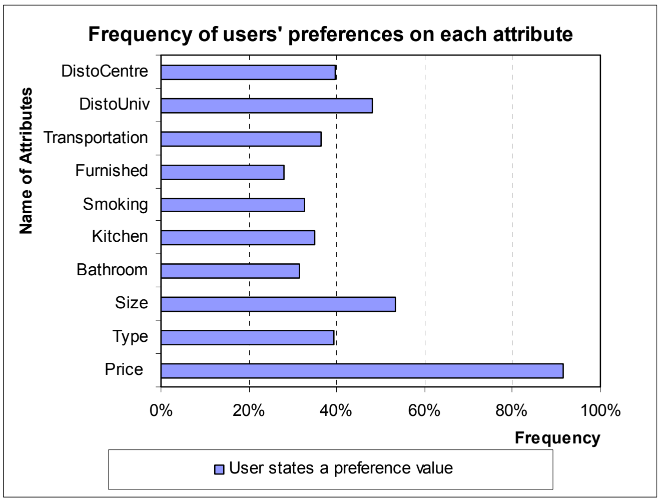

Fig. 6. The frequency of users' preferences on each attribute.

During the experiments we observed that quite often users input only part of their preferences into the system. Here we measured the result of the frequency of the user's preference on each attribute, which is shown in figure 6. The result shows users are more likely to specify preference on *Price* attribute than other attributes. In 92% cases users have specified their preferences on *Price*, while other attributes are in a much lower rate. This result generally shows that *Price* is more important than other preferences from users' point of view and suggests that the decision support system should allow users specify preference on *Price* easily. Besides *Price*, *Size* and *DistoUniv* are the next two attributes that users like to state preference values.

We also calculated the co-occurrence frequency of two and more attributes that users are likely to input preference values. As shown in table 2, The *Price* attribute is frequently appeared together with other attributes. Especially, in 51% of all the sessions users have stated attribute *Price* and *Size* together,and in 45% of all the sessions users have stated *Price* and *DistoUniv* together. Besides the attribute*Price*, <Type, Size> and <Size, DistoUniv> are two attribute pairs frequently appeared together(co-occurrence frequency: 25%). There are also a few patterns of three attributes that users are likely to state together. For example, in 25% of all the sessions users have stated <Price, Type, Size> together.

Figure 7 shows a distribution of users' different input values on the *Price* attribute. It shows that the price values they specified are concentrated around *800CHF*. The distribution of all the values on the *Price* attribute is likely to be a Gaussian distribution.

Table 2. The co-occurrence frequency of two or more attributes

Pattern	**Frequency**
Co-occurrence of 2 attributes:	
⟨Price,Size⟩	51%
⟨Price,DistoUniv⟩	45%
⟨Price,Transportation⟩	36%
⟨Price,DistoCentre⟩	36%
⟨Price,Type⟩	35%
⟨Price,Kitchen⟩	32%
⟨Price,Smoking⟩	32%
⟨Price,Bathroom⟩	29%
⟨Price,Furnished⟩	26%
⟨Type,Size⟩	25%
⟨Size,DistoUniv⟩	25%
Co-occurrence of 3 attributes:	
⟨Price,Type,Size⟩	25%
⟨Price,Size,DistoUniv⟩	25%

5 Conclusions and Future Work

In this report, we studied user's online decision behavior in terms of preferences input, interaction cycles, and the frequency of preferences on each attribute. Our results show that users have bounded rationality and erratic in stating preferences into the system. For example, at the beginning the user will only input a few (around 3) initial preferences into the system, and then she will interact with the system only a few cycles (typically less than 10) before stopping the interaction process. Our results are in line with the earlier researches that users have adaptive decision behavior.

According to [5], in online environments users mainly concerns the following three factors during the decision process: the decision accuracy, the interactive effort and the cognitive effort. These three factors are mutually related each other. For example, if a user wants a more accurate decision, she can spend some more interaction effort with the system. In this work, we provide evidences to show that a decision maker has a certain threshold for the interaction effort, and it is not expected that she can spend more interaction effort beyond this threshold to achieve a higher decision accuracy. And due to the adaptive nature of the decision maker, the decision support system is expected to allow users to input their preferences as they wish to save the cognitive effort and interactive effort.

In the IT community, currently many people still tend to regard that users can always make rational choices. For example, some websites still show a lot of information at one page to users, expecting they can reveal all preferences correctly. Also, some systems are designed in such a way that users have to answer some tough questions. This research provides some guidance for system designers to know what exactly the system should be designed. If we can model

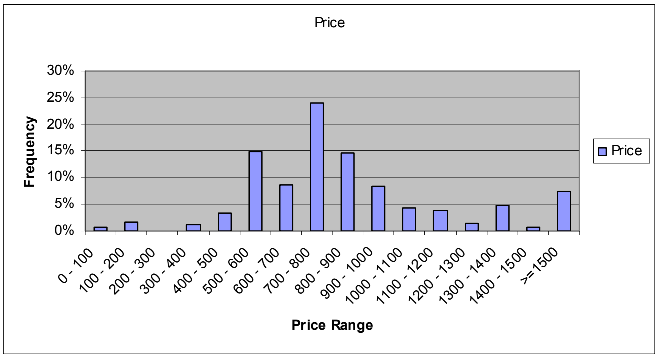

Fig. 7. Distribution of the *Price* attribute.

the user's online decision behavior precisely, then most of the decision support systems can be designed and tested before the real users actually try it and the development time and expenses can be greatly saved.

There are some limitations for our current results. One limitation is that the results we obtained are mainly dependent on the specific decision problem (finding an apartment from a list of candidates) and the specific system we developed (the flat finder system based on example-critiquing interaction). As mentioned earlier, in online situations the decision process is a tri-nary relationship among the decision maker, the decision support system and the decision problem. The user's online decision behavior is affected by both the nature of the decision problem and the design of the decision support system. When the decision problem or the decision support system changes, the user's online decision behavior inevitably will be altered. Another limitation is that so far we only have tens of subjects participating our experiment, and because of this, it is still unknown the certainty level of these results we obtained. Moreover, currently we only provided evidences to show how users act during the interaction process, it is lack of a systematic behavior model to predict the decision behavior of other unknown users or the known users in the future time.

In the future we plan to collect more interaction log files by organizing more subjects to participate online experiments. We believe that more samples of the interaction records can help us develop a more accurate user decision behavior model. We will also try to integrate the extended accuracy-effort framework with some mathematical modelling techniques like Markov Model, Bayesian Decision Network to systematically model the user's online decision behavior. The long-term goal of our research is to develop a more general behavior theory to model user's online decision making behavior and to predict the user's future decision behavior.

6 Acknowledgements

Funding for this research was provided by Swiss National Science Foundation under grant 200020-103490. The authors thank Vincent Schickel-Zuber for significant contribution in the development of the web based interface of FlatFinder system.

References

1. von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press, Princeton (1944)
2. Keeney, R., Raiffa, H.: Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley and Sons, New York (1976)
3. Simon, H.A.: A behavioral model of rational choice. Quarterly Journal of Economics **69** (1955) 99–118
4. Payne, J., Bettman, J., Johnson, E.: The Adaptive Decision Maker. Cambridge University Press (1993)
5. Zhang, J., Pu, P.: Effort and accuracy analysis of choice strategies for electronic product catalogs. In: Proceedings of The 20th ACM Symposium on Applied Computing (SAC-2005). (2005)
6. Pu, P., Faltings, B., Torrens, M.: Effective interaction principles for online product search environments. In: Proceedings of the 3rd ACM/IEEE International Conference on Web Intelligence, IEEE Press (2004)
7. Keeney, R.: Value-Focused Thinking: A Path to Creative Decision Making. Harvard University Press, Cambridge (1992)
8. Burke, R.D., Hammond, K.J., Young, B.C.: The FindMe approach to assisted browsing. IEEE Expert **12**(4) (1997) 32–40
9. Linden, G., Hanks, S., Lesh, N.: Interactive assessment of user preference models: The automated travel assistant. In: Proceedings of the 6th International Conference on User Modeling (UM97). (1997)
10. Shearin, S., Lieberman, H.: Intelligent profiling by example. In: Proceedings of the Conference on Intelligent User Interfaces, ACM Press New York, NY, USA (2001) 145–151
11. Boutilier, C., Brafman, R.I., Domshlak, C., Hoos, H., Poole, D.: Preference-based constrained optimization with CP-nets. Computational Intelligence, Special Issue on Preferences in AI and CP, to appear (2005)
12. Reilly, J., McCarthy, K., McGinty, L., Smyth, B.: Dynamic critiquing. In Funk, P., González-Calero, P.A., eds.: Proceedings of the European Conference on Case-Based Reasoning (ECCBR-04). Volume 3155 of Lecture Notes in Computer Science., Springer (2004) 763–777
13. Reilly, J., McCarthy, K., McGinty, L., Smyth, B.: Incremental critiquing. Knowledge Based Systems **18**(4-5) (2005) 143–151
14. Zhang, J., Pu, P.: A comparative study of compound critique generation in conversational recommender systems. In Wade, V.P., Ashman, H., Smyth, B., eds.: AH. Volume 4018 of Lecture Notes in Computer Science., Springer (2006) 234–243
15. Pu, P., Viappiani, P., Faltings, B.: Increasing user decision accuracy using suggestions. In: SIGCHI conference on Human factors in computing systems (CHI06), Montreal, Canada (2006)