Towards Enhancing Trust and Accuracy in Preference Elicitation

Janna-Lynn Weber

December 1, 2008

Abstract

Systems which elicit preference from users are presented with a unique set of problems based on their nature. This paper examines user trust as a possible reason why preference elicitation systems are met with resistance from users. With the example of a personalized news recommendation system in mind, we examine ideas control, understanding and personalization. We make recommendations for improvement to this system and the subsequent study that should see increased improvement in user satisfaction and system accuracy. The underlying goal is that increasing the user's trust in the system, which will in turn increase the accuracy of preference elicitation.

1 Introduction

Researchers agree that preference elicitation systems must take care to properly interact with their users. As with any system that works with human using computers, there can be many unexpected reactions. While these systems continue to improve their accuracy by way of the algorithms and complexity they are still, in some cases, met with user resistance. Some users feel out of control and at risk to have their privacy violated[11]; they have trouble trusting these systems and their outputs. This can cause users to withhold information which, may then cause the system to perform poorly, decreasing the user's trust even further.

Personalization has been suggested as ways to increase trust [7, 8, 10, 14]. This concept allow the user to gain a sense of control and understanding. This paper will explore the ideas of trust in preference elicitation systems. Including the role personalization can play in gathering and displaying the user's preferences or making recommendations. What techniques can be employed to increase trust and at the same time improve the system's accuracy?

The motivating work of [1] used an adaptive news recommender system with open user models. Their study suggested that open user models improved trust but harmed the accuracy of the system. We will take a critical look at this work and extend it, in order to improve the results. As with any system that works with human users it is critically important to find the right solution for the right problem. We can then extract ideas observed in this instance to make recommendations for more preference elicitation systems.

2 Motivations

As the world of options becomes more complex, people are likely to become more reliant on recommender systems to be able to easily survey all their available options. A recommender system are generally classified in two ways; Users creating recommendations and the system coordinating them for other users, or the system providing recommendations to the user based on their specific preferences. For purpose of this paper we will focus the latter system and discuss it in more detail. Generally as we begin to turn to these systems for recommendations, they need to learn about us in order to make appropriate recommendations. Regardless of their accuracy, they face interesting challenges when it comes to being accepted and relied on by users[11]. Some obstacles that stand in the way of acceptance include: trust, privacy, control, persuasion, believability, accuracy, security, etc.

On a more specific level we will examine the paper "Open User Profiles in Adaptive News Systems: Help or Harm?" [1]. The group, Ahn et al., created and studied an adaptive news systems with an open user profile. The

YourNews system is explained in greater detail in their paper. YourNews consists of implicit preference gathering based on keywords within the news articles that the user clicks on to read. The preference profile that is created is revealed to the user in an understandable way; The user is then allowed to add and remove keywords from the profile as they wish, in hope of creating a better profile. They report in this study that, the open profile as designed, may have increased some trust with its users but overall negatively effected the system's ability to provide the most accurate recommended results. This outcome seems counterintuitive of the way a recommender system should work. If there is more trust from the user, they should be more willing to provide their accurate preferences and work with the system. If you have a more complete set of preferences, the system should perform better. As we explore the ideas of trust and preference elicitation systems we will refer back to the YourNews system as an example.

3 Concepts of Trust

In preference elicitation systems we ask the users to trust the system's recommendations. We also ask the user to trust the system with their personal preferences. Preference elicitation system must trust that the users are providing accurate information about themselves in order to produce accurate results. It is easy to see that trust is an important topic in preference elicitation systems. In general, researchers have been studying the concepts of trust for decades. Much of the computer science literature on trust borrows from the psychology and sociology disciplines.

While trust and its manifestations are always being researched, trust has generally been referred to as a risk assumption in the absence of complete information. Each person has their own trusting style, their own predisposition to trust[10]. Social research suggest that factors such as religious affiliations, culture, birth order and economic level combine to influence the individual's ability to trust. Psychology suggests that trust is based on perceived control in a situation. Research into trust formations have suggested many bases that people use to build their trust[10]. Two of these bases

that are critically looked at in computer systems are identification-based and knowledge-based. Identification-based trust discusses the trust build between parties who genuinely care about the other and believe in the intrinsic value of the relationship. Knowledge-based trust or cognition-based trust, is trust based on evidence of trustworthiness. This evidence of trustworthiness comes from the ability of the truster to predict and understand the trustee. By increasing the user understanding and control we aim to increase the user's knowledge-based trust in the system.

Trust in the age of the World Wide Web (WWW) has a new set of dimensions. Users are no longer able to use non-verbal communication to assess the trustworthiness of a person or company. This constraint makes identification-based trust more difficult to assess. In the case of knowledge-based trust there are many possibilities to expand the user's understanding and control in a system, in order increase the trust. The inherent complexity of computer systems makes the average user unsure of how they work and who to trust. Additional dimensions include the user's sense of privacy and security. The vastness of the WWW, has significantly raised the risk of identity theft to the user's awareness. Web services must work harder to secure information so that the user's privacy is protected. Users also claim they wish to be in control of which private information is expressed and used within the company or service[3]. From the companies perspective, the user's information is essential to personalize the experience.

Personalization has a been shown to have a reciprocal relationship on user's trust[4]."Trust is not only a prerequisite for good personalization, good personalization is also generates trust" [11]. There are three general concepts relating to personalization. These concepts can be used in conjunction to effect the user's perception of a system. The first is that personalization can be seen as the system emulating a customer relationship[10]. This typically occurs when a Web site individually greets the user and is used in place of a face to face conversation or greeting. This type of personalization can be seen as a way to increase identification-based trust. The second concept of personalization is also referred to as customization, where the users is given the ability to tailor their experience[10]. This personalization

helps the user feel in control of the system. Finally, personalization can be the system adapting to the user, by considering the user's preferences[3, 7]. This concept leads into preference elicitation systems being used for adaptive interfaces. Each concept of personalization uses personal or preference information about the user to tailor the experience for them. In the next section we describe some concepts of preference elicitation systems.

4 Preference Elicitation

Preference elicitation systems seek to gather, interpret and apply the personal preferences of the user in order to solve a problem. The process has become popular with online services as they are potentially serving a large number of people and catering to each person is ideal. Recommendation services are used to search over large sets of products and select the best matches based on the user profile[5]. A user profile (model) in this case, is the set of preferences for the individual user, organized in a coherent way. Often this organization will become complicated as the constraints on the problem accumulate. For example, in the YourNews system there are many different articles, from many different sources that the user might be interested in; The system's goal is to provide accurate news article recommendations based on the previously generated preferences for that user.

There are two ways to gather preferences: explicitly or implicitly. Explicit preference gathering is the process of having the user explain their preferences. Users are able to state their preferences in the form of survey responses, questionnaires, or even in the way they rate items on the Web site. With the previously mentioned rise of identity theft, many people are being more careful about sharing their personal information. Implicit preference gathering is done by observing the users behaviours in some way and basing the user model on this information. Some Web services consider the item the user is looking at currently, what they have looked at in the past and sometimes what they have purchased. Using this information they then recommend similar items the user should be interested in. Many Web services use a combination of explicit and implicit preference gathering tech-

niques. Users tend to prefer to be explicit with their preferences as they feel it gives them more control[3]. Implicit gathering techniques, can be seen as less bothersome to the user experience. In the study, [3] users explain that implicit preference gathering has suffered from two problems. First, the preferences gathered have produced the wrong recommendations for the user, reducing their trust in the system. The second problem they found was that when implicit preference elicitation was done correctly, it felt like a violation of the user's privacy. The occurrence of the second problem also caused the user's trust to decrease, as they felt out of control and confused by the system.

5 Trust and Preference Elicitation

Preference elicitation systems represent very novel work, and with any sort of new work they are going to be met with some resistance. People are hesitant to give up their control and grasp of the situation to trust the computer system to work for them. Based on the information already presented in this paper, we want to blend control, understanding and personalization into the system in order to increase user's trust and the overall effectiveness of the system. In the next section we look at the YourNews system presented by Ahn et al. [1] and discuss how it fell short of this goal and ways to improve the system and the study method.

6 Personalized News

6.1 YourNews System

The YourNews system was created and presented by Ahn et al. [1]. The system creates a profile for the user based on which news articles the user clicks on within a particular subject. The profile consists of keywords from these articles based on scraping the article and removing stop words. Keywords that occur more often are given more weight in the user profile. Using this profile, the system recommends news articles that the user may find inter-

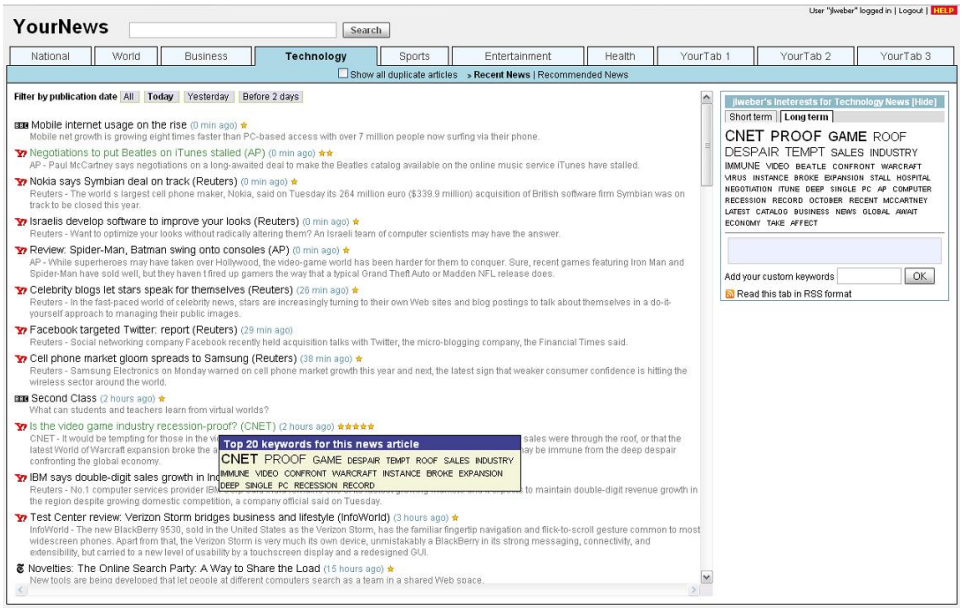

Figure 1: A screenshot taken from the YourNews system.

esting or relevant. The profile is available for the user to see and manipulate. The keywords are arranged in order of their weight in the profile and for visual feedback the font is increased. Users are able to add or remove words to personalize their profile. Users are also able to see the keywords that will be added to their profile, if they click on a particular article. The user is given a short term and a long term profile for each of the 8 categories of news. In addition, for any of these profiles the user can view recent news or recommended news. Figure 1 shows the YourNews system.

In the study, the accuracy of the short term profile was tested. They were interested in the effect of allowing the user to manipulate the profile. Participants were asked to research about a topic in recent news. Then the recommended articles for each profile was ranked and compared against the "ground truth". They discovered that when the user manipulated the profile the recommendations were generally not as accurate for this task. However, the ability to personalize the profile was well received. Based on exit interviews the participants said they trusted the experimental YourNews

system more because of the open profile. While the YourNews system was able to improve some trust the loss of accuracy was unexpected. As users continued to work with the system, in the second half of the study, they learned that editing the profile was not effective and stop doing so. With the following improvements to the system and the study we believe the results will be much more positive.

6.2 System Improvements

Allowing the user to edit the profile, must be done carefully in a way that compliments the user's mental model of the system. The YourNews system presented a clear way to convey to the user, which keywords were being used and which keywords carried the most weight. This presentation allowed the user a clear understanding of how the system works. In order to increase trust by user control, the system should allow for mostly explicit preferences. The paper [3], demonstrated that users prefer explicit preference gathering. In the YourNews system, the user should be able to initiate the profile based on their own keywords or suggestions of common keywords. Additionally, instead of automatically adding the keywords from any article that the user reads, the system would suggest these keywords for the profile. In this case, the user has control over what is added into their profile. Any changes in the profile should be immediately reflected in recommended articles, so the user continues to develop their understanding of the system and can correct the profile as needed. The much earlier work of [12] demonstrates the need for this immediate feedback. In addition to keywords pulled directly from the article, synonyms of words should be suggested as they may be relevant to the user but not necessarily mentioned within the article. In the current system, keywords can be added or removed from the profile; in the new version, the user should also be able to easily increase or decrease the weight of keywords. Because all preferences that are being added to the profile are explicit, the ability to arrange keywords is important.

Finally, to reduce the user's mental load and increase the usefulness of the system, each user should be provided with only one profile. Asking the

user to balance 16 profiles across 32 views, is unfair to the user. The system's goal is to make reading relevant news articles easier for the user. It should be able to recommend news articles regardless of the category. However, if implicit information shows that a user is consistently choosing articles from one particular category they should be able to represent this in their profile.

6.3 Study Improvements

The study of the YourNews system, measured accuracy based on the recommended articles for a specific research task. While this is a valid use of the system and the short term user profile, we believe it is not the best use of the system. The news recommendation service should provide accurate and interesting news articles to the user, so that the user does not have to sort through all the news articles. The measures of accuracy are less directly related to the content and more towards the user's happiness. To measure this type of accuracy and the user trust in the system, we recommend a longer study where the participants are avid news readers. The length of the study should be at least a week of allowing the participants to work with the system at their leisure. This allows the user to become comfortable with the system, and it allows the system time to learn about the user. The length of study also considers the results of the [10] paper; Which suggests the effect of the an individual's predisposition to trust has less impact the more the user works with the system. Secondly, we suggest avid news readers because they should be genuinely interested in the product. Also they should be more aware that the system is recommending the type of news articles they want.

The specific measures in a study like this include: the amount of personalizing they do to their profile, the amount of recommended articles they read, and qualitative data from surveys or interviews with each participant. The amount of personalization to the profile indicates the reliance on the system. If we study this over time we may see the users either constantly editing or setting and forgetting it. The amount of recommended articles read indicates their trust in the system. The supporting hypothesis would

be that the group of participants using the experimental system read more articles than the control group. Finally, surveys and interviews will allow the participant to express deeper feeling about the system, the profile and the recommendations.

7 Discussion

The paper presented by Ahn et al.[1] on the YourNews system used a number of techniques to increase the user's trust while using the system. They wanted to study the effects on trust and accuracy when the user is able to see and edit the user profile. We believe with the proposed improvements, the study results would show an increase in trust and accuracy. Following the paper [1], Ahn et al. presented a paper on a new system called TaskSieve[2]. This system used short term, editable task models, to assist in researching activities. The task models were created in a similar fashion as the user profiles in the YourNews system. They were able to test the system with information analysis and found three main findings. First, searches on their system returned a larger number of relevant documents for the participant. Second, TaskSieve allowed participants to be more productive in their research related task. This was shown by the participants of the experimental system collecting more notes in the first ten minutes than the control system. Third, users were generally happy using the experimental system, though statistically their enthusiasm had not increase from that of the control system. This finding is still important as the experimental system was considerably more complicated and the interface was more cluttered. These positive study results partly reflect matching the right test to the right solution to the right problem.

There are other solutions that have been posed to improve trust and accuracy in preference elicitation systems. Here we briefly highlight three such solutions. First, personalized elicitation process[7]. Where the process of elicitation is customized by means of language or interface. The aim of this system is increase trust by making the user feel comfortable with the system, thus effecting identification-based trust. Second, providing

explanations through behavior metamodels[13]. This follows the idea that the profiles themselves can become complicated and a higher level abstraction could be a better explanation for the user. Third, the use interactive explanations[9]. In this system the user creates a dialog with the system and at any point can ask for further explanations. Unfortunately, many of these solution have not been tested with actual users.

The work outlined in this paper does not just apply to preference elicitation systems. For example, work towards intelligent assistant agents may consider increasing control and understanding in the same ways. Trust mechanisms may continue to change from product to product or site to site, because humans can be inconsistent. However, the idea to convey control and understanding to the user will go far to increase trust. Personalization is important to both trust and preference gathering. Preference elicitation systems should equally aim to improve their accuracy and their influence on trust, in order to increase their acceptability and utility in the future.

8 Future Work and Conclusion

There are two possible areas for future research in this area. First, much of the discussion in this paper focuses on enhancing trust between user and the system. Possible future work can expand into creating trust withing multiple user preference elicitation system; Where trust must be additionally fostered user-to-user. The risk of one user manipulating the system increases, as well as the risk of privacy violations.

The second possible area for future work would examine security risks. By using a preference profile that the user can see and edit are there new security risks? Does exposing preference information make the user likely to be identified? In addition it is possible that the company producing the system would be jeopardizing their business model by exposing how they create preference profiles.

In this paper, we reviewed work regarding the effects of users working with preference elicitation system. Allowing the user to be comfortable with the system, will increase the acceptance and reliance on these systems.

However, making users feel comfortable is a difficult task, where earning the user's trust is crucial. We discussed concepts of trust including control, understanding and personalization. Using these concepts we critically reviewed the YourNews system of Ahn et al.[1] to offer improvement that will increase the trust and accuracy of this system. We discussed how this work can relate to the larger issue of trust in preference elicitation. Including a brief overview of other possible solutions. Future work focuses on introducing additional levels of unpredictability, including multiple users and security risks. Working with human users is difficult as we are unpredictable creatures and success in one implementation does not guarantee success in the next.

References

[1] Ahn, J., Brusilovsky, P., Grady, J., He, D., and Syn, S. Open user profiles for adaptive news systems: help or harm? In *WWW '07: Proceedings of the 16th international conference on World Wide Web* (New York, NY, USA, 2007), ACM, pp. 11–20.

[2] Ahn, J., Brusilovsky, P., He, D., Grady, J., and Li, Q. Personalized web exploration with task models. In *WWW '08: Proceeding of the 17th international conference on World Wide Web* (New York, NY, USA, 2008), ACM, pp. 1–10.

[3] Alpert, S., Karat, J., Karat, C., Brodie, C., and Vergo, J. User Attitudes Regarding a User-Adaptive eCommerce Web Site. *User Modeling and User-Adapted Interaction 13*, 4 (2003), 373–396.

[4] Briggs, P., Simpson, B., and De Angeli, A. Personalisation and trust: a reciprocal relationship? *Designing Personalized user experiences in eCommerce* (2004).

[5] Chen, L., and Pu, P. Survey of Preference Elicitation Methods. *Ecole Politechnique Federale de Lausanne, Tech. Rep. IC/2004/67* (2004).

[6] Dzindolet, M., Peterson, S., Pomranky, R., Pierce, L., and Beck, H. The role of trust in automation reliance. *International Journal of Human-Computer Studies 58*, 6 (2003), 697–718.

[7] Jannach, D., and Kreutler, G. Personalized user preference elicitation for e-services. In *The IEEE International Conference on e-Technology, e-Commerce and e-Service, 2005* (2005), pp. 604–611.

[8] Lam, S., Frankowski, D., and Riedl, J. Do You Trust Your Recommendations? An Exploration of Security and Privacy Issues in Recommender Systems. *LECTURE NOTES IN COMPUTER SCIENCE 3995* (2006), 14.

[9] McGuinness, D., Glass, A., Wolverton, M., and da Silva, P. Explaining Task Processing in Cognitive Assistants That Learn. In *Proceedings of the 20th International FLAIRS Conference (FLAIRS-20)* (2007).

[10] Serino, C., Furner, C., and Smatt, C. Making it Personal: How Personalization Affects Trust Over Time. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05)-Track 7-Volume 07* (2005), IEEE Computer Society Washington, DC, USA.

[11] van der Geest, T., Pieterson, W., and de Vries, P. Informed Consent to Address Trust, Control, and Privacy Concerns in User Profiling. *Privacy Enhanced Personalisation, PEP* (2005), 23–34.

[12] Wærn, A. User Involvement in Automatic Filtering: An Experimental Study. *User Modeling and User-Adapted Interaction 14*, 2 (2004), 201–237.

[13] Wallace, S. Enabling Trust with Behavior Metamodels. In *Proceedings of the 20th International FLAIRS Conference (FLAIRS-20)* (2007).

[14] Weiss, D., Scheuerer, J., Wenleder, M., Erk, A., Gülbahar, M., and Linnhoff-Popien, C. A user profile-based personalization system for digital multimedia content. In *DIMEA '08: Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts* (New York, NY, USA, 2008), ACM, pp. 281–288.