Opportunities for User Involvement within Interface Personalization

Andrea Bunt [1,2] and Michael Terry[2]
[1]Department of Computer Science
University of Manitoba
bunt@cs.umanitoba.ca
[2]David R Cheriton School of Computer Science
University of Waterloo
{abunt, mterry}@cs.uwaterloo.ca

Abstract

Personalization provides users with the opportunity to work in feature-reduced interfaces that uniquely suit their needs. While adding system intelligence to the process has the potential to make personalization more efficient and effective, doing so in a useful and usable manner typically requires a high degree of user involvement. In this paper we discuss two projects, MICA and Ingimp, each of which highlights a different type of user involvement. In MICA, user involvement occurs at run time: Users are given the opportunity to decide how much to rely on intelligent recommendations and are given access to the system's underlying reasoning. In Ingimp, user involvement occurs at design time: A community of active users is supplying usage data that we are using to decide what types of personalization schemes would be most beneficial. We also discuss open issues and plans for future work within each project description.

1 Introduction

Interface personalization is the process of taking a full-featured interface (or high-functionality application -- HFA [Fisher, 2004]) designed to suit the masses and transforming it to more effectively support the individual. Example personalization schemes include hiding infrequently or rarely used features [McGrenere *et al.*, 2002], creating additional personalized interface structures [Gajos *et al.*, 2005], and ordering features so that the more frequently used features appear in more easily accessible locations (e.g., at the top of a menu vs. the bottom) [Findlater and McGrenere, 2002].

One way to achieve a personalized interface is to augment the interface with a special mechanism that allows the user to personalize autonomously (i.e., an *adaptable* approach). These mechanisms have been positively received by some users (e.g., McGrenere *et al.*'s two-interface model for MsWord [2002]). Fully user-controlled personalization, however, can be both time-consuming and, at times, ineffective [Bunt *et al.*, 2004]. Consequently, there has been considerable interest in adding intelligence to the interface to assist with personalization, through either *mixed-initiative* or purely *adaptive* approaches (e.g., [Bunt *et al.*, 2007a; Gajos and Weld, 2004]).

Despite theoretical benefits of intelligence-enhanced interface personalization, effective personalization based solely on system intelligence, and without a high degree of user involvement, is difficult to achieve. First, there is the issue of system accuracy; even with the most sophisticated AI techniques, perfect assessment and prediction of user needs are unattainable goals. Second, fully automated and opaque personalization can have negative usability side effects, including feelings of loss of user control, lack of transparency and lack of predictability [Jameson, 2003].

In this paper we discuss two projects towards intelligence-enhanced personalization, illustrating a number of ways to involve users in the process. The first is the MICA project, where users can choose the level of autonomy with respect to personalization decisions and view the inner-workings of the system's intelligence. We then present work-in-progress on the Ingimp project, an instrumented open-source application, where a community of over 200 regular international users is currently supplying detailed data on their usage of the application. We discuss how this type of user involvement has the potential to inform the design of any personalization schemes, and in particular, task-based personalization.

2 The MICA System

The overall goal of the MICA (Mixed-Initiative Customization Assistance) project [Bunt *et al.*, 2007a; Bunt *et al.*, 2007b] was to explore a middle ground between two opposing approaches to personalization: (1) an *adaptable* approach, where personalization is fully user controlled and (2) an *adaptive* approach, where personalization is fully system controlled.

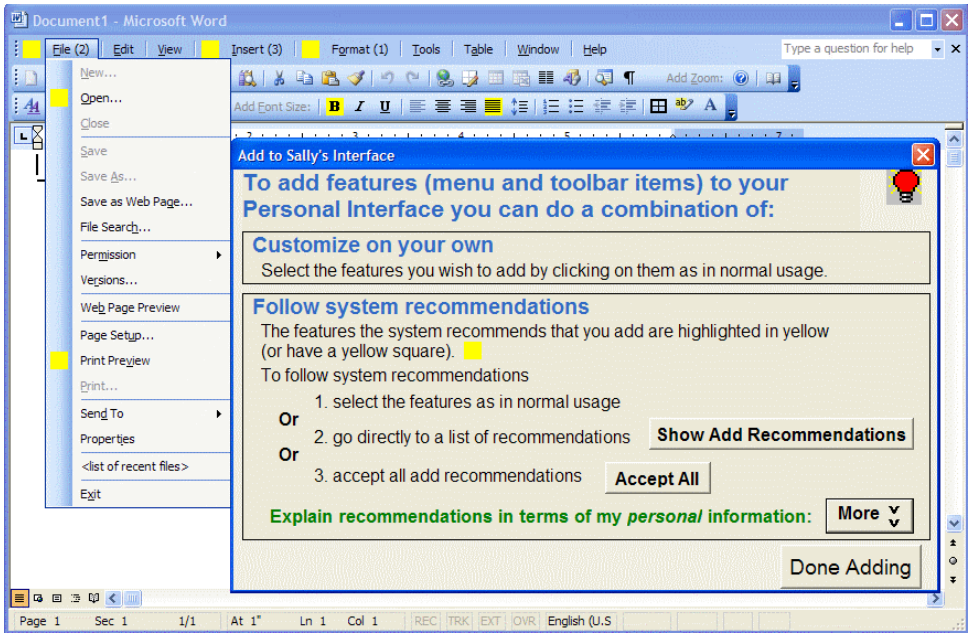

Figure 1: MICA's personalization interface for adding features to the user's Personal Interface. Users can add features by (1) selecting them from the menus and toolbars, where MICA's recommendations are highlighted with yellow squares, (2) selecting them from the list accessible through the "Show Add Recommendations" button, or (3) accepting all recommendations using the "Accept All" button. MICA's rationale component is accessible through the "More" button.

Figure 1 shows a screenshot of MICA's interface, which is currently implemented for MSWord 2003. Personalization within MICA relies on McGrenere *et al.*'s two-interface model [2002], where users can create and maintain a feature-reduced Personal Interface, with the Full Interface a button click away. When users choose to personalize, MICA dynamically generates recommendations, consisting of features to add to, or remove from, their Personal Interfaces.

Within MICA, we explored user involvement at two levels. The first was allowing the user to choose where on the adaptive/adaptable spectrum s/he wishes to reside. We also explored involving the user in MICA's strategy for generating intelligent personalization recommendations.

2.1 Choosing the Level of Autonomy

Prior work, such as that by McGrenere *et al.* [2002] and Jameson and Shwarzkopf [2002], has indicated that different users have different preferences with respect to user-controlled versus system-controlled personalization. Therefore, we designed MICA's personalization interface to allow users to choose the extent they wished to follow system recommendations. By selecting features from the menus and toolbars, users can personalize almost completely autonomously (perhaps using system recommendations as a visual guide). Users can increase their reliance on system intelligence by picking and choosing from a list of recommendations. Finally, they can automatically accept all recommendations with a single button click.

Similar to prior research, qualitative feedback from our evaluations indicated that users did, in fact, have strong preferences in terms of their desired levels of autonomy. Therefore, giving users the choice of how to much to make use of system intelligence appears to be both a natural and promising direction for user involvement in intelligence-enhanced personalization.

While the above approach was favourably received by the users in our evaluations, we note that neither the purely adaptable approach nor the purely adaptive approach is currently available as an alternative within MICA's interface. In other words, intelligent recommendations are always present, and the interface never adapts without user input. Providing the user with choices along this full spectrum is a potential area of future work. We are also interested in exploring whether users' preferences along this spectrum change over time as familiarity with both personalization in general, and MICA's recommendations in particular, increases.

2.2 Access to and Involvement with System Rationale

The second form of user involvement within MICA consists of involving the user in MICA's decision-making process. Currently this consists of providing information on MICA's rationale including information on why and how MICA makes recommendations, along with any relevant User Model assessments (see Figure 2). Providing a window into the system's underlying reasoning serves the purpose of increasing system transparency and predictability,

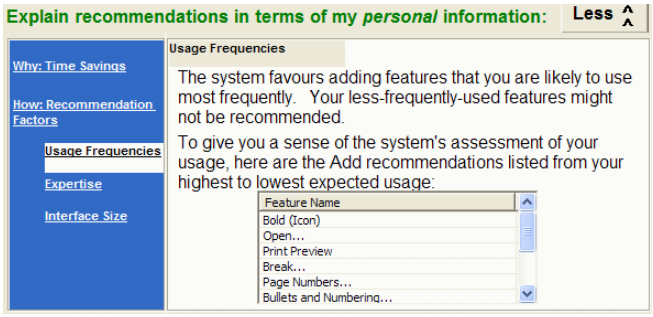

Figure 2: A portion of MICA's rationale, which describes why and how recommendations are generated. The portion shown here describes the Usage Frequencies factor (other factors include Expertise and Interface Size) and displays the system's User Model assessments for this factor.

and also allows the users to make an informed decision as to the extent to which to follow recommendations. For example, if the user disagrees with any of the rationale or assessments, s/he could (and perhaps rightfully so) choose to personalize autonomously.

Our evaluations showed that for some users (but not all) this type of transparency is an important component to feelings of trust and predictability in the system. Further study, however, is required to understand the manner in which system accuracy impacts users' reliance on the rationale and the nature of their subsequent behaviour towards the system intelligence. For example, users might spend more time viewing the rationale when they receive recommendations that they feel are inappropriate. The number of recommendations that they then choose to follow might depend on the perceived accuracy of the recommendations.

While not currently implemented within MICA, a way to further increase user involvement in system decision making would be to make the rationale component more interactive, for example, by allowing the user to change parameters in the User Model. The hope is that this type of open interaction (also referred to as *scrutable* [Kay, 2006]) would lead to more accurate system assessments and consequently more accurate personalization recommendations. Designing and evaluating a modifiable rationale component is another planned avenue of future research.

3 Ingimp

We now shift gears to discuss a project that is currently in its initial stages -- exploring personalization opportunities for an instrumented open source application known as Ingimp. Ingimp is an instrumented version of the open-source drawing application GNU Image Manipulation Program (GIMP) [Terry *et al.*, 2008]. As part of its instrumentation, Ingimp collects a large amount of data concerning users' interactions with the system, including the following: system characteristics (e.g., number of monitors); command usage; interface events; and document characteristics (e.g., number of layers in an image). Users are also given the opportunity to provide a description of their current tasks by inputting an activity tag (see Figure 3).

Figure 3: The Ingimp startup screen, where users can enter an activity tag describing the task(s) they are about to perform.

Currently Ingimp has been installed over 800 times, with over 5,000 log files collected in the 18 months since it was released. An original motivation of the project was to collect data that could help improve system usability. We are now exploring how we can use this large amount of data and high degree of user involvement to provide users with personalized interfaces.

While some initial proof-of-concept analysis was performed when designing MICA [Bunt *et al.*, 2004], user involvement in MICA occurs primarily at run time, a classic and direct form of user involvement in an intelligent system. In contrast, Ingimp's user community is involved at design time. Specifically, we are using the wealth of data that they have provided (and continue to provide) to first understand the extent to which personalization would be beneficial and second to explore the feasibility and suitability of different personalization schemes. In our opinion, this type of feasibility analysis, and/or proof-of-concept exploration is far too often missing in the intelligent interaction community, with researchers eager to explore new AI techniques (or novel applications of existing ones) without first carefully and thoroughly examining user needs.

Our use of the data for these purposes is still in the early stages of analysis. Therefore, our intent for the remainder of this section is to give a sense of what types of analysis could be beneficial and what other user involvement might be appropriate as work in this area progresses. We conclude the section by discussing how others in the community can access the data.

3.1 Potential for Personalization

Our first goal was to understand whether Ingimp users could potentially benefit from personalized interfaces. Therefore, we began our analysis by exploring the degree to which users are making use of the available features or commands within the interface (of which there are several hundred). Using one-month's worth of data, for 194 users

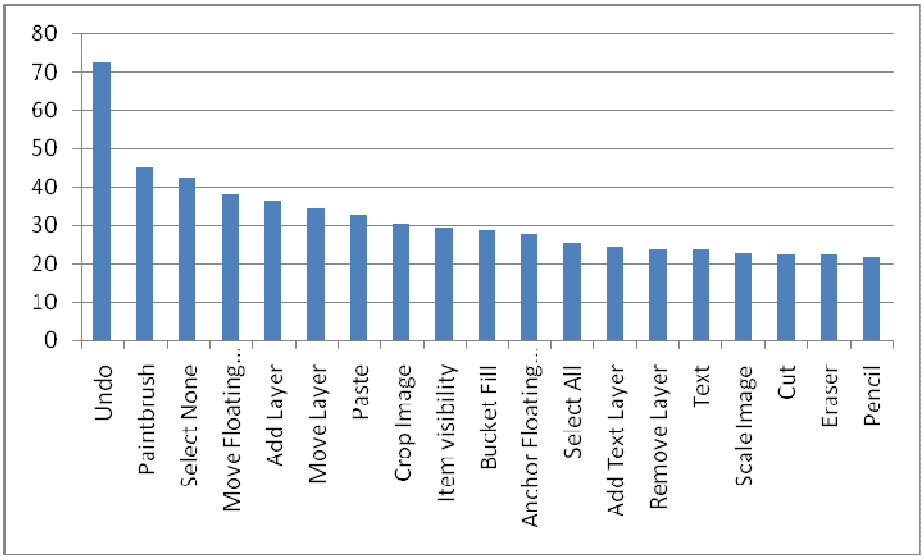

Figure 4: The percentage of Ingimp users who used each of the top 20 most popular commands.

who used at least one command during that time period, we found that, on average 14.6 commands were used, ranging from 1 to 149 (stdev: 19.2).

We next looked at the degree of overlap between our users' command sets. A total of 326 different commands were used during the one-month period being studied. Figure 4 shows the percentage of people that used each of the most popular 20 commands in terms of number of users who invoked them (as opposed to the number of total invocations). Apart from Undo, which was used by 75% of users, even the most popular commands were used by less than 50% of users. Further analysis showed that there were 208 commands used by 5 users or fewer. In other words, our users' command vocabularies are fairly distinct.

Combining the above two analyses, we see that a) Ingimp's interface is much larger than these users need and b) it would be difficult to design a feature-reduced interface that would suit all users. Consequently personalization would likely be beneficial for these users.

3.2 Feasibility of Task-Based Personalization

Most personalization schemes, including the type supported by MICA, assume that the user will have only one personalized interface that will gradually evolve to suit his/her needs. An alternative is to have more of a task-based approach, where the nature of the personalized interface changes based on the current task. We are currently analyzing the data to explore the feasibility of this latter alternative.

We have discovered that even how to find support for task-based personalization within the data is not trivial in that it is not immediately obvious which types of analyzes are appropriate. Consequently, we will likely have to examine the data from several different angles. First, we are interested in examining both the extent to which command usage changes for a given user from session to session. We are also looking for sets of features that tend to cluster together across users. Finally, we are interested in analyzing the extent to which the user-supplied tags correspond to meaningful command clusters.

While clustering the data according to activity tags sounds particularly appealing, unfortunately, not all users supplied meaningful tags when starting a new session, which might be at least partially due to the fact that the benefit in doing so is not necessarily apparent at the present time. If we are able to find enough initial support for task-based personalization from the current data, our hope is that once this type of personalization scheme is made available to the community, the amount of user involvement along these lines would increase. In particular, we hope that for both personal benefit and to benefit the community at large, users will be willing to take the time to carefully label their tasks. Given the culture of the open source community and the fact that these users have already indicated willingness to help the community by sharing their usage data, there is reason to be optimistic that this type of involvement might be possible. With a single-user, commercial application such as MSWord (i.e., MICA's current domain), such user involvement would be far less likely.

3.3 StatsJam

The data collected through the Ingimp project is available not only for our experimentation, but in keeping with spirit of the open-source application, the data is also openly accessible through a web-based SQL query interface known as Stats Jam [Stats Jam, 2009]. Apart from the data collected as part of the OWL project, which examined MSWord usage for users within a single organization [Linton and Schaefer, 2000], this type of usage data has rarely been

made this easily accessible to those interested in developing intelligent interactions.

4 Summary and Future Work

MICA and Ingimp illustrate a number of ways to involve users in intelligence-enhanced personalization. The involvement can occur at run time, by providing the users with differing levels of personalization autonomy or by making the details of the system's rationale available. Users can also participate at design time by supplying usage data to motivate and guide the development of personalization schemes.

One avenue of future work involves allowing users to not only view the system's rationale but also to actively manipulate it, to help create system behaviour that better suits their needs. We also plan to continue exploring user involvement in task-based personalization, including incentives for users to effectively label and describe their tasks, so that the system can generate sets of task-specific personalizations that are of benefit to the entire community.

Acknowledgements

The MICA project was done in collaboration with Cristina Conati and Joanna McGrenere at the University of British Columbia. Ingimp project collaborators include Matthew Kay, Brad Van Vugt, Brandon Slack, Terry Park, and Jon Pipitone.

References

[Bunt *et al.*, 2004] Andrea Bunt, Cristina Conati, and Joanna McGrenere. What role can adaptivity play in an adaptable system? In *Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2004),* pages 117-124, 2004

[Bunt *et al.*, 2007a] Andrea Bunt, Cristina Conati, and Joanna McGrenere. Supporting interface customization using a mixed-initiative approach. In *Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2007),* pages 92- 101, 2007

[Bunt *et al.*, 2007b] Andrea Bunt, Joanna McGrenere, and Cristina Conati. Understanding the utility of rationale in a mixed-initiative system for GUI customization. In *Proceedings of the International Conference on User Modeling (UM 2007),* pages 147-156, 2007.

[Findlater and McGrenere, 2004] Leah Findlater and Joanna McGrenere. A comparison of static, adaptive, and adaptable menus. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2004),* pages 89-96, 2004.

[Fisher, 2004] Gerhard Fischer. User modeling in human–computer interaction. *International Journal of User Modeling and User-Adapted Interactions*, 11(1-2):65-86, 2004.

[Gajos and Weld, 2004] Krzysztof Gajos and Daniel Weld. SUPPLE: automatically generating user interfaces. In *Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2004)*, pages 93-100, 2004.

[Gajos *et al.*, 2005] Krzysztof Gajos, David Christianson, Raphael Hoffmann, Tal Shaked, Kiera Henning, Jing Jing Long, and Daniel S. Weld. Fast and robust interface generation for ubiquitous applications. In *Proceedings of the International Conference on Ubiquitous Computing (UBICOMP'05)*, pages 37-55, 2005.

[Jameson, 2003] Anthony Jameson. Adaptive interfaces and agents. In J. Jacko and A. Sears, editors, Human-Computer Interaction Handbook, pages 305–330. Erlbaum, Mahwah, NJ, 2003.

[Jameson and Schwarzkopf, 2002] Anthony Jameson and Eric Schwarzkopf. Pros and cons of controllability: An empirical study. In *Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2002)*, 2002.

[Kay, 2006] Judy Kay. Judy Kay: Scrutable Adaptation: Because We Can and Must. In *Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems* (AH 2006), pages 11-19, 2006.

[Linton and Schaefer, 2000] Frank Linton and Hans-Peter Schaefer. Recommender systems for learning: Building user and expert models through long-term observation of application use, *International Journal of User Modeling and User-Adapted Interactions*, 10(2-3): 181-208, 2000.

[McGrenere *et al.*, 2002] Joanna McGrenere, Ron Baecker, and Kellog Booth. An evaluation of a multiple interface design solution for bloated software. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2002),* pages 163-170, 2002.

[Stats Jam, 2009] Stats Jam.http://www.ingimp.org/statsjam Accessed May 2009.

[Terry *et al.*, 2008] Michael Terry, Matthew Kay, Brad Van Vugt, Brandon Slack, and Terry Park. Ingimp: introducing instrumentation to an end-user open source application. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2008)*, pages 607-616, 2008.

.