

# Personalized User Preference Elicitation for e-Services

Dietmar Jannach and Gerold Kreutler

*Institute for Business Informatics and Application Systems, University of Klagenfurt*  
{dietmar.jannach, gerold.kreutler}@uni-klu.ac.at

## Abstract

*The quality of the results produced by personalized e-service applications like product recommenders, buying advisory applications, or product configurators is strongly determined by the accuracy of the system's estimate of the individual customer's real needs and preferences. In particular in domains where customers cannot be classified automatically, e.g., based on past buying behavior, these needs have to be interactively elicited by questioning the user. In many existing systems only a "one-style-fits-all" approach based on static fill-out forms is chosen. However, this does not take the user's background or capabilities into account, which consequently leads to a poor quality of the acquired user model.*

*In this paper, we show how extensive personalization of the user preference elicitation process itself can significantly improve the accuracy of interactively acquired user models. A comprehensive view on adaptation and personalization opportunities in the elicitation process is developed and corresponding examples for the domain of interactive buying advisory are given.*

*The presented personalization and adaptation techniques are implemented in a domain-independent software framework for building interactive advisory applications. We describe specific architectural requirements for such a system and discuss results from various real-world applications.*

## 1. Introduction

The provision of high-quality online services is a key differentiation factor for companies in today's highly competitive and transparent business-to-consumer markets. In particular, value-adding electronic services like personalized product recommendation, help desks, buying advisory, or online product configuration are proven means to attract new customers, increase customer loyalty, and turn visitors into buyers. The quality of the results produced by such systems, however, strongly depends on the accuracy of the system's picture of the customer's problem, indi-

vidual background, and his/her real needs and preferences.

Although there are domains where good results can be achieved by automatic classification of the user, for instance based on collaborative filtering techniques or long-term user models, there are also application domains where we have to deal with one-time or first-time users whose preferences must be interactively acquired. In most of today's systems, however, this preference elicitation process is solely based on direct questioning and static fill-out forms. The problem of such a "one-style-fits-all" approach is that in reality users are different with respect to the way they (are able to) frame their requirements. If we think of a system that recommends digital cameras, there will be users who want to specify the technical details of the desired model, while others will only be able to express what they want to take pictures of; others again only want to compare individual models and decide by themselves.

Therefore, we argue that the quality of the results and the end-user acceptance of such e-service applications can be significantly improved when the system interacts with the customer in a personalized way. Extensive personalization of the interaction process can bring us one step closer to real-world face-to-face communication where communication partners adapt their communication style to their vis-à-vis: Users are enabled to express their requirements in a natural way and their confidence in the answers of the system increases when they have the feeling that their requirements are taken adequately into account.

In this paper, we describe a comprehensive view on adaptation and personalization opportunities in the preference elicitation process, which is necessary in a multitude of electronic services that are delivered over the Internet (i.e. e-Services in the sense of [4]). Throughout, we will give examples from the domain of interactive buying advisory. Most of the described techniques were implemented in a domain-independent software framework [10], for which we describe architectural requirements and shortly discuss results from real-world applications.

## 2. Personalization and Adaptation in the Elicitation Process

As a basic classification scheme, we follow the structure from Kobsa et al. [16] and Klein et al. [15] who identified basic categories of personalization opportunities for general hypermedia applications<sup>1</sup>. According to the definition of Kobsa et al., personalized applications are systems that adapt to each individual user's characteristic or usage behavior on several levels [16]. In this paper, we will strictly focus on the *personalization of the preference elicitation phase* and give examples from the field of interactive sales advisory. The described techniques were implemented in the knowledge-based sales advisory system ADVISOR SUITE and are described in detail in [10], [12], and [13], where also a detailed evaluation can be found..

### 2.1 Personalization on the Content Level

**Questions and Answers.** Interactive preference elicitation is based on asking the user a set of questions where in most cases answer possibilities are predefined. Non-adaptive approaches where every user is asked the same set of questions, however, are problematic because they can lead to a poor or even wrong model of the user's real needs and preferences. The typical problems are, e.g.,

- the user does not understand the question because of missing background knowledge. So the question remains unanswered, the default answer is chosen, or even a wrong answer is given, because no suitable answer alternative is presented.
- the user is annoyed by too many questions that are already irrelevant in the current situation because of the previous answers that were given.
- the user is frustrated by a non-natural interaction style in which the system statically poses further questions without reacting situatively on the user's current answers.

Beside the poor user model that is caused by such situations, the user's confidence in the system's results is also reduced if he has the feeling that he was not able to clearly express his real requirements.

In the ADVISOR SUITE system, we address these problems on different levels by personalization of the presentation of questions and answers, as well as personalization of the dialog flow itself in order to achieve a more natural *conversational* interaction style

<sup>1</sup> Note that in general a strict separation is not always possible and there are overlaps in the categorization.

([3], [6], [12]). Figure 1 shows an example of a personalized dialog page. In our system, the individual contents of the page are dynamically constructed on the basis of a declarative knowledge base that contains the required text fragments as well as the personalization rules that determine the page content based on the current user's characteristics [10].

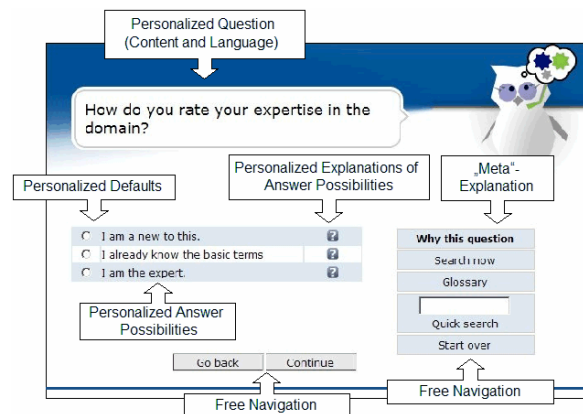


Figure 1. A personalized dialog page.

The personalization capabilities comprise, e.g.,

- the selection of a *language* or the selection of one of several *jargons* that the different user groups might be used to or feel comfortable with.
- the dynamic construction of the answers, i.e., depending on the user's previous answers we can add or remove some alternatives, that are, for instance, irrelevant in the current situation or too complex to answer for an individual user group (variation of *degrees of freedom*).
- the automatic selection of appropriate situation-dependent *defaults* (system proposals), such that the number of required clicks is minimized, which is particularly important in longer dialogs.
- the amount of *optional detailed information* [16] for a question or the answer alternatives, depending on the user's estimated domain expertise.

On the level of the dialog itself, ADVISOR SUITE allows us to design personalized *dialogs* (see Figure 2). In the system, the web-based conversation is modeled in terms of a sequence of *dialog pages* [12]. Each page contains a set of questions and an associated presentation style that determines how the question should be displayed graphically. The personalization rules, i.e., which page should be displayed under what conditions, are modeled as successor relations. These relations are annotated with declarative *transition conditions* over the current user's characteristics. At runtime, the system automatically evaluates these condi-

tions and chooses the appropriate successor page. The system also manages the user's inputs intelligently, such that obsolete user inputs are removed in case of backward navigation and cached for reuse when a page is re-visited later on.

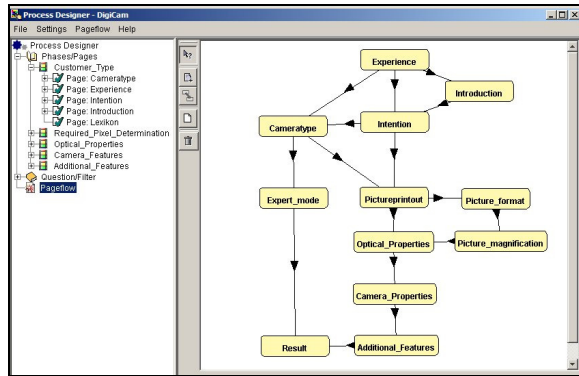


Figure 2. Modeling the dialog flow.

As a further means of personalization, the dialog can also be structured into *phases*. During the dialog, these phases are used to provide the user some feedback on the progress of the dialog. Furthermore, phases allow varying the degrees of freedom with respect to navigation, i.e., whether we allow a user to freely navigate to a certain set of requirements to be elicited.

**Hints.** Another general possibility of adaptation described by Kobsa et al. is the provision of “*optional opportunistic hints*” [16]. Particularly in the domain of web-based requirements elicitation, such hints are a major means to enrich the otherwise mostly system-driven dialogs. With ADVISOR SUITE, hints can be used on different levels, for which the personalization rules are again modeled as conditions over the current user's characteristics. First, they can be used to provide additional information about specific answer alternatives, e.g., for cross-selling or up-selling purposes or in order to provide detailed technical information for non-experts. Second, we also use hints to actively interrupt the dialog, in particular in cases where the user has to be informed about possible inconsistencies in his requirements. Finally, they allow for personalization of the *result* page, i.e., the page where the product proposals are displayed. There, we typically provide supplementary information on the proposal or explain what additional inferences about the user requirements have been made in cases where we do internal reasoning about user characteristics that cannot be directly acquired.

Our experiences show that these hints are well-appreciated by the end-users because they get immediate and personalized feedback on their inputs, i.e., they get the feeling that the system actively monitors the user inputs and participates in the dialog.

**Explanations and Reasoning.** Providing *understandable explanations* for the results<sup>2</sup> produced by the system is a key opportunity to increase the user's confidence in the quality of the system's output [9]. In order to be understandable and useful for the different kinds of end-users, we argue that also the explanations that are provided by the system should be personalized. Depending on the current user's capabilities and interests, we can for instance vary

- the language used in the explanations (e.g., technical or non-technical terms),
- the level of details of the underlying reasoning process that are presented, and
- the *argumentation* style, depending on the user's interests<sup>3</sup>.

A specific form of personalized user preference elicitation with respect to explanations is to enable the user to override the outcomes of the reasoning process to some extent. A typical example for this is related to *derived customer characteristics*. In many applications, the system derives some estimate of customer properties which cannot be asked directly, e.g., the risk class of a customer in an investment advisory scenario. The further reasoning process is then based on the outcome of that classification which is typically also part of the explanation that the system provides. Allowing the (advanced) user to override these *estimates* can therefore also lead to a more accurate user model and consequently better results.

Another particular example where the user can state his preferences in terms of overriding the system's default behavior can also be found in the domain of interactive buying advisory. In the algorithm described in [11], the advisory system uses filter rules to determine suitable product proposals. In situations where no product fulfills all the requirements, the system tries to remove low-priority filter rules until a product is found that fulfills as many requirements as possible. In the explanation phase, the system then both uses a list of requirements that were fulfilled as

<sup>2</sup> This could be, e.g., a product proposal, a configured product variant, or some advice given by a help-desk system.

<sup>3</sup> This option is primarily interesting in the advisory domain, where a certain advisory (product selection) rule will be applied for different kinds of users. The explanation of the rule, however, can be adapted such that it takes the current user's particular interests into account and thus provides personalized argumentations.

well as a list of requirements that were dropped. The initial priorities for relaxing requirements are defined by a domain expert who knows the typical trade-offs in the domain and estimates the average user's preferences. Thus, the possibility for the advanced user to manipulate and override these priorities interactively (see Figure 3), adds another form of how users can state their preferences. Consequently, this reduces the risk of an inaccurate estimate of the user's interests. Finally, hyperlinks in the explanations can be used to reveal to the user which of his inputs caused the application of a rule and to enable the navigation to the corresponding input page to possibly revise his decision.

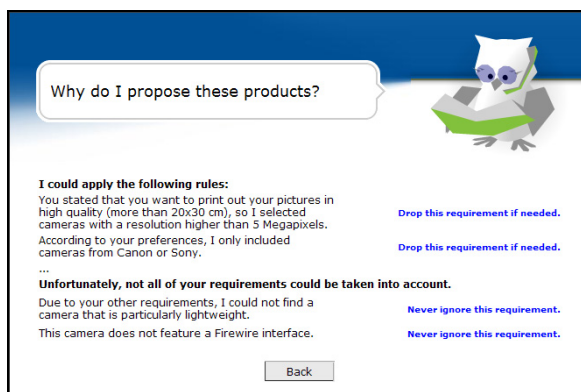


Figure 3. Overriding advisory rules.

**Result Presentation<sup>4</sup>.** Even in the phase when the user requirements have already been elicited and the corresponding results are presented, there are still different opportunities for letting the user refine his requirements, in particular in the advisory domain, where the result is a list of recommended products.

One option implemented in ADVISOR SUITE is to present multiple alternatives, both relatively similar ones as well as *reference products* from different classes of products. In fact, many customers only realize what they really want when they see what alternatives there are. For each of the proposals, there is a corresponding explanation and the option to search for "more like this". Therefore, different similarity measures can be incorporated. For *reference products*, the user can ask for an explanation whether a particular product fits his needs or not. The computation of these explanations consisting of arguments and counter-arguments is again based on the current user model and serves as means for the user to better understand the proposal and how his inputs are taken into account.

<sup>4</sup> We do not focus on the manifold ways of personalizing the presentation of the result itself, as this is not primarily related with the preference elicitation process.

In addition to this, ADVISOR SUITE is capable of providing another way of capturing preferences in specific situations where no proposal can be made (e.g., due to unrealistic user requirements). With the help of a novel technique [11], the system can compute a list of *action alternatives* the user can choose from. Each alternative corresponds to a set of slight changes in the requirements, e.g., revising a specific preference from "very important" to "important", such that a proposal can be made. Again, the key point is that the user can decide by himself about the relative importance of his requirements and possibly revise his initial wishes.

In general, the result presentation phase can be also used to monitor the overall quality of the elicitation process over time. This can be done either by letting the user submit a rating whether he found the proposal useful or not, or by partially monitoring the click-behavior of the user. In ADVISOR SUITE applications, we often present multiple proposals and monitor which items the user did click on, e.g., for viewing a details screen. Furthermore, we can record which requirements were typically given up by users in situations where no proposals were possible and thus learn about the relative importance of individual features. Fully automatic evaluation of these logs and the detection of possible flaws in the knowledge base or the elicitation process are part of our future work.

## 2.2 Personalization on the Interaction and Presentation Level

On the interaction level, we consider to aspects of personalization, the *interaction style* and degrees of freedom in the *navigation*. Regarding the interaction style, up to now we only considered one basic form – a system-driven personalized dialog with dynamically adapted fill-out forms. This style is chosen as the basic form of interaction in ADVISOR SUITE applications because online users are well-acquainted with such forms and in addition feel comfortable when the system actively guides them through the process. Nonetheless, depending on the current user and on the application domain, other forms of interaction or variants thereof can be more intuitive for the user and finally lead to better results in the elicitation process.

**Natural Language Interaction.** The most important aspect in that context is to decide whether a more "user driven" dialog can be provided, e.g., whether the user should be allowed to actively ask questions. In theory, the ultimate solution for this would be a full natural-language interface with an intelligent agent in the background that has both the full domain knowledge

as well as the required conversational knowledge, i.e., how to lead dialogs or react on specific situations. Existing work in that direction (e.g., [22], [23]) shows the general applicability of such natural-language style approaches in e-commerce settings, but still there are open problems. Beside the unresolved problem of understanding general user utterances beside from domain knowledge, today's systems still require huge knowledge acquisition and dialog modeling efforts to reach an acceptable dialog quality. In particular, the dialog cannot be fully user-driven in many domains, because the user's background knowledge is too limited, i.e., they can not properly articulate their questions<sup>5</sup>. Still, there are other ways beside natural-language interfaces of how we can vary the interaction style according to the current user's preferences and capabilities.

**Degrees of Freedom in Navigation.** One important means in that context is the variation of the "degrees of freedom" with respect to navigation. While some users might prefer strong guidance in the dialog, others feel more comfortable when they can more freely steer the dialog, e.g., decide on the order they answer the questions, navigate back and forth, revise answers, or try alternatives. In addition, the amount of additional functionality that is visible to the user can also be personalized, such that beginners are not overwhelmed and frustrated by the complexity of the application, as well as experts do not feel too restricted in their possibilities on the other hand. The same holds for the additional surrounding information that is provided, like a glossary or *Frequently Asked Questions* that the user could be pointed to during the dialog.

**Domain-specific Interaction Styles.** In real life, customers prefer different communication styles with their vis-à-vis depending on the domain in which they make use of a service. For instance, in the financial domain clients are used to be presented a product proposal from their sales person after an intensive requirement elicitation dialog, whereas in domains of consumer goods, e.g., digital cameras, expert customers could expect support in comparing several products from the sales person. In domains of quality and taste, customers might also like to know about opinions of other people before they decide on a product.

Therefore, we claim that supporting the specific interaction styles of a certain domain is a key means to enhance the elicitation process when a service is

brought to the online channel. Consider the following examples: In typical product configuration applications, some users like to specify every detail of the desired configuration by themselves; others would like to start with a basic, preconfigured model and adapt one or the other part; others again only want to specify some key components and functionality and let the system decide on the rest. In interactive buying advisory, some users like to go through a dialog and answer questions; others want to specify requirements on a technical level. A third group might not want to answer questions at all but rather rate different proposals that the system brings up and compare a product with other reference products or evaluate system-generated comparisons (like in ADVISOR SUITE).

Overall, we argue that the quality of the results produced by such online applications significantly increases when we can provide several alternative ways of stating the requirements. The selection of the appropriate interaction style itself can either be done explicitly by the user at the beginning or implicitly by an intelligent system that determines the most suitable interaction form with a few questions.

**Presentation Style.** Adaptation on the presentation level [16] is in our context strongly related with the interaction and content levels. The variation of how things are presented is mostly used to contribute to a comfortable user experience. Personalized presentation variants could be provided with respect to the following dimensions.

First, we can decide whether the elicitation dialog is performed in an own, wizard-style window or is part of a surrounding web-site or portal. While the first variant has the advantage that the user is focused on a small area of interest, the second option allows us to provide additional, surrounding information (glossaries, further hyperlinks to additional information, frequently asked questions) for users that want to actively look for more information during the elicitation process.

Other forms of personalization of the presentation style are related with the language and interaction style used in the elicitation process. Depending on the language style (e.g., more formal or informal and entertaining), an appropriate user interface layout can be chosen; an animated "avatar" (see Figure 1) can serve as virtual conversation partner in guided dialogs that increases the user's buying experience.

In general, also the standard personalization possibilities for handicapped users should be provided in such applications, like, e.g., increasing the font-size or varying the contrast.

---

<sup>5</sup> Users can be even frustrated by poor dialog quality in natural-language systems, because they attribute more intelligence to the system than there actually is.

### 3. Architectural Requirements and Implementation.

There are two major aspects that have to be addressed when developing extensively personalized web applications. First, personalization is a *knowledge-intensive* task [16]. Therefore, such systems have to provide adequate means for acquiring, representing, and – in particular – maintaining the required personalization knowledge. On the other hand, personalized user interfaces have to be extremely flexible, because both the content as well as the navigation options have to be dynamically determined and displayed based on the underlying personalization rules. Moreover, the strong interdependencies between user interface, reasoning, and the knowledge base are also challenging from an engineering perspective, because a clear separation between the application components in the sense of the Model-View-Controller approach [17] can be difficult. A concrete implementation of such a framework for building extensively personalized web applications is the fully knowledge-based ADVISOR SUITE system ([10], [12], [13]). The application area of this framework is *interactive buying advisory* and up to now, several advisory applications in different domains have been successfully implemented with our industrial partners.

Figure 4 gives an overview of the architectural logic of the ADVISOR SUITE system. A major design goal was to adequately support people in their different roles in the development process and also minimize interdependencies between the model, control, and view layers. In our approach, the domain expert (assisted by a knowledge engineer in the early phases) models the elicitation dialog, i.e., questions and an-

swers, hints, the dialog flow, as well as personalization rules with the help of a set of graphical tools and a simple, high-level expression language. For the development of the dynamic web pages, ADVISOR SUITE comprises a set of predefined page templates and user interface (UI) fragments that are automatically assembled by a GUI generator. In these fragments which can be adapted by a web developer, *Java Custom Tags* are used to transparently implement the communication with the advisory server, such that changes in the knowledge base are immediately reflected in the user interface without requiring any changes in the HTML pages. At run-time, the personalized interaction with the clients is managed by a generic controller component that continuously evaluates the current user inputs and steers the personalization and the dialog based on the definitions in the knowledge repository (see [12]).

The experiences from several instantiations show that the strict separation of tasks and roles can help us to significantly reduce development and maintenance costs for highly generic, adaptive web applications. The evaluation of around 1.000 feedback forms and around 100.000 recorded user sessions from one application indicates that the personalized requirements elicitation feature was highly appreciated by the end users. From the recorded sessions we learned that more than 85% stepped through the whole, rather long dialog, which suggests that users felt comfortable and were not overwhelmed during the dialog (the unfinished sessions mostly ended without answering a single question; probably incidental visitors). The evaluation of the questionnaire showed that around 90% liked the personalized service and stated that they would appreciate an expansion of the personalized service on the particular web shop.

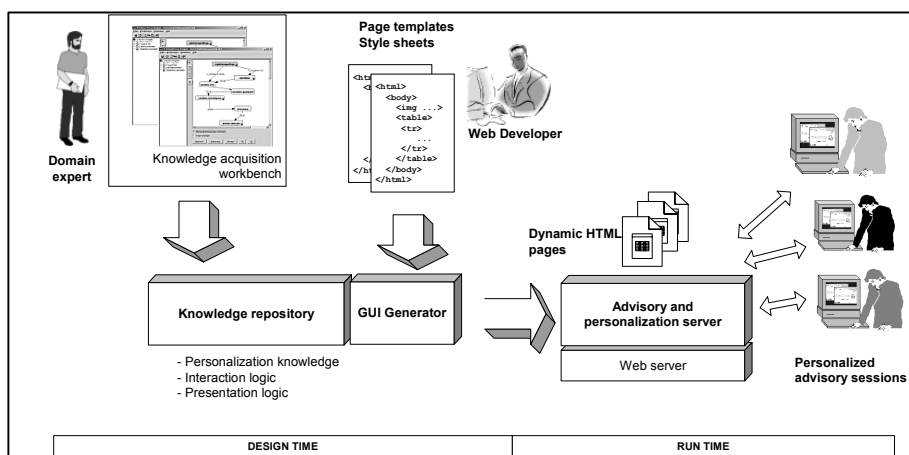


Figure 4. Architecture overview of ADVISOR SUITE.



## 4. Related Work

With the increasing acceptance of the Web as effective distribution channel and the complexity of nowadays' interactive e-Services, the acquisition of high-quality user models and user preference elicitation processes continuously gain in importance [14]. Over the last years, several approaches have been developed that address this problem in application domains such different as Decision Support Systems [8], recommender systems ([3], [6], [24]), or product configuration ([1], [2]). In general, we can distinguish between two basic strategies, how an intelligent system can obtain a good picture of the current user's needs and preferences for its task. Either the system explicitly asks the user about the requirements or it performs indirect reasoning based on e.g., the user's (past) behavior or some stereotype classification mechanism<sup>6</sup>. Many systems also follow a hybrid elicitation process and employ a combination of these strategies.

In the CAWICOMS project ([1], [2]), for instance, such a hybrid approach was chosen for the domain of configuration of complex telecommunication switches: The CAWICOMS system's user interface automatically adapts to the current user's skills by varying the complexity level of the configuration process. The personalization process in the system is both driven by a continuously updated long-term user model (cf. [23]), as well as by a rule-based system that decides on a more coarse level on the overall personalization "strategy". One of the major challenges of such an approach, however, lies in the fine-tuning of the interoperation between the two techniques. The explicit personalization rules in such a hybrid system are in many cases driven by the estimates contained in user model based on Bayesian Networks (like in the POET tool [19]). As a side effect, small changes in the user model can cause a threshold to be exceeded such that an expert rule causing unexpected major effects is activated.

On the other hand, there are systems that base the personalization solely on long-term user models, reasoning on past behavior, or stereotypes. While these systems have shown to be able to produce promising results in domains with re-visiting users ([21], [23]), they typically face the *new user problem* [18], i.e., a poor user model for first-time visitors. In addition, the user models that are maintained by these systems in many cases only capture high-level characteristics like the customer's price sensitivity. In theory, such general user models then can be re-used across different ser-

vices or applications, e.g., in recommenders for different product categories. In practice, however, this portability may be limited because in most applications, domain-specific user interests have to be acquired.

Several forms of (indirect) user preference elicitation are described [7]. The *Analytic Hierarchy Process* (AHP) [20], for instance, is a mathematical decision support tool that uses pair-wise comparisons to determine the relative importance of decision criteria. Such an approach can also be used as a means for preference elicitation in e-Service applications like buying advisory. The advantage of that approach lies in the fact that no complex interactions with the customer are required, i.e., the customer only has to express which one of the two presented products he prefers.

In [8], a critique-based approach for the domain of interactive travel planning is proposed, which acquires the preferences incrementally in a mixed-initiative interaction with the goal to enable users to state hidden preferences, revising conflicting preference values, and weighing objectives in trade-off situations. In their approach, the authors also claim that such systems offer a better user experience and acceptance both because the user is put in control of the process while on the other hand is provided with additional information during the interaction. In general, we claim that the quality of the results of the preference elicitation process also strongly depends on this user experience, i.e., whether the user feels comfortable and well-informed when stating his requirements.

Finally, we argue that the acceptance of personalized applications in general strongly depends on whether the system is capable of explaining the personalized behavior, e.g., why a certain product was proposed due to an internal classification, or why certain questions were asked in the elicitation process.

## 5. Conclusions

E-services are a valuable means to improve customer relations on the online channel. The quality of the results produced by such services strongly depends on the system's capability to elicit the customer's *real* preferences. In this paper, we have argued that personalization of the interaction process is a key factor to accurately acquire the customers' needs. Therefore, we have proposed a comprehensive view on personalization opportunities in this context. We have given examples from the domain of interactive buying advisory and most of the described techniques are implemented in ADVISOR SUITE, a domain-independent software tool. Finally, a short view on the system's architecture, experiences from practical settings and related work have been presented.

---

<sup>6</sup> For a general overview on user preference elicitation techniques, see, e.g., [7].

## Acknowledgements

This work is partly funded by grants from the Austrian Central Bank, OeNB, No. 9706, and from the Kärntner Wirtschaftsförderungsfonds.

## References

- [1] Ardissono, L., Felfernig, A., Friedrich, G., Goy, A., Jannach, D., Petrone, G., Schäfer, R., and Zanker, M.: A Framework for the Development of Personalized, Distributed Web-Based Configuration Systems, *AI Magazine*, 24 (3), Fall 2003, pp. 93-110.
- [2] Ardissono, L., Felfernig, A., Friedrich, G., Goy, A., Jannach, D., Petrone, G., Schäfer, R., and Zanker, M.: Personalizing on-line configuration of products and services, *Proceedings of the 15<sup>th</sup> European Conference on Artificial Intelligence*, Lyon, France, IOS Press, 2000.
- [3] Bridge, D.: Towards Conversational Recommender Systems: A Dialogue Grammar Approach, *Proceedings of the Workshop in Mixed-Initiative Case-Based Reasoning*, Workshop Programme at the 6<sup>th</sup> European Conference in Case-Based Reasoning, 2002, pp. 9-22.
- [4] Bruhn, M.: E-Services – eine Einführung in die theoretischen und praktischen Problem, in: Bruhn, M., Stauss, B. (Eds.): *Electronic Services*, Gabler, 2002, pp. 3-41.
- [5] Burke, R.: Hybrid recommender systems: Survey and experiments, *User Modeling and User-Adapted Interaction*, 4, 2002, pp. 331-370.
- [6] Carenini, G., Smith, J., Poole, D.: Towards more Conversational and Collaborative Recommender Systems, *Proceedings of the Intelligent User Interfaces 2003 (IUI'03)*, Jan. 12-15, Miami, Florida, USA, 2003, pp. 12-18.
- [7] Chen, L., Pu, P.: Survey of Preference Elicitation Methods, Technical Report No. IC/200467, Swiss Federal Institute of Technology in Lausanne (EPFL), Lausanne, Switzerland, July, 2004.
- [8] Pu, P., Faltings, B., Torrens, M.: User-Involved Preference Elicitation, *Workshop Notes, Workshop on Configuration, Eighteenth International Joint Conference on Artificial Intelligence (IJCAI'03)*, 2003.
- [9] Friedrich, G.: Elimination of spurious explanations, *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004*, Valencia, Spain, August 22-27, 2004. IOS Press 2004, pp. 813-817.
- [10] Jannach, D.: Advisor Suite – A knowledge-based sales advisory system, *Proceedings of the 16<sup>th</sup> European Conference on Artificial Intelligence (ECAI 2004)*, Amsterdam, 2004, pp. 720-724.
- [11] Jannach, D.: Preference-based Treatment of Empty Result Sets in Product Finders and Knowledge-based Recommenders, *Poster Proceedings of the 27th Annual German Conference on Artificial Intelligence, KI 2004*, Ulm, September 2004, 145-159.
- [12] Jannach, D., Kreutler, G.: A Knowledge-Based Framework for the Rapid Development of Conversational Recommenders, *Proceedings of 5th Intl. Conference on Web Information Systems Engineering (WISE 2004)*, Brisbane, Australia, November, 2004.
- [13] Jannach, D., Kreutler, G.: Building On-line Sales Assistance Systems with Advisor Suite, *Proceedings of the 16<sup>th</sup> International Conference on Software Engineering & Knowledge Engineering (SEKE 2004)*, Alberta, Banff, Canada, June 2004, pp. 110-117.
- [14] Kiang, M.Y., Gilsdorf, J., Chi, R.T.: Understand User Preference of Online Shoppers, *Proceedings of the 2004 IEEE Conference on e-Technology, e-Commerce and e-Service (EEE'04)*.
- [15] Klein, S., Gueler, S., Lederbogen, K.: Personalisierung im elektronischen Handel, *WISU* 29 (1), 2000, pp. 88-94.
- [16] Kobsa, A., Koenemann, J., and Pohl, W.: Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships, *The Knowledge Engineering Review*, 16 (2), 2001, pp. 11-155.
- [17] Krasner, G.E., Pope S.T.: A Description of the Model-View-Controller User Interface Paradigm in the Smalltalk-80 System. ParcPlace Systems Inc., Mountain View, 1988.
- [18] Rashid, A.M., Albert, I., Cosley, D., Lam, S.K., McNee, S.M., Konstan, J.A., Riedl, J.: Getting to Know You: Learning New User Preferences in Recommender Systems, *Proceedings of the Intelligent User Interfaces 2002 (IUI'02)*, Jan. 13-16, San Francisco, USA, 2002, pp. 127-134.
- [19] Royalty, J., Holland, R., Goldsmith, J., Dekhtyar, A.: POET: The Online Preference Elicitation Tool, AAAI Workshop on Preferences in AI and CP: Symbolic Approaches, 2002.
- [20] Satty T.L.: *The Analytic Hierarchy Process*, McGraw-Hill International, New York, U.S.A., 1980.
- [21] Sung, H.H.: Helping Customers Decide through Web Personalization, *IEEE Intelligent Systems*, 17 (6), November/December 2002, pp. 34-43.
- [22] Thiel, U., Abbate, M.L., Paradiso, A., Stein, A., Semeraro, G., Abbattista, F.: Intelligent E-commerce with guiding agents based on Personalized Interaction Tools, in: Gasos, J., Thoben K.-D. (Eds.): *E-Business Applications*, Springer Verlag, 2002, pp. 61-76.
- [23] Thompson, C.A., Göker, M.H., and Langley, P.: A Personalized System for Conversational Recommendations, *Journal of Artificial Intelligence Research*, 21, 2004, pp. 393-428.
- [24] Weng, S.-S., Liu, M.-J.: Personalized Product Recommendation in E-Commerce, *Proceedings of the 2004 IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE'04)*, Taipei, Taiwan, pp. 413-420.