

Developing a conversational travel advisor with ADVISOR SUITE

Dietmar Jannach^a, Markus Zanker^a,
Markus Jessenitschnig^{a,b}, and
Oskar Seidler^c

^a Institute for Intelligent Systems and Business Informatics
University Klagenfurt, Austria
{dietmar,markus}@ifit.uni-klu.ac.at

^b e-tourism competence center Austria (ECCA)
Innsbruck, Austria
markus.jessenitschnig@etourism-austria.at

^c ThermenResort Warmbad-Villach
Villach, Austria, oskar.seidler@warmbad.at

Abstract

Due to the inherent complexity of building highly-interactive and personalized web applications, the development of a web-based travel advisory system can be a costly and time-consuming task. We see this as one of the major obstacles to a more widespread adoption of such systems in particular with respect to small and medium-sized companies and e-Tourism platforms. The goal of the ADVISOR SUITE project discussed in this paper is thus to provide an off-the-shelf framework and development environment that allows us to build intelligent and easy-to-maintain advisory applications in a cost-efficient way: The main pillars of the presented system are therefore an integrated graphical modelling-environment, the provision of different domain-independent recommendation algorithms, as well as model-based mechanisms to fully generate functional web applications based on declarative definitions in a central knowledge repository. The paper discusses the core concepts and main functionalities of the system by means of an example of an interactive travel advisor developed for an Austrian spa resort.

Keywords: Recommendation system, interactive travel advisory, consumer decision support systems.

1 Introduction

According to a variety of recent studies (e.g. European Travel Commission – ETC) at least in Western countries the Web is the primary source of information for people when searching for suitable travel destinations or planning a trip. In addition, these

studies show that also the share of people that actually book their arrangements via the online channel is constantly taking up. In parallel, the number and variety of information offerings on the Web is also growing, which increases consumer confusion. Thus, it is becoming increasingly important for platform providers to differentiate among each other and come up with value-adding (electronic) services to turn lookers into bookers. Recommender systems can be seen as one of these services as they have proven to be a valuable means to help online customers to overcome information overload and to support them in the decision-making process. Thus, such systems are supposed to play a vital role in current and future e-tourism (see, e.g. Werthner, 2003; Ricci & del Missier, 2004).

However, developing recommender applications and fielding them successfully is a challenging task, in particular w.r.t. small and medium-sized companies and platform providers: On the one hand, established collaborative-filtering based approaches can be quickly deployed but require the existence of a relatively large user community and suffer from cold-start problems. Conversational and knowledge-based approaches on the other hand typically help us to overcome and minimize these ramp-up problems as they base their recommendations - at least partially - on general domain knowledge. However, the development of such a knowledge-based system may induce significant costs for initial knowledge acquisition as well as for application maintenance. Furthermore, the development of a *conversational* user interface (Bridge, 2004; Carenini et al., 2003; Thompson et al., 2004) that aims at simulating the behaviour of an experienced sales agent or travel assistant by leading a personalized dialog with the customer is a cost-intensive task.

The ADVISOR SUITE system discussed in this paper has been designed with the goal of providing an off-the-shelf software system for the cost-efficient development of such advanced, conversational recommender systems. The key contributions lie in the provisioning of a graphical modelling environment for such knowledge-intensive applications, mechanisms for the automated generation of functional web applications, as well as in establishing a more comprehensive approach to product recommendation ('advisory'): The task of *virtual advisor* built with the system shall thus not be limited to retrieving a suitable set of products for a given set of customer requirements, it shall for instance also be able to elicit the real customer requirements in a personalized way, be capable of explaining the proposal, finding alternatives, explaining differences and similarities between products, or providing additional help and guidance depending on the user's background.

In the paper we will first give an overview of a deployed application in the e-tourism domain (a virtual advisor developed for an Austrian spa resort) and will subsequently discuss on the system's architecture and the design rationale behind selected

components. The paper finishes with a discussion of experiences gained from several projects and gives an outlook on future work.

2 VIBE – The virtual spa advisor

The Warmbad-Villach resort is one of Austria's most renowned spas offering a variety of different tourism products that ranges from hotel accommodation in four and five-star categories, a variety of recreational and sporting facilities to health and beauty therapies. Being part of the resort's Web portal, the goal of the virtual advisor VIBE (signifying the name of an ancient goddess of fonts) in that context is to welcome online visitors and provide them a single point of contact for multi-lingual guidance and preference elicitation.

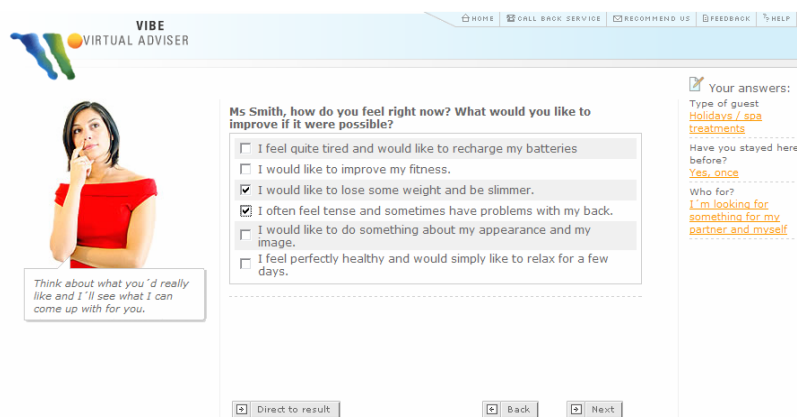


Fig. 1: Virtual advisor screenshot¹

Fig. 1 depicts an exemplary screenshot of the application and gives an impression of the basic question/answer communication style of ADVISOR SUITE: Rather than allowing free-text interaction or using static fill-out forms, the virtual advisor guides the online user through a personalized series of questions in which the customer preferences are incrementally elicited; depending on the current situation, the advisor will also provide additional explanations or give situation-dependent hints and tips.

At the end of the dialog or when the user decides to have provided enough details, VIBE comes up with suitable recommendations (Fig. 2) which also include a personalized list of supplementary add-on packages as well as additional information and explanations which are again based on the customer's preferences.

¹ The virtual advisor can be reached via <http://www.warmbad.at>.

The screenshot displays the VIBE Virtual Adviser web interface. At the top, there is a navigation bar with 'HOME', 'CALL BACK SERVICE', and 'RECOMMEND' links. The main content area features a virtual adviser, a woman in a red dress, with a speech bubble that reads: "Wonderful, we've now got to your final selection. Here's my recommendation for you ...". To the right of the adviser is a table for the 'Ladys Week' package:

Ladys Week	
Length of stay:	per week (7 nights) per person
Meals:	Full board
Accommodation:	The Warmbaderhof
Dates:	01.01.2006-31.12.2006
Rate in single room:	Euro 994
Rate in double room:	Euro 931

Links for 'Details' and 'Why?' are provided next to the table. Below the table, a section titled 'I can also recommend the following packages:' lists several options:

- Our Beauty Special gives you lots of lovely treatments and pampering. Find out all about it at our VIBE beauty salon.
- You will receive information on our biofeedback in the spa health centre.
- Jacobson's relaxation is part of the leisure programme at the Warmbaderhof.
- Health-oriented fitness training, aerobics, stretching, running seminars, Nordic walking and much more.
- We provide medical and health checks at the fitness studio. Our specialist Dr. Markus Kristen, dermatologist and

At the bottom of the interface, there are buttons for 'Back', 'Restart', 'Print', and 'Online-request'.

Fig. 2: VIBE's personalized recommendation

3 Application development with ADVISOR SUITE

In order to build an interactive, personalized advisory service like VIBE, background knowledge has to be elicited from domain experts and formalized in the system. Moreover, the development of a highly-adaptive and maintainable user interface, which reflects the underlying personalization knowledge, can also be a challenging task due to the typically strong interdependencies between recommendation-, personalization-, and presentation logic. ADVISOR SUITE's model-driven development process thus aims at minimizing the costs associated with the construction of such applications (Fig. 1): All of the required knowledge, i.e., both the recommendation- and personalization knowledge can be modeled and maintained with the help of a user-oriented, graphical modeling environment and are stored in a central, declarative knowledge base. A novel template-mechanism and a GUI-generation module are subsequently used to compile a working web application from these definitions; layout changes can be done by a Web developer whose typical task is to adapt the appearance of the application to be coherent with the layout of the e-Tourism platform². At run-time, an intelligent domain-independent personalization agent evaluates the content of the knowledge base and manages the interaction with the individual users accordingly. Personalization itself is implemented based on a rule-based approach as described for instance in (Ardissono et al., 2003); however, the

² For technical details of the implementation and underlying methodology and framework, see for instance (Jannach & Kreutler, 2004) and (Felfernig et al., 2006).

scope of personalization is much broader here and allows for personalization in various dimensions: content, dialog, and presentation (Jannach & Kreutler, 2005).

Within the application-modelling phase, the following different pieces of knowledge (models) have to be designed: The *product model*, which is developed in the first phase, comprises the characteristics of the offered products and services. In the tourism domain, the product model typically comprises a service identifier, prices, availability dates, and so forth. Note that this initial phase is in most cases straight forward, as data about products and services already exists in some electronic form. Developing a *customer model* in the next phase is not that easy anymore, as the domain expert already has to bear in mind, which customer characteristics are of relevance for computation of recommendations and, in addition, which questions need to be posed in order to elicit these potentially latent preferences. The *recommendation model*, in which the logic for generating actual proposals is represented, consists of two separate sub-models: In the *preference-and-filter* model the relationship between customer properties and desired product features can be expressed in terms of *if-then*-style recommendation rules like "If the customer is interested in a family vacation with children, prefer packages that include a childcare service or a petting zoo". It is important to note that with the help of these rules we can also devise a more customer-oriented interaction style, i.e., the virtual advisor will rather ask the customer about preferences and desires rather than on product or service features.

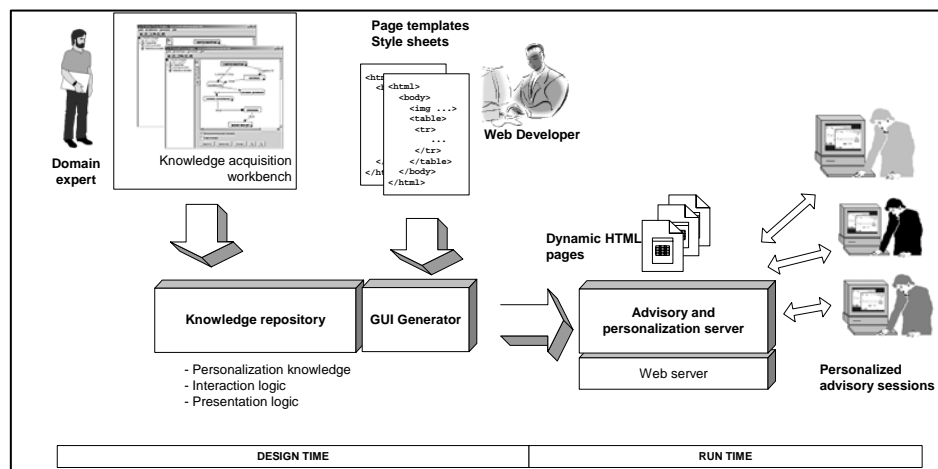


Fig. 3: Process overview (Jannach and Kreutler, 2004)

At run-time, the system will then – based on a new matching and relaxation procedure (Jannach, 2006) – determine those offerings that fulfil as many of the requirements as possible; this particular techniques are also used to derive human-readable explanations that can be presented to the user. In the *utility model*, finally, one can define the rules, according to which the remaining products shall be ranked: In ADVISOR SUITE, the ranking of products is also personalized, i.e., it depends on the user's preferences and is based on the Multi-Attribute-Utility Theory (for more details, see e.g., Felfernig et al., 2006).

One of the major differences between the ADVISOR SUITE system and other content- and knowledge-based approaches to product recommendation lies in the existence of a *dialog and preference-elicitation* model. In this particular modelling view, the domain expert can define how the conversation between the system and the customer shall proceed: Here, the expert can for instance define which questions should be asked in which order or under which circumstances additional hints and explanations should be given to the customer. Like all other modelling views implemented by the system, the dialog is designed on a graphical level, in that case with the help of a sort of flowgraph-based representation. Compared with a natural-language interaction style, the number of conversational paths is of course more restricted. On the other hand, this restriction allows us to keep the modelling complexity at a level, which is still manageable by a domain expert. Furthermore, this rather simple approach of dialog modelling allows us to develop interactive recommenders whose functionality goes far beyond what is available in current web-based recommender applications on the Web.

4 Experiences & further directions

Experiences: Up to now, we have developed more than two dozen commercial advisory applications for different domains with the now-commercialized ADVISOR SUITE system. The broad spectrum of application domains besides e-tourism (from consumer electronics over financial services to fine wine and cigars) indicates that on the one hand our overall system design and the implemented algorithms are domain-independent and sufficiently general; on the other hand, the framework has also proven to be flexible and extensible enough such that domain-specific extensions and characteristics can be easily taken into account. From the perspective of time and costs required for developing a virtual advisor and keeping it up to date, the benefits of using the presented off-the-shelf system as opposed to creating a new solution from scratch have also met our expectations: Quite interestingly, in most cases, the initial sketch of the knowledge base including product data, questions, basic recommendation rules, and a first dialog model can be already drafted already within a first full-day workshop assuming lively participation of stakeholders. In this first

project phase, the ability to quickly generate functional prototype applications proved to be specifically advantageous: This feature helped us on the one hand to establish a common understanding of the projected project outcomes in the early project phases; on the other hand, it also increased the stakeholders' involvement as they see their suggestions and ideas to be immediately taken into account. In addition, we noticed that the number of core recommendation rules in most domains (and also in the travelling domain) remains at a manageable size of a few dozens; the overall costs are mainly determined by the amount of additional information to be provided by the system, the fine-tuning of the different personalization rules, as well as the amount of development times required for setting up the layout or integrating external data sources.

The design of ADVISOR SUITE's user-oriented conceptualisations and graphical modelling tools is driven by the idea that the domain expert herself shall be able to formalize her expert knowledge into the system. However, in most projects up to now, the help of a knowledge engineer was required at least in the early phases or for designing more complex business rules. After that initial phase, however, many of the domain experts were – supported by a comprehensive test and debugging environment (Felfernig et al., 2006) – capable to modify and update the system's knowledge base by themselves. Although these experiences are basically encouraging, further improvements in that direction are required and are part of our future work.

In detail, this specific e-tourism application is characterized as follows: About thirty possible questions have been modelled, out of which always only a personalized subset is asked in order to keep dialog lengths at a size which is acceptable by end users. About the same number of rules has been designed, which – together with a specific relaxation rule for searching for suitable arrival and departure dates – were sufficient to generate personalized recommendations. Overall, one particularity of the domain lies in the limited number of available 'products': In the consumer electronics domain, for instance, sometimes hundreds of different products are available, whereas typical in this domain only a smaller number of packages is available. These limitations could be overcome by using the relaxation mechanism described above ('near matches') together with a possibility to propose 'bundles' consisting of a basic package and a personalized set of add-on activities and side-offers. However, a detailed evaluation of the interaction data of currently several thousands of online users having successfully completed the advisory dialogues has not been done yet; an analysis in the line of (Zanker et al., 2006) is part of our current work.

Future directions: In the particular context of e-tourism applications, we see the following extensions to our system to be highly relevant for future applications:

a) *Intelligent bundling*: In the current implementation of VIBE, the products to be recommended are already bundled packages consisting of hotel accommodation, a set of therapies and so forth. In addition to each package, VIBE also proposes a suitable set of individual add-on services (Fig. 2), whose selection is based on rather simple business rules. However, current trends suggest that in the future the demand for individualized touristic arrangements will be further increasing. Thus, we are currently implementing mechanisms that shall allow us to 'configure' and ensure the validity of such individual combinations (Zanker et al., 2006b). As a basic mechanism we employ a constraint-based configuration approach (see, e.g., Mailharro, 1998), which supports more complex dependencies and bundling rules.

b) *Mobile applications*: VIBE is a pre-trip information and decision support tool. However, travel advisory and recommendation could also be an added value for the customer during and even after his or her journey: The etPlanner project (Höpken et al., 2006), which is partially based on the ADVISOR SUITE technology, is an approach towards more comprehensive electronic service support, which addresses issues of 'context' in the recommendation process as well as content adaptation for mobile devices. Integration of geographical information systems will thus greatly improve the perceived utility and quality of recommendations.

c) *Data analysis*: The ADVISOR SUITE system comprises a built-in logging component which can be configured to record the full history of interactions of each individual user, which means that ADVISOR SUITE's interaction logs are much more detailed and structured than standard Web server logs. Thus, from such logs important types of information about the customer and his preferences can be obtained. One can, for instance, analyse whether there are typical customer profiles or combinations of services which are requested by customers but are not part of any package. This additional knowledge could be subsequently used to adapt the range of offered services accordingly. A second type of information contained in the logs is related to the advisory service itself and pertains to the effectiveness of the application: Typical questions in that context are, e.g., whether the virtual advisor is a) accepted and used by customers and b) whether the system is actually capable of influencing the buying behaviour, which is of particular importance for determining the business value of providing such an add-on service. Although ADVISOR SUITE in its current state already comprises a 'statistics' module which documents the distribution of user answers to specific questions or the ratio of successfully completed user sessions, both further engineering as well as research are still required and part of our current work: On the engineering side, we currently aim at integrating a data mining component with which we can for instance detect hidden correlations between user preferences. On the research side, we are both interested in new metrics that allow us to, e.g., measure changes in actual sales when using such a system and on the other hand on finding basic concepts and guiding principles that are capable of allowing

recommender systems to become an even more persuasive type of technology (Zanker et al., 2006).

d) *Self-adaptation and learning*: One shortcoming of the described system compared with other approaches to recommendation lies in the *static* nature of the system's behaviour: While for instance in systems based on collaborative filtering or case-based reasoning recommendations continuously improve over time, there is up to now no means in our system that resembles such a learning behaviour. In the context of our system, such an improvement has to be based on advanced analysis capabilities as described above. Currently, we work on extending our framework to accommodate such learning behaviour by taking community ratings into account when generating proposals. The question of how such hybrid system can be fine-tuned (e.g., balancing the weight of a community rating vs. an expert opinion) is still in the scope of further research.

5 Conclusions

The cost-efficient development of an interactive travel recommendation service can be a challenging task, in particular for small and medium-sized e-tourism platforms or service providers as standard collaborative approaches do not work well when there is no broad user-community and on the other hand the costs and complexity of building conversational, knowledge-based systems are comparably high. The ADVISOR SUITE presented in this paper is thus an approach to providing an off-the-shelf software framework for the rapid and cost-efficient development of an online pre-trip travel advisory service, the main novelties being the model-driven development approach and the support for the automatic generation of functional web applications. By using an example of a commercially deployed travel advisory application for an Austrian spa resort, we have shown how the development process is supported by the tool and how the typical knowledge-acquisition bottleneck of knowledge-based approaches can be overcome, when adequate modelling environments and user-oriented conceptualisations are provided.

Our further work primarily aims at developing general methods and algorithms, with which the information in the detailed interaction logs can be further exploited, both for learning more about the customers as well as for improving the effectiveness of the advisory application itself.

References

- Ardissono L., Felfernig A., Friedrich G., Goy A., Jannach D., Petrone G., Schäfer R., Zanker M. (2003): A Framework for the Development of Personalized, Distributed Web-Based Configuration Systems. In: *AI Magazine*, 24 (3), pp. 2003, 93-110.
- Bridge, D., (2002) Towards Conversational Recommender Systems: A Dialogue Grammar Approach, Proceedings of the Workshop in Mixed-Initiative Case-Based Reasoning, Workshop Programme at the Sixth European Conference in Case-Based Reasoning, pp. 9-22, 2002.
- Carenini, G., Smith, J., and Poole, D. (2003), Towards more Conversational and Collaborative Recommender Systems, Proceedings of the 8th international conference on Intelligent user Interfaces, Miami, Florida, 12-18.
- Ricci, F. & Del Missier, F. (2004). Supporting Travel Decision Making through Personalized Recommendation. In Clare-Marie Karat, Jan Blom, and John Karat (Eds.), *Designing Personalized User Experiences for eCommerce*, (pp. 221-251), Kluwer Academic Publisher.
- Felfernig, A., Friedrich, G., Jannach, D., Zanker, M. (2006), An Integrated Environment for the Development of Knowledge-Based Recommender Applications, *Intl. Journal of Electronic Commerce*, Special issue on Recommender Systems, to appear.
- Höpken, W. et al. (2006), etPlanner: An IT framework for comprehensive and integrative travel guidance, in: Hitz, M. et al.: *Information and Communication Technologies in Tourism, ENTER 2006*, (pp. 125-134), Springer Wien/New York.
- Jannach, D., and Kreutler, G. (2004), A Knowledge-Based Framework for the Rapid Development of Conversational Recommenders. In: X. Zhou, S. Su, M. Papazoglou, M. Orłowska, K. Jeffery (Eds.): *Web Information Systems - WISE 2004*, Springer New York, LNCS 3306, 390-402.
- Jannach, D., and Kreutler G. (2005), Personalized User Preference Elicitation for e-Services. In: Cheung W., Hsu J. (Eds.): *IEEE International conference on e-Technology, e-Commerce and e-Service*, Hong Kong, IEEE Computer Society, 604-611.
- Jannach D. (2006), Finding Preferred Query Relaxations in Content-based Recommenders, Proceedings of IEEE Intelligent Systems Conference IS'2006, Westminster, UK, IEEE Press, 2006, pp. 355-360.
- Mailharro, D. (1998), A classification and constraint-based framework for configuration. *AI in Engineering, Design and Manufacturing*, 12:383-397, Cambridge University Press.
- Thompson, C.A., Göker, M.H., and Langley, P. (2004), A Personalized System for Conversational Recommendations, *Journal of Artificial Intelligence Research*, 21, pp. 393-428.
- Werthner, H. (2003), *Intelligent Systems in Travel and Tourism*. Proceeding of the 18th International Joint Conference on Artificial Intelligence, IJCAI-03, Acapulco, Mexico, 2003
- Zanker, M., Bricman, M., Gordea, S., Jannach, D., Jessenitschnig, M. (2006): Persuasive online-selling in quality & taste domains, Proceedings of 7th International Conference on Electronic Commerce and Web Technologies - EC-Web '06, Krakow, Poland, pp. 51-60.
- Zanker, M., Aschinger, M., Jessenitschnig, M. (2006b): Knowledge-based composition of Recommendations – Extended Abstract, Workshop Notes of the Configuration Workshop held in conjunction with the European Conference on Artificial Intelligence (ECAI-06), Riva del Garda, Italy.