Zanker, M., Jessenitschnig, M.: ISeller - A Generic and Hybrid Recommendation System for Interactive Selling Scenarios,
In: 15th European Conference on Information Systems, St. Gallen, 2007, pp. 70-80.

ISELLER - A GENERIC AND HYBRID RECOMMENDATION SYSTEM FOR INTERACTIVE SELLING SCENARIOS

Zanker, Markus, Institute of Applied Informatics, University Klagenfurt, 9020 Klagenfurt, Austria, markus.zanker@uni-klu.ac.at

Jessenitschnig, Markus, eTourism Competence Center Austria, 6020 Innsbruck, Austria, markus.jessenitschnig@etourism-austria.at

Abstract

We present `ISeller` an industry-strength recommendation system for online shopping platforms. The system supports several recommendation paradigms like collaborative, content-based and knowledge-based filtering as well as one-shot and conversational interaction modes. A generic user modelling component allows different forms of hybrid reasoning strategies as well as enables the system to support a process-oriented way of interactive selling in various product domains. This application paper contributes a comprehensive scenario for interactive selling on commercial platforms and thus motivates central user modelling services for recommendation systems. Furthermore, we give an outline on the technical architecture and the implemented system. Our presentation will be illustrated with actually fielded examples from domains of luxury-articles like coffee and cigars.

Keywords: Interactive Selling, Hybrid Recommender System, User Modelling

1 INTRODUCTION

Recommender systems (RS) are deemed to help users when they are confronted with large choices and to refer them to those items they might be most probably interested in. Nowadays, many large online shopping platforms employ RS as a convenient service feature as well as a means to increase buying conversion rates. Although, these systems are well appreciated by online shoppers, the vision of a *virtual salesman* is still far away from being realized. Many times RS are solely deployed on platforms as self-contained and somewhat 'isolated' tools relying on their own data collections, but not being integrated with other online and off-line applications such as online booking/selling functionality or customer relationship management (CRM) systems. Furthermore, they are mostly very limited in terms of interactivity when following a one-shot recommendation protocol, while conversational RS engage in a preference elicitation dialogue with users either by explicitly asking questions or by proposing items to choose or critique.

The term *Interactive Selling* is used for information systems that support the sales process in electronic commerce as well as help salespersons in becoming more efficient. An empirical study on the interpersonal relationships between salespersons and customers from a marketing perspective and the implications of different transaction oriented behaviours in real-world sales situations was presented by Tam and Wong (2001). The `ISeller` (**I**nteractive **Seller**) recommendation system supports such online interactive selling scenarios, where different (recommendation) services are dynamically composed. They successively interact with the online customer to create an enhanced shopping experience. Depending on the actual context of the user different sales dialogues and interaction styles are applied. Furthermore, based on `ISeller` hybrid recommendation strategies can be developed that also exploit user information collected by a preceding recommendation service in the current online visit of the user. For instance, if during a product search dialogue the user revealed her/his preference

Zanker, M., Jessenitschnig, M.: ISeller - A Generic and Hybrid Recommendation System for Interactive Selling Scenarios,
In: 15th European Conference on Information Systems, St. Gallen, 2007, pp. 70-80.

for *Robusta* coffee blends due to their higher amount of caffeine and better formation of *crema*[1], a recommender service that promotes accessories will thus exclude standard coffee makers from its selection and restrict its search space on premium espresso machines.

Our contribution in this paper is therefore a scenario of such interplay between different recommendation strategies in an interactive selling scenario. We will give a more formal view on such interactive recommendation problems based on the definition of Adomavicius and Tuzhilin (2005). Consequently, we discuss the resulting requirements on the implementation of recommendation applications as well as present the `ISeller` system.

The paper is structured as follows: Next, we will sketch a motivating example from an actually fielded online shop for premium coffee. In Section 3 we discuss related research that leads into an extension of a recommendation problem to include interactivity aspects (Section 4). Furthermore, we present the `ISeller` system, give an overview on its architecture and explain some of its design rationale.

2 MOTIVATING EXAMPLE

Over the last decade electronic marketplaces established themselves as a popular alternative to traditional brick & mortar shops. While online platforms have serious advantages such as low access costs for computer literates and a bigger choice, its biggest disadvantages besides trust and security issues are still a lack of shopping experience and high search complexity. Commonly recommender systems address these issues by making personalized product propositions to their users. However, for most online-shopping situations the assumption that users invoke a recommendation service once and then decide on buying one of the proposed items is too simplistic. Therefore we propose a more flexible and process oriented interaction paradigm for recommendation applications. First, depending on their intent, users require different types of services to support and guide them through the information acquisition and decision process. For instance advisory services for getting informed on a specific product category or content-based navigation mechanisms to reach similar product items. Furthermore, online shoppers may actually visit a site several times before actually deciding on buying an item or not.

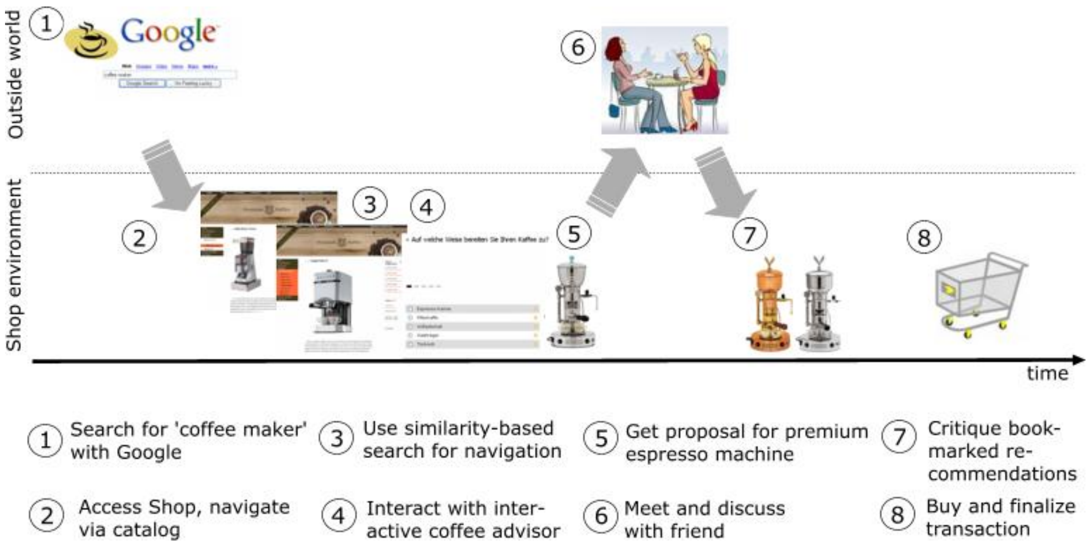

Figure 1: Interactive selling scenario

[1] I.e. reddish brown foam of espresso brewings.

In Figure 1 we therefore depict the shopping process of a fictitious user Sue for illustration. Sue needs to buy a new coffee maker as her old one does not work any longer. A search engine refers her to our online shop 'premium-coffee' (1). After accessing the shop she navigates the product catalogue and carefully studies the detailed product descriptions (2). The shop offers a similarity-based browsing feature that allows users to retrieve quickly those product items that are most similar to the one they are currently looking at according to technical product knowledge. As Sue wants to make sure that she does not miss a better deal, she invokes this recommendation service (3).

After Sue spent already considerable time in accessing different coffee maker descriptions, the shop environment concludes that - though she exhibits considerable purchase interest - she cannot make up her mind. Therefore, the shop environment offers a link to a virtual coffee advisor to her. This catches her attention and Sue conversationally interacts in a question/answer style dialogue with the knowledge-based recommendation service (4).

The system tries to find out about her preferences and makes her curios to different coffee preparation techniques. Finally, it turns out that she would be interested in the espresso brewing method. Therefore, the recommendation service proposes her an espresso machine and explains its advantages as opposed to standard coffee makers which she was looking at before (5). Sue is quite excited about the beautiful chrome-plated machine, but it does not fit within the budget restrictions she had in mind when entering the shop. She leaves and meets her friend (6).

When Sue revisits 'premium-coffee' after several days, her mind has changed and she thinks she would deserve some luxury. The shop environment recognizes her again and makes its past recommendations easily accessible to her in a sidebar. Now Sue applies the critiquing feature of the shopping platform (7). She can formulate some predefined critiques on a reference product such as 'lower price', 'more user-friendly' or 'smaller space dimensions' to get alternative proposals. After few critiquing rounds she is finally satisfied with the proposal and may actually finish the buying transaction (8). Though this scenario is quite simplified, it clearly outlines some of the key characteristics an integrated environment for the interactive selling of products and services requires:

- The buying process spans over a potentially long period of time, which may include several visits of users.
- Users have full autonomy on their buying process, but an interactive selling environment can reason on different states of this process and dynamically offer recommendation services and/or helpful features without annoying them by being too pushy.
- According to the situational context recommendation services need to be parameterized and hybridized to ensure a consistent argumentation chain towards the online customer.

In the following we will discuss these requirements w.r.t. to related work on recommendation systems.

3 DISCUSSION

According to Burke (2002) five basic recommendation techniques can be distinguished: collaborative, content-based, demographic, utility-based and knowledge-based. Among those, collaborative and content-based filtering are the most prominent representatives. Shopping platforms like *amazon.com* exploit collaborative filtering; they suggest items to their customers that other users with similar viewing or buying preferences also liked. Demographic RS are comparable to collaborative filtering except for the fact that the similarity between users is not determined based on similar preferences in the past but on demographic characteristics such as age, address or income. Content-based recommender systems require detailed product descriptions and defined similarity measures in order to identify recommendable items based on their descriptive proximity to those items the user liked in the past. Comparably, knowledge- and utility-based variants require detailed product descriptions, too. In addition both approaches first interact with users to identify their abstract requirements. Then knowledge-based RS reason on explicit domain knowledge in the form of mapping rules between customer requirements and product characteristics. Contrastingly, utility-based RS require domain knowledge encoded in utility functions to compute the recommendable items with highest rank.

Adomavicius and Tuzhilin (2005) identified the common characteristics of these recommendation approaches and thus formalized a recommendation problem by assuming a utility function u_{rec} that measures the usefulness of an item $p \in P$ to users $c \in C$: i.e. $u_{rec} : C \times P \mapsto R$, where R is a totally ordered set of numbers within a certain range. Therefore, they define a recommendation task as finding the item p_i that maximizes the user's utility.

However, this definition is only valid for one-shot recommendation modes where the system does not engage into several rounds of interaction. In order to understand interactive RS that guide their online users and continuously refine their suggestions, a more process oriented view becomes necessary. There exists already a considerable amount of work on so-called conversational variants of RS: Case-based recommendation systems converse with users to identify similar interaction cases from the past to derive recommendable items (Burke 1999, Thompson & Langley 2004, Ricci 2002). Critiquing systems elicit explicit feedback on suggested items and exploit it to narrow down search results in subsequent iterations (Smyth et al., 2004). Shimazu (2001) presented a system that combines explicit asking about customer requirements in a first phase with critiquing of product suggestions in a second phase. Jannach and Kreutler (2005) refined the question/answer style of dialogue to reveal user preferences by identifying several ways of personalizing this process. A mixed initiative model on a dialogue grammar approach was developed by Bridge (2000), who demonstrated the feasibility with a prototype system developed in Prolog.

Nevertheless, these forms of interaction must not only be seen as means towards preference elicitation but also as important functions w.r.t. sales psychology in order to create more persuasive systems (Fogg 1999), let alone being able to purposefully influence users' behaviour in order to increase sales. However research is still in its infancy such that we are not able to understand the cognitive processes and psychological laws that influence online shoppers to their full extend. Users being explicitly asked about the importance of a specific feature of a product might start a cognitive process that lets them end up with a different mindset than before. For instance when users are forced to being explicit about the primary motives for their holiday trip (Ricci 2002), the importance of design or size when choosing an electronic consumer good (Jannach & Kreutler 2005) or defining the level of risk they can take when deciding on a financial investment product (Felfernig 2005) might influence their viewpoint when assessing the utility value a recommended item would provide them. Next, we will have a more formal view on these processes.

4 INTERACTIVE RECOMMENDATION PROBLEM

Thus when moving from pure one-shot interaction modes to interactive recommendation and selling, an additional process-oriented perspective comes into play. The recommendation problem does not solely consist of computing a ranked list of products for a given user, but a series of process steps and conversational moves such as posing a question, presenting an item, allowing the user to critique it or coming up with some specific sales arguments need to be dynamically determined. In addition, RS are operating on some additional information that is neither part of the user nor of the product model. This residual data might name additional contextual information such as the season of the year, current weather conditions (especially in an ubiquitous computing context) or tuneable strategic decisions of the operator of the RS, e.g. to promote a specific set of products for a limited period of time. In many cases a sharp distinction between characteristics of such an environmental model and the user model does not exist. For instance, there might be good arguments for attributing the current location of the user or the technical characteristics of the device she is using to the user model as well as for seeing them as part of the environmental system context. However, a parameter that describes the current workload of the system or the current weather conditions will be for sure related to the latter one.

Therefore, we extend the definition of Adomavicius and Tuzhilin (2005) of a recommendation problem to include the aforementioned interactivity aspects as well as contextual information as follows:

An Interactive Recommendation Problem (*IRecP*) consists of sets of users *C*, items *P*, process steps *S* and system parameters *E* as well as additional relations representing the user model *UM*, the product model *PM* and the process step model *PSM*.

For every user $c \in C$, time points t and path expressions $path_{UM} \in paths(UM)$, the relation UM holds a quadruple $\langle c, path_{UM}, t, val_{UM} \rangle$, where $val_{UM} \neq NULL$ is the evaluation of the expression $path_{UM}$ at a time point t on the instantiated user model of c. $path_{UM}$ contains fully qualified symbolic path expressions, such that deeply structured user characteristics (i.e. domain features). Comparably, knowledge on every item $p \in P$ is structured by path expressions $path_{PM} \in paths(PM)$ and encoded as triples $\langle p, path_{PM}, val_{PM} \rangle$ in the relation PM. Finally, all potential process steps $s \in S$ of the interactive recommendation problem are characterized by $\langle s, path_{PSM}, val_{PSM} \rangle$ in the model of process steps PSM. Note, that val_{UM}, val_{PM} and val_{PSM} in the `ISeller` system are structured objects themselves to allow encodings of sets of values as well as probabilistic inferences.

For illustration of the formalism we sketch a short example based on Section 2:

$$C = \{Sue, \ldots\},\ P = \{CoffeeMaker_1, \ldots\},\ S = \{q_1, \ldots\}$$
$$UM = \{\langle Sue,\ questions.q_1,\ t_1,\ 'replace\ my\ old\ one' \rangle, \ldots\}$$
$$PM = \{\langle CoffeeMaker_1,\ price.domestic,\ 150 \rangle, \ldots\}$$
$$PSM = \{\langle questions.q_1,\ text.english,\ 'Why\ are\ you\ looking\ \ for\ a\ coffee\ maker'? \rangle, \ldots\}$$

Example 1: Illustration of the formalism

Therefore the interactive recommendation task is to determine for a given user a sequence of process steps $s_1^v, \ldots, s_t^v, \ldots, s_n^v$ that reaches the highest overall utility value v according to an utility function ut. This means that $\forall v \neq w : ut(\langle s_1^v, \ldots, s_t^v, \ldots, s_n^v \rangle) \geq ut(\langle s_1^w, \ldots, s_t^w, \ldots, s_n^w \rangle)$.

Note, that each process step s_z^t is determined either by a recommender system or the user in a mixed-initiative model. Depending on its strategy the recommender system exploits more or less information from the three models UM, PM and PSM at a specific point of time to compute its next step.

Compared to the one-shot recommendation case, interactive recommendation does not only suggest products but also questions to pose, sales arguments to consider and any other conversational moves that might be appropriate in specific situations. Since the utility of the functioning of the system is determined ex post by assessing the whole interaction sequence with a user, predicting useful processes is non trivial. Situation dependent deciding on the next move and thinking several steps in advance spans even larger search spaces than chess and thus makes optimal sales processes impossible.

Nevertheless, some first works have been reported on introducing learning based methods for adaptivity of recommender systems. Ricci and Mahmood (2006) applied Markov Decision Processes (MDP) to specific subproblems of interactive recommendation such as query tightening to evaluate their applicability. At the current state, the presented `ISeller` application uses explicit process knowledge to reason on follow up steps and adapts its behaviour accordingly.

5 IMPLEMENTATION

We developed `ISeller` starting from an already existing knowledge-based conversational RS - Advisor Suite (Jannach 2004). Furthermore, our industrial experiences gained when fielding more than thirty recommendation applications in diverse fields such as electronic consumer goods, online tourism or financial services helped us to understand the flexibility and integration requirements that let to the development of the `ISeller` system over the past two years. While most recommender systems are monolithic in the sense of inseparable implementing user modelling and recommendation functionalities within a single application, `ISeller` was conceptualized with a component-based plug-in system architecture following principles of extensibility and modularity. Figure 2 sketches the system architecture having in its core the encapsulated knowledge-based recommender system (KB recommender). The main technical advancements of `ISeller` w.r.t. its preceding Advisor Suite system are: a centralized User Modelling Service supporting a variety of different recommendation strategies and thus enabling interactive selling scenarios, an abstract recommendation service that is instantiated by several different implementations as well as an industry-strength integration layer ensuring short deployment times. Starting from the right side Figure 2 depicts the Data Object Service component that accesses external data sources and performs necessary pre-processing steps. The User Modelling Service initializes itself with already existing user related information and further warehouses all additionally collected profiling and interaction data. The different recommendation services implement a common service interface to support the required polymorphism when the selection strategy needs to be dynamically switched. Users of the system are on the one hand side the domain experts performing setup and maintenance tasks and on the other hand side online shoppers themselves who interact with a commercial platform like sketched in the scenario in Section 2.

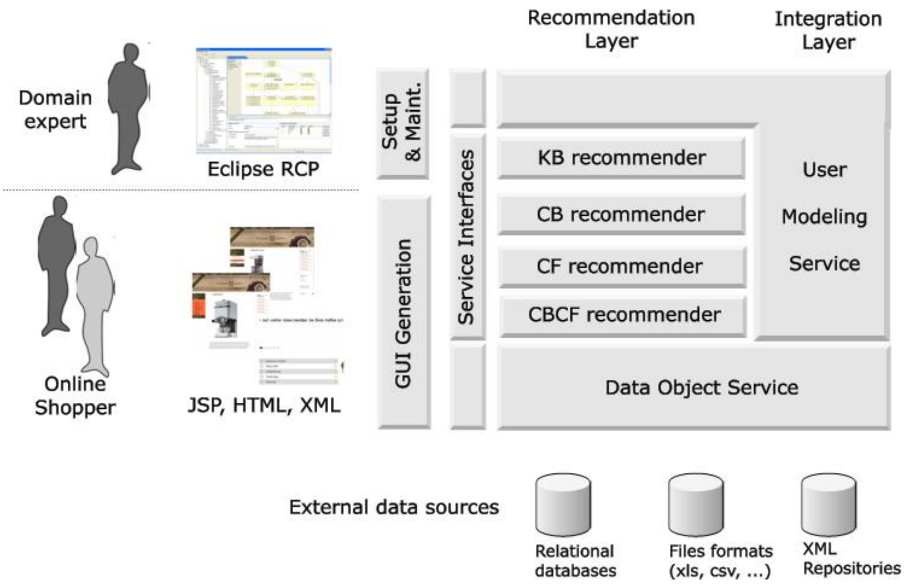

Figure 2: System architecture

The system is conceptualized such that domain experts, like the shop owner or experienced sales staff, are capable to parameterize the functioning of the system and its recommendation strategies in order to keep software engineering efforts low and deployment cycle times short. Online customers interact with Web interfaces that are dynamically generated from the dialogue and sales process definitions due to a template mechanism (Jannach 2004). In the following we will therefore focus on the above mentioned main technical advances: Central user modelling component for different recommendation approaches, abstract recommendation service and a comprehensive integration and pre-processing layer in separate subsections.

5.1 Data integration and pre-processing

Deployment costs as well as the quality of recommendations greatly depend on an efficient as well as encompassing access of already existing product and user data. Automated access to various types of data sources like relational databases, XML repositories and some popular file formats (.xls, .csv[2]) without additional programming effort is realized by providing highly configurable data source connectors from which the user can choose when retrieving from an existing data store. As mentioned above, the system is capable to handle different types of data sources. Especially in the e-commerce domain many different database management systems are in use like MySQL, PostgreSQL or Oracle. In order to cope with the variety of provider-specific extensions to the standard query language we use Hibernate[3] as an object/relational persistency and querying service.

Based on a description how the system can access existing domain data, additional data structures can be defined to model complex relationships between domain objects, like products, user buying histories or interaction data in general. The definition of these data structures is externalized from the application code and provided in the form of domain ontologies, including besides structural information also semantic information (e.g. relationships between data elements). Thus, the structure of existing domain objects can easily be modified by modifying the abstract representation instead of the application code itself. Besides flexibility and extensibility, our ontology-based description enables the implementation of efficient and powerful search mechanisms as well as highly personalized search interfaces for the end user. Furthermore, by using such an ontology-based approach the system benefits from flexible data storage and import mechanism utilizing semantic mapping techniques (Missikoff et al., 2003).

Additional data transformation and pre-processing steps (e.g. definition of relationships, de-normalization or value mappings) are necessary to enable different recommendation strategies. Starting from the integration and selection of existing data stores, a pre-processing step de-normalizes and aggregates source data while the succeeding transformation step maps strings onto numeric values if required or applies functions that normalize domain values onto intervals. Performing these steps results in efficient internal data structures thus that algorithms can operate optimal during runtime.

Furthermore, an extensible scheduling component using the Quartz Job Scheduler[4] was developed that performs periodic tasks to ensure data integrity. The `ISeller` system includes further a modular editor environment based on the Eclipse RCP (Rich-Client Platform) technology[5] such that the traditional knowledge acquisition bottleneck is reduced and domain experts themselves can define and maintain these tasks on a graphical level.

5.2 User modelling

In most recommendation applications, the user model is proprietary to the implemented recommendation technique. However, when having a selling functionality in mind where different recommendation and customer care applications have to work seamlessly together in order to realize a virtual sales assistant, user modelling must not be seen recommendation technique centric. Comparable to CRM tools (customer relationship management) where a central repository maintains user profiles and historic customer contacts, we designed a central user modelling component following the idea of service oriented architectures (SOA). Like all other components of the `ISeller`, it can be deployed in a distributed environment and integrated with other system components using some form of remote method invocation. Its purpose is to generically warehouse all

[2] Microsoft Excel, comma-separated values files.
[3] See http://www.hibernate.org/ for reference.
[4] See http://www.opensymphony.com/quartz/ for reference.
[5] See http://www.eclipse.org/rcp for reference.

Zanker, M., Jessenitschnig, M.: ISeller - A Generic and Hybrid Recommendation System for Interactive Selling Scenarios,
In: 15th European Conference on Information Systems, St. Gallen, 2007, pp. 70-80.

forms of preferences and transaction data and enable means-end oriented personalization and user interaction (Kobsa 2001). Via a query interface recommender services can evaluate arbitrary expressions on the model of a single user as well as retrieve a set of users satisfying a given expression.

Figure 3 depicts the definition of the process flow between different preference elicitation steps (q_1,q_2,…) and the formulation of transition conditions between these steps.

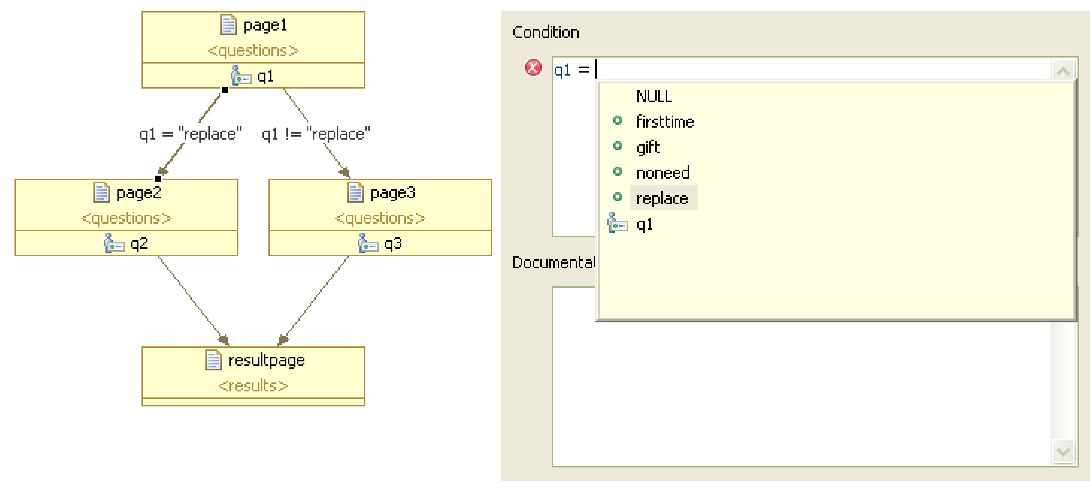

Figure 3: Process flow between interaction steps and their transition conditions

Following a generic User Modeling Shell approach from Kobsa (2001), the user modelling system supports assumptions on system beliefs about user beliefs (SBUB), where system as well as user beliefs are represented by probability values. Based on the data objects identified in the underlying data integration component domain experts can define structured user profiles as well as transaction series on a graphical level. Using a sliding time window approach past transaction data (e.g. sales records, Web logs or product ratings) are extracted from their original external sources, related to the corresponding users and aggregated depending on the informational needs.

Besides the data collection and organization aspects of the user modeling service, it owns strong inferential capabilities: Using a modular extension mechanism, values of profile attributes are derived according to an associated algorithm implementation. They reach from simple aggregations and explicit personalization rules to some forms of ontological user profiling (Middleton et al., 2004). In the latter case additional user preferences are derived from exploiting dependencies between abstract preference categories encoded in ontologies. Both, profile as well as transaction data serve recommender systems depending on their strategy.

5.3 Recommendation strategies

There exists an extensible set of different recommendation functions that converse with the user in problem focused dialogues and make knowledgeable recommendations (knowledge-based/KB recommender), derive cross-selling opportunities (collaborative filtering/CF recommender, content-boosted collaborative filtering/CBCF hybrid recommender (Melville et al., 2002)) and offer similarity-based browsing of product items (content-based/CB recommender).

The knowledge-based approach is using deep domain knowledge and bases its recommendations on an accurate understanding of the user’s current needs, comparable to real-life sales situations. Knowledge-based recommender systems elicit user preferences explicitly by allowing the user to specify the ideal case in an interactively environment using dynamic personalized and potentially persuasive sales dialogs (Jannach 2004). Figure 4 depicts the definition of a question *q1* as already sketched in the small example in Section 4. Recommendations are the result of a reasoning process on the domain knowledge. This reasoning process also allows such systems to explain why recommended items are proposed. However, knowledge-based recommendation requires the existence of domain knowledge that can be exploited and has high setup costs for knowledge acquisition.

Due to the commercial success of e-commerce sites like *amazon.com*, the most well known recommendation technique is collaborative or social filtering and variants thereof. This recommendation technique exploits similarities between users based on the past behaviour (e.g. past ratings of items, purchase histories etc.) and recommends those items to a user that his/her nearest-neighbours liked (Burke 2002). However to the fact that this approach does not need deep domain knowledge, it only leads to good recommendation results, if the users past behaviour is dense enough. For user with a small number of statements on specific items, the algorithm is not able to compute good neighbourhoods and therefore leads to unsatisfying recommendations. On the other hand, also items which were not rated in the past will not be recommended at all. These problems are known as *first-rater* and *sparsity* problem (Melville et al., 2002). The longer the recommender system is in use and active users participates the more accurate results can be achieved.

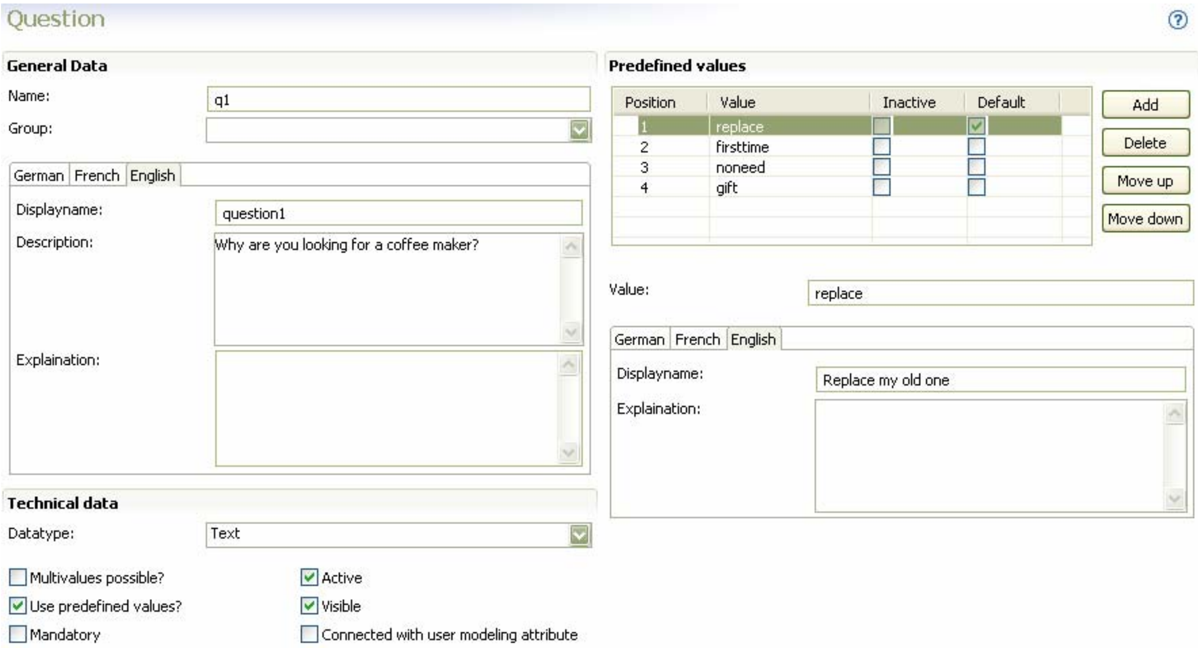

Figure 4: recommendation process

Content-based filtering is based on the hypothesis that the preferred items of a single user can be extrapolated from his/her preferences in the past, thus recommending more of the same. In contrast to the collaborative filtering approach, this approach is quite powerful making recommendations to new users or recommending new items, but is weak in terms of serendipity, i.e. making non-straight-forward recommendations.

As outlined, each recommendation technique considered alone has some shortcomings. Hybridisation of recommendation techniques helps to cope with some of the aforementioned shortcomings. E.g. a knowledge-based recommendation system can achieve good results from the very beginning and could even improve when complemented with a collaborative or social filtering system. Also combinations of collaborative and content-based algorithms exist to overcome the start-up problems of the pure collaborative approach (Melville et al., 2002).

By operating on the same model of the user state-full behaviour in the sense of considering past user actions when computing a recommendation becomes possible. For instance, the user can be asked if she enjoyed the cigar model she bought last time and if she wants a similar recommendation based on product content information. Further hybrids of different recommendation strategies become possible such as a cascading strategy (Burke 2002) where the set of recommendable items according to a knowledge-based recommender serves as input for a collaborative filtering strategy. Content-based recommendation and browsing strategies build on knowledgeable similarity definitions described by Zanker et al. (2006).

6 CONCLUSIONS

The paper presented the `ISeller` recommender system that supports interactive e-commerce scenarios due to its generic user modelling server and polymorphic recommendation components. We contributed a description of an interactive recommendation and selling scenario and discussed related work in this area. Furthermore, we formalized an interactive recommendation problem starting from Adomavicius and Tuzhilin (2005) and presented an implemented system that supports such interactive recommendation and selling interactions on e-commerce platforms. We outlined the system architecture and gave some implementation details.

References

Adomavicius, G. and Tuzhilin, A. (2005). Towards the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. IEEE Transactions on Knowledge and Data Engineering, 17 (6).

Bridge, D. (2000) Towards Conversational Recommender Systems: A Dialogue Grammar Approach. In Aha, D., ed.: Workshop in Mixed-Initiative Case-Based Reasoning at European Conference on Case-based Reasoning (ECCBR).

Burke, R (1999). The Wasabi Personal Shopper: A Case-Based Recommender System. In: 11th Conference on Innovative Applications of Artificial Intelligence (IAAI), 844-849.

Burke, R (2002). Hybrid recommender systems: Survey and experiments. User Modeling and User-Adapted Interaction. 12 (4), 331-370.

Felfernig, A. and Kiener, A. (2005): Knowledge-based interactive selling of financial services using FSAdvisor. In: 17th Innovative Applications of Artificial Intelligence Conference (IAAI), 1475-1482.

Fogg, B.J. (1999). Persuasive Technologies. Communications of the ACM, 42(5), 27-29.

Jannach, D. (2004). Advisor suite - a knowledge-based sales advisory system. In: Proceedings of the European Conference on Artificial Intelligence (ECAI).

Jannach, D. and Kreutler, G. (2005). Personalized User Preference Elicitation for e-Services. In: Proceedings of the IEEE International Conference on e-Technology, e-Commerce and e-Service.

Kobsa, A. (2001). Generic user modeling systems. User Modeling and User-Adapted Interaction. 11 (1-2), 49-63.

Melville, P., Mooney R. and Nagarajan, R. (2002). Content-boosted collaborative Filtering for improved recommendations. In: Proceedings of the 18th National Conference on Artificial Intelligence (AAAI), 187-192.

Middleton, S.E., Shadbolt, N.R. and de Roure, D.C. (2004). Ontological user profiling in recommender systems. ACM Transactions on Information Systems. 22 (1), 54-88.

Missikoff, M., Werthner, H., Höpken, W., Dell'Erba, M., Fodor, O., Formica, A. and Taglino, F. (2003). Harmonise: Towards Interoperability in the Tourism Domain. In: Information and Communication Technologies in Tourism (ENTER), 58-66.

Ricci, F. (2002). Travel recommender systems. IEEE Intelligent Systems. 17 (6), 55-57.

Ricci, F. and Mahmood, T. (2006). Learning and adaptivity in interactive recommender systems. In: Proceedings of the Workshop on Recommender Systems, 111-116.

Shimazu, H. (2001). Expert clerk: Navigating shoppers buying process with the combination of asking and proposing. In: Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI), 1443-1448.

Smyth, B., McGinty, L., Reilly, J. and McCarthy, K. (2004). Compound critiques for conversational recommender systems. In: IEEE/WIC/ACM International Conference on Web Intelligence (WI), Washington, DC, USA, IEEE Computer Society, 145-151.

Tam, J.L. and Wong, Y (2001). Interactive selling: a dynamic framework for services. Journal of Services Marketing. 15 (5), 379-396.

Thompson, C. A. and Langley, P (2004). A personalized system for conversational recommendations. Journal of AI Research. 21, 393-428.

Zanker, M., Gordea, S., Jessenitschnig, M. and Schnabl, M. (2006). A Hybrid Similarity Concept for Browsing Semi-structured Product Items. In Proceedings of the 7th International Conference on E-Commerce and Web Technologies (EC-Web), 21-30.