

An Integrated Environment for the Development of Knowledge-Based Recommender Applications

Alexander Felfernig*, Gerhard Friedrich, Dietmar Jannach, and Markus Zanker

Computer Science and Manufacturing, University Klagenfurt

*Contact author: alexander.felfernig@uni-klu.ac.at.

Abstract—The complexity of product assortments offered by online selling platforms makes the selection of appropriate items a challenging task. Customers can differ significantly in their expertise and level of knowledge regarding such product assortments. Consequently, intelligent recommender systems are required which provide personalized dialogues supporting the customer in the product selection process. In this paper we present the domain-independent, knowledge-based recommender environment CWAdvisor which assists users by guaranteeing the consistency and appropriateness of solutions, by identifying additional selling opportunities, and by providing explanations for solutions. Using examples from different application domains, we show how model-based diagnosis, personalization, and intuitive knowledge acquisition techniques support the effective implementation of customer-oriented sales dialogues. In this context, we report our experiences gained in industrial projects and present an evaluation of successfully deployed recommender applications.

Keywords: Recommender Systems, Personalization, Knowledge Acquisition, End-User Programming, Knowledge-based Recommenders, Testing Recommender Knowledge Bases, Model-based Diagnosis.

I. INTRODUCTION

The selection of products from a complex assortment such as financial services or digital cameras is still a challenging task for customers interacting with online selling environments. In many cases, those environments only offer simple query interfaces based on the assumption that customers know technical product details.

Recommender technologies [1], [31], [37], [39], [41] improve this situation by generating product proposals which are automatically derived from a set of given customer requirements. The three most well known approaches to the implementation of recommender applications are the following: collaborative filtering [20], [36] relies on product preferences of a large set of customers. Recommendations are derived from preferences of a group of customers with similar purchasing patterns. Consequently, no deep knowledge about the product domain is needed in this context. When using content-based filtering [33], recommendations are based on similarities between product descriptions and the preferences of the current customer. When a customer interacts with the recommender application, products are proposed that are similar to those the customer has liked in the past. Finally, knowledge-based recommender applications (advisors) [6], [5], [14], [24], [45] exploit deep knowledge about the product domain in order to determine recommendations. When selling, for instance, financial services, recommendations must adhere to legal regulations and suit the customers financial restrictions as well as their needs and wishes. Compared to customers purchasing simple products such as books or compact disks, customers purchasing financial services are much more in the need of information and in the need of intelligent interaction mechanisms which support the selection of appropriate solutions. Therefore, an explicit representation of product, marketing and sales knowledge [14] is needed. Such deep knowledge allows us (a)

to calculate solutions which adhere to legal regulations, which are in the line with a company's marketing and sales strategy, and which suit the requirements of the customer, (b) to explain solutions to a customer, and (c) to support customers in situations in which no solution can be found. Primarily, knowledge-based recommender technologies can provide the formalisms that are needed in this context.

In this paper, we present the knowledge-based recommender environment CWAdvisor. The major innovations of this environment compared to other knowledge-based approaches [6], [5], [24], [45] are the following:

- Graphical knowledge acquisition. CWAdvisor provides a graphical development and test environment for knowledge-based recommender applications [14], [16]. This environment supports the design of knowledge bases (product, marketing and sales knowledge) and process definitions (intended behavior of the recommender user interface) on a graphical level (see Section II). Such an approach allows rapid prototyping processes by automatically translating graphical models into an executable application.
- Relaxations of filter constraints and repairs of inconsistent customer requirements. Model-based diagnosis techniques [11], [35] can be used to actively support customers in situations where no solution could be found (see Section III): either by identifying minimal relaxations of filter constraints or by the determination of repair actions for customer requirements. Both can guarantee the identification of at least one solution.
- Personalization of sales dialogues and recommendation results. Recommender dialogues in CWAdvisor are based on a finite state model [16] which describes possible interaction sequences of a recommender application on a graphical level. Using such representations, the formulation of questions, answer alternatives, and explanations can be automatically adapted to the domain knowledge level and preferences of a customer. Furthermore, recommendations can be ranked using concepts of multi-attribute object rating [2].

CWAdvisor has been successfully applied in domains such as financial services or electrical equipment. All these applications are based on one of the following basic scenarios:

- Consumer/Customer support. Similar to traditional sales channels, improved sales assistance generates added value for online customers. On the one hand, knowledge-based advisors allow an intuitive access to products for customers. On the other hand, sales representatives are relieved from routine advisory tasks which are taken over by recommender applications.
- Support of sales representatives. Sales representatives interact with advisors when preparing a sales dialogue or when talking to a customer. In this context, they are supported by guided dialogues which consist of questions and explanations focusing on the customers needs and wishes. Consequently, customer-oriented sales dialogues are supported in this context.

Within the scope of these scenarios, we present two example showcases (see Section IV) which clearly demonstrate the utility of our approach. First, a digital camera advisor has been deployed for the largest Austrian online price comparison platform. Effective preference elicitation is a challenge in this context since customers can differ significantly in their product domain knowledge and in their expectations. Second, financial service advisors have been deployed for an Austrian financial service provider. Financial service advisory is a knowledge-intensive task which in many cases overwhelms customers as well as sales representatives. Since the product assortment is quite manifold (investment, financing, pension, life insurance, etc.), many sales representatives focus on selling only a restricted set of products which leads to sub-optimal offers. The goal here is to identify a solution which suits the needs, wishes,

and financial restrictions of the customer and which is in the line with the sales strategy of the company.

The remainder of this paper is organized as follows. In Section II we present the CWAdvisor environment. In Section III we give an overview of Artificial Intelligence (AI) technologies which ease the development of and improve the interaction with knowledge-based recommender applications. Section IV presents an evaluation of successfully deployed recommender applications. Thus, we provide an overview of the CWAdvisor environment which is the result of long-term research in the areas of knowledge-based configuration [2], [12], [18], model-based diagnosis [11], [16] and knowledge-based recommender systems [14], [16], [22]. We focus our discussions (a) on basic principles of testing knowledge-based recommender applications, (b) on mechanisms supporting users in situations where no solution can be found, and (c) on experiences gained from commercial recommender application development projects.

II. CWADVISOR DESIGNER

CWAdvisor technologies can be used for a variety of tasks: the formalization of product, marketing and sales knowledge by non-programmers; the test of recommender knowledge bases; the mapping of customer requirements to a set of product properties; the repair of inconsistent customer requirements, and the explanation of solutions in order to increase a customers trust. The major reasons for the application of CWAdvisor technologies in the financial services domain are the following:

- Solutions must be objective, correct, and explainable. This requirement makes approaches such as collaborative or content-based filtering not the best choice since, explanations in the form of repair proposals, for example, rely on the existence of a knowledge base.
- Typically, financial service providers want to develop and test recommender applications autonomously. Therefore, knowledge representations

are needed which allow both an effective knowledge base design and the validation of calculated results by domain experts.

- Financial services recommendation is a complex task with a large number of constraints and possible solutions [14]. In such scenarios, a knowledge-based approach can significantly reduce the development efforts for recommender applications [17]. Similar reasons motivate the application of CWAdvisor technologies in other domains such as digital cameras where the following additional aspects have to be taken into account:

- In the consumer electronics market we observe a fast technological progress. Knowledge bases have to be updated on a regular basis when new features have to be considered in the recommendation process, or when increased performance characteristics become standard. Such change requests related to the knowledge base happen very frequently and must be fulfilled fast and without extensive programming efforts.
- Personal preferences that cannot be captured in advance play an important role in the domain of digital cameras. In this context, recommender applications should allow customers to experiment with decision alternatives (undoing choices, sorting proposals by different criteria, changing the priority of preferences, etc.).

A. CWAdvisor Architecture

The overall CWAdvisor architecture is shown in Figure 1. This architecture consists of two major parts: the development and test environment, which supports the graphical design of recommender applications, and the runtime environment, which is responsible for controlling the execution of recommendation processes. CWAdvisor knowledge bases and process definitions are developed and maintained using the development and test environment (CWAdvisor Designer, Process Designer, and Test Designer) and are stored in an underlying

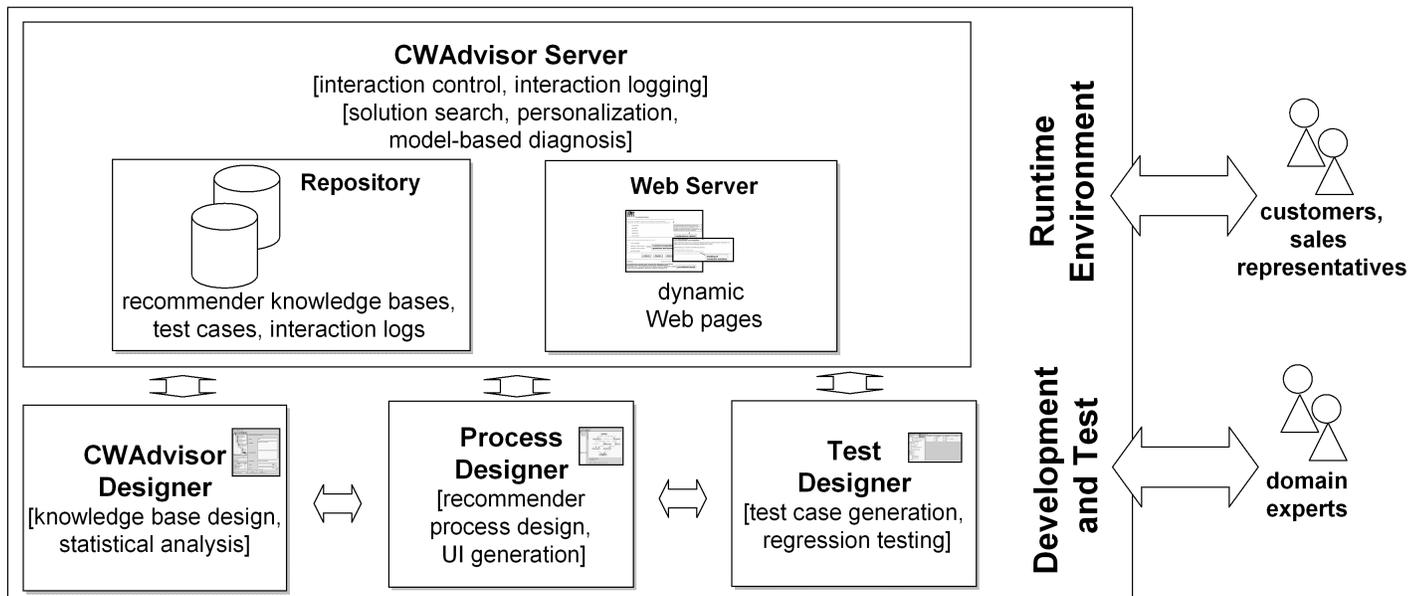


Fig. 1. CWAdvisor architecture.

relational database (Repository). Product structures and instances are defined within CWAdvisor Designer or imported from external systems using an XML interface. CWAdvisor Server supports the execution of advisory processes. Based on given user inputs, the server determines and executes the personalized dialogue flow, initiates result computations and retrieves explanations, such as explanations as to why a certain product suits the needs and wishes of the customer.

B. CWAdvisor Designer and Process Designer

CWAdvisor Designer is a graphical development environment for knowledge-based recommender applications (see Figure 2). It is based on Java Web Start technology which provides an environment for deploying Java-based applications on a Web server and executing those applications on a corresponding client. Using CWAdvisor Designer, the relevant set of product and customer properties and constraints is identified and transformed into the formal representation of a recommender knowledge base [11], [14].

A recommender knowledge base consists of the following parts:

- 1) Customer properties describe possible customer requirements. Customer requirements are collected by posing questions to the customer. The question *under the assumption that your investment of 10,000 EURO decreases in value, at which value would you sell your investment?* is related to the customer property *willingness to take risks (low, medium, high)*. Examples from the digital camera domain are *preferred motif (landscapes, portraits, sport scenes)* or maximum price.
- 2) Product properties describe possible product instances. Life insurances, for instance, can be characterized by the possible *length of life insurance policies, premiums of life insurance policies, or product type*. Furthermore, digital cameras can be described by their *resolution, brand name, or weight* of the camera.
- 3) Compatibility constraints are restricting possible combinations of customer requirements: *return rates above 9% require the willingness to take risks or high resolution cameras cannot be provided in a low-price segment*. Furthermore, filter constraints define the relationship between customer

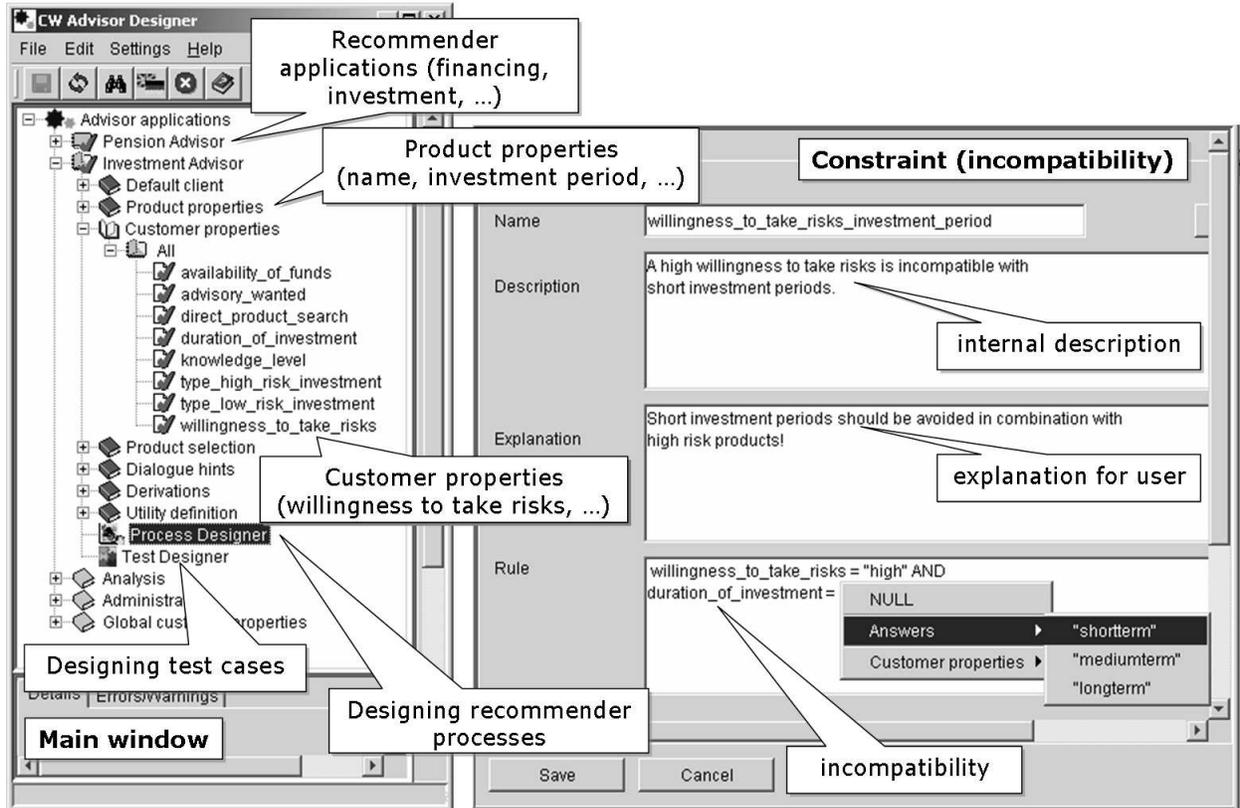


Fig. 2. Definition of customer properties and constraints in CWAdvisor Designer.

properties and product properties: *customers without experiences in the financial services domain should not receive recommendations which include high risk products*. It is important to note that compatibility as well as filter constraints can also be represented on a tabular level.

Apart from the description of the offered products, customer properties and constraints, a process definition is needed to complete the model of a recommender application (see Figure 3).

A recommender process definition represents possible navigation paths which define the way the system adapts the formulation of questions, explanations, error messages, etc. to the knowledge level and interests of the customer. Such process definitions are based on the formalism of a predicate augmented finite state recognizer (PFSR) [16]. In such a PFSR, constraints describe transitions between different states (see Figure 3). Based on layout template definitions [22], knowledge bases and

process definitions can be automatically translated into an executable recommender application. Figure 4 depicts example screenshots of a financial service recommender application. First, the advisor poses questions to the customer (a). Answers provided by the customer (requirements) serve as input for the calculation of a solution (b). If no solution can be found, alternative repair actions are presented to the customer (c). Selecting such a repair action guarantees the identification of a solution (see Section III).

C. CWAdvisor Test Designer

The increasing complexity of recommender knowledge bases makes quality assurance a critical task [26], [34]. Mechanisms have to be provided which allow us to guarantee the correctness of calculated recommendations. In order to avoid situations in which inconsistencies are detected by a customer who interacts with the recommender application, the recommender development

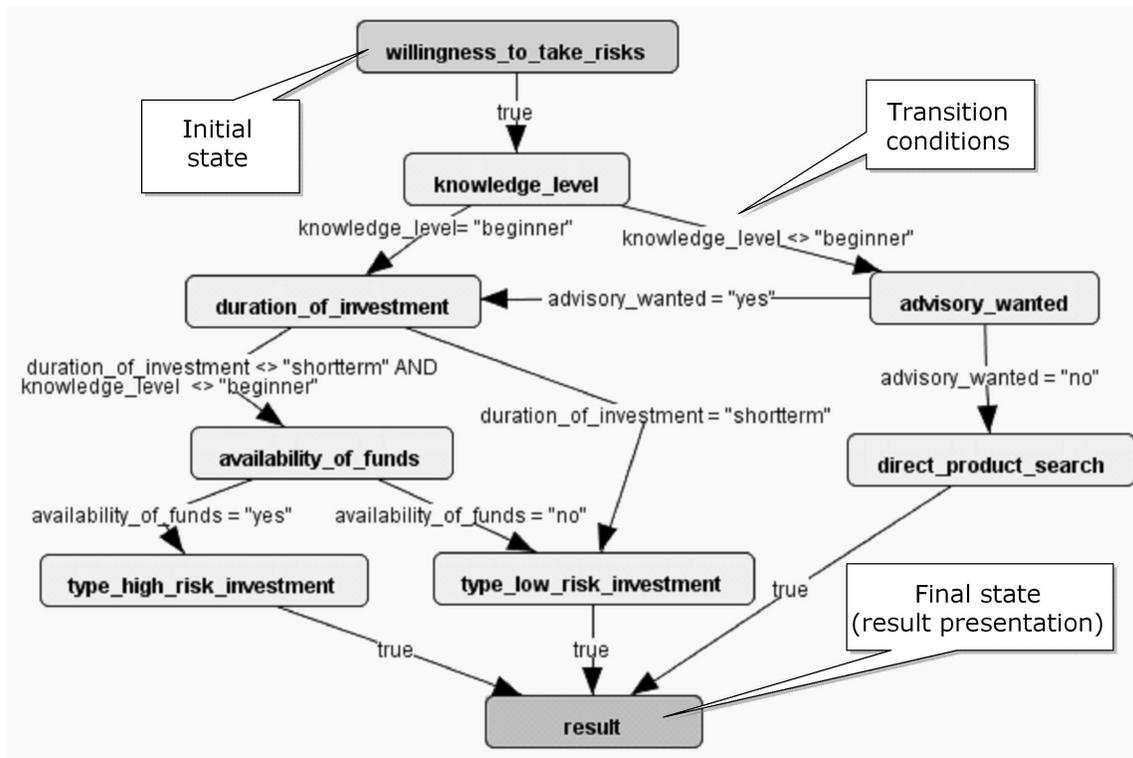


Fig. 3. Recommender process definition.

environment must indicate potential sources of inconsistencies as early as possible. Therefore, automated testing functionalities have been integrated into the CWAdvisor development environment. In Test Designer, process definitions (see Section II.B) are the basis for automatically generating test cases for recommender knowledge bases. A test case is the combination of an input sequence and a set of corresponding results calculated by the current version of the recommender knowledge base. An input sequence represents a set of customer requirements (see Section II) for which a corresponding result has to be calculated by the knowledge base. Typically, domain experts are responsible for checking the correctness of results calculated for a given set of input sequences. This validation step is necessary in order to make the test cases applicable for regression testing. Figure 5 depicts a simple example of the generation of input sequences. The basis for the generation of input sequences is a set of paths which have been selected from a given recommender process definition (in our case only one

path is selected). For each selected path, we can derive a query which generates a set of input sequences. For the purpose of automated input sequence generation, customer properties are represented as tables, their possible instantiations are represented as tuples of the table (see Figure 5). If we execute the query of Figure 5 by only using path transition conditions as selection criteria, we receive 12 different input sequences as query result (see Figure 5). The following concepts allow us to further reduce the number of input sequences:

- Equivalence partitioning. Variable domains can be split up into equivalence classes out of which we can select a representative subset of input values. When selecting equivalence classes, the interchangeability of the values within a class needs to be taken into account because different variable values inside a class should not change a corresponding

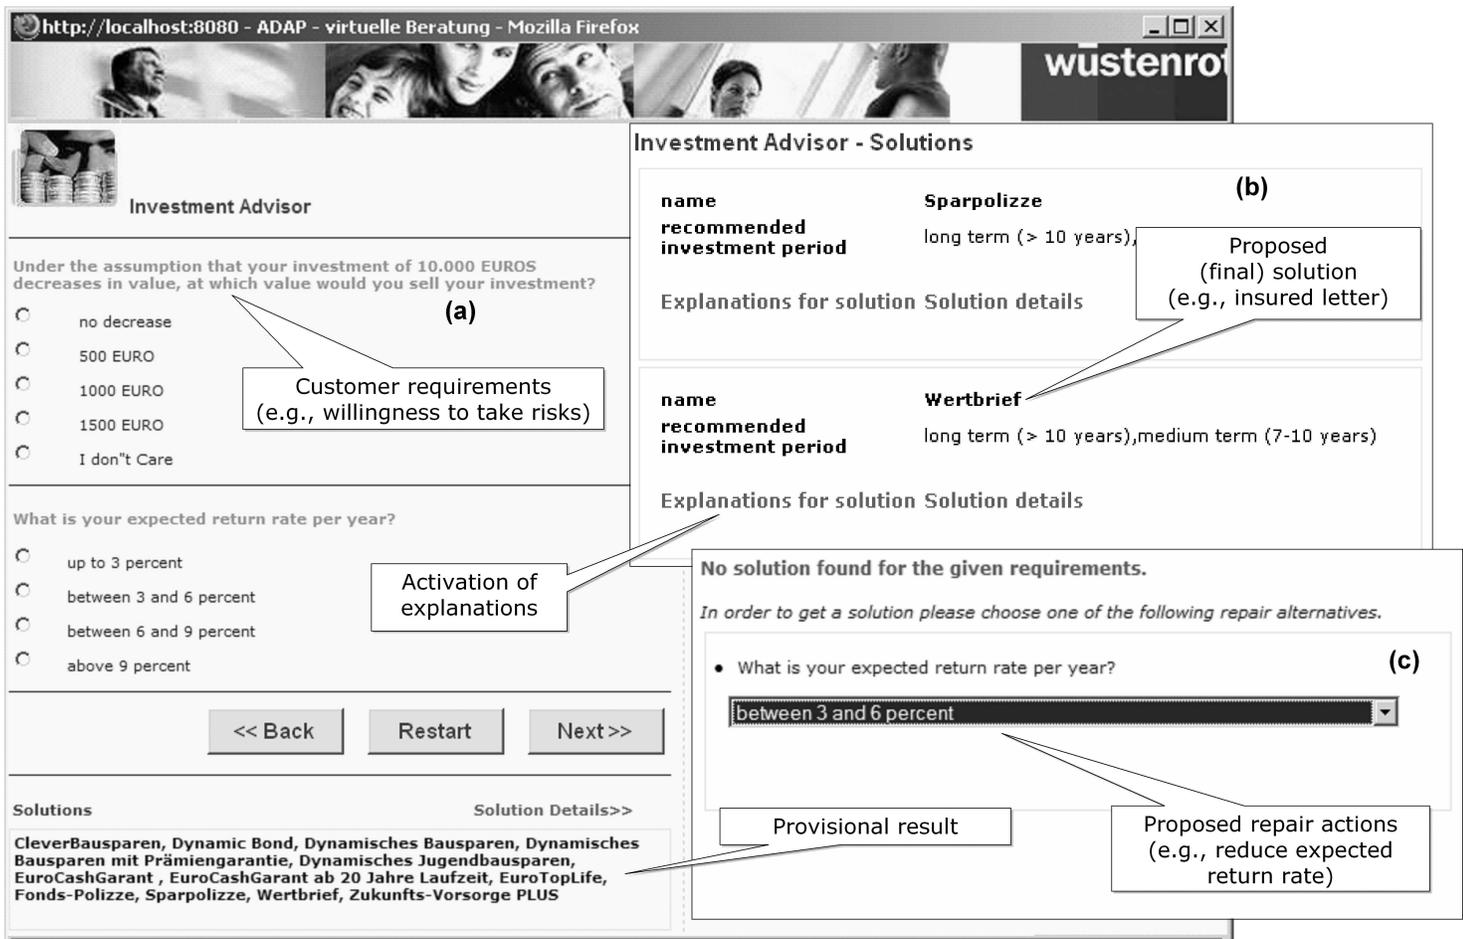


Fig. 4. Example user interface of a financial service recommender application.

result.¹ In the example of Figure 5 we assume that the domain expert selects *knowledge level (average, expert)* as single equivalence class and *knowledge level = average* as representative input value. This input value is included as additional restriction criterion in the input sequence generation query of Figure 5.

- **Constraint Certification.** By the inclusion of incompatibility constraints which have been considered as correct by domain experts (certified constraints), additional input sequences can be eliminated. Those sequences which satisfy the negated incompatibility constraint are filtered out. If the constraint *a*

¹Currently, the selection of equivalence classes is the task of domain experts. In future versions of the test environment we will include mechanisms supporting the automated derivation of equivalence classes.

high willingness to take risks is incompatible with short term investment periods is certified, we can filter out all sequences with related assignments (*willingness to take risks=high* \Rightarrow *duration of investment=shortterm*). If we apply the concepts of constraint certification and equivalence partitioning to our example, the number of relevant input sequences can be reduced to 4 (see Figure 5).

- **Variables with no effects.** In some situations, the advisor asks questions which have no influence on the recommendation. When recommending pension products, a customer may be asked to make a decision concerning returns on investment (annuity payment, singular payment). Since all pension products allow a decision to be made at the end of the investment period, the given answer does not have any effect on the recommendation. Therefore,

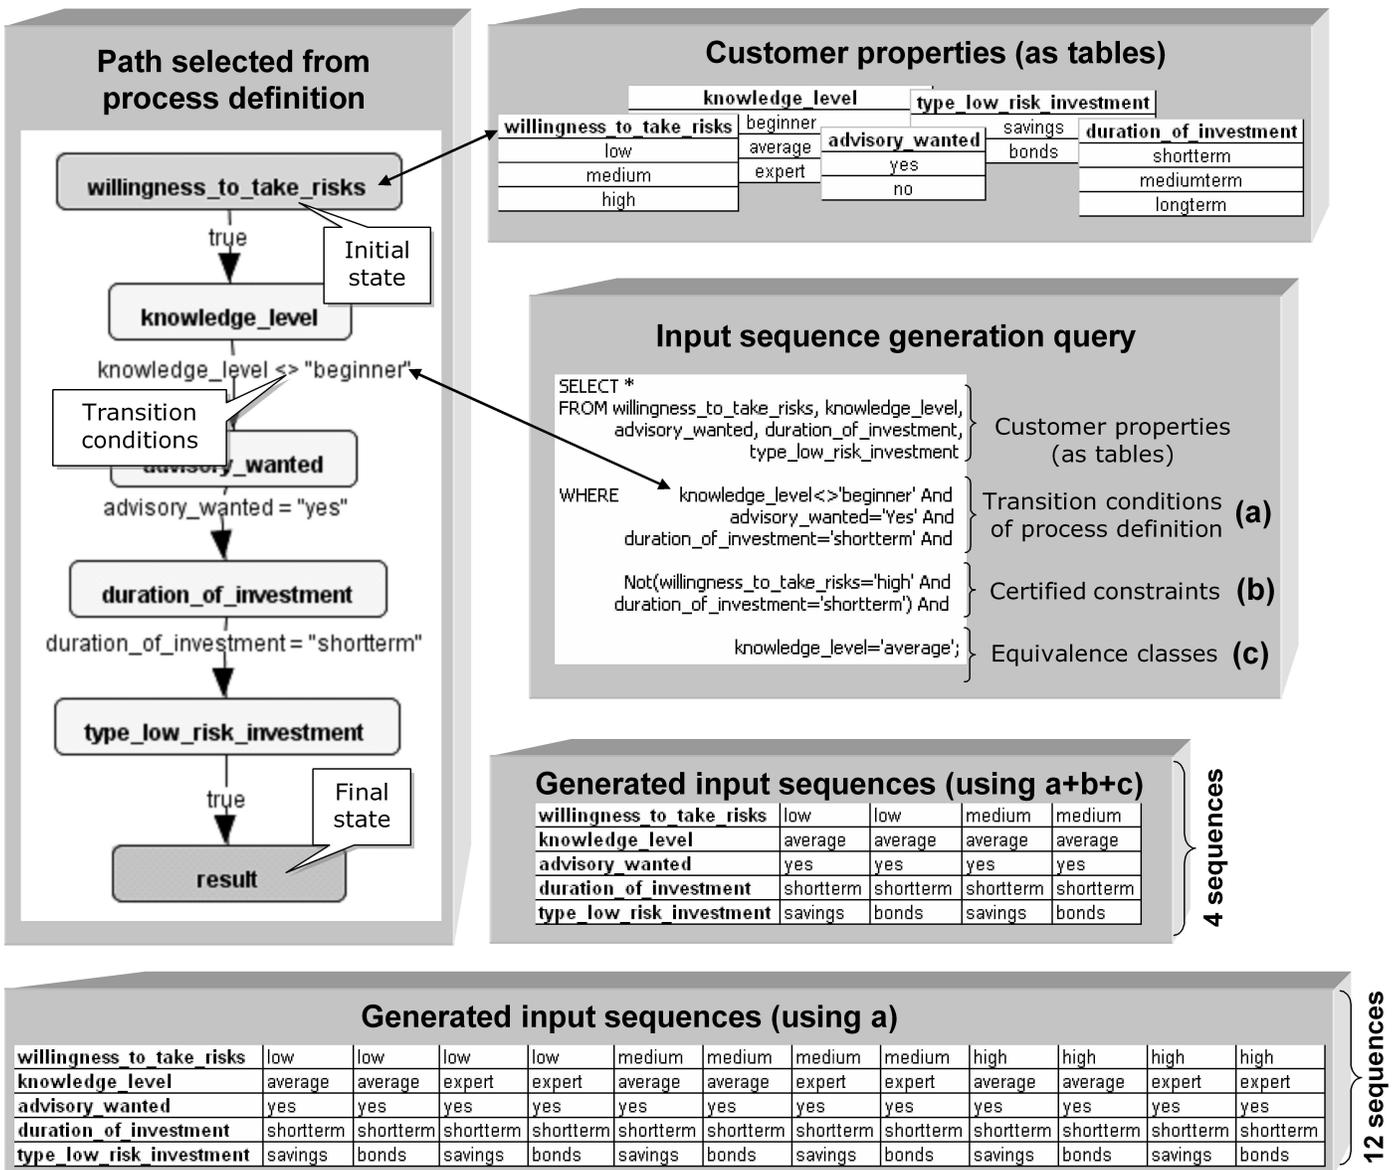


Fig. 5. Generation of input sequences (for test cases).

all values of the corresponding customer property can be assigned to a single equivalence class.

The example depicted in Figure 5 is a simple one. A typical recommender application such as an investment advisor comprises about 20 possible paths. Each path is defined by about 7 variables and 5 possible values per variable. Without including any additional restrictions, the input sequence generation queries for such a setting would generate 1.5mio sequences. This does not allow an effective result validation by domain experts.

However, additional restrictions (transition conditions of paths, equivalence partitioning, constraint certification, variables with no effects) can significantly reduce the number of test cases: our experiences from the financial services domain indicate a potential reduction of 1.5mio test cases to the number of about 500-1000.

III. AI TECHNOLOGIES IN CWADVISOR

Using a knowledge-based approach, the relationship between customer requirements and products can be explicitly modeled [12], [14]. Such a representation is

the precondition for applying the technologies discussed in the following subsections.

A. Representing Recommendation Knowledge and Calculating Solutions

The first step when building an advisor is the construction of a recommender knowledge base which consists of two variable sets (V_C , V_{PROD}) and three sets of constraints (C_R , C_F , C_{PROD}).

- Customer Properties (V_C) describe possible customer requirements, such as *willingness to take risks* (*low*, *medium*, *high*) or *knowledge level* (*beginner*, *average*, *expert*).
- Compatibility Constraints (C_R) are restricting the possible combinations of customer requirements: *not(willingness to take risks = low \wedge expected return rate = >9%)*.
- Product Properties (V_{PROD}) describe possible product instances: *product type* (*savings*, *bond*, *equity fund*) or *risk level* (*low*, *medium*, *high*).
- Filter Constraints (C_F) establish the relationship between customer properties and product properties: *knowledge level = beginner \Rightarrow risk level $\langle \rangle$ high*.
- Allowed instantiations of product properties (offered set of products) are represented by constraints (C_{PROD}) which define restrictions on the possible instantiations of variables in V_{PROD} .

In order to calculate recommendations (solutions), we have implemented a relational query-based approach², in which a set of customer requirements makes a conjunctive query. The query is constructed from the consequent part of those filter constraints of C_F whose condition is consistent with the given customer requirements (active filter constraints). The consequent part of our example filter constraint *knowledge level = beginner \Rightarrow risk level $\langle \rangle$ high* is translated into the expression *risk level $\langle \rangle$ high* as part of the corresponding conjunctive query (in

the case of *knowledge level = beginner* being a customer requirement). Accordingly, V_{PROD} is represented by a set of table attribute definitions (the product table) and C_{PROD} is represented by tuples whose values represent instantiations of the attributes defined in V_{PROD} . Furthermore, customer properties (V_C) are represented as input variables where the compatibility (C_R) of the corresponding instantiations is ensured by a consistency checker. The execution of the conjunctive query on a product table results in a set of recommendations which are presented to the customer.

Relaxing Filter Constraints. Filter-based approaches have well-known limitations. In the case of inconsistent customer requirements, all products in the catalog are filtered out and no recommendation can be given. This problem is addressed in CWAdvisor using the concept of query relaxations in the sense of [19], [29], [30]. If none of the available products fulfills the conjunctive query that is constructed from active filter constraints, we trigger the calculation of a Maximum Succeeding Subquery (XSS): individual atoms are deleted from the conjunctive query in order to identify products that fulfill as many of the active filter constraints as possible. CWAdvisor implements a conflict-directed approach to the calculation of minimal relaxations of a given query. The computation of conflicts is based on the QuickXPlain algorithm [25] which solves a conflict detection problem using a divide and conquer strategy. Preferences related to the elimination of filter constraints (priorities of filter constraints) can be defined a priori by domain experts. They know that most digital camera users would rather settle on a camera of a brand that is not their first choice than accept a camera with a different resolution. The second approach for determining preferences is to directly ask the customer during an advisory session. XSS calculation is based on the resolution of minimal conflicts [25] detected in a conjunctive query (implemented as a version of the standard algorithm for calculating hitting sets [35]). All of our advisory applications follow the strategy of immediately

²A recommendation task can be represented as a Constraint Satisfaction Problem (CSP) as well [14], [45].

computing an (optimal) relaxation when no product fulfills all customer requirements. An optimal relaxation can be determined by selecting those filter constraints with the minimum relaxation costs. Relaxation costs are determined by a corresponding cost function [23]. A simple cost function is, for example, the sum of filter priorities in a relaxation candidate, where higher values represent higher priorities.

Repairing Customer Requirements. There are domains in which filter relaxation is not possible or desirable. In many cases, filter constraints in the financial services domain cannot be relaxed for reasons such as legal regulations or quality aspects. Simply reporting retrieval failures without making further suggestions of how to recover from such a situation is not acceptable, however. Therefore, we need to find out which requirements the customer is willing to change. We aim to find possible compromises (repair actions) that can be presented to the customer. Similar to the computation of filter relaxations, the identification of a minimal set of repair actions is based on the calculation of hitting sets [11], [35]. The goal is to identify a minimal set of variable settings in a set of requirements that should be changed in order to find a solution. A simple example of the calculation of repair actions is shown in Figure 6.

In this example, $S \cup C_R$ has no solution because $\{r_1, r_2\} \cup C_R$ and $\{r_1, r_3\} \cup C_R$ are inconsistent and therefore both $\{r_1, r_2\}$ and $\{r_1, r_3\}$ induce a conflict [25] with the given compatibility constraints. Based on the hitting set algorithm [11], [35], we have to resolve each of the given conflicts: in our example we simply have to change the setting of the expected return rate. Consequently, one possible repair for S is to change the requirement r_1 : *expected return rate* \Rightarrow 9% to r_1 : *expected return rate* = 7-9%. This makes $S' \cup C_R$ consistent ($S' = \{r_1$: *expected return rate* = 7-9%, r_2 : *willingness to take risks* = low, r_3 : *preferred investment period* = *shortterm*}). Another repair alternative would be $S' = \{r_1$: *expected return rate* \Rightarrow 9%, r_2 : *willingness to take risks* = *medium*, r_3 : *preferred investment period* = *mediumterm*}).

B. Personalization

Personalized Repair Proposals. In cases where no solution can be found for a given set of customer requirements, CWAdvisor calculates a set of possible repair actions which allow the calculation of a solution. Different customer properties (variables) have an assigned priority (represented by a utility value) which indicates the utility of the variable for the customer. The lower the utility of the variable the higher is the probability that the variable is considered as the focus of repair actions, e.g., if the return rate is less important for a customer, this property is primarily considered as a potential candidate for repair actions, i.e., repair actions are adapted to the customer's preferences. Table I and Table II depict a simple example for personalizing repair actions depending on defined utilities for variables (customer requirements). Table I shows the utility of the variables *rr* (*expected return rate*), *ip* (*preferred investment period*), and *wr* (*willingness to take risks*) for customer₁ and customer₂, e.g., the utility of the variable *rr* is higher for customer₂. Table II depicts the application of the different utility values to proposed repair actions (repair₁, repair₂). The sum of the utility values of repair₂ for customer₁ is 7 (utility(*ip*) + utility(*wr*)), of repair₂ for customer₂ it is 3. Consequently, repair₁ is proposed to customer₁, since it has the lowest utility for customer₁. Vice versa, repair₂ is proposed to customer₂.

Dialogue Control. Customers have different preferences of how to specify their requirements: from the direct specification of product parameters (such as a certain savings account running for three years) to a general specification of their personal goals (for example, financing their childrens education). Depending on the answers provided, the dialogue with the customer can be adapted as follows (the knowledge level of a customer is derived either from answers to a set of test questions or from a self-assessment at the beginning of an advisory session):

- Alternative formulation of questions. Questions posed to expert users can be differentiated from

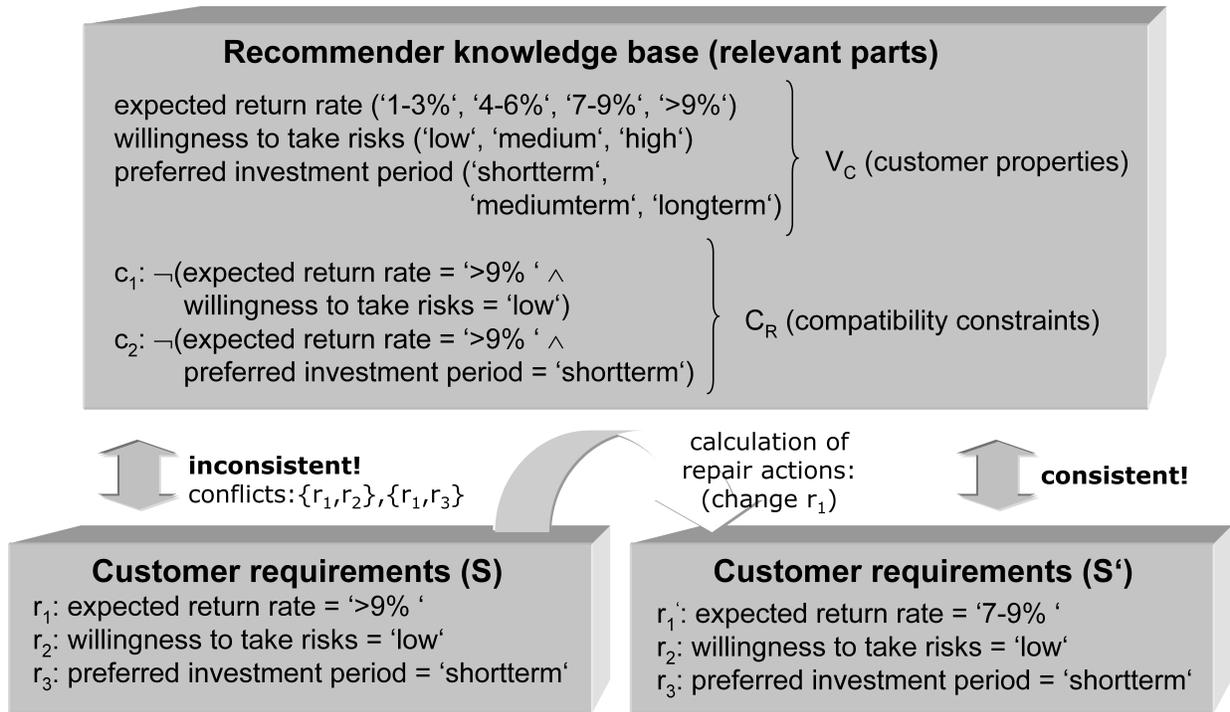


Fig. 6. Simple example for the repair of customer requirements.

customer	utility(rr)	utility(ip)	utility(wr)
customer ₁	2	3	4
customer ₂	5	2	1

TABLE I
 CUSTOMER-INDIVIDUAL UTILITY VALUES FOR REQUIREMENTS.

repair	utility customer ₁	utility customer ₂
repair ₁ : {rr=7-9%}	2	5
repair ₂ : {ip=mediumterm, wr=medium}	7	3

TABLE II
 PERSONALIZED DETERMINATION OF REPAIR ACTIONS.

those posed to customers with less knowledge about the product domain. Questions are defined during design time; they are not dynamically composed during program execution. The selection of questions strictly depends on the process definition (see Section II).

- Rule-based formulation of default answers. If the goal of the customer is to put money away for a rainy day the default answer to a question related

to the maximum accepted decrease in value of the investment is *no value decrease* accepted.

- Alternative explanations for constraint violations. If the customer has no detailed knowledge about the product domain (novice), a very general explanation about changes in the pension law is given. More detailed information can be included for experts.

Utility/Ranking of Solutions. A solution for a recommendation task is a set of products or services. The order

of products within a solution should strictly reflect the degree the products correspond to the wishes of a customer. CWAdvisor supports multi-attribute object rating [2], where each product is evaluated according to a predefined set of interest dimensions. *Profit, availability* and *risk* are examples for such dimensions in the financial services domain. If a customer is strongly interested in products with high return rates, compared to availability and risk, profit is a very important dimension. Consequently, the order of products depends on the importance the individual dimensions have for the customer. The set of products that are part of a solution is ordered using the formula $g(x) = \sum_{(i=1..n)} e_i s_i(x)$, where n denotes the number of dimensions, $g(x)$ represents the utility of a product x , e_i represents the interest of the customer in dimension i , and $s_i(x)$ is the contribution of product x to dimension i . Values for e_i can be automatically derived from customer requirements on the basis of scoring rules. The score 9 for the dimension profit in Table III can be derived from the customer requirement expected return rate = >9% where the corresponding scoring rule for the dimension profit can be defined as (return rate = >9% \Rightarrow 9).³ Another alternative is to weigh the interest dimensions conforming to an assessment provided by the customer. Values for $s_i(x)$ have to be defined at design time and manually specified for each product. Financial services can be ranked as follows (see Table III)⁴: for customer₁, $g(\text{savings}) = 9*1 + 4*8 + 7*8 = 97$, whereas for customer₂, $g(\text{savings}) = 6*1 + 5*8 + 1*8 = 54$. The utility of savings is higher for customer₁ than for customer₂.

Presentation of Solutions. For each solution, a set of immediate explanations [18] is calculated. These explanations are derived from active filter constraints. An explanation related to our example filter condition *knowledge level = beginner \Rightarrow risk level $\langle \rangle$ high* could be *this product assures adequate return rates with a lower level of related risks*. Consequently, explanations

³Note that the score related to an interest dimension (e.g., profit) is not necessarily based on a single requirement.

⁴In this example, scores are taken from a scale between 0 and 10.

related to filter conditions are predefined at design time.

Product Comparison. Product comparisons provide basic mechanisms to compare different products that are part of a recommendation result. Product comparison is based on the definition of rules stating under which conditions an argumentation for/against a certain product should be displayed. If the price of product A is significantly higher than the price of product B, the product comparison component should display a corresponding hint. In this context, the level of significance has to be defined by the domain expert. Product comparisons are currently used in our digital camera recommender application of geizhals.at (see Section IV).

Handling of Profiles. The user interface relies on a user model that describes capabilities and preferences of individual customers. Some of these properties are directly provided by the customer (e.g., name or personal goals). Other properties are derived using scoring mechanisms which relate user answers to abstract dimensions [2] such as preparedness to take risks or knowledge about financial services. This information is stored in the profile and can be exploited in future advisory sessions.

IV. EXPERIENCES FROM INDUSTRIAL PROJECTS

A. Digital Camera Advisor

In the domain of consumer electronics, customers are typically confronted with a large variety of products. In our case (digital cameras of the price comparison platform geizhals.at), customers can choose among more than 600 different models. The platform geizhals.at informs approximately one million unique clients per month about the best price of products from domains such as computer hardware, consumer electronics or household machines (about 85,000 products offered by more than 700 online shops). Customers can navigate in a product catalog which is organized in hierarchical categories. For each product, an ordered list of prices and additional user ratings of the corresponding online shops are presented to the customer. The business model of geizhals.at is built on paid advertisements and pay-

product	profit	availability	risk	dimension	customer ₁	customer ₂
savings	1	8	8	profit	9	6
bonds	4	2	2	availability	4	5
equity funds	9	2	0	risk	7	1

TABLE III
UTILITY OF SOLUTIONS (OBJECT RATINGS AND CUSTOMER PREFERENCES).

per-click commissions of the enlisted online shops. The digital camera advisor project was not only initiated to provide a better service for experts, but also to support new user groups with less product domain knowledge. Therefore, the dialogue includes interaction paths that allow the specification of technical features such as optical zoom as well as answers to problem-oriented questions such as *what is your preferred motif?* The digicam advisor was introduced in 2003. Since then, more than 200,000 users have successfully completed their advisory sessions.

Our evaluation of the digital camera advisor consisted of $n=1,253$ online users. An announced lottery ensured that participants identified themselves with their genuine names and addresses and no duplicate questionnaires were counted. The sample consisted of arbitrary online users of the digital camera section of geizhals.at, who did not necessarily buy a digital camera. The goal of the questionnaire was to evaluate to what extent the advisor helped users to find what they were looking for and which of the provided advisory features were most useful. Therefore, the key hypothesis of the evaluation was that advisor applications help users to better orientate themselves when being confronted with a large product assortment. We asked users whether they had noticed and used the advisor, and whether they had found the product they were looking for. The answers verified our hypothesis: a significantly higher share of users had successfully completed their product search among those who employed the advisor ($\chi^2(1) = 9.39$; $p < 0.01$) which demonstrates an interdependence between the observed variables (see Table IV). Furthermore, we wanted to know which features of a recommender application users

especially liked. Interviewees were asked to rate their subjective benefit from the different features of the application. Product comparisons on the result page and the easy and quick way to access the right products were rated highest (see Table V). None of the results changed significantly when allocating the sample by gender, age or technical expertise. A study on consumer behavior in the interaction with knowledge-based recommender applications confirmed above results (for details see [13]). This study was conducted within the scope of the COHAVE project⁵. Examples of results of this study are the following:

- Advisors supporting product comparisons clearly outperform advisors without this functionality (trust in recommended products was one of the relevant features herein). The reason for this lies in the lower mental workload of users when differences between alternatives are clearly presented.
- Knowledge-based advisors clearly outperform simple product lists regarding dimensions such as overall satisfaction with the advisory support or trust in recommended products.
- Finally, explanations of recommendation results significantly contribute to an increased perceived conformance between expected and recommended products.

B. Financial Services Advisor

CWAdvisor has been deployed for 1,400 sales representatives (insurance brokers and employees of the company) of one of the largest Austrian building and

⁵Consumer Buying Behavior & Decision Modeling for Recommender Systems (Austrian Research Fund 810996).

Found the right products	Did use advisor	Did not use advisor
Yes	276 (76%)	604 (67%)
No	85 (24%)	288 (33%)
Total	361	892

TABLE IV

FEEDBACK OF USERS ACCESSING THE DIGITAL CAMERA SECTION OF GEIZHALS.AT.

Application features	(very) high usefulness
Product comparisons	74.4%
Easy and quick access to the <i>right</i> products	72.2 %
Explanation of product properties	57.6%

TABLE V

CWADVISOR FUNCTIONALITIES RATED *high* OR *very high* REGARDING USEFULNESS.

loan associations [14]⁶. In this context, financial service recommender applications have been integrated into an existing Customer Relationship Management (CRM) environment that supports sales representatives by providing functionalities such as calculations, presentation of technical product details, quotation generation or scheduling of meetings. The recommender applications support sales representatives in the preparation, conduction, and summary of sales dialogues. In this context, they help to explain recommendation results. The motivation for the development of knowledge-based financial services recommender applications was twofold. First, time effort related to the preparation, conduction and completion of sales dialogues should be reduced. Second, an automated explanation of advisory results should be supported in order to take into account regulations of the European Union [47] relating to an improved documentation of advisory sessions in the financial services domain. On average, a sales representative sells about 60-70 products per year, experts sell up to 500. One and a half years after the initial deployment of the advisor applications we interviewed sales representatives (n=52) with the goal to evaluate the acceptance of the provided advisory applications and

to detect potentials for further improvements. The most important topics/questions of the survey are contained in Table VI.

Expertise level. About 23.1% of the interviewees were beginners (focus is on selling 1-3 different types of products), 28.9% were advanced sales representatives (selling more than 3 different types of products) and about 48.0% were experts (selling all product types).

Importance of advisory support. The interviewees strongly agreed on the importance of the provided advisory functionality (67.3% high importance, 21.1% average importance). Major arguments for the application of knowledge-based recommender technologies were (a) automated consistency checks which improve the correctness of offers (this was significantly more important for beginners and representatives selling less than 50 products per year), and (b) the automated generation of advisory protocols which allow for time savings in the introductory dialogue as well as in the quotation phase.

Major application area. Most of the sales representatives (81.8%) reported that they used CWAdvisor functionalities throughout the sales dialogue or for the generation of documentations for completed advisory dialogues (protocols). Each such protocol includes a summary of customer preferences, a listing of recommended products, an argumentation as to why the

⁶See Figure 4 for an example screenshot of a deployed investment recommender application.

recommendations suit the needs and wishes of the customer (explanations), and a listing of further issues to be discussed with the customer. Furthermore, 53.8% of the sales representatives reported that they also used the advisors at home in order to prepare for a sales dialogue.

Time savings. On average, the time savings caused by the application of financial service advisors were specified with 11.73 minutes per advisory session. Most savings are explained by the generation of advisory protocols at the end of a session and available summaries of previous dialogues at the beginning of a new session. Assuming that an average sales representative conducts about 170 advisory sessions per year, time savings can amount up to 33 hours per year for each person. For a sales expert, time savings can amount up to 100 hours per year.

Importance for new representatives. Most of the sales representatives (97%) definitely agreed on the potential of knowledge-based recommender technologies to provide e-learning support. Due to this feedback, a corresponding project has already been initiated which exploits recommender technologies in order to support learning processes for new sales representatives. The software will be applied in the context of sales courses.

C. The Role of Knowledge Acquisition in Recommender Projects

Due to a lack of programming skills [21], there is a significant discrepancy between technical experts who are able to create software artifacts and domain experts who are not. Consequently, domain experts without technical knowledge are in the position of being solely responsible for providing domain knowledge and knowledge engineers transform this knowledge into the formal representation of an underlying knowledge base. This process is quite error-prone and is also known as the knowledge acquisition bottleneck. In order to reduce this bottleneck, our goal was to empower autonomous development and maintenance processes for end users. Such an integration of end users has great

potential to significantly improve the effectiveness of recommender application development. Within the scope of the development of the CWAdvisor environment we took into account the following principles which are crucial for supporting effective knowledge acquisition and maintenance processes [7], [8]:

- Rapid prototyping follows the principle of concreteness. The user immediately sees the effects of changing explanation texts, product properties, images, transition conditions, etc. in the recommender knowledge base or process definition. In CWAdvisor, such changes are performed on a graphical level (see Section II); the corresponding recommender applications can be automatically generated from graphical representations. Such an automated application generation is extremely important to make the development of advisors feasible for domain experts (end users) without a special IT education [4].
- A flexible and simple graphical design of the intended behavior of the user interface protects the user from programming details. CWAdvisor provides simple modeling concepts that allow the explicit specification of the intended behavior of the recommender user interface (principle of explicitness). This approach also follows the principle of separating application logic design from implementation details.
- Effective test support follows the principle of immediate feedback. Potential sources of inconsistencies in the knowledge base caused by the creation of a new knowledge base version are indicated promptly. Especially automated test case generation and regression testing have been identified as prerequisites for effective end-user debugging [38]. Following this approach, errors in the recommender knowledge base can be detected immediately and are not forwarded to the productive environment where, in the worst case, those errors are detected by customers. Thus, domain experts can establish a

Topic	Description
Expertise level	Expertise level of the sales representative
Importance of advisory support	Importance of advisory support for sales representatives
Major application area	Major application area of CWAdvisor
Time savings	Time savings in sales dialogues related to CWAdvisor
Importance for new representatives	Importance for educating sales representatives

TABLE VI
MAJOR TOPICS/QUESTIONS OF THE QUESTIONNAIRE.

higher level of trust in recommendations, which is extremely important for our projects.

Our experiences show that all these principles are important for the development and maintenance of recommender knowledge bases. We compared time efforts needed for the development of advisors before and after CWAdvisor was available. Due to the graphical development support, the overall development efforts were reduced significantly. In the case of our financial services advisor project, three domain experts are responsible for the development of knowledge bases (investment & financing, pension & life insurance, property insurances). The first knowledge base versions were developed in cooperation with knowledge engineers. Now, most of the maintenance activities, including updates of products, (in)compatibility tables, filter conditions, and process definitions, are conducted by domain experts. The time efforts related to knowledge base maintenance done by knowledge engineers are now negligible. Changes to current knowledge base versions are conducted in a test environment. After the completion of changes in the knowledge base, a set of pre-generated test cases (see Section II) is used for regression testing. Thereafter, the new version of the knowledge base is synchronized with the productive environment (synchronization with the client applications of sales representatives). In the case of geizhals.at, knowledge engineers who are knowledgeable about digital cameras are responsible for maintaining the knowledge base. In this context, knowledge bases must be designed as generic as possible in order to be stable regarding the frequent changes in the product assortment

(technical information such as memory size or resolution is constantly changing).

V. RELATED WORK

Recommender Technologies. Over the last years, several technologies supporting personalized product recommendations have been successfully applied in the context of commercial systems (see, e.g., the online store of www.amazon.com). Recommender technologies relieve customers and sales representatives from a tedious search within a potentially large set of complex products and services. An overview of different applications of recommender technologies can be found in [31], [39], overviews of different technological approaches to the implementation of recommender applications can be found in [1], [5], [44]. Basically, there are three widely known technological approaches to the implementation of recommender applications. Collaborative filtering systems [20] automate the recommendation process on the basis of user opinions about items and generate new recommendations based on inter-user comparisons. If two customers have bought similar CDs in the past and have rated those CDs in a similar way, CDs (with a positive rating) bought by only one of them are recommended to the other one. Because recommendations follow from a comparison between the target user and other users, a user with few ratings becomes difficult to categorize which is also known as the new user problem. Similarly, a new item that does not have many ratings also cannot be easily recommended this is known as the new item problem. Content-based filtering [33] is a special type

of information filtering that uses features of items the user has liked in the past to infer new recommendations. In contrast to collaborative approaches, content-based filtering is limited in that it cannot provide serendipitous recommendations, because it selects and recommends all products based on content available from the current user [40]. In the case of collaborative filtering as well as in the case of content-based filtering, user profiles are long-term models. Both approaches [20], [33] do not exploit deep knowledge about the product domain and therefore are excellent techniques supporting recommendation processes for simple products such as books or movies. A major strength of these approaches is that no additional knowledge acquisition and maintenance efforts are needed. Knowledge-based recommender technologies receive increasing attention in research [6], [5], [14], [24], [45]. In contrast to collaborative and content-based filtering approaches, knowledge-based approaches are based on an explicit representation of product, marketing and sales knowledge. Using such representations, new item and new user problems are avoided since recommendations are directly derived from user preferences identified within the scope of the requirements elicitation phase. The main reason for choosing a knowledge-based recommendation approach stems from the requirements of domains such as financial services where deep product knowledge is needed in order to identify and explain solutions. Jiang et al. [24] present an approach to multimedia-enhanced product recommendation where basic recommendation technologies are additionally equipped with a component supporting the visualization of results. Thompson et al. [45] focus on the integration of conversational natural language interfaces with personalized recommender systems with the major goal of reducing system-user interactions. Long-term user preferences are obtained in the course of normal recommendation dialogues and are used to adapt future dialogs with the same user [45]. A study conducted in the restaurant recommendation domain found that users needed fewer interactions to find a good restaurant when

using a personalized conversational recommender than when using a non-personalized one [45]. Both natural language interaction [45] and result visualization [24] are not provided by CWAdvisor but are within the scope of future work. Compared to the approaches of [6], [5], [24], [45], CWAdvisor additionally provides a graphical recommender application development and test environment and supports the calculation of repair actions for customer requirements in situations in which no solution can be found. Both functionalities are extremely important when deploying knowledge-based advisors in industrial environments.

As mentioned in [5], case-based reasoning (CBR) [27], [48] can be interpreted as a specific type of knowledge-based recommendation approach. CBR is based on the concept of solving new problems on the basis of retrieving old problems likely to have similar solutions [5]. Case-based product retrieval searches in a database of product descriptions [3], where the retrieval of product recommendations is based on the evaluation of similarity metrics, i.e., case-based approaches are primarily not based on an explicit definition of a recommender knowledge base. An example for the application of case-based reasoning technologies in the financial services domain (life insurance policies) is presented in [43]. In this context, advanced concepts of case-based reasoning [27] are taking into account generalized cases. A generalized case does not cover only a point of the case space, but a whole subspace of it [43].

Recommender Technologies in the Financial Services Domain. [32] points out the potential of Internet applications to move information technology in the financial sector from a pure transaction-oriented paradigm towards income generating applications. In order to achieve this goal, a flexible IT architecture is needed which allows flexible adaptations to continuously changing products and sales processes. In the context of such an architecture, knowledge-based recommender technologies are a building block supporting the effective

mapping of product, marketing and sales knowledge to the formal representation of a knowledge base.

There already exist a number of approaches to apply knowledge-based technologies for supporting sales processes in the financial services domain. [42] show the application of knowledge-based configuration and personalization technologies for the selection of insurance products, where the integration of feature-based search and interactive product configuration is discussed within the context of application examples. Using retirement planning as application scenario [9] present their financial services advisory system. Similar to CWAdvisor, the system is based on a process-oriented approach, where the user is guided through the whole financial planning process. Although [9] provide a good impression of the underlying advisory process and the requirements related to the representation of customer requirements (hard and soft constraints should be supported), no details are provided about the underlying reasoning mechanisms of the system. [28] motivate the application of knowledge-based configuration technologies for the configuration of insurance services. A meta-model for generic product structures is presented which takes into account structural properties of configurable products as well as additional restrictions related to the combination of different sub-components. Similarly, the meta-model of CWAdvisor supports the representation of product properties and related constraints. Additionally, our meta-model takes into account the design of recommender processes which represent personalized navigation paths in a recommender application. Based on the scenario of composing insurance and investment services with the goal to provide an integrated (configured) solution for young families with children, the concept of customer- and needs-oriented product configuration is presented in [15]. In the context of this scenario different types of constraint representations are discussed (e.g., constraints between customer requirements or constraints establishing the relationship between customer requirements and product properties) these types of constraints are

included in the CWAdvisor environment. [10] discusses the state-of-the art in tool support for financial services selection and points out that, although needed for the improvement of the accessibility of financial services for the average customer, recommender technologies are seldom provided. This insight is based on a study conducted for financial service providers in German-speaking countries. The study has been repeated in the following years one of the major results for the year 2005 was that the task of improving the quality of advisory is one of the major focuses of the interviewed financial services providers.

Knowledge Base Validation. Within the context of knowledge-based systems development, validation technologies have, in many cases, not been adopted by practitioners; ad hoc techniques are still dominating [34]. The literature on test case generation in knowledge-based systems development is still largely directed at rule-based types of systems [34]. It is not directly applicable to knowledge-based recommendation. Tiihonen et al. [46] present an approach to the testing of configurator applications, in which the configuration model is considered as consisting of a set of local requirement groups representing a set of potential inputs provided by the user. A test case is represented by a group of requirement items, and test case generation is based on randomly selecting requirement groups. In contrast to our work, [46] directly deal with the generation of test cases from paronomies, whereas our approach defines test cases on the basis of results of conjunctive queries that are related to paths in recommender process definitions. Furthermore, no mechanisms for reducing the number of test cases are presented in [46]. This is, however, crucial for the validation of knowledge bases by domain experts.

VI. CONCLUSIONS

Effective knowledge acquisition processes and intuitive interaction mechanisms that support the user in the product and service selection process are major preconditions for successful recommender applications. Both

aspects are supported by the recommender technologies discussed in this paper. First, the CWAdvisor knowledge acquisition environment provides innovative development, maintenance and testing techniques for recommender knowledge bases. Second, model-based diagnosis techniques actively support customers in situations in which no solution could be found. Our goal for the future is to further increase the acceptance of knowledge-based recommender technologies. In order to reach this goal, we will integrate psychological theories from the area of consumer buying behavior into the design of our recommender environment. Following this strategy, CWAdvisor technologies continuously improve recommender application development processes and contribute to an improved accessibility of complex product and service assortments.

VII. ACKNOWLEDGEMENTS

Parts of the work presented in this paper have been done within the scope of the project Koba4MS (Knowledge-based Advisors for Marketing and Sales FFG-808479).

REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. Toward 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):734–749, 2005.
- [2] L. Ardissano, A. Felfernig, G. Friedrich, D. Jannach, G. Petrone, R. Schäfer, and M. Zanker. A Framework for the development of personalized, distributed web-based configuration systems. *AI Magazine*, 24(3):93–108, 2003.
- [3] R. Bergmann, R. Traphoner, R. Schmitt, P. Cunningham, and B. Smyth. Knowledge-intensive Product Search and Customisation in Electronic Commerce. *E-Business Applications*, pages 89–103, 2002.
- [4] H. Bhargava, S. Sridhar, and C. Herrick. Beyond spreadsheets: tools for building decision support systems. *IEEE Computer*, 32(3):31–39, 1999.
- [5] R. Burke. Knowledge-based Recommender Systems. *Encyclopedia of Library and Information Systems*, 69(32), 2000.
- [6] R. Burke, K. Hammond, and B. Young. The FindMe Approach to Assisted Browsing. *IEEE Expert*, 12(4):32–40, 1997.
- [7] M. Burnett. HCI research regarding end-user requirement specification: a tutorial. *Knowledge-based Systems*, 16:341–349, 2003.
- [8] C. DiGiano, K. Kahn, A. Cypher, and D. Smith. Integrating learning supports into the design of visual programming systems. *Visual Languages and Computing*, 12:501–524, 2001.
- [9] C. Dziarstek, F. Farnschläder, S. Gilleßen, I. Süßmilch-Walther, and V. Winkler. A User-Aware Financial Advisory System. In *Multikonferenz Wirtschaftsinformatik (MKWI04)*, pages 217–229, Duisburg-Essen, Germany, 2004.
- [10] V. Eckert-Niemeyer. Innovative Tools zur Realisierung der virtuellen Beratung. *Banking and Information Technology*, 1(1):23–31, 2000.
- [11] A. Felfernig, G. Friedrich, D. Jannach, and M. Stumptner. Consistency-based Diagnosis of Configuration Knowledge Bases. *AI Journal*, 2(152):213–234, 2004.
- [12] A. Felfernig, G. Friedrich, D. Jannach, M. Stumptner, and M. Zanker. Configuration knowledge representations for Semantic Web applications. *AI Engineering Design, Analysis and Manufacturing Journal*, 17:31–50, 2003.
- [13] A. Felfernig and B. Gula. An Empirical Study on Consumer Behavior in the Interaction with Knowledge-based Recommender Applications. In *IEEE Conference on e-Commerce Technology (CEC06)*, page to appear, San Francisco, USA, 2006.
- [14] A. Felfernig and A. Kiener. Knowledge-based Interactive Selling of Financial Services using FSAdvisor. In *17th Innovative Applications of Artificial Intelligence Conference (IAAI'05)*, pages 1475–1482, Pittsburgh, Pennsylvania, 2005.
- [15] A. Felfernig, J. Mehla, A. Wimmer, M. Zanker, and C. Russ. Konzepte zur flexiblen Konfiguration von Finanzdienstleistungen. *Banking and Information Technology, Sonderheft zur Multi-Konferenz Wirtschaftsinformatik 2004*, 5(1):5–17, 2004.
- [16] A. Felfernig and K. Shchekotykhin. Debugging User Interface Descriptions of Knowledge-based Recommender Applications. In C. Paris and C. Sidner, editors, *ACM International Conference on Intelligent User Interfaces (IUI'06)*, pages 234–241, Sydney, Australia, 2006.
- [17] G. Fleischanderl, G. Friedrich, A. Haselböck, H. Schreiner, and M. Stumptner. Configuring Large Systems Using Generative Constraint Satisfaction. *IEEE Intelligent Systems*, 13(4):59–68, 1998.
- [18] G. Friedrich. Elimination of Spurious Explanations. In G. Müller and K. Lin, editors, *16th European Conference on AI (ECAI 2004)*, pages 813–817, Valencia, Spain, 2004.
- [19] P. Godfrey. Minimization in Cooperative Response to Failing Database Queries. *International Journal of Cooperative Information Systems*, 6(2):95–149, 1997.

- [20] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. Riedl. Evaluating Collaborative Filtering Recommender Systems. *ACM Trans. on Inf. Systems*, 22(1):5–53, 2004.
- [21] A. Ioannidou. Programmorphosis: a Knowledge-Based Approach to End-User Programming. In *Interact 2003*, pages 152–159, Zuerich, Switzerland, 2003.
- [22] D. Jannach and G. Kreutler. A Knowledge-Based Framework for the Rapid Development of Conversational Recommenders. In X. Zhou, M. Papazoglou S. Su, M. Orłowska, and K. Jeffery, editors, *Web Information Systems - WISE 2004*, pages 390–402. Springer LNCS 3306, 2004.
- [23] D. Jannach and J. Liegl. Conflict-directed Relaxation of Constraints in Content-based Recommender Systems. In *IEAAIE*, pages 819–829, Annency, France, 2006.
- [24] B. Jiang, W. Wang, and I. Benbasat. Multimedia-Based Interactive Advising Technology for Online Consumer Decision Support. *Comm. of the ACM*, 48(9):93–98, 2005.
- [25] U. Junker. QUICKXPLAIN: Preferred Explanations and Relaxations for Over-Constrained Problems. *19th National Conf. on AI (AAAI04)*, pages 167–172, 2004.
- [26] S. H. Kirani, I. A. Zualkernan, and W. T. Tsai. Evaluation of Expert System Testing Methods. *Communications of the ACM*, 37(11):71–81, 1994.
- [27] J. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, 1993.
- [28] S. Leist and R. Winter. Konfiguration von Versicherungsleistungen. *Wirtschaftsinformatik*, 36(1):45–56, 1994.
- [29] D. McSherry. Maximally Successful Relaxations of Unsuccessful Queries. In *15th Conference on Artificial Intelligence and Cognitive Science*, pages 127–136, 2004.
- [30] N. Mirzadeh, F. Ricci, and M. Bansal. Supporting User Query Relaxation in a Recommender System. In *EC-WEB*, pages 31–40, Zaragoza, Spain, 2004.
- [31] M. Montaner, B. Lopez, and J. De la Rose. A Taxonomy of Recommender Agents on the Internet. *Artificial Intelligence Review*, 19:285–330, 2003.
- [32] E. Papathanassiou. Mass customization: management approaches and internet opportunities in the financial sector in the UK. *International Journal of Information Management*, 24:387–399, 2004.
- [33] M. Pazzani and D. Billsus. Learning and Revising User Profiles: The Identification of Interesting Web Sites. *Machine Learning*, (27):313–331, 1997.
- [34] A. Preece, S. Talbot, and L. Vignollet. Evaluation of Verification Tools for Knowledge-Based Systems. *International Journal of Human-Computer Studies*, 47(3):629–658, 1997.
- [35] R. Reiter. A theory of diagnosis from first principles. *AI Journal*, 23(1):57–95, 1987.
- [36] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *ACM Conf. on Computer Supported Cooperative Work*, pages 175–186, 1994.
- [37] F. Ricci, A. Venturini, D. Cavada, N. Mirzadeh, D. Blaas, and M. Nones. Product Recommendation with Interactive Query Management and Twofold Similarity. In *5th International Conference on Case-Based Reasoning (ICCBR 2003)*, pages 479–493, Trondheim, Norway, 2003.
- [38] J. Ruthruff and M. Burnett. Six Challenges in Supporting End-User Debugging. In *Workshop on End-User Software Engineering*, Saint Louis, Missouri, USA, 2005.
- [39] J. Schafer, J. Konstan, and J. Riedl. Electronic Commerce Recommender Applications. *Journal of Data Mining and Knowledge Discovery*, 5(1/2):115–152, 2000.
- [40] Y. Shih and D. Liu. Hybrid recommendation approaches: collaborative filtering via valuable content information. In *38th Hawaii International Conference on System Sciences (HICSS05)*, page 217b, Hawaii, USA, 1999.
- [41] B. Smyth, E. Balfe, O. Boydell, K. Bradley, P. Briggs, M. Coyle, and J. Freyne. A Live User Evaluation of Collaborative Web Search. In *19th International Joint Conference on AI*, pages 1419–1424, Edinburgh, Scotland, 2005.
- [42] M. Stolze, S. Field, and P. Kleijer. Combining Configuration and Evaluation Mechanisms to Support to Selection of Modular Insurance Products. In *8th European Conference on Information Systems (ECIS2000)*, pages 858–865, Vienna, Austria, 2000.
- [43] A. Tartakovski, M. Schaaf, and R. Bergmann. Retrieval and Configuration of Life Insurance Policies. In *6th International Conference on Case-Based Reasoning (ICCBR2005)*, pages 552–565, Chicago, Illinois (USA), 2005. Lecture Notes in Artificial Intelligence 3620.
- [44] L. Terveen and W. Hill. *Beyond recommender systems: Helping people help each other, HCI in the New Millennium*. Addison Wesley, 2001.
- [45] C. Thompson, M. Göker, and P. Langley. A Personalized System for Conversational Recommendations. *Journal of Artificial Intelligence Research*, 21:393–428, 2004.
- [46] J. Tiihonen, T. Soininen, I. Niemelä, and R. Sulonen. Empirical Testing of a Weight Constraint Rule Based Configurator. In *ECAI 2002 Workshop on Configuration*, pages 17–22, Lyon, France, 2002.
- [47] European Union. *Richtlinie 2002/92/EG des Europäischen Parlaments und des Rates vom 9. Dezember 2002 über Versicherungsvermittlung*. Amtsblatt der EU, 2002.
- [48] I. Vollrath, W. Wilke, and R. Bergmann. Case-based reasoning support for online catalog sales. *IEEE Internet Computing*,

2(4):47-54, 1998.