# Combining Case-Based and Similarity-Based Product Recommendation

Armin Stahl

German Research Center for Artificial Intelligence (DFKI) GmbH
Research Group Image Understanding and Pattern Recognition (IUPR)
Technical University of Kaiserslautern
Erwin-Schrödinger-Str. 57, 67663 Kaiserslautern, Germany
**Armin.Stahl@dfki.de**

**Abstract.** Product recommender systems are a popular application and research field of CBR for several years now. However, almost all CBR-based recommender systems are not case-based in the original view of CBR, but just perform a similarity-based retrieval of product descriptions. Here, a predefined similarity measure is used as a heuristic for estimating the customers' product preferences. In this paper we propose an extension of these systems, which enables case-based learning of customer preferences. Further, we show how this approach can be combined with existing approaches for learning the similarity measure directly. The presented results of a first experimental evaluation demonstrate the feasibility of our novel approach in an example test domain.

## 1 Introduction

With the increasing success of e-Commerce web-sites, the development of intelligent *recommender systems* has become a popular field of research. Today, many e-Commerce sites are already deploying recommender systems to support their customers during the selection of a product that best matches their requirements and preferences. Depending on the type of offered products, the desire for such support can be explained by different issues:

- When being confronted with huge product databases, the search for a suitable product can become very time consuming.
- When purchasing complex products (e.g. technical products like PCs, travels [1], insurance products [2]) customers often do not possess the expertise to select the optimal product with respect to their requirements.
- Some products cannot be described sufficiently by explicit and objective properties (e.g. books, music [3], videos [4]) but are selected on the basis of subtle aspects like personal taste. Without a recommendation a customer cannot estimate the personal value of such a product until purchasing it.

Since the requirements on the actual recommendation process are varying between different business scenarios, many different recommendation techniques

have been developed during the last few years (for an overview see [5, 6]). In principle, two major approaches can be distinguished: *content-based recommendation* and *collaborative filtering (CF)*.

Content-based recommendation can deal with the first two issues, i.e. finding suitable products in large databases or advising customers when purchasing complex products. The customer has to define his product requirements, e.g. by filling out a predefined query form. This information is then compared with the descriptions of the available products in order to identify a set of potential product candidates. If the comparison is based on exact match (e.g. by performing a simple SQL query), this is called *filter-based recommendation (FBR)*. FBR often leads to unsatisfactory results. An alternative is *similarity-based recommendation (SBR)*. Here, the comparison between the query and the product descriptions is based on a specific similarity measure which also allows to rank retrieved products. As *Case-Based Reasoning (CBR)* provides powerful techniques for realizing similarity-based retrieval it has become a popular technique for building SBR systems [7].

Collaborative filtering [8], on the other hand, is typically used to deal with the third issue, i.e. to provide recommendations for products that cannot be described sufficiently by explicit properties. The basic idea of CF is to collect user ratings about seen or bought products and to use rating correlations between different users and products in order to recommend products. Hence, CF relies on a vast amount of user feedback before producing satisfactory recommendations.

In recent years several hybrid recommendation techniques which incorporate content-based and collaborative approaches have been developed [6], and some of them apply also CBR techniques [9, 10, 1, 11].

In principle, a recommender system must possess knowledge about the customers' requirements and preferences and their relationship to the offered products. Generally, the following types of user needs can be distinguished [1]:

- hard requirements vs. preferences
- explicit vs. implicit preferences, i.e. is the preference explicitly expressed in the query or not
- general vs. individual preferences, i.e. is it a general preference of almost all customers or is it customer specific
- short-term vs. long-term preferences, i.e. is the preference only valid for the actual recommendation process or durable

FBR can only treat the hard, explicit, individual and short-term preferences encoded in the query. SBR allows a much wider consideration of customer preferences. Similarity-measures (which may be customer specific) can be used to model almost all kind of preferences. Only implicit subtle preferences, that are difficult to express formally, can be treated exclusively by CF techniques.

The most challenging task when building a recommender system is the acquisition of knowledge about the different kinds of preferences. While CF strictly relies on user feedback, SBR is applicable without any feedback by using a predefined similarity measure as a heuristic. However, the quality of this heuristic influences the recommendation quality dramatically.

In this paper we propose a novel approach for learning customer preferences in content-based recommender systems. The approach combines SBR with the original idea of CBR, i.e. the reuse of collected experience knowledge. Therefore, it incorporates knowledge about successful recommendations of the past into the similarity-based product retrieval. We show that an additional optimization of the underlying similarity measure results in further improvements to the recommendation quality.

The advantage of our approach is its easy integration into state-of-the-art SBR systems. At the beginning the system can be applied with a standard similarity measure without relying on any user feedback. If feedback becomes available during usage it will enable the system to learn its users' preferences automatically over time leading to improved recommendation results. Moreover, our approach also provides the possibility to consider more subtle product properties which are not explicitly described by existing product descriptions.

Section 2 starts with a short review of the functionality of SBR systems and existing approaches towards learning customer preferences. Section 3 then describes our novel approach which combines case-based learning with existing techniques for learning similarity measures. The results of an experimental evaluation presented in Section 4 demonstrate the feasibility of our approach. After discussing related work in Section 5, we conclude with a summary and an outlook on future work.

## 2 Similarity-Based Product Recommendation Systems

Similarity-based recommendation systems have become a very popular CBR research area and numerous successful commercial applications are in use today. Surprisingly, on a closer look, most of these systems are not at all CBR systems in the traditional view of CBR since the used "cases" do not represent problem-solution pairs of the past but are typically just product descriptions.

### 2.1 Utility-Oriented Matching

The basic functionality of a SBR system is illustrated in Figure 1. Given a customer *query* which describes the desired product properties, a CBR system applies a predefined similarity measure for comparing the query with the descriptions of all available products which are stored in a *product database (PB)*. Finally, a ranked set of the $s$ most similar products (a typical value of $s$ is 10) is presented as the *result set (RS)* to the customer[1].

If we look at this scenario, it becomes obvious that the system does not compare two problem descriptions as assumed by the traditional idea of CBR. Instead, it compares a problem—the query—directly with potential solutions—the products. This works well for product recommendation because here problems and solutions can be described by using the same vocabulary. However, it also

---

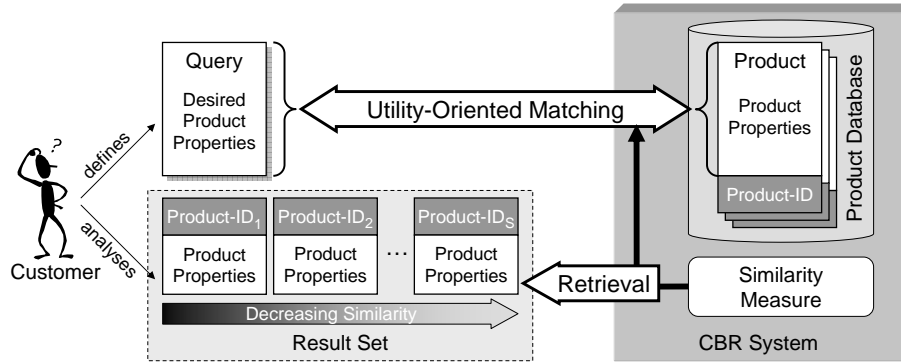[1] In this paper we do not consider adaptation.

**Fig. 1.** Product Recommendation by Utility-Oriented Matching

restricts the features that can be used during the recommendation process to the information contained in existing product descriptions.

Traditional cases consisting of problem-solution pairs are not used at all in this scenario. Some authors have characterized this approach as *utility-oriented matching* because the similarity measure is directly used to approximate the utility of known solutions—here product descriptions—for a given problem [12].

In principle, the utility of a product description $p_i$ with respect to a given query $q$ can be characterized as the conditional probability that the product will be accepted by the customer—we denote this event as $\omega_i$—given $q$, i.e. we may define a *utility function* $u$ as follows by applying Bayes rule:

$$u(q, p_i) = P(\omega_i|q) = \frac{P(q|\omega_i) \cdot P(\omega_i)}{P(q)} \tag{1}$$

In a SBR system, a predefined similarity measure *sim* is used to approximate this unknown utility function $u$. Since it does not possess any other knowledge about the customers' preferences, the recommendation quality of such a system depends completely on the accuracy of this approximation.

### 2.2  Dealing with Customer Preferences

However, because of the complexity of customer preferences, in practice standard similarity measures such as the Euclidean Distance will result in a poor approximation of $u$. In principle, $u$ will be determined by different kinds of preferences with different locality in the problem space:

1. the different importance of general product properties, e.g. "the price is usually much more important than the color"
2. preferences concerning different values of product properties
   (a) independent from $q$ and other properties, e.g. "black cars are generally preferred over white cars"

(b) depending on $q$ but independent from other properties, e.g. "if a black car is desired, a dark blue car will likely be preferred over a yellow car"

(c) depending on other properties, e.g. "black BMWs are mostly preferred over red BMWs"

3. product specific preferences that are independent from $q$ (in the probabilistic view this is the prior probability $P(\omega_i)$ of class $\omega_i$ in formula (1)), e.g. "the silver BMW 320i is a very popular car and is generally preferred over many other similar cars"

With similarity measures commonly supported by CBR tools [13], influence 1 can be modeled with global *feature weights* and influences 2a) and 2b) can be modeled with *local similarity measures*. However, in particular the definition of accurate local similarity measures is a very time consuming task.

The consideration of influences 2c) and 3) would require more sophisticated measures requiring a modeling effort that is usually not tolerable in practice. But the more serious problem of defining an accurate similarity measure is the fact that knowledge about the customers' preferences is a priori completely missing or only partially known.

In our previous work we have proposed to apply a machine learning approach which allows automatically learning of feature weights and local similarity measures based on user feedback [14, 15, 13]. We have shown that this approach also allows the incorporation of partially known background knowledge into the learning process [16]. However, in particular the learning of local similarity measures is generally susceptible to overfitting if not enough user feedback is available. Moreover, the approach does not provide a solution for the consideration of all above enumerated kinds of preferences.

## 3  Case-Based Learning of Customer Preferences

In this section we present an alternative approach for learning customer preferences which avoids some of the problems of the previously described approaches.

The basic idea of this approach is illustrated in Figure 2. At the beginning of its life cycle, the extended recommender system will behave like a standard SBR system, i.e. it will perform utility-oriented matching on the given $PB$. However, any time a customer has selected a product that is acceptable for him (e.g. if he orders the product), his query (optionally together with additional information, see Section 3.2) will be stored in the case base $CB$ together with the product-ID of the selected product. These records now represent actual cases in the traditional view of CBR; the combination of a problem description—the query—and a corresponding solution—the accepted product. Such cases do not necessarily represent optimal cases, because in general it cannot be guaranteed that the customer orders the optimal product regarding to his query. Nevertheless, such cases contain some implicit knowledge about the customers' preferences, i.e. the relationship between certain queries and products that are at least acceptable for the customer.
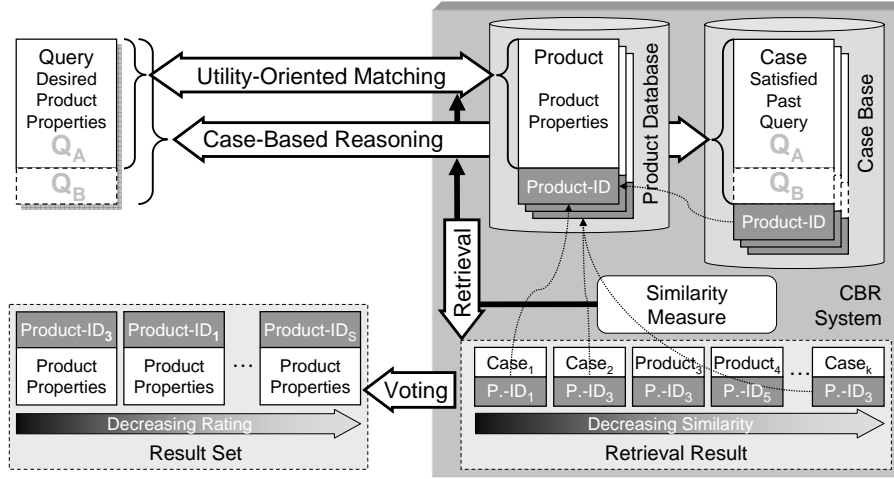
**Fig. 2.** Case-Based Learning of Customer Preferences

During subsequent recommendation sessions, this knowledge can be used to estimate $u$ more accurately than possible with a predefined similarity measure $sim$ alone. Therefore, the current query $q$ is not only matched against $PB$, but also against $CB$ by using $sim$. This is possible because both $PB$ and $CB$ contain product descriptions ($Q_A$ in Figure 2). This procedure results in the corresponding *retrieval result* $RR := (r_1, \ldots, r_k)$ where $r_i \in PB \cup CB$ with $sim(q, r_i) \geq sim(q, r_j)$ for all $i < j$. $RR$ cannot be used directly as result set because different $r_i$ may refer to the same product $p_i$. In order to generate a unique result set, some post-processing of $RR$, e.g. by applying a voting strategy, is required.

### 3.1 Voting Strategy

Because in the recommendation scenario one is not only interested in the probably most useful product, but in a set of the $s$ most useful products, the definition of a fixed size $k$ of $RR$ is not suited for our approach. If $RR$ contains only $r_i$ that correspond to $t < s$ different products, the system will not be able to recommend $s$ different products. Hence, $k$ has to be determined dynamically after each retrieval process. It will be set to the smallest possible value $k$, so that $RR$ includes exactly $s$ different products.

Now $RR$ can be used to generate the required result set $RS$ consisting of $s$ different product proposals. The ranking of these products within $RS$ will be determined by some voting strategy. Different voting strategies are commonly used in $k$-NN classification, e.g. simple majority voting, weighted voting [17] or advanced and adaptive approaches [18, 19]. In our experiments (see Section 4) product $p_x$ is ranked over product $p_y$ (written as $p_x \succ p_y$) according to the

following weighted majority voting rule, where all $r_i^{p_x}$ correspond to the same product $p_x$:

$$p_x \succ p_y \Leftrightarrow \sum_{r_i^{p_x} \in RR} sim(q, r_i^{p_x}) > \sum_{r_i^{p_y} \in RR} sim(q, r_i^{p_y}) \qquad (2)$$

This means, that $r_i$ which are very similar to $q$ will have a higher impact on the majority voting. Such a simple weighted voting strategy may lead to overfitting problems as long as only a few cases have been collected because the corresponding products will then have a much higher probability to be recommended compared with still unbought products. More adaptive strategies which take the cases' distribution into account might outperform the proposed strategy.

Finally, the actual product descriptions have to be retrieved from the product database in order to generate the final result set to be presented to the customer. This is necessary because the retrieval result might contain only cases from $CB$ which do not contain the product description itself but only the product-ID.

### 3.2 Learning Additional Case Indexes

Up to now, we have assumed that queries consist of the same attributes that are used in the original product descriptions (denoted as $Q_A$ in Figure 2). As already described, this is a precondition for applying utility-oriented matching. However, by learning cases of successful recommendation sessions, this is no longer strictly necessary. One may enable the customer to ask also for additional product properties that are not contained at all in the original product descriptions (part $Q_B$ of the query). Typical examples of such additional query items are more subtle (e.g. "I want a very sporty car") or functional requirements (e.g. "I want to use my PC mainly for gaming"). In principle, a fixed set of such additional features may be considered explicitly in the query interface or the interface may provide the option to enter some free text to be processed by textual CBR techniques. In order to handle this additional information during the retrieval step, the used similarity measure $sim$ has to be extended appropriately. This extension will only influence the retrieval of cases in $CB$ but not the retrieval of product descriptions in $PB$.

The additional desires of the customer cannot improve the recommendations if the case base is still empty. However, the more cases that contain such information are stored in the case base, the higher will be the influence of this information on the recommendation results. In principle, the queries of the customers are then used to implicitly index the products automatically by using additional features which would be too expensive to be done manually by domain experts. At some point, the collected information could also be extracted automatically from the case base by applying statistical techniques in order to explicitly extend the product descriptions stored in the product database. Another possibility is the incorporation of collaborative features (e.g. user profiles) [1] in part $Q_B$ of the query.

### 3.3 Acquisition of Training Data

As typical for a supervised learning approach, the acquisition of accurate training data is crucial. In our approach, we assume that the customer states some query (this query might also be the result of a sales dialog [20]) and in the case that he accepts one of the proposed products (e.g. because he orders it), this data is used to create a new training example, i.e. a new case. To control the learning process, one may choose one of the CBL algorithms [21]. For example, when applying the CBL2 algorithm, one would store a new case only if the ordered product was not recommended as the optimal product.

However, in general it cannot be guaranteed, that the resulting case represents an optimal query-product pair. Maybe there are other products in $PB$ that the customer has not seen, for example, because they were not included in the result set, but that he would definitely prefer. Moreover, changes in $PB$ may also influence the quality of the collected cases. Maybe a customer would now prefer a newer product which was not available at the time the case was created.

This means, we will only get *relative utility feedback* [22] about the utility of the products included in the original result sets[2] of past recommendation sessions. If the system proposes $s$ different products $p_1, \ldots, p_s$ and the customer orders $p_3$ we only get evidence that $u(q, p_3) \geq u(q, p_i)$ for all $i \leq s$, but we do not obtain reliable information about the absolute value of $u(q, p_3)$. However, this information would be necessary in order to ensure that the learned case alone represents accurate knowledge about $u$.

This situation is less problematic if the retrieval set contains the optimal case with high probability, even if it is not ranked correctly. Hence, the quality of the predefined similarity measure which determines the initial result sets is crucial in order to restrict the noise in the training data required for case-based learning. However, learning of extremely noisy training examples is generally unlikely because then the customer would not have ordered the product.

To guarantee a minimal quality of the used similarity measure it is possible to apply machine learning, too. In [14, 15, 13] we have presented an algorithm for learning similarity measures which can handle the kind of relative utility feedback that we obtain in the recommendation scenario. Hence, this feedback can also be used to optimize the similarity measure in parallel or a priori to learning new cases in order to reduce the noise in the training data.

## 4 Experimental Evaluation

In order to evaluate our novel approach we have performed some first experiments in a simulated product recommendation scenario.

### 4.1 Test Domain

As test scenario we have chosen a used cars domain consisting of 100 descriptions of different used cars which we have extracted from a real world online

---

[2] Here we assume that the customer analyzes each product contained in the result set.

used cars market. Each car is described by 4 numeric and 4 symbolic attributes, such as price, power, color, year of construction, etc. For a more detailed description of the used test domain see [13]. In the described evaluation we have not investigated the possibility to learn additional product features during the recommendation process (cf. Section 3.2), i.e. the 8 mentioned features solely correspond to part $Q_A$ of the query illustrated in Figure 2.

Since we were not able to perform an experiment with real world customers, we have simulated imaginable average customer preferences with a manually defined similarity measure $sim_U$ consisting of specific feature weights and specific local similarity measures for each attribute. Of course, such a model is not sufficient for simulating the actual behaviour of real world customers. On the one hand, it does not simulate the inconsistencies between the individual preferences of different customers that would occur in the real world. On the other hand, it also does not model all the kinds of preferences discussed in Section 2.2 (2c and 3 cannot be modeled with such a kind of similarity measure). However, it is sufficient for a first proof of concept of our approach.

## 4.2  Experiments

In order to evaluate the capability to learn the simulated customer preferences in the described test domain, we have performed several experiments where we have applied the proposed case-based learning approach and/or our previous algorithm for learning feature weights [14]. In principle, each experiment consisted of the following steps:

1. create empty case base $CB$, empty feedback set $FB$, and initialize standard similarity measure $sim$ with uniform weights

2. select a set of training queries $Q_{train} := (q_1, \ldots, q_{10000})$

3. for each $q_i \in Q_{train}$ do

   (a) generate result set $RS_i := (p_1, \ldots, p_{10})$ consisting of 10 product descriptions $p_j$ by following the procedure described in Section 3 and by using $q_i$, $sim$, $CB$ and the static product database $PB$

   (b) determine *preferred product* $pp_i := \arg\max_{p_j \in RS_i} sim_U(q_i, p_j)$

   (c) generate feedback $FB_i := (pp_i, (\bar{p}_1, \ldots, \bar{p}_9))$ where $\bar{p}_l \in RS_i \setminus pp_i$

   (d) store feedback, i.e $FB := FB \cup FB_i$

   (e) optional: learn feature weights from $FB$ and update $sim$ accordingly

   (f) create a new case $c_i$ from $q_i$ and the product-ID of $pp_i$

   (g) optional: insert $c_i$ into $CB$ by applying CBL1 or CBL2

   (h) if $i \in \{5, 10, 25, 50, 100, 250, 500, 1000, 2500, 5000, 10000\}$ then evaluate the recommendation accuracy on query test set $Q_{test}$ using $sim_U$

By different combinations of the optional learning steps 3(e) and 3(g) we have generated the following five experiments:

**SIM:** Exclusively learning of feature weights by using the relative utility feedback $FB$ which only expresses that the $pp_i$ are more useful than all other $p_j$ contained in the respective result sets $RS_i$.

**CBL1/2:** Exclusively applying case-based learning algorithms CBL1 (each $c_i$ is stored) or CBL2 ($c_i$ is stored only if $pp_i \neq p_1$ holds) (cf. [21]).

**SIM-CBL1/2:** A-priori learning of feature weights using the feedback of the first 5/10/25/50 queries and activation of CBL1/2 starting from query 51.

Each experiment was repeated with 5 different, a priori randomly generated training query sets where each attribute value of the individual queries was selected randomly. For the evaluation of the achieved recommendation accuracy a static set of 250 independent randomly generated test queries $Q_{test}$ was used to compute 4 different quality measures:

$mpp$**-in-**$x$**:** The average percentage of recommendation sessions, where the theoretically *most preferred product* $mpp = \arg\max_{p_i \in PB} sim_U(q, p_i)$ was contained in the first $x \in \{1, 3, 10\}$ recommended products.

**avg-**$mpp$**:** The average position of $mpp$ in the result sets.

### 4.3   Results

Figure 3 summarizes the results of the experiments SIM and CBL1/2. The left chart shows the achieved improvements concerning the $mpp$-in-$x$ measures. For the exclusive optimization of $sim$ one observes a rapid ascent of all learning curves where about 10 training queries are sufficient to achieve the maximal improvements, e.g. for the $mpp$-in-10 measure an increase from 52% to 81%.
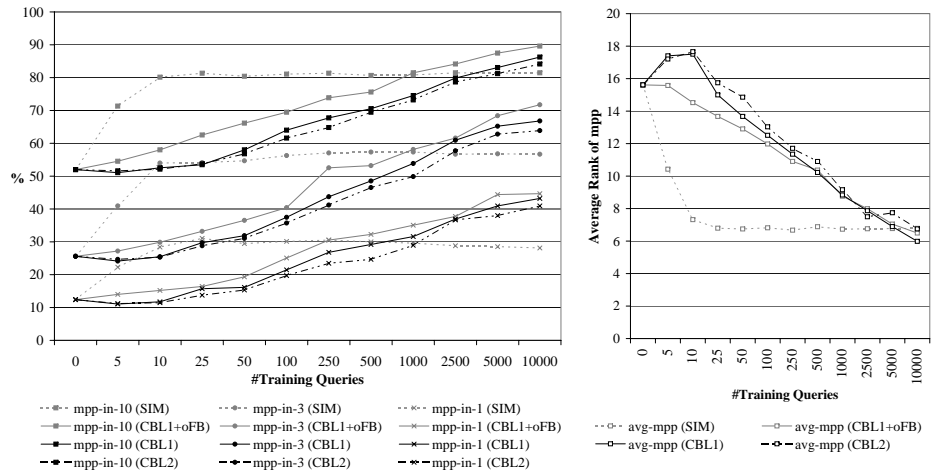


**Fig. 3.** Results of Experiments SIM and CBL1/2

In contrast, the learning curves of the CBL experiments show much slower improvements of the recommendation quality. However, after 1000-2500 training queries case-based learning starts to outperform similarity measure learning and achieves significantly better results after 10000 queries, e.g. for the $mpp$-in-1 measure over 40% (compared to about 29%). This is not surprising since the case-based learning approach is able to learn the preferences encoded in the local similarity measures of $sim_U$ which cannot be modeled with feature weights. However, surprisingly the differences between the CBL1 and CBL2 are very small, even though the average number of stored cases is significantly lower in the CBL2 experiment (6032 compared to 10000 in CBL1).

In order to be able to evaluate the impact of noisy feedback, we have performed an additional CBL1 experiment with optimal feedback (CBL1+oFB) by using $mpp$ instead of $pp$ in step 3(b). The CBL1+oFB learning curve shows continuous recommendation improvements from the beginning. In the more realistic experiments CBL1/2 the improvements achieved with less then 50 training queries are quite small or even negative. This can be explained by overfitting which will be amplified by the noisy training data in CBL1/2 and becomes more obvious in the avg-$mpp$ measure (right chart of Figure 3).

Figure 4 shows the results of the experiments SIM-CBL1/2. Here, the first 50 queries[3] were used to exclusively optimize the similarity measure in order to improve the feedback quality for the subsequent case-based learning process.



**Fig. 4.** Results of Experiments SIM-CBL1/2

The achieved results clearly show the advantage of the combination of both learning techniques. On the one hand, a priori optimization of $sim$ ensures much faster performance gains compared with applying case-based learning alone. Now

---

[3] according to the results of SIM even 25 queries would be sufficient

only 250 training queries are sufficient to outperform the results of SIM. This also leads to increased robustness against overfitting, since the negative impact of a too small case base is compensated by the preliminary improvements achieved by optimizing *sim*. However, the avg-*mpp* curves clearly show that the overfitting effect is still present (see peak at 100 training queries) as long as the case base contains less than 50 cases[4]. Although overfitting is more prominent in SIM-CBL2, in general the differences between SIM-CBL1 and SIM-CBL2 are almost not recognizable. This is all the more surprising because the average number of learned cases is further reduced (4930) compared to experiment CBL2.

On the other hand, the finally achieved recommendation accuracy is significantly higher compared with the results that can be achieved with each learning technique alone. This becomes particularly obvious in the avg-*mpp* quality measure. While each learning technique alone was able to decrease the average rank of *mpp* from about 16 to 6, the combination of both approaches finally achieves an average rank smaller than 3. This would allow to decrease the size of the returned result sets significantly, e.g. if displayed on mobile devices [23].

## 5 Related Work

The work that is most related to the approach presented in this paper are the results of the DIETORECS project [1, 24]. In this project, an advanced travel recommendation system which combines case-based and CF-based techniques has been developed. This system also uses the combination of a raw product database and a case base containing so-called *session-cases*. These cases describe recommendation sessions of the past, containing stated queries, selected travel components, and also collaborative features.

The major difference compared to our work is the way the two databases are used. In the DieToRecs system, on the one hand, the product database is used for an initial filter-based retrieval which requires conversational techniques in order to obtain useful result sets. On the other hand, the case base is used only to determine the ranking of the previously selected cases by using collaborative techniques. Moreover, the system does not optimize the similarity measure required to retrieve session-cases.

An early algorithm which integrates case-based learning with optimization of the required similarity measure is the CBL4 algorithm [21]. However, this algorithm is designed for simple classification tasks and requires absolute utility feedback [22] about the correctness of the solution proposed by the CBR system. Hence, it is not applicable in the product recommendation scenario.

Other work which deals with learning of user preferences is described in [25–27]. However, none of these approaches applies a combination of case-based and similarity measure learning as proposed in this paper.

---

[4] Note, that the first 50 queries were not used for learning cases.

# 6 Conclusion and Future Work

In this paper we have presented a novel approach for learning customer preferences in content-based recommender systems. This approach extends the functionality of existing similarity-based recommender systems by applying case-based learning in combination with similarity measure learning. On the one hand, optimizing the similarity measure directly improves the accuracy of the approximation of the unknown utility function. On the other hand, it also improves the quality of the absolute utility feedback required by the case-based learner.

The best suited target applications for our approach are product recommendation systems which deal with relatively static product databases. Very dynamic product databases will lead to problems because then the risk that stored cases become obsolete is very high. In this paper we have only considered learning of general or average preferences of all or at least a certain class of customers. In principle, our approach can also be used to support personalized recommendation. However, in practice this will only be feasible in domains where individual customers frequently buy the same type of products.

Although in this paper we have focused on product recommendation, the approach is not restricted to this application scenario. It is also suited to learn other types of user preferences, e.g. like those occurring in knowledge management domains where users are interested in getting advice about available knowledge resources (e.g., documents, web sites, pictures) with respect to their individual information needs [28].

The advantage of our approach is its broad applicability and its compatibility with already successfully applied SBR systems. Moreover it allows an automatic extension of the set of features used to characterize products or information resources. The results of the presented experimental evaluation show the principal ability of our approach to learn customer preferences from easy to acquire customer feedback.

For future work we are planning to perform a more realistic evaluation by using a more sophisticated model of the customers' preferences including all kinds of preferences discussed in Section 2.2 and also nondeterministic behavior. In such a scenario the learning task is generally more challenging. However, here the case-based approach should also outperform solely learning of the similarity measure more clearly due to its less restricted hypotheses space. In such an extended evaluation it would also be interesting to investigate the impact of learning additional product features (cf. Section 3.2).

We also plan to further improve our approach. On the one hand, advanced voting strategies which incorporate statistical information about the learned cases might allow to model the prior probabilities $P(\omega_i)$ explicitly in order to improve the recommendation accuracy. Moreover, this might also help to reduce the risk of overfitting for small case bases. Generally, we plan to investigate the potential of the generation and incorporation of statistical models into the recommendation process with the increasing number of collected cases. On the other hand, smarter learning policies than CBL2 (e.g. such as CBL3 [21]) can

help to reduce the size of the case base while maintaining or even improving the recommendation accuracy. This is important in order to minimize retrieval times. Another interesting issue would be the application of our advanced similarity measure learning algorithm which allows an optimization of local similarity measures [15, 13, 16].

Last but not least, we want to investigate whether our approach is also suited to be used in domains where products can customized [29, 30].

## Acknowledgements

## References

1. Ricci, F., Venturini, A., Cavada, D., Mirzadeh, N., Blaas, D., Nones, M.: Product Recommendation with Interactive Query Management and Twofold Similarity. In: Proceedings of the 5th International Conference on CBR, Springer (2003)
2. Tartakovski, A., Schaaf, M., Bergmann, R.: Retrieval and Configuration of Life Insurance Policies. In: Proceedings of the 6th International Conference on CBR, Springer (2005)
3. Hayes, C., Avesani, P., Baldo, E., Cunningham, P.: Re-using Implicit Knowledge in Short-Term Information Profiles for Context-Sensitive Tasks. In: Proceedings of the 6th International Conference on CBR, Springer (2005)
4. Miller, B.N., Albert, I., Lam, S.K., Konstan, J.A., Riedl, J.: MovieLens Unplugged: Experiences with an Occasionally Connected Recommender System. In: Intelligent User Interfaces. (2003)
5. Schafer, J.B., Konstan, J., Riedi, J.: Recommender systems in e-commerce. In: Proceedings of the 1st ACM Conference on Electronic commerce, ACM Press (1999)
6. Burke, R.D.: Hybrid Recommender Systems: Survey and Experiments. User Modeling and User-Adapted Interaction **12**(4) (2002)
7. Bergmann, R., Schmitt, S., Stahl, A.: Intelligent Customer Support for Product Selection with Case-Based Reasoning. In: E-Commerce and Intelligent Methods. Physica-Verlag (2002)
8. Breese, J.S., Heckerman, D., Kadie, C.M.: Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In: Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence. (1998) 43–52
9. Hayes, C., Cunningham, P., Smyth, B.: A Case-Based Reasoning View of Automated Collaborative Filterning. In: Proceedings of the 4th International Conference on Case-Based Reasoning (ICCBR'2001), Springer (2001)
10. O'Sullivan, D., Wilson, D., Smyth, B.: Improving Case-Based Recommendation: A Collaborative Filtering Approach. In: Proceedings of the 6th European Conference on Case-Based Reasoning (ECCBR'2002), Springer (2002)
11. Burke, R.D.: Hybrid Recommender Systems with Case-Based Components. In: Proceedings of the 7th European Conference on CBR. (2004) 91–105

12. Bergmann, R., Richter, M.M., Schmitt, S., Stahl, A., Vollrath, I.: Utility-Oriented Matching: A New Research Direction for Case-Based Reasoning. In: Professionelles Wissensmanagement: Erfahrungen und Visionen. Proceedings of the 1st Conference on Professional Knowledge Management, Shaker (2001)
13. Stahl, A.: Learning of Knowledge-Intensive Similarity Measures in Case-Based Reasoning. Volume 986. dissertation.de (2004)
14. Stahl, A.: Learning Feature Weights from Case Order Feedback. In: Proceedings of the 4th International Conference on Case-Based Reasoning (ICCBR'2001), Springer (2001)
15. Stahl, A., Gabel, T.: Using Evolution Programs to Learn Local Similarity Measures. In: Proceedings of the 5th International Conference on Case-Based Reasoning (ICCBR'2003), Springer (2003)
16. Gabel, T., Stahl, A.: Exploiting background knowledge when learning similarity measures. In: Proceedings of the 7th European Conference on Case-Based Reasoning (ECCBR'2004), Springer (2004)
17. Dudani, S.A.: The Distance-Weighted k-Nearest Neighbor Rule. IEEE Transactions on Systems **6** (1981)
18. Rovatti, R., Tagazzoni, R., Kovcs, Z.M., Guerrieri, R.: Adaptive Voting Rules for k-NN Classifiers. Neural Computation **7**(3) (1995)
19. Wang, H., Bell, D.: Extended k-Nearest Neighbours based on Evidence Theory. The Computer Journal **47**(6) (2004)
20. Schmitt, S.: Dialog Tailoring for Similarity-Based Electronic Commerce Systems. Ph.D. Thesis, Kaiserslautern University of Technology (2003)
21. Aha, D.: Case-Based Learning Algorithms. In: Proceedings of the DARPA Case-Based Reasoning Workshop, Morgan Kaufmann (1991) 147–158
22. Stahl, A.: Learning Similarity Measures: A Formal View Based on a Generalized CBR Model. In: Proceedings of the 6th International Conference on Case-Based Reasoning, Springer (2005)
23. Smyth, B., McClave P.: Similarity vs. Diversity. In: Proceedings of the 4th International Conference on Case-Based Reasoning, Springer (2001)
24. Ricci, F., Arslan, B., Mirzadeh, N., Venturini, A.: Itr: A case-based travel advisory system. In: Proceedings of the 6th European Conference on CBR, Springer (2002)
25. Coyle, L., Cunningham, P.: Exploiting Re-ranking Information in a Case-Based Personal Travel Assistent. In: Workshop on Mixed-Initiative Case-Based Reasoning at the 5th International Conference on Case-Based Reasoning (ICCBR'2003), Springer (2003)
26. Branting, K.: Acquiring Customer Preferences from Return-Set Selections. In: Proceedings of the 4th International Conference on Case-Based Reasoning (ICCBR'2001), Springer (2001)
27. Gomes, P., Bento, C.: Learning User Preferences in Case-Based Software Reuse. In: Proceedings of the 5th European Workshop on Case-Based Reasoning (EWCBR'2000), Springer (2000)
28. Holz, H.: An Incremental Approach to Task-Specific Information Delivery in SE Processe. In: 18th IEEE International Conference on Automated Software Engineering. (2003)
29. Stahl, A., Bergmann, R.: Applying Recursive CBR for the Customization of Structured Products in an Electronic Shop. In: Proceedings of the 5th European Workshop on Case-Based Reasoning (EWCBR'2000), Springer (2000)
30. Schmitt, S., Bergmann, R.: Product Customization in an Electronic Commerce Environment Using Adaptation Operators. In: Proceedings of the 7th German Workshop on Case-Based Reasoning (GWCBR'99). (1999)