

Andreas Thor · Nick Golovin · Erhard Rahm

Adaptive website recommendations with AWESOME

Received: 30 November 2004/ Accepted: 12 April 2005/ Published online: 20 October 2005
© Springer-Verlag 2005

Abstract Recommendations are crucial for the success of large websites. While there are many ways to determine recommendations, the relative quality of these recommenders depends on many factors and is largely unknown. We present the architecture and implementation of AWESOME (Adaptive website recommendations), a data warehouse-based recommendation system. It allows the coordinated use of a large number of recommenders to automatically generate website recommendations. Recommendations are dynamically selected by efficient rule-based approaches utilizing continuously measured user feedback on presented recommendations. AWESOME supports a completely automatic generation and optimization of selection rules to minimize website administration overhead and quickly adapt to changing situations. We propose a classification of recommenders and use AWESOME to comparatively evaluate the relative quality of several recommenders for a sample website. Furthermore, we propose and evaluate several rule-based schemes for dynamically selecting the most promising recommendations. In particular, we investigate two-step selection approaches that first determine the most promising recommenders and then apply their recommendations for the current situation. We also evaluate one-step schemes that try to directly determine the most promising recommendations.

Keywords Adaptive web recommendations · Web data warehouse · Classification of recommendation algorithms · Web usage mining

1 Introduction

To guide web users, presented website pages frequently contain web recommendations, i.e. links to related or potentially useful information. Especially for large websites recommendations are crucial for the usefulness and success of

the website. For instance, e-commerce sites offering thousands of products cannot solely rely on standard navigation and search features but need to apply recommendations to help users quickly find “interesting” products or services. Given many users and products, manual generation of recommendations becomes much too laborious and ineffective. Hence, a key question is how recommendations should be generated automatically to optimally serve the users of a website.

There are many ways to automatically generate recommendations taking into account different types of information (e.g. product characteristics, user characteristics, or buying history) and applying different statistical or data mining approaches [1, 2]. Sample approaches include recommendations of top-selling products (overall or per product category), new products, similar products, products bought together by customers, products viewed together in the same web session, or products bought by similar customers. Obviously, the relative utility of these recommendation approaches (*recommenders* for short) depends on the website, its users and other factors; hence, there cannot be a single best approach. Website developers thus have to decide about which approaches they should support and where and when these approaches should be applied. Surprisingly, little information is available in the open literature on the relative quality of different recommenders. Hence, one focus of our work is an approach for comparative quantitative evaluations of different recommenders.

Advanced websites, such as Amazon [3], support many recommenders but apparently are unable to select the most effective approach per user or product. They overwhelm the user with many different types of recommendations leading to huge web pages and reduced usability. While commercial websites often consider the buying behavior for generating recommendations, the navigation behavior on the website remains largely unexploited. We believe this a major shortcoming, since the navigation behavior contains detailed information on the users’ interests not reflected in the purchase data. Moreover, the navigation behavior contains valuable user feedback not only on products or other content but

A. Thor (✉) · N. Golovin · E. Rahm
University of Leipzig, Leipzig, Germany
E-mail: {thor, golovin, rahm}@informatik.uni-leipzig.de

also on the presented recommendations. The utilization of this feedback to automatically and adaptively improve recommendation quality is a major goal of our work.

AWESOME (Adaptive website recommendations) is a new data warehouse-based website evaluation and recommendation system under development at the University of Leipzig. It contains an extensible library of recommender algorithms that can be comparatively evaluated for real websites based on user feedback. Moreover, AWESOME can perform an automatic closed-loop website optimization by dynamically selecting the most promising recommendations for a website access. Adaptiveness is important to automatically react to changing user interests and changing content. In *closed-loop optimization* we use the system's own past efficiency measurements to optimize its future behavior. In particular, we continuously measure user feedback on presented recommendations and use the observed recommendation quality to select new recommendations. Thus, AWESOME can be instrumental in creating so-called adaptive web sites: web sites that automatically improve their organization and presentation by learning from user access patterns [4]. To support high performance and scalability, recommendation quality characteristics and recommendations are largely pre-computed. AWESOME is fully operational and in continuous use at a sample website; adoption to further sites is in preparation.

The main contributions of this paper are as follows:

- We outline the AWESOME architecture for warehouse-based recommender evaluation and for scalable adaptive website recommendations. We provide details on the data warehouse implementation to illustrate how good data quality is supported and how recommendation feedback is maintained
- We present a new classification of recommenders for websites supporting a comparison of different approaches. We show how sample approaches fit the classification and propose a new recommender for users

coming from search engines. We also provide a comparative quantitative evaluation of several recommenders for a sample website. The considered recommenders cover a large part of our classification's design space.

- We propose two-step and one-step approaches for dynamically determining recommendations based on different types of selection rules. One-step approaches try to directly determine the most suitable recommendations, while two-step approaches first determine the most promising recommenders and then apply their recommendations for the current situation. We describe implementation alternatives for two-step and one-step recommendation selection including machine-learning approaches for taking recommendation feedback into account. Finally, we provide a comparative evaluation of the approaches for dynamic recommender and recommendation selection for a sample website.

In the next two sections, we present the AWESOME architecture and the underlying data warehouse approach. In Sect. 4 we outline our recommender classification and sample recommenders. Section 5 contains the comparative evaluation of several recommenders for a non-commercial website. We describe the feedback-based selection of recommenders and recommendations in Sect. 6 and give a comparative evaluation of both selection approaches in Sect. 7. In Sect. 8, we briefly review related work. Section 9 contains the conclusions.

2 Architecture

Figure 1 illustrates the overall architecture of AWESOME, which is closely integrated with the application server running the website. AWESOME is invoked for every website access, specified by a so-called context. A *context* is defined by a set of parameters, which describe the current state of the interaction between the web user and the website,

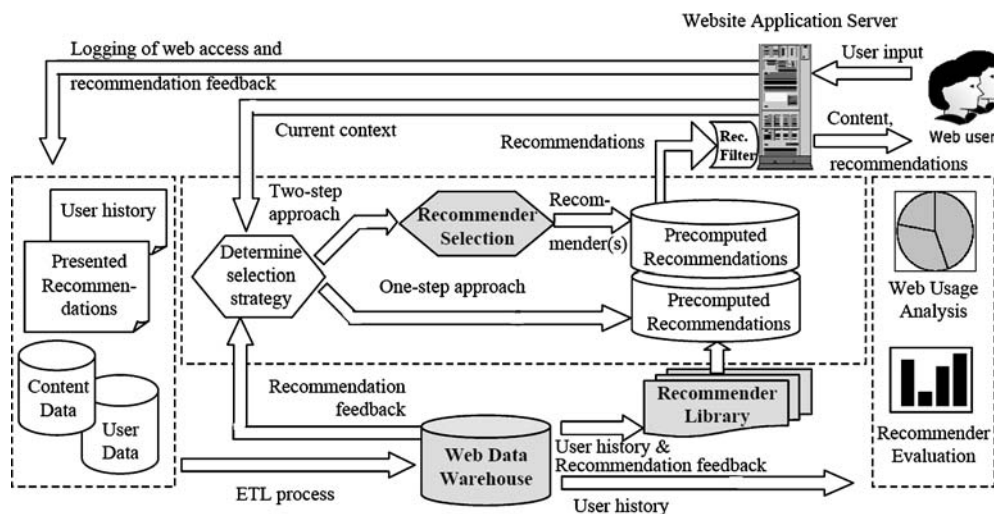


Fig. 1 AWESOME architecture

including information from the current HTTP request such as URL, timestamp and user-related data. For such a context, AWESOME dynamically generates a list of *recommendations*, i.e links to website content, to be displayed together with the requested website content. Recommendations are automatically determined by a variety of algorithms (*recommenders*) from an extensible *recommender library*. *Recommenders* may use usage history of the website and additional information maintained in a *web data warehouse* to generate recommendations. For performance reasons, recommendations are pre-computed and periodically refreshed and can thus be quickly looked up at runtime. The recommendations are subject to a final filter step to avoid the presentation of unsuitable or irrelevant recommendations, e.g., recommendation of the current page or the homepage.

AWESOME supports two-step and one-step selection of recommendations at runtime. The former approach first selects the most appropriate recommender(s) for the current context and then applies their recommendations. The one-step approach directly determines the most promising recommendations for the current context. Both approaches are controlled by selection rules which may be determined manually or automatically. The automatic approaches utilize user feedback on previously presented recommendations which is recorded in the web data warehouse.

Two-step selection approaches typically use a small or moderate number of *recommender rules* to determine the most promising recommenders per context. Feedback is associated to recommenders, which is simpler and more stable than trying to use feedback for individual recommendations, e.g. specific web pages or products. Hence, this approach supports a complete manual specification of recommender rules. Adding and deleting web pages or products does not directly influence the aggregated feedback and can thus easily be supported. Moreover, adapting the rules can be performed less frequently than in one-step approaches, since significant changes in the selection rules can only be expected when there is a certain amount of new feedback for most recommenders.

On the other hand, one-step approaches use so-called *recommendation rules* to directly determine the most promising recommendations per context. This approach is more fine-grained and allows selection of good recommendations irrespective of by which algorithm they were determined. One problem with the one-step approach is that individual pages/products may be frequently added and that there is no feedback available for such new content. Conversely, removing content results in a loss of the associated recommendation feedback.

Since there are typically many more possible recommendations than recommenders, the number of recommendation rules needs to be much larger than the number of recommender rules. This not only may make it more difficult to achieve fast recommendation selection but also to ensure that all potential recommendations are actually presented and thus obtain a fair chance to collect recommendation feedback. On the other hand, associating feedback to

Table 1 Strategies for dynamic selection of recommendations

	Strategy type	
	Two-step	One-step
Rule type/granularity	Recommender	Recommendation
No. of rules	Small to moderate	Many
Rule adaptation	Offline	Offline and online

individual recommendations may enable a more responsive approach by adapting the selection rules more frequently than with the two-step scheme. We have thus implemented approaches that not only re-compute recommendations periodically (offline) but also use short-term feedback on presented recommendations for an online adaptation of recommendation rules.

Table 1 summarizes some of the characteristics of the two strategy types for dynamic selection of recommendations. AWESOME implements several two-step and one-step selection strategies including machine-learning approaches for taking recommendation feedback into account. We will present the approaches and evaluate their effectiveness in Sects. 6 and 7, respectively.

AWESOME is based on a comprehensive web data warehouse integrating information on the website structure and content (e.g., product catalog), website users and customers, the website usage history and recommendation feedback. The application server continuously records the users' web accesses and which presented recommendations have been and which ones have NOT been followed. During an extensive ETL (extract, transform, load) process (including data cleaning, session and user identification) the usage data and recommendation feedback is added to the warehouse.

The warehouse serves several purposes. Most importantly it is the common data platform for all recommenders and keeps feedback for the dynamic recommender /recommendation selection thus enabling an automatic closed-loop website optimization. However, it can also be used for extensive offline evaluations, e.g., using OLAP tools, not only for web usage analysis but also for a comparative evaluation of different recommenders and of different selection strategies.

This functionality of AWESOME allows us to systematically evaluate the various approaches under a large range of conditions. It is also an important feature for website designers to fine-tune the recommendation system, e.g. to deactivate or improve less effective recommenders. Therefore, the recommender library is strictly separated from the selection process. New recommenders can be added to AWESOME without changing other parts of the system. Furthermore, the functionality of all recommenders is encapsulated so that further development of recommenders is not restricted.

2.1 Prototype implementation

The current AWESOME implementation runs on different servers. The recommendation engine is largely implemented by PHP programs on a Unix-based application server running the Apache web server. The pre-computed

Table 2 Log file formats

(a) Web usage log	
Page	Requested page
Date, time	Date and time of the request
Client	IP address of the user's computer
Referrer	Referring URL
Session ID	Session identifier
User ID	User identifier
(b) Recommendation log	
Page view ID	Page view where recommendation has been presented
Recommendation	Recommended content
Position	Position of this recommendation inside a recommendation list
Recommender	Recommender that generated the recommendation (only for two-step selection)
Strategy	Name and type (one-step or two-step) of selection strategy

Table 3 Fraction of web usage log file containing three sessions

Session ID	Host	Time	Page	Referrer	Session
–	cw12.H1.srv.t-online.de	01:15:02	/DBS1/ch4-1.html	www.google.com/search?q=sql+statements	A
–	pD9E1621B.dip.t-dialin.net	01:15:10	/ADS1/ch3-5.html	www.google.com/search?q=linked+list	B
1	vc05cop7001.cpmc.columbia.edu	01:15:12	/People/rahm.html	research.microsoft.com/~philbe/	C
–	cw12.H1.srv.t-online.de	01:15:20	/DBS1/ch4-2.html	/DBS1/ch4-1.html	A
1	vc05cop7001.cpmc.columbia.edu	01:15:21	/index.html	/People/rahm.html	C
–	cw08.H1.srv.t-online.de	01:15:23	/DBS1/ch4-3.html	/DBS1/ch4-2.html	A
–	pD9E1621B.dip.t-dialin.net	01:15:43	/ADS1/ch3-4.html	/ADS1/ch3-5.html	B

recommendations and selection rules are maintained in a MySQL database. During the two-step selection process, for a given context and a set of selected recommenders a dynamic SQL statement is generated to select recommendations at runtime.

The warehouse is on a dedicated machine running MS SQL server. Most recommenders are implemented in SQL, but Java and Perl programs are also utilized. All recommenders are periodically (currently, once a day) executed to re-compute the recommendations.

3 Web data warehouse

To illustrate the value of a data warehouse approach for generating web recommendations we provide some details on its implementation and design. We present the log formats of the application server and explain some ETL functions important for achieving good data quality. Furthermore, we sketch the warehouse schema used for web accesses and recommendation feedback.

3.1 Web log files

The standard log files of web servers are not sufficient for our purposes because to obtain sufficient recommendation feedback we need to record all presented recommendations and whether or not they have been followed. We thus decided to use tailored application server logging to record this information. Application server logging also enables us to apply effective approaches for session and user identification and early elimination of crawler accesses, thus supporting good data quality (see later discussion).

We use two log files: a web usage and a recommendation log file with the formats shown in Table 2. The recommendation log file records all presented recommendations and is required for our recommender evaluation. It allows us to determine positive and negative user feedback, i.e. whether or not a presented recommendation was clicked. For each presented recommendation, we also record the relative position of the recommendation on the page, the generating recommender and the used selection strategy.

An example of the web usage log file is shown in Table 3. User IDs are stored inside permanent cookies and are used for user identification. If the user does not accept permanent cookies, user recognition is not done. The session ID is generated by the application server and stored inside a temporary cookie on the user's computer (if enabled). It is used for data preparation inside the ETL process, which is described in the following section.

3.2 ETL process

The ETL process periodically processes the log files of the application server and other data sources to refresh the data warehouse. Like in other data warehouse projects the main ETL tasks are data cleaning and transformation to ensure good data quality for the warehouse. The specifics of web usage data require tailored data preparation steps which have been investigated thoroughly in recent years; [5] gives a good overview. We illustrate the AWESOME approaches for two data preparation problems: crawler detection and session reconstruction.

Crawler detection deals with the identification of page requests not originating from human users but programs




Fig. 2 Number of monthly page views (2004)

such as search engine crawlers or website copiers. Obviously, we do not want to generate recommendations for such programs and only consider human navigation behavior. We use a combination of several heuristics to detect crawler accesses and avoid their logging or eliminate them from the log. First, we utilize a list of IP addresses of known crawlers (e.g., `crawl.googlebot.com` as a crawler from the Google search engine). Similarly, we provide a list of known agent names of website copiers (e.g., `HTTrack`¹) to identify and eliminate their requests. In addition, on every delivered page the application server integrates a link that is invisible to humans. Sessions containing a request to this linked page are deleted. Finally, we analyze the navigation behavior by measuring attributes like session length, average time between two requests and number of requests with blank referrer to distinguish between human user sessions and web crawler sessions [6]. To illustrate the effectiveness of our crawler detection we performed a comparison with a freely available web log file analysis tool² which currently does not eliminate crawler accesses. Figure 2 compares the number of all requested pages (measured by the free web server log file analysis program Webalizer) and the number of page views determined with this tool compared with the number of page views determined by AWESOME after elimination of crawler accesses for the same website and time period. The comparison illustrates that about 70% of all page views are identified as non-human requests for the considered website and that this share may actually vary significantly. This demonstrates the need and importance of an exhaustive crawler detection and that not performing this kind of data cleaning leads to irrelevant data and analysis results.

Another important data preparation task is session reconstruction to group page views from the same user during a website visit. Knowledge about sessions is obviously needed to analyze the navigation behavior of users as well as to determine whether presented recommendations are accepted. In the terminology of [7], AWESOME supports both a proactive and reactive session reconstruction. Proactive session reconstruction is done by a session ID generated by the application server and stored within temporary cookies. If temporary cookies are enabled (which is the case for about 85% of the users of our prototype website) they provide a safe method to session construction. The reactive ap-

proaches are based on heuristics and come into play when temporary cookies are disabled. In a first heuristic, we use the combination of the user's host name (actually only the last four parts of it to deal with varying host names, e.g., from proxy servers) and agent name as a temporary identifier. All accesses with the same identifier and a temporal distance within a threshold are considered to belong to the same session. Additionally, we use the referrer information to reconstruct sequences of page views. Given the high share of proactive session construction and the typically high accuracy of the heuristics we expect that the vast majority of sessions can be correctly constructed.

3.3 Data warehouse design

The web data warehouse is a relational database with a “galaxy” schema consisting of several fact tables sharing several dimensions. Like in previous approaches on web usage analysis [8] we use separate fact tables for page views, sessions, and – for commercial sites – purchases. In addition we use a recommendation fact table. The simplified schema in Fig. 3 shows the page view and recommendation fact tables and their dimension tables which are referenced by the respective ID attributes. The details of these tables depend on the website, e.g. on how the content (e.g., products), users or customers are categorized. In the example of Fig. 3 there are two content dimensions for different hierarchical categorizations. Such hierarchies may represent different kinds of product catalogs, e.g., organized by topic or by media type. Other dimensions such as user, customer, region, and date are also hierarchically organized and thus allow evaluations and recommendation decisions at different levels of detail.

The page view fact table represents the web usage history and contains a record for each requested page. The duration until the next page view of the same session is stored in the view time attribute. The recommendation fact table represents the positive and negative user feedback on recommendations. Each record in this table refers to one presented recommendation. The recommended content is characterized by the content dimension tables, whereas the utilized recommender and the applied selection strategy are described using additional dimension tables. Three Boolean measures are used to derive recommendation quality metrics (see Sect. 5). *Accepted* indicates whether or not the recommendation was directly accepted (clicked) by the user, while *Viewed* specifies whether the recommended content was viewed later during the respective session (i.e., the recommendation was a useful hint). *Purchased* is only used for e-commerce websites to indicate whether or not the recommended product was purchased during the current session.

4 Recommenders

A recommender generates for a given web page request, specified by a context, an ordered list of recommendations.

¹ <http://www.httrack.com/>

² <http://www.mrunix.net/webalizer/>

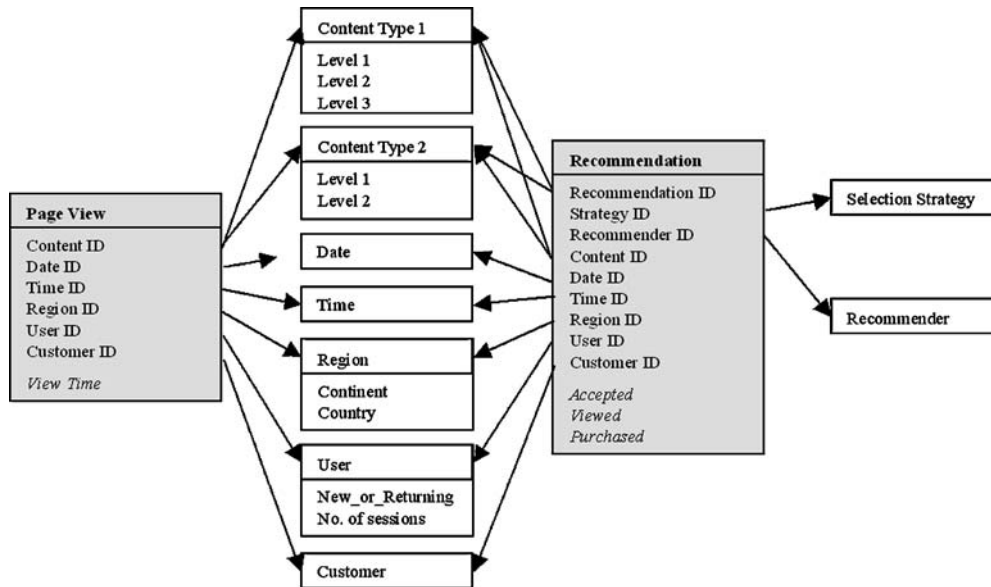


Fig. 3 Simplified data warehouse galaxy schema

Such recommendations link to current website content, e.g. pages describing a product or providing other information or services.

To calculate recommendations, recommenders can make use of the information available in the context as well as additional input, e.g. recorded purchase and web usage data. We distinguish between three types of context information relevant for determining recommendations:

- Current content, i.e. the currently viewed content (page view, product, etc.) and its related information such as content categories.
- Current user, e.g. identified by a cookie, and associated information, e.g. her previous purchases, previous web usage, interest preferences, or current session.
- Additional information available from the HTTP request (current date and time, user’s referrer, etc.)

4.1 Recommender classification

Given the many possibilities to determine recommendations, there have been several attempts to classify recommenders [2, 9, 10, 11]. These classifications typically started from a given set of recommenders and tried to come up with a set of criteria covering all considered recommenders. This led to rather complex and specialized classifications with criteria that are only relevant for a subset of recommenders. Moreover, new recommenders can easily require additional criteria to keep the classification complete. For example, [10] introduces a large number of specialized criteria for e-commerce recommenders such as input from target customers, community inputs, degree of personalization, etc.

To avoid these problems we, propose a general top-level classification of website recommenders focusing on

the usable input data, in particular the context information. This classification may be refined by taking additional aspects into account, but already leads to a distinction of major recommender types thus illustrating the design space. Moreover, the classification helps to compare different recommenders and guides us in the evaluation of different approaches.

Figure 4 illustrates our recommender classification and indicates where sample approaches fit in. We classify recommenders based on three binary criteria, namely whether or not they use information on the current *content*, the current *user*, and recorded usage (or purchase) *history* of users. This leads to a distinction of eight types of recommenders (Fig. 4). We specify each recommender type by a three-character-code describing whether (+) or not (–) each of the three types of information is used. For instance, type [+,-,-] holds for recommenders that use information on the current content and current user, but do not take into account user history.

The first classification criteria considers whether or not a recommender uses the current content, i.e. the currently requested page or product. A sample content-based approach (type [+,-,-]) is to recommend content that is most *similar* to the current content, e.g. based on text-based similarity metrics such as TF/IDF. Content-based recommenders may also use generalized information on the content category (e.g., to recommend products within the current content category). Sample content-insensitive recommenders (type [-,-,-]) are to recommend the *most recent* content, e.g. added within the last week, or to give a fixed recommendation at each page, e.g. for a special offer.

At the second level, we consider whether or not a recommender utilizes information on the current user. User-based approaches could thus provide special recommendations for specific user subsets, e.g. returning users or customers, or

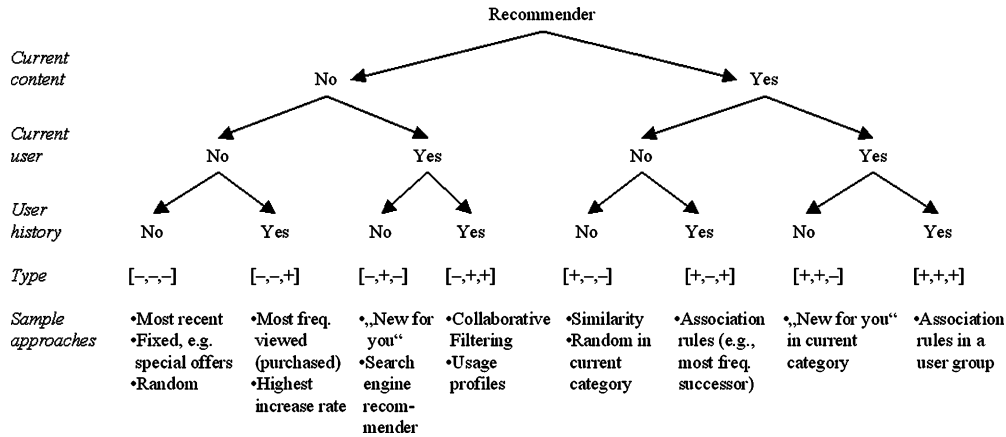


Fig. 4 Top-level classification of recommenders

based on personal interest profiles. Recommenders could also recommend content for individuals, e.g. new additions since a user’s last visit (“*New for you*”). We developed a new recommender of type $[-,+,-]$ for users coming from a search engine such as Google. This *search engine recommender (SER)* utilizes that the HTTP referrer information typically contains the search terms (keywords) of the user [12]. SER recommends the website content (different from the current page that was reached from the search engine) that best matches these keywords. The SER implementation in AWESOME utilizes a predetermined search index of the website to quickly provide the recommendations at runtime.

With the third classification criteria, we differentiate recommenders by their use of *user history* information. For commercial sites, recommenders can consider information on previous product purchases of customers. Another example is the evaluation of the previous navigation patterns of website users. Simple recommenders of type $[-,-,+]$ recommend the *most frequently* purchased/viewed content (top-seller) or the content with the *highest recent increase* of interest.

While not made explicit in the classification, recommenders can utilize additional information than on current content, current user or history, e.g. the current date or time. Furthermore, additional classification criteria could be considered, such as metrics used for ranking recommendations (e.g. similarity metrics, relative or absolute access/purchase frequencies, recency, monetary metrics, etc.) or the type of analysis algorithm (simple statistics, association rules, clustering, etc.).

4.2 Additional approaches

Interesting recommenders often consider more than one of the three main types of user input. We briefly describe some examples to further illustrate the power and flexibility of our classification and to introduce approaches that are considered in our evaluation.

$[+,-,+]$: *Association rule* based recommenders such as “Users who bought this item also bought...” made famous by Amazon [3], consider the current content (item)

and purchase history but are independent of the current user (i.e. every user sees the same recommendations for an item). Association rules can also be applied on web usage history to recommend content which is frequently viewed together within a session.

$[-,+,-]$: Information on navigation/purchase history can be used to determine *usage profiles* [13] or groups of similar users, e.g. by *collaborative filtering* approaches. Recommenders can assign the current user to a user group (either based on previous sessions or the current session) and recommend content most popular for this group.

In our evaluation, we test a *personal interests* recommender, which is applicable to returning users. It determines the most frequently accessed content categories per user as an indication of her personal interests. When the user returns to the website, the most frequently accessed content of the respective categories is recommended.

$[+,+,-]$: A recommender of this type could use both user groups (as discussed for $[-,+,-]$) and association rules to recommend the current user those items that were frequently accessed (purchased) by similar users in addition to the current content.

5 Recommender evaluation

The AWESOME prototype presented in Sect. 2 allows us to systematically evaluate recommenders for a given website. In Sect. 5.2, we demonstrate this for a sample non-commercial website. Before that, we introduce several metrics for measuring recommendation quality which are needed for our evaluation of recommenders and selection strategies.

5.1 Evaluation metrics

To evaluate the quality of presented recommendations, we utilize the Accepted, Viewed, and Purchased measures recorded in the recommendation fact table (Sect. 3.3). The first two are always applicable, while the last one only applies for commercial websites. We further differentiate

between metrics at two levels of granularity, namely with respect to page views and with respect to user sessions.

Acceptance rate is a straightforward, domain-independent metric for recommendation quality. It indicates the share of page views for which at least one presented recommendation was accepted, i.e. clicked. The definition thus is

$$AcceptanceRate = |P_A|/|P|$$

where P is the set of all page views containing a recommendation and P_A the subset of page views with an accepted recommendation.

Analogously, we define a session-oriented quality metric

$$SessionAcceptanceRate = |S_A|/|S|$$

where S is the set of all user sessions and S_A the set of sessions for which at least one of the presented recommendations was accepted.

Recommendations can also be considered of good quality if the user does not directly click them but reaches the associated content later in the session (hence, the recommendation was a correct prediction of user interests). Let P_V be the set of all page views for which any of the presented recommendations was reached later in the user session. We define

$$ViewRate = |P_V|/|P|$$

The corresponding metric at the session level is

$$SessionViewRate = |S_V|/|S|$$

where S_V is the set of all user sessions with at least one page view in P_V . Obviously, every accepted recommendation is also a viewed recommendation, i.e. $P_A \subseteq P_V \subseteq P$ and $S_A \subseteq S_V \subseteq S$, so that view rates are always larger than or equal to the acceptance rates.

In commercial sites, product purchases are of primary interest. Note that purchase metrics should be session oriented because the number of page views needed to finally purchase a product is of minor interest. A useful metric for recommendation quality is the share of sessions S_{AP} containing a purchase that followed an accepted recommendation of the product. Hence, we define the following metric:

$$RecommendedPurchaseRate = |S_{AP}|/|S|$$

Obviously, it holds $S_{AP} \subseteq S_A \subseteq S$.

5.2 Sample evaluation

We implemented and tested the AWESOME approach for recommender evaluation for a sample website, namely the website of our database group (<http://dbs.uni-leipzig.de>). We use two content hierarchies and Fig. 5 shows a fragment of one of them together with some numbers on the relative size and access frequencies. The website contains more than

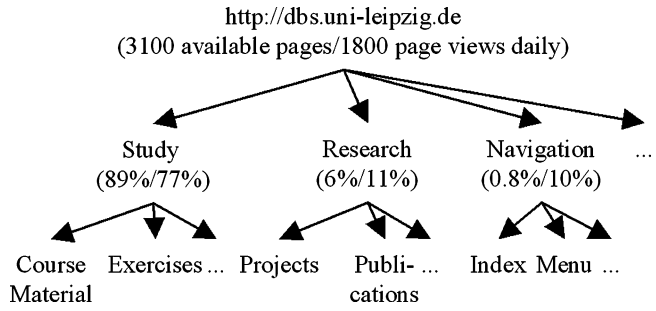


Fig. 5 Example of content hierarchy



Fig. 6 Recommendation screenshot

3,100 pages and receives about 1,800 human page views per day (excluding accesses from members of our database group and from crawlers). As indicated in Fig. 5, about 89% of the content is educational study material, which receives about 77% of the page views.

We changed the existing website to show two recommendations on each page so that approximately 3,600 recommendations are presented every day. For each page view AWESOME dynamically selects one recommender and presents its two top recommendations (see example in Fig. 6) for the respective context as described in Sect. 2. We implemented and included more than 100 recommenders in our recommender library. Many of them are variations of other approaches, e.g. considering different user categories or utilizing history data for different periods of time. Here we present results for the six representative recommenders of different types listed in Table 4, which were already introduced in Sect. 4. The presented results refer to the period from December 1, 2003 until April 30, 2004.

Table 4 Acceptance rate vs. user type


Recommender		User type		
Type	Name	New users (%)	Returning users (%)	Σ (%)
[-,-,-]	Most recent	(0.72)	(2.25)	(0.93)
[-,-,+]	Most frequent	1.05	1.19	1.09
[-,+,-]	SER	2.83	1.42	2.77
[-,+,+]	Personal interests	-	1.43	1.43
[+,-,-]	Similarity	1.86	0.76	1.73
[+,-,+]	Association rules	1.61	1.49	1.59
	Σ	1.85	1.29	1.74

The AWESOME warehouse infrastructure allows us to aggregate and evaluate recommendation quality metrics for a huge variety of constellations (combination of dimension attributes), in particular for different recommenders. For our evaluation we primarily use (page view) acceptance rates as the most precise metric.³ The average acceptance rate for all considered recommenders was 1.44%; the average view rate was 15.89%. For sessions containing more than one page view the average session acceptance rate was 9.03%, and the session view rate was 25.24%. These rather low acceptance rates are influenced by the fact that every single web page contains a full navigation menu with 78 links (partly nested in sub-menus) and that we consciously do not highlight recommendations using big fonts or the like. Note however, that reported “click-tru” metrics are in a comparable range than our acceptance rates [14]. Furthermore, the absolute values are less relevant for our evaluation than the relative differences between recommenders.

Table 4 shows the observed acceptance rates for the six recommenders differentiating between new and returning users. Figure 7 compares the recommenders w.r.t. the current page type. Please note that some recommenders are not always applicable. For instance, the *personal interests* recommender is only applicable for returning users (about 15% for our website). Similarly, SER can only be applied for users coming from a search engine, 95% of which turned out to be new users of the website. We only show results for recommenders with a minimum support of 5% i.e. they were applied for at least 5% of all page views of the respective page type. In Table 4 the results for the *most recent* recommender are shown in parentheses because the minimal support could not be achieved due to relatively few content additions during the considered time period.

As expected, there are significant differences between recommenders. As Table 4 shows, for our website the *search engine* recommender (SER) achieved the best average acceptance rates (2.77%), followed by the *similarity* and *association rules* recommenders. On the other hand, simple approaches such as recommending the *most frequently* accessed or *most recent* content achieved only poor average re-

³ The recommendations presented during a session typically come from different recommenders making the session-oriented quality metrics unsuitable for evaluating recommenders. Session acceptance rates will be used in Sect. 7. The *RecommendedPurchaseRate* does not apply for non-commercial sites.

**Fig. 7** Acceptance rate vs. page type

sults. Table 4 shows that new users are more likely to accept recommendations than returning users. An obvious explanation is that returning users (e.g., students for our website) often know where to find relevant information on the website. We also observed that the first page view of a session has a much higher acceptance rate (5.77%) than later page views in a session (1.35%). In the latter value, the last page view of a session is not considered, because its acceptance rate obviously equals 0.⁴

Figure 7 illustrates that the relative quality of recommenders differs for different contexts. While the SER recommender achieved the best average results for study pages, the *association rules* and *personal interests* recommenders received the best user feedback on navigation pages. For study pages and non-search engine users (when SER is not applicable), either the *similarity* or *association rules* promise the best recommendation quality.

Although we proposed a general top-level classification of recommenders, we can consider additional (recommender type specific) attributes to characterize recommenders. Such attributes are helpful for recommender fine-tuning, i.e. manual optimization of single recommenders regarding their

⁴ Layout aspects and other factors also influence acceptance rates. For instance, from the two recommendations shown per page the acceptance rate of the first one was about 50% higher compared to the second recommendation.




Fig. 8 Page view acceptance rate for different types of most frequent and association rule recommenders.

functionality. Figure 8 illustrates fine-tuning experiments for the *most frequent* and *association rules* recommenders. For the *most frequent* recommender we varied the time window of the considered user history from 24 h to 7 days. The results obtained indicate that for our website a user history of the last 7 days is to be preferred. For the *association rule* recommender we considered the influence of the relative page view position. For a given page (*direct*) successor only takes into account the (directly) following page views, whereas *normal* utilizes all page views inside the same session. The results shown in Fig. 8 indicate that the differences between the approaches are small for our website and that it suffices to only consider the direct successors by the association rule recommender.

While these observations are site specific, they illustrate that the best recommender (and the best recommendation respectively) depends on context attributes such as the current content or current user. A careful OLAP analysis may help to determine manually which recommender should be selected in which situation. However, for larger and highly dynamic websites this is difficult and labor intensive, so that recommender and recommendation selection should be automatically optimized.

6 Adaptive website recommendations

AWESOME supports a dynamic selection of recommendations for every website access. As discussed in Sect. 2, this selection is either a two-step or one-step process controlled by recommender or recommendation rules, respectively. Such rules may either be manually defined or automatically generated. We first present the structure and use of these two types of selection rules. In Sect. 6.2, we propose two approaches to automatically generate recommender rules that utilize recommendation feedback to adapt to changing conditions. In Sect. 6.3, we present an approach to automatically generate and adapt recommendation rules.

{ Usertype='new user' AND ContentCategory1='Navigation' }	→	'Most frequent'	{0.6}
{ Referrertype='search engine' }	→	'SER'	{0.8}
{ Clienttype='university' AND Usertype='returning user' }	→	'Personal interest'	{0.4}

Fig. 9 Examples of recommender rules

6.1 Rule-based selection of recommendations

Determination of adaptive website recommendations entails the dynamic selection of the most promising recommendations. Since AWESOME supports two ways of recommendation selection, it utilizes two types of selection rules of the following structure:

1. Recommender rule:

$$ContextPattern \rightarrow recommender [weight]$$

2. Recommendation rule:

$$ContextPattern \rightarrow recommendation [weight]$$

Here context pattern is a sequence of values from different context attributes which are represented as dimension attributes in our warehouse. Typically, only a subset of attributes is specified, implying that there is no value restriction for the unspecified attributes. The examples in Fig. 9 illustrate this for recommender rules.

On the right-hand side of recommender rules, *recommender* uniquely identifies an algorithm of the recommender library. Conversely, recommendation rules directly point to single recommendations of the set of pre-computed recommendations. *Weight* is a real number specifying the relative importance of the rule.

In AWESOME, we maintain all selection rules in a *Recommender-Rules* and *Recommendation-Rules* table respectively. In fact, for the one-step selection, the selection rules and pre-computed recommendations fall together. Hence, the recommendation-rule table corresponds to the set of pre-computed recommendations extended with the weight attribute which may be dynamically adapted. To simultaneously allow one-step and two-step recommendation selection, we keep the recommendation-rule table separate from the pre-computed recommender-specific recommendations.

Selection rules allow a straightforward and efficient implementation of the selection process. It entails a *match* step to find all rules with a context pattern matching the current context. The rules with the highest weights then indicate the recommenders/recommendations to be applied. The number of recommenders/recommendations to choose is typically fixed, say *k*, i.e., we choose the *k* rules with the highest weight. Note, that *k* is usually greater for recommendation rules than for recommender rules, since each recommender can provide multiple recommendations per context.

The SQL query of Fig. 10 illustrates how the sketched selection process can be done for the case of recommender rules. Since there may be several matching rules per recommender, the ranking could also be based on the average instead of the maximal weight per recommender.

```

SELECT TOP k Recommender, MAX (Weight)
FROM RecommenderRules
WHERE ((RuleContextAttribute1 = CurrentContextAttribute1)
      OR (RuleContextAttribute1 IS NULL))
AND   ((RuleContextAttribute2 = CurrentContextAttribute2)
      OR (RuleContextAttribute2 IS NULL))
AND ...
GROUP BY Recommender
ORDER BY MAX(Weight) DESC

```

Fig. 10 Simplified SQL query for recommender selection strategy execution

Example 1 Consider a new user who reaches the website from a search engine. If her current page belongs to the navigation category, only the first two rules in Fig. 9 match. For $k = 1$, we select the recommender with the highest weight – SER.

The rule-based selection is highly flexible. Selection rules allow the dynamic consideration of different parts of the current context, and the weights can be used to indicate different degrees of certainty. Rules can easily be added, deleted or modified independently from other rules. Moreover, rules can be specified manually, e.g. by website editors, or be generated automatically. Another option is a hybrid strategy with automatically generated rules that are subsequently modified or extended manually, e.g. to enforce specific considerations.

6.2 Generating recommender rules

We present two approaches to automatically generate recommender rules, which have been implemented in AWESOME. Both approaches use the positive and negative feedback on previously presented recommendations. The first approach uses the aggregation and query functionality of the data warehouse to determine selection rules. The second approach is more complex and uses a machine-learning algorithm to learn the most promising recommender for different context constellations.

6.2.1 Query-based top recommender

This approach takes advantage of the data warehouse query functionality. It generates recommender rules as follows:

1. Find all relevant context patterns in the recommendation fact table, i.e. context patterns exceeding a minimal support.
2. For every such context pattern P do
 - (a) find recommender R with highest acceptance rate A ,
 - (b) add recommender rule $P \rightarrow R [A]$.
3. Delete inapplicable rules.

The first step ensures that only context constellations with a minimal number of occurrences are considered. This is important to avoid generalization of very rare and special situations (overfitting problem). Note that step 1 checks all

possible context patterns, i.e. any of the content attributes may be unspecified, which is efficiently supported by the CUBE operator (SQL extension: GROUP BY CUBE) [15]. AWESOME is based on a commercial RDBMS providing this operator. For every such context pattern, we run a query to determine the recommender with the highest acceptance rate and produce a corresponding selection rule.

Finally, we perform a rule pruning taking into account that we only want to determine the top recommender per context. We observe that for a rule A with a more general context pattern and a higher weight than rule B, the latter will never be applied (every context that matches rule B also matches rule A, but A will be selected due to its higher weight). Hence, we eliminate all such inapplicable rules in step 3 to limit the total number of rules.

6.2.2 Machine-learning approach

Recommender selection can be interpreted as a classifier selecting one recommender from a predefined set of recommenders. Hence, machine-learning (classification) algorithms can be applied to generate recommender rules. Our approach utilizes a well-known classification algorithm constructing a decision tree based on training instances (Weka J48 algorithm [16]). To apply this approach, we thus have to transform recommendation feedback into training instances. An important requirement is that the generation of training data must be completely automatic so that the periodic recalculation of recommender rules to incorporate new recommendation feedback is not delayed by the need of human intervention.

The stored recommendation feedback indicates for each presented recommendation, its associated context attributes, the used recommender, and whether or not the recommendation was accepted. A naïve approach to generate training instances would simply select a random sample from the recommendation fact table (Fig. 3), e.g. in the format (*context, recommender, accepted*). However, classifiers using such training instances would rarely predict a successful recommendation, since the vast majority of the instances may represent negative feedback (>98% for the sample website). Ignoring negative feedback is also no solution, since the number of accepted recommendations is heavily influenced by the different applicability of recommenders and not only by their recommendation quality. Therefore, we propose a more sophisticated approach that determines the number of training instances according to the acceptance rates:

1. Find all relevant feedback combinations (*context, recommender*).
2. For every combination c do the following:
 - (a) Determine acceptance rate for c . Scale and round it to compute integer weight n_c .
 - (b) Add instance (*context, recommender*) n_c times to training data.
3. Apply decision tree algorithm.
4. Rewrite decision tree into recommender rules.

In step 1, we do not evaluate context patterns (as in the previous approach), which may leave some context attributes unspecified. We only consider fully specified context attributes and select those combinations exceeding a minimal number of recommendation presentations. For each such relevant combination c (*context, recommender*), we use its acceptance rate to determine the number of training instances n_c . To determine n_c , we linearly scale the respective acceptance rate from the 0 to 1 range by multiplying it with a constant f and rounding to an integer value. For example, assume 50 page views for the combination of context (“returning user,” “search engine,” “navigation,” etc.) and recommender “most frequent.” If there are seven accepted recommendations for this combination (i.e. acceptance rate 0.14) and $f = 100$, we add $n_c = 14$ identical instances of the combination to the training data. This procedure ensures that recommenders with a high acceptance rate produce more training instances than less effective recommenders and therefore have a higher chance to be predicted.

The resulting set of training instances is the input for the classification algorithm producing a decision tree. With the help of cross-validation, all training instances are simultaneously used as test instances. The final decision tree (see Fig. 11 for an example) can easily be rewritten into recommender rules. Every path from the root to a leaf defines a context pattern where all unspecified context attributes are set to NULL. Each leaf specifies a recommender and the rule weight is set to the relative fraction of correctly classified instances provided by the classification algorithm.

6.3 Generating and adapting the recommendation rules

AWESOME also supports a dynamic one-step selection of recommendations, independent of the recommenders. This process is controlled by recommendation rules, which may be continuously adapted. We introduced the underlying approach already in [17, 18]. We integrated the one-step approach with the two-step approach in one system to allow their comparative evaluation.

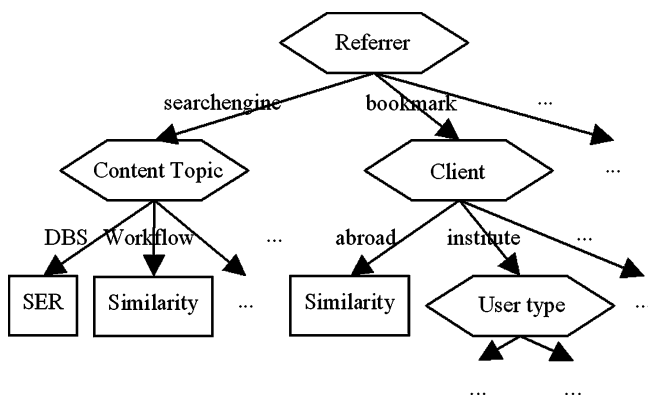


Fig. 11 Fraction of decision tree constructed by the machine-learning approach to generate recommender rules

The recommendation rules are created by the same set of recommenders from the recommender library and are stored in a *recommendation rule table*. When several recommenders make the same recommendation for the same context, the recommendation rule is stored in the recommendation rule table only once. In contrast to the recommender rules, we support both an offline and online adaptation of recommendation rules. During the periodic offline computation of recommendations we do not completely regenerate all recommendation rules but only add new recommendations to preserve the dynamically accumulated recommendation feedback. We also delete recommendations for deleted content and with low weights. The online adaptation of the recommendation rules dynamically adjusts the rule weights according to the recommendation feedback. The adaptation technique is based on reinforcement learning and is described later.

We have explored two possibilities of setting the *initial weights* of newly generated recommendation rules. In the first approach, we simply set all initial weights to zero (ZeroStart). The second approach uses normalized recommender-specific weights or relative priorities for the respective contexts. When several recommenders generate the same recommendation we use the maximum of their weights. The initial weights are expected to be relevant primarily for new recommendations, since the weights for presented recommendations are continuously adapted.

To adapt the shown recommendations to the users’ interests, we adjust the weights of the recommendation rules in such a way that the more useful recommendations are shown more often than less useful ones. This is achieved by a machine-learning approach based on *reinforcement learning* [19]. The adaptation process evaluates whether or not presented recommendations have been accepted and adjusts the weights of the participating recommendation rules according to the obtained feedback. When some presented recommendation r is clicked, r receives positive feedback and all other recommendations shown together with r receive negative feedback. When no recommendation is clicked, all presented recommendations receive negative feedback. To prevent the weights from sliding into extreme values, the feedback values should be chosen in such a way that an approximate equilibrium is maintained throughout the process:

$$\sum (\text{positive feedback}) \approx - \sum (\text{negative feedback})$$

For this purpose, we set for a presented recommendation r

$$\text{Feedback}(r) = 1 \text{ if } r \text{ was clicked}$$

$$\text{Feedback}(r) = -p \text{ if } r \text{ was not clicked}$$

where p is the overall acceptance rate of the website, i.e. the probability that a recommendation is clicked (default $p = 0, 01$).

After each presentation, for every presented recommendation r we adapt its weight $W(r)$ by the following formula:

$$W(r) = (1 - 1/T) \times W(r) + \text{Feedback}(r)/T.$$

Table 5 Tested selection strategies

Name	Description
Top-Rec	Automatic strategy of Sect. 6.2.1 (query based)
Decision tree	Automatic strategy of Sect. 6.2.2 (machine learning)
Reinforcement learning	Automatic strategy of Sect. 6.3 (reinforcement learning)
Reinforcement learning zero	Automatic strategy of Sect. 6.3 (reinforcement learning, with initial weights set to zero)
Manual	For search engine users, the search engine recommender is applied. Otherwise the content similarity recommender (for course material pages) or association rule recommender (for other pages) is selected
Random	Random selection of a recommender

Table 6 Comparison of selection strategies

Strategy	No. of rules	Acceptance rate (%)	Session acceptance rate (%)
Top-Rec	~2,000	1.35	10.27
Decision tree	~250	1.74	11.13
Reinforcement learning	~60,000	1.92	11.22
Reinforcement learning zero	~60,000	1.87	10.35
Manual	5	1.97	12.54
Random	137	0.96	6.98

The parameter T is an aging parameter reducing the weight of the old value of the weight compared to the more recent feedback.

7 Evaluation of selection strategies

To evaluate the effectiveness of the presented selection strategies, we tested them with AWESOME on the sample website introduced in Sect. 5.2. For comparison purposes, we also evaluated one set of manually specified rules and a random recommender selection giving a total of six approaches (see Table 5). For every user session AWESOME randomly selected one of the strategies; the chosen strategy is additionally recorded in the recommendation log file for evaluation purposes. We applied the recommender selection strategies from January 1 until June 30, 2004 and the recommendation selection strategies from April 1 until September 30, 2004.

Table 6 shows the average number of rules per selection strategy. For the automatic recommender selection approaches the number of rules is moderate (250–2000), whereas both recommendation selection approaches generate significantly more rules. Nevertheless, all tested selection strategies still lead to acceptable execution times (10–110 ms).

Table 6 also shows the average (page view) acceptance rates and session acceptance rates for the six selection strategies. All automatic feedback-based strategies showed significantly better average quality than random. The decision tree strategy and both reinforcement learning strategies nearly obtain the acceptance rates of the Manual strategy. Note that the very effective strategy Manual utilizes background knowledge about the website structure and typical user groups (students, researchers) as well as evaluation results obtained after an extensive manual OLAP analysis

(partially presented in Sect. 5.2), such as the effectiveness of the search engine recommender.

The fact that the completely automatic machine-learning algorithms achieves comparable effectiveness is thus a very positive result. It indicates the feasibility of the automatic closed-loop optimization for generating recommendations and the high value of using feedback to significantly improve recommendation quality without manual effort. In our experiments, the one-step approaches based on reinforcement learning slightly outperformed the two-step recommender-based strategies. Even the simple approach starting with weights 0 achieved surprisingly good results, indicating that the initial weights are not crucial for the success of the adaptation algorithm. The effectiveness of the one-step approach is mainly due to the online consideration of recommendation feedback providing a high responsiveness to current user behavior.

The comparison of the two automatic recommender selection strategies shows that the machine-learning approach performs much better than the query-based top recommender scheme. The decision tree approach uses significantly fewer rules and was able to order the context attributes according to their relevance. The most significant attributes appear in the upper part of the decision tree and therefore have a big influence on the selection process. On the other hand, Top-Rec handles all context attributes equally and uses many more rules. So recommender selection was frequently based on less relevant attributes, resulting in poorer acceptance rates.

The warehouse infrastructure of AWESOME allows us to analyze the recommendation quality of selection strategies for many conditions, similar to the evaluation of individual recommenders (Sect. 5.2). Figure 12 shows the session acceptance rates of three selection strategies w.r.t. user type, referrer, and entry page type, i.e. the page type of the first session page view. We observe that the Manual strategy is more effective for search engine users by always

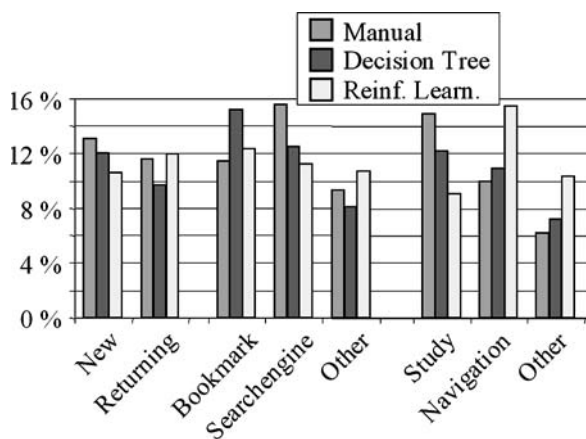


Fig. 12 Session acceptance rate w.r.t. user type, referrer, and session entry page type.

applying the SER recommender to them. This helped to also get slightly better results for new users and sessions starting with an access to study material. On the other hand, both automatic approaches were more effective for users using a bookmark to reach the website. These results indicate that the automatically generated selection rules help generate good recommendations in many cases without the need of extensive manual evaluations, e.g. using OLAP tools.

The two automatic approaches show significant differences for sessions starting with a study page and a navigation page. As indicated in Fig. 5, our prototype website contains a large number of study pages, whereas the number of navigation pages is rather small. However, navigation pages receive almost 15 times more feedback per page than study pages. Therefore, the recommendation-based selection approach can easily identify the best recommendations for navigation pages, but not for study pages. On the other hand, the feedback aggregation of the recommender-based approach can better handle this lack of feedback, but loses information to generate better recommendations for navigation pages.

The evaluation also identifies the capabilities of the one-step and two-step approach. The usage of recommender rules makes it possible to completely define a selection strategy by manually specified rules. Due to the huge number of rules this is not feasible for the one-step approach. Nevertheless, recommendation rules can be added manually to directly influence the recommendations for single products or pages. This can be done to promote certain products or to incorporate reasonable recommendations that are not detected by any of the available recommenders. Finally, note that all presented results are site specific. For example, the frequency of content updates on the website may influence acceptance rates to an unequal extent for the one-step and two-step approach.

8 Related work

An overview of previous recommendation systems and the applied techniques can be found in [1, 2, 20]. Reference [3]

describes the Amazon recommendation algorithms, which are primarily content (item) based and also heavily use pre-computation to achieve scalability to many users. Reference [9] surveys and classifies so-called hybrid recommendation systems which combine several recommenders. To improve hybrid recommendation systems, reference [21] proposes to manually assign weights to recommenders to influence recommendations. Reference [22] presents a hybrid recommendation system switching between different recommenders based on the current page's position within a website. The Yoda system [23] uses information on the current session of a user to dynamically select recommendations from several predefined recommendation lists. In contrast to AWESOME, these previous hybrid recommendation systems do not evaluate or use recommendation feedback.

Reference [24] sketches a simple hybrid recommendation system using recommendation feedback to a limited extent. They measure which recommendations produced by three different recommenders are clicked to determine a weight per recommender (with a metric corresponding to our view rate). These weights are used to combine and rank recommendations from the individual recommenders. In contrast to AWESOME negative recommendation feedback and the current context are not considered for recommender evaluation. Moreover, there is no automatic closed-loop adaptation but the recommender weights are determined by an offline evaluation.

The evaluation of recommendation systems and quantitative comparison of recommenders has received little attention so far. Reference [25] monitored users that were told to solve certain tasks on a website, e.g. to find specific information. By splitting users in two groups (with recommendations vs. without) the influence of the recommendation system is measured. Other studies [26, 27] asked users to explicitly rate the quality of recommendations. This approach obviously is labor intensive and cannot be applied to compare many different recommenders.

Refs [26, 28] discuss several metrics for recommendation quality, in particular the use of the information retrieval metrics precision and recall. The studies determine recommendations based on an offline evaluation of web log or purchase data; the precision metric, for instance, indicates how many of the recommendations were reached within the same session (thus corresponding to our view rate). In contrast to our evaluation, these studies are not based on really presented recommendations and measured recommendation feedback so that the predicted recommendation quality remains unverified.

In [29] a methodology is presented for evaluating two competing recommenders. It underlines the importance of such an online evaluation and discusses different evaluation aspects. Cosley et al. [14] developed the REFEREE framework to compare different recommenders for the CiteSeer website. Click metrics (e.g., how often a user followed a link or downloaded a paper), which are similar to the acceptance rates used in our study, are used to measure recommendation quality.

9 Summary

We presented AWESOME, a new data warehouse-based website evaluation and recommendation system. It allows the coordinated use of a large number of recommenders to automatically generate website recommendations. Recommendations are dynamically determined by a flexible rule-based approach selecting the most promising recommender/recommendations for the respective context. AWESOME supports a completely automatic generation and optimization of selection rules to minimize website administration overhead and quickly adapt to changing situations. This optimization is based on a continuous measurement of user feedback on presented recommendations. To our knowledge, AWESOME is the first system enabling such a completely automatic closed-loop website optimization. The use of data warehouse technology and pre-computation of recommendations support scalability, high data quality, and fast web access times.

We presented a simple but general recommender classification. It distinguishes eight types of recommenders based on whether or not they consider input information on the current content, current user and users history. To evaluate the quality of recommendations and recommenders, we proposed the use of several acceptance rate metrics based on measured recommendation feedback. We used these metrics for a detailed comparative evaluation of different recommenders and different recommendation selection strategies for a sample website.

We have described and evaluated several rule-based strategies for dynamically selecting the most promising recommender or recommendations for a given context. We differentiated between one-step and two-step approaches depending on whether recommendations are directly selected or whether the most promising recommenders are determined first. Our results so far indicate that in both cases the use of machine learning is very effective for considering recommendation feedback. For a direct (one-step) selection of recommendations the described use of reinforcement learning effectively allows an online adaptation of individual recommendations regardless of the used recommenders. For automatic recommender selection, we presented a powerful decision tree approach. The proposed policy is able to automatically determine suitable training data so that its periodic re-execution to consider new feedback does not require human intervention.

We have begun to adopt AWESOME to additional websites, in particular e-shops, to further verify and fine-tune the presented approach. We also plan to explore domain-specific recommendation opportunities such as selecting the best recommender for product bundling (cross-selling).

Acknowledgements We thank Robert Lokaiczkyk for his help with the implementation. The first two authors are funded by the German Research Foundation within the Graduiertenkolleg "Knowledge Representation".

References

1. Jameson, A., Konstan, J., Riedl, J.: AI techniques for personalized recommendation. In: Tutorial at 18th National Conference on Artificial Intelligence (AAAI) (2002)
2. Koutri, M., Daskalaki, S., Avouris, N.: Adaptive interaction with web sites: an overview of methods and techniques. In: Proceedings of the 4th International Workshop on Computer Science and Information Technologies (CSIT) (2002)
3. Linden, G., Smith, B., York, J.: Amazon.com Recommendations: item-to-item collaborative filtering. *IEEE Distrib. Syst. Online* **4**(1) (2003)
4. Perkowitz, M., Etzioni, O.: Adaptive web sites: an AI challenge. In: Proceedings of the 15th International Joint Conference on Artificial Intelligence. Morgan Kaufmann (1997)
5. Cooley, R., Mobasher, B., Srivastava, J.: Data preparation for mining world wide web browsing patterns. *Knowledge Inform. Syst.* **1**(1) (1999)
6. Tan, P., Kumar, V.: Modeling of web robot navigational patterns. In: Proceedings of the ACM WebKDD Workshop (2000)
7. Spiliopoulou, M., Mobasher, B., Berendt, B., Nakagawa, M.: A framework for the evaluation of session reconstruction heuristics in web usage analysis. Special issue on mining web-based data for e-business applications. *INFORMS J. Comput.* **15**(2) (2003)
8. Kimball, R., Merz, R.: *The Data Warehouse Toolkit – Building Web-Enabled Data Warehouse*. Wiley Computer Publishing, New York (2000)
9. Burke, R.: Hybrid recommender systems: survey and experiments. *User Model. User-Adapt. Interact.* **12**(4) (2002)
10. Schafer, J.B., Konstan, J.A., Riedl, J.: Electronic commerce recommender applications. *J. Data Min. Knowledge Discov.* **5**(1/2) (2001)
11. Terveen, L., Hill, W.: Human-computer collaboration in recommender systems. In: Carroll, J. (ed.) *Human Computer Interaction in the New Millennium*. Addison-Wesley, New York (2001)
12. Kushmerick, N., McKee, J., Toolan, F.: Toward zero-input personalization: referrer-based page recommendation. In: Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems (2000)
13. Mobasher, B., Dai, H., Luo, T., Nakagawa, M.: Discovery and evaluation of aggregate usage profiles for web personalization. *Data Min. Knowledge Discov.* **6**(1) (2002)
14. Cosley, D., Lawrence, S., Pennock, D.M.: REFEREE: an open framework for practical testing of recommender systems using ResearchIndex. In: Proceedings of the 28th VLDB Conference (2002)
15. Gray, J., Bosworth, A., Layman, A., Pirahesh, H.: Data cube: a relational aggregation operator generalizing groupby, cross-tab, and sub-total. In: Proceedings of the 12th IEEE International Conference on Data Engineering (ICDE) (1995)
16. Witten, I.H., Frank, E.: *Data Mining. Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann (2000)
17. Golovin, N., Rahm, E.: Reinforcement learning architecture for web recommendations. In: Proceedings of the International Conference on Information Technology (ITCC), Vol. 1, Las Vegas (2004)
18. Golovin, N., Rahm, E.: Automatic optimization of web recommendations using feedback and ontology graphs. In: Proceedings of the International Conference on Web Engineering (ICWE), Sydney, LNCS 3579, Springer Verlag (2005), submitted for publication
19. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA (1998)
20. Srivastava, J., Cooley, R., Deshpande, M., Tan, P.-T.: Web usage mining: discovery and applications of usage patterns from web Data. *SIGKDD Explor.* **1**(2) (2000)

21. Schafer, J.B., Konstan, J.A., Riedl, J.: Meta-recommendation systems: user-controlled integration of diverse recommendations. In: Proceedings of the 11th International Conference on Information and Knowledge Management (CIKM) (2002)
22. Mobasher, B., Nakagawa, M.: A hybrid web personalization model based on site connectivity. In: Proceedings of the ACM WebKDD Workshop (2003)
23. Shahabi, C., Chen, Y.: An adaptive recommendation system without explicit acquisition of user relevance feedback. *Distribut. Parallel Databases* **14**(2) (2003)
24. Lim, M., Kim, J.: An adaptive recommendation system with a coordinator agent. In: Proceedings of the First Asia-Pacific Conference on Web Intelligence: Research and Development (2001)
25. Kim, K., Carroll, J.M., Rosson, M.B.: An empirical study of web personalization assistants: supporting end-users in web information systems. In: Proceedings of the IEEE 2002 Symposium on Human Centric Computing Languages and Environments (2002)
26. Geyer-Schulz, A., Hahsler, M.: Evaluation of recommender algorithms for an internet information broker based on simple association rules and on the repeat-buying theory. In: Proceedings of the ACM WebKDD Workshop (2002)
27. Heer, J., Chi, E.H.: Separating the swarm: categorization methods for user sessions on the web. In: Proceedings of the Conference on Human Factors in Computing Systems (2002)
28. Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Analysis of recommendation algorithms for e-commerce. In: Proceedings of ACM E-Commerce (2000)
29. Hayes, C., Massa, P., Avesani, P., Cunningham, P.: An on-line evaluation framework for recommender systems. In: Proceedings of Workshop on Personalization and Recommendation in E-Commerce (2002)
30. Shahabi, C., Banaei-Kashani, F., Faruque, J.: A reliable, efficient, and scalable system for web usage data acquisition. In: Proceedings of the ACM WebKDD Workshop (2001)