Hybrid Personalization For Recommendations

Eelco Herder and Philipp Kärger
L3S Research Center
Hannover,Germany
{herder,kaerger}@L3S.de

Abstract

In this paper we present the concept of hybrid personalization, the combination of multiple atomic personalization mechanisms. The idea of hybrid personalization is related to hybrid recommender systems, but works on a conceptual level—it is decoupled from the actual adaptation in the user interface. This has as an advantage that one can optimize the adaptation 'behind the screens' or—conversely—attach a new visualization mechanism to the personalization technique. We show the practical benefits of this layered, hybrid adaptation mechanisms by means of a case study on personalized curriculum planning where it is recommended which course could or should be followed at which state in the learning process.

1 Introduction

Adaptive hypermedia systems are commonly defined as systems that reflect some features of the user in a *user model* and apply this model to *adapt* various visible aspects of the system to the user [2]. Typical adaptation techniques include text modification, the addition, removal or annotation of hyperlinks, personalization of search results and recommendations. These visible end results of the personalization process are based on a more conceptual *adaptation decision*, which on its turn is based on information on the users' preferences, needs, interests, location and background.

A conceptual adaptation decision, such as 'I want to provide my users with book recommendations' can be visualized in many different ways. Many online stores provide their users with an explicit list of 'Items Recommended For You' on the personalized portal page. Based on the same model of their customers, they often customize product listings, to satisfy their customers and to increase their sales. Conversely, one could base a list of recommendations on various different techniques, varying from semantic matching of book titles, topics or contents to collaborative techniques. Whether an adaptation is successful depends on both the adaptation decision process (in this example the selection of books to recommend) and the way it is shown to the user (explicit recommendations, reordering of results, emphasis on selected items) [12].

Whereas the concept of *layered evaluation* [22]—which separates the various steps that lead to the actual adaptation—is often discussed in the literature, in most adaptive systems there is no separation between the adaptation concept and the way it is visualized. We think this is a missed opportunity. In particular if there are many aspects that one can personalize for, it would be useful if one could easily switch between concepts or algorithms—or combine them—to optimize the actual personalization in the user interface. One could also allow the user to *adapt the adaptation*, by adapting weights or selecting aspects to take into account.

In particular in the field of e-learning, there is a need for more flexibility and control on the adaptation decision process. The relevance of courses or course elements, and the order in which they could or should be followed, depends on many factors. To name a few:

- the learners' knowledge, interests, goals and tasks, background, individual traits, context of work [4]
- pedagogic and formal constraints, as is often the case in formal curricula [21]
- preferences concerning, for example, where the course is given, on which day, in which language, by which teacher, kind of examination, costs, etc. [15]

The variety in approaches, assumptions and techniques explored by researchers in the field of adaptive educational hypermedia—as can be observed in the proceedings of recent conferences, such as EC-TEL[1], UM[2], and AH[3]—shows that there is no formula for success. In this paper we present a *hybrid* personalization approach, which allows flexible selection, combination and configuration of various *atomic* adaptation techniques. Our implementation of a Hybrid Personalization Service is used for positioning course offers in an interactive visualization pane, which allows students to effectively plan their curricula. Four atomic personalization services—of which two work on the course's metadata, and the remaining two use collaborative and statistical techniques—are used for creating an initial, personalized advice for a curriculum that the students can further adapt to their needs.

The remainder of this paper is structured as follows. In the next section we explain the theories and techniques related to hybrid personalization. In Section 3 we describe the issues associated with curriculum planning, which are addressed by our system. We argue that it is essential that the learners plan their curricula themselves—rather than using an automatic curriculum generator—and that they need the system to provide personalized feedback on the relevance of each option, from various perspectives. In Section 4 we present the Hybrid Personalization Service,

[1] European Conference on Technology Enhanced Learning
[2] Internation Conference on User Modeling
[3] International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems

its constituent atomic personalization mechanisms, and its user interface, the Graphical Curriculum Planning Tool. We end the paper with a discussion and concluding remarks.

2 Theoretical and Technical Background

Hybrid Web Recommender Systems is a relatively new and active research area [5]. A hybrid recommender system is one that combines multiple recommendation techniques, which may be collaborative, content-based, demographic or knowledge-based. By combining these various techniques one can leverage the problem of not having sufficient data on the user or on the content. There are various ways in which atomic recommender systems can be combined, among which [5]:

- the input of all atomic recommender systems is combined, with a certain *weighting scheme*
- the system *switches* between the available recommender systems
- the different recommendations are presented *next to one another*
- the recommender systems are used *consecutively*

Recommender systems are typically associated with the creation of top-*n* lists and the reordering of search results. However, one can do various other things with a ranked list of items. The ranking can be used for common personalization techniques, such as annotation, highlighting, cross-linking and personalized graphical overviews [3] [19]. Another advantage of decoupling the view from the model is that one can use multiple personalization techniques simultaneously—which implies that one does not have to summarize the (possibly orthogonal) outcomes of the atomic personalizers into one weighted average. Furthermore, by allowing the user to *adapt* the view, one can get different perspectives on the same set of objects. In earlier work, we created a system for the visualization of user navigation the Web [13], in which one could visualize various aspects of each page visit (duration, frequency, page size, Web site) with adaptable color coding, labels and markers.

In the past decade, several *layered frameworks* for adaptive systems have been developed, among which the AHA! Framework [9] and the LAOS Model [8]. E-learning systems that are meant for deployment increasingly support the SCORM standard [10], which allows the same content to be used in different systems ('to be played in different players'). Whereas these architectures are quite different from a conceptual point of view, frameworks that were meant for the *evaluation* of adaptive systems [22] differ mainly in level of granularity. In essence, they separate the adaptation process in the following different phases:

- in the *data acquisition phase* the information on user and context is gathered, for example by monitoring the interaction
- in the *knowledge inference phase* the data is transformed into a user model; in this phase, the data becomes 'meaningful'
- in the *adaptation decision making phase* the information in the user model is used for deciding on a (conceptual) adaptation; in most systems this phase is fairly straightforward in the form of simple if-then rules
- in the *adaptation execution phase* the actual adaptation is carried out

In most systems the adaptation execution phase is merged with the decision making phase. However, it has been shown in several studies (for example in [18]) that the way an item, such as a link, is visualized, has a strong impact on the user. Conversely, several studies—among which [7]—report that similar-looking lists, but based on different algorithms, may perform radically differently.

Within e-learning, we can see a shift from author-predefined adaptation rules to collaborative filtering techniques and the use of Web 2.0 interaction mechanisms [6]. With a huge pool of data, many candidate user groups to compare the user with, and several methods at hand, it becomes even more important to experiment with and optimize the conceptual adaptation decisions—in an iterative process, including user studies [17]—while keeping the interface itself constant.

3 Case Study: Curriculum Planning

Later in the paper we will present our prototype Hybrid Personalizer, which provides a flexible personalization solution for the TenCompetence Graphical Planning Tool for Competence Development Programs (CDP). In this section we provide some details on curriculum planning in order to appreciate the underlying personalization concept and the reason why we considered a hybrid combination of several atomic personalization services.

Whereas most current standards and tools for e-learning provide relatively sophisticated functionality for the management of learning activity, only little support is provided on the level of curricula. Universities and other educational institutes do have overviews of the courses that are given, or provide standard curricula. However, if a learner has goals that are somehow not standard, she needs to resort to the course descriptions and make a planning on her own. Mentors as well have little support for providing advice and need to resort to their experience. This problem becomes even more apparent if we leave the structured environment of (higher) education and concentrate on corporate learning and lifelong learning.

In corporate and lifelong learning, employers and employees can choose between potentially many course offerings, varying in:

- the competences that they provide
- the level of knowledge and skills that one learns (from beginner to advanced)
- the domain in which they are applicable
- the nature of courses and examination (from formal learning to informal workplace learning)
- availability, schedule, planning, costs
- etcetera

In the remainder of this paper, we call these kinds of curricula *competence development programs (CDPs)*, to indicate that they need not be (predefined) curricula and need not consist of units that were envisaged as courses (spending an afternoon with a colleague, or reading a book on a certain topic might do the trick as well).

The amount of available learning activities and their scopes may be overwhelming—in particular if several institutions combine their offerings (regional, national, international). For course designers this implies that they

need to indicate in which situations and for what goals their offerings are suitable. For curriculum designers this implies that they need to interrelate courses, offer several alternatives, indicate benefits and drawbacks, create a realistic time schedule for the potential learners. For learners this implies that they should be advised on the learning possibilities that match their *current* competence level and that work toward their *desired* competence level (learning goals), taking into account their restrictions and preferences. In order to accomplish this, we need editors, visualizations and interactive tools to work with a large amount of offerings. Further, selection and structuring mechanisms (we call them positioning and navigation) should be available to find the right stuff in the huge basket. As these tools rely on the underlying (meta)data, we need standardized descriptions of learning activities and their role within competence development programs (CDPs).

The aim of our research and development activities is to provide tools for supporting stakeholders in the field of lifelong learning in their activities related to the design, creation, selection, personalization and usage of competence development programs. One issue is that we do not exactly know what these tools should look like, as support for the creation of curricula is lacking in virtually all e-learning systems. Apparently, this is something that still is being done (or rather needs to be done) by hand. There are some good reasons why:

- various failed attempts in the field of adaptive educational hypermedia have shown that you can't predict with 100% certainty the learners' goals
- as one almost always needs to trade-off, a system just can't come up with the curriculum; best it can do, is to provide several options
- the process of creating a CDP, based on an initial idea of a goal, helps in making the goal more concrete. The further one is in the process, the more context one has and the better one can decide whether a learning activity is relevant/fun/interesting
- curricula are not just planned completely beforehand

At schools or Universities, students are offered curricula that typically consist of a fixed, obligatory part and some space that they can fill in themselves. In practice, students revise their study plan each year, based on their past year's experiences, available courses and personal factors, such as time constraints, focus in study goals and preferences. Often they also discuss their choices with their peers and with their mentors, who have to approve the plan. In lifelong learning, we see the same effect. On the one hand, many companies offer fixed, possibly obligatory, training programs. On the other hand, lifelong learners plan their learning activities based on current, possibly short-term needs ('I need to learn how to work with PhotoShop for this project.'), career planning and personal interests (the latter often also in the learner's private time).

Which learning opportunities learners take up, depends on a large number of factors. Apart from current training needs and wishes, factors as balance between private life and work, current offerings in the neighborhood or region, support from their chefs and personal motivation play an important role. Instead of (automatically) generating a supposedly 'perfect' competence development program, we need to support the learners in planning their learning activities themselves.

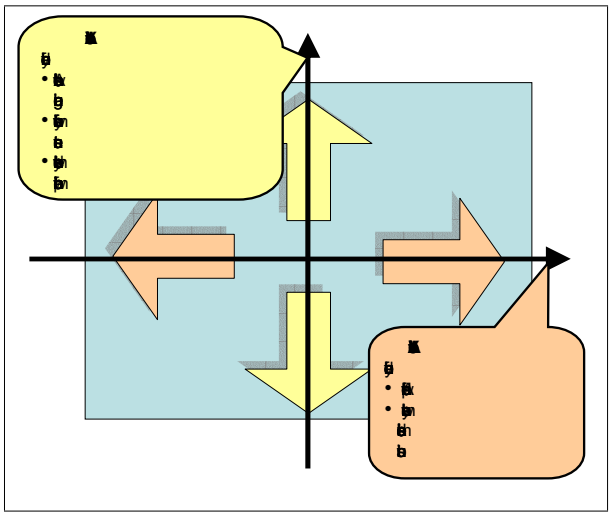

Figure 1: How information about the learner, the learning activity (here abbreviated as LA), and about other learners influence the learning activity's location on the screen.

4 Supporting Curriculum Planning with the Hybrid Personalizer

In this section we decribe our approach and its implementation in the system called *Hybrid Personalizer*. We provide details about how we combine simple recommendation services and therefore exploit their complementary features in order to build up a single hybrid personalization service for curriculum planning.

4.1 A personalization approach to Curriculum Planning

The Hybrid Personalizer serves as a recommender system for learning activities, such as courses. The recommendation is computed based on information available about the learner, the learning activities, and the behavior of other (successful) learners. The recommendation provided by our system is two-fold: on the one hand, it ranks learning activities based on how close they are to the learner's current knowledge level; that is, learning activities that are still too advanced are scheduled at a later point in the initial recommended visualization of the curriculum (or rather, competence development program). On the other hand, as a second dimension, our system ranks learning activities based on to what extent they match the learners' preferences—as explicitly indicated in their profiles and as estimated from the behavior of similar users. These two orthogonal aspects—at what point to plan a learning activity in the CDP and the extent to which learning activities are preferred to one another—are visualized in a two-dimensional diagram, in which each learning activity gets an initial position on the horizontal and vertical axes. In Figure 1 we depict the two axes, of which the values are determined as follows: learning activities that are more advanced or that typically appear later on in the curriculum are located higher (further away from the learner's initial point) in the proposed plan—which means that it is not recommended to involve in this activity right away; learning activities that match best with the learner's preferences or that are selected most by peer learners are placed in the middle, whereas less preferred activities are located in the periphery. As a result, the initial graphical overview can be

Figure 2: The TENCompetence Planning Tool empowered with the Hybrid Personalizer. Circles represent learning activities. Activities located in the lower middle are recommended. Hovering over the circles activates surrogates detailing the learning activity.

read from the bottom to the top, with the middle line representing the recommended activities in the recommended order.

By placing the learning activities in this way, the learner is provided with a personalized view on the learning network. This results in a better overview of the opportunities for the learner: the learning activities that best match the learner's current knowledge level are placed in the lower middle of the screen and can be considered the recommended next steps. Figure 2 shows a screenshot of the prototype Graphical Planning Tool, which provides the actual visualization of the two different rankings. Following Ben Shneiderman's 'overview first, details on demand' [20] the learning activities are initially only depicted with their titles and a color coding that indicates what learning goal the course serves. Upon hovering over the object with the mouse, more details are displayed in a popup window. The learner can actively engage with the learning activities, move them around, reschedule items, exchange one alternative for another, and eventually create her personal development plan. The complete interaction design is beyond the scope of this paper, which focuses on the hybrid personalization technologies that work in the background.

4.2 System Features and Architecture

The Hybrid Personalizer consists of two features to provide a personalized view for curriculum planning tools:

1. computation of recommendations for learning activities
2. computation of learning paths through the whole set of learning activities.

We briefly detail these two features in the following.

Compute recommendation values. For recommendation purposes, our system computes two dimensional recommendation values for each learning activity by taking into account information about the learner, other learner's behaviour, and the learning activity.

Compute recommended learning paths. A learning path is a sequence of learning activities typically guiding a learner from an initial state of knowledge to a learning goal. Computing the learning paths follows the ideas presented in [1]. Given a learner's current state of knowledge and her learning goal, the Hybrid Personalizer is able to create a learning path to be suggested.

Figure 3 depicts the architecture of our system. It is partitioned into three layers. The bottom layer comprises the so-called Atomic Personalization Services. Each of these services acts as a recommendation service on its own. The middle layer takes care of calling the Atomic Services, gathering information about the learner, and deciding which atomic service's output to combine and in what way. The upper layer comprises the graphical user interface actually showing the result of the hybrid personalization service by retrieving the personalized information from the middle layer.

4.3 The Atomic Personalization Services

As input for our hybrid personalization system, four Atomic Personalization Services are exploited (see the first layer in Figure 3). Each of them provides a complementary aspect of the final recommendation by returning a numerical value representing how much a certain learning activity fits the learner's current situation. For this computation, each Atomic Service exploits a certain facet of the information available describing the learner and the learning activities. In the following we briefly introduce each Atomic Service.

Positioning Service. The positioning service estimates the relevance of a learning activity by applying Latent Semantic Analysis (statistical comparison of textual contents) on the learning activities and the learner model [14]. Intuitively, the positioning service recommends learning activities that are more similar to the learner model. The learner model is supposed to contain a portfolio, a set of documents describing a learner's current state of knowledge. The positioning service does not require any learner metadata to exist, but has the disadvantage that it inherently introduces some uncertainty.

Navigation Service. The navigation service uses collaborative filtering techniques for determining the most popular followed steps after having completed a learning activity [11]. Intuitively, this service provides a higher ranking to learning activities that were successfully attended by other learners in the same situation. Similar to the positioning service, the navigation service does not require any metadata and makes use of the 'wisdom of the crowds'. As a disadvantage, this service requires a relatively large user base and might not properly take envisaged didactics into account.

Curriculum-based Service. The curriculum-based service imposes an order on learning activities (in this case typically courses), by comparing their prerequisites and learning outcomes: if a learning activity B requires competences that can be learned from learning activity A, B is placed after A; that is, B is ranked less high than A. We based our implementation for this service on the approach described in [1]. This top-down technique requires metadata in terms of prerequisites and learning outcomes on the learning activities (which is an authoring effort), but is able to generate suggested paths that follow didactic principles.

Preference-based Service. The preference-based selection of learning resources analyzes to what extent a learning

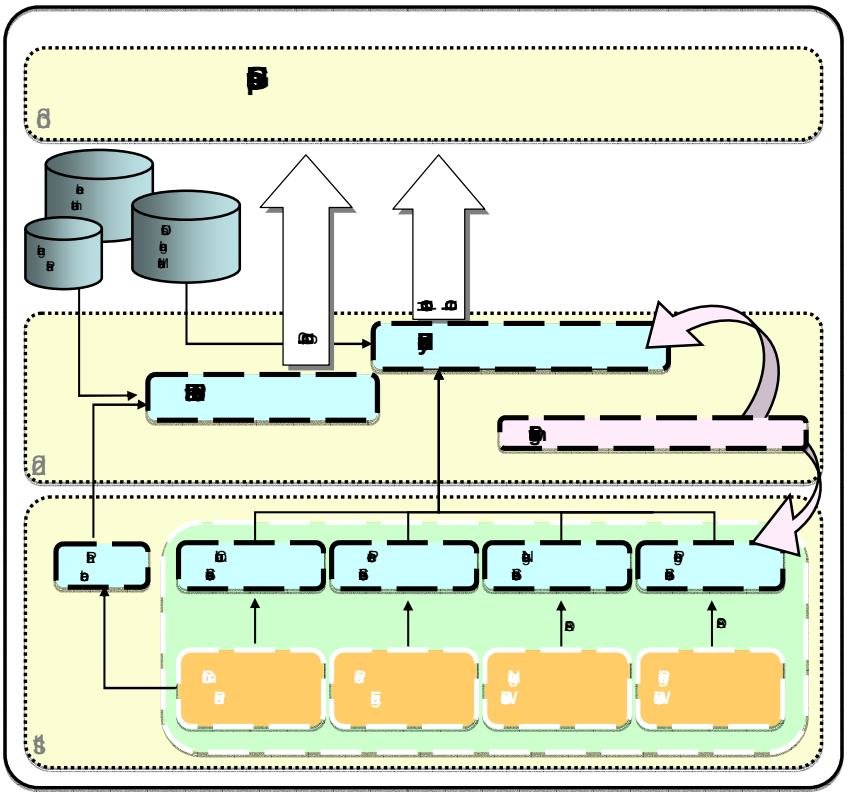

Figure 3: The three layer architecture of the Hybrid Personalizer.

activity matches the learner's preferences [15]. Preferences may concern anything ranging from type of assessment and study load to where the course is given and the costs associated with it. Therefore, this service requires preferences to be specified by the learner, such as 'I prefer oral exams to written ones.'.

These four services take various aspects of the learner's situation into account and combine these aspects with the information available about the learning activities. The four services are complementary in terms of the information they use: the first two follow a bottom-up approach while the last two compute recommendations in a top-down fashion. In other words, the former two extract implicit information from available data via Information Retrieval techniques (Latent Semantic Analysis for the Positioning Service and Collaborative Filtering for the Navigation Service). The latter two services exploit metadata information in order to compute recommendation. In the following section we show how we merge these two complementary approaches by combining the four Atomic Services.

4.4 Becoming Hybrid—Combining the Atomic Services

The atomic Personalization Services—as introduced in the previous section—form the constituent parts of the advanced integrated Personalization Service, which we call the Hybrid Personalizer. Each of the atomic services provides the middle layer with complementary information on which learning activity suits best a learner's needs and preferences in her current situation. Combining these complementary input values is the challenge that the Hybrid Personalizer is dealing with. As it has been stated previously, the four atomic services can be divided into two bottom up approaches, namely the navigation and positioning Web Services, and into two top-down approaches, namely Preference-based and Curriculum-based personalization.

Merging the output of the four atomic services in order to provide a single personalized view on the learning space is a challenging task. There are arbitrary many ways of combining the output of the atomic services. This may depend on the information available for computing the personalized locations; if there are no portfolios available for the current learner, the Positioning Service cannot be applied. Depending on the data available one may also want to put different weights onto the results of the atomic services, leave some of them out, or even let the user decide how to configure the Hybrid Personalizer. Since any decision on how to combine the services has to be taken carefully and may be tuned according to the mentioned conditions, we developed and integrated a configuration component that allows a fine-grained tuning and adoption of how the returned values of the atomic services are used to compute a single personalization value. By this means, the strategy of the hybrid personalization can be modified easily.

5 Conclusions and Future Work

In this paper we have introduced the concept of hybrid personalization, which is based on the idea of hybrid recommender systems and allows for a flexible combination of various adaptation techniques, ranging from knowledge-driven to content-driven approaches. We instantiated this approach for supporting curriculum planning in a learning context. We implemented the Hybrid Personalizer, which provides conceptual adaptation decisions, based on a number of atomic personalization services. These services can be combined in various different ways and the output of

the Hybrid Personalizer can be attached to several visualizations or hypermedia personalization techniques.

The concepts of hybrid personalization are inspired by the fields of hybrid recommender systems and layered (evaluation of) adaptive systems. In theories and frameworks on adaptive (educational) hypermedia systems, the separation of model, view and controller are well taken into account, but in the actual implementation they often end up being merged. With our prototype implementation as presented and discussed in this paper, we haven shown the practical benefits of keeping these layers separated.

Currently, the combination of the four personalization services and their weights is based on 'heuristics and intuition'. We plan to evaluate the graphical curriculum planner with students, who will use the planner for planning their study activities for the upcoming year(s). The observed usage data, as well as transcriptions of the user comments, will be used for an informed configuration of the Hybrid Personalizer. Further, the user interface of the configuration tool for the wrappers is currently not suitable for end users. As it is important that users can *scrutinize* [16] the configuration of the hybrid personalization service—to adapt the output to their needs and to be provided different perspectives—we aim to develop a user interaction paradigm that hides the complexity from the user, but that still provides a wide range of configuration possibilities.

We have applied hybrid personalization to the field of e-learning, making use of e-learning oriented personalization mechanisms. To conclude this paper, we would like to stress that the same principle can be applied to other fields, including personalized news sites, desktop search and e-commerce.

Acknowledgments

The authors efforts were (partly) funded by the European Commission in the TENCompetence project (IST-2004-02787). The Graphical Curriculum Planning Tool, which displays the results of the Hybrid Personalizer, has been developed by our project partners at the Universitat Pompeu Fabra.

References

[1] Matteo Baldoni, Cristina Baroglio, Ingo Brunkhorst, Elisa Marengo, and Viviana Patti. Reasoning-based curriculum sequencing and validation: Integration in a service-oriented architecture. In *Proc. ECTEL 2007*, 2007.

[2] P. Brusilovsky. Methods and techniques of adaptive hypermedia. *User Modeling and User Adapted Interaction*, 6 (2-3):87–129, 1996.

[3] P. Brusilovsky. Adaptive hypermedia. *User Modeling and User-Adapted Interaction*, 11:87–110, 2001.

[4] P. Brusilovsky and E. Millán. *The Adaptive Web - Methods and Strategies of Web Personalization*, chapter User Models for Adaptive Hypermedia and Adaptive Educational Systems, pages 3–53. Springer, 2007.

[5] R. Burke. *The Adaptive Web - Methods and Strategies of Web Personalization*, chapter Hybrid Web Recommender Systems, pages 377–408. Springer, 2007.

[6] M. A. Chatti and M. Jarke. The future of e-learning: a shift to knowledge networking and social software. *Int. J. Knowledge and Learning*, 3 (4/5):404–420, 2007.

[7] M. Coyle and B. Smyth. (web search)shared - social aspects of a collaborative, community-based search network. In *Proc. Adaptive Hypermedia 2008*, 2008.

[8] A. I. Cristea and A. De Mooij. Laos: Layered www ahs authoring model and their corresponding algebraic operators. In *Proc. WWW2003*, 2003.

[9] P. De Bra, A. Aerts, D. Smits, and N. Stash. Aha! version 2.0, more adaptation flexibility for authors. In *Proc. AACE ELearn'2002 Conference*, pages 240–246, 2002.

[10] P. Dodds and S.e. Thropp. Scorm 2004 - overview. Web site, July 2004.

[11] H. Drachsler, H. Hummel, B. Van den Berg, J. Eshuis, W. Waterink, R. Nadolski, A. Berlanga, N. Boers, and R. Koper. Effects of the isis recommender system for navigation support in self-organised learning networks. In *Proc. Workshop on Technology Support for Self-Organized Learners at Edumedia 2008*, 2008.

[12] E. Herder. Utility-based evaluation of adaptive systems. In *Proc. Workshop Evaluation of Adaptive Systems at UM2003*, 2003.

[13] E. Herder and H. Weinreich. Interactive web usage mining with the navigation visualizer. In *Proc. Chi 2005*, 2005.

[14] M. Kalz, J. Van Bruggen, B. Giesbers, W. Waterink, J. Eshuis, and R. Koper. A model for new linkages for prior learning assessment. *Campus Wide Information Systems*, 25 (3), 2008.

[15] P. Kärger, D. Olmedilla, F. Abel, E. Herder, and W. Siberski. What do you prefer? using preferences to enhance learning technology. *IEEE Trans. Learning Technologies*, (to appear), 2008.

[16] J. Kay, R.J. Kummerfeld, and P. Lauder. Personis: a server for user models. In *Proc. Adaptive Hypermedia 2002*, pages 203–212, 2002.

[17] J. Nielsen. Personalization is over-rated. Alertbox, October 1998.

[18] H. Obendorf and H. Weinreich. Comparing link marker visualiztion techniques - changes in reading behavior. In *Proc. WWW Conference*, pages 736–745, 2003.

[19] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. IEEE Visual Languages*, 1996.

[20] B. Shneiderman. Designing information-abundant web sites: Issues and recommendations. *Intl. J. Human-Computer Studies*, 47 (1):5–29, 1997.

[21] C. Tattersall, J. Janssen, B. van den Berg, and R. Koper. Using ims learning design to model curricula. *Interactive Learning Environments*, 15 (2), 2007.

[22] S. Weibelzahl, A. Paramythis, and A. Totter. A layered framework for the evaluation of interactive adaptive systems. In *Proc. 2nd Workshop on Empirical Evaluation of Adaptive Systems (at UM2003)*, page supplement, 2003.