The Knowledge Engineering Review, Vol. 16:2, 111–155. © 2001, Cambridge University Press
DOI: 10.1017/S0269888901000108 Printed in the United Kingdom

Personalised hypermedia presentation techniques for improving online customer relationships*

ALFRED KOBSA,[1] JÜRGEN KOENEMANN[2] and WOLFGANG POHL[3]

[1] *Dept. of Information and Computer Science, University of California, Irvine, CA 92697–3425, USA*
[2] *humanIT Human Information Technologies, 53757 St. Augustin, Germany*
[3] *German Informatics Society, Ahrstr. 45, 53175 Bonn, Germany*

Abstract

This article gives a comprehensive overview of techniques for personalised hypermedia presentation. It describes the data about the computer user, the computer usage and the physical environment that can be taken into account when adapting hypermedia pages to the needs of the current user. Methods for acquiring these data, for representing them as models in formal systems and for making generalisations and predictions about the user based thereon are discussed. Different types of hypermedia adaptation to the individual user's needs are distinguished and recommendations for further research and applications given. While the focus of the article is on hypermedia adaptation for improving customer relationship management utilising the World Wide Web, many of the techniques and distinctions also apply to other types of personalised hypermedia applications within and outside the World Wide Web, like adaptive educational systems.

1 Introduction

In today's competitive business environment, providing value to the customer is of paramount importance for businesses to survive. The single most important way to provide value is to know the customers and serve them as individuals. Customers need to feel they have a unique personal relationship with the business. The terms *personalisation*, *micro-marketing* and *one-to-one marketing* are used to describe this business model (Peppers & Rogers, 1993, 1997).

The World Wide Web (WWW or 'web' for short) can be fruitfully employed to support the entire sales cycle. In the pre-sales phase, it is used to establish and strengthen corporate and brand identities and to draw customers' attention to new products and services. In the sales phase, web-based online storefronts with electronic multimedia catalogues and online ordering/purchasing facilities enable a customer to select the desired product and to purchase it directly. In the post-sales phase the web is used to reassure customers of his or her purchase decision and to deliver additional values through services like product support, support for user groups and loyalty programs, thus creating opportunities

* This paper is based on research that has been supported by the European Commission in its ACTS programme (projects AVANTI and COBRA), by DaimlerChrysler AG, and by T-Nova Berkom. Contributions by Josef Fink, Igor Jaceniak and Andreas Nill are greatly appreciated. Comments, on a previous version, by Peter Brusilovsky, Eric Horvitz, Tony Jameson, Judy Kay and the anonymous reviewers of the *Knowledge Engineering Review* considerably helped improve the style and the content of this article.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

for long-term customer retention and eventually cross- and up-selling (Allen *et al.*, 1998; Hanson, 1999; Konstan, 1997; Pitt *et al.*, 1996).

Until recently, electronic commerce on the WWW was mass-oriented. Business offers and services were uniform and geared to suit the largest possible buyership. As is the case in traditional commerce, one can, however, expect that in the long run those companies will survive more easily that can establish, maintain and extend a customer base by delivering tailored products and services and thereby create stable long-term relationships with repeat customers.

Customer relationship software has supported sales and marketing divisions in their task of providing such individualised customer care and of integrating and utilising information from various sources (sales, marketing, and recently e-commerce) to create targeted information, services and product offers. Web-based customer relationship software empowers websites with this functionality directly, offering a number of key advantages at reasonable cost, compared to traditional channels:

- it facilitates the collection of information about large numbers of customers (their interests, online purchase behaviour, and support needs);
- it offers the opportunity for rapid changes and updates of content and presentation to quickly react to new opportunities and challenges;
- it enables a global around-the-clock presence independent of its locality; and
- it offers the opportunity for dynamic creation of content and presentation formats for narrowly targeted and/or personalised information delivery.

Current personalisation techniques on the web are relatively simple. They have been developed on the basis of concrete business demands and the experiences of the sales and marketing departments. Early signs show that even this simple personalisation based on purchase data and personal data provided by the user has a considerable payoff. Jupiter Communications reports that personalisation at 25 consumer e-commerce sites boosted the number of new customers by 47%, and revenues by 52%, in the first year (Hof *et al.*, 1998). Nielsen NetRatings reports that registered visitors to portal sites (who obtain the privilege of adapting the displayed information to their interests) spend over 3 times longer at their home portal than other users and view 3 to 4 times more pages at their portal (Thompson, 1999). In Nielsen (1999) they also report that e-commerce sites offering personalised services convert approximately twice as many visitors into buyers than e-commerce sites that do not offer personalised services.[1]

This paper provides a detailed overview of techniques that can be used to realise significantly more far-reaching personalised applications for improving customer relationships than exist today, specifically on the World Wide Web. Many of these techniques are knowledge-based, allowing one to incorporate, e.g. experiences with regard to customers and their characteristics, customer types, products and their characteristics, and sales strategies. We will use the notion of *customer relationship application* as denoting any (web-based) application that is aimed at establishing a bond with web visitors through the provision of useful personalised services, with the ultimate goal of turning visitors into loyal customers.

Interactive web-based systems belong to the group of hypermedia systems. We define a hypermedia system as an interactive system that allows users to navigate a network of linked *hypermedia objects*. In the case of the WWW, these hypermedia objects are web pages. Hypermedia objects contain a set of related content-bearing elements of different media types such as text, images, video-clips, audio-clips, small applications (applets) and interaction elements (e.g., menus, buttons or checkboxes).[2]

[1] The economic expectations for personalisation are very high: Appian estimates that the revenues made by the online personalisation industry, including custom development and independent consulting, will reach $1.3 billion in 2000, and $5.3 billion by 2003 (Appian, 2000). Ovum forecasts that the world-wide revenues for personalisation software will rise from $10.85 million in 2000 to $93.4 million in 2005 (Millhouse *et al.*, 2000). Gartner predicts that "by 2003, nearly 85 percent of global 1.000 Web sites will use some form of personalisation (0.7 probability)" (Abrams *et al.*, 1999).

[2] The notion of *hypertext* denotes hypermedia that has text as the sole media type.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

These elements are presented in a particular format and in particular spatial relation to each other (layout). Elements are linked to hypermedia objects, or specific elements within these objects. Links can exist between elements within the same hypermedia object (page links), between hypermedia objects that belong to the same website (internal links), and between hypermedia objects of the local website and site-external hypermedia objects (e.g. those that belong to the website of another corporation). In the latter case, the local site only has control over the hypermedia object containing the link anchor, not over the target object.

We define a *personalised hypermedia application* as a hypermedia system which adapts the content, structure and/or presentation of the networked hypermedia objects to each individual user's characteristics, usage behaviour and/or usage environment. There are different basic types of adaptation depending on the amount of control a user has over the adaptation. We thereby differentiate the roles of the adaptation *initiator*, adaptation *proposer*, adaptation *selector* and adaptation *producer*.[3]

Systems where the user is in control of initiation, proposal, selection and production of adaptation ("in control" thereby meaning that the user can perform these functions, but can also opt to let the system perform some of them) are called *adaptable* (Oppermann, 1994). For instance, a user of a website may see a need to introduce a shortcut for a web page that is frequently visited by him or her but deeply buried in the site hierarchy (adaptation initiation). The user then "proposes" to himself or herself to introduce a new link on the lateral navigation bar of the site pages (assume this is possible) or to define a bookmark in the browser, selects the shortcut link and performs the necessary steps to produce this adaptation. In contrast, systems that perform all steps autonomously are called *adaptive*.[4] For instance, the AVANTI system (Fink *et al.*, 1998) automatically inserts such personalised shortcut links for pages that a user frequently visits. Intermediate forms are also possible, as when

(a) the system signals the need for shortcuts, proposes pages for which they should be introduced, lets the user make the selection and performs the selected adaptations ("user-controlled adaptivity"); or
(b) the user requests the introduction of shortcut links and lets the system decide which are best ("user-initiated adaptivity").

Adaptability and adaptivity can coexist in the same application. Which form to choose has to be carefully weighed for each class of adaptation, taking convenience for the user, demands on the user, irritation of the user and the consequences of false adaptation into account. User control may be provided on a general level (users can allow or disallow adaptation at large), on a type level (users can approve or disapprove that certain types of adaptation take place) or on a case-by-case basis. For example, a system that monitors transmission rates may: (1) suggest forgoing the loading of a single video and offer a still image instead, (2) suggest forgoing all videos, or (3) forgo all videos without asking and thereby leave the user unaware that videos are available.

The *volatility* of the adaptation addresses the timing and the dynamics of adaptation production. One extreme is the one-time specification of a user or usage model (see below) which occurs typically at the time when an application is installed or a session initialised (some authors refer to such systems as *configurable* systems). On the other extreme lies continuous adaptation throughout use, based on changes in the user's behaviour or dynamic environmental factors. Most hypermedia systems that will be described in this paper need to acquire a certain amount of data before they can start adapting to the user. They are therefore only useful in application domains where users engage in extended (and

[3] The distinction between adaptation initiation, proposal, selection ("decision") and production ("execution") is due to Dieterich *et al.* (1993).
[4] Stephanidis *et al.* (1998) use the terms "adaptive" and "adaptable" in a different sense. In their terminology, "adaptive" denotes adaptations that occur at runtime (and that may be produced both by the system and by the user), and "adaptable" denotes adaptations before runtime (e.g. when the system is first installed). Other authors use the term "configurable" or "customisable" to refer to this latter kind of adaptation.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

in most cases even repeated) sessions of system use. They may not be appropriate for infrequent users with typically short sessions.

Finally, the *complexity* of adaptation is a relative measure of the complexity and directness of the process that produces the adaptation based on a number of input data. If there is a direct relationship between the input and the adaptation, then the complexity is low. This is, for example, the case when a user selects the "text-only" mode in his or her web browser, and when subsequently no images are displayed by the system. An example of a complex adaptation is the inference of likely user interests based on the (long-term) observation of interaction behaviour in an information resource. In general, adaptations can be assumed to be the more appropriate (i.e. in accordance with user needs) the lower the complexity is for producing them. However, complex adaptations can considerably gain in appropriateness with respect to users' needs if several (independent) pieces of information about the user are taken into account in producing them.

The personalisation process in hypermedia applications can be divided into three major tasks that are usually performed by different system components:

Acquisition is the task

- of identifying information that is available about users' characteristics and computer usage behaviour as well as about the usage environment, either by monitoring computer usage or by obtaining this information from external sources;
- of making this information accessible to the adaptation component of the application; and
- of constructing initial models of the user, computer usage, and/or the usage environment (the so-called *user model*, *usage model* and *environment model*).

Representation and *secondary inference* are the tasks

- of expressing the content of the user and usage models appropriately in a formal system, of allowing access and further processing; and
- of drawing further ("secondary") assumptions about users and/or user groups, their behaviour, and their environment, thereby integrating information from various sources.

Production is the task

- of generating the adaptation of content, presentation and modality, and structure, based on a given user, usage and environment model.

The paper is roughly structured along these tasks. Section 2 describes the data about the user, computer usage and the physical environment that adaptive hypermedia systems take into account when adapting the presentation of hypermedia pages to individual users' needs. Section 3 describes how these data can be acquired by a personalised hypermedia application. In Section 4 we discuss methods for representing these data in formal systems, and for making generalisations and predictions about the user based upon them. Section 5 presents methods for adapting the content, the presentation, the modality and the structure of hypermedia pages to users' needs. Different personalisation functions of hypermedia adaptation are thereby distinguished. Section 6, finally, discusses current and likely future application scenarios for personalised hypermedia in the broad context of customer relationship management, and gives recommendations concerning data acquisition and representation.

Customer relationship improvement in electronic commerce through personalisation is a very topical application area for hypermedia personalisation that can be expected to be fully adopted in the near future. Personalisation of hypermedia presentation is beneficial for several other purposes as well, most notably for improving the learning progress in educational software (see, for example most papers in Brusilovsky *et al.*, 1998 and the evaluations in Eklund & Brusilovsky, 1998; Specht, 1998; and Specht & Kobsa, 1999). Many of the methods described in this paper can also be fruitfully applied to these other purposes, but there clearly exist adaptation methods that are unique to each application purpose (e.g. personalised pedagogical strategies for hypermedia tutoring systems).

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

2 Input data for personalised systems

Personalisation is often a data-intensive task. Some data can typically be observed by the system directly, while most others may require one or more additional acquisition steps. This section deals with the different kinds of data that user-adaptive systems may need to consider, namely data about the user, data about computer usage and data about the user's hardware, software and physical environment.

2.1 User data

We introduce the term *user data* to denote information about personal characteristics of the user, while *usage data* (see Section 2.2) is related to a user's (interactive) behaviour. However, we will see that there are potential overlaps between these two categories. While some user data can be directly supplied by the user (see Section 3.1.1), most data must normally be inferred from usage observations. In many cases, the complexity of acquisition will be high and thus the reliability of the result less than certain. We will therefore often use the term "assumption" in combination with user data (Pohl, 1998; Wahlster & Kobsa, 1986). In the following subsections we describe several categories of user data that have been the basis for adaptation in a number of systems developed so far.

2.1.1 Demographic data

Demographic data about the user are "objective facts" like the following:

- record data (e.g., name, address, phone number),
- geographic data (area code, city, state, country),
- user characteristics (e.g., age, sex, education, disposable income),
- psychographic data (data indicating lifestyle),
- customer qualifying data (e.g., frequency of product/service usage),
- registration for information offerings, participation in raffles and so on.

Today's personalised websites in many cases operate on the basis of such demographic data and purchase data only. The value of these data can be high when combined with high-quality statistical data, such as about the purchase behaviour of different customer groups.

2.1.2 User knowledge

Assumptions about users' knowledge (or more general beliefs) about concepts, relationships between concepts and facts and rules with regard to the domain of the application system have always been among the most important sources for personalisation. Early user modelling research, which was strongly focused on natural-language systems, considered user domain knowledge in both understanding user utterances and generating appropriate system responses (Kobsa & Wahlster, 1989; Kok, 1991; McTear 1993). Early research also distinguished different types of nested and mutual beliefs (Ballim & Wilks, 1991; Kobsa, 1989; Taylor *et al.*, 1996) which until now have been of only minor importance outside natural-language systems.

More recently, many approaches to building adaptive hypermedia systems have taken individual user knowledge into account. The rationale is to adjust the presentation of hypermedia material in such a way that users become neither bored by unnecessary explanations nor confused by details they cannot understand. Examples include

- restricting or increasing the number of explanatory pages to be presented to the user depending on his or her expertise (Sales Assistant, Popp & Lödel, 1996);
- presenting expertise-dependent explanations and technical details (Metadoc, Boyle & Encarnacion, 1994; KN-AHS, Kobsa *et al.* 1994); and
- generating expertise-dependent product descriptions (SETA, Ardissono & Goy 1999, 2000b; Ardissono *et al.*, 1999).

Adaptation to a user's knowledge of domain concepts, of rules and of other items is also a typical feature of intelligent tutoring systems. In many such systems, user knowledge is taken into account

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

when guiding the user through the learning material. Examples are the ISIS-Tutor (Brusilovsky & Pesin, 1994), Hypadapter (Hohl *et al.*, 1996) and ELM-ART II (Weber & Specht, 1997) systems. Good surveys of such systems are given in McCalla and Greer (1994) and Brusilovsky (1996, 2001).

2.1.3 User skills and capabilities

The systems mentioned in the previous section mainly deal with users' knowledge about concepts, and relationships between concepts, in an application domain. Besides this "knowing what", a user's "knowing how" can also play an important role in adapting systems to user needs. Adaptive help systems are typical representatives of this approach. For instance, the Unix Consultant (Chin, 1989) tailors its help messages and explanations to the user's familiarity with UNIX commands. Peter and Rösner (1994) tailor repair instructions to the user's familiarity with the operations involved in the suggested repair plan.

Küpper and Kobsa (1999) go further and distinguish between the actions a user is familiar with and the actions he or she is actually able to perform. It is possible that a user knows how to do something but is not able to perform the action due to lack of required permissions or to some physical handicap. The tourist information system AVANTI (Fink *et al.*, 1998), which takes the needs of different kinds of disabled people (wheelchair-bound, motor-impaired and vision-impaired) into account, therefore only recommends actions that these users are actually able to perform.

2.1.4 User interests and preferences

Considering user interests and preferences is vital for many personalised web-based systems. Interests among users of the same application often vary considerably, and information and product offerings that are targeted to one particular interest group may not only be of no interest to another group but may even also have adverse effects. An example is the promotion of cars to different audiences for which different and maybe even conflicting sets of attributes (speed, sex-appeal, safety, family-friendliness) must be emphasised (see Jameson *et al.*, 1995, for an example in the used-car domain and Ardissono & Goy, 1999, for the domain of telephony devices).

User interest is also a central notion for so-called recommender systems (Resnick & Varian, 1997). Such systems attempt to recommend items to users. These items may be products, services, documents, news and so on. In this case, the user's interest is generally expressed in terms of features (which in the case of documents may be characteristic words of these documents). Recommendations can also be made by asking users to rate items with which they are already familiar (and thereby express their interests indirectly). The system will then try to find other users with similar ratings and recommend items that these other users have also rated highly ("clique-based" or "collaborative" filtering). Both types of filtering will be explained in more detail in Sections 4.2 and 4.3.1.

2.1.5 User goals and plans

When interacting with web-based systems, users will often browse and visit pages without a specific goal in mind. However, it is also possible that users pursue concrete goals and plans when visiting a site. Typical goals may be to find information on a certain topic or to shop for some kind of product. A system that supports users in achieving their goals would be very helpful. For instance, Lesh *et al.* (1999) demonstrate that plan-recognition facilitates and speeds up interaction considerably since the system has expectations about the next user actions and can therefore interpret them in a more flexible way. PUSH (Höök *et al.*, 1996) limits the information presented to the user to those types that are relevant to the user's current goals (which are specified by the user or deduced by the system). HYPERFLEX (Kaplan *et al.*, 1993) also allows users to specify their goals and presents those paths that earlier users with the same goals had most frequently visited.

Methods for recognising goals and plans will be briefly discussed in Section 3.1.3.

2.2 Usage data

Usage data may be directly observed and recorded, or acquired by analysing observable data. The extent to which observation is technically possible varies considerably. Hypermedia systems that are

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

exclusively based on HTML will only be able to record what pages and files have been requested from the server (which includes a trace of the navigation paths for links on the same site). Systems that have more control about the interaction by, for example, using Java applets, can record usage data on the level of mouse clicks and movements. For instance, ELFI (Pohl & Nick, 1999) logs all user interaction with a hierarchical navigator, and TELLIM (Joerding *et al.*, 1998) monitors users' interaction at the widget level.

In addition to interaction behaviour, the usage context may also be considered as a source for adaptation. Among the relevant items are the current task and the interaction history. Typically, merely application-specific context is currently taken into account, although it might be desirable to also consider the user's general situation, such as other applications that the user currently or frequently employs (like the newsgroups that he or she is reading in a separate newsreader – cf. Sakagami & Kamba, 1997; and Sivadas *et al.*, 1998[5]).

In this section we first describe usage data that can be directly observed (Section 2.2.1). Such data may directly lead to adaptation. They are often, however, only or additionally the basis for inferences aimed at discovering more general regularities in observable usage data (Section 2.2.2). These in turn can be used by the system to adapt to user preferences, habits and levels of expertise that become manifest in these regularities.

2.2.1 Observable usage

There are many ways in which users may interact with a system. We identify different types of interactions that are relevant in the discussion of adaptive systems.

2.2.1.1 Selective actions. The most frequent kind of interaction with web-based systems is clicking on a link. When doing so, the user makes a choice if competitive links are available on the current page. For example, on an e-commerce site the user may select one of the products offered on an overview page to take a look at a more detailed description. Such selective actions can be regarded as indicators for several types of user data:

Interest. There are many systems that use characteristic selective actions as indicators for user interest. Among them is WebWatcher (Joachims *et al.*, 1997), which monitors link-selection on web pages to annotate the most relevant links on each page (with respect to a given user goal). While WebWatcher makes its annotations based on the selections of all visitors of the page, Personal WebWatcher (Mladenic, 1996) is a variant that only takes individual preferences into account. Letizia (Lieberman, 1995) learns about users' interests from page selections and presents related pages in smaller page frames. Another example is the system HIPS (Oppermann & Specht, 1999, 2000), a mobile web-based museum guide. HIPS lets users select from a number of paintings that have been detected to be close to the user's current location. Based on such selective actions, users' interests in certain painters, themes and periods are then predicted.

Selective actions need not only be mouse clicks but may also be scrolling and enlarging operations for hypermedia objects (Sakagami & Kamba, 1997), document expansion operations (Sakagami *et al.*, 1998), movie and audio control operations (Joerding *et al.*, 1998) and other actions at the user interface (Encarnação & Stoev, 1999).

However, selective actions may only be weak indicators for user interests since the user may not have seen or understood all available choices, or may have simply been inattentive. In addition, selective actions yield positive evidence of interest only: not following a link when competitive links are present is in general not a reliable indicator of users' non-interest in these other items.

Unfamiliarity. Selective actions can also be used to recognise users' unfamiliarity with technical terms. For instance, MetaDoc (Boyle & Encarnacion, 1994), KN-AHS (Kobsa *et al.*, 1994) and Sales Assistant (Popp & Lödel, 1996) allow users to obtain explanations for technical concepts mentioned

[5] Sivadas *et al.* (1998) showed that in the case of music, membership in newsgroups that discuss specific music styles is a better predictor of consumption and purchases related to this music style than demographic data.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

on a page by simply clicking on them. When users perform these actions, the assumption is made that they had not been familiar with the respective technical term. In this case, again, selective actions lead to positive evidence only: not following an explanatory link is in general not a reliable indicator of users' familiarity with the item explained. Moreover, it should be noted that explanatory links may also be selected for reasons other than unfamiliarity, like curiosity, assuring oneself that one has full understanding and so on.

Preferences. A special way of using selective actions is to present prototypical objects of potential interest to the user and then let the user select from these objects. The selections made by the user are taken to infer user preferences with respect to the type of presented objects. For instance, the Adaptive Graphics Analyser (Holynski, 1988) attempts to acquire user preferences concerning computer-generated graphic images. The system presents a set of default images to the user so that the variable settings for the graphics generation process are covered by these images. Users are then asked to tell the system for each image whether they like it or not. (Actually, users may choose from five different degrees of preference, resulting in a selective process that is graduated.) This approach is related to user ratings that will be described in Section 2.2.1.3.

2.2.1.2 Temporal viewing behaviour. Viewing time has been discussed and experimented with as a potential indicator for user characteristics, particularly user interest (Joerding,1999; Joerding *et al.*, 1998; Konstan *et al.*, 1997; Morita & Shinoda, 1994). Measurement of effective viewing time is difficult, though. It is often impossible to tell whether the user has been present in front of the computer screen and looked at a specific item within a specific time interval. In multi-window environments, a user may not even perceive a particular item because the window is covered by other windows, or an item may be outside the visible window area. In most cases viewing time will therefore be a weak indicator.

However, viewing time can serve as negative evidence. If the presentation time (and thus the maximal viewing time) of a hypermedia page is below a certain threshold, then the information on that page is most likely not interesting to the user. If the download of a hypermedia page is aborted or if the user presses the back button shortly after the page download commenced, this may be regarded as an even stronger indicator that the user realised that he or she is actually not interested in the item just selected, provided that the download time was within an acceptable limit. (The usage of the back button is, however, normally not recognisable by the web application.)

Determining the viewing time of hypermedia objects is easier if they are streamed (such as video or audio) and if a user reaction (such as a new page request) was noted shortly after the termination of the streaming. The fact that the user waited until the end of the presentation can also be regarded as a relatively strong indicator for his or her interest in this streaming object (Joerding, 1999). In 3D webstores (Chittaro & Ranon, 2000), the spatial and temporal navigation behaviour of users' avatars can be taken into account. Logging of the micro-interaction level (such as the tracking of mouse and slider movements, as in Encarnação & Stoev, 1999; Joerding 1999; Joerding *et al.*, 1998; Sakagami *et al.*, 1998; Sakagami & Kamba, 1997) and usage of eye-tracking for obtaining positive evidence of user interest still await further research.

2.2.1.3 Ratings. In several systems, users are required to explicitly rate objects (like documents, news articles, products). These ratings indicate how relevant or interesting this object is to the user, or how relevant or interesting the user thinks a document is to other users. Typically, users can apply either a binary rating scale (i.e. classify an object as, for example, "interesting" or "not interesting") or a limited, discrete scale. The latter kind of rating scale is typically numeric (for example, all integers from 0 to 5) or symbolic (with a mapping to a numeric scale). In recent years, ratings have become very popular. Examples include Firefly (Shardanand & Maes, 1995), Syskill and Webert (Pazzani & Billsus, 1997) and GroupLens (Konstan *et al.*, 1997).

In Syskill and Webert (Pazzani *et al.*, 1996), users have the possibility of rating a web page as "hot", "lukewarm", or "cold" (the "lukewarm" rating was abandoned in later versions of the system; Pazzani & Billsus, 1997). The web bookstore Amazon.com offered users the opportunity to rate books in

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

various categories on a 5-point scale and uses this input for a clique-based filtering algorithm (see Section) once a minimum of four ratings has been obtained.

Two problems with ratings are that relevance of information is always relative to the (changing) information need of a user, and that in information environments relevance judgements of individual items are typically assumed to be independent when in fact they are not (for example, the third article on the same topic may simply be rated lower because the first two items satisfied the information need and the user is judging incremental relevance at this point). For an extended discussion on the issue of relevance see, for example, Saracevic (1975) and Harter (1992).

Another problem is that computer users are known not to supply very many ratings anyway, particularly not negative ones.[6] Users are generally very reluctant to perform actions that are not directed towards their immediate goals if they do not receive immediate benefits, even when they would profit in the long run (Carroll & Rosson, 1987). Acquisition methods should therefore be preferred that do not exclusively rely on such explicit user input, and particularly not on negative ratings (Schwab *et al.*, 2000).

2.2.1.4 Purchases and purchase-related actions. Purchases made at an e-commerce site are generally regarded as strong indicators for user interest. Several commercial websites (the most well known being Amazon.com) react adaptively to a purchase by suggesting similar or related goods. If the customer relationship application uses an underlying feature-based model (see Section 4.3.1), the assumption is made that a purchase is a strong indicator of interest in some of the features of the purchased product. Of course, there is no one-to-one mapping of purchases and interests since, for example, customers purchase items for other people (for example, as gifts) and because people may already own an available item. Amazon.com attempts to address this issue by disregarding purchases with shipping addresses that are different from the user's address, and by encouraging customers to indicate that they already own a particular item.

Actions that are related to purchases but do not indicate user interest with the same strength include transferring products into the virtual shopping cart (without eventually buying them), answering quizzes, registering products, volunteering address data for information about product updates and company news.

2.2.1.5 Other confirmatory and disconfirmatory actions. Purchases and ratings can be examples of actions that confirm (or disconfirm) previous selections. Such (dis-)confirmatory actions are often stronger indicators than the preceding selective actions themselves. Even when a confirmation does not have strong indicative power by itself, it may still be useful to strengthen an assumption in concert with a preceding selection.

There are many other examples of confirmatory actions. For documents like web pages, news articles or email messages, it is interesting to monitor whether the user does any further processing like saving a document, printing a document, bookmarking a web page and so on (Konstan *et al.* 1997). Again, the type of web interface will determine whether or not this type of action is observable by the system. For example, sites like ZDNET.com give users the option of forwarding news stories by email, formatting them for printing and saving them in a server-side personal repository. Since these actions are performed under the control of the web application, they can be registered and evaluated.

2.2.2 Usage regularities

In many cases, observed user interactions will not directly lead to adaptations. Further processing of usage data is necessary to acquire information about users' preferences, habits, and levels of expertise upon which adaptive system behaviour can be based. In this subsection, we describe typical examples of usage information that is acquired from observed usage data.

Usage frequency. Perhaps the most obvious way to process usage data is to categorize events and count their frequencies. Debevc *et al.* (1996) describe a mechanism that controls the configuration of

[6] Mike Pazzani (personal communication) reports that only 15% of the users of Syskill and Webert would supply interest ratings even though they were encouraged to do so.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

an adaptive icon toolbar for Microsoft Word. Icons are added to or removed from the toolbar, based mainly on the usage frequency of the underlying commands. Linton *et al.* (1999) and Linton and Schaefer (2000) compare individual long-term usage of Word commands with group command usage data. Based on this comparison, commands that are not used by an individual but frequently employed by the group could be pointed out to this individual. Flexcel (Krogsæter *et al.*, 1994; Thomas & Krogsæter, 1993) tracks and analyses command usage on a deeper level of detail (such as whether command dialogues are often used with identical values). If this is the case, Flexcel suggests to the user the definition of new shortcuts or menu entries for these complex commands. AVANTI (Fink *et al.*, 1998) monitors how often individual users visit certain web pages and introduces shortcut links in the lateral navigation bar that accompanies all pages.

Situation-action correlations. In the early 1990s the notion of interface agents or personal assistants arose, mainly through the work of Mitchell *et al.* (1994) and Maes (1994). A common property of at least the initial work in this area was that situation-action correlations form the basis for adaptive behaviour. Both Mitchell and Maes describe email assistants that generate suggestions on how the user should deal with incoming email messages. These suggestions are based on statistics that express generalised correlations between previous emails (situations) and how the user processed them (actions). Maes also describes a meeting-scheduling agent that works in a similar way, but processes meeting requests instead of general emails (Kozierok & Maes, 1993).

Action sequences. Action sequences are mostly analysed to: (a) recommend the generation of macros for frequently used action sequences, (b) predict future user actions on the basis of past actions, and (c) recommend actions based on frequent action sequences of other users. Various methods are being used to determine regularities and similarities in action sequences (see Section 4.2).

2.3 Environment data

On the client side of web-based systems, the range of different hardware and software used is extremely wide (GVU, 1998). The spectrum is currently even becoming broader since the number of web-capable appliances with limited abilities (such as PDAs and mobile telephones) is rapidly increasing. Thus the chances of building one fixed web-based application that suits the large majority of possible usage environments become increasingly smaller. Web usage may be influenced by both the software and the hardware of the individual user, and by the characteristics of the user's current locale. We will consider each of these constraints in the next sections.

2.3.1 Software environment

Not every piece of web software supports all existing features and, particularly, older browser versions have very limited capabilities. New types of mobile, web-capable devices will add to this diversity. The following information is of particular interest when catering to the different software environments of users:

Browser version and platform. Web browsers differ in the range of features they support. The variety of versions in use is substantial, especially if one considers a global audience. Users increasingly become slow in upgrading their browsers.[7] The determination of a minimum set of capabilities and an assessment of what alternative versions are needed for presentation therefore becomes more and more important.

Availability of plug-ins. Sites that want to deliver rich media content need to consider whether the user has the appropriate plug-ins available. If this is not the case and if the user also is not willing to install them, the system has to take care that the website is still useful for such a user as well.

Java and JavaScript. A Java Virtual Machine (VM) for running Java applets is built into most modern Web browsers. Applets are often employed to allow for a higher degree of interaction with a website.

[7] For example, Nielsen (1999) predicts that not until the year 2003 will advanced browsers be sufficiently widely used to allow sites to go beyond the basics.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

However, not all Java versions have identical capabilities. Moreover, users may simply have Java disabled, e.g. because of fear of security risks or since corporate firewalls will block any applet. Similarly, JavaScript may also be disabled. Website developers need to take care that their sites will not become dysfunctional without Java and JavaScript.

In Section 3.3 we will discuss how these data can by acquired by an adaptive hypermedia application. It should be noted that data about the software environment can also give rise to assumptions about user data (Section 2.1). For example, the presence of the newest beta version of a browser on a home PC may indicate that the user is an early adopter and is experienced in using computers.

2.3.2 Hardware environment

There is often a huge gap between the hardware which web developers and designers use, and the hardware that is available to visitors of a website. For intranet applications, precise information about hardware constraints may be available; in the Internet case, only guesses and estimates are possible. The major differences include the following:

Bandwidth. Users may be connected to the Internet via analog modems, ISDN, DSL, Cable Modem, T1 connections and so on, which differ in bandwidth by a factor of up to 1000. Available bandwidth influences download time, and short download times are a significant factor for user satisfaction with a website and for the interpretation of abnormalities in temporal viewing behaviour (see Section 2.2.1.2).

Processing speed. Many users, particularly home users, keep their computers for many years. This implies that even if a user has a high-bandwidth Internet connection, there may be processor or memory constraints that preclude the viewing of web pages with numerous large images.

Display devices. With the advent of small mobile computing devices like handheld computers and smart mobile phones, the range of possible web displays has increased. The kind of content that can be delivered to certain client software is highly constrained (standard resolution, size, colour) and may not even conform to the "lowermost common denominator" of 640×480 VGA. But even between "standard" computers, there may be many differences in screen size and resolution. Optimised display for one specific resolution may be less than sub-optimal for different resolutions.

Input devices. With modern small devices, it may become difficult to provide input to interactive websites, for example if the target-area for selecting a navigation element is very small. Non-traditional input devices such as telephones may only provide a limited set of keys for interaction (restricting, for example, textual input, like in WAP). In the future, exclusively speech-driven interaction will also have to be considered (e.g. in hands-free situations like car-driving).

2.3.3 Locale

Information about the usage locale can be used to filter content, to adapt e.g. presentation formats and media, and to make recommendations based on geographical knowledge and past observations concerning user regularities (Fink, 2001; Jameson, 2001; Langley, 1999; Oppermann & Specht, 1999, 2000). Information about the usage locale includes:

Users' current location. The required granularity of location information varies widely, ranging from country level (e.g. as determined by the top-level domain of the web client) to meters or even centimeters in the case of location-sensitive applications. Additional information such as direction of gaze and direction of movement may be relevant as well, e.g. in order to identify the object a person is looking at or to provide directions.

Characteristics of usage locale. Information that can be taken into account includes the noise level and brightness of the surroundings, and information about places and objects in the immediate environment.

3 Acquisition methods

The previous section described the types of user, usage and environment data that are potential input to a personalised hypermedia system. We now describe methods that can be used to obtain these data.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

As we mentioned, the majority of these data cannot be used directly for adaptation but must be processed to obtain the initial contents of a user model or usage model.

In the first subsection, we will describe methods that are typically employed for acquiring explicit assumptions about user data. Collections of such explicit information are generally called user models. In Section 3.2, we will discuss methods that construct aggregated information about a user's interactive behaviour from individual observations. The results of such analyses, which often serve to predict future user behaviour, are sometimes called usage models. The last subsection (3.3) presents methods for acquiring information about the user's environment.

3.1 User model acquisition methods

This section describes methods to acquire user data as characterised in Section, and to store these data in an initial user model. We first describe methods that are based on explicit user input (and in most cases on the active acquisition behaviour of the system). In Sections 3.1.2 to 3.1.4 we then discuss methods for passive acquisition, namely domain-dependent acquisition rules, plan and goal recognition, and stereotypes for user classification.

3.1.1 User-supplied information

An obvious strategy for acquiring information about a user is to let the user supply the necessary data. For some user data (particularly demographic data, see Section 2.1.1) provision by the user is currently the only possible source of information. User data provision may take place through questions posed by the system, typically in an initial phase of system usage. Examples of the use of initial interviews are described in Sleeman (1985), Rich (1979), Boyle and Encarnacion (1994), and Fink *et al.* (1998). Most current websites that provide personalisation also employ interviews extensively. Initial interviews often lead to an assignment of the user to one of a predefined set of user subgroups (see Section 3.1.4 on "stereotype reasoning").

Figure 1 shows the initial interview of the AVANTI tourist information system (Fink *et al.*, 1998). Three general questions concern the user's familiarity with computers, with AVANTI itself, and with the location of interest (here, the Italian city of Siena). Three other questions try to elicit information about specific information interests related to disabilities.

Note that there is a danger in building upon explicit user statements: self-assessment is error-prone since users are often not correctly aware of things like their own capabilities (see e.g. Hothi & Hall, 1998; McGuire & Padawer-Singer, 1976; Nisbett & Wilson, 1977; Rich, 1983).[8] Some systems therefore present controlled queries, tests, exercises and so on that are aimed at a more objective assessment of the user. For example, the system SATELIT (Akoulchina & Ganascia, 1997) systematically asks users questions about concepts in a classification domain, following a specific knowledge acquisition method, and evaluates their responses to construct a model of the user's domain expertise. The web-based Lisp tutor ELM-ART II (Weber & Specht, 1997) makes students work on exercises and evaluates solutions to assess what has already been mastered. Other systems (e.g. Gates *et al.*, 1998) rely on users' self-reports about their usual activities rather than users' self-assessment. Websites directed at consumers can incorporate such data by masking quizzes as entertainment and offering incentives such as free merchandise.

In the terminology of Chin (1993), all kinds of asking the user can be classified as *active* acquisition, which is controlled by the interactive system. It is important to smoothly integrate active user model acquisition into the system-user dialogue. Decision strategies for initiating active acquisition processes have been investigated by, for example, Wu (1991).

One of the serious downsides of using initial interviews or similar initialization interaction activities is the *Paradox of the Active User* (Carroll & Rosson, 1987): users are motivated to get started and the

[8] In contrast, Srinivasan and Park (1997) found that people's self-explications concerning decision criteria are a good predictor for actual choices.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

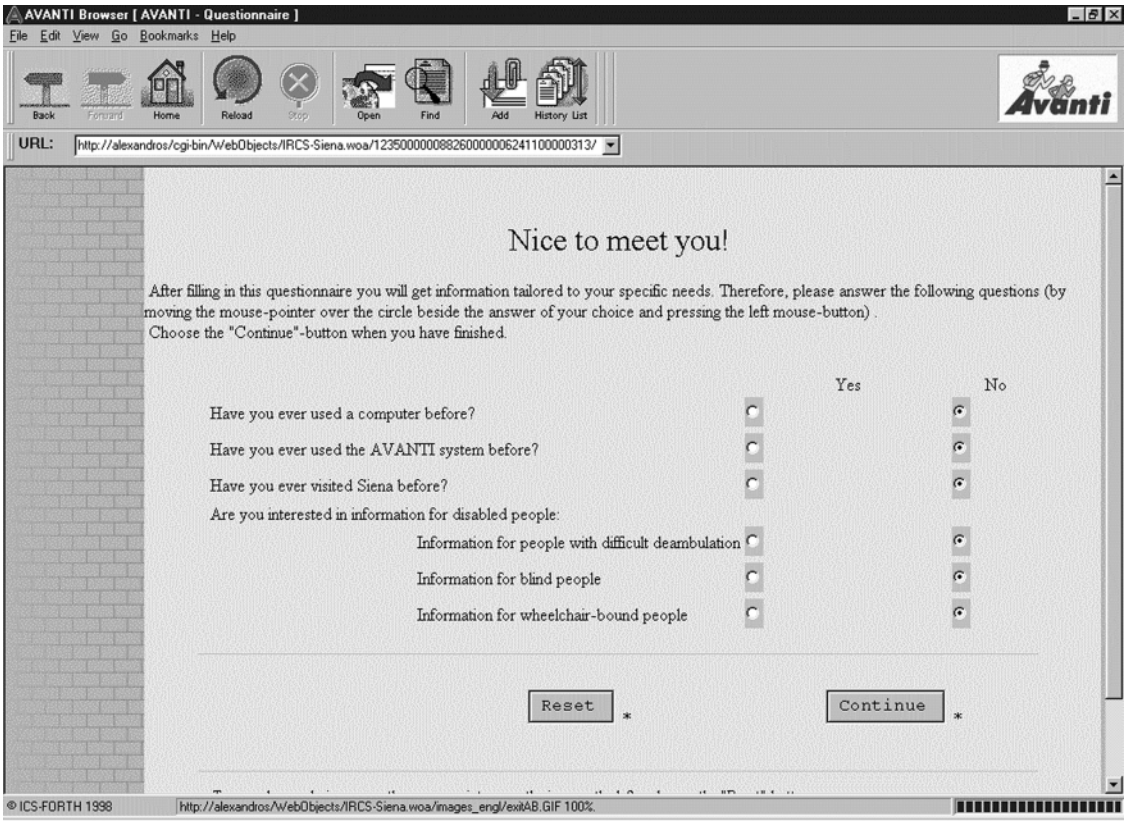


Figure 1 Initial interview in AVANTI

need to get their immediate task done or their immediate information need met. They do not spend time up front on set-up, manuals or online help. This is a paradox because users would save time in the long term by taking some initial time to optimise the system. In cases of competing information sources, users may simply refuse to visit the site if they have to respond to an interview first. If done at all, the acquisition phase should therefore be minimised and ideally be administered only after the user has already obtained some impression about the benefits that the site offers.

Users should therefore be given the opportunity to provide information about personal characteristics in a more self-initiated way, e.g. as part of preference dialogues (as described in Strachan *et al.*, 1997, 2000 and elsewhere) or at arbitrary moments of interaction with the system (Bares & Lester, 1997). The WWW standard P3P (Reagle & Cranor, 1999) will allow users to automatically tell sites with trusted privacy policies certain frequently requested personal data (for instance, name, address and phone and credit card numbers) without having to type it in for each site. User model servers (Kobsa & Pohl, 1995; Machado *et al.*, 1999; Orwant, 1995) also allow authorised networked applications from different sites to access data about the user from a central user model. Machine-readable data-carriers, particularly contactless chips that can be read from a distance and can be integrated into smartcards, watches or jewellery, allow users to wear and to easily supply personal data wherever they go (Fink *et al.*, 1997). The fact that user data can be read at kiosks, ATMs and so on before users start interacting with these devices is particularly useful for people with disabilities since applications can then adapt their presentation modality (see Section 5.2) to the specific user needs without any user intervention (van Schaik *et al.*, 1996).

3.1.2 Acquisition rules

Passive or implicit acquisition methods promise to be less disturbing or annoying to the user than active acquisition. By definition, passive acquisition methods do not initiate any interaction with the user. This and the following sections will mainly describe techniques for passive acquisition.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

```
(define-d-act :name REQUEST-EXPLANATION
  :parameters (topic)
  :presuppositions ((B S (not (B U topic)))))
```

Figure 2 Sample definition dialogue act type in modal logic (BGP-MS)

A frequently used means for generating assumptions about the user are acquisition rules, i.e. inference rules that are typically executed when new information about the user is available. In most cases, acquisition rules refer to observed user actions or to a more or less straightforward interpretation of user behaviour. Acquisition rules may be specific to a given application domain, or may be application-independent. The system KNOME (Chin, 1989) is a classic example of domain-independent acquisition. It uses heuristics like "if the user wants to know X, then the user does not know X". Other application-independent acquisition rules were investigated in GUMAC (Kass, 1991).

Pohl *et al.* (1995) discuss so-called dialogue acts, i.e. pragmatic elements in the interaction between user and system which often have implications for the user model. In the user modelling shell system BGP-MS (Kobsa & Pohl, 1995), dialogue act types were introduced that permit the definition of acquisition rules as classes of dialogue acts with identical implications for the user model.[9] Like acquisition rules, such dialog act types may be application-independent, but are typically more useful if embedded into some application scenario. For instance, in a hypertext system that offers explanations about specific terms in the text, a user can perform actions that can be interpreted as "requests for explanation" for these terms. In this case, it can be assumed that the user is not familiar with the underlying concepts (Boyle & Encarnacion, 1994; Kobsa *et al.*, 1994; Pohl *et al.*, 1995). Figure 2 shows the definition of the corresponding dialogue act type in the notation used in BGP-MS. Its parameter (namely the topic to be explained) also appears in the presupposition pattern of the dialogue type (which describes in a modal logic notation that system believes that the user is assumed not to know this topic). When an instance of this type is processed, the actual value of the topic parameter will be substituted into the pattern, so that finally a specific user model content is constructed.

Acquisition rules that are completely application-specific are perhaps the most popular means for user model acquisition, since they are easy to implement. However, they are not very flexible to use, and their properties typically cannot be described in a formal way (see Pohl, 1998, for a more detailed argumentation). Examples for the use of application-specific acquisition heuristics are discussed by Strachan *et al.* (1997) and Fink *et al.* (1998). In many cases, application-specific rules or procedures are applied to derive further conclusions from interpretations of user reactions to test or practice items (Akoulchina & Ganascia, 1997; Nwana, 1991; Weber & Specht, 1997).

One detailed example for the use of application-specific acquisition rules is the system TIMS (Tax and Investment Management Strategizer) of Strachan *et al.* (1997, 2000). The user model in TIMS consists of three variables that represent the user's level of expertise with respect to financial planning (which is the application domain of TIMS), familiarity with TIMS, and operating system (Windows) usage. Each variable may take one of the values "novice", "intermediate" or "expert". TIMS updates the values of these variables regularly, employing specific heuristic acquisition rules. Two rules that control the assessment of the user's familiarity with TIMS system usage are given in Figure 3.

3.1.3 Plan recognition

Plan recognition deals with reasoning about the goals that the user may pursue and the action sequence (plan) he or she performs to achieve them. A plan recognition system consists of a task knowledge base

[9] These implications can be derived from the so-called "presuppositions" of the dialogue acts that describe their applicability conditions.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

```
If the user has been away too long,
    Downgrade TIMS experience level by 1 category.
If the user has used the system long enough since the last update
    Upgrade TIMS experience level by 1 category.
```

Figure 3 Domain-specific acquisition rule in TIMS

that models possible user actions and relationships between these actions, and of a mechanism that identifies the current plan (and the associated goals) of the user from the observed interactions. In early work on plan recognition, symbolic methods played a dominant role (Allen & Perrault, 1980; Carberry, 1989, 1990; Weida & Litman, 1992). During the last decade, there has been an increasing amount of work on applying numeric (Albrecht *et al.*, 1997; Bauer, 1996) or graph-based (Lesh & Etzioni, 1995) techniques to plan recognition.

Plan recognition is especially promising for applications with a small number of possible goals and a small number of possible ways to achieve these goals. For example, in message centres and information systems, users often have specific goals such as listening to new messages, getting billing information or receiving weather forecast information for a local region. If the system recognises such goals (over time), it can provide shortcuts for these types of routine goals. Lesh *et al.* (1999) showed that plan recognition speeds up interaction considerably in a message-handling application.

3.1.4 Stereotype reasoning

A simple method for making a first assessment of others is to classify them into categories and to then make predictions about them based on a *stereotype* that is associated with each category. Stereotypes thus contain the standard assumptions that one makes about members of that category (McCauley *et al.*, 1980). The use of stereotypes in computer systems that maintain models of their users was introduced by Rich with the system GRUNDY (Rich, 1979, 1983). The main components of a stereotype are

(a) a body, which contains information that is typically true of users to whom the stereotype applies and
(b) a set of activation conditions ("triggers") for applying the stereotype to a user.

Stereotype reasoning, then, means to regularly evaluate the activation conditions of available stereotypes and, if a condition is satisfied for the current user, to integrate the contents of the corresponding stereotype as assumptions into the current user model for this particular user. The conditions may refer to current user data ("if the user model shows that the user is interested in childcare, activate the stereotype 'parent' "), usage data ("if the user bought at least two books on childcare, activate the stereotype 'parent' "), or environment data ("if the user's browser is a recent beta version, activate the stereotype 'computer wizard' ").

Besides explicit triggers, other mechanisms may be used for stereotype activation, like neural networks (Ambrosini *et al.*, 1997). A further development of stereotype reasoning was led by the idea that one can also group the required preconditions for the activation of stereotypes into disjoint categories, thus arriving at more general stereotype models. This "double stereotype" approach was introduced by Chin (1989) and further developed by Jameson (1992). Currently, stereotypes are always "hand-crafted", based on empirical observations like user type analysis or sales data. Section 4.3.2 describes beginning work on using machine-learning techniques for determining the structure and the contents of stereotypes.

The stereotype approach has been used in many systems (Ambrosini *et al.*, 1997; Ardissono *et al.*, 1999; Fink *et al.*, 1998; Kobsa *et al.*, 1994; Moore & Paris, 1992, to name but a few). The importance of the stereotype approach is reflected by the fact that stereotyping is one of the few common features of the so-called "user modelling shell systems" (Kobsa, 1995) that were built in the mid-1990s as toolboxes for the development of user modelling systems. The effectiveness of stereotyping depends

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

on the quality of the stereotypes, for example the number of different stereotypes known to the system, the accuracy of the assignment of users to stereotypes, and the quality of inferences that are drawn from stereotype membership. These, in return, depend on the quality of information about the user population and the extent to which subgroups with different application-relevant characteristics can be empirically distinguished (see Bushey, 1999, for a methodology to gain such information).

In customer relationship applications, stereotypes can be used to describe typical characteristics of customer classes (Ardissono & Goy, 1999; Ardissono *et al.*, 1999). Existing data about customers and their characteristics which are typically available from marketing databases may reduce the effort of creating a set of valid stereotypes of interest to a company, for instance about customer groups exhibiting different purchase behaviour.

3.2 Usage model acquisition methods

In several recent systems, user behaviour has not only been *observed* but also *modelled* as a direct basis for system personalisation. The terms "usage modelling" and "usage profiles" have been coined for systems like Flexcel (Krogsæter *et al.*, 1994) and Basar (Thomas & Fischer, 1996), which record user actions in order to obtain information about user behaviour. More sophisticated and ambitious techniques are employed by "interface agents" and "personal assistants" (Maes, 1994; Mitchell *et al.*, 1994), following the earlier idea of learning apprentices (Mitchell *et al.*, 1985). Such a system "becomes gradually more effective as it learns the user's interests, habits, and preferences" (Maes, 1994).

Currently interface agents mainly learn correlations between situations that the user encounters and the actions he or she performs. These data are used to predict user behaviour in future situations, to suggest appropriate actions to the user and perhaps to automatically perform actions on the user's behalf. For instance, the "calendar agent" (Kozierok & Maes, 1993) suggests reactions to meeting requests, while Maxim (Maes, 1994) proposes how to react to incoming e-mail messages.

Several machine learning algorithms (Dutton & Conroy, 1997) have been applied to interface agents. Memory-based learning (a kind of case-based learning), reinforcement learning and induction of decision trees (ID3) are used in scheduling agents (Kozierok & Maes, 1993; Mitchell *et al.*, 1994). Armstrong *et al.* (1995) compare different algorithms for classifying feature vectors with respect to their suitability for a WWW advisor.

Machine learning has also been used in genuine user-modelling systems to construct behaviour-oriented user models. Webb and Kuzmycz (1996) present the feature-based modelling approach to learning situation-action correlations (like interface agents) and demonstrate its use in tutoring systems. In the user modelling shell system Doppelgänger (Orwant, 1995), statistical and machine-learning methods are employed to collect evidence about user behaviour patterns.

In addition to learning to predict actions that a user may perform in a particular situation, one can also attempt to predict action sequences. The acquisition of typical action sequences of users goes beyond the task of learning situation-action relationships, since a (short) history of actions may be involved. The discovery of typical action sequences may lead to suggesting the introduction of new macro-operators to the user, or to predicting future actions and perhaps even performing some of these actions on behalf of the user, and to recommending actions based on action sequences of other users.

The behaviour patterns learned by Doppelgänger also include such typical (inter-)action sequences of the user. Doppelgänger uses hidden Markov models to acquire a model of usage sequences. Graph-based induction is another method for finding such patterns by processing observations and pre-defined domain knowledge (Yoshida & Motoda, 1996). In earlier work, attempts to learn grammar rules that express interaction regularities have been made by Hoppe and Plötzner (1991). Pohl (1992) describes special heuristics, which are in part motivated by results from psychology, for identifying candidates for typical action sequences within usage logs. In more recent work, several statistical methods for predicting sequences of user actions have been investigated and compared (Davison & Hirsh, 1998). A number of early systems that observe user behaviours and learn action sequences are described in Cypher (1993) along with descriptions of the tasks they perform.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

```
GET /index.html HTTP/1.0
Accept: */*
User-Agent: Mozilla/4.04 [en] (WinNT; I)
```

Figure 4 Example of an HTTP header field

3.3 Environment data acquisition methods

This section discusses specific methods for acquiring data about the software and hardware environments and the locale of the current user (these data were discussed in more detail in Section 2.3). A key problem is the mapping of a (typically identifiable) physical device to an individual person. In some contexts (e.g. office PCs) it is reasonable to assume that the system is a single-user machine. In such a case, environment data can be stored with the other data about this user. This approach will not be possible in other settings, such as information kiosks in showrooms, since these are inherently multi-user. In such a case, environment data must be represented separately from user data.

3.3.1 Software environment

Many websites take the software constraints of the browser into account. Information about the web client can be obtained from the header of HTTP requests that are received by the server. Each request carries values for a number of variables (also called header fields). The example in Figure 4 shows an HTTP GET request. The value of the "user-agent" variable indicates that the request is issued by a Netscape browser, version 4.04. The fact that the version is English suggests that the user is able to read English. Note, however, that the client can cheat and deliver fake information. Less common browsers quite often pretend to be frequently used browsers when they have the same capabilities.

3.3.2 Hardware environment

Hardware constraints are often difficult to assess. Some aspects may be guessed from the type of browser used (e.g. small mobile devices need special software for Web browsing or display). Other hardware characteristics, like bandwidth and processor speed, are much more difficult to determine. In particular, the effective amount of available bandwidth depends on the whole Internet route that web pages take to reach the client machine, and this amount may change significantly within short periods of time.

Systems like Hyperspace Agent (Fuller & de Graaff, 1996), AVANTI (Fink *et al.*, 1998), TELLIM (Joerding *et al.*, 1998) and AWCD (Chen *et al.*, 2000) predict the download time of hypermedia objects using information about their sizes and feedback from network probes that become injected into requested HTML pages. This prediction can then be used as one factor in determining the adaptive composition of a page, e.g. by replacing high-resolution images and videos by less bandwidth-intensive, but nevertheless appropriate, equivalents.

3.3.3 Locale

Locality information can be acquired by a wide range of methods. A detailed discussion of these methods is beyond the scope of this article. We only give a short list of possibilities.

For *stationary* networked devices, information about the relative location with respect to other objects is currently typically recorded in a database (e.g. the mapping of a device such as a printer to a particular room or building). Developing standards and interfaces (Jini, XML) will enable the querying of such devices directly. *Mobile* devices can provide locality information actively or passively either through dedicated technology such as infrared sensors and transmitters in closed environments, or general technology such as global positioning systems (GPS).

Technologies that are available today for determining users' positions include

- electromagnetic fields (GPS, DGPS, BlueTooth, contactless chips, radio, cellular phone systems, DECT and so on),
- ultrasound,

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

- optical recognition (laser, digital cameras with pattern recognition),
- infrared (IR-systems based on IrDa or RCF, laser, barcode reader and so on), and
- datagloves, force feedback devices, and sensors registering tilting, acceleration or movement of body parts to determine users' orientation and the position of body parts.

Typically, a combination of methods has to be used for best results, e.g. to enable both indoor use and outdoor use and to cover the range of resolutions needed for a service. For example, the accuracy needed to determine which direction a vehicle is moving on a highway to provide targeted traffic alerts differs from the level of detail and speed of update needed for supporting navigation in an unfamiliar building such as a museum or public administration complex.

4 Representation and secondary inferences

Once user and usage models have been acquired, they need to be represented in order to be available for further exploitation. Some personalised systems can get along with fairly simple model structures. They contain mostly feature-value pairs like "CRYPTOCOCCUS KNOWN" (Sleeman, 1985), which expresses the system's assumption that the user is familiar with the concept "cryptococcus". Quite a few systems can also base their content adaptations (Section 5) directly on the results of the user, usage and environment model acquisition process (Section 3) without very much need for intermediate steps.

Many personalised hypermedia systems, however, have higher demands with respect to user and usage model representation, and also need to employ inferences to further augment the user and usage model based on initial acquisition results and, e.g., domain knowledge and knowledge about other users. Such inferences do not consider the current user input any more. They have been called *secondary*, in contrast to the primary inferences of user model acquisition. Since secondary inferences operate on user or usage model contents, methods for user model representation and secondary inferences are often closely related.

In this section we will describe the most common representation approaches and the inference techniques associated with them. From an epistemological point of view, several types of reasoning can be distinguished, namely deductive reasoning (from the more general to the more specific), inductive reasoning (from specific cases to the general case) and analogical reasoning (from similar cases to the present case). All these methods are being employed in the area of personalised hypermedia systems and will be described in the next three subsections.[10] More than one type can in principle be employed in a system, but in practice this is rarely the case.

4.1 Deductive reasoning

4.1.1 Logic-based representation and inference

The use of logic-based methods in user modelling systems has been analysed in detail in Pohl (1998). Most user modelling shell systems (Kobsa, 1995) offer logic-based methods along with an access function interface.

For an example, let us take a look at the adaptive hypertext system KN-AHS (Kobsa *et al.*, 1994), which makes use of the user modelling shell system BGP-MS (Kobsa & Pohl, 1995) to maintain assumptions about user beliefs (namely the user's familiarity with the technical concepts in a hypertext). BGP-MS represents both assumptions and domain knowledge using a concept formalism (namely a KL-ONE derivative, see Kobsa, 1991), with each textual link anchor of the hypertext corresponding to one concept. For instance, the assumption that the user knows the concept X is represented by entering a representation of that concept into the user modelling knowledge base. Once added to the user model, the assumption may trigger further meta-level reasoning based on concept

[10] A fourth type, abductive reasoning (from the consequences to the premises), is also sometimes employed (e.g. in plan recognition, see Section 2.1.5), but will not be discussed here.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

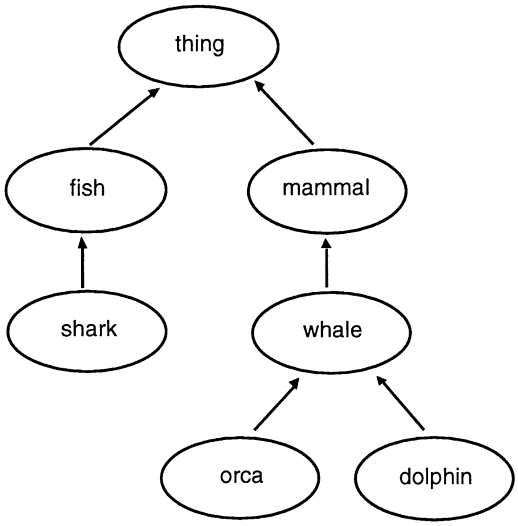


Figure 5 A concept hierarchy in graphical notation

relationships represented as domain knowledge. Figure 5 shows a graphical notation of a concept hierarchy (i.e. an ontology) from the animal kingdom, as it might be used for representing user knowledge about this domain (cf. Akoulchina & Ganascia, 1997; Ardissono & Goy, 1999, 2000b; Ardissono *et al.*, 1999; Milosavljevic, 1997).

While concept formalisms have been employed quite frequently for representing user models, there are numerous other logic-based formalisms that have been utilised as well, ranging from propositional calculus to modal logics. A comprehensive overview which also discusses the application of inference techniques in logic-based user modelling is given in Pohl (1998).

The two crucial shortcomings of standard logic-based approaches to user model representation is their limited ability to deal with uncertainty (see the next section), and with changes to the user model. In many applications, relevant characteristics of users change over time, so that a user model representation method should be able to deal with such changes. One possible solution to this problem is provided by the arsenal of so-called non-standard logic-based methods. Non-monotonic logics have been developed to deal with changes and revisions in knowledge bases that are represented in a formal logic. One possibility for implementing non-monotonic logic and the associated non-monotonic reasoning procedures is to employ so-called truth-maintenance techniques. Since this is a difficult task, it is not surprising that truth-maintenance methods have mostly been implemented in generic user modelling tools (Brajnik & Tasso, 1994; Ikeda & Mizoguchi, 1994; Paiva & Self, 1995).

4.1.2 Representation and reasoning with uncertainty

In order to cope with the uncertainty that can be regarded as inherent in user modelling, many systems have used numerical methods to rate the validity of user model contents by evidence values (for a detailed overview see Jameson, 1996). One example is the system HYDRIVE (Mislevy & Gitomer, 1996), which models a student's competence at troubleshooting an aircraft hydraulic system. HYDRIVE employs Bayesian networks (BNs). The probability distribution of network nodes is used to explicitly represent variables like "knowledge in electronics" or "strategic knowledge", which are organised into different levels of abstraction. At the most concrete level, nodes represent "interpreted actions". Such nodes store the probabilities of an observed sequence of user activities belonging to one of a fixed number of action categories. For input into the user model, HYDRIVE creates a network node and adds it to the net with the probability of the observed action category set to 1. This is the acquisition process of HYDRIVE. Secondary *inferences* are made by propagating probabilities from the interpreted action nodes to higher level nodes like "strategic knowledge". That is, possible inferences are mainly determined by predefined domain knowledge as represented in the structure of the BN and in the conditional probabilities associated with the network links.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Recently, Bayesian networks have been among the most popular methods for representing uncertainty in user models in very different domains. Other examples include, e.g. the consultation system for data analysis EPIAIM (De Rosis *et al.*, 1992), the natural-language dialogue system PRACMA (Jameson *et al.*, 1995), and the Office Assistant by Microsoft (Horvitz, 1997; Horvitz *et al.*, 1998).

Other evidence-based techniques which were employed for probabilistic user model representation are linear parameters, fuzzy logic, and Dempster-Shafer theory[11]. Linear parameters (see, e.g. Sleeman, 1985) are usually used together with feature-value pairs that represent assumptions about user characteristics; the additional parameter simply adds the degree of uncertainty of this assumption.

Fuzzy logic offers means for representing vague concepts and rules of thumb. Arguments supporting the use of fuzzy logic are: (1) that people (users) reason in terms of vague concepts when experiencing uncertainty, and (2) that the information that users may supply about themselves is vague. An example for a system based on fuzzy logic is the Sales Assistant (Popp & Lödel, 1996), which aims at predicting the suitability for the user of products in electronic product catalogues (e.g. personal computers). The system has to correctly interpret user requirements like "the computer should have 64MB RAM" (which probably does not mean that any other amount of RAM is completely unacceptable). In this vein, fuzzy mechanisms are introduced that evaluate products with the desired value as ideal, and products with other values as less suitable with respect to the RAM-size attribute (but not as completely unsuitable, so that a sub-optimal value for one attribute will not overly disturb the positive assessment of a product with optimal values for all other attributes).

4.2 Inductive reasoning: learning

Inductive reasoning about the user involves monitoring users' interaction with the application system and drawing general conclusions based on a series of observations. This type of reasoning is often called "learning about the user". Learning algorithms can in principle be employed for inferring any type of assumption about the user. Specifically in customer-relationship management they are, however, currently mostly used for constructing users' explicit or implicit "interest profiles".[12] Interest profiles represent users' interest in or affinity to objects, based on an assessment of their interest in specific features of these objects. Such a profile can be regarded as a user model that allows for personalised recommendation or other kind of filtering (cf. Oard, 1997). This kind of filtering is called feature-based, since it relies on features that describe characteristics of the objects of possible user interest.[13]

For example, an application developer may make the assumption that user interest in movies is determined by preferences about the genre, the actors, the director, and other movie features. The system attempts to assess the user's preferences with respect to each of these features. For each movie, the feature values can then be compared with this user interest profile and the movie be rated as being more or less interesting to the user. The ratings can be used for personalised filtering, e.g. recommending movies that are rated highly based on the learned interest profile.

There are several techniques to acquire interest profiles. An early system for filtering Internet news was developed by Jennings and Higuchi (1993). It employed a neural network for constructing an interest profile, but no comparison with more traditional information-retrieval (IR) methods was performed. During the training period, users rate documents as being interesting or not interesting for them. For each content-bearing word that occurs at least twice in the set of training documents, a node is introduced into the neural network whose initial activity corresponds to its frequency in the positively rated documents. The link weights correspond to the co-occurrence frequency of the linked

[11] Applications of Dempster-Shafer theory will not be discussed here. See Bauer (1996) for an example.

[12] Following the terminological tradition in this area of user modelling research, we will use the terms "user profile" and "interest profile" instead of "user model" and "interest model" in the remainder of Section. Both terms have their origins in information retrieval research.

[13] The term "content-based filtering" is used in those cases in which the content of documents is expressed using those words as features that are most representative for the content of these documents.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

words within the same documents. When new documents are presented to the trained neural network, the nodes that correspond to the meaning-bearing words in the document become activated with their initial activity, and propagate their activity via the differently weighted links to other nodes. After a certain period of time, the overall network activity ("temperature") is measured and the new document rated as interesting for the user if the temperature exceeds a given threshold.

Syskill and Webert (Pazzani & Billsus, 1997) demonstrates the use of machine-learning techniques for acquiring interest profiles from explicit user ratings of a selected set of documents. Based on binary ratings, supervised learning algorithms can be employed, which take the documents as positive and negative training examples. One of the most important questions when applying machine learning is how to code the training examples. In the case of Syskill and Webert, a document is represented as a vector of length *n* that represents *n* features of the document. The 128 most informative words in the rated documents are determined based on a certain information-gain measure. Each of these words becomes a feature with two possible values: the word is either present or absent in the current document.

Pazzani and Billsus applied several learning algorithms to the Boolean feature vectors that describe the rated documents, namely the nearest-neighbour algorithm (Duda & Hart, 1973) and its PEBLS variant (Cost & Salzberg, 1993), induction of decision trees (e.g. ID3; Quinlan, 1986), two neural network approaches, and the naive Bayesian classifier (Duda & Hart, 1973). Several of these algorithms were found to be similarly useful. Decision-tree induction was less accurate; this method is preferable for data with a smaller number of features. The nearest-neighbour algorithm also did not work in a satisfactory way. Finally, the Bayes classifier was chosen as the default algorithm in Syskill & Webert. Improvements to learning accuracy were achieved by selecting features that are particularly relevant to the classification task, e.g. revising a user-defined profile (where the user selects relevant features and pre-rates the features) can increase the accuracy of the classification process.

Table 1 gives an example of the user data inferred by Syskill and Webert from usage data (i.e. user ratings). The features are a pre-selected set of words from the biomedical domain. The numeric values express the probability that the corresponding word occurred in previously seen documents that the user rated as "hot" and "cold", respectively. They are used by the Bayesian formula which controls the classification of web pages that have not been rated by the user before. That is, this data is used for classification in a "closed-loop" fashion only; it does not really become entered into a user model. However, since probabilities have a clear semantics in the case of Bayesian classification, one can consider reusing such learning results or transforming them into reusable user-profile contents. This makes sense if the features (the words in this case) are relevant to additional adaptivity tasks besides web-page recommendation. For instance, if the words are related to the domain of programming in Java, learning results of the web page classifier could influence other recommendation processes where the same features are relevant (e.g. news articles or emails related to Java programming). Section 4.4 describes a general approach to reusing learning results via user profiles.

In many other systems that use inductive learning for personalisation, interest profiles also refer to the information content of documents, where the features are words that are considered to be more or less relevant to the user (based on their occurrence in the documents). Major examples for interest-based adaptive information filtering/retrieval systems besides Syskill and Webert are Fab (Balabanovic, 1997; Balabanovic & Shoham, 1997) and Letizia (Lieberman, 1995). For describing users' interest in text documents, classical approaches of information retrieval can be used as well, as has been shown by, e.g. Balabanovic (1998). Both the documents and the interest profile can be described using the

Table 1 Syskill & Webert user data

grants .7.2	database .8.1	genome .6.3
molecular .6.1	protein .5.2	prediction .9.1
classification .9.1	structure .6.2	function .6.1
webmaster .05.1	com .1.4	

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

vector-space model of IR. That is, they are vectors relative to some dictionary vector, with a slot and a specific weight for each word of the dictionary. In a document vector, the weight expresses the importance of the word to the document; in a profile vector, the weight expresses the relevance of the word to the user. In this respect, the user profile can be considered a query. Moreover, as is done for query processing, profile vectors are simply multiplied with document vectors to yield an assessment of document relevance to the user. The acquisition of such interest profiles can be achieved by IR techniques as well. The user must score documents to give relevance-feedback to the system. These scores are used to update the profile vector.

As mentioned before, numerous other examples of inductive learning in adaptive systems exist in the literature. Which learning technique to use is determined by many factors. For instance, Joerding (1998, 1999) employs the CDL4 algorithm (Shen, 1996) to learn user interest in product characteristics since CDL4 can process set-valued features.

4.3 Analogical reasoning

This section describes two related approaches to user profiling that exploit the fact that web-based systems have a large number of users, and use analogical reasoning based on recognised similarities between users. Section 4.3.1 describes methods for clique-based filtering (matching a single profile with profiles of similar users) while Section 4.3.2 describes clustering for grouping user profiles.

4.3.1 Clique-based filtering

As has been described in Section 4.2, there are systems that acquire models of user interest by analysing the features of the objects in which the user has expressed an interest (these features are often representative words of the information content of these objects). This approach has shortcomings, however. The content of objects may not be easy to analyse by computers (e.g. in case of multimedia objects), content may not be the only aspect that determines a user's interest in an object (e.g. an interesting web page may not be visited if it takes several minutes to download), content may be hard to express in feature vectors (e.g. in the case of books and films) and, finally, users' interests may not even be based on features of objects. Feature-based approaches also suffer from the requirement that a certain minimum number of example objects that were rated by the particular user must be available and need to be analysed in order to find regularities across objects.

These problems are addressed by systems that for a given user try to find other users who show similar interaction behaviour. Such systems adapt to the individual user based on the behaviour of these "interest neighbours". The set of similar users constitutes an implicit profile of the individual. Following Alspector *et al.* (1997), we will call this approach "clique-based filtering".[14]

For instance, the system GroupLens (Konstan *et al.*, 1997) computes correlations between readers of Usenet newsgroups by comparing their ratings of news articles. The ratings of an individual user are used to find related users with similar ratings, and their ratings are processed to predict the user's interest in new articles. Siteseer (Rucker & Polanco, 1997) builds virtual communities of users based on their bookmark folders. For individual users, the system tries to find "recommenders", i.e. other users with overlapping (sub-)folders in the bookmark hierarchy. The degree of overlap determines the degree of qualification as a recommender. As in GroupLens, these correlations between users are exploited to make recommendations: Siteseer provides as recommendation those pages which have been bookmarked by the user's virtual neighbours, giving preference to pages drawn from folders with the highest overlap as well as those held within multiple folders in the neighbourhood.

In general, the process of computing a recommendation has three steps.

(a) **Find similar neighbours.** Standard similarity measures (such as the vector-space model) are used to compute the distance between the current user's representation (such as a feature vector) and the

[14] The more frequently used term "collaborative filtering" does not seem to be appropriate any more since this approach has only very few collaborative characteristics.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

representation of a set of users. In smaller applications these may all be users; in larger systems statistical sampling methods are used to find a representative subset for which similarity is computed.

(b) **Select a comparison group of neighbours.** Once the distances to other users are available, a set of "closest" users must be selected. One can either choose a fixed threshold approach, where all users within a given distance are selected (at the possible risk of obtaining too few or too many closest users), or one can select a set of fixed-n closest users (risking that some of these may be poor matches or that users are ignored that could add to the quality of the prediction).

(c) **Compute prediction based on (weighted) representations of selected neighbours.** Different approaches are possible to determine whether a particular neighbour is likely to be a good predictor. Most of these evaluations are based on statistical measures such as the overall rating average of the person, the person's average deviation from the group mean and the correlation of deviations between the target user and the neighbour.

Breese *et al.* (1998) discuss and compare a number of possible clique-based filtering algorithms. In general, the performance of these methods is difficult to predict and highly dependent on the rating distributions in the user population. A core problem is the sparse population of the matrix of users and rated objects, unless there is a large common ground (as is the case with popular movies). Because of this sparsity, there may be high correlations between two users in their ratings based on very few common ratings, while the predictive power of these high correlations is low for the rating of new items. Alspector *et al.* (1997) compare feature-based and clique-based filtering in the area of movie selection.

Clique-based filtering is becoming increasingly used in customer-relationship applications. Major current sites which employ this technique include Amazon.com, the Music Advisor of musicmaker.com, bol.com (in combination with feature-based filtering), the Album Advisor™ of CDNOW, moviefinder.com, the personal television guide ptv.ie, LetsEatOut.com and many others.[15]

4.3.2 Clustering User Profiles

As we have seen above, users can be characterised by determining a set of similar users. No explicit user profile is thereby constructed. Yet, when explicit profiles are available and necessary, there also exists the possibility of exploiting similarities between users. In Section 3.1.4, we described stereotype reasoning as a quick way of assessing the user by classifying him or her into one or more from a predefined set of user categories, assuming that the user shares characteristics typical of this group (which are explicitly given). A problem with this approach is that the classification is rigid: in most systems, reclassification is not well supported, and it is difficult to specify exceptions, i.e. to respect individual characteristics that a specific user is already known not to share with the group.

The system Doppelgänger (Orwant, 1995) used a related, but different, approach for quicker user assessments. The system forms explicit user profiles using machine-learning methods and statistics. Since Doppelgänger is a user modelling server, it stores profiles of different users. The system applies a clustering algorithm to the available profiles, to find similar users and to form group profiles. Individual profiles can be compared with these group profiles. Similar to stereotype methods, information from related group profiles can be employed if the corresponding information is not available in the individual profile. In contrast to stereotype reasoning, a degree of validity can be associated with the group values: it can be computed from the strength of the value in the group model (as computed by the clustering algorithm) and the distance between the individual user profile and the group profile. Through regular application of the clustering procedure, changes in individual profiles are taken into account.

Orwant referred to such automatically acquired group data as "communities". More recently, Paliouras *et al.* (1999) explored a mixed approach to reasoning about user similarities. They use

[15] Amazon.com and CDNOW report a repeat customer rate of 58% and 55% respectively, which is well above the average of 35–40% for e-commerce websites reported by Forrester Research (http://www.iconocast.com/icono-archive/ICONO.070198.html)

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

learning techniques both to determine the contents of given stereotypes (supervised learning) and to build communities from individual users' interest profiles (unsupervised learning).

4.4 A hybrid approach: user profiles as learning results

In Section 4.1 we described several methods for representing user profiles explicitly. These explicit representations are independent of the acquisition processes that yield the contents to be represented in the profile. In contrast, the learning methods described in Section 4.2 only create very simple user profiles that are specific to the given learning method, or no profiles at all. In these systems a kind of "closed-loop adaptation" can be observed: a learning method processes one specific kind of input as training data, stores the results in a proprietary format and uses these results for a specific adaptation task that is related to the original input information.

A typical example of closed-loop adaptivity is feature-based filtering. Imagine a system that aims at recommending music CDs based on an interest profile. As training input, feature-based descriptions are needed of CDs in which the user has shown an interest and (for most learning methods) of CDs in which the user has shown no interest. A supervised learning algorithm can be employed to learn to classify CDs "interesting" and "not interesting". It then depends completely on the learning algorithm how the results of the learning process are represented and stored. Algorithms for decision-tree induction create decision trees, neural networks modify the weights of the connections between network nodes, and a Bayesian classifier computes a number of probabilities for each feature of the CD descriptions. The primary purpose of learning is not to acquire a user profile that may be widely employed but to allow for the classification of further CD descriptions into "interesting" and "not interesting", and subsequent recommendations based thereon. An example of such a feature-based system with closed-loop adaptivity is Syskill & Webert (Pazzani & Billsus, 1997), which was discussed in Section 4.2. Another, perhaps more typical, example is the system of Jennings and Higuchi (1993) in which the user's "interest profile" is distributed over a whole connectionist classification network.

LaboUr (Pohl & Nick, 1999) defined a framework for a hybrid approach that exploits the advantages of both explicit and implicit user profiles. A LaboUr system (see Figure 6) accepts and stores observations about the user. Learning components (LCs) and acquisition components (ACs) can select appropriate user observations. LCs generate usage statistics and patterns, which will if possible

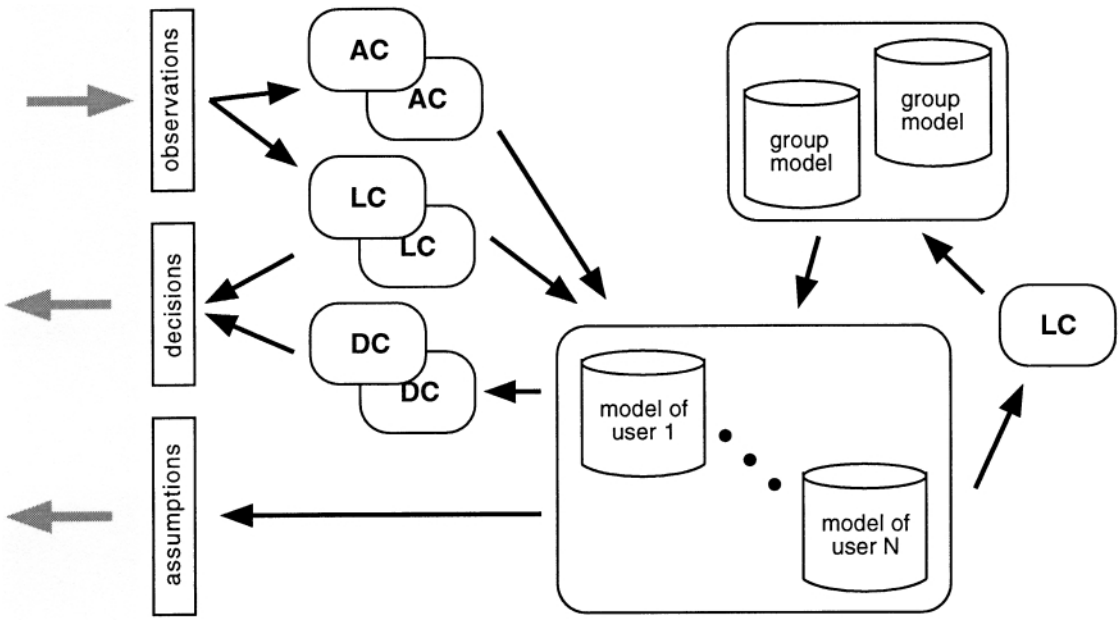


Figure 6 The LaboUr approach using machine learning for user modelling

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

subsequently be transformed into explicit assumptions. ACs support heuristic rule-based acquisition from isolated observations (see Section **?**). They directly generate user-model contents, which may be usage-related or user-related. LCs can be consulted in adaptivity decisions and will base their decisions on learning results. Additional decision components (DCs) that refer to explicit user models may be present in a LaboUr system. Besides supporting decisions, a LaboUr system can also provide direct access to user model contents.

A LaboUr system can maintain several user models. If this is the case, further learning components can be used for clustering user models into user group models. User models can be associated with group models which may contain complementary information (see the parallels to the stereotype approach described in Section 3.1.4). LaboUr is an open user-modelling architecture in the sense that several acquisition components may contribute to the user model, which in turn can support several kinds of adaptation or adaptive applications.

5 Adaptation production

The previous sections described the user, usage and environment data on which to base an adaptation, the methods that can be used to acquire them, the methods used to represent them and the methods used to draw initial inferences based on observations as well as secondary inferences based on user-model contents. This section explains possible types of adaptation to the user, usage and environment data of an individual user.[16]

Adaptations in hypermedia systems can take place at different levels that we will discuss in the following sections:

Adaptation of content: Section 5.1 describes methods for personalising the content of hypermedia objects and pages in accordance with user, usage and environment data.

Adaptation of presentation and modality: Section 5.2 describes methods for changing the presentation and media format, and the interaction elements.

Adaptation of structure: Section 5.3 describes methods for personalising the presentation of links based on user, usage and environment data.

5.1 Adaptation of content

We will first describe the major personalisation functions that content adaptation serves in current personalised hypermedia systems, and then discuss the techniques that have been developed to date to adapt hypermedia content to user needs.

5.1.1 Some personalisation functions of content adaptation

The adaptation of content for hypermedia-based customer-relationship applications to user, usage and environment data can in principle serve any communicative function and take on any presentation form. Frequently used personalisation functions that are interesting from the point of view of customer-relationship management are given below.

Optional explanations can help users who lack the necessary background knowledge better understand certain items in a hypermedia page. Such explanations would be redundant and even distracting for users who already have this background knowledge and they should therefore be omitted for such users (Boyle & Encarnacion, 1994; de Rosis *et al.*, 1992; Höök *et al.*, 1996; Kobsa *et al.*, 1994).

[16] To some extent we will therefore follow the classification of Brusilovsky (1996) that focuses on adaptation in hypermedia-based tutoring systems. The focus of the present paper, namely adaptation in hypermedia-based customer-relationship applications, and the consideration of usage and environment data in addition to user data, make it necessary, however, to considerably augment and revise this previous classification.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Optional detailed information can improve the relevance of a hypermedia page for users who are interested in the items on which detailed information can be given. It can also help keep a hypermedia page interesting and challenging for users who are already basically familiar with these items (Boyle & Encarnacion, 1994; Kobsa *et al.*, 1994).

Personalised recommendations inform users about available offerings in which they may be interested. So far, adaptive suggestive arguments consisted mainly in presenting or highlighting those products in which the user is probably interested (Smyth & Cotter, 2000), or emphasising those product characteristics in which the user is presumably interested (Ardissono & Goy, 1999; Jameson, 1989; Jameson *et al.*, 1995; Joerding *et al.*, 1998).

Theory-driven presentation caters to users following more general didactic, rhetorical or psychological principles, often combined with restricted natural-language generation (De Carolis *et al.*, 1998; de Rosis *et al.*, 1993; Hohl *et al.*, 1996; Jameson, 1989; Paris, 1989). A special case is the adaptive comparison that relates new information to a user's knowledge or specific interests (e.g. Fischer *et al.*, 1990; Kay & Kummerfeld, 1994; McCoy, 1989; Milosavljevic & Oberlander, 1998).

Optional opportunistic hints are based on users' presumed interests and on current circumstances. Examples include news flashes, pointers to special discounts, pointers to interesting objects that are currently in the user's proximity and so on.

In the next few sections we will describe the techniques that have been employed so far for realising content adaptation in hypermedia systems.

5.1.2 Techniques for content adaptation

A number of techniques on different levels of granularity and localisation have been developed so far for adapting hypermedia content to different user needs. These techniques will be described below.

5.1.2.1 Page variants. In the page-variant approach, adaptive system behaviour is realised by authoring different versions of all pages in which adaptation occurs. From a technical point of view, this approach is relatively simple since adaptation at runtime is confined to adaptive page selection. The approach is, however, also cumbersome since a completely new page has to be written for each variation of local adaptations that may occur on a page, and inflexible since many pages have to be modified if a single local adaptation is changed. Page variants for users pertaining to different stereotypes have been used, for example, in SYPROS (Gonschorek & Herzog, 1995), ORIMUHS (Encarnação & Stoev, 1999), and KBS Hyperbook (Henze & Nejdl, 2000).

5.1.2.2 Fragment variants. In the fragment-variant approach, variants are authored for each adaptive page fragment. At runtime, the appropriate fragments are included into a static page frame. The granularity of a page fragment corresponds roughly to a paragraph of text, an image, a table, a video and so on. In many cases, a fragment corresponds to a content-bearing hypermedia element (see Section 1). This approach is technically more difficult than page variants since it requires the dynamic generation of web pages at runtime. Major techniques employed to date for this purpose include

- proprietary enhancement of HTML to define fragment variants and conditions for their usage (de Bra & Calvi, 1998; Kay & Kummerfeld, 1994),
- standardized enhancements of HTML, like CGI scripts (De Carolis *et al.*, 1998) and Java Servlets (Gates *et al.*, 1998),
- the composition of pages using middleware (like WebObjects (Apple, 1999) in the case of AVANTI; Fink *et al.*, 1998),
- XML (Koike *et al.*, 1999),
- programmable web servers (like CL-HTTP (CL-HTTP homepage) in the case of Weber and Specht, 1997),
- rules in the case of KN-AHS (Kobsa *et al.*, 1994), AVANTI (Fink *et al.*, 1998), BroadVision (2001) and ATG (2001).

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Fragment variants can be used together with page variants, e.g. in Anatom-Tutor (Beaumont, 1994) for further adaptation to the user's knowledge after selecting the most suitable page variant according to the user's background. An intermediate kind are frame variants where personalisation is confined to one or more frames on a page. In each frame, then, a choice among complete frame variants is made (see e.g. Joachims *et al.*, 1997).

A frequently used special case of fragment variants are optional fragments which may or may not be inserted into a page frame depending on, for example, the user's knowledge and interests (see e.g. Fink *et al.*, 1998). A simple technique to implement optional text is to attach a condition to each optional chunk of text which determines whether to include or rather omit the textual element. This was used in early systems like ITEM/IP (Brusilovsky *et al.*, 1993) and LISP-Critic (Fischer *et al.*, 1990). Similar in spirit but more sophisticated is AVANTI (Fink *et al.*, 1998) which separates the condition rules from the text elements to avoid low-level authoring when writing adaptive hypermedia pages.

5.1.2.3 Fragment colouring. In the fragment-colouring approach, the content of the hypermedia remains unchanged across all users. For each individual user, though, certain elements of the hypermedia presentation may be marked out, e.g. as being important, irrelevant or too demanding for him or her. The major advantage of fragment colouring lies in the fact that all users can see all available information. Errors in the assessment of users (like underestimating or overestimating their levels of expertise) is therefore likely to have less critical effects as when information is hidden from the user (like with the page or fragment variant approach). Fragment colouring can, however, only be applied in such areas where content can be presented in the same formulation to all users, and where the variability of adaptation across all users is relatively low (since every user always sees the superset of the individual adaptations to each user). Fragment colouring, therefore, has so far not been very much applied (see e.g. Hothi *et al.*, 2000).

5.1.2.4 Adaptive stretchtext. Stretchtext is "elastic" text that the user can extend or collapse by clicking on it with the mouse. The length of the adaptable text is usually very small (in the range of a few words or one sentence). In personalised hypermedia systems, stretchtext can additionally be automatically expanded and collapsed by the system, taking the user model into account. The advantage of adaptive stretchtext is that users can still adapt the page content manually if the adaptation that was automatically performed by the system is not appropriate or desired.[17] This user reaction, then, is a good source of further assumptions about the user that can augment or refine the user model (see Section 2.2.1).

MetaDoc (Boyle & Encarnacion, 1994) and KN-AHS (Kobsa *et al.*, 1994) are examples of hypermedia-based online information systems which use stretchtext to personalise the hyperdocument. In the KN-AHS page shown in Figure 7, the user is assumed not to know the technical term "physiologische Krümmungen". Therefore, additional explanatory information is automatically inserted, namely an optional page fragment containing a thumbnail graphic and stretchtext that expands to an explanation (the text paragraph marked with **).

5.1.2.5 Adaptive natural-language generation. Some hypertext systems use natural language generation techniques to create alternative text descriptions for different users. A simple approach are text templates with slots that can be filled with, for example, descriptions of different complexity based on the user's level of expertise (Ardissono & Goy, 2000a). PEBA-II (Milosavljevic & Oberlander, 1998) produces different descriptions and comparisons of animals based on a taxonomic knowledge base depending on whether the user is assumed to be a layperson or an expert. ILEX (Milosavljevic & Oberlander, 1998) generates interest-tailored descriptions of objects from the National Museums of Scotland's 20th Century Jewellery Gallery. Figures 8 and 9 show two descriptions which are tuned to different user interests. ARIANNE (De Carolis *et al.*, 1998) employs natural-language generation to

[17] Stretchtext therefore can also be regarded as a special case of text fragment variants that are both under the control of the system and that of the user.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

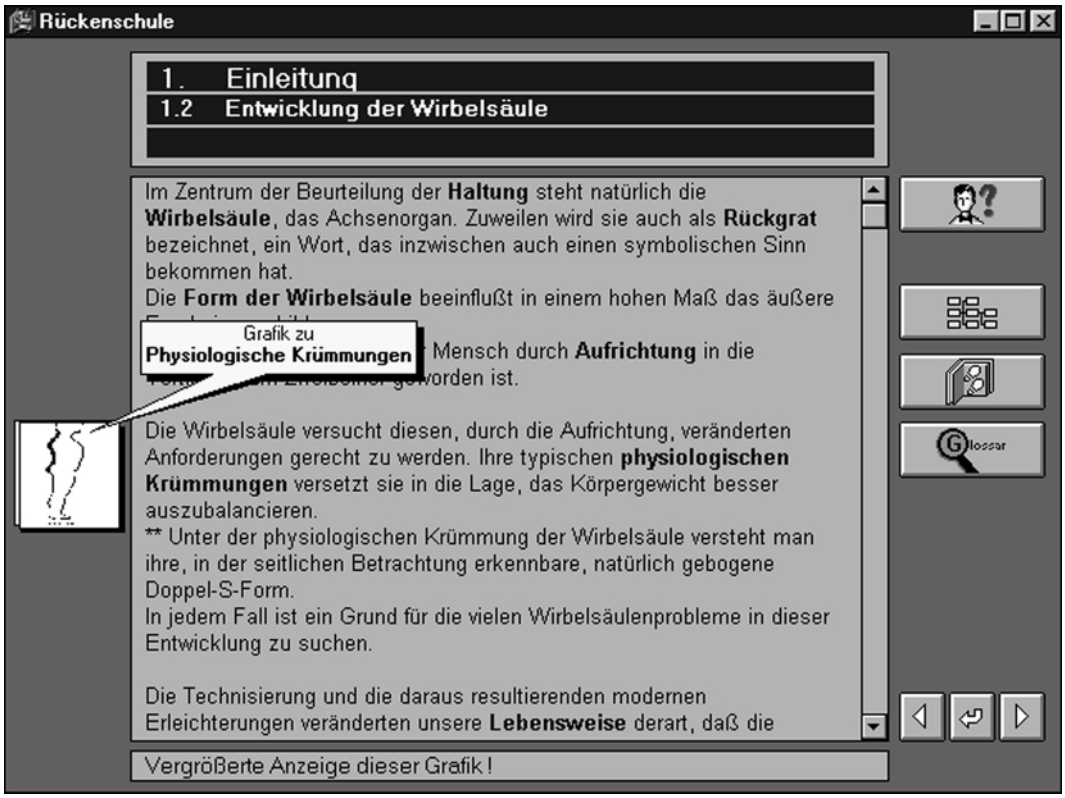


Figure 7 KN-AHS adaptive hypermedia course "Rückenschule"

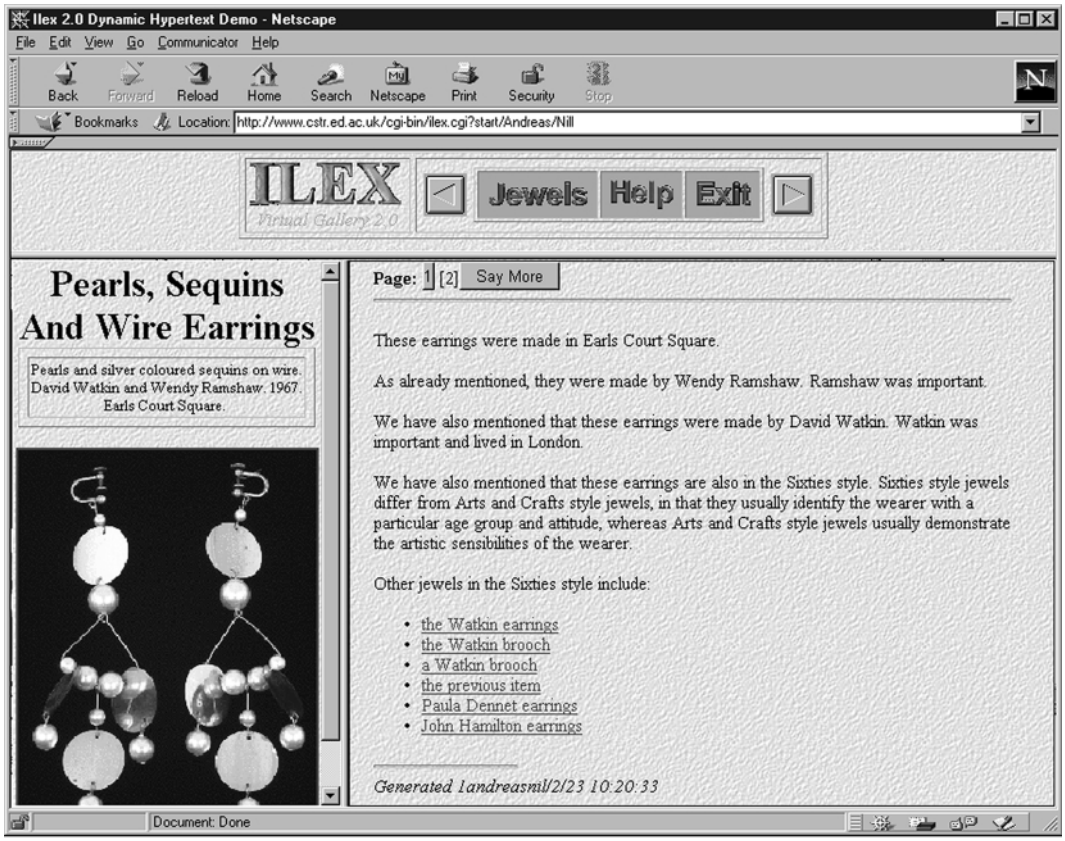


Figure 8 Text presentation to a user with an interest in styles (from ILEX, 1999; reprinted with permission)

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

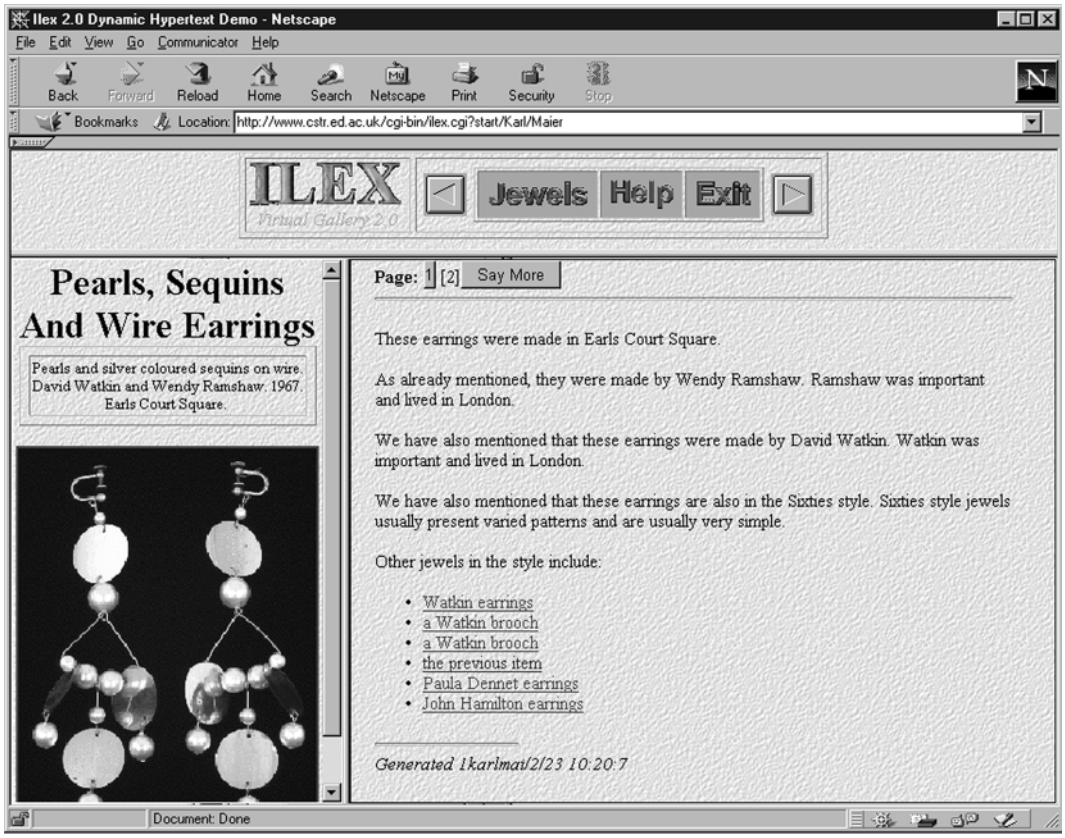


Figure 9 Text presentation to a visitor with no specific interests from ILEX, 1999; reprinted with permission)

dynamically produce hypertext pages that are adapted to the current user's goal and knowledge as well as his or her professional role.

Limited natural-language generation is frequently employed for theory-driven presentation that caters to users following more general didactic, rhetorical or psychological principles (see Section 5.1.1). Natural-language generation also seems to be a promising complement to stretchtext. Hypermedia pages with local stretchtext need to be linguistically well formed in all possible combinations of stretchtext expansion. Smoothing of the expansion results by a global natural-language generator would be a feasible way to achieve this.

5.2 Adaptation of presentation and modality

Adaptations of the hypermedia presentation are adaptations where the information content of the hypermedia objects ideally stays the same (in contrast to the changes discussed in the previous section) while the format and layout of the objects change. A special type of presentation change is the change of modality, for example from images to text, from text to audio, or from video to still images. In existing applications we often find that presentation adaptation is used in conjunction with content and structure adaptation.

Adaptation in hypermedia systems concerning multimedia presentations are often based on the user's preferences. These preferences can be communicated to the system at the beginning of the interaction in a short interview or on a machine-readable data carrier (see Section 3.1.1). The system can, however, also monitor which modalities the user selected in the past and under what circumstances, and present these modalities with preference in the future (Joerding *et al.*, 1998).

Based on user preferences and the users' interaction, some intelligent multimedia systems generate multimodal explanations in technical domains. In COMET (Feiner & McKeown, 1993), for instance, a combination of text and three-dimensional graphics is produced to support the maintenance and

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

repair of military radio receiver-transmitters. The system knows the user's actual status and the integrated text generator can select context-dependent verbs like "install", "reinstall" and "return". In problem cases, detailed information, including additional text or other multimodal data, will be presented to the user. One method often used is the so-called adaptive focusing in the context of generating multimodal explanations: cross-references between different media are used to support users in their current situation (Goodman, 1993; Graf, 1998; Wahlster *et al.*, 1992).

The AVANTI system (Fink *et al.*, 1998) bases the selection of different modalities on the user's physical abilities. In Figure 10, AVANTI presents a map of the city of Siena, Italy which includes all interesting places like museums, churches, restaurants and hotels. The "same" page for blind people is shown in Figure 11. Here the modality "image" has been changed to "text". (In addition, content adaptation also has taken place: while the visual map of Siena gives an overview of the different locations and allows users to figure out how to get there, the "map" for the visually impaired contains textual descriptions of all important locations relative to the "Piazza del Campo", the central square and reference point in Siena).

A second group of modality adaptations in AVANTI is based on system performance: videos and large images are automatically replaced by a link and the estimated transmission time if the latter exceeds a certain threshold value. This value is configurable by the user.

5.3 Adaptation of structure

While content adaptation changes the information that is presented in hypermedia pages, and adaptation of presentation and modality changes the way in which information is conveyed to the user, adaptation of structure refers to changes in the way in which the link structure of hypermedia documents or its presentation to users is changed. In the following two sections, we will discuss the major presentation techniques for structure adaptation that have been used in systems developed so far. Some of them have already been subject to initial empirical evaluation. In Section 5.3.2 we will discuss the main purposes for which structure adaptation has been used so far in personalised hypermedia systems.

5.3.1 Techniques for structure adaptation

When discussing the techniques for structure adaptation that have been employed to date, we have to distinguish whether or not a link anchor is part of a context, like in the case of a word in a sentence or an element in a pattern. (A link anchor thereby is the "visible cue" of the link, i.e. the word, symbol or icon that symbolises the link in the source page). Apparently there is more freedom manipulating non-contextual links than there is manipulating contextual links since the overall context needs to be changed as well when a contextual link is altered.

On an implementation level, all of the described techniques for structure adaptation can be realised with the methods described in Section 5.1.2.3.

5.3.1.1 "Collateral" structure adaptation. Content adaptations can implicitly contain link adaptations if fragment variants (see Section 5.1.2.3) contain links which are not presented to the user in one of the variants. Such "collateral" structure adaptation occurred, for example, in Boyle and Encarnacion (1994) and Kobsa *et al.* (1994). These adaptations are only (desired) side effects of content adaptation. In this section we rather focus on explicit adaptive link presentation where links are presented in different forms to indicate their assumed relevance or significance according to user characteristics stored in the user model.

5.3.1.2 Adaptive link sorting. The sorting of link lists is a simple technique that can be used for non-contextual links only. It is primarily employed for ranking the target web pages based on their relevance to users' interests and goals, and their appropriateness for the user based on the user's background knowledge. In Mathé and Chen (1994), user input about, for example, his or her background, task, search goal, and interest is taken into account to generate a link list ordered by relevance to the given input. The user can rate the relevance of the presented links, which in turn is

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press



Figure 10 AVANTI Siena map in graphical form

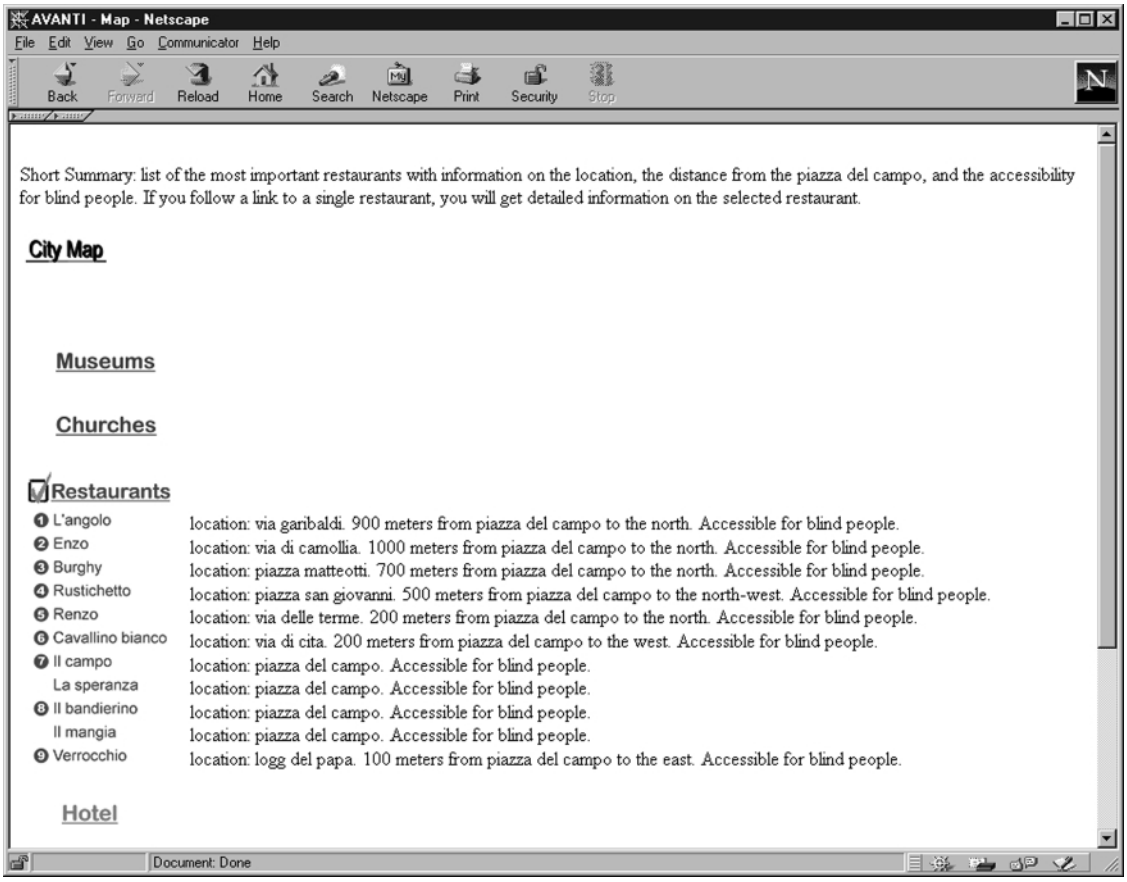


Figure 11 AVANTI Siena map for blind people (visual representation of auditory output)

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

used by the system to update the user model. Similar relevance feedback is used in information-retrieval systems with ranked output, e.g. by all major search engines.

Link lists in HYPERFLEX (Kaplan *et al.*, 1993) indicate the relevance of the links to the current page and the specified goal in descending order. User feedback can be given by sorting the links to explicate a user's preferred ordering. Users' rejection and modification of the system's ranking is used to refine the user model in order to gradually improve the search results (see Section 2.2.1). The authors demonstrate empirically that navigation suggestions through sorting can significantly accelerate users' search for information and reduce their errors.

Adaptive link sorting is also frequently used for presenting ranked lists of recommended items, such as movies (Alspector *et al.*, 1997) or technical equipment (Ardissono & Goy, 1999, 2000b; Ardissono *et al.*, 1999). Another possible use is link sorting based on frequency of use, e.g. in personalised views and spaces (see Section 5.3.2.3). Caution should be exercised, however, since automatic item sorting in menus based on usage frequency has been found to potentially confuse users (Mitchell & Shneiderman, 1988).

5.3.1.3 Adaptive link annotation. Link annotation can be used for both contextual and non-contextual links. Non-adaptive link annotation is well known from all major web browsers: links that have already been visited change their colours. Adaptive hypermedia systems use different colours and symbol codes to annotate links in a personalised manner. For instance, the adaptive hypermedia systems ITEM-IP (Brusilovsky & Pesin, 1994, 1995) and AHA (de Bra & Calvi, 1998) use different backgrounds, font colours and font styles, respectively, to mark link anchors that are recommended by the system, and those that are recommended against. In the educational hypermedia system ELM-ART II (Weber & Specht, 1997), coloured bullets annotate links to chapters or concepts based on the users' current knowledge level. There are up to six different annotations with the following meanings: *suggested*, *ready*, *inferred*, *known*, *done* and *not-ready*. Hypadapter (Hohl *et al.*, 1996) uses different font sizes to annotate links as very relevant, relevant, and less relevant for the current user. Syskill and Webert (Pazzani & Billsus, 1997) uses "thumb up" and "thumb down" icons to denote the ratings of the user and the suggestions of the system. WebWatcher (Joachims *et al.*, 1997) and AVANTI (Fink *et al.*, 1998) annotate links on a page in which the user can be assumed to be interested, based on an explication of his or her interests or a categorization of the previously visited pages. In many systems, annotations are customisable, i.e. users can change the colours and icons.

5.3.1.4 Adaptive link hiding and "unhiding". Adaptive link hiding removes the "visible cue" of a link in such a way that the link anchor looks like normal text or a normal icon. The functionality, however, is still present (and can be discovered through a different cursor shape or through status line information when hovering the cursor over the link). The idea is to visually reduce the hyperspace to support users' navigation, and to guide users to those pages which the system assumes to be the currently most relevant ones or that are probably comprehensible to the user given his or her presumed level of knowledge. The system AHA (de Bra & Calvi, 1998) "hides" links to pages for which the user is not deemed to be ready yet, and later "unhides" them after the user has visited prerequisite pages.[18]

5.3.1.5 Adaptive link disabling and enabling. While link hiding removes all visual indicators of a link, link disabling removes the functionality of a link but leaves its visual appearance nearly untouched (i.e. the link anchor still looks like a link, but nothing happens when one clicks on it). Since such behaviour of a link considerably violates the principle of expectation conformance in human-computer interaction, link disabling and enabling is currently used together with link hiding and unhiding only.

Even this is not unproblematic, though. De Bra and Calvi (1998) report that users of a previous version of their system that used link disabling in combination with link hiding were still dissatisfied

[18] However, the fact that these "forbidden" links are written in a regular underlined font distinguishes them both from normal text and from recommended links (which are set in the regular font and in a blue underlined font respectively).

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

since they could guess the presence of many links due to the fact that these frequently occurred in bulleted lists. Specht and Kobsa (1999) found that learners of a hypermedia course that used "undiscoverable" link disabling and hiding, and revealed them only after students had mastered prerequisite material, learned more slowly (sometimes significantly) and issued significantly more page requests than students in experimental conditions where learning recommendations were presented differently. A likely explanation is that learners had constantly to look for new hyperlinks that possibly appeared on some pages after learners had viewed their prerequisite pages.

On a more positive note, Specht and Kobsa (1999) also report that the only one of four tested presentation methods for page recommendations that could be shown to be significantly beneficial for learners with low prior knowledge was one where non-recommended links were disabled and also marked as such.

5.3.1.6 Adaptive link removal/addition. Adaptive link removal deletes the link anchors completely and is a technique that can only be applied to non-contextual links (unless the whole context is removed as well, as is occasionally done in de Bra & Calvi, 1998). The underlying motivation is to reduce the hypermedia space by removing links to non-relevant information. Examples include the removal of links to pages which are not yet appropriate for a learner (ISIS-Tutor; Brusilovsky & Pesin, 1994), removing items in a product listing that are probably of no interest to the user (Raskutti *et al.*, 1997), and removing links to irrelevant subtasks (Vassileva, 1996).

According to a study with ISIS-Tutor, adaptive link removal seems to be a very effective way to support users' navigation by reducing the number of navigation steps to achieve a certain goal and by reducing the user's cognitive overload (Brusilovsky & Pesin, 1994). In contrast, de Bra and Calvi (1998) report that their users did not like link removal since "they cannot see what lies ahead (but is still unreachable)". The above-mentioned findings of Specht and Kobsa (1999) concerning link disabling plus hiding probably also affect adaptive link removal. Moreover, if a stable listing of links is used frequently, removal of individual links would also violate the constancy principle of human-computer interaction and should therefore be used with caution.

Another usage of link addition has been found in the mobile web-based museum guide HIPS (Oppermann & Specht, 1999, 2000). This system automatically introduces links to nearby paintings (see Figure 12) and also links to more distant paintings if it inferred a user interest in their topic, painter or time period based on the topics, painters and time periods of paintings that the user visited before.

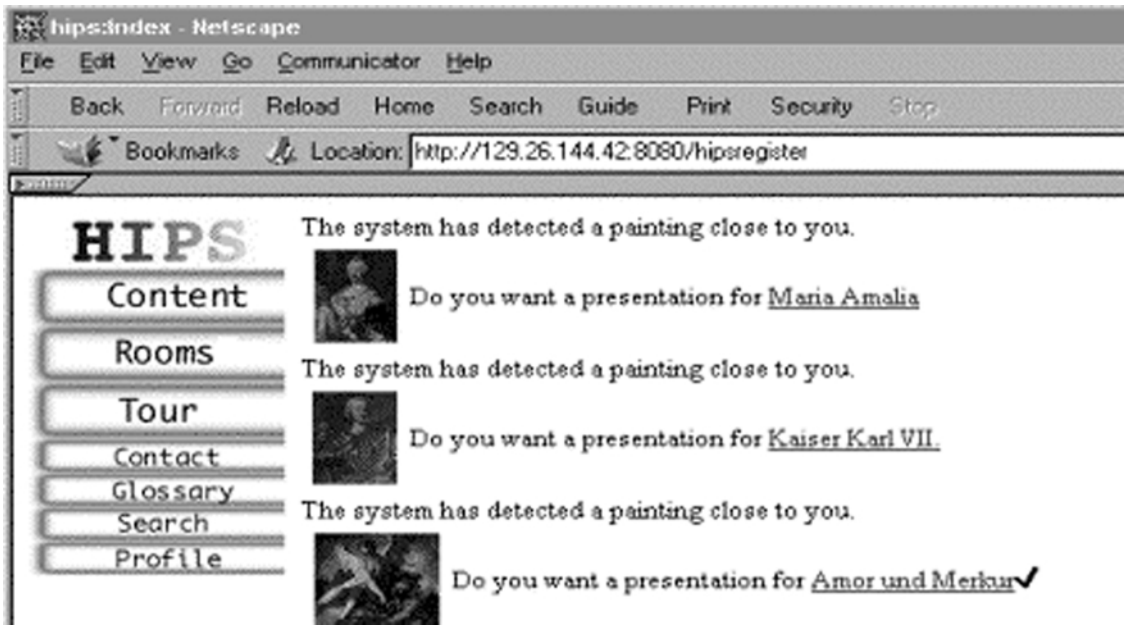


Figure 12 Recommendation of nearby paintings in HIPS

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

5.3.2 Some personalisation functions of structure adaptation

While structure adaptation in adaptive hypermedia systems can in principle again be used for arbitrary purposes, a few major personalisation functions have evolved in the last few years. These include adaptive recommendations, adaptive guidance and orientation, and personal views and spaces, all of which are useful for customer relationship management. In all three cases, personalisation can be based on user, usage and environment data.

5.3.2.1 Adaptive recommendations. Adaptive recommendations can be presented to the user with any of the techniques discussed in Section 5.3.1. The most frequently used types of recommendation are given below.

Recommendations concerning products. Lists of links to "products" and services are filtered or ranked based on user data, and presented to the user (see e.g. Alspector *et al.*, 1997; Ardissono & Goy, 1999, 2000b; Ardissono *et al.*, 1999; Amazon.com, the Music Advisor of musicmaker.com, bol.com, moviefinder.com, the Album Advisor™ of CDNOW, and the recommendation of paintings in HIPS.)

Recommendation concerning information. Lists of links to documents or other pieces of information are ranked based on user and usage data (see e.g. Mathé & Chen, 1994; Pazzani & Billsus, 1997).

Navigation recommendations. Links to hypermedia pages (usually at the same site) are filtered or ranked based on user, usage and environment data (see e.g. de Bra & Calvi, 1998; Joachims *et al.*, 1997; Kaplan *et al.*, 1993; Weber & Specht, 1997). Filtering may also mean that "inappropriate" links are removed and "recommendations" thereby enforced.

5.3.2.2 Adaptive orientation and guidance. One of the main characteristics of the hypertext and hypermedia approach is its interaction style, namely free navigation through the hyperspace by arbitrarily following any links. Unfortunately, the freedom of navigation often leaves users uncertain as to how to reach what they are looking for or how to achieve what they want to do.

Overview maps have become a standard means for helping users figure out the contents and the structure of a website. Personalised overview maps (e.g. Rössel, 1998) can, among other things, mark those pages that users visited or bookmarked in the past.

Guided site tours have become another popular way of familiarising first-time users with the basic offerings of a website. Personalised guided tours can take user data into account (which users usually have to explicate in a questionnaire) and modify the tour so that it caters better to users' presumable interests (see e.g. Rössel, 1998). Users just have to press a "next" button and will be taken to the next overview page that the system deems to be interesting for them.

Personalised next buttons are sometimes also used in educational hypermedia systems (e.g. Brusilovsky & Pesin, 1994; Gutiérrez *et al.*, 1996; Weber & Specht, 1997). This is a very flexible method for the presentation of learning material because the destination node of the next button is not directly connected to the current node but can be dynamically computed at runtime, taking even the very last actions of the user into account. Direct guidance might, however, give users the feeling of scaffolding since there is no other support for them other than following the system's suggestions. Nevertheless, linear navigation via direct guidance can be an effective way for novices to achieve the learning goal in educational hypermedia systems (cf. Reed *et al.*, 1995).

Less restrictive guidance can be given with navigation recommendations as described in Section 5.3.2.2. Brusilovsky (1996) makes finer distinctions between global and local guidance and global and local orientation. The boundaries are, however, not clearly cut and most systems fall into more than one category.

5.3.2.3 Personal views and spaces. The bookmarking facilities that are integrated into most current web browsers provide users with personalised access to web resources. Hierarchical bookmark lists give them the possibility to organize their own view of the available information. In traditional browsers, responsibility for creating and maintaining the bookmark lists lies exclusively on the side of

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

the user and there is no support of the system available. Personalisation techniques can support users in creating personalised views and more comprehensive "information spaces", based mostly on usage data.

The system of Debevc *et al.* (1997) collects the URLs that were recently most frequently used into an "adaptive short list" (from where they can be copied into the bookmark list, if so desired). The adaptive bookmarking service WebTagger (Keller *et al.*, 1997) enables individual users and user groups to store, access and rate URLs, including automatic categorisation of URLs based on initial manual categorisation. Users' feedback concerning the usefulness of URLs is taken into account to modify the ranking of the presented link. The AVANTI system (Fink *et al.*, 1998) automatically generates shortcuts for frequently followed links. BASAR (Thomas & Fischer, 1996) informs users about obsolete and never or rarely used links, repeats searches at pre-specified intervals and compares the results with the bookmark list.

Many portals and other major websites recently started encouraging users to additionally create personalised spaces on their sites that allow them, for example, to

- view histories of their past actions (e.g. purchases and reservations),
- set markers (e.g. for books to buy in the future),
- define shortcuts to site-specific resources they frequently access,
- specify information they want to have forwarded to them automatically (e.g. quotes for certain stocks, news from certain categories) and
- save documents and news in a personal repository.

While there is currently not very much system assistance offered for creating such personalised spaces, user and usage modelling techniques as described in Sections 3.1 and 3.2 could be used to realise system-supported adaptability to users' needs.

6 Conclusions and prospects

This article reviewed methods for the construction of personalised hypermedia systems. The majority of these methods have been developed in the context of research projects, and their implementation and evaluation has been typically restricted to small prototypes. Any of these methods can be applied in customer relationship applications, and there are now commercial tools available to realise individual methods (see Fink & Kobsa, 2000).

In the following we would like to characterise the main objectives of personalisation in what will be the likely main application areas of personalised customer relationship applications in a wide sense. Thereafter we make recommendations with respect to several methodological issues in personalised customer-relationship applications, specifically as far as the acquisition and representation of assumptions about the user is concerned.

6.1 Application areas and their requirements

One major application area for adaptive hypermedia-based customer-relationship software will be public websites. It is generally not difficult to drive traffic to a website through appropriate advertising of the site and through measures that improve the ranking of the site by the major search engines. What is more difficult is to keep visitors at the site, to turn them into customers and to make them come back. Different types of support must be given to first-time visitors, returning visitors, and infrequent and frequent users of the website.

Personalisation is not necessarily useful for all kinds of products sold over the web. Hanson (1999) distinguishes several criteria (e.g. the variation of customers' lifetime values, the variations of needs among customers and the nature of the product attributes) that help determine whether and what kind of personalisation should be implemented. Personalisation performed by the computer also does not

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

necessarily put human sales assistants and advisors out of the loop. Åberg and Shahmehri (1999) describe a web shop where customers can ask for human assistance when needed. First evaluations with a restricted mock-up system yielded very positive user reactions.

A second application area is information kiosks, such as at fairs, exhibitions, showrooms, and public places. Here we have a similar situation in that the system needs to support "walk up and use" by first-time users or infrequent users. In addition to user and usage modelling, adaptation to the locale (noise, reduced privacy due to proximity of other people) is specifically important. As soon as means become available to transfer user data between kiosk systems (see Section 3.1.1), longer-term user modelling will also be possible. Rather than developing two different applications for web use and kiosk use, an integrated approach should be pursued where the adaptation of content and presentation format is also based on the different hardware and software environments and the different locales.

A third challenging application area in the near future comes from the dramatically increasing number of web-capable appliances and web-capable mobile devices that can be utilised to establish ubiquitous access opportunities for customers, ranging from Internet kiosks to car-mounted displays and portable digital assistants integrated with telephone functionality. These developments create new challenges and opportunities for information delivery: on the one hand, one has to serve an increasing variety of output devices with different capabilities and connection speeds in a larger variety of settings (which, for example, demand different presentation modalities such as video and touch input for information kiosks and audio-based hands-free interaction in driving situations). On the other hand, mobile devices linked with position information and embedded devices that have access to data from their physical environments can provide the basis for adapting information to the current needs of individual customers, such as providing a list of dealers or maintenance facilities closest to their current location.

A fourth application area, finally, comes from the requirements of "universal access", "design for all" and "user interfaces for all" with respect to computer systems (Shneiderman, 2000; Stephanidis, 2001). For a few years, there have been increasing public awareness and pressure both from legislation (e.g. the Americans with Disabilities Act), standards (e.g. ISO norms) and guidelines (e.g. Jacobs *et al.*, 2001; White *et al.*, 2001) that software systems should be designed in such a way that they pose no access barriers to people with special needs, such as users with disabilities, elderly users and users with different cultural backgrounds. While this requirement may in principle be met by developing specialised access systems (such as screen readers or web browsers for the blind) and customized versions (such as localized software) for each of these user subgroups, it seems to be economically more feasible in the long run to develop more generic applications that can be adapted or can adapt themselves to the needs of these different special audiences.

6.2 Acquisition and representation

As users visit a site, the paradox of the active user and the task orientation of users should be taken into account. This means that the user should be exposed immediately to content, and not to lengthy registration procedures or initial interviews. If a user interview method (see Section 3.1.1) is chosen, it should be restricted to 2 or 3 core questions that are embedded on the main page. As a result, adaptivity has to be primarily based on user data from external sources (such as marketing data) and stereotypes based thereon (see Section 3.1.4), and on usage data (i.e. observation of user behaviour). For example, two questions about a currently owned vehicle and postal code (in the web context) may provide enough input to select a very rough stereotype culled from demographic marketing data that can be used as an initial model until further usage data become available.

Other active adaptation methods can be used by offering adaptivity to the user as an option. For example, users can be offered a short quiz and promised better recommendations for products and services in return. Users should always feel in control and should have the option of correcting, undoing or simply ignoring adaptive modifications. For instance, they should be able to re-sort adaptively generated result lists (and this should update the system's user model, see Sections 2.2.1.1 and 2.2.1.3). Personalised systems should allow the user to view all items (with recommendatory

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

annotations, see Section 5.3.1.3) and not just the recommended subset (link removal, see Section 5.3.1.6).[19]

With respect to passive acquisition, we do recommend the logging of user interaction and navigation behaviour on the page level. Standard tools exist to collect these data and they are needed anyway for the iterative design of websites and kiosks (e.g. to detect usability problems). Analysis beyond frequency distributions and especially inferences drawn upon statistical analyses for adaptivity purposes are an open field for active research. Especially interesting is the combination of results from a usage data analysis with user data from other customer-relationship software components like marketing databases and customer sales records.

At the current state of the art, we do not recommend a universal logging of usage data on the micro-interaction level, such as the tracking of mouse movements within applets, unless the purpose of the logging has already been specified (e.g. for determining users' interest in page segments, like in the systems of Sakagami and Kamba, 1997; Sakagami *et al.*, 1998; Joerding *et al.*, 1998; Joerding, 1999 described in Section 2.2.1). The amount of data collected is very large, the computation needed to derive recommendations for adaptations is extensive, and confidence in the suitability of these adaptations is likely to be relatively low. However, it seems promising to experiment with such data in smaller, laboratory contexts to drive the development of new methods in this area.

In the not-so-distant future one can imagine that the same technology that will be used to identify and authorise users based on physiological characteristics (e.g. voice, fingerprint and retina analysis in keyless entry systems) can also be utilised to access things like information kiosks in public areas. If this is the case and if a central or mobile repository for user data exists (see Section 3.1.1), the adaptive system will have immediate access to stored user data and can provide rich content adaptation. Within the context of electronic commerce, this can be used, for example, to prominently display news and product items closely related to the product(s) a user already owns with the goal of creating up-selling and cross-selling opportunities, to provide safety-related information such as product-specific recalls or upcoming product maintenance needs, or to ask the user to provide feedback on recent new purchases or services.

In all logging activities, privacy concerns of users have to be taken very seriously and have to be addressed proactively even if the actual risks of privacy invasion are low. Key aspects of a privacy policy are to tell users what is done and to demonstrate the added value of providing personal information, as well as to use collected data only for the intended purposes in order to establish trust and comply with privacy laws in many countries (Kobsa, 2001). If possible, users should also be given the opportunity to inspect their user model (Bull, 1997; Cook & Kay, 1994; Cotter & Smyth, 2000; Czarnowski & Kay, 2000) and to opt out of logging and thus personalisation. Security and privacy infrastructures should be used that both secure user data from possible attacks and allow users to remain highly anonymous with respect to the personalised customer relationship application while still enjoying personalisation (see Schreck, 2001).

With regard to the debate about the use of implicit and explicit representations of user models that was discussed in Section 4.4 we suggest a dual approach. "Implicit models" that produce specific adaptations, such as product recommendations based on statistical similarity measures, are strong, effective and efficient methods that are being used in commercial settings today. In contrast, many of the explicit representation and reasoning methods are just converting from research prototypes to products, and one has to carefully test individual solutions to determine whether the robustness and speed will hold up in large-scale commercial use. We expect, however, that the further rapid increase in computing power especially on the server side will allow for more knowledge-intensive processing in the future. Explicit user models and usage models have advantages that make it desirable to continue research and development in this area. The two main points in favour of them are given below.

Explanation. Pure statistical methods that find "nearest neighbours" in n-dimensional spaces are capable of producing recommendations, but users have no opportunity to inspect the reasoning that

[19] See Smyth and Cotter (2000) for an example in the area of personalised television show recommendations.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

resulted in these recommendations (and thus have no chance to correct reasoning with which they do not concur). The reasoning apparatus associated with explicit models can be used for also implementing explanation components that upon a user's request provide a rationale as to why a particular item or product has been recommended. This will increase the trust users have in those recommendations (given a satisfactory level of recommendation and explanation quality).

Integration with other customer-relationship applications and data. Users' "implicit profiles" cannot be directly transferred to other systems. In current customer-relationship components, many data exist in explicit form (e.g. demographic marketing data). Exploiting these data will be easier with explicit models. Users' specification of their interests and knowledge also comes in an explicit form and often not on the low level that is used to compute nearest neighbours (e.g. users indicate the topics in which they are interested and do not list the significant words that are contained in documents in which they are interested).

Since implicit and explicit user models each have their own merits, it also seems advantageous to integrate both into hybrid systems. Research in this direction has, however, not started until recently (see e.g. Billsus & Pazzani, 1999, 2000; Pohl & Nick, 1999).

References

Åberg, J and Shahmehri, N, 1999 "Web assistants: towards an intelligent and personal web shop" *Proceedings of the 2nd Workshop on Adaptive Systems and User Modeling on the WWW*, WWW-8, Toronto, Canada and UM99, Banff, Canada, 5–12. http://www.contrib.andrew.cmu.edu/~plb/WWWUM99_workshop/aberg/aberg.html

Abrams, C, Bernstein, M, de Sisto, R, Drobik, A and Herschel, G, 1999, "E-business: the business tsunami" *Proceedings of Gartner Group Symposium/ITxpo*, Cannes, France 1–18.

Akoulchina, I and Ganascia, J-G, 1997, "SATELIT-Agent: an adaptive interface based on learning interface agents technology" in, A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag 21–32. http://www.um.org/um_97/gz/AkoulchinaIps.gz

Albrecht, DW, Zukerman, I, Nicholson, AE and Bud, A, 1997, "Towards a Bayesian model for keyhole plan recognition in large domains" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag, 365–376. http://www.um.org/um_97/gz/AlbrechtDps.gz

Allen, C, Kania, D and Yaeckel, B, 1998, *Internet World Guide to One-To-One Web Marketing* John Wiley & Sons.

Allen, JF and Perrault, CR, 1980, "Analyzing intentions in utterances" *Artificial Intelligence* **15**(3) 143–178.

Alspector, J, Kolcz, A and Karunanithi, N, 1997, "Feature-based and clique-based user models for movie selection: a comparative study" *User Modeling and User-Adapted Interaction* **7**(4) 279–304.

Ambrosini, L, Cirillo, V and Micarelli, A, 1997, "A hybrid architecture for user-adapted information filtering on the World Wide Web" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag., 59–61.
http://www.um.org/um_97/gz/AmbrosiniLps.gz

Appian, 2000, Appian. http://www.appiancorp.com

Apple Computer, Inc. WebObjects 4 Developer Documentation. http://developer.apple.com/techpubs/webobjects/

Ardissono, L and Goy, A, 1999, "Tailoring the interaction with users in electronic shops" in J Kay (ed) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag, 35–44. http://www.cs.usask.ca/UM99/Proc/ardissono.pdf

Ardissono, L and Goy, A, 2000a, "Dynamic generation of adaptive web catalogs" in: P Brusilivsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer 5–16.

Ardissono, L and Goy, A, 2000b, "Tailoring the interaction with users in web stores" *User Modeling and User-Adapted Interaction* **10**(4) 251–303.

Ardissono, L, Goy, A, Meo, R and Petrone, G, 1999, "A configurable system for the construction of adaptive virtual stores" *World Wide Web* **2**(3) 143–159.

Armstrong, R, Freitag, D, Joachims, T and Mitchell, T, 1995, "Webwatcher: a learning apprentice for the World Wide Web" *Proceedings of the 1995 AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments* Stanford University, Stanford, CA. http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-6/web-agent/www/webagent-plus/webagent-plus.html

Art Technology Group, ATG Dynamo Personalization Server. http://www.atg.com/products/dps/dps_main.html

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Balabanovic, M and Shoham, Y, 1997, "Fab: content-based, collaborative recommendation". *Communications of the ACM* **40**(3) 66–72.

Balabanovic, M, 1997, "An adaptive web page recommendation service" in *Proceedings of the 1st International Conference on Autonomous Agents*, Marina del Rey, CA, 378–385.

Ballim, A and Wilks, Y, 1991, "Beliefs, stereotypes and dynamic agent modeling" *User Modeling and User-Adapted Interaction* **1**(1) 33–65.

Bares, WH and Lester, JC, 1997, "Cinematographic user models for automated realtime camera control in dynamic 3D environments" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference*, 215–226 Springer-Verlag. `http://www.um.org/um_97/gz/BaresWHps.gz`

Bauer, M, 1996, "A Dempster-Shafer approach to modeling agent preferences for plan recognition" *User Modeling and User-Adapted Interaction* **5**(3–4) 317–348.

Beaumont, IH, 1994, "User modelling in the interactive anatomy tutoring system ANATOM-TUTOR" *User Modeling and User-Adapted Interaction*, **4**(1) 21–45.

Billsus, D and Pazzani, M, 1999, "A hybrid user model for news story classification" in J Kay (ed) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag 99–108.

Billsus, D and Pazzani, MJ, 2000, "User modeling for adaptive news access" *User Modeling and User-Adapted Interaction* **10**(2–3) 147–180. `http://www.um.org/um_99/Proc/billsus.pdf`

Boyle, C and Encarnacion, AO, 1994, "Metadoc: an adaptive hypertext reading system" *User Modeling and User-Adapted Interaction* **4**(1) 1–19.

Brajnik, G and Tasso, C, 1994, "A shell for developing non-monotonic user modeling systems" *International Journal of Human-Computer Studies* **40** 31–62.

Breese, J, Heckerman, D and Kadie, C, 1998, "Empirical analysis of predictive algorithms for collaborative filtering" *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-98)* Morgan Kaufmann 43–52.

Broadvision, 2001, BroadVision One-to-One Enterprise. `http://www.broadvision.com`

Brusilovsky, P, Pesin, L and Zyryanov M, 1993, "Towards an adaptive hypermedia component for an intelligent learning environment" in LJ Bass, J Gornostaev and Claus Unger (eds) *Human-Computer Interaction, Third International Conference, EWHCI'93* Springer-Verlag 348–358.

Brusilovsky, P and Pesin, L, 1994, "ISIS-Tutor: an intelligent learning environment for CDS/ISIS users" *Online Proceedings of CLCE'94*. `http://cs.joensuu.fi/~mtuki/www_clce.270296/Brusilov.html`

Brusilovsky, P and Pesin, L, 1995, "Visual annotation of links in adaptive hypermedia" *Human Factors in Computing Systems Conference CHI'95, Conference Companion: Short Papers*. `http://www.acm.org/sigchi/chi95/Electronic/documnts/shortppr/plb_bdy.htm`

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

Brusilovsky, P, Kobsa, A and Vassileva, J (eds), 1998, *Adaptive Hypertext and Hypermedia* Kluwer Academic Publishers.

Brusilovsky, P, 2001, "Adaptive hypermedia" *User Modeling and User-Adapted Interaction*, **11**(1–2), 87–110.

Bull, S, 2000, "See yourself write: a simple student model to make students think" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference, UM97* Springer 315–326. `http://www.um.org/um_97/gz/BullSps.gz`

Bushey, R, Mauney, JM and Deelman, T, 1999, "The development of behavior-based user models for a computer system" in J Kay (ed) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer. `http://www.cs.usask.ca/UM99/Proc/bushey.pdf`

Carberry, S, 1989, "Plan recognition and its use in understanding dialog" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer 133–162.

Carberry, S, 1990, *Plan Recognition in Natural Language Dialogue* MIT Press.

Carroll, J and Rosson, MB, 1987, "The paradox of the active user" in JM Carroll (ed.) *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction* MIT Press.

Chen, J, Yang, Y and Zhang, H, 2000, "An adaptive web content delivery system" in P Brusilivsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer 284–288.

Chin, DN, 1989, "KNOME: modeling what the user knows in UC" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer Verlag 74–107.

Chin, DN, 1993, "Acquiring user models" *Artificial Intelligence Review* **7** 185–197.

Chittaro, L and R Ranon, 2000, "Adding adaptive features to virtual reality interfaces for e-commerce" in P Brusilivsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer 86–91.

CL-HTTP homepage `http://www.ai.mit.edu/projects/iiip/doc/cl-http/home-page.html`

Cook, R and J Kay, 1994, "The justified user model: a viewable, explained user model" *Fourth International Conference on User Modeling* 145–150.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Cost, S and Salzberg, S, 1993, "A weighted nearest neighbor algorithm for learning with symbolic features" *Machine Learning* **10** 57–78.

Cotter, P and B Smyth, 2000, "WAPing the web: content personalization for WAP-enabled devices" in P Brusilivsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer 98–108.

Cypher, A (ed.), 1993, *Watch What I Do: Programming by Demonstration* MIT Press.

Czarkowski, M and J Kay, 2000, "Bringing scrutability to adaptive hypertext teaching" in G Gauthier, C Frasson and K Van Lehn (eds) *Intelligent Tutoring Systems* Springer 423–433.

Davison, BD and Hirsh, H, 1998, "Predicting sequences of user actions" in *Proceedings of the AAAI/ICML 1998 Workshop on "Predicting the Future: AI Approaches to Time-Series Analysis"*. http://www.cs.rutgers.edu/~davison/pubs/aaai98ws.ps

De Bra, P and Calvi, L, 1998, "AHA! an open adaptive hypermedia architecture" *The New Review of Hypermedia and Multimedia* **4** 115–139.

De Carolis, B, de Rosis, F, Andreoli, C, Cavallo, V and De Cicco, ML, 1998, "The dynamic generation of hypertext presentations of medical guidelines" *The New Review of Hypermedia and Multimedia* **4** 67–88.

De Rosis, F, Pizzutilo, S, Russo, A, Berry, DC and Molina, FJN, 1992, "Modeling the user knowledge by belief networks" *User Modeling and User-Adapted Interaction*, **2**(4) 367–388.

Debevc, M, Meyer, B and Svecko, R, 1997, "An adaptive short list for documents on the World Wide Web" *Proceedings of the 1997 International Conference on Intelligent User Interfaces* 209–211.

Debevc, M, Meyer, B, Donlagic, D and Svecko, R, 1996, "Design and evaluation of an adaptive icon toolbar" *User Modeling and User-Adapted Interaction* **6**(1) 1–21.

Dieterich, H, Malinowski, U, Kühme, T and Schneider-Hufschmidt, M, 1993, "State of the art in adaptive user interfaces" in M Schneider-Hufschmidt, T Kühme and U Malinowski (eds) *Adaptive User Interfaces: Principle and Practice* North Holland.

Duda, R and Hart, P, 1973, *Pattern Classification and Scene Analysis.* John Wiley & Sons.

Dutton, DM and Conroy, GV, 1997, "A review of machine learning" *The Knowledge Engineering Review* **12**(4) 341–367.

Eklund, J and Brusilovsky, P, 1998, "The value of adaptivity in hypermedia learning environments: a short review of empirical evidence" in P Brusilovsky and P De Bra (eds) *Proceedings of Second Adaptive Hypertext and Hypermedia Workshop at the Ninth ACM International Hypertext Conference Hypertext'98* 11–17. http://wwwis.win.tue.nl/ah98/Eklund.html

Encarnação, LM and Stoev, SL, 1999, "An application-independent user support system exploiting action-sequence based user modeling" in J Kay (ed) *UM99 User Modeling: Proceedings of the Seventh International Conference* 245–254. http://www.cs.usask.ca/UM99/Proc/encarnacao.pdf

Feiner, SK and McKeown, KR, 1993, "Automating the generation of coordinated multimedia explanations" in MT Maybury (ed.) *Intelligent Multimedia Interfaces* MIT Press 117–138.

Fink, J and Kobsa, A, 2000, "A review and analysis of commercial user modeling servers for personalization on the World Wide Web" *User Modeling and User-Adapted Interaction* **9**(3–4) special issue on deployed user modeling 209–249.

Fink, J, 2001, *User Modeling Servers: Requirements, Design, and Evaluation* Department of Mathematics and Computer Science, University of Essen, Germany, forthcoming.

Fink, J, Kobsa, A and Jaceniak, I, 1997, "Individualisierung von Benutzerschnittstellen mit Hilfe von Datenchips für Personalisierungsinformation" *GMD-Spiegel* **1**/1997, 16–17. http://ics.uci.edu/~kobsa/papers/1997-GMD-kobsa.ps

Fink, J, Kobsa, A and Nill, A, 1998, "Adaptable and adaptive information provision for all users, including disabled and elderly people" *The New Review of Hypermedia and Multimedia* **4** 163–188.

Fischer, G, Mastaglio, T, Reeves, B and Riemann, J, 1990, "Minimalist explanations in knowledge-based systems" *23rd Annual Hawaii International Conference on System Sciences* 309–317.

Fuller, R and de Graaff, JJ, 1996, "Measuring user motivation from server log file" in *Proceedings of the 'Designing for the Web: Empirical Studies'* Microsoft Usability Group. http://www.microsoft.com/usability/webconf/fuller/fuller.htm

Gates, KF, Lawhead, PB and Wilkins, DE, 1998, "Toward an adaptive WWW: a case study in customized hypermedia" *The New Review of Hypermedia and Multimedia* **4** 89–113.

Gonschorek, M and Herzog, C, 1995, "Using hypertext for an adaptive help system in an intelligent tutoring system" in J Greer (ed.) *Proceedings of the World Conference on Artificial Intelligence in Education*, August 1995 274–281.

Goodman, BA, 1993, "Multimedia explanations for intelligent training systems" in MT Maybury (ed.) *Intelligent Multimedia Interfaces* MIT Press 148–171.

Graf, WH, 1998, "Constraint-based graphical layout of multimodal presentations" in MT Maybury and W Wahlster (eds) *Readings in Intelligent User Interfaces* Morgan Kaufmann 263–283. (Originally in T Catarci,

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

MF Costabile and S Levialdi (eds.) *Proceedings of Advanced Visual Interfaces (AVI'92)* World Scientific Press 365–385).

Gutiérrez, J, Pérez, TA, Usandizaga, I and Lopistéguy, P, 1996, "HyperTutor: adapting hypermedia systems to the user" *Proceedings of the Fifth International Conference on User Modeling, UM-96* 210–212.

GVU, 1998, GVU's 10th WWW User Survey. `http://www.cc.gatech.edu/gvu/user_surveys/survey-1998-10`

Hanson, W, 1999, "Personalization" Chapter 7 of *idem*, *Principles of Internet Marketing* Thomson Publishing.

Harter, SP, 1992, "Psychological relevance and information science" *Journal of the American Society for Information Science*. **43**(9) 602–615.

Henze, N and Nejdl, W, 2000, "Extendible adaptive hypermedia courseware: integrating different courses and web material" in P Brusilovsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer 109–120.

Hof, RD, Green, H and Himmelstein, L, 1998, "Now it's YOUR WEB" *Business Week*, October 5, 68–74.

Hohl, H, Böcker, H-D and Gunzenhäuser, R, 1996, "HYPADAPTER: an adaptive hypertext system for exploratory learning and programming" *User Modeling and User-Adapted Interaction*, **6**(2–3) 131–155.

Holynski, M, 1988, "User-adaptive computer graphics" *International Journal of Man-Machine Studies* **29** 539–548.

Höök, K, Karlgren, J, Waern, A, Dahlbäck, N, Jansson, C, Karlgren, K and Lemaire, B, 1996, "A glass box approach to adaptive hypermedia" *User Modeling and User-Adapted Interaction* **6**(2–3) 157–184.

Hoppe, HU and Plötzner, R, 1991, "Inductive knowledge acquisition for a UNIX coach" in *Mental Models and Human-Computer Interaction 2*. Elsevier Science Publishers.

Horvitz, E, 1997, "Agents with beliefs: reflections on Bayesian methods for user modeling" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag, 441–442. `http://www.um.org/um_97/gz/HorvitzEps.gz`

Horvitz, E, Breese, J, Heckerman, D, Hovel, D and Rommelse, K, 1998, "The Lumière Project: Bayesian user modeling for inferring the goals and needs of software users" *Fourteenth Conference on Uncertainty in Artificial Intelligence* 256–265.

Hothi, J and Hall, W, 1998, "An evaluation of adapted hypermedia techniques using static user modelling" *Proceedings of the 2nd Workshop on Adaptive Hypertext and Hypermedia at HYPERTEXT'98* 45–55. `http://wwwis.win.tue.nl/ah98/Hothi/Hothi.html`

Hothi, J, Hall, W and Sly, T, 2000, "A study comparing the use of shaded text and adaptive navigational support in adaptive hypermedia" in P Brusilivsky, O Stock and C Strappavara (eds) *Adaptive Hypermedia and Adaptive Web-Based Systems* Springer, 335–342.

Ikeda, M and Mizoguchi, R, 1994, "Fits: a framework for ITS – a computational model of tutoring" *Journal of Artificial Intelligence in Education* **5**(3) 319–348.

ILEX Virtual Gallerie 2.0. `http://www.cstr.ed.ac.uk/cgi-bin/ilex.cgi`

Jacobs, I, Gunderson, J and Hansen, E, 2001, "User Agent Accessibility Guidelines. W3C Working Draft 26 January 2001". http://www.w3.org/WAI/UA/WD-UAAG10–20010126/

Jameson, A, 1989, "But what will the listener think? belief ascription and image maintenance in dialog" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer-Verlag.

Jameson, A, 1992, "Generalizing the double-stereotype approach: a psychological perspective" in E André, R Cohen, W Graf, B Kass, C Paris and W Wahlster (eds) *UM92: Third International Workshop on User Modeling* 69–83.

Jameson, A, 1996, "Numerical uncertainty management in user and student modeling: an overview of systems and issues" *User Modeling and User-Adapted Interaction* **5**(3–4) 193–251.

Jameson, A, 2001, "Modeling both the context and the user" *Personal Technologies* **5**(1), 29–33.

Jameson, A, Schäfer, R, Simons, J and Weis, T, 1995, "Adaptive provision of evaluation-oriented information: tasks and techniques" *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* Morgan Kaufmann, 1886–1893.

Jennings, A and Higuchi, H, 1993, "A user model neural network for a personal news service" *User Modeling and User-Adapted Interaction* **3**(1) 1–25.

Joachims, T, Freitag, D and Mitchell, T, 1997, "Webwatcher: a tour guide for the World Wide Web" in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* Morgan Kaufmann Publishers.

Joerding, T, 1999, "A temporary user modeling approach for adaptive shopping on the web" *Proceedings of the 2nd Workshop on Adaptive Systems and User Modeling on the WWW, WWW-8, Toronto, Canada and UM99, Banff, Canada*. `http://www.contrib.andrew.cmu.edu/~plb/WWWUM99_workshop/joerding/joerding.html`

Joerding, T, Michel, S and Popella, M, 1998, "Tellim – ein System für adaptive Multimediale Produktpräsentationen im World Wide Web" in UJ Timm and M Rössel (eds) *ABIS-98 – 6. Workshop Adaptivität und*

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Benutzer-modellierung in interaktiven Softwaresystemen, 29–40. FORWISS http://www-mmt.inf.tudresden.de/joerding/abis98/abis98.html

Kaplan, C, Fenwick, J and Chen, J, 1993, "Adaptive hypertext navigation based on user goals and context" *User Modeling and User-Adapted Interaction* **3**(3) 193–220.

Kass, R, 1991, "Building a user model implicitly from a cooperative advisory dialog" *User Modeling and User-Adapted Interaction* **1**(3) 203–258.

Kay, J and Kummerfeld RJ, 1994, "An individualised course for the C programming language" *Second International World Wide Web Conference '94: Mosaic and the Web*, Chicago, IL. http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/Educ/kummerfeld/kummerfeld.html

Keller, RM, Wolfe, SR, Chen, JR, Rabinowitz, JL and Mathe, N, 1997, "A bookmarking service for organizing and sharing URLs" *Proceedings of the Sixth International WWW Conference*. http://ic-www.arc.nasa.gov/ic/projects/aim/papers/www6/paper.html

Kobsa, A (ed.), 1995, Special Issue on User Modeling Shell Systems. *User Modeling and User-Adapted Interaction* **4**(2) and **4**(3).

Kobsa, A and Pohl, W, 1995, "The user modeling shell system BGP-MS" *User Modeling and User-Adapted Interaction* **4**(2) 59–106.

Kobsa, A and Wahlster, W (eds), 1989, *User Models in Dialog Systems* Springer-Verlag.

Kobsa, A, 1989, "A taxanomy of beliefs and goals for user models in dialog systems" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer-Verlag.

Kobsa, A, 1991, "Utilizing knowledge: the components of the SB-ONE knowledge representation workbench" in J Sowa (ed.) *Principles of Semantic Networks: Explorations in the Representation of Knowledge* Morgan Kaufmann, 457–486.

Kobsa, A, Müller, D and Nill, A, 1994, "KN-AHS: an adaptive hypertext client of the user modeling system BGP-MS" in *Proceedings of the Fourth International Conference on User Modeling* 99–105. Reprinted in MT Marbury and W Wahlster (eds), 1998, *Intelligent User Interfaces* Morgan Kaufman, 372–378. http://ics.uci.edu/~kobsa/papers/1994-UM94-kobsa.ps

Kobsa, A, 2001, "Tailoring Privacy to Users' Needs" (Invited keynote) in M Bauer, PJ Gmytrasiewicz and J Vassileva (eds.) User Modeling 2001: 8th International Conference. Berlin-Heidelberg, Springer Verlag, 303–313.

Koike, Y, Kamba, T and Langheinrich, M, 1999, "PIDL: Personalized Information Description Language" W3C. http://www.w3.org/TR/NOTE-PIDL

Kok, AJ, 1991, "A review and synthesis of user modelling in intelligent systems" *The Knowledge Engineering Review* **6**(1) 21–47.

Konstan, JA, 1997, "Added Value in the Digital Age: Bringing One-to-One Marketing to the Internet" White Paper. http://www.netperceptions.com/

Konstan, JA, Miller, BN, Maltz, D, Herlocker, JL, Gordon, LR and Riedl, J, 1997, "GroupLens: applying collaborative filtering to Usenet news" *Communications of the ACM* **40**(3) 77–87.

Kozierok, R and Maes, P, 1993, "A learning interface agent for scheduling meetings" in WD Gray, WE Hefley and D Murray (eds) *Proceedings of the International Workshop on Intelligent User Interfaces* ACM Press 81–88.

Krogsæter, M, Oppermann, R and Thomas, CG, 1994, "A user interface integrating adaptability and adaptivity" in R Oppermann (ed.) *Adaptive User Support: Ergonomic Design of Manually and Automatically Adaptable Software* Lawrence Erlbaum 97–125.

Küpper, D and Kobsa, A, 1999, "User-tailored plan generation" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag 45–54. http://www.cs.usask.ca/UM99/Proc/kuepperd.pdf

Langley, P, 1999, "User modeling in adaptive interfaces" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer 357–370.

Lesh, N and Etzioni, O, 1995, "A sound and fast goal recognizer" in *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* Morgan Kaufmann 1704–1710.

Lesh, N, Rich, C and Sidner, CL, 1999, "Using plan recognition in human-computer collaboration" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag 23–32. http://www.cs.usask.ca/UM99/Proc/lesh.pdf

Lieberman, H, 1995, "Letizia: An agent that assists web browsing" in *Proceedings of the International Joint Conference on Artificial Intelligence, Montreal, Canada* Morgan Kaufmann.

Linton, F and Schaefer, H-P, 2000, "Recommender systems for learning: building user and expert models through long-term observation of application use" *User Modeling and User-Adapted Interaction* **10**(2–3) 181–208.

Linton, F, Joy, D and Schaefer, H-P, 1999, "Building user and expert models by long-term observation of application usage" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag, 129–138. http://www.cs.usask.ca/UM99/Proc/linton.pdf

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Machado, I, Martins, A and Paiva, A, 1999, "One for all and all in one: a learner modelling server in a multi-agent platform" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag 211–221. http://www.cs.usask.ca/UM99/Proc/paiva.pdf

Maes, P, 1994, "Agents that reduce work and information overload" *Communications of the ACM* **37**(7) 31–40.

Mathé, N and Chen, J, 1994, "A user-centered approach to adaptive hypertext based on an information relevance model" in *Proceedings of the Fourth International Conference on User Modeling* 107–114. http://ic-www.arc.nasa.gov/ic/projects/aim/papers/um94/cidum94.html

McCalla, GI and Greer, JE, 1994, *Student Modelling: The Key to Individualized Knowledge-Based Instruction* Springer Verlag.

McCauley, C, Stitt, CL and Segal, M, 1980, "Stereotyping: from prejudice to prediction" *Psychological Bulletin* **87** 195–208.

McCoy, C, 1989, "Highlighting a user model to respond to misconceptions" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer-Verlag.

McGuire, WJ and Padawer-Singer, A, 1976, "Trait salience in the spontaneous self-concept" *Journal of Personality and Social Psychology* **33** 743–754.

McTear, MF, 1993, "User modelling for adaptive computer systems: a survey" *Artificial Intelligence Review* **7**(3–4) 157–184.

Millhouse, C, Brash, C and Chapple, D, 2000, *E-CRM: Personalization Technologies for the Web* Ovum.

Milosavljevic, M and Oberlander, J, 1998, "Dynamic Electronic catalogues: helping users to help themselves" in *Proceedings of Hypertext 98* 123–131.

Milosavljevic, M, 1997, "Augmenting the user's knowledge via comparison" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag. http://www.um.org/um_97/gz/MilosavljevicMps.gz

Mislevy, RJ and Gitomer, DH, 1996, "The role of probability-based inference in an intelligent tutoring system" *User Modeling and User-Adapted Interaction* **5**(3–4) 253–282.

Mitchell, J and Shneiderman, B, 1988, "Dynamic versus static menus: an experimental comparison" *ACM SIGCHI Bulletin* **20**(4) 33–36.

Mitchell, T, Caruana, R, Freitag, D, McDermott, J and Zabowski, D, 1994, "Experience with a learning personal assistant" *Communications of the ACM* **37**(7) 81–91.

Mitchell, TM, Mahadevan, S and Steinberg, L, 1985, "Leap: A learning apprentice for VLSI design" *Proceedings of the 9th International Joint Conference on Artificial Intelligence* 573–580.

Mladenic, D, 1996, *Personal WebWatcher: Implementation and Design* Technical Report, Institut Josef Stefan, Ljubljana, Slovenia.

Moore, JD and Paris, CL, 1992, "Exploiting user feedback to compensate for the unreliability of user models" *User Modeling and User-Adapted Interaction* **2**(4) 331–365.

Morita, M and Shinoda, Y, 1994, "Information filtering based on user behavior analysis and best match text retrieval" *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval* 272–281. http://www.acm.org/pubs/articles/proceedings/ir/188490/p272-morita/p272-morita.pdf

Nielsen, J, 1999, Jakob Nielsen's Alertbox, April 18, 1999. http://www.useit.com/alertbox/990418.html

Nisbett, R and Wilson, T, 1977, "Telling more than we can know: verbal reports on mental processes" *Psychological Review* **84** 231–259.

Nwana, HS, 1991, "User modelling and user adapted interaction in an intelligent tutoring system" *User Modeling and User-Adapted Interaction* **5**(1) 1–32.

Oard, DW, 1997, "The state of the art in text filtering" *User Modeling and User-Adapted Interaction* **7**(3) 141–178.

Oppermann, R (ed.), 1994, *Adaptive User Support: Ergonomic Design of Manually and Automatically Adaptable Software* Lawrence Erlbaum.

Oppermann, R and Specht, M, 1999, "Adaptive information for nomadic activities: a process oriented approach" in U Arend, E Eberleh and K Pitschke (eds.) *Software-Ergonomie '99. Design von Informationswelten* B G Teubner 256–264.

Oppermann, R and Specht, M, 2000, "A context-sensitive nomadic information system as an exhibition guide" *Proceedings of the Handheld and Ubiquitous Computing Second International Symposium, HUC 2000* 127–142.

Orwant, J, 1995, "Heterogeneous learning in the Doppelgänger user modeling system" *User Modeling and User-Adapted Interaction* **4**(2) 107–130.

Paiva, A and Self, J, 1995, "Tagus: a user and learner modeling workbench" *User Modeling and User-Adapted Interaction* **4**(3) 197–226.

Paliouras, G, Karkaletsis, V, Papatheodorou, C and Spyropoulos, CD, 1999, "Exploiting learning techniques for the acquisition of user stereotypes and communities" in J Kay (ed.) *UM99 User Modeling: Proceedings of the*

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Seventh International Conference Springer-Verlag 45–54. http://www.cs.usask.ca/UM99/Proc/karkaletsis.pdf

Paris, C, 1989, "The use of explicit user models in a generation system" in A Kobsa and W Wahlster (eds) *User Models in Dialog Systems* Springer-Verlag.

Pazzani, M and Billsus, D, 1997, "Learning and revising user profiles: the identification of interesting web sites" *Machine Learning* **27** 313–331.

Pazzani, M, Muramatsu, J and Billsus, D, 1996, "Syskill & Webert: identifying interesting web sites" *Proceedings of AAAI'96/IAAI'96* AAAI Press 54–61.

Peppers, D and Rogers, M, 1997, *Enterprise One to One: Tools for Competing in the Interactive Age* Currency Doubleday.

Peppers, D and Rogers, M, 1997, *The One to One Future. Building Relationships One Customer at a Time* Currency Doubleday.

Peter, G and Rösner, D, 1994, "User-model-driven generation of instructions" *User Modeling and User-Adapted Interaction* **3**(4) 289–319.

Pitt, L, Berthon, P and Watson, R, 1996, "From surfer to buyer on the WWW: what marketing managers might want to know" *Journal of General Management* **22**(1) 1–13.

Pohl, W and Nick, A, 1999, "Machine learning and knowledge-based user modeling in the LaboUr approach" in J Kay (ed.) *UM99 User Modeling: Proceedings of the Seventh International Conference* Springer-Verlag 179–188. http://www.cs.usask.ca/UM99/Proc/pohl.pdf

Pohl, W, 1992, *Beispielerkennung für die induktive Generierung von Benutzermakros* GMD.

Pohl, W, 1998, *Logic-Based Representation and Reasoning for User Modeling Shell Systems*. St. Augustin, Germany: infix.

Pohl, W, Kobsa, A and Kutter, O, 1995, "User model acquisition heuristics based on dialogue acts" *International Workshop on the Design of Cooperative Systems* 471–486.

Popp, H and Lödel, D, 1996, "Fuzzy techniques and user modeling in sales assistants" *User Modeling and User-Adapted Interaction* **5**(3–4) 349–370.

Quinlan, JR, 1986, "Induction of decision trees" *Machine Learning* **1**(1) 81–106.

Raskutti, B, Beitz, A and Ward, B, 1997, "A feature-based approach to recommending selections based on past preferences" *User Modeling and User-Adapted Interaction* **7**(3) 179–218.

Reagle, J and Cranor, LF, 1999, "The platform for privacy preferences" *Communications of the ACM* **42**(2) 48–55.

Reed, WM, Oughton, M, Ayersman, DJ, Giessler, SF and Ervin, JR, 1995, "Computer experience and learning style: linear versus nonlinear navigation in a hypermedia environment" in H Maurer (ed.) *Proceedings of Educational Multimedia and Hypermedia (ED-MEDIA) 95*.

Resnick, P and Varian, HR, 1997, "Recommender systems" *Communications of the ACM*, **40**(3) 56–58.

Rich, E, 1979, "User modeling via stereotypes" *Cognitive Science* **3** 329–354.

Rich, E, 1983, "Users are individuals: individualizing user models" *International Journal of Man-Machine Studies* **18** 199–214.

Rössel, M, 1998, "Pragmatische Benutzermodellierung im adaptiven Mulitmedialen Präsentationssystem AMPreS" In UJ Timm and M Rössel (eds) *ABIS-98: 6. Workshop Adaptivität und Benutzermodellierung in interaktiven Softwaresystemen* FORWISS, 29–40.

Rucker, J and Polanco, MJ, 1997, "Siteseer: personalized navigation for the web" *Communications of the ACM* **40**(3) 73–75.

Sakagami, H and Kamba, T, 1997, "Learning personal preferences on online newspaper articles from user behaviors" *Proceedings of the 6th World Wide Web Conference* 291–300.

Sakagami, H, Kamba, T, Sugiura, A and Koseki, Y, 1998, "Effective personalization of push-type systems: visualizing information freshness" *Proceedings of the 7th World Wide Web Conference* http://www7.scu.edu.au/programme/fullpapers/1871/com1871.htm

Saracevic, T, 1975, "Relevance: a review of and a framework for the thinking on the notion in information science" *Journal of the American Society for Information Science* **26**(6) 321–343.

Schreck, J, 2001, "Security and privacy in user modeling" doctoral dissertation, Department of Mathematics and Computer Science, University of Essen, Germany.

Schwab, I, Kobsa, A and Koychev, I, 2000, Learning about users from observation. In *Adaptive User Interfaces: Papers from the 2000 AAAI Spring Symposium*. Menlo Park, CA: AAAI Press, 102–106.

Shardanand, U and Maes, P, 1995, "Social information filtering: algorithms for automating word of mouth" *Proceedings of the Human Factors in Computing Systems Conference (CHI-95)* ACM Press 210–217.

Shen, WM, 1996, "An efficient algorithm for incremental learning of decision lists" Technical Report, USC-ISI-96-012, Information Sciences Institute, University of Southern California. http://www.isi.edu/~shen/CDL4.ps.

Shneiderman, B, 2000, "Universal usability: pushing human-computer interaction to empower every citizen" *Communications of the ACM* **43**(5) 85–91.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press

Sivdas, E, Grewal, R and Kellaris, J, 1998, "The internet as a micro marketing tool: targeting consumers through preferences revealed in music newsgroup usage" *Journal of Business Research* **41** 179–186.

Sleeman, D, 1985, "A user modelling front-end subsystem" *International Journal of Man-Machine Studies* **23** 71–88.

Smyth, B and Cotter, P, 2000, "A personalised TV listings service for the digital TV age" *Knowledge-Based Systems* **13** 53–59.

Specht, M and Kobsa, A, 1999, "Interaction of domain expertise and interface design in adaptive educational hypermedia" *Second Workshop on Adaptive Systems and User Modeling on the World Wide Web at WWW-8, Toronto, Canada and UM-99, Banff, Canada* 89–93. `http://ics.uci.edu/~kobsa/papers/1999-WWW8UM99-kobsa.pdf`

Specht, M, 1998, "Empirical evaluation of adaptive annotation in hypermedia" *Proceedings of the ED-MEDIA98* 1327–1332.

Srinivasan VS and Park, CS, 1997, "Surprising robustness of the self-explicated approach to customer preference structure measurement" *Journal of Marketing Research* **34**(2) 286–291.

Stephanidis, C (ed.), 2001, *User Interfaces for All: Concepts, Methods, and Tools* Lawrence Erlbaum.

Stephanidis, C, Paramythis, A, Sfyrakis, M, Stergiou, A, Maou, N, Leventis, A, Paparoulis, G and Karagiandidis, C, 1998, "Adaptable and adaptive user interfaces for disabled users in AVANTI Project" in S Triglia, A Mullery, M Campolargo, H Vanderstraeten and M Mampaey (eds) *Intelligence in Services and Networks: Technology for Ubiquitous Telecom Services* Springer 153–166.

Strachan, L, Anderson, J, Sneesby, M and Evans, M, 1997, "Pragmatic user modelling in a commercial software system" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag 189–200. `http://www.um.org/um_97/gz/StrachanLps.gz`

Strachan, L, Anderson, J, Sneesby, M and Evans, M, 2000, "Minimalist user modelling in a complex commercial software system *User Modeling and User-Adapted Interaction* **10**(2–3) 109–146.

Taylor, JA, Carletta, J and Mellish, C, 1996, "Requirements for belief models in cooperative dialogue" *User Modeling and User-Adapted Interaction* **6**(1) 23–68.

Thomas, CG and Fischer, G, 1996, "Using agents to improve the usability and the usefulness of the World-Wide Web" *Fifth International Conference on User Modeling* 5–12.

Thomas, CG and Krogsæter, M, 1993, "An adaptive environment for the user interface of excel" *Proceedings of Intelligent User Interfaces '93* ACM Press 123–130.

Thompson, M, 1999, "Registered visitors are a portal's best friend" *The Industry Standard*, 7 June 1999. `http://www.thestandard.net`

van Schaik, P, Petrie, H, Japp, J and Edwards, J, 1996, "The design and evaluation of self-service terminals and smart cards for disabled and elderly users" in J Klaus, E Auff, W Kremser, WL Zagler (eds.) *Interdisciplinary Aspects on Computers: Helping People with Special Needs* R Oldenbourg 65–67.

Vassileva, J, 1996, "A task-centred approach for user modeling a hypermedia office documentation system" *User Modeling and User-Adapted Interaction* **6**(2/3) 185–223.

Wahlster, W and Kobsa, A, 1986, "Dialogue-based user models" *Proceedings of the IEEE* **74**(7) 948–960.

Wahlster, W, Andre, E, Bandyopadhyay, S, Graf, W and Rist, T, 1992, "WIP: the coordinated generation of multimodal presentations from a common representation" in A Ortony, J Slack and O Stock (eds) *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues* Springer 121–144.

Webb, GI and Kuzmycz, M, 1996, "Feature based modelling: a methodology for producing coherent, consistent, dynamically changing models of agent's competencies" *User Modeling and User-Adapted Interaction* **5**(2) 117–150.

Weber, G and Specht, M, 1997, "User modeling and adaptive navigation support in WWW-based tutoring systems" in A Jameson, C Paris and C Tasso (eds) *User Modeling: Proceedings of the Sixth International Conference* Springer-Verlag 289–300. `http://www.um.org/um_97/gz/WeberGps.gz`

Weida, R and Litman, D, 1992, "Terminological reasoning with constraint networks and an application to plan recognition" in B Nebel, C Rich and W Swartout (eds) *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR'92)* Morgan Kaufmann 282–293.

White, J, Chisholm, W and Vanderheiden, G (eds), 2001, "Web content accessibility guidelines 2.0" W3C Working Draft 25 January 2001. `http://www.w3.org/TR/WCAG20/`

Wu, D, 1991, "Active acquisition of user models: implications for decision-theoretic dialog planning and plan recognition" *User Modeling and User-Adapted Interaction* **1**(2) 149–172.

Yoshida, K and Motoda, H, 1996, "Automated user modeling for intelligent interface" *International Journal of Human-Computer Interaction* **8**(3) 237–258.

https://doi.org/10.1017/S0269888901000108 Published online by Cambridge University Press