Contents

X-XXXX-XXXX-X/XX/$X.XX+$X.XX

1

Adaptive Hypermedia and Adaptive Web

Peter Brusilovsky[1] **and Wolfgang Nejdl**[2], [1] *University of Pittsburg* [2] *L3S, University of Hannover*

CONTENTS

Abstract Adaptive Systems use explicit user models representing user knowledge, goals, interests, etc. that enable them to tailor interaction to different users. Adaptive hypermedia and Adaptive Web have used this paradigm to allow personalization in hypertext systems and the WWW, with diverse applications ranging from museum guides to web-based education. The goal of this chapter is to present the history of adaptive hypermedia, introduce a number of classic but popular techniques, and discuss emerging research directions in the context of the Adaptive and Semantic Web, that challenge the adaptive hypermedia researchers in the new Millennium.

1.1 Introduction

Web systems suffer from an inability to satisfy the heterogeneous needs of many users. For example, Web courses present the same static learning material to students with widely differing knowledge of the subject. Web stores offer the same selection of "featured items" to customers with different needs and preferences. Virtual museums on the web offer the same "guided tour" to visitors with different goals and interests. Health information sites present the same information to readers with different health problems. Adaptive hypermedia offers an alternative to the traditional "one-size-fits-all" approach. The use of adaptive hypermedia techniques allows Web-based systems to adapt their behavior to the goals, tasks, interests, and other features of individual users.

X-XXXX-XXXX-X/XX/$X.XX+$X.XX

Adaptive hypermedia systems belong to the class of user-adaptive software systems [Schneider-Hufschmidt et al., 1993]. A distinctive feature of an adaptive system is an explicit user model that represents user knowledge, goals, interests, and other features that enable the system to distinguish among different users. An adaptive system collects data for the user model from various sources that can include implicitly observing user interaction and explicitly requesting direct input from the user. The user model is used to provide an adaptation effect, i.e., tailor interaction to different users in the same context. Adaptive systems often use intelligent technologies for user modeling and adaptation.

Adaptive hypermedia is a relatively young research area. Starting with a few pioneering works on adaptive hypertext in early 1990, it now attracts many researchers from different communities such as hypertext, user modeling, machine learning, natural language generation, information retrieval, intelligent tutoring systems, cognitive science, and Web-based education. Nowadays adaptive hypermedia techniques are used almost exclusively for developing various adaptive Web-based systems. The goal of this chapter is to present the history of adaptive hypermedia, introduce a number of classic but popular techniques, and discuss emerging research directions in the context of the Adaptive Web, that challenge the adaptive hypermedia researchers in the new Millennium.

1.2 Adaptive Hypermedia

Adaptive Hypermedia research can be traced back to the early 1990s. At that time, a number of research teams had begun to explore various ways to adapt the output and behavior of hypertext systems to individual users. By the year 1996 several innovative adaptive hypermedia techniques had been developed, and several research-level adaptive hypermedia systems had been built and evaluated. A collection of papers presenting early adaptive hypermedia systems is available in [Brusilovsky et al., 1998b]. A review of early adaptive hypermedia systems, methods and techniques is provided by [Brusilovsky, 1996].

The year of 1996 can be considered a turning point in adaptive hypermedia research. Before this time, research in this area was performed by a few isolated teams. However, since 1996, adaptive hypermedia has gone through a period of rapid growth. In 2000 a series of international conferences on adaptive hypermedia and adaptive Web-based systems was established. The most recent event in this series, AH'2002, has assembled more than 200 of researchers.

Two major factors accounts for this growth of research activity: the maturity of adaptive hypermedia as a research field and the maturity of the Word Wide Web as an application platform.

The early researchers were generally not aware about each other's work. In contrast, many papers published since 1996 cite earlier work, and usually suggest an elaboration or an extension of techniques suggested earlier. Almost all adaptive hypermedia systems reported by 1996 were "classic hypertext" laboratory systems developed to demonstrate and explore innovative ideas. In contrast, almost all systems developed since 1996 are Web-based adaptive hypermedia systems with many of them being either practical systems, or research systems developed for real world settings.

The change of the platform from classic hypertext and hypermedia to the Web has also gradually caused a change both in used techniques and typical application areas. The first "pre-Web" generation of adaptive hypermedia systems explored mainly adaptive presentation and adaptive navigation support and concentrated on modeling user knowledge and goals [Brusilovsky et al., 1998b]. Empirical studies have shown that adaptive navigation support can increase the speed of navigation [Kaplan et al., 1993] and learning [Brusilovsky and Pesin, 1998], whereas adaptive presentation can improve content understanding [Boyle and Encarnacion, 1994]. The second "Web" generation brought classic technologies to the Web and explored a number of new technologies based on modeling user interests such as adaptive content selection and adaptive recommendation [Brusilovsky et al., 2000]. The first empirical studies report the benefits of using these technologies [Billsus et al., 2002]. The third "New Adaptive Web" generation strives to move adaptive hypermedia beyond traditional borders of closed corpus desktop hypermedia systems embracing such modern Web trends as “mobile Web”, “open Web”, and “Semantic Web”.

Early adaptive hypermedia systems were focusing almost exclusively on such academic areas as education or information retrieval. While these areas are still popular application areas for adaptive hypermedia techniques, most recent systems are exploring new, promising application areas such as kiosk-style information systems, e-commerce, medicine, and tourism. A few successful industrial systems [Billsus et al., 2002; Fink et al., 2002; Weber et al., 2001] show the commercial potential of the field.

1.2.1 What Can be Adapted in Adaptive Web and Adaptive Hypermedia

In different kinds of adaptive systems, adaptation effects could be different. Adaptive Web systems are essentially Webs of connected information items that allow users to navigate from one item to another and search for relevant items. The adaptation effect in this reasonably rigid context is limited to three major adaptation technologies - adaptive content selection, adaptive navigation support, and adaptive presentation. The first of these three technologies comes from the field of adaptive information retrieval (IR) and is associated with a search-based access to information. When the user searches for relevant information, the system can adaptively select and prioritize the most relevant items. The second technology was introduced by adaptive hypermedia systems [Brusilovsky, 1996] and is associated with a browsing-based access to information. When the user navigates from one item to another, the system can manipulate the links (e.g., hide, sort, annotate) to guide the user adaptively to most relevant information items. The third technology has some deep roots in the research on adaptive explanation and adaptive presentation in intelligent systems [Moore and Swartout, 1989; Paris, 1988]. It deals with presentation, not access to information. When the user gets to a particular page, the system can present its content adaptively.

Both adaptive presentation (content-level adaptation) and adaptive navigation support (link-level adaptation) have been extensively explored in a number of adaptive hypermedia projects. Early works on adaptive hypermedia were focused more on adaptive text presentation [Beaumont, 1994; Boyle and Encarnacion, 1994]. Later, the gradual growth of the number of nodes managed by a typical adaptive hypermedia systems (especially, Web hypermedia) have shifted the focus of research to adaptive navigation support techniques. Since adaptive navigation support is the kind of adaptation that is most specific to hypertext context, we provide a detailed review of several major techniques in the next section.

1.2.2 Adaptive Navigation Support

The idea of adaptive navigation support techniques is to help users to find their paths in hyperspace by adapting link presentation to the goals, knowledge, and other characteristics of an individual user. These techniques can be classified in several groups according to the way they adapt presentation of links. These groups of techniques are traditionally considered as different technologies for adapting link presentation. The most popular technologies are direct guidance, sorting, hiding, annotation, and generation.

Direct guidance is the simplest technology of adaptive navigation support. Direct guidance can be applied in any system which can suggest the “next best” node for the user to visit according user’s goals, knowledge, or/and other parameters represented in the user model. To provide direct guidance, the system can outline visually the link to the “best” node as done in Web Watcher [Armstrong et al., 1995], or present an additional dynamic link (usually called “next”) which is connected to the “next best” node as done in the InterBook [Brusilovsky et al., 1998a] or ELM-ART [Weber and Brusilovsky, 2001] systems. The former way is clearer; the latter is more flexible, because it can be used to recommend the node that is not connected directly to the current one (and not represented on the current page). A problem of direct guidance is that it provides no support for the users who would not like to follow the system’s suggestions. Direct guidance is useful but it should be used

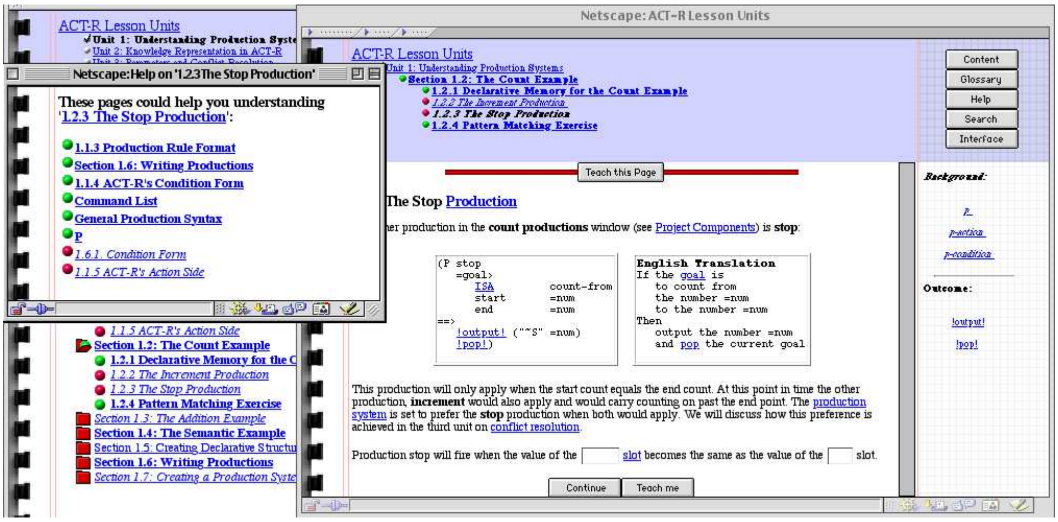

Figure 1.1: Adaptive guidance with "Teach Me" button and adaptive annotation with colored bullets in InterBook system

together with one of the more "supportive" technologies that are listed below. An example of an InterBook page with direct guidance is shown on Figure 1.1.

The idea of adaptive ordering technology is to order all the links of a particular page according to the user model and to some user-valuable criteria: the closer to the top, the more relevant the link is. Adaptive ordering has a limited applicability: it can be used with non-contextual (free-standing) links, but it can hardly be used for indexes and content pages (which usually have a stable order of links), and can never be used with contextual links (hot words in text) and maps. Another problem with adaptive ordering is that this technology makes the order of links unstable: it may change each time the user enters the page. For both reasons this technology is most often used now for showing new links to the user in conjunction with the link generation. Experimental research [Kaplan et al., 1993] showed that adaptive ordering can significantly reduce navigation time in IR hypermedia applications.

The idea of navigation support by hiding is to restrict the navigation space by hiding, removing, or disabling links to irrelevant pages. A page can be considered as not relevant for several reasons: for example, if it is not related to the user's current goal or if it presents materials which the user is not yet prepared to understand. Hiding protects users from the complexity of the unrestricted hyperspace and reduces their cognitive overload. Early adaptive hypermedia systems have used a simple way of hiding - essentially removing the link together with the anchor from a page. De Bra and Calvi [De Bra and Calvi, 1998] called this way link removal and have suggested and implemented several other variants for link hiding. A number of studies of link hiding demonstrated that users are unhappy when previously available links become invisible or disabled. Nowadays link hiding is mostly used in "reverse order" - as gradual link enabling when more and more links become visible for the user.

The idea of adaptive annotation technology is to augment links with some form of comments, which can tell the user more about the current state of the nodes behind the annotated links. These annotations can be provided in textual form or in the form of visual cues using, for example, different font colors [De Bra and Calvi, 1998], font sizes [Hohl et al., 1996], font types [Brusilovsky et al., 1998a] for the link anchor or different icons next to the anchor [Brusilovsky et al., 1998a; Henze and Nejdl, 2001; Weber and Brusilovsky, 2001]. Several studies have shown that adaptive link annotation is an effective way of navigation support. For example, [Brusilovsky and Pesin, 1998] has compared the performance of students who were attempting to achieve the same educational goal using ISIS-Tutor with and without adaptive annotation. The groups working with enabled adaptive navigation support were able to achieve this educational goal almost twice as faster and with significantly smaller navigation overhead. Another study [Weber and Brusilovsky, 2001] reported that advanced users of a Web-based educational system have stayed with the system significantly longer if provided with annotation-based adaptive navigation support.

Annotation can be naturally used with all possible forms of links. This technique supports stable order of links and avoids problems with incorrect mental maps. For all the above reasons, adaptive annotation has gradually grown into the most often used adaptive annotation technology.

One of the most popular methods of adaptive link annotation is the traffic light metaphor that is used primary in educational hypermedia systems. Green bullet in front of a link indicates recommended readings, while red bullet indicates that the student might not have enough knowledge to understand the information behind the link yet. Other colors like yellow or white may indicate other educational states. This approach was pioneered in 1996 in ELM-ART and InterBook systems [Brusilovsky et al., 1998a; Weber and Brusilovsky, 2001] and used later in numerous other adaptive educational hypermedia systems. Figure 1.1 shows adaptive annotation in InterBook [Brusilovsky et al., 1998a] and Figure 1.2 in KBS-HyperBook system [Henze and Nejdl, 2001].

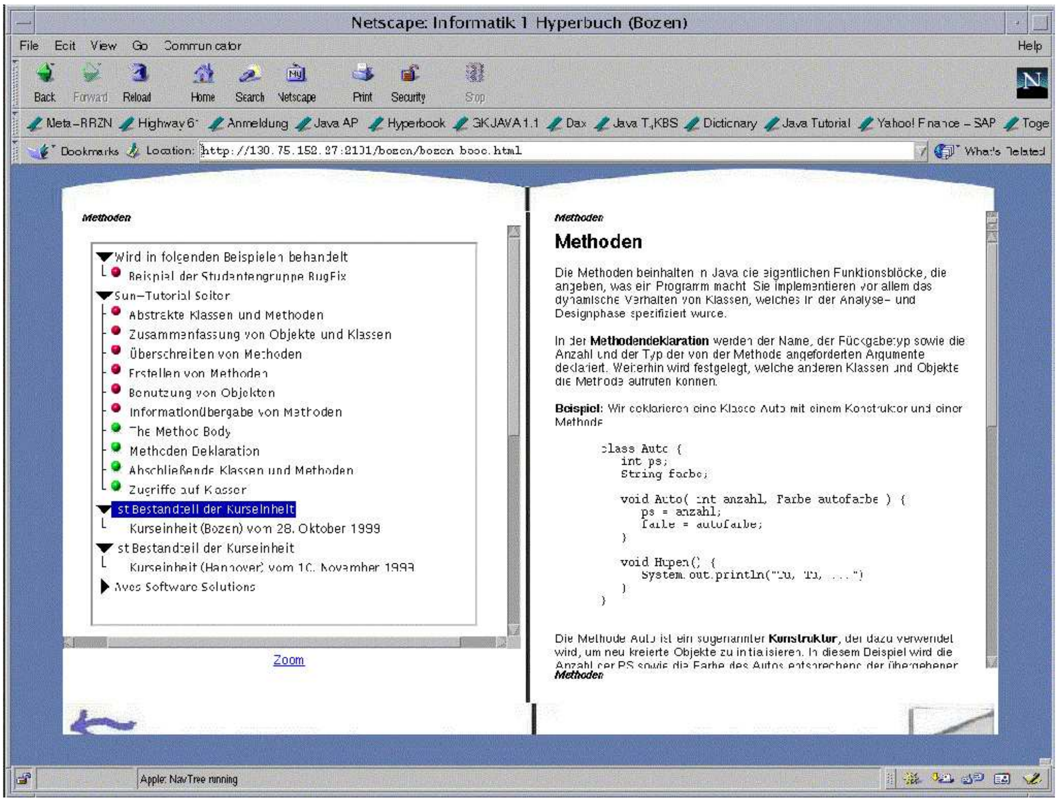

Figure 1.2: Adaptive annotation in KBS-HyperBook system

The last of major adaptive navigation support technologies is link generation. It became popular in Web hypermedia in the context of recommender systems. Unlike pure annotation, sorting and hiding technologies that adapt the way to present pre-authored links, link generation creates new, non-authored links for a page. There are three popular kinds of link generation: discovering new useful links between documents and adding them permanently to the set of existing links; generating links for similarity-based navigation between items; and dynamic recommendation of relevant links. The first two kinds have been present in the neighboring research area of intelligent hypertext for years. The third kind is relatively new, but already well explored in the areas of IR hypermedia, on-line information systems and even educational hypermedia.

Direct guidance, ordering, hiding, annotation, and generation are the primary technologies for adaptive navigation support. While most existing systems use exactly one of these ways to provide adaptive navigation support, these technologies are not mutually exclusive and can be used in combinations. For example, InterBook [Brusilovsky et al., 1998a] uses direct guidance, generation, and annotation. Hypadapter [Hohl et al., 1996] uses sorting, hiding, and annotation. Link generation is used almost exclusively with link sorting. Direct guidance technology can be naturally used in combination with any of the other technologies.

1.3 Adaptive Web

The traditional direction of adaptive hypermedia research - bringing adaptivity to classic hypermedia systems - is being quite well explored. The recent two to three years have added few new methods, techniques, and ideas. Most of the work being performed in this direction is now centered on developing specific variations of known methods and techniques and in developing practical systems. While this kind of work is important, some researchers may be more interested in expanding adaptive hypermedia beyond its traditional borders. We are now witnessing at least three exciting new directions of work towards an Adaptive Web, focused on mobile devices, open hypermedia and semantic web technologies. In the following three subsections we provide a brief overview of research currently being performed in these emerging areas. Most space is allocated to the Semantic Web direction, which is the most recent and probably the most challenging of the three.

1.3.1 Adaptive Hypermedia and Mobile Web

The work on adaptive hypermedia on handheld and mobile device that was not originally connected to Web hypermedia is quickly moving towards an Adaptive Mobile Web. From one side, various handheld and mobile devices such as portable computers or personal information managers (PIM) provide an attractive platform to run a number of hypermedia applications such as aircraft maintenance support systems [Brusilovsky and Cooper, 1999], museum guides [Not et al., 1998], or news delivery systems [Billsus et al., 2002]. From another side, the need for adaptation is especially evident in Mobile Web applications. Technologies such as adaptive presentation, adaptive content selection, adaptive navigation support that were an attractive luxury for desktop platform with large screens, high bandwidth and rich interface become a necessity for mobile handheld devices [Billsus et al., 2002].

Mobile Web has brought two major research challenges to the adaptive hypermedia community. First, most of mobile devices have relatively small screens. Advanced adaptive presentation and adaptive navigation support techniques have to be developed to make a “small-screen interface” more useable. Second, user location and movement in a real space becomes an important and easy-to-get (with the help of such devices as GPS) part of a user model. A meaningful adaptation to user position in space (and time) is a new opportunity that has to be explored.

More generally, mobile devices have introduced a clear need to extend the borders of adaptation. In addition to adaptation to the personal characteristics of users they demanded adaptation to user’s environment. Since users of the same server-side Web application can reside virtually everywhere and use different equipment, adaptation to user’s environment (location, time, computing platform, bandwidth) has become an important issue. A number of current adaptive hypermedia systems suggested some techniques to adapt to both the user location and the user platform. Simple adaptation to the platform (hardware, software, network bandwidth) usually involves selecting the type of material and media (i.e., still picture vs. movie) to present the content [Joerding, 1999]. More advanced technologies can provide considerably different interface to the users with different platforms and even use platform limitation to the benefits of user modeling. For example, a Palm Pilot version of AIS [Billsus and Pazzani, 2000] requires the user to explicitly request the following pages of a news story – thus sending a message to a system that the story is of interest. This direction of adaptation will certainly remain important and will likely provoke new interesting techniques. Adaptation to user location may be successfully used by many on-line information systems: SWAN [Garlatti and Iksal, 2000] demonstrates a successful use of user location for information filtering in a marine information system.

The currently most exciting kind of adaptive mobile web applications are mobile handheld guides. Mobile adaptive guides were pioneered by HYPERUDIO project [Not et al., 1998] well before the emergence of the mobile Web and now becoming very popular. Various recent projects explore a number of interesting adaptation techniques that take into account user location, direction of sight and movements in both "museum guide" [Oppermann and Specht, 1999] and "city guide" [Cheverst et al., 2002] contexts.

1.3.2 Open Corpus Adaptive Hypermedia

Currently, almost all adaptive hypermedia systems work with closed corpus set of documents assembled together at design time. This closed corpus is known to the system: knowledge about documents and links are traditionally obtained by manual indexing of documents and fragments with the user's possible knowledge, goals, background, etc. This approach cannot be applied to an open corpus such as open Web. To deal with the open Web an adaptive hypermedia system should be able to extend its set of documents with minimal efforts from the human side. A simple approach to do it is manually "extendable" hypermedia that allows an adaptive system to take into account documents that has not been indexed at the design time. The real challenge is to develop systems that are able to extract some meaning from an open corpus of documents and work with the open Web without the help of a human indexer. The research on open corpus adaptive hypermedia has benefited from the existing streams of work on open hypermedia and ontologies.

Open Hypermedia research for a long time evolved parallel to adaptive hypermedia research as well as to the World Wide Web, and have focused on hypermedia architectures which separate links from documents, and allow to process navigational structures independent of the content objects served by the hypermedia system (for example in the Microcosm system [Fountain et al., 1990], Chimera [Anderson et al., 1994] and Hyper-G [Andrews et al., 1995]).

Recently, this has lead to approaches incorporating adaptive functionalities in such an open hypermedia environment. [Bailey et al., 2002] for example build on the Auld Linky system [Michaelides et al., 2001], a contextual link server, which stores and serves the appropriate data structures for expressing information about content (data objects, together with context and behavior objects) and navigational structures (link structures, together with association and reference objects), among others. This makes it possible to provide basic adaptive functionalities (including link annotation or link hiding), and serve hypermedia content based on distributed content. Some pieces, however, still remain centralized in this architecture, e.g. the main piece of the hypermedia engine, the link server.

Once we want to integrate materials from different authors / heterogeneous sources, it becomes important to use commonly agreed sets of topics to index and characterize the content of the hypermedia pages integrated in the system [Henze and Nejdl, 2001, 2002]. This is addressed through the use of *ontologies*, which are "formal explicit specifications of shared conceptualizations" [Gruber, 1993]. In the process of ontology construction communities of users / authors agree on a topic hierarchy, possibly with additional constraints expressed in first order logic, which enables interoperability and exchangability between different sources.

Furthermore, for really open adaptive hypermedia systems, which "operate on an open corpus of documents" [Henze and Nejdl, 2002], the data structures and metadata should be compatible with those defined by current web standards. Therefore, the next step is to investigate, which metadata standards and representation languages should be used in the context of the World Wide Web, and whether centralized link servers can be substituted by decentralized solutions.

1.3.3 Adaptive Hypermedia and the Semantic Web

The basic idea of the hypermedia / hypertext paradigm is that information is interconnected by links, and different information items can be accessed by navigating through this link structure. The World Wide Web, by implementing this basic paradigm in a simple and efficient manner, has made this model the standard way for information access on the Internet. Obviously, in an open environment like the World Wide Web, adaptive functionalities like navigational hints and other personalization features would arguable be even more useful. To do this, however, we have to extend adaptation functionalities from the closed architectures of conventional systems to an open environment, and we have to investigate the possibilities of providing additional metadata based on Semantic Web formalisms in this open environment as input to make these adaptation functionalities possible.

In the previous section we discussed how hypermedia system architectures can be extended into an open hypermedia environment. This allows us to accommodate distributed content, but still relies on a central server and central data structures to integrate and serve distributed content. Peer-to-Peer infrastructures go a step further, and allow the provision of distributed services and content by a set of distributed peers, based on decentralized algorithms and data structures. An example for such a peer-to-peer infrastructure is the Edutella network (see e.g. [Nejdl et al., 2002a,b]), which implements a peer-to-peer infrastructure based on Semantic Web technologies. In this network, information is provided by independent peers, who can interchange information with others. Both data and metadata can be distributed in an arbitrary manner, data can be arbitrary digital resources, including educational content.

The crucial questions in such an environment are how to use standardized metadata to describe and classify information and to describe knowledge, preferences and experiences of users accessing this information. Last, but not least, adaptive functionalities as described in the previous sections have now to be implemented as queries in this open environment. Though a lot of questions still remain to be solved in this area, we'll sketch some possible starting points below (see also [Dolog et al., 2003]).

1.3.3.1 Describing Educational Resources

Describing our digital resources is the first step in providing a (distributed) hypermedia system. One of the most common metadata schemas on the web today is the "Dublin Core Schema" (DC) by the Dublin Core Metadata Initiative (DCMI). DCMI is an organization dedicated to promoting the widespread adoption of interoperable metadata standards and developing specialized metadata vocabularies for describing resources, that enable more intelligent information discovery for digital resources.

Each Dublin Core element is defined using a set of 15 attributes for the description of data elements, including Title, Identifier, Language and Comment. To annotate the author of a learning resource DC suggests to use the element creator, and thus we write for example *dc:creator(Resource) = nejdl*. Whereas "Simple Dublin Core" uses only the elements from the Dublin Core metadata set as attribute-value-pairs, "Qualified Dublin Core" (DCQ) employs additional qualifiers to further refine the meaning of a resource. Since Dublin Core is designed for metadata describing any kind of (digital) resource, it pays no heed to the specific needs we encounter in describing learning resources. The "Learning Objects Metadata Standard" (LOM) [IEEE-LTSC] by the IEEE Learning Technology Standards Committee (LTSC) was therefore established as an extension of Dublin Core.

These metadata can be encoded in RDF [Lassila and Swick, 1999; Brickley and Guha, 2003], which makes distributed annotation of resources possible. Using RDF Schema, we can represent the schemas as discussed above, i.e. the vocabulary to describe our resources. Specific properties are then represented as RDF triples $<$*subject, property, value*$>$, where *subject* identifies the resource we want to describe (using a URI), *property* specifies what property we use (e.g. *dc:creator*), and *value* the specific value, expressed as a string (e.g. "Nejdl") or another URI. We can then describe resources on the Web as shown in the following example:

```
<rdf:Description rdf:about="http://www.xyz.org/ai-2.html">
    <dc:title>Artificial Intelligence, Part 2</dc:title>
    <dc:author>Wolfgang Nejdl</dc:author>
    <dcq:requires resource="http://www.xyz.org/ai-1.html"/>
    <dcq:hasPart resource="http://www.xyz.org/ai-22.html"/>
</rdf:Description>
```

We can use any properties defined in the schemas we use, possibly mix different schemas without any problem, and also relate different resources to each other, for example when we want to express interdependencies between these resources, hierarchical relationships, or others.

1.3.3.2 Topic Ontologies for Content Classification

Personalized access means that resources are tailored according to some relevant aspects of the user. Which aspects of the user are important or not depends on the personalization domain. For educational scenarios it is important to take into account aspects like whether the user is student or a teacher, whether he wants to obtain a certain qualification, has specific preferences, and, of course, which is his knowledge level for the topics covered in the course.

Preferences about learning materials can be easily exploited, especially if they coincide directly with the metadata and metadata values used. For users preferring Powerpoint presentations for example, we can add the constraint *dc:format(Resource) = powerpoint* to queries searching appropriate learning materials.

Taking user knowledge about topics covered in the course into account is more tricky. The general idea is that we annotate each document by the topics covered in this document. Topics can be covered by sets of documents, and we will assume that a user fully knows a topic if he understands all documents annotated with this topic. However, though the standards we have just explored only provide one attribute (*dc:subject*) for annotating resources with topics, in reality we might want to have different kinds of annotations, to distinguish between just mentioning a topic, introducing a topic, and covering a topic. In the following we will simply assume that *dc:subject* is used for "covered" topics, but additional properties for these annotations might be useful in other contexts. Furthermore, we have to define which sets of documents for a given subject are necessary to "fully cover" a topic.

Additionally, it is obvious that self-defined keywords cannot be used, and we have to use an ontology for annotating documents and describing user knowledge (see also [Henze and Nejdl, 2002]). Defining a private ontology for a specific field works only in the closed microworld of a single university, so we have to use shared ontologies. One such ontology is the ACM Computer Classification system ([ACM2002]) which has been used by the Association for Computing Machinery since several decades to classify scientific publications in the field of computer science. This ontology can be described in RDF such that each entry in the ontology can be referenced by a URI and can be used with the *dc:subject* property as follows:

```
<rdf:Description rdf:about="http://www.xyz.org/ai-2.html">
  <dc:subject resource=
   "http://www.xyz.org/acm/ccs.rdf#I.1.2.4_Semantic_Networks"/>
</rdf:Description>
```

1.3.3.3 Describing Users

Though user profile standardization is not yet as advanced as learning object metadata standards, there are two main ongoing efforts to standardize metadata for user profiles, the IEEE Personal and Private Information (PAPI) [IEEE] project and the IMS Learner Information Package (LIP) [IMS]. If we compare these standards, we realize that they have been developed from different points of view.

IMS LIP provides us with richer structures and aspects. Categories are rather independent and the relationships between different records which instantiate different categories can be accomplished via the instances of the relationships category of the LIP standard. The structure of IMS LIP standard was derived from best practices in writing resumes. The IMS standard does not explicitly consider relations to other people though these can be represented by relationships between different records of the identification category. Accessibility policies to the data about different learners are not defined.

PAPI on the other hand has been developed from the perspective of a learner's performance during his study. The main categories are thus performance, portfolio, certificates and relations to other people (classmate, teacher and so on). This overlaps with the IMS activity category. However, IMS LIP defines activity category as a slot for any activity somehow related to a learner. To reflect this, IMS activity involves fields, which are related more to information required from management perspectives than from personalization based on level of knowledge. This can be solved in PAPI by introducing extensions and type of performance or by considering activity at the portfolio level, because any portfolio item is the result of some activity related to learning. PAPI do not cover the goal category at all, which can be used for recommendation and filtering techniques, and does not deal with transcript category explicitly. IMS LIP defines transcript as a record that is used to provide an institutionally-based summary of academic achievements. In PAPI, portfolio can be used, which will refer to an external document where the transcript is stored.

Using RDF's ability to mix features from more than one schema, we can use schema elements of both standards and also elements of other schemas. These RDF models can be accessible by different peers, and different and overlapping models are possible. Such distributed learner models were already discussed in [Vassileva et al., 2003] in the context of distributed learner modelling, though not in the context of RDF-based environments.

1.3.3.4 Adaptive Functionalities as Queries in a Peer-to-Peer Network

Based on the assumption that all resources managed within the network are described by RDF metadata, the Edutella peer-to-peer network [Nejdl et al., 2002a] provides a standardized query exchange mechanism for RDF metadata stored in distributed RDF repositories using arbitrary RDFS schemata.

To enable different repositories to participate in the Edutella network, Edutella wrappers are used to translate queries and results from a common Edutella query and result exchange format to the local format of the peer and vice versa, and to connect the peer to the Edutella network by a JXTA-based P2P library [Gong, 2001]. For communication with the Edutella network the wrapper translates the local data model into the Edutella Common Data Model (ECDM) and vice versa, and connects to the Edutella Network using the JXTA P2P primitives, transmitting the queries based on ECDM in RDF/XML form. The ECDM is based on Datalog, see e.g. [Garcia-Molina et al., 2002], which is a well-known non-procedural query language based on Horn clauses without function symbols. Datalog queries, which are a subset of Prolog programs and of predicate logic, easily map to relations and relational query languages like relational algebra or SQL or to logic programming languages like Prolog. In terms of relational algebra Datalog is capable of expressing selection, union, join and projection and hence is a relationally complete query language. Additional features include transitive closure and other recursive definitions.

Based on the RDF metadata managed within the Edutella network, we can now cast adaptive functionalities as Datalog queries over these resources, which are then distributed through the network to retrieve the appropriate learning resources. Personalization queries are then sent not only to the local repository, but to the entire Edutella network. In the following we use Prolog and first order predicate logic notation to express these queries, and use binary predicates to represent RDF statements.

In this way we can start to implement different adaptive hypermedia techniques, as described in the first part of this chapter. *Link annotation* for example can be implemented by an `annotate(+Page, +User, -Color)` predicate. We use the traffic light metaphor to express the suitability of the resources for the user, taking into account the user profile. A green icon represents a document that is recommended for reading, for example. We can formalize that a document is recommended for the user if it has not been understood yet and if all its prerequisites have already been understood:

```
forall Page, User, Prereq:
annotate(Page, User, green) <--
not_understood_page(Page, User),
prerequisites(Page, Prereq),
forall P in Prereq understood_page(P,User).
```

In Prolog, the criterion above and a query asking for recommended pages for nejdl then looks as follows:

```
annotated(Page, User, green) :-
    not_understood_page(Page, User),
    prerequisites(Page, Prereq),
    not (member(P, Prereq),
    not_understood_page(P, User)).

?- recommended (Page, nejdl, green)
```

Similar logic programs and queries have to be written for other adaptive functionalities as well. This not only leads to increased flexibility and openness of the adaptive hypermedia system, but also allows us to logically characterize adaptive hypermedia systems without restricting the means for their actual implementation [Henze and Nejdl, 2003].

1.4 Conclusion

Adaptive Hypermedia systems have progressed a lot since their early days. We now have a large range of possibilities available for implementing adaptive hypermedia systems. Special purpose and educational adaptive hypermedia system can be implemented on top of adaptive hypermedia engines or link servers with a large array of adaptive functionalities. In the World Wide Web context, adaptive functionalities and personalization features are gaining ground as well, and will extend the current Web to a more advanced Adaptive Web.

References

ACM2002. The ACM computing classification system. http://www.acm.org/class/1998/, 2002.

Kenneth M. Anderson, Richard N. Taylor, and E. James Whitehead. Chimera: Hypertext for heterogeneous software environments. In *European Conference on Hypermedia Technology*, September 1994.

Keith Andrews, Frank Kappe, and Hermann Maurer. Serving information to the Web with Hyper-G. In *Proceedings of the 3rd Intl. World Wide Web Conference*, Darmstadt, Germany, April 1995.

Robert Armstrong, Dayne Freitag, Thorsten Joachims, and Tom Mitchell. WebWatcher: A learning apprentice for the world wide web. In C. Knoblock and A. Levy, editors, *AAAI Spring Symposium on Information Gathering from Distributed, Heterogeneous Environments*, pages 6–12, Menlo Park, CA, 1995. AAAI Press.

Christopher Bailey, Wendy Hall, David Millard, and Mark Weal. Towards open adaptive hypermedia. In *Proccedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2002)*, Malaga, Spain, May 2002.

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

Daniel Billsus, Clifford A. Brunk, Craig Evans, Brian Gladish, and Michael Pazzani. Adaptive interfaces for ubiquitous web access. *Commun. ACM*, 45(5):34–38, 2002.

Daniel Billsus and Michael J. Pazzani. A learning agent for wireless news access. In Henry Lieberman, editor, *2000 International Conference on Intelligent User Interfaces*, pages 94–97, New Orleans, LA, 2000. ACM Press.

Craig Boyle and Antonio O. Encarnacion. MetaDoc: an adaptive hypertext reading system. *User Modeling and User-Adapted Interaction*, 4(1):1–19, 1994.

Dan Brickley and Ramanathan V. Guha. RDF Vocabulary Description Language 1.0: RDF Schema, January 2003. http://www.w3.org/TR/rdf-schema/.

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

Peter Brusilovsky and David W. Cooper. ADAPTS: Adaptive hypermedia for a web-based performance support system. In Peter Brusilovsky and Paul De Bra, editors, *Second Workshop on Adaptive Systems and User Modeling on WWW at 8th International Word Wide Web Conference and 7-th International Conference on User Modeling*, pages 41–47, Toronto and Banff, Canada, 1999. Eindhoven University of Technology.

Peter Brusilovsky, John Eklund, and Elmar Schwarz. Web-based education for all: A tool for developing adaptive courseware. *Computer Networks and ISDN Systems*, 30(1-7):291–300, 1998a.

Peter Brusilovsky, Alfred Kobsa, and Julita Vassileva, editors. *Adaptive Hypertext and Hypermedia*. Kluwer Academic Publishers, Dordrecht, 1998b.

Peter Brusilovsky and Leonid Pesin. Adaptive navigation support in educational hypermedia: An evaluation of the ISIS-Tutor. *Journal of Computing and Information Technology*, 6(1):27–38, 1998.

Peter Brusilovsky, Olivero Stock, and Carlo Strapparava, editors. *Adaptive Hypermedia and Adaptive Web-based Systems, AH2000*, volume 1892 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 2000.

Keith Cheverst, Keith Mitchell, and Nigel Davies. The role of adaptive hypermedia in a context-aware tourist guide. *Communications of the ACM*, 45(5):47–51, 2002.

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

Peter Dolog, Rita Gavriloaie, and Wolfgang Nejdl. Integrating adaptive hypermedia techniques and open rdf-based environments. In *Proceedings of the 12th World Wide Web Conference*, pages 88–98, Budapest, Hungary, May 2003.

Josef Fink, Jürgen Koenemann, Stephan Noller, and Ingo Schwab. Putting personalization into practice. *Communications of the ACM*, 45(5):41–42, 2002.

Andrew Fountain, Wendy Hall, Ian Heath, and Hugh Davis. MICROCOSM: An open model for hypermedia with dynamic linking. In *Proceedings of the ACM Conference on Hypertext*, 1990.

Hector Garcia-Molina, Jeffrey Ullman, and Jennifer Widom. *Database Systems - The Complete Book*. Prentice Hall, 2002.

Serge Garlatti and Sébastien Iksal. Context filtering and spacial filtering in an adaptive information system. In Peter Brusilovsky, Olivero Stock, and Strapparava Carlo, editors, *Int. Conference on Adaptive Hypermedia and Adaptive Web-based systems*, pages 315–318, Berlin, 2000. Springer-Verlag.

Li Gong. Project JXTA: A technology overview. Technical report, SUN Microsystems, April 2001. http://www.jxta.org/project/www/docs/TechOverview.pdf.

Tom Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5: 199–220, 1993.

Nicola Henze and Wolfgang Nejdl. Adaptation in open corpus hypermedia. *International Journal of Artificial Intelligence in Education*, 12(4):325–350, 2001.

Nicola Henze and Wolfgang Nejdl. Knowledge modeling for open adaptive hypermedia. In *Proccedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH 2002)*, pages 174–183, Malaga, Spain, May 2002.

Nicola Henze and Wolfgang Nejdl. Logically characterizing adaptive educational hypermedia systems. In *Proc. of AH2003 Workshop, 13th World Wide Web Conference*, Budapest, Hungary, May 2003.

Hubertus Hohl, Heinz-Dieter Böcker, and Rul Gunzenhäuser. Hypadapter: An adaptive hypertext system for exploratory learning and programming. *User Modeling and User-Adapted Interaction*, 6(2-3):131–156, 1996.

IEEE. IEEE P1484.2/D7, 2000-11-28. draft standard for learning technology. public and private information (papi) for learners. Available at: http://ltsc.ieee.org/wg2/.

IEEE-LTSC. IEEE LOM working draft 6.1. Available at: http://ltsc.ieee.org/wg12/index.html.

IMS. IMS learner information package specification. Available at: http://www.imsproject.org/profiles/index.cfm.

Tanja Joerding. A temporary user modeling approach for adaptive shopping on the web. In Peter Brusilovsky and Paul De Bra, editors, *Second Workshop on Adaptive Systems and User Modeling on the World Wide Web*, pages 75–79, Toronto and Banff, Canada, 1999. Eindhoven University of Technology.

Craig Kaplan, Justine Fenwick, and James Chen. Adaptive hypertext navigation based on user goals and context. *User Modeling and User-Adapted Interaction*, 3(3):193–220, 1993.

Ora Lassila and Ralph R. Swick. W3C Resource Description Framework model and syntax specification, February 1999. http://www.w3.org/TR/REC-rdf-syntax/.

Danius T. Michaelides, David E. Millard, Mark J. Weal, and David De Roure. Auld Leaky: A contextual open hypermedia link server. In *Proceedings of the 7th Workshop on Open Hypermedia Systems, ACM Hypertext 2001 Conference*, Aarhus, Denmark, 2001.

Johanna D. Moore and William R. Swartout. Pointing: A way toward explanation dialogue. In *Eight National Conference on Artificial Intelligence*, pages 457–464, 1989.

Wolfgang Nejdl, Boris Wolf, Changtao Qu, Stefan Decker, Michael Sintek, Ambjoern Naeve, Mikhael Nilsson, Matthias Palmr, and Tore Risch. EDUTELLA: A P2P networking infrastructure based on RDF. In *Proceedings of the 11th International World Wide Web Conference (WWW2002)*, Hawaii, USA, June 2002a.

Wolfgang Nejdl, Boris Wolf, Steffen Staab, and Julien Tane. EDUTELLA: Searching and annotating resources within an RDF-based P2P network. In *Proceedings of the Semantic Web Workshop*, Honolulu, Hawaii, May 2002b.

Elena Not, Daniela Petrelli, Marcello Sarini, Olivero Stock, Carlo Strapparava, and Massimo Zancanaro. Hypernavigation in the physical space: adapting presentation to the user and to the situational context. *New Review of Multimedia and Hypermedia*, 4:33–45, 1998.

Reinhard Oppermann and Marcus Specht. Adaptive information for nomadic activities. a process oriented approach. In *Software Ergonomie '99*, pages 255–264, Walldorf , Germany, 1999. Stuttgart: Teubner.

Cécile Paris. Tailoring object description to a user's level of expertise. *Computational Linguistics*, 14(3):64–78, 1988.

Matthias Schneider-Hufschmidt, Thomas Kühme, and Uwe Malinowski, editors. *Adaptive user interfaces: Principles and practice*. Human Factors in Information Technology. North-Holland, Amsterdam, 1993.

Julita Vassileva, Gordon McCalla, and Jim Greer. Multi-agent multi-user modelling in I-Help. *User Modeling and User-Adapted Interaction*, 13(1):179–210, February 2003. Kluwer Academic Press.

Gerhard Weber and Peter Brusilovsky. ELM-ART: An adaptive versatile system for web-based instruction. *International Journal of Artificial Intelligence in Education*, 12(4):351–384, 2001.

Gerhard Weber, Hans-Christian Kuhl, and Stephan Weibelzahl. Developing adaptive internet based courses with the authoring system NetCoach. In Paul De Bra, Peter Brusilovsky, and Alfred Kobsa, editors, *Third workshop on Adaptive Hypertext and Hypermedia*, pages 35–48, Sonthofen, Germany, 2001. Technical University Eindhoven.

Index

X-XXXX-XXXX-X/XX/$X.XX+$X.XX