

Adaptive Content Presentation for the Web

Andrea Bunt, Giuseppe Carenini, Cristina Conati

Department of Computer Science
University of British Columbia
{bunt, carenini, conati}@cs.ubc.ca

Abstract. In this chapter we describe techniques for adaptive presentation of content on the Web. We first describe techniques to select and structure the content deemed to be most relevant for the current user in the current interaction context. We then illustrate approaches that deal with the problem of how to adaptively deliver this content.

1. Introduction

Previous chapters in this book have described types of adaptation for Web-based systems that include adaptive navigation support (see Ch X), adaptive search (see Ch XX) and personalized recommendation of items of interest (see Ch XXX). In this chapter, we will focus on an additional type of adaptation widely known as *adaptive presentation of content*: how to present Web-based content in a manner that best suits individual users' needs. This type of adaptation can complement several of the adaptation types discussed in previous chapters. For instance, the content of Web pages pointed to by a tailored link in a system that provides adaptive navigation support (Ch X), or returned by adaptive search (Ch XX) can be modified to highlight the parts that are more interesting for the current user. Similarly, the description of the items returned by a recommender system (Ch XXX) can be adapted to play up the items' features that are more relevant to the user's needs, or changed to be more suitable to the user's level of familiarity with the items.

The focus of this chapter will be on computational techniques necessary to provide the user with a tailored presentation of content, rather than implementation details and technologies. Also, the chapter is not limited to techniques used in adaptive Web-based applications. It aims to suggest areas of future research by discussing alternative approaches that have a strong potential to augment the set of existing techniques for adaptive presentation on the Web.

The chapter is structured as follows. In section 2, we address techniques for content adaptation, i.e., techniques to decide what content is most relevant to the current user and how to structure this content in a coherent way, *before* presenting it to the user. These techniques range from those that require the existence of pre-crafted versions of the relevant content, to those that can automatically adapt content from abstract knowledge sources. Our discussion focuses more on the latter, given their relative novelty in Web-based applications and their higher level of sophistication. In section 3, we discuss techniques for adaptively presenting the

selected relevant content to the user. In section 3.1, we introduce techniques that deal with the problem of how to present this content so that user focus/attention is drawn to the most relevant information (possibly defined by using any of the techniques described in section 2) while still preserving the contextual information that can often be provided by content of secondary importance. In section 3.2, we discuss techniques to decide which media/modality to use to best convey the selected content.

2. Techniques for Content Adaptation

The process of adapting content to specific user needs can be thought of as two main sub processes. The first sub process, which we will refer to as *content adaptation*, involves understanding what content can be most relevant to the current user's interests, and how this content should be organized. The second sub process, which we will refer to as *content presentation* involves deciding how to most effectively present the selected content to the user. This section reviews techniques for content adaptation, while section 3 will focus on content presentation. Content adaptation of Web pages can be characterized along the following key dimension: the nature of the content provided as input. Along this dimension, we first briefly describe two rather simple approaches in which adaptation is achieved by selecting appropriate canned pages or page fragments. These approaches are referred to in the literature as *page* and *fragment variants* respectively, and they have been extensively discussed in previous surveys (e.g., [30]). After a brief description of page and fragment variants, we provide an in-depth discussion of more sophisticated approaches to content adaptation in which the input is abstract information. To the best of our knowledge, these approaches have never been covered in detail in any previous survey on adaptive hypermedia.

2.1 Approaches Based on Page and Fragment Variants

The simplest form of content adaptation is the page-variant approach [30]. Here, the input of the adaptation process consists of different versions of each adaptable page along with a model of the user/context (we refer to this combination as the interaction context from now on). These versions have to be written in advance. At runtime, the adaptation mechanism selects one that is most appropriate to the current interaction context for presentation. Clearly, this approach does not scale up to complex adaptation. If several aspects of the page must be adapted in many different ways, an unmanageable large number of variants need to be written. Nevertheless, in some domains, where only high-level adaptation is needed, this approach has been effectively applied. For instance, in the ORIMUHS system [15] page variants are applied to support user interaction in two complex software systems: a CAD modeler and a medical application. Page variants are also applied in the KBS Hyperbook system [24] to develop educational courseware on Java programming.

Moving up in the ladder of adaptation complexity, we have the fragment-variant approach. In this approach, the adaptation is performed at a finer level of granularity. More specifically, the page presented to the user is not selected from a pool of fixed pages. Rather, it is constructed by selecting and combining an appropriate set of

fragments, where each fragment typically corresponds to a self-contained information element, such as a text paragraph or a picture. As with the page-variant approach, these fragments are written in advance. Two common strategies for fragment variants are: *optional fragments* and *altering fragments*. In optional fragments, a page is specified as a set of fragments, where each fragment is associated with a set of applicability conditions. At runtime, the page is generated by selecting only those fragments whose conditions are satisfied in the current interaction context. For instance in [17], different optional fragments are selected depending on the user's knowledge, interests and abilities. Altering fragments are rather different from optional fragments. In altering fragments, a page is specified as a set of constituents, and for each constituent there is a corresponding set of fragments. At runtime, the page is created by selecting for each constituent the fragment that is most appropriate in the current interaction context. Altering fragments are applied, for instance, in the AHA system [14], in which different presentations of the same entity can be selected depending on whether the target user has the necessary background knowledge.

In general, a noticeable disadvantage of fragment variants is that the selection and assembly of a suitable set of fragments may involve a substantial overhead at runtime. Furthermore, it may sometimes be difficult to combine the set of independently selected fragments into a coherent whole. On the other hand, the key advantage of this approach is that, once a set of fragments and their applicability conditions have been written, a large number of pages can be automatically generated to cover a corresponding large number of situations. For pointers to specific techniques to implement the fragment-variant approach the reader should refer to [30].

Note that because in the two approaches above the units of content adaptation are either whole pages or predefined page components, the two sub processes of content adaptation and presentation actually coincide. That is, the decision of what content is most relevant to the user (i.e. the page to be displayed) uniquely identifies what will be presented to the user. On the one hand, this simplifies and speeds up the complete adaptation process. On the other hand, it reduces flexibility because it eliminates the possibility to further tailor the information through adaptive presentation techniques once the first level of adaptive content presentation, content selection, has been achieved, as we will see in section 3.

2.2 Approaches Based on Abstract Information

Although many adaptive Web systems have been designed in recent years by relying only on page or fragment variants, in this section we describe more sophisticated adaptation techniques that require as input abstract information the system can reason about. Notice that part of the research on sophisticated content adaptation has been developed in the field of Natural Language Generation (NLG) [39], which investigates how natural language text can be generated from abstract non-linguistic information.

Sophisticated content adaptation, also called tailoring in NLG, requires an abstract representation of the domain from which the content is selected, as well as the features of the interaction context to which the content is tailored. Several formalisms have been used in the literature, including:

- *Traditional knowledge bases* [42] expressing domain entities and relationships between them. For instance, one application of the ILEX system [37] generates tailored jewel labels by relying on a large object-centered knowledge base about jewelers, materials, designers, etc. This knowledge base includes both abstract propositions, such as the fact that a necklace is a jewel, and specific propositions, such as the fact that a particular jewel was made in Birmingham in 1905.
- *Bayesian Networks* [42] expressing probabilistic relationships between random variables representing the domain. For instance, one application of the NAG system [49] [31] generates arguments about the expected rate of a researcher's future publications by relying on a Bayesian Network. This network specifies probabilistic relationships between the publication rate of a researcher and the factors that influence it, such as the strength of the institution from which the researcher graduated (e.g., the stronger is the institution the higher the likelihood of a high publication rate).
- *Preference models* [42] expressing the user's preferences about different aspects of the domain. For instance, one application of the GEA system [8] generates user-tailored arguments on whether the user will like/dislike a given house by relying on a model specifying what aspects of a house the user cares most about (e.g., location, amenities).

Depending on the application, the same or different formalisms can be used to represent the domain model and the user/context. For instance, in NAG [49], a system for generating factual arguments (claiming that something is or is not the case), both the domain and user model are represented as Bayesian Networks. Similarly, in HYLITE+ [5], a system for generating adaptive hypertext encyclopedia-style explanations, both the domain and the user models are expressed as traditional knowledge bases, more precisely as conceptual graphs [44]. In contrast, in GEA, a system for generating evaluative arguments (claiming that something is good vs. bad), the domain model is represented as a traditional knowledge base while the user model is expressed as a value tree [42], which is a preference model commonly used in decision theory.

The process of sophisticated content adaptation involves the two conceptually distinct phases of *content selection/determination* and *content structuring*, also jointly referred to as content planning. Although we will describe them separately to simplify the presentation, it should be noted that content selection and structuring are often implemented as one single process that simultaneously performs both phases [39].

2.2.1 Content Selection

During content selection, a subset of the domain knowledge is identified, possibly through some reasoning mechanism, as relevant for the current user and situation. Strategies for content selection rely on domain-specific knowledge to different degrees. For instance, the content selection strategy used in STOP, a system for generating smoking cessation letters, is quite domain specific as it refers to psychological knowledge about addictive behavior and smoking [40]. In contrast, the content selection strategy used by the GEA system does not rely on any domain-specific knowledge (as we will see later in this section) and can be therefore applied in any domain [10]. Because of their generality, in this section we focus on strategies

that are primarily domain-independent. For a discussion of more domain-specific strategies and in particular of how they can be acquired, the reader should refer to [39].

In practice, most domain-independent strategies for content selection compute a measure of relevance for each content element (i.e., fact) and then use this measure to select an appropriate subset of the available content. Content adaptation is achieved by having this measure of relevance take into account features of the current user and context. For illustration, let's consider three systems that cover the key design options.

The Intelligent Labeling Explorer (ILEX) - We start with ILEX [37], a system for generating contextually-relevant hypertext descriptions of objects (e.g., museum artifacts, computer components). In ILEX, the measure of relevance for content selection combines a measure of structural relevance of a knowledge element/fact with its intrinsic score. Structural relevance takes into account the structure of the domain knowledge base - a semantic net. More specifically, structural relevance is computed starting from the focal entity (i.e., the entity being described) by considering two basic heuristics: (i) information becomes less relevant the more distant it is from the focal object, in term of semantic links; (ii) different semantic link types (e.g., GENERALISE) maintain relevance to different degrees. The intrinsic score of a knowledge element combines numerical estimates of three factors: (i) the potential interest of the information to the current user, (ii) the importance of the information to the system's informational goals and (iii) to what degree the user may already know this information. Once the two measures of structural relevance and intrinsic score have been computed, they are combined in a single measure of relevance by straight multiplication.

In ILEX, the content selection strategy is then to return the n most relevant knowledge elements. However, if the selection process based on relevance cannot find a sufficient number of knowledge elements, additional content selection routines are activated. For instance, one technique applied by ILEX is to identify an entity which is sufficiently similar to the focal entity, so that an interesting comparison between the two can be also selected for presentation. Notice that this introduces a point of general interest. When the goal of the content selection component is to return a fixed amount of content, it may be necessary to supplement the main selection strategy with a set of ancillary strategies.

The Generator of Evaluative Arguments (GEA) - As mentioned before, GEA is a system for generating evaluative arguments (claiming that something is good vs. bad). Here the measure of relevance is computed by applying a quantitative model of the user preferences to the entity being evaluated. Generally speaking, GEA's user model relies on the notion that if something is valued, it is valued for multiple reasons. More specifically, for each user, the model specifies a decomposition of the user's overall assessment of entities in a given class (e.g., houses) into a hierarchy of aspects of the entities (e.g., location, number-of-bedrooms). The model hierarchy is annotated with numerical weights and functions that specify the relative importance/preferability of each attribute and domain-value (e.g., two bedrooms) for the particular user. Once the model is applied to an entity, it is possible to compute for each attribute how much its evaluation contributes to the overall evaluation of that entity for the current user. Based on this, a measure of relevance is defined by assuming that an attribute is relevant either because of its strength or because of its

weakness in contributing to the value of the entity. For instance, if distance-from-work is an important attribute for the current user, this attribute will have high relevance with respect to the evaluation of a house that is very close to the user's workplace (because of its strong contribution), as well as to the evaluation of a house that is very far from it (because of its weak contribution).

Once the relevance of all attributes is assessed, the content selection strategy in GEA is to return for each level in the user model hierarchy, only those attributes whose relevance is greater than a (customizable) threshold. By setting this threshold to different values it is possible to generate, in a principled way, arguments that contain different amounts of user-tailored relevant content, as shown in Figure 1 (see [9] for details).

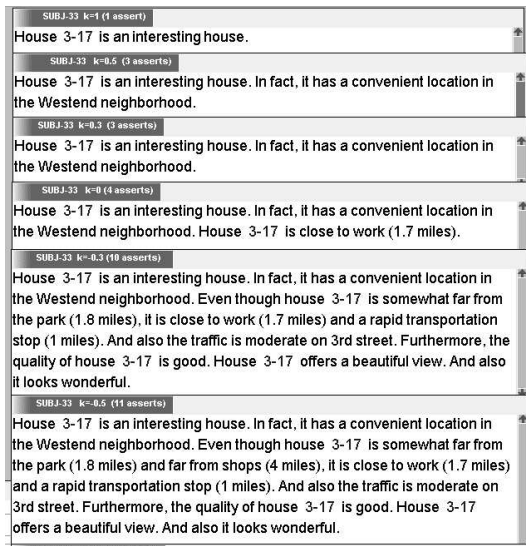


Figure 1: GEA system - Arguments about the same house tailored to the same user containing an increasing (or equal) amount of content

The Responsive Information Architect (RIA) - In RIA [47], a multimedia conversation system to support information seeking tasks, content selection is formalized as an optimization problem. The goal is to identify the most desirable subset of data dimensions (e.g., price and style in the real-estate domain) in the current interaction context. The desirability of each data dimension is computed as the linear combination of a large set of feature-based metrics that characterize how important the dimension is with respect to the interaction context. Most of these features are labeled as content relevance features and include features of the data (e.g., inherent importance of the dimension in the domain, such as the importance of price in the real-estate domain), features of the user (e.g., relation of the dimension to the user's interests, such as how much the user cares about price), as well as features relating the dimension to the user request (e.g., the user may be specifically

requesting information on a particular dimension, such as explicitly requesting information on price) and the interaction history, which tries to maintain coherence in which dimensions are presented across multiple queries. So, RIA desirability is fundamentally what we have been referring to as a measure of content relevance. Once all the available data dimensions have been assigned their desirability, RIA's content selection strategy returns the set of data dimensions such that their overall desirability is maximized and their cost is within the given space and time allocated for the target presentation. Special techniques are introduced to take care of dependencies between data dimensions. For instance, some groups of dimensions are considered as a single bundle (e.g., number of bedrooms and number of bathrooms) and are either all included or all excluded from the final presentation.

Figure 2 shows how RIA, given the same query, selects different data dimensions for two users with different preferences in the real estate domain.

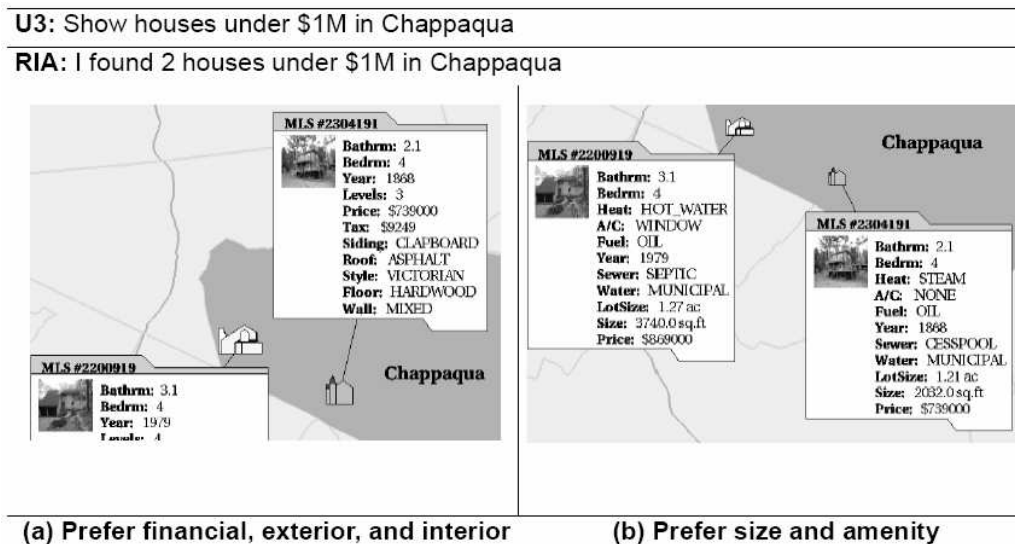


Figure 2: RIA system: examples of queries and user-tailored responses for two users with different preferences

2.2.2 Content Structuring

Once the most relevant content elements are selected they must be organized in order to be effectively communicated/presented. This involves not only ordering and grouping them, but also specifying what discourse relations (e.g., contrast, evidence) [29] must hold between the resulting groups. Schemas [39] are the method of choice to accomplish all these tasks and are commonly implemented with task-decomposition planners (technically referred to as HTN planners [42]).

With respect to content selection, we provide a much more limited treatment of content structuring because adaptation of the latter is rather less common than adaptation of the former. One form of structure adaptation relies on the measure of relevance used in content selection for ordering the selected content elements. For instance, in GEA the selected elements are ordered according to the measure of relevance by following principles from argumentation theory [10]. Another form of content structuring adaptation is the selection of the discourse relations. For instance, in GEA a given fact can be selected as supporting or contrasting evidence depending on the user's preference for that fact, as determined by the evaluation of the user model.

Note that any structuring information derived from this phase of content adaptation can serve as a guide to decide how to actually present the selected content to the user. The various techniques for adaptive content presentation that we will overview in the next section can be used to adaptively render the information defined by content structuring. For instance, techniques for content emphasis can be used to express relevance information implicit in content ordering. Discourse structure can be used to identify portions of texts that can be made available on demand instead of being displayed up-front [38]. Finally, relationship information can be used to define which medium to use to display the related content (e.g. use graphics to highlight quantitative relationships such as a comparison between two sets of numbers [28]). In practice, most previous and current work on adaptive content presentation relies on simpler text-based fragments, which are selected according to a simple user model. However, we believe that work on integrating adaptive content structuring with adaptive content presentation promises to be of great value for the advancement of adaptive content presentation on the Web.

We conclude this section on content adaptation from abstract information with a brief discussion of issues related to knowledge acquisition and evaluation. While all of the techniques for content adaptation from abstract information discussed so far have been engineered by researchers, very recent and promising work is exploring how content adaptation strategies can be learned from user feedback on sample (multimedia) presentations [20]. This work relies on machine learning techniques that have been already successfully used in other NLG adaptation tasks [45]. As for evaluation, recent years have witnessed a surge of interest in empirically testing techniques for content adaptation from abstract information. Human judges [12], human designers [47] and task efficacy (e.g., [5], [11]) are the three basic methods that have been applied to evaluate such techniques. These methods are described under the category "controlled experiments" in Chapter X, which provides a comprehensive overview of all empirical and non-empirical evaluation methods for adaptive Web systems.

3. Techniques for Content Presentation

In the previous section we discussed techniques to identify and structure the content most relevant to the interaction context. Here, we focus on techniques to

effectively present content once its degree of relevance and its structure have been determined by content adaptation. In particular, we present techniques that decide how to present content based on its relevance, and techniques to select the type of media most appropriate to deliver the content, given the interaction context.

3.1 Relevance-Based Techniques

Most of the techniques that we categorize here as relevance based were introduced in [7] and [30] as ways to manage canned content fragments. Following [46], in this section we will discuss them along two general dimensions, which we see as critical to both canned and generated content:

- Maintaining *focus*, i.e., how these techniques emphasize the content that has been classified as most relevant for the current user
- Maintaining *context*, i.e., if and how they allow for access of the less relevant content so as to preserve the contextual information that it may provide

There is an obvious tradeoff between these two dimensions: context is more easily maintained if much of the original content is visible to the user. However, the more content is shown, the higher the chance of generating information overload and reducing attention to the most relevant information, defeating one of the very reasons for having adaptive hypermedia in the first place.

The techniques we address here can be grouped in two main categories, depending upon how they address the context-focus tradeoff: Priority on Focus and Priority on Context.

3.1.1 Priority on Focus

All of the techniques in the Priority on Focus category choose to maximize focus by (a) showing the user only the content that is deemed most relevant, and (b) precluding access to the rest of the content. They include not only the fragment-variant techniques that we discussed in the previous section (i.e., optional fragments and altering fragments), but also any strategy for sophisticated content selection in which only the most relevant knowledge elements are presented to the user (see for instance Figure 1 and Figure 2 in section 2).

In addition to potentially losing contextual information when limiting the content the user can see, Priority on Focus techniques suffer from two main drawbacks:

- They are highly impacted by the validity of the adaptation mechanism, as the user has no way to recover from bad adaptation.
- They do not allow for user control, one of the dimensions that defines usability in human-computer interaction and that significantly influences acceptance of adaptive interfaces ([26] and [27]).

Priority on Focus techniques seem to be mostly used in Adaptive Educational Hypermedia (see [14] and [36]), possibly because these systems are less subject to the above drawbacks. The pedagogical nature of the interaction makes it both easier for the system to create an accurate user model of relevant user traits, and possibly more

acceptable for the user to have limited control over the computer tutor's adaptation decisions.

3.1.2 Priority on Context

The Priority on Context category includes the techniques known as (1) *stretchtext*, (2) *dimming fragments* and (3) *scaling fragments*. The first two techniques have been around for a while and were previously addressed in adaptive hypermedia reviews (e.g., [7]). The scaling fragments technique is a more recent attempt to adapt the well known *fisheye* visualization technique [18] to content organization in adaptive hypermedia. For this reason, we will describe it in slightly more detail than the other two. In general, all three techniques try to preserve the context around the most relevant content by providing different ways to make the non-selected less relevant information visible without distracting the user from the primary content.

Stretchtext, arguably the most well established of the three techniques, relies on placeholders to signal the presence of and allow access to information of secondary importance (see for instance [6] and [25]). Usually the place holders are short headers summarizing the hidden information, as shown in the screen shot of the Push system [25] in Figure 3 (see items labeled as IE – Information Entities – and Hotlist). Note that, while this technique has been mostly used with text (hence the name), researchers have started generalizing it to the adaptive presentation of multimedia content, for instance in Interactive TV applications [34].

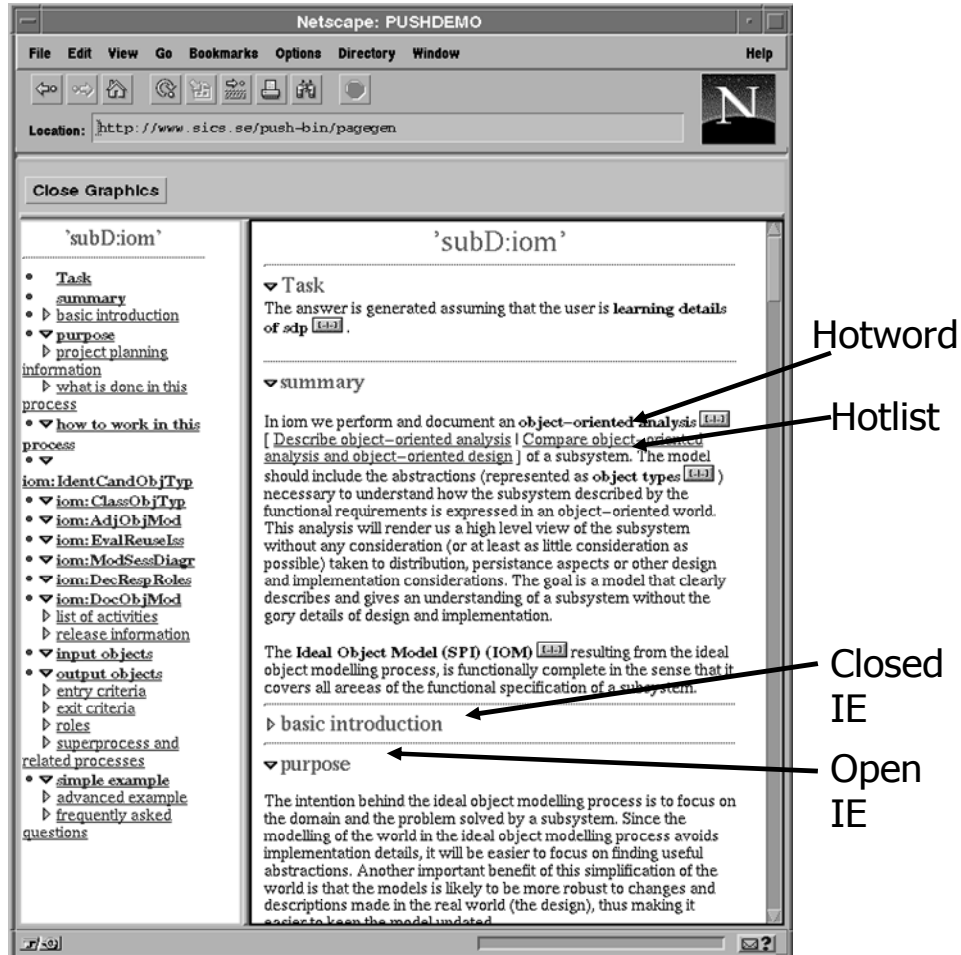


Figure 3: Example of stretchtext based adaptation in the Push system [25]

Stretchtext preserves focus by hiding the less relevant content when a page is first presented. In contrast, dimming and scaling fragments preserve focus by *deemphasizing* rather than hiding the less relevant content. Dimming deemphasizes content by fading its color [6]. Scaling deemphasizes by reducing size. While the existing incarnations of dimming equally deemphasize all the less relevant content, in scaling, size increases as a function of the Degree of Interest (DOI) of each content fragment for the user [46], as assessed by a similarity measure between vectors representing the user's focus of interest and the content of each available fragment. In this respect, scaling is a variation of the fisheye visualization technique, with the following differences:

- In a traditional fisheye visualization, there is a unique focal point, based on the user's current focus of attention. In scaling, there can be multiple focal points based on the user's focus of interest.
- In traditional fisheye views, DOI (and thus content size) decreases with *geometrical* distance from the focal point, while in scaling, DOI decreases with *semantic* distance.

Figure 4 shows an example of scaling from [46]. Here the user's focus of interest is assessed to be theater, thus paragraphs with different degree of relation with this topic are presented in different font sizes.

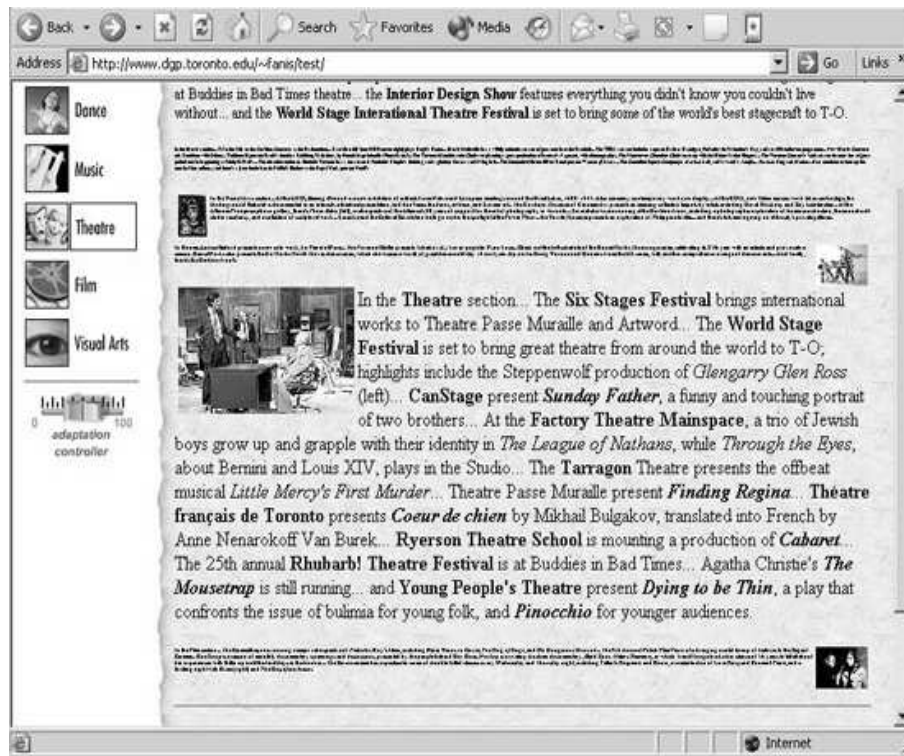


Figure 4: Example of scaling based adaptation from [46]

Although all three techniques in this category do provide more contextual information than those in the Priority on Focus category, their relative placement in the focus-context space is still an open question.

The scaling technique preserves distinctive structural elements of the deemphasized information (such as pictures, layout, number and relative length of paragraphs), and thus has better potential to provide contextual cues than stretchtext. Tsandilas and Shraefel [46] conducted a preliminary study to compare stretchtext and scaling, but the small number of subjects does not allow drawing general conclusions.

They found no difference in terms of task completion time, but identified a potential interaction between technique effectiveness and page size. Because stretchtext generally presents less content, it performed better on large pages, where scaling required more scrolling to access relevant content. However, 4 out of the 6 subjects in the study gave a higher overall score to scaling because they felt it provided better information on the content of the deemphasized paragraphs.

Like scaling, dimming also preserves distinctive structural elements of the available content, but it does not reduce page size (like stretchtext and scaling do). Thus, it may require additional scrolling to access the content of interest, potentially reducing focus. Unfortunately there are currently no studies that directly compare this technique with the other two.

Even as more empirical results become available, practitioners interested in adopting these techniques should be aware that their effectiveness depends on a number of design elements that should be carefully tested before drawing general conclusions on each technique's overall effectiveness. These elements include:

1. Quality of the headers used to indicate the presence of stretchtext ([25] and [43])
2. Cost associated with reading the deemphasized content (e.g., fatigue generated by reading small or faded text)
3. Presence and quality of ways to summarize deemphasized content so that a user can get contextual information without reading it (see for instance the mouse-over glosses described in [46])
4. Presence and effectiveness of mechanisms provided to the user to change content emphasis (e.g., double clicking on deemphasized content to bring it in focus [46])

In relation to the first element above, we describe a technique that, although currently not used in adaptive hypermedia, has an interesting potential for enriching the Priority on Context category. This technique, known as *summary thumbnail* [33] has been devised to address the problem of how to display Web pages designed for desktop-sized monitors on small screen devices. One common way to address this problem is to rescale the page to fit the width of the small screen, thus creating what is known as a *thumbnail view*. The idea is that the user should use this view to rapidly identify the content of interest and then use provided zooming mechanisms to view it. The problem is that often text in the thumbnail view becomes unreadable, forcing the user to resort to zooming to browse the content. Summary thumbnail addresses this problem by adding fragments of readable text to the thumbnails. The text is automatically generated by either removing common words (as defined in a standard word frequency list) or by cropping paragraphs, so as to maintain the total number of lines and overall page layout.

Lam and Baudisch [33] showed that this technique generated better user performance and satisfaction in a browsing task, compared to two other standard techniques for page reduction. This indicates that having readable, although incomplete and possibly not fully coherent text is an effective place holder for hidden content. Thus, this technique could be adapted to be a variation of stretchtext for adaptive content presentation, where place holders for hidden text are summary thumbnails. However, it should be noted that in Lam and Baudisch's experiment [33], all content was equally reduced. In order to verify the applicability of this

technique to adaptive approaches, it will be necessary to test whether it remains effective when summary thumbnails are used in combination with fully-displayed content during adaptive presentation.

3.2 Techniques for Media Adaptation

The previous sections discussed techniques to adaptively select, structure, and present relevant information. Presentation has been addressed, however, only with respect to the problem of how to highlight relevant information and how to possibly allow the user access to relevant context. Here we will address a different form of presentation tailoring: adapting the medium (e.g., text, graphics, spoken language) through which the selected information is conveyed to the user. We begin by discussing factors that can influence a system's choice of media. We then provide illustrative examples of adaptive hypermedia systems that adapt the medium through which information is presented. Finally, we discuss at a conceptual level two common approaches to media adaptation: the rule-based approach and the optimization approach.

3.2.1 Factors Relevant for Media Adaptation

The following is a description of the types of factors that a system may want to consider when deciding how to adapt the media:

- *User-Specific Features*: Relevant user features include preferences, abilities and accessibility issues. Users may have preferences for receiving information in different modalities, for example, a user may explicitly request the information to be presented in a graphical way [48]. In terms of abilities, the user may be better able to reason about information presented using a given medium. For example, if the user has poor language abilities, the material would be better presented visually [21]. As a final example, a relevant accessibility issue would be that a visually-impaired user should not be presented with information visually, but rather through speech (e.g., [17]).
- *Information Features*: Given the presentation goals, certain types of information are better presented using specific media. For example, graphics should be used to highlight quantitative relationships [28], while text should be used if the system wants to convey a precise value, such as the name of a city [48].
- *Contextual Information*: Under certain circumstances it may be appropriate for the system to consider properties of the user's environment when deciding how to best present information. As an example, with systems designed for use in vehicles, variables such as the current weather conditions and other relevant operating conditions (e.g., speed and traffic) should influence both the appropriate media type and quantity of information to be presented [13]. For example, in situations where the weather is poor, or the user is driving in a high speed zone, designers are exploring the use of other media such as haptics and speech to present information to avoid overloading the already loaded visual channel (e.g., [13] and [4]).

- *Media Constraints:* This factor refers to taking into account how the different media should cooperate. For example, a mix of compatible media (e.g., text and speech) can increase the recallability of the information presented [48], and contrasting pieces of information should be presented using the same medium (e.g., [1]and [48]).
- *Limitations of Technical Resources:* This factor relates to the device on which the information will eventually be displayed and the possible media limitations of this device. Examples of resource limitations include the available bandwidth (e.g., large images should not be displayed if there is low bandwidth [17, 32]), the general availability of different media (e.g., whether or not speech is available [41]) or the available screen real estate.

3.2.2. Example Applications

We now describe some examples of systems that aim to perform some sort of media adaptation in an adaptive hypermedia system or adaptive Web site. The first two examples are forms of tourist information systems. The remaining examples are: a mobile navigation system, a learning environment and a pointer to a more general framework for generating multimedia presentations on the Web.

The AVANTI system [17], adapts the media through which the relevant information is presented to the user according to 1) accessibility issues, for example, using spoken language for visually-impaired users, and 2) resources issues, such as not presenting too many graphics in low-bandwidth situations. In [17], to illustrate this type of media adaptation, the authors present two example versions of a page containing information on a city map: one for sighted users and another for visually-impaired users. For visually impaired users, the map is described using text (which could presumably be read by a speech synthesizer), as an alternative to a regular visual map which would be displayed for sighted users.

The MASTROCARONTE system [13], an in-car tourist information system, adapts the medium through which recommendations are made to the user by considering 1) the user-specific factors, such as preferences and the user's current level of fatigue, and 2) contextual factors, such as speed and traffic volume. For example, if the user is requesting restaurant recommendations, the system determines the current level of risk as indicated by the contextual factors (e.g., traffic or visibility) and the user's current level of fatigue to decide a) whether to present the recommendations visually or using speech, and b) how many recommendations to present at one time. If the contextual factors indicate a high-level of risk (e.g., high traffic or low visibility) or if the user is fatigued, the system will elect to present the recommendations visually, and will present only one recommendation at a time (requiring the user to press a button to retrieve additional recommendations). If there is both a high-level of risk and the user is fatigued the system will present the recommendations using speech. If neither condition holds, the system will present several recommendations at once and will present them visually.

In [32], the authors describe guidelines to be used by a mobile navigation system to decide which medium to use when presenting route information. These guidelines

take into account technical resource limitations (the last factor described in the previous section), as well as the user's context and cognitive resources. Route descriptions using text and speech are only appropriate if the system knows the user's location and are useful when the amount of available bandwidth is small (since no images have to be sent to the mobile device). Speech is preferred to text in this scenario if the user is also experiencing a high level of cognitive load since it does not require visual attention. Graphical route instructions, on the other hand, do not require location information, but require more cognitive and technical resources.

The CUMAPH adaptive hypermedia environment [21] adapts hypermedia documents according to a user profile that describes the user's cognitive abilities, such as the user's ability to explore visual and spatial information, and the user's visual and auditory memory. A description of how the profile is acquired can be found in [22]. The environment has so far been applied to an instructional domain involving a course about the human brain. For example, when adapting a web page on the topic of "memory", the system uses the optimization approach described in the next section to decide whether to display the content of two relevant sub items ("definition" and "mechanism") using text, graphics, sound or a combination of two of the three media.

Finally, the hypermedia document formatter described in [41] includes capabilities to adapt media presentation to device characteristics and user preferences. The paper does not illustrate an application of the framework. Instead, the authors focus primarily on the architecture necessary to realize multimedia presentations on the Web.

3.2.3 General Approaches

The most common techniques for media adaptation fall into 2 general categories: 1) *rule-based* or *planning* approaches and 2) *optimization* approaches. We will now describe each type of technique. Once again, we extend our discussion to include systems that are not adaptive web-based systems since we feel that the techniques could certainly be generalized to adaptive Web applications.

Rule-Based

The vast majority of systems that perform media adaptation do so using rules that describe how to best convey the target information given subsets of the factors described in section 3.2.1. Examples of such systems include [2], [3], [16], [17], [19], [23], [35] and [41].

Andre [1] discusses some of the differences between several of the above systems. One primary difference is how integrated the content selection process (as discussed in section 2.2.1) is with the media allocation process (i.e., determining which medium to use for a given information element). The systems also differ in the type and amount of communication among the components in charge of realizing the media-specific information to ensure a coherent and feasible overall result. To illustrate these differences and to provide a more concrete understanding of how media allocation using rules can work, we will now elaborate on two examples: the work by Arens et al. [3] and the WIP system [2].

In [3], Arens et al. describe a system that can adapt the media based on characteristics of the information to be conveyed, media constraints, the user's interests and abilities, and the overall goals of the information presentation (referred to as the "presenter's goals"). The system begins by selecting and structuring the content (using NLG techniques similar those discussed in sections 2.2.1 and 2.2.2), which produces a tree that represents the discourse structure. The system then applies media allocation rules to the discourse structure to obtain a presentation structure. It applies the media allocation rules by traversing the tree in a bottom-up manner. As the rule-application process moves up the tree, earlier media decisions can be reconsidered based on information about the more global structure of the presentation, which is only visible at the higher-level nodes. Such a situation could occur if, for example, content item A is assigned to be presented using text and content item B graphics, but the higher-level discourse goal is to compare the two items. With this higher-level goal, both items should be presented using the same medium. Once the presentation structure is complete, each element is then sent to the appropriate generator, and all of the results are sent to a final layout specialist, which decides how to arrange the results on the target display.

The WIP system, described in [2], interleaves content planning and media assignment. WIP focuses primarily on adapting the media based on characteristics of the information to be displayed, but the authors indicate that their technique could be extended to take additional factors into consideration. Generating the presentation involves a task-decomposition schema-based approach similar to what is often used for content structuring in NLG (see section 2.2.2). To permit media allocation, WIP's schemas (referred to as "presentation strategies") have an additional slot for the medium through which the information should be conveyed. Some schemas specify the appropriate medium, other schemas leave this slot open, to be filled later in the planning process. Using the schemas, the system engages in a top-down planning process, starting from the high-level presentation goal. When more than one schema is applicable, but they have different media assignments, the system selects which schema to use based on meta-rules relating to how well the media accomplishes the current presentation goal. As soon as a media slot is filled, the given presentation goal is sent to the generator in charge of realizing information through that medium. This has the advantage that a given presentation goal can be refined by a generator, if it is not able to fulfill the request. Such a situation may occur if, for example, a generator is asked to display a piece of information using an image, but doing so would require an image that is larger than the amount of remaining screen real estate.

Andre [1] argues for content selection/structuring to be interleaved with media allocation because it allows the media allocation process to inform the discourse structure. However, if one does not want media selection to influence content selection/structuring, or wants to use an existing content selection/structuring process as is, it makes sense to employ the sequential approach. Arens et al. [3] also argue for the simplicity of the sequential approach.

Optimization Approach

An alternative to a rule-based approach is to formulate the media adaptation process as an optimization problem. That is, given information on the relevant factors, the goal is to find the media combination that produces the best overall result. Examples of systems that follow this approach are the RIA system [48] and the CUMAPH system [21]. As discussed in section 2.2.1, RIA optimizes a set of feature-based metrics to perform tailored content selection. In [48] the authors follow a similar optimization approach to adapt the media once the relevant information elements have been selected. The authors present a comprehensive list of feature-based metrics that they use to find the optimal media allocation. These metrics fall into two categories: *media selection* metrics that indicate which media best suits a particular content item (based on user preferences, properties of the user's task, and properties of the content items themselves) and *presentation coordination* metrics that describe the level of coordination among the different media. These metrics are combined in an overall objective function and the media allocation process then becomes the task of maximizing this objective function.

U2 Speech: How much are those?

U3 Speech: How much is this one?

Gesture: Point to a house on the screen

RIA The prices are shown on the screen. **RIA** The asking price is \$499,000.

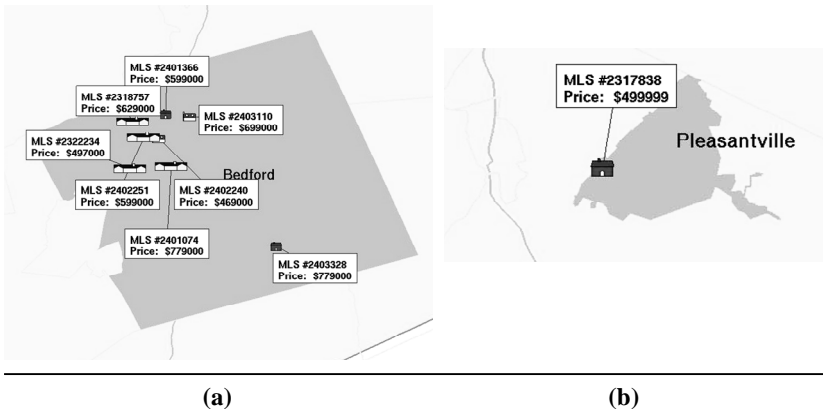


Figure 5: Example of media adaptation in the RIA system.

Figure 5 shows the RIA system responding to two similar user requests (U2 and U3) with different media allocation decisions. In (a), the user is requesting the prices of a number of houses, whereas in (b) the user is requesting the price of a particular house. We use this example to show how the system's optimization procedure is able to deal with three metrics: the suitability of the information to the media, the desire to increase recallability, and the desire to maintain presentation consistency (the first is a media selection metric and the latter are two presentation coordination metrics). In (a) because of the high-volume of data, the price information should be displayed visually and speech can be used only to focus the user's attention to the screen. In

contrast, the low data volume in (b) allows the system to display the price information visually and also to mention the price using speech, which increases recallability. In both examples, because house location, city name and city boundary are considered to be inter-related, all three are presented visually (the house location and city boundary would likely have an especially high compatibility with graphics).

The CUMAPH system [21] follows a similar, but scaled down approach. They have 2 metrics: one assigns the highest value to the media combination that best fits the user profile; the other favors combining multiple media (e.g., using both graphics and text as opposed to just graphics). The system generates all possible combinations of media assignments to information items and picks the one whose sum of the two metrics is the highest. This type of exhaustive generation and testing is likely only possible when the number of combinations is small.

Both of these systems perform media allocation after the content has been selected. Future work in the area could involve investigating ways to use the optimization approach to interleave the two processes.

Comparison of Approaches

According to [47] and [48], the advantages of the optimization approach are: 1) it does not require a large set of rules and/or plans to be authored, 2) it allows the system to handle issues with conflicting or interdependent factors without a large amount of communication among different system components, 3) it is more easily extended, and 4) it is more easily transferred to different domains. A possible downside of the optimization approach is that there do appear to be certain cases where the system is required to repair the media allocation after the optimization process is complete [48]. In the case of [48], these situations are detected and repaired by essentially relying on rules.

We are not aware of any evaluations that have directly compared the rule-based and the optimization approaches. In fact, there has been very little evaluation of systems that perform media adaptation. One exception can be found in [48], where RIA's multimedia presentations were evaluated using the human designer approach (described in Chapter X). The results of the experiment were promising. Out of 50 test cases, RIA was rated (or co-rated) highest in 17 cases. In the 28 of the remaining cases, the difference between the winner and RIA was minor. For the other 5 cases, the authors felt that the problem was due to content selection rather than media selection.

4. Conclusions

In this chapter we discussed techniques for providing adaptive content presentation for the Web. We structured the discussion by first introducing techniques to select and coherently structure the content deemed to be most relevant for the current interaction context, and then by illustrating approaches to further adapt the selected content by tailoring the way in which it is actually presented to the user. In particular, we have

described techniques to adapt the presentation based on content relevance, as well as techniques to adapt the type of media for optimal content delivery.

Adaptive presentation of Web content can complement several of the other types of adaptations presented in this book. For instance, it can be integrated with both adaptive navigation support and adaptive search to further adapt the content of the web pages returned by these two techniques to a user's needs; or it can be coupled with recommender systems to provide tailored descriptions of the recommended items. Although there has been some work on exploiting the integration of adaptive presentation with other types of adaptations (e.g., [8]) we see this as an area that still yields a high potential for innovative research in the adaptive Web. Another area that calls for strong research efforts is the validation of techniques for adaptive presentation. In this chapter we have discussed the results of some evaluations; however, the actual effectiveness of the many of the techniques is still largely open to investigation.

References

1. Andre, E.: The Generation of Multimedia Documents. In: R. Dale, H. Moisl, and H. Somers, (eds.): *A Handbook of Natural Language Processing: Techniques and Applications for the Processing of Language as Text*. Marcel Dekker Inc (2000) 305-327
2. Andre, E., Rist, T.: Generating Coherent Presentations Employing Textual and Visual Material. *AI Review* 9 (1995) 147-165
3. Arens, Y., Hovy, E. H., van Mulken, H.: Structure and Rules in Automated Multimedia Presentation Planning. In: *Proc. of the 13th. International Joint Conference on Artificial Intelligence* (1993) 1253-1259
4. Bellotti, F., Gloria, A. D., Montanari, R., Dosio, N., Morreale, D.: COMUNICAR: Designing a Multimedia, Context-Aware Human-Machine Interface for Cars. *Cognition, Technology & Work* 7 (2005) 36-45
5. Bontcheva, K., Wilks, Y.: Tailoring Automatically Generated Hypertext. *User Modeling and User-Adapted Interaction. Special Issue on Language-Based Interaction*. 15 (2005) 135-168
6. Boyle, C., Encarnacion, A. O.: Metadoc: An Adaptive Hypertext Reading System. *User Modeling and User Adapted Interaction* 4(1) (1994) 1-19
7. Brusilovsky, P.: Methods and Techniques of Adaptive Hypermedia. *User Modeling and User-Adapted Interaction* 6(2-3) (1996) 87-129.
8. Carenini, G.: GEA: A Complete, Modular System for Generating Evaluative Arguments. In: *Proc. of the International workshop on Computational Models of Natural Language Argument* (2001) 959-968
9. Carenini, G., Moore, J.: An Empirical Study of the Influence of Argument Conciseness on Argument Effectiveness. In: *Proc. of the 38th Meeting of the Association for Computational Linguistics* (2000)
10. Carenini, G., Moore, J.: A Strategy for Generating Evaluative Arguments. In: *Proc. of the First International Conference on Natural Language Generation* (2000) 47-54
11. Carenini, G., Moore, J.: An Empirical Study of the Influence of User Tailoring on Evaluative Argument Effectiveness. In: *Proc. of the International Joint Conference on Artificial Intelligence* (2001) 1307-1314
12. Chu-Carroll, J., Carberry, S.: Collaborative Response Generation in Planning Dialogues. *Computational Linguistics* 24(2) (1998) 355-400

13. Console, L., Torre, I., Lombardi, I., Gioria, S., Surano, V.: Personalized and Adaptive Services on Board a Car: An Application for Tourist Information. *Journal of Intelligent Information Systems* 21(3) (2003) 249-285
14. De Bra, P., Aerts, A., Berden, B., Lange, B. d., Rousseau, B., Santic, T., Smits, D., Stash, N.: AHA! The Adaptive Hypermedia Architecture. In: *Proc. of the Fourteenth ACM Conference on Hypertext and Hypermedia* (2003) 81-84
15. Encarnacao, L. M., Stoev, S. L.: An Application-Independent Intelligent User Support System Exploiting Action-Sequence Based User Modelling. In: *Proc. of UM'99, International Conference on User Modeling* (1999)
16. Feiner, S., McKeown, K.: Automating the Generation of Coordinated Multimedia. *IEEE Computer* 24(10) (1994) 33-41
17. Fink, J., Kobsa, A., Nill, A.: Adaptable and Adaptive Information Provision for All Users, Including Disabled and Elderly People. *The New Review of Hypermedia and Multimedia* 4 (1998) 163-188
18. Furnas, G. W.: Generalized Fisheye Views. In: *Proc. of CHI '86, Conference on Human Factors in Computing Systems* (1986) 16-23
19. Green, N. L., Carenini, G., Kerpedjiev, S., Mattis, J., Moore, J. D., Roth, S. F.: AutoBrief: An Experimental System for the Automatic Generation of Briefings in Integrated Text and Information Graphics. *International Journal on Human-Computer Studies* 61(1) (2004) 32-70
20. Guo, H., Stent, A.: Trainable Adaptable Multimedia Presentation Generation. In: *Proc. of the 7th International Conference on Multimodal Interfaces* (2005)
21. Habieb-Mammar, H., Tarpin-Bernard, F.: CUMAPH: Cognitive User Modeling for Adaptive Presentation of Hyper-Documents. An Experimental Study. In: *Proc. of AH 2004, Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (2004) 136-145
22. Habied-Mammar, H., Tarpin-Bernard, F., Prevot, P.: Adaptive Presentation of Multimedia Interface Case Study: "Brain Story" Course. In: *Proc. of UM'03, International Conference on User Modelling* (2003) 15-24
23. Han, Y., Zukerman, I.: Using Cooperative Agents to Plan Multimodal Presentations. In: *Proc. of Multimodal Human-Computer Communication* (1995) 122-157
24. Henze, N., Nejd, W.: Extendible Adaptive Hypermedia Courseware: Integrating Different Courses and Web Material. In: *Proc. of the International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (2000) 109-120
25. Hook, K.: Evaluating the Utility and Usability of an Adaptive Hypermedia System. In: *Proc. of IUI'97, International Conference on Intelligent User Interfaces* (1997) 179-186
26. Hook, K.: Steps to Take before Intelligent User Interfaces Become Real. *Interacting with Computers* 12 (2000) 409-426
27. Jameson, A.: Adaptive Interfaces and Agents. In: J. Jacko and A. Sears, (eds.): *Human-Computer Interaction Handbook*. Erlbaum (2003) 305-330
28. Kerpedjiev, S., Carenini, G., Green, N., Moore, J., Roth, S.: Saying It in Graphics: From Intentions to Visualizations. In: *Proc. of the IEEE Symposium on Information Visualization* (1998) 97-101
29. Knott, A., Dale, R.: Choosing a Set of Coherence Relations for Text Generation: A Data-Driven Approach. In: *Trends in Natural Language Generation: An Artificial Intelligence Perspective* (G. Adorni & M. Zock, Eds.). Springer-Verlag (1996) 47-67
30. Kobsa, A., Koenemann, J., Pohl, W.: Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships. *The Knowledge Engineering Review* 16(2) (2001) 111-155

31. Korb, K., McConachy, R., Zukerman, I.: A Cognitive Model of Argumentation. In: *Proc. of Proc. of the Nineteenth Annual Conference of the Cognitive Science Society* (1997) 400-405
32. Kray, C., Laakso, K., Elting, C., Coors, V.: Presenting Route Instructions on Mobile Devices. In: *Proc. of IUI'03, International Conference on Intelligent User Interfaces* (2003) 117-124
33. Lam, H., Baudisch, P.: Summary Thumbnails: Readable Overviews for Small Screen Web Browsers. In: *Proc. of CHI 2005, Conference on Human Factors in Computing Systems* (2005) 681-690
34. Masthoff, J., Pemberton, L.: Adaptive Hypermedia for Personalized TV. In: S. Chen and G. Magoulas, (eds.): *Adaptable and Adaptive Hypermedia Systems*. IDEA group publishing (In Press)
35. Maybury, M.: Planning Multimedia Explanations Using Communicative Acts. In: M. Maybury, (ed.) *Intelligent Multimedia Interfaces*. AAAI Press ' The MIT Press (1993) 60-74
36. Melis, E., Andres, E., Franke, A., Frischauf, A., Goguadse, G., Libbrecht, P., Pollet, M., Ullrich, C.: Activemath: A Web-Based Learning Environment. *International Journal of AI in Education* 12 (2001) 385-407.
37. O'Donnell, M., Mellish, C., Oberlander, J., Knott, A.: ILEX: An Architecture for a Dynamic Hypertext Generation System. *Journal of Natural Language Engineering* 7(3) (2003) 225-250
38. Paris, C., Wan, S., Wilkinson, R., Wu, M.: Generating Personal Travel Guides - and Who Wants Them? In: *Proc. of UM'01, International Conference on User Modeling* (2001) 251-253
39. Reiter, E., Dale, R.: *Building Natural Language Generation Systems*. Cambridge University Press (2000)
40. Reiter, E., Robertson, R., Osman, L.: Types of Knowledge Required to Personalise Smoking Cessation Letters. In: *Proc. of the Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making (AIMDM'99)* (1999) 389-399
41. Rodrigues, R. F., Rodrigues, P. S. L., Feijo, B., Velho, L., Soares, L. F. G.: Cross-Media and Elastic Time Adaptive Presentations: The Integration of a Talking Head Tool into a Hypermedia Formatter. In: *Proc. of AH 2004, Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (2004) 215-224
42. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. second edn. Morgan-Kaufman, Los Altos, CA (2003)
43. Schraefel, M. C.: Contexts: Adaptable Hypermedia. In: *Proc. of AH 2000, International Conference on Adaptive Hypermedia and Adaptive Web-based Systems* (2000) 369-374
44. Sowa, J.: *Conceptual Structures: Information Processing in Mind and Machine*. Addison Wesley (1984)
45. Stent, A., Prasad, R., Walker, M.: Trainable Sentence Planning for Complex Information Presentation in Spoken Dialog Systems. In: *Proc. of 42nd Annual Meeting of the Association for Computational Linguistics* (2004)
46. Tsandilas, T., schraefel, m.: Usable Adaptive Hypermedia. *The New Review of Hypermedia and Multimedia* 61(6) (2004) 5-29
47. Zhou, M. X., Aggarwal, V.: An Optimization-Based Approach to Dynamic Data Content Selection in Intelligent Multimedia Interfaces. In: *Proc. of UIST'04, the 17th Annual ACM Symposium on User Interface Software* (2004) 227-236
48. Zhou, M. X., Wen, Z., Aggarwal, V.: A Graph-Matching Approach to Dynamic Media Allocation in Intelligent Multimedia Interfaces. In: *Proc. of IUI'05, International Conference on Intelligent User Interfaces* (2005) 114-121

49. Zukerman, I., McConachy, R., Korb, K.: Using Argumentation Strategies in Automated Argument Generation. In: *Proc. of the First International Natural Language Generation Conference* (2000) pp. 55-62