THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Dynamic Hypertext Catalogues: Helping Users to Help Themselves

Citation for published version:

Milosavljevic, M & Oberlander, J 1998, Dynamic Hypertext Catalogues: Helping Users to Help Themselves. in *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia: Links, Objects, Time and Space---structure in Hypermedia Systems: Links, Objects, Time and Space---structure in Hypermedia Systems*. HYPERTEXT '98, ACM, New York, NY, USA, pp. 123-131, The Ninth ACM Conference on Hypertext and Hypermedia (Hypertext '98), Pittsburgh, PA, United States, 2/06/98.
<https://doi.org/10.1145/276627.276641>

Digital Object Identifier (DOI):

[10.1145/276627.276641](https://doi.org/10.1145/276627.276641)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Dynamic Hypertext Catalogues: Helping Users to Help Themselves

Maria Milosavljevic

CSIRO Mathematical and Information Sciences
Locked Bag 17
Sydney NSW 2113 Australia
Email: Maria.Milosavljevic@cmis.csiro.au

Jon Oberlander

Human Communication Research Centre
University of Edinburgh
Edinburgh EH8 9LW Scotland
Email: J.Oberlander@ed.ac.uk

ABSTRACT

Electronic hypertext catalogues provide an important channel for information provision. However, static hypertext documents cannot be dynamically adapted to help the user find what he/she is looking for. We demonstrate that natural language generation techniques can be used to produce tailored hypertext documents, and we focus on two key benefits of the resulting DYNAMIC HYPERTEXT. First, documents can be tailored more precisely to an individual's needs and background, thus aiding the search process. Secondly, the incorporation of techniques for comparing catalogue items allows the user to search still more effectively. We describe the automatic generation of hypertext documents containing comparisons, with illustrations from two implemented systems.

KEYWORDS: adaptive hypertext, dynamic hypertext, natural language generation, user modelling, discourse history

INTRODUCTION

The advent of on-line distributed hypertext systems and the world wide web (WWW) has led to the extensive popularisation of electronic hypertext documents. As a result, we no longer need to leave home in order to browse through a library, buy our groceries or visit a museum. However, the static nature of most hypertext documents limits their usefulness, since they cannot be tailored to any particular user's requirements or abilities. Instead, document authors construct general-purpose documents which are written according to a wide audience model, or else construct (and update) multiple documents for users' anticipated needs.

By its very nature, hypertext invites user interaction with documents. It is possible for hypertext systems to alter what information is available to a user, by exploiting information about the user-type or session history. Nonetheless, there are limits to the flexibility that current methods afford; we overcome these by incorporating natural language generation techniques into such systems, allowing each document to

be created with an individual user in mind. The resulting DYNAMIC HYPERTEXT documents are tailored on the basis of a model of the user's knowledge or the user's previous interactions with the system.

Comparison plays a central role in both *learning* and *choosing*, two of the central affordances of electronic catalogues. Shopping catalogues allow choices to be made by the buyer; other catalogues, such as encyclopædias or museum guides, help the user to learn something. The effective support of comparisons is therefore essential in any dynamic hypertext system which aims to help users to help themselves, in searching for information or in learning something new.

By utilising techniques from natural language generation systems, dynamic hypertext systems automatically construct entire hypertext networks and the nodes (or documents) of that network at run-time, and adapt these to an individual's knowledge and needs. The textual content of the documents is constructed from an underlying knowledge base of facts, rather than from larger pieces of canned text. Thus, each user may be presented with a highly personalised catalogue which can make references to, or comparisons with, other catalogue items if these are relevant in the current discourse context.

In this paper, we outline the architecture and benefits of dynamic hypertext systems, and emphasise the importance of comparison in these systems; we argue that, by making more effective use of a user model and the discourse history, NLG techniques can offer highly flexible hypertext documents. We illustrate the advantages of such dynamic hypertext techniques through examples from two implemented systems: the PEBA-II and the ILEX text generation systems, which dynamically produce descriptions of entities as WWW pages.

DYNAMIC HYPERTEXT GENERATION

Natural Language Generation

NATURAL LANGUAGE GENERATION aims to produce coherent natural language text from an underlying representation of knowledge. It can be viewed as a goal-driven planning process, involving the formulation of texts that satisfy some communicative goal. Many of the ideas here are borrowed from conventional planning techniques developed within artificial intelligence; so, for example, a top level communicative goal such as 'instruct the user how to operate a telescope' may be decomposed into a number of constituent sub-goals such as 'tell the user what a lens is', 'tell the user where

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HyperText 98 Pittsburgh PA USA

Copyright ACM 1998 0-89791-972-6/98/6...\$5.00

Figure 1: System architectures: (a) traditional NLG; (b) dynamic hypertext.

the focusing mechanism is' and so on. This decompositional process iterates until the resulting goals can each be realised by means of a natural language utterance.

Figure 1a shows the traditional architecture of NLG systems. NLG systems embody two main processing components: the text planner and the surface realisation component.

The text planning stage typically encapsulates all those decisions involving choices of *what to say*. Based on the discourse goals, the text planner must decide what is relevant in a particular situation (*content selection*), and then organise this content in a way that allows realisation of a coherent discourse that guides the reader's inferences (*text organisation*). The text planning component achieves this by composing a discourse plan using facts from the knowledge base (KB). For example, McKeown's schema-based approach stores a number of plan outlines in a plan library and fills in the appropriate information from the KB [13].

A model of the user's knowledge can be used by an NLG system to tailor the text to the individual, as in Paris's work [21, 22]. In addition, Moore has shown that the ongoing discourse with each particular user can also be recorded in the discourse history component to enable the system to adapt future texts to what has been said before [17, 18].

The discourse plan is realised as natural language utterances by the surface realisation component. This makes use of knowledge of the language's grammar and lexicon to produce well-formed utterances conveying the required content.

Dynamic Hypertext

DYNAMIC HYPERTEXT is an area of research within NLG which takes advantage of hypertextual interaction to give the user the freedom to perform high-level discourse planning, thereby

reducing the burden on the NLG system of having to reason more deeply about his/her goals. A key element in any dynamic hypertext system is that the hypertext network and the nodes of this network (the documents themselves) are *ynamically* created at run-time when the user requests them; there are no existing hypertext documents, and there may not even be any pre-existing representations of what *could* be documents within the system.¹ This is in contrast both to Ted Nelson's original notion of "stretch text", and to recent work on ADAPTIVE HYPERTEXT systems.

Instead of presenting the user with a new document each time a link is selected, systems utilising "stretch text" (such as [10]) allow for the insertion of additional text into the *current document*, immediately following (or replacing) the selected hyperlink. In this case, however, there is still a fixed network of hypertext documents, and what is displayed at any time is under direct user control.

Adaptive hypertext systems require a fixed network of hypertext documents, but utilise a user model in order to control what is presented to users, according to their knowledge of the concepts within any particular document. For instance, the various versions of Trellis [25] are based on a Petri-net model of hypertext, making the presence of links conditional on user actions; this controls the accessibility of individual nodes, or entire classes of nodes. Some adaptive approaches have relied more directly on artificial intelligence techniques. For example, Simon and Erdmann's SIROG [24] utilised an expert system to select contextually appropriate sections of on-line help. In Kaindl and Snaprud's [9] work, each hyper-

¹Such systems are therefore instances of the type of integration between hypermedia and artificial intelligence foreseen by Halasz [5], and indeed, they move further towards virtual structures than then seemed feasible, since static structures are entirely replaced.

The Echidna

The Echidna, also known as the spiny Anteater, is a type of Monotreme that is covered in stiff, sharp spines mixed with long, coarse hairs.

The Echidna has the following subtypes:

- the short-beaked Echidna and
- the long-beaked Echidna.

The Echidna is about the same length as a domestic cat. It ranges from 2 kg to 7 kg in weight. It has a brownish black coat and paler-coloured spines. It has a small head. It has a prolonged, slender snout. It has no teeth. It uses its extensible, sticky tongue for catching ants, termites and other small insects. It is a carnivore and eats ants, termites and earthworms. It has powerful claws allowing for rapid digging of hard ground. It is found in Australia. It is active at dawn and dusk. It lives by itself. It has an average lifespan in captivity of 50 years.

This text is generated for the novice user level. If you would like the text for the expert user level click [here](#).

Figure 2: A description of the echidna produced by PEBA-II.

text node is represented as an AI frame, and the text of each node is stored in a slot of its frame. The textual component may contain links to other hypertext nodes, and can be partitioned into different sections which may be revealed to the user as required, giving rise to “active text”.

For a survey of existing adaptive hypertext systems and further elaboration of the concepts involved, see [1]. Using adaptive techniques, a system can present different views of the same document to different users. However, this flexibility is still fundamentally limited to choosing among the author’s pre-written text segments.

By contrast, a dynamic hypertext system goes further, because it operates in a similar fashion to traditional NLG systems. Figure 1b shows the architecture of a dynamic hypertext system; the traditional NLG architecture shown in Figure 1a is augmented with some additional components which are required within a hypertext environment. A KB contains information about those concepts in the domain, and the system selects which elements of the KB are important for creating the required hypertext document. The surface realisation component of a dynamic hypertext system must also mark the hypertext anchors into the text in order to produce a document which can be viewed using the target hypertext interface. The hyperlinks represent follow-up questions which the user can ask, and are generally concepts (or other entities) that can be described by the system. In operation, the user can effectively perform high-level discourse planning, driving the system by selecting hypertext links. Each hyperlink indicates a new discourse goal to the system.

A survey and comparison of existing dynamic hypertext systems can be found in [6], and further discussion on the advantages of such systems is provided elsewhere [12, 3]. In the next section, we look more closely at two particular systems we have been involved with.

TWO DYNAMIC HYPERTEXT SYSTEMS

The PEBA-II System

PEBA-II² is an NLG system which produces descriptions and comparisons of entities represented in a taxonomic knowledge base. An overview of PEBA-II will be presented in this section; more details are available in [15]. Figure 2³ contains an example description produced by PEBA-II; note that the underlined words or concepts indicate hypertext links which, when clicked on, will result in new hypertext documents being dynamically produced.

The architecture of the PEBA-II system is essentially that shown in Figure 1b; the system has two types of discourse goals: to describe a single entity or to compare two entities. In realising these goals, the system makes use of a user-type, a user model and the discourse history. At the beginning of each interaction with the system, the user is permitted to classify him/herself as either a naïve or an expert user-type; this choice results in the system taking different views of the structure and content of the KB of animal facts. The user can also enter her details into a more specific user-model which is used to make inferences about his/her specific knowledge and to draw comparisons with similar or familiar entities. The discourse history is used to tailor the output to take account of the previous discourse, as described further below.

PEBA-II uses a phrasal lexicon. That is, the mapping from KB elements to surface form varies from being single words (such as *Yak*) to short-phrases (such as *lifespan in captivity*) to full phrases (such as *has a long shaggy coat which hangs to the ground like a fringe*). In other words, we only decompose those concepts which require linguistic variation. Our aim here is two-fold: to populate a KB while avoiding difficult representational issues; and to produce a system which is efficient, not rebuilding the same surface form many times (ILEX achieves a similar goal by using annotated stories).

²A version of the PEBA-II system is available on the WWW at URL: <http://www.mri.mq.edu.au/~peba/>.

³In all figures, the underlined words are hypertext links.

Pair of brooches on mount

Silver, gold, mahogany, walnut and perspex

This item was made in 1979 and is made of silver, gold, mahogany, walnut and perspex. It was designed by Martin Page who was English. Like the necklace designed by Flockinger, this item is in the Organic style. Organic jewels tend to be coarsely textured. However, this item has smooth surfaces.

With a piece like this, the boundary between 'jewellery' and 'sculpture' or 'art' starts to become quite indistinct. One important theme across 20th Century jewellery has been what to do with a piece of jewellery when it is not being worn. From the 1970s onwards, jewellers have started exploring the idea of turning jewellery into sculpture—so that you can hang it on the wall, or prop it on the mantelpiece when you are not wearing it. This piece works equally well whether it is being worn or being displayed (as at present).

Other jewels in the organic style include:

- a pendant necklace designed by Björn Weckstrom
- the necklace designed by Flockinger
- a bracelet designed by Flockinger
- a finger ring designed by Frances Beck
- a finger ring designed by Jacqueline Mina
- the previous item
- a finger ring designed by Ernest Blyth

Figure 3: A description produced by ILEX-1.2.

The ILEX System

The Intelligent Labelling Explorer project has built the ILEX system,⁴ which uses NLG technology to generate descriptions of objects displayed in a museum gallery. An overview of ILEX is given here; for more details, see [6, 8]. To date, three versions have been implemented (ILEX-0, -1, and -2); these systems describe objects in the National Museums of Scotland's 20th Century Jewellery Gallery. Both field and laboratory evaluation are currently underway.

Visitors start from a visual index, composed of thumbnail images. Clicking an image causes a description of the relevant object to be generated. There is no separate introductory article, since background material is incorporated into the descriptions generated on demand (on the relative usability of indices and introductions in more conventional museum hypertexts, see Shneiderman et al. [23]). Figure 3 shows an example output from ILEX-1.2. Although a description contains suggested onward links, the visitor can at any time return to the visual index, and choose a new object from there.

As well as being accurate artefact descriptions, ILEX's labels must convey information which is: *important*, in the sense of helping educate the visitor more broadly; and *interesting*, since when the descriptions are boring, the visitor can just walk away. To help meet these criteria, ILEX uses simple user-types, a discourse history, and its own SYSTEM AGENDA of communicative goals. Thus, the user has freedom to explore any object in the gallery at any time, potentially

making completely unanticipated connections between objects; however, the descriptions produced are constrained via the system's agenda of educational goals, which it strives to achieve when the opportunity arises.

ILEX-2's user-typing is very simple, allowing descriptions to be tailored to different lengths of visit, and different special interests. Figure 4 shows two descriptions of a necklace which are tuned to different user interests by ILEX-2.0 (note that images and titles have been omitted to save space). The first is for a user interested in styles; the second for a user interested in designers. The obvious differences lie in the initial paragraphs: the first description does not even mention the designer by name.

The architecture of the ILEX system differs from that in Figure 1b in only minor respects. First, the system agenda is an additional resource encoding a set of communicative goals, separate from those posted by the user. Secondly, the KB has two main parts: information parsed from the museum's (very large) database, and text either entered by hand, or acquired from interviews. Hand-entered information includes type hierarchies for jewels and designers, and this is used for fully-fledged text generation, along with the database information. Interview-based information starts out with canned text STORIES, extracted from interviews with the curator. These stories typically concern individual jewels or classes of jewels, and can be used as part of, among other things, ARGUMENTS and MISCONCEPTION CORRECTIONS. Significantly, however, the canned text used for these stories can be marked up with various degrees of ANNOTATION, wherever linguistic variation may be required, as in PEBA-II.

⁴Versions of the ILEX system are available on the WWW via URL: <http://cirrus.dai.ed.ac.uk:8000>.

This jewel is a necklace and is in the Organic style. It was made in 1976. It is made from opals, diamonds and pearls.

Organic style jewels usually draw on natural themes for inspiration (for instance, this jewel uses natural pearls). Organic style jewels are usually encrusted with jewels. To take an example, this jewel has silver links encrusted asymmetrically with pearls and diamonds.

This jewel is a necklace and was made by Gerda Flockinger, who was a designer and was English. The jewel, which is in the Organic style, was made in 1976.

Organic style jewels usually draw on natural themes for inspiration; for instance, this jewel uses natural pearls. Organic style jewels are usually encrusted with jewels; for instance, this jewel has silver links encrusted asymmetrically with pearls and diamonds.

Figure 4: Text from two user-tailored descriptions of a necklace produced by ILEX-2.0.

DISCOURSE CONTEXT AND COMPARISON

Dynamic hypertext generation can be viewed as a limited form of dialog between a system and the user where the user's click on a hypertext link represents a *follow-up question* [19]. However, the system must reason carefully about what information should be pointed to from the current discourse context (the current document) in order to allow the user freedom of exploration. The system must also decide what information to include in the hypertext document and what information to collapse into a link. Such systems can provide particularly close couplings between the anchors and destinations of links, dynamically changing the text at both the departure and arrival points to maximise the coherence between them. They thus effectively address the issues raised by Landow concerning the rhetoric of departure and arrival [11]; in fact, for us, navigation really is conversation.

Electronic catalogues consist of descriptions of *concepts* and *entities*. Comparison is an important part of the description process, as well as an effective tool for introducing new concepts or entities to a reader. It is also an important aspect of choice: for example, in deciding which car to purchase. It is therefore important in electronic catalogue descriptions, since the purpose of these is to either teach the user something new (as in on-line encyclopædias or museum guides) or to help the user make an informed choice (as in shopping guides). By investigating the use of comparisons in electronic encyclopædia articles, we have been able to build the same functionality into our systems.

This section next outlines three types of comparison which were identified in an analysis of on-line hypertext encyclopædia articles. We then turn to ways in which comparison can be harnessed within the current discourse context, allowing dynamic hypertext systems to help the user to help themselves, and bypassing the shortcomings of the overload or insufficiency of information in a hypertext network.

The Types of Comparison

Milosavljevic and Dale describe three types of comparisons found in electronic encyclopædia systems [14]:

- A *direct comparison* is *bi-focal*: neither of the two entities being compared is more central to the discourse. For example, we might distinguish the two subtypes of a class such as: *There are two kinds of camels: the dromedary, or Arabian camel, which has one hump, and the Bactrian camel, which has two humps* [4]. We also find more lengthy forms of direct comparison when two entities are particularly similar,

such as the rabbit and hare. The PEBA-II system produces direct comparisons of animals when a user specifically asks for a comparison of two similar entities, by comparing each of the property types (eg. length, colouring) which the two animals share, in a point-by-point manner. Figure 5 provides an example direct comparison produced by PEBA-II.

- A *clarificatory comparison* is produced when, in describing an entity (the *focused entity*), there is another entity which is either: (i) a *potential confusor* of the focused entity and hence needs to be distinguished; or (ii) known by the reader and shares a number of salient features with the focused entity and hence makes a good *comparator* by which to describe the new entity. A clarificatory comparison is *uni-focal* since the focused entity is more important than the potential confusor. Figure 6 shows a description of the echidna produced by the PEBA-II system which contains a clarificatory comparison with the porcupine. This comparison is used for users who are either very familiar with the porcupine (and therefore their existing knowledge is used to better enable them to understand the description of the echidna), or who may confuse the echidna with the porcupine. ILEX currently produces comparisons such as: *Like the Jesse King Necklace, this item is in the arts-and-crafts style. However, it is made from silver and enamel rather than gold and enamel.*

- An *illustrative comparison* is useful when there is a commonly known entity (or an entity specifically known by this user) which shares a particular property with the focused entity, and hence can be used to illustrate the property. An example illustrative comparison produced by the PEBA-II system is found in Figure 2: *The Echidna is about the same length as a domestic cat.* ILEX produces comparisons with previously described items such as: *Like the previous item, this necklace is in the Organic style.*

Comparisons in Context

In order to meet the user's needs, a dynamic hypertext system must provide the user with full freedom to explore a hyperspace. In addition, by POINTING BACKWARDS, a document can relate the current context to previously-covered information (see the next section); by POINTING FORWARDS to information which might be relevant or interesting to the particular user, the system can guide the user's traversal of information in the physical or virtual space.

When learning about unfamiliar items in the domain, the user will augment their existing knowledge with new knowledge. Consequently, when describing a new concept to the user, we should build on their existing knowledge and refer to

The Rabbit and the Hare

The Rabbit is a member of the Leporidae Family that has very helpless young which are born naked and with closed eyes. The Hare is a member of the Leporidae Family that has young which are born furred and with open eyes.

Some comparisons:

- Like the Hare, the Rabbit has long ears.
 - Like the Hare, the Rabbit feeds on herbs, tree bark and vegetables.
 - Like the Hare, the Rabbit has a short, upturned tail.
 - The Hare is longer than the Rabbit.
 - The Hare weighs twice as much as the Rabbit.
 - The Rabbit lives in underground burrows whereas the Hare lives in a simple nest.
 - The Rabbit lives in colonies whereas the Hare rarely lives socially.
-

Figure 5: A direct comparison of the rabbit and hare produced by PEBA-II.

concepts with which he/she is familiar, in order to ease both the description task for the system and the understanding of the new concept for the user. Note, however, that within a hypertext document, the compared concept can be a hypertext link, and hence if the user is interested, he/she can request further information about the comparator. The important question here, then, is whether the user has to have knowledge of the comparator in order to make a comparison.

In order to provide hypertext links from the current hypertext document (describing a focused entity) to other entities (or to other potential hypertext documents describing other entities), we need to assess object similarity.⁵ However, *similarity* does not operate without a context (or respect): the statement *a is similar to b* does not simply mean that entities *a* and *b* share a large number of properties, since every entity shares an infinite number of properties with every other entity;⁶ rather, this statement means that *a* is similar to *b* in a certain respect, *r* [7]. For example, *children* and *photos* do not appear to be very similar concepts; however, under a context such as *items to be saved from a fire*, they are highly related. When computing the similarity of concepts in a particular setting, we need to specify what *r* is: for example, in domains consisting of concrete objects (such as animals or jewellery), *r* might be simple attributes such as length or colour or more complex notions such as appearance, elaborateness or size.

In any particular domain, the current context can be taken to be either the user's interest (at the beginning of an interaction, this might be all we can rely on), the currently focused item, or a combination of the two. A dynamic hypertext system can then incorporate into the current document links to related concepts or entities in the domain by taking account of this context. The context is mapped to a representation of *r* which is then used to compute object similarity with other items of potential interest.

The nature of *r* will vary depending on the current context.

⁵Recall that we are only considering documents describing one particular entity or concept.

⁶For example, if entity *a* is a necklace and entity *b* is a ring, then both entities weigh less than 1kg, 2kg, 3kg and so on.

That is, the computation of object similarity might boil down to a simple attribute match in cases where we consider such things as *animals with scales*, *black sleeveless dresses* or *jewels made by Jesse King*, and the comparisons formed from such instances will thus be simple illustrative comparisons such as: *In the next case there is another jewel which was made by Jesse King*. On the other hand, in situations where we require more complex object similarity matching according to such things as appearance or elaborateness, the resulting comparison will be clarificatory since we will need to specify *why* the objects are similar (it is not a case of simple attribute matching) and then *delimit* where the similarity ends [16]. Any particular context will reveal similarities between objects which might not otherwise be considered similar. For example, under the context *animals with scales*, we might consider *fish* to be similar to *armadillos*.

We have developed a similarity metric which takes into account the different types—and the non-binary nature—of similarity. This allows us to determine when properties which are not identical are, in fact, similar: for example, the material *silver* is more similar to *gold* than it is to *paper*, and a body covering of *soft fur* is more similar to *wool* than it is to *spines*. This allows us to more accurately determine the similarity of objects than simple feature-matching metrics such as Tversky's [26].

To illustrate the utility of this, imagine a person approaches an electronic catalogue system and enters their interest as something like *I am interested in music*. Such a system might be in a museum, in a department store or in a library. A museum-based system can then determine which items in the museum might be relevant or interesting to the user, and propose a path and associated story through the museum; that user's tour would be personalised. Similarly in a library or department store, the system can direct a user to items of potential interest. As the visitor "walks" through the information space (whether physical or virtual), the system can adapt the tour by taking into account those items which are, in fact, interesting to him/her, assessing the similarity of other items to a changing model of his/her interests. Concepts and entities will be represented in a *virtual space* for use by the system, and may or may not also be located in some *physical space* in

The Echidna

The Echidna, also known as the spiny Anteater, is a type of Monotreme that is covered in stiff, sharp spines mixed with long, coarse hairs. Although it is similar in appearance to the African Porcupine it is not closely related. The African Porcupine is a type of Rodent that has long sharp spines, up to 50cm long, which cover its whole back and can be raised by muscles under the skin. Like the African Porcupine, the Echidna has a brownish black coat and paler-coloured spines. The African Porcupine is twice the length of the Echidna (80.0 cm vs 47.5 cm). The Echidna has an average weight of 4.5 kg whereas the African Porcupine has an average weight of 25.0 kg. The Echidna is a carnivore and eats ants, termites and earthworms whereas the African Porcupine is a herbivore and eats leaves, roots and fruit.

The Echidna has the following subtypes:

- the short-beaked Echidna and
- the long-beaked Echidna.

The Echidna has a small head. It has a prolonged, slender snout. It has no teeth. It uses its extensible, sticky tongue for catching ants, termites and other small insects. It has powerful claws allowing for rapid digging of hard ground. It is found in Australia. It is active at dawn and dusk. It lives by itself. It has an average lifespan in captivity of 50 years.

This text is generated for the novice user level. If you would like the text for the expert user level click here.

Figure 6: A description with clarificatory comparison produced by PEBA-II.

a museum or department store. The interaction between these spaces is particularly interesting since we can lead the visitor or buyer to items of potential interest which are in close proximity [20]. The constraints posed by a physical layout will impact the choice of comparators which are linked into the unfolding discourse, both those to which we point back (the most recently visited items) and those to which we point forward, which should be the items in closest proximity. Additionally, for the purposes of illustrative comparison, we might make use of items which are currently perceivable by the user—instead of items visited previously—in order to make the description more immediately accessible.

Maximising Coherence Through Past Discourse

The importance of being able to refer back (or POINT BACKWARDS) to previous conversations with a listener, or more specifically, to concepts previously described to a listener, is obvious when we enter a conversation or debate which has been going for some time, or when we try reading a book from the middle, or start watching a movie which is already half-way through. Although initially the new environment will be thoroughly confusing, we can often pick up much of what we missed by listening for the use of discourse history in the conversation, debate, book or movie. The discourse history plays a central role in NLG systems, particularly when constructing anaphoric referring expressions such as *it* or *red ones* [2], or when trying to refer to concepts mentioned in past communication with a user [17].

A dynamic hypertext system which is capable of producing comparisons can make considerable use of the discourse history. Within the local textual environment, the system can reinforce the relationship between the previous and current documents by linking the ways in which concepts are described in each; we refer to this as TEXTUAL COHERENCE.

In the more global discourse context, the system can build on those concepts previously described, in a similar way to employing a user model; we refer to this as CONCEPTUAL COHERENCE.

Textual Coherence. The path by which a user might arrive at a particular hypertext document cannot always be predicted in advance. Thus, making comparisons with the concepts described in the most recent document is an effective way of making more coherent the transition from one document to the next—it is part of the rhetoric of arrival, in Landow's terms [11]. This cannot be achieved by means of the simpler annotations used to mark up documents in straightforward adaptive hypertext systems.

As described earlier, the automatic generation of hypertext documents effectively constitutes an ongoing discourse with the user, and hence the description should provide a connection to the most recently described concept. Both the PEBA-II and the ILEX systems link the focused entity to the most recently described entity. For example, the description of the Echidna in Figure 2 is discourse-initial (that is, the first description for the user); but if the user had arrived at the Echidna document from the *short-beaked Echidna* document, then the PEBA-II system would have produced the following linking text within the resulting description of the Echidna: *Apart from the short-beaked Echidna, the Echidna has the following subtype: the long-beaked Echidna.* The ILEX system currently generates linking sentences such as: *Like the brooch designed by Page, this item is in the Organic style.*

ILEX uses these linking sentences to liken the current piece of jewellery to a previously described piece in order to bring the similarities to the user's attention. There are several al-

ternative ways of making the discourse more connected, but the difficulty lies in finding the most relevant relationship between the entities described in the hypertext documents. Except where a document has been reached from the top of the system (both PEBA-II and ILEX have introductory indices where the user can select any object), in most instances there should be some characteristic by which to link the two documents, since the current document is likely to have been listed in the text of the previous document for some reason. It is often this very reason which provides us with some clue as to what to use in the linking sentence.

Conceptual Coherence. A dynamic hypertext system can make use of the longer-term discourse history by making comparisons with those concepts which have been described to the user in the past discourse; recall that we call this *pointing backward*. For example, in describing the porpoise to a listener, it is often easier to compare it to the dolphin using a clarificatory comparison. If the system has previously described the dolphin to the user then this comparison can easily be made.⁷

Dynamically producing hypertext documents allows for the production of shorter texts, since the system can omit links to concepts it no longer considers relevant, or include links to those it determines to be relevant—but not sufficiently so to realise fully within the current document. If the user requires further information about any of those concepts, then he/she can ask follow-up questions by selecting the links. This helps to alleviate the problem of overwhelming the user with too much text, and increases the likelihood that the user will read the whole document.⁸

In the PEBA-II system, it is assumed that the user reads all the text displayed but that their knowledge of concepts diminishes over time. In ILEX, some pieces of information are re-iterated; how often depends both on the item's assimilation score (which indicates the extent to which it is assumed to be known), and on the user-type's assimilation rate (which indicates the ease with which the item becomes assimilated). It remains to be seen which strategy will prove to be the more effective; it is possible that different strategies will be appropriate in different domains.

Coherence Arising from the System's Own Goals. A final issue related to conceptual coherence is worth noting. In the system-user dialog, the main aim is to let the user achieve his/her goals. However, it is still feasible for the system to pursue its own agenda while servicing the user's needs.

In a sense, the user's freedom to explore information is partially restricted by the hypertext links the system chooses to include. Educational systems (such as on-line encyclopædias or museum guides), and shopping catalogue systems, might have in-built goals which they try to achieve, pointing the user to items which are determined as important (for learning or selling purposes), interesting or useful. Of course, in both our systems, the user can at any time return to the front page,

⁷Of course, it is not always entirely clear which aspects of the concepts in the past discourse will be remembered by the user.

⁸Many users of hypertext systems not only browse through documents; they skip through the paragraphs of text within them.

and select any object to be described. They are still “really” free.

Furthermore, in producing descriptions of objects which the user requests, the ILEX system ensures that each description contributes maximally towards the user's goals, while building up a coherent picture for the individual user. For example, suppose we are describing an item of jewellery which happens to be made of plastic; then (while this may not be important in itself), it allows the system the opportunity to make one of the points on its own agenda: *although we now regard plastic as a cheap material (and thus generally use it in mass-produced jewellery), early this century, it was difficult to make, and hence used in one-off designer pieces*. In this way, the system's own goals are opportunistically satisfied.

CONCLUSION

We have argued that adding natural language generation techniques to on-line hypertext catalogues purchases important leverage, particularly given the wide range of users of such systems. In particular, a hypertextual network and the nodes of that network (the documents) can be automatically constructed, and customised to the individual user's knowledge and needs; we term the result DYNAMIC HYPERTEXT. By using NLG techniques, the production of on-line texts is highly flexible and can be tailored to a level of detail beyond that possible using multiple canned texts; we can introduce a significant amount of variation, presenting each user with a genuinely personalised electronic catalogue which builds on their knowledge by making references to known concepts. Using comparison in dynamic hypertext systems helps users find what they are searching for in a maze of hypertext documents, and to acquire the information in an understandable way. Throughout this paper, we have illustrated these ideas with examples from two implemented systems on the WWW.

ACKNOWLEDGEMENTS

We would like to thank the members of the ILEX project team, the Microsoft Research Institute (MRI), and the Intelligent Interactive Technology group at CSIRO MIS for their support and input into the contents of this paper; particular thanks to Chris Mellish, Robert Dale, Cécile Paris, Mick O'Donnell, Alistair Knott and Steven Green. We are also grateful to the four anonymous reviewers for their very constructive suggestions. The first author would like to thank HCRC and MRI for funding a trip to Edinburgh to implement comparisons in the ILEX system. The second author is an EPSRC Advanced Fellow, and his work is supported through EPSRC project grant GR/K53321.

REFERENCES

1. Brusilovsky, P. [1996] Methods and techniques of adaptive hypermedia. *User Modeling and User Adapted Interaction*, 6(2–3).
2. Dale, R. [1992] *Generating Referring Expressions*. Cambridge, MA: MIT Press.
3. Dale, R., Oberlander, J., Milosavljevic, M. and Knott, A. [In Press] Integrating natural language generation and hypertext to produce dynamic documents. *Interacting with Computers*, 10.

4. Encarta [1995] *Microsoft (R) Encarta'95 Encyclopedia*. Copyright ©1994 Microsoft Corporation. Copyright ©1994 Funk and Wagnall's Corporation.
5. Halasz, F. G. [1988] Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, **31**, 836–852.
6. Knott, A., Mellish, C., Oberlander, J. and O'Donnell, M. [1996] Sources of flexibility in dynamic hypertext generation. In *Proceedings of the 8th International Workshop on Natural Language Generation*, pp. 151–160. Herstmonceux, Sussex, UK.
7. Goodman, N. [1972] Seven strictures on similarity. In Goodman, N. (ed), *Problems and Projects*, pp. 437–447. New York: Bobbs-Merrill.
8. Hitzeman, J., Mellish, C. and Oberlander, J. [1997] Dynamic generation of museum web pages: The intelligent labelling explorer. *Archives and Museum Informatics*, **11**, 105–112.
9. Kaindl, H. and Snaprud, M. [1991] Hypertext and structured object representation: A unifying view. In *Proceedings of the 3rd ACM Conference on Hypertext*, pp. 345–358.
10. Kobsa, A., Muller, D. and Nill, A. [1994] KN-AHS: An adaptive hypertext client of the user modeling system BGP-MS. In *Proceedings of the 4th International Conference on User Modelling*, pp. 99–105.
11. Landow, G. [1989] The rhetoric of hypermedia: Some rules for authors. *Journal of Computing in Higher Education*, **1**, 39–64. Reprinted in P. Delany and G. Landow (eds.) *Hypermedia and Literary Studies*, pp. 81–104. Cambridge, MA: MIT Press, 1991.
12. Levine, J., Cawsey, A., Mellish, C., Poynter, L., Reiter, E., Tyson, P. and Walker, J. [1991] IDAS: Combining hypertext and natural language generation. In *Proceedings of the Third European Workshop on Natural Language Generation*, pp. 55–62. Innsbruck, Austria.
13. McKeown, K. R. [1985] Discourse strategies for generating natural-language text. *Artificial Intelligence*, **27**, 1–41.
14. Milosavljevic, M. and Dale, R. [1996] Strategies for comparison in encyclopædia descriptions. In *Proceedings of the 8th International Workshop on Natural Language Generation*, pp. 161–170. Herstmonceux, Sussex, UK.
15. Milosavljevic, M., Tulloch, A. and Dale, R. [1996] Text generation in a dynamic hypertext environment. In *Proceedings of the Nineteenth Australasian Computer Science Conference*, pp. 417–426. Melbourne, Australia.
16. Milosavljevic, M. [1997] Content selection in comparison generation. In *Proceedings of the 6th European Workshop on Natural Language Generation*, pp. 72–81. Duisburg, Germany.
17. Moore, J. D. [1989] A Reactive Approach to Explanation in Expert and Advice-Giving Systems. *PhD Thesis*, University of California, Los Angeles.
18. Moore, J. D. [1995] *Participating in Explanatory Dialogs: Interpreting and Responding to Questions in Context*. Cambridge, MA: MIT Press.
19. Moore, J. D. and Swartout, W. R. [1990] Pointing: a way toward explanation dialogue. In *Proceedings of AAAI90*.
20. Not, E., Petrelli, D., Stock, O., Strapparava, C., and Zancanaro, M. [1997] Augmented space: Bringing the physical dimension into play. In *Proceedings of the Flexible Hypertext Workshop*, held in conjunction with the 8th ACM International Hypertext Conference. Southampton, UK, April 1997.
21. Paris, C. L. [1987] The Use of Explicit User Models in Text Generation: Tailoring to a User's Level of Expertise. *PhD Thesis*, Columbia University.
22. Paris, C. L. [1993] *User Modeling in Text Generation*. London: Pinter Publishers.
23. Shneiderman, B., Brethauer, D., Plaisant, C. and Potter, R. [1989] The Hyperties electronic encyclopedia: An evaluation based on three museum installations. *Journal of the American Society for Information Science*, **40**, 172–182.
24. Simon, L. and Erdmann, J. [1994] SIROG: a responsive hypertext manual. In *Proceedings of the ACM European Conference on Hypermedia Technology*, pp. 108–116, Edinburgh.
25. Stotts, D. and Furuta, R. [1989] Petri-net based hypertext: document structure with browsing semantics. *ACM Transactions on Information Systems*, **7**, 3–29.
26. Tversky, A. [1977] Features of similarity. *Psychological Review*, **84**, 327–352.