Virtual Museums on the Information Superhighway: Prospects and Potholes

Maria Milosavljevic†, Robert Dale‡, Stephen J Green‡, Cécile Paris† and Sandra Williams‡

†Intelligent Interactive Technologies Group
CSIRO Mathematical and Information Sciences
Locked Bag 17
Sydney NSW 1670 Australia
Email: {Maria.Milosavljevic,Cecile.Paris}@cmis.csiro.au

‡MRI Language Technology Group
Macquarie University
Sydney NSW 2109 Australia
Email: {rdale,sjgreen,swilliam}@mri.mq.edu.au

Abstract

Museum labels are written according to the assumed knowledge and needs of a single restricted audience model. However, each visitor has unique reasons for visiting a museum or for being attracted to a particular exhibit. In addition, each visitor has unique knowledge, abilities, needs and limitations. In order to make labels more informative, interesting and useful to the individual visitor, labels must become *dynamic*; that is, a label must be able to be different things to different people, according to their needs. In this paper, we describe the use of natural language generation technology in the construction of personalised virtual museum labels which are constructed on-the-fly by a system with a specific person's needs in mind. The key point here, is that the description of a museum object can automatically adapt to the individual user. We call the result of this technology DYNAMIC HYPERTEXT. We highlight the benefits which arise from using this approach in the domain of museums: labels can be tailored more precisely to an individual's needs and background, not only in terms of the links to other interesting items in the museum, but more importantly, the descriptions of items themselves can be tailored to the user, thus making their visit a far more interesting and rewarding experience.

1. Introduction

With the fast uptake of Internet-based technologies, many businesses are grappling with a major paradigm shift in order to firm up their future business models. By embracing the use of the Internet within and outside of the museum, museums can (and often do) lead the way in this new world. In particular, we are already seeing extensive use of the web for providing virtual museums. However, most web-based documents are static in nature, and are not able to be tailored to the individual user's knowledge, needs, abilities and limitations. This severely limits their usefulness, as users may become lost in hyperspace (Conklin, 1987), may not know how to find items of interest, and may not find the descriptions of items in the museum useful. Like other businesses, it is essential that the visitor has a rewarding experience, or they will not return to the museum. In order to overcome this, we, as document authors, may provide multiple versions of documents, or documents with varying levels of descriptions for all the conceivable questions a user might have. However, this task becomes impossible when we are dealing with more than a handful of objects, and more than a few user types. Furthermore, even with such careful authoring, we would not be able to author documents that could take the visitor's past interactions with the system into account.

To overcome the problems of information overload, we need systems which allow a more interactive, personalised, dialogue-based means of exploring information on-line. By utilising natural language generation (NLG) technology, we can build systems which provide users with personalised electronic catalogues (Milosavljevic and Oberlander, 1998; Milosavljevic 1998). NLG is concerned with the development of techniques for producing linguistic output, either written or spoken, from some underlying information source. These systems can utilise information about the users (such as their goals, knowledge, abilities and limitations) in order to provide information which is most relevant and understandable. We thus overcome the limitations of static museum labels by exploiting natural language generation techniques to create labels that can be tailored on the basis of a model of the user's knowledge or the user's previous interactions with the system.

In this paper, we outline the architecture and benefits of dynamic hypertext systems, and argue that by making more effective use of a user model and the discourse history, NLG techniques can produce highly personalised museum labels. We illustrate the advantages of these techniques through examples from some implemented systems which dynamically produce descriptions of entities as World Wide Web (WWW) pages.

2. Dynamically Generating Museum Labels

If we were to imagine our ideal museum system, a person would be able to approach it (or use a hand-held version of the same) and enter their interest: e.g., *I am interested in music*. The system would then determine which items in the museum might be relevant or interesting to that person, and propose a path and associated story through the museum. That person's tour would be entirely personalised. As the visitor "walks" through the information space (whether physical or virtual), the system would be able to 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 would 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 museum. The interaction between these spaces is particularly interesting since such a system would be able to lead the visitor to items of potential interest in close proximity. The constraints posed by a physical layout could also impact the choice of other items linked into the unfolding discourse, both those referred back to (the most recently visited items) and those toward which the visitor is pointed (those items most interesting to the visitor and in the closest proximity). Additionally, for the purposes of illustration, the system would be able to make use of items which are currently perceivable by the user in order to make the description more immediately accessible, for example "*Notice the other item in this cabinet was also made in this era.*"

As well as incorporating some of the above user modelling ideas, we also aim to make our systems conversational in nature. Ideally, the user's interaction with such systems should be more like that of a person-to-person dialogue. In this dialogue, the user's actions (such as clicking on a hypertext link) represent *follow-up questions* to which the system must respond (see Moore, 1990). In responding, however, the system must maintain an equilibrium between the user's freedom of exploration in the physical or virtual museum, and the goals of the system, in educating the user about some concepts. The system must therefore be selective about what information to include in a description of a museum object, and what other information should be pointed to from the context of the current description.

Comparison is an effective tool for introducing new concepts or entities to a reader. Hence, since the purpose of museum descriptions is to teach the user something about the objects in the museum, as well as more general concepts, comparison must play an important role in such systems (Milosavljevic, 1997). Just like the physical museum space, hypertext is, by its nature, non-linear in form, meaning that the path by which a user might arrive at a particular item cannot always be predicted in advance. Hence, we can make effective use of comparison in order to smooth the transition from one description to the next. This cannot be achieved using simple annotations in existing documents. For example, the system might produce a sentence such as "*Here we have yet another example of …*" or "*This item, although it appears similar* *to the previous one, in fact has …*" This also creates a more dialogue-like interaction between the system and user.

Bear in mind that the objectives described above are our end goals, and our current systems are early prototypes in which we eventually plan to demonstrate these capabilities. We have developed two systems which produce descriptions of items in the Powerhouse Museum in Sydney, Australia. For our first system, Power, we constructed the information source (the KNOWLEDGE BASE) by hand. This involved encoding abstract information about museum objects such as date of manufacture, designer, dimensions of the item and so on. Because we included a wide range of information for a small number of objects, the resulting descriptions are rich in content. However, encoding information by hand is a very time consuming task, and for the entire 200,000 objects owned by the museum it is simply not feasible. So for our second system, PowerTNG, we extracted information automatically from a text version of part of the Powerhouse's Collection Information System (CIS) database. Our system constructs descriptions of museum objects from a knowledge base built with this extracted information. As a result, the descriptions in PowerTNG are shorter, since the amount of information that we could easily extract was limited, but our coverage is extremely wide; the system can describe approximately 15000 objects in English, French, Spanish and Chinese (to a greater or lesser degree). In the sections that follow, we first highlight some advantages of our approach, and then describe the two systems we implemented, including the prospects and pitfalls of both methods. We then turn to some lessons learned in order to aid others in avoiding similar pitfalls.

2.1 Natural Language Generation and Dynamic Hypertext

NATURAL LANGUAGE GENERATION systems aim to produce coherent written or spoken language from underlying representations of knowledge. Some of the advantages of using NLG technology in general include:

- **Description of existing data:** given the appropriate linguistic resources, an NLG system can automatically produce a text to describe a wide range of existing data. For example, an NLG system can be employed to describe the contents of a database, to produce reports describing numbers (e.g., stock reports, weather reports, etc.), to explain the reasoning of an expert system, or to provide documentation for a program.
- **Contextual tailoring:** the generation process can make use of information only available at the point of use (such as characteristics of the particular reader, or information about the content of recent interactions the user has had with the system) to create texts that are tailored to specific requirements. Since the process does not need to be constrained by existing human-written textual fragments, the range of texts that can be produced is potentially very large and not easily producible by other means.
- **Up-to-date reporting and documentation:** if descriptions of the information source are created automatically and dynamically, there is no requirement to update such descriptions manually, with the attendant problems of errors and time lag.
- **Multilinguality:** if the underlying information source is not expressed in terms of a particular natural language, then it is possible to generate descriptions of the same information in different languages automatically.

DYNAMIC HYPERTEXT systems build on conventional NLG technology, but focus on the production of hypertext documents. There are several differences which need to be taken into account when producing hypertext rather than other types of text; see Milosavljevic *et al* (1996) and Dale *et al* (In press) for a discussion of these. The most significant advantage of hypertextual interaction is that the user is given the freedom to perform high-level discourse planning. That is, by selecting hypertext links in an electronic document, the user effectively drives the system, and the system can make assumptions 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 **dynamically** 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. A survey and comparison of existing dynamic hypertext systems can be found in Knott (1996), and further discussion on the advantages of such systems is provided elsewhere (Levine *et al*, 1991, Milosavljevic *et al*, 1996, Dale, *et al*, In press).

2.2 Stage 1: The Power System

As we've mentioned, NLG systems provide mechanisms for producing linguistic output, whether written or spoken, from some underlying information source. Figure 1(a) shows the typical architecture of conventional NLG systems, and Figure 1(b) shows our architecture for generating virtual museum labels on the Web. The Power system is based on the PEBA-II system (Milosavljevic *et al*, 1996), which produces descriptions and comparisons of animals as WWW pages. The system begins with a DISCOURSE GOAL; this indicates the *purpose* of the text to be generated. It is a user request, and in this system, can be to either describe a single museum object or to compare two museum objects.

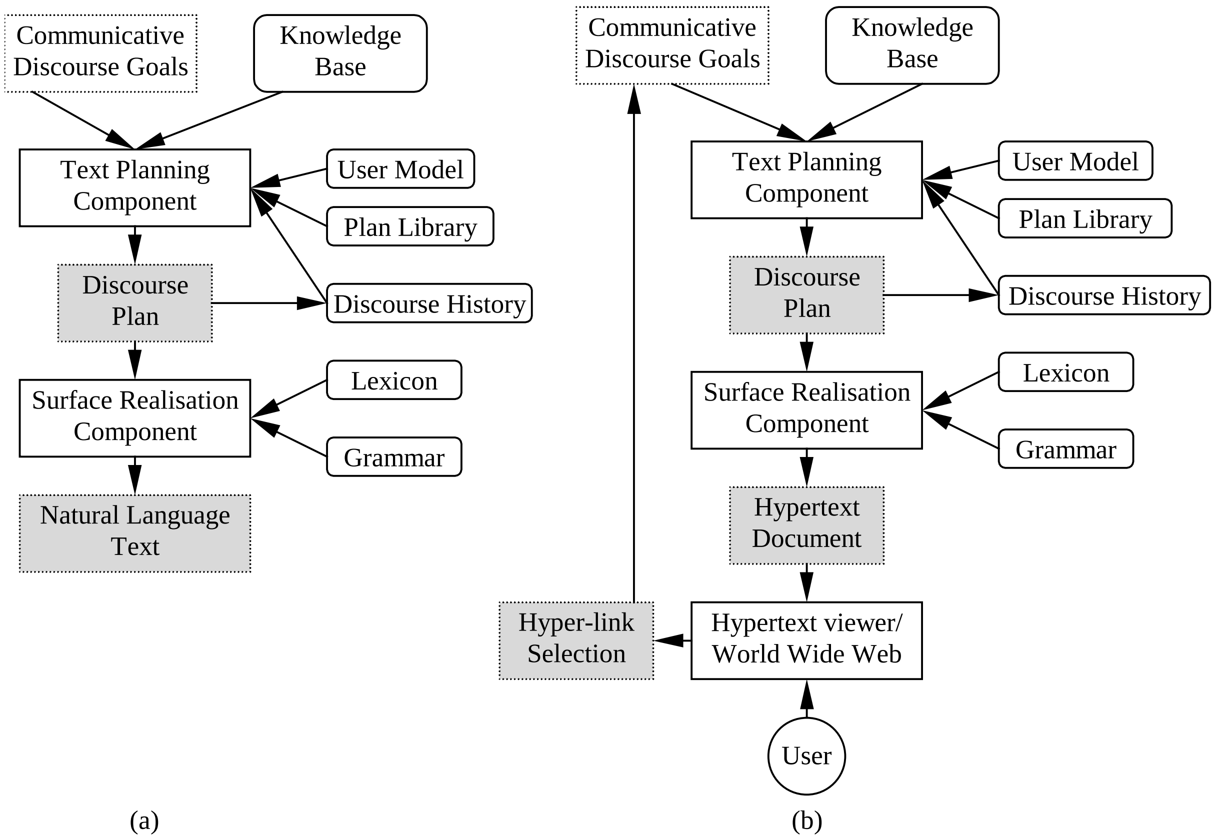

Figure 1: System architectures: (a) traditional NLG; (b) dynamic hypertext.

Based on the discourse goal, the system selects a DISCOURSE PLAN from the plan library which can be used to satisfy the goal. The plan library is part of the system's linguistic resources, and contains standard ways of describing or comparing objects. The discourse plans in the Power system are based on the notion of discourse schemas introduced by McKeown (1982), but modified for use in a hypertext environment. For example, the system needs to determine in which circumstances a hyperlink is appropriate, and what discourse goal is to be given to the system when clicking on the corresponding hypertext link.

After selecting a discourse plan, the text planning component fills it with *facts* from the knowledge base. These facts constitute the information which will be contained in the resulting text. For the Power system, the knowledge base was hand-constructed as is common in many other NLG systems.

A user model stores specialised information about particular users. This allows the tailoring of texts to the requirements of individual users. In the Power system, the user model is exploited in two ways: to provide two different views of the knowledge base, one for naïve and one for expert users; and to modify how descriptions and comparisons are presented (see Milosavljevic *et al*, 1996).

A record of the discourse is also maintained for each user, and is exploited in combination with the user model in order to improve the coherence of descriptions. For example, if the user has knowledge of an entity (as recorded in the user model) or has been told about an entity (as recorded in the discourse history), then the entity can be referred to in later descriptions where a comparison can be made with that entity (Milosavljevic, 1997). Figure 2 shows a comparison between the Difference Engine and the Analytical Engine produced by this mechanism. It is important to note that this text does not exist as text anywhere, but was generated on the fly from the underlying knowledge base.

Figure 2: A text generated by Power including a comparison

Once the text planning component has pulled together all the information about the entity (or entities) to be described in the document according to the user model and the discourse history, the filled discourse plan is passed to the surface realisation component. Here, the discourse plan is realised as natural language sentences, and HTML tags are positioned within the text to allow the user to request follow-up questions by selecting them. When the user selects a hypertext link within the description, a new discourse goal is posted to the text planning component, and the cycle repeats.

2.3 Stage 2: The PowerTNG System

Having demonstrated the idea of using NLG to produce virtual museum labels on the Web, we then examined the issue of using an existing database, as opposed to our carefully hand-crafted knowledge base in which we encoded precisely the kinds of information we needed to generate the texts we were aiming for. Our aim here is two-fold. On the one hand, we want to see the types of texts that can be generated when using a real database. On the other hand, we also want to draw some conclusions as to how the databases should be constructed in order to take advantage of what NLG can offer in terms of the production of virtual documents on demand.

In the PowerTNG system, our implementation uses the same basic NLG architecture. It produces descriptions of the Powerhouse artefacts on the Web, starting with their Collection Information System (CIS). The Powerhouse Museum's CIS is a database of the 200,000 objects the museum owns, although for our pilot study we narrowed our focus to the approximately 5,000 objects that are actually on display on the museum floor (as with many museums, most of the collection is in storage). We then supplemented these 5,000 objects with any objects that are part of a display object, and any objects that have a display object as a part. Our current system contains information on 15,483 objects.


```
<rec num=12798 id="H4448-513">
OID: H4448-513
INT: Part
LOC: TH2.STEP.6A
LOD: 27/11/1997
OBN: Boots
OBS: Balmoral boots, elastic sided, pair, women's, patent/kid/leather/elasticised
fabric/wood,/brass prize work, [Gundry & Sons], England, c.1851; 1862-1869. DES: Balmoral boots,
elastic sided, pair, women's, patent / kid / leather / elasticised fabric / wood /brass, prize
work, [Gundry & Sons], England, c.1851; 1862-1869. Pair of women's elastic sided boots
(Balmoral), with wooden filler, of welted construction with rounded toes featuring peaked caps
and stacked heels.  The uppers consist of a patent golosh, seamed at the back, glace kid leg,
seamed at front and back, and elastic sides extending to the golosh. The uppers are decorated
with oval stitching at the edge of caps and scallops at the throat of golosh.  The leather heel
is fine wheeled, featuring a top piece with brass nailed edge. The black leather sole features a
sueded forepart with brass nails, as well as an internal clump and brass hinged section for extra
strength and a brown polished ridged waist with black edge.  Reputed to have been made by Gundry
& Sons.  (See object file for specialist report by June Swann)
MDE: Gundry & Sons; London, England
MDN: 1965 list says "made by Gundry & Sons, Soho Square."  Swann says hinged device to increase
flexibility is unusual. Similar screws on H4448-515.  Note hinged sole in 1862 exhibition.  She
finds no information about Box in information she has about the 1851 exhibition, though William
Walsh is mentioned in connection with a pair of shoes.  Patent 558, 5 March 1861, granted to J.M.
Carter, a similar sole with 2 cuts across the tread and 4 rows of screws "for soldiers, riflemen,
sportsmen.  The inner sole is whole and contains pitch."  It is not possible to confirm whether
these boots contain pitch.
DAT: c 1851 - 1869
MAR: Interior obscured by last, no marks on exterior
DIM: Length 248 mm Height 31 mm Overall Height 160 mm Width 58 mm
</rec>
```


Figure 3: A CIS database record

Figure 3 shows a typical record from the CIS database. A great deal of the information contained in the CIS records is of a *textual* nature. Most NLG systems do not start with text as input, as they are not always capable of understanding the *meaning* of a text fragment. Instead, information is most often represented as *facts* or *semantic* units of information which form the input to the system. However, many systems, including our Power system, make use of a combination of factual and textual information. Given our set of 15,483 records like the one shown in the figure, we processed these to build a hierarchical semantic knowledge base that can then be used by the generation module. Figure 4 shows a text generated from the information obtained automatically from the above database record. It is clearly of a less sophisticated nature than the text shown in Figure 2. This is largely because the knowledge base created automatically from the database record is not as sophisticated or rich as the knowledge base created by hand for the purpose of generating descriptions. Yet, even with less information, some of the benefits of producing texts automatically on demand using NLG techniques can be seen. For example, the text can be generated in different languages on demand, as shown in Figure 5, provided the appropriate linguistic resources are available.[1]

We now look in more detail at how we extracted information automatically from the database. The Powerhouse museum provided us with a dump of their database in ASCII format with the fields in the database records indicated by tags at the beginning of each field. They also provided us with a thesaurus

[1] In this figure, the words in single quotes indicate that the word is in English, as the appropriate lexical item for the language of the generated text has not yet been provided.

of object types. This is the only information we had available to build a structured knowledge base from which to produce text. To be able to obtain a knowledge base that can serve as input to the generation process, we processed the data file as follows:

1. **Normalisation of the database:** This is to ensure that each record is surrounded by an SGML-style `rec` tag, and that each field of an entry is on a single line.

2. **Extraction of dimensions:** In this step, a Perl script extracts the dimensions of the objects. This information resides in easily identifiable fields (e.g., the DIM field in Figure 3) and the information in that field is structured and can be decomposed into its subfields (e.g., length, height).

3. **Extraction of thesaurus categories:** This step involves trying to identify the thesaurus category that applies to each of the objects in the database. This is normally found in the OBN (Object Name) field and corresponds to an entry in the Powerhouse's thesaurus.

Figure 4: A text generated by PowerTNG

4. **Extraction of names, materials, makers, locations, and dates of construction:** This involves extracting information from the textual information contained in the database records. Most of our work here so far has focussed on the OBS (Object Statement) field. This field is supposed to include information encoded in a standardised and rigorous way. However, in practice, not all the information that is supposed to be included is present, or it is present in a different order, or format, from the norm. Yet, in many cases, we were able to automatically identify information such as date of manufacturing or purchasing, materials and location.

5. **Extraction of PART-OF and A-KIND-OF information:** We use the OID (Object ID) field to determine the PART-OF hierarchy for the database. For example, in the database record shown in Figure 3, the OID `H4448-513` indicates that this object is the 513th part of the object with OID `H4448` (in this case the Balmoral boots are part of a large collection of footwear). According to the database specifications, an object may have parts, sub-parts, and sub-sub-parts.

‘Balmoral boots’

Estos son ‘Balmoral boots’. Fueron hechos entre los años 1870 y 1875. Fue producido en `London', Inglaterra. Están hechos de cuero, patent leather, glace kid, lino y madera. Los `Balmoral boots' tienen 45 mm de altura, 255 mm de largo, 55 mm de anchura total, 150 mm de altura total y 30 mm de ancho.

`Balmoral boots'

Ces objects sont des `Balmoral boots'. Ils ont ete fabriques en entre 1870 et 1875. Ils ont ete fabriques a `London'. Ils sont en `leather', patent leather, glace kid, `linen' et `wood'. Les `Balmoral boots' sont 45 mm de hauteur, 255 mm de longeur, 55 mm de largeur total, 150 mm de hauteur total et 30 mm de largeur.

Figure 5: Texts in different languages: Spanish and French

The result of this processing is the construction of an information record such as the one shown in Figure 6, from which our knowledge base containing 15,483 objects used by the PowerTNG NLG system was automatically generated.

```
OBS.original: Balmoral boots, elastic sided, pair, women’s,
patent/kid/leather/elasticised fabric/wood,/brass prize work,
[Gundry & Sons], England, c.1851; 1862-1869.
OBS.object: Balmoral boots
OBS.object.number: plural
OBS.material.1: patent
OBS.material.2: kid
OBS.material.3: leather
OBS.material.4: elasticised fabric
OBS.material.5: wood
OBS.production.country: England
OBS.create: 1851
OBS.create.inexact: 1
```

Figure 6: Data extracted from an Object Statement field.

3. Discussion

From our experiments with the systems Power and PowerTNG, we conclude that NLG techniques are appropriate to produce virtual museum labels on demand, and that in fact they offer a number of advantages over techniques dealing with already existing documents, in specific situations where underlying information is available. In order for this technology to be effective, however, we would like to make some recommendations on issues related to the source of information. In our experiments with an existing database which was not built for this specific purpose, we learnt some valuable lessons.

To obtain sufficient information for the production of sophisticated texts, special care must be taken when the database is designed and populated. Lack of structure and consistency in the database, and the inclusion of large amount of unstructured textual information[2] in the fields restricted the amount of high quality information we were able to extract. Yet, we believe these obstacles are not insurmountable in many cases. Indeed, the features required (e.g., consistency and structure) are important for well-designed databases. It is important to note that complying with them would not necessarily impose constraints on the end-users, given appropriate interfaces and tools. In fact the database would be greatly improved for all.

It is also quite possible that there will be fewer problems of this kind in the future: as application programs become more sophisticated, it is likely that their underlying representations will have the

[2] Note that even the physical description section in the CIDOC Guidelines for Museum Object Information does not cater for information extraction possibilities.

characteristics required and that their content will move closer to the kinds of rich symbolic structures expected in AI systems. It is also possible that increasingly sophisticated data input tools will be developed to enable the construction of such knowledge bases (see for example, Paris and Vander Linden, 1996 and Paris *et al*, 1998), so that database entry clerks do not have to acquire the skills of knowledge engineers in order to do their jobs.

4. Some Other Hypertext Systems in the Museum Domain

In this section, we briefly describe some other dynamic hypertext systems in the museum domain.

4.1 ILEX

The ILEX system[3] (Hitzeman *et al*, 1997) produces descriptions of museum artefacts as WWW documents. The system currently describes objects in the National Museums of Scotland's 20th Century Jewellery Gallery. ILEX strives to produce descriptions which are accurate, and which include information that is both interesting to the visitor and important in terms of educating the visitor. The system maintains a SYSTEM AGENDA which contains those communicative goals the system should attempt to achieve, while at the same time allowing the visitor to perform high-level discourse planning. Hence, when the user asks for a description of a particular necklace, the system might produce the following response if the user's interest lies in the styles of jewellery: "*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.*" Note, in particular, how the system attempts to educate the user more broadly by describing the organic style in more detail, and relating this new information to the focused item (Knott *et al*, 1997).

4.2 HyperAudio

The HyperAudio system[4] (Not *et al*, 1997) has been developed at IRST in Trento, Italy in conjection with the Civic Museum of Roverto. Instead of using NLG techniques to automatically construct documents, the system constructs speech files from a collection of speech segments. It operates in a similar way to NLG systems, by assessing which information is relevant to the user at any given point. However, the speech files are "canned" rather than generated on the fly. What is interesting about this project is the investigation of the interaction between a physical world (the museum itself) and the virtual world (the representation of the museum and objects in the system). The question is whether and how these worlds impact each other, for example, should a user in the virtual world be informed of the physical layout of the museum, and/or should the descriptions be constrained by this. In Milosavljevic and Oberlander (1998), we discuss this issue in terms of the similarity of items in a particular context. So, for example, if the user is in the physical setting of a museum, then his/her current position should impact not only the way items are described, but also those items which (s)he is directed to which are similar to others which (s)he has expressed interest in.

4.3 AVANTI

The AVANTI system[5] (Fink *et al*, 1997) is being developed by the German National Research Center for Information Technology (GMD). The purpose of this system is to cater for the heterogeneous needs of the users of museum systems (as well as more general systems). In particular, they are concerned with the needs of elderly and disabled visitors. The system does not use NLG techniques, but instead controls

---

[3] ILEX is available on the WWW at: `http://cirrus.dai.ed.ac.uk:8000/`

[4] The HIPS project home page can be found at `http://ecate.itc.it:1024/projects.html`

[5] The AVANTI home page is at `http://www.avanti-acts.org/`

which parts of a document are shown to a particular user, including sections of the document content, links to other documents, the prominence of particular parts, the navigation aids and so on. The system also modifies the display mechanisms used in the interface, incorporating a macro mouse, Braille display and speech synthesiser are appropriate for the blind or deaf.

4.4 Virtual Spurlock Museum

The Spurlock Museum[6] (Marty, 1997) of world history and culture at the University of Illinois, is currently closed while its new building is being constructed. In the meantime, a virtual counterpart is currently being investigated by studying the profiles of the museum's visitors, and designing the system to react to the user's interests. The museum is currently developing virtual representations of the new museum's galleries, and users of the system will be able to tour the new facility, visit the virtual exhibitions and view artifacts before the new building is complete. The exhibitions are aimed at different classes of users, including the general public, students and teachers. The museum is also currently investigating the educational uses of the virtual museum in schools, and it is anticipated that over 5000 students will visit the site per year.

5. Conclusion

In this paper, we have described the use of natural language generation techniques to automatically produce personalised descriptions of museum artifacts which address the user's individual knowledge, needs, abilities or limitations. These descriptions are virtual; they do not exist in any form prior to the user's request, but are created on-the-fly by the system, and customised to the individual user's needs; we call the result DYNAMIC HYPERTEXT. The benefits of this technology is especially important for museums, given the wide range of different goals visitors may have; 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. In this way, we can introduce a significant amount of variation, presenting each user with a genuinely personalised tour of the museum which builds on their knowledge and makes references to known concepts. In order to build realistic systems rather than research prototypes, we need to access existing museum databases. However, as evidenced by our experience, there will undoubtedly be problems with existing data, since it was not constructed for this purpose. If we are to successfully bring museums into the virtual world where the visitor's visit is entirely personalised, then data representation issues will become considerably more important.

References

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

Conklin J. (1987) Hypertext: an introduction and survey. *IEEE Computer* 20(9): 17-41.

Dale R., Green S J., Milosavljevic M., Paris C., Verspoor C. and Williams S. (1998) The Realities of Generating Natural Language from Databases. In *Proceedings of the 11th Australian Joint Conference on Artificial Intelligence*. 13-17 July 1998. Brisbane, Australia.

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.

Fink J., Kobsa A. A. and Schreck J. (1997) Personalised Hypermedia Information through Adaptive and Adaptable System Features: User Modeling, Privacy and Security Issues. In Mullery A., Besson M., Campolargo R. and Reed R (eds), *Intelligence in Services and Networks: Technology for Cooperative Competition*. Pages 459-467. Berlin Heidelberg: Springer.

[6] The Spurlock Museum can be found at `http://spurlock.lang.uiuc.edu/`

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.

Knott A., O'Donnell M., Oberlander J. and Mellish C. (1997) Defeasible Rules in Content Selection and Text Structuring. In *Proceedings of the 6th European Workshop on Natural Language Generation*. 24-26 March 1997. Duisburg, Germany.

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.

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.

Marty P. F. (1997) Database publishing over the Internet: Building a museum over the WWW. Presented at the *Third Annual Advanced Inforamtion Technology (AIT) Faculty Seminar.* University of Illinois.

McKeown K. R. (1985) *Discourse strategies for generating natural-language text*. Artificial Intelligence, 27:1--41.

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.

Milosavljevic M. (1997) Augmenting the User's Knowledge via Comparison. In *Proceedings of the 6th International Conference on User Modelling.* Sardinia, Italy.

Milosavljevic M. and Oberlander J. Dynamic Hypertext Catalogues: Helping Users to Help Themselves. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia* (Hypertext'98). June 20-24, 1998. Pittsburg, Pennsylvania.

Milosavljevic M. Electronic Commerce via Personalised Virtual Electronic Catalogues. 1998. To appear in *Proceedings of The 2nd Annual CollECTeR Workshop on Electronic Commerce* (CollECTeR'98). 29th September 1998. Sydney, Australia.

Moore J. D. (1989) *A Reactive Approach to Explanation in Expert and Advice-Giving Systems.* PhD Thesis, University of California, Los Angeles.

Moore J. D. (1995) *Participating in Explanatory Dialogs: Interpreting and Responding to Questions in Context*. Cambridge, MA: MIT Press.

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.

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.

Paris C. L. (1993) *User Modeling in Text Generation*. London: Pinter Publishers.

Paris C.L. and Vander Linden K. (1996). Building Knowledge Bases for the Generation of Software Documentation. In *Proceedings of the 1996 Meeting of the International Association for Computational Linguistics* (COLING-96).

Paris C.L., Vander Linden K. and Lu, S. (1998). Automatic Document Creation from Software Specifications. In *Proceedings of the Third Australian Document Computing Symposium* (ADCS'98). 21 August 1998. Sydney, Australia.