Intelligent Agents for Information Presentation: Dynamic Description of Knowledge Base Objects

Cornelia M. Verspoor*, Robert Dale*, Stephen J. Green*,
Maria Milosavljevic*, Cecile Paris†, and Sandra Williams*

* Microsoft Research Institute
Macquarie University
Sydney NSW 2109
Australia
{kversp,rdale,sjgreen,mariam,swilliam}@mri.mq.edu.au
Tel: +61 (2) 9850 9528, Fax: +61 (2) 9850 9529

†CSIRO Mathematical and Information Sciences
Locked Bag 17
North Ryde NSW 2113
Australia
Cecile.Paris@cmis.csiro.au
Tel: +61 (2) 9325 3158, Fax: +61 (2) 9325 3200

Keywords: natural language generation, user tailoring, multilingual information presentation

Abstract

Users of the World Wide Web have needs and interests which can help to determine what of the vast quantities of information available might be relevant to them. Intelligent agents might be used to select content for a particular user. However, it is also important to consider *how* that content is provided to a user. We suggest that this *information presentation* must also take into consideration the needs of a user, and discuss a set of agents which utilizes natural language generation techniques to present information in an appropriate way.

In this paper we describe two systems we have built which dynamically generate descriptions of knowledge base entities, and consider the extension of the techniques used there for multilingual information presentation. We describe the notion of a *phrasal lexicon* as a basis for dynamic object description, and propose a model for dynamic multilingual description which builds on that notion.

1 Introduction

In the context of the World Wide Web, intelligent agents are often discussed as a means of finding information that is relevant to a particular user's needs at a particular point in time. Those needs, however, also influence the most effective presentation of the relevant information, as might the needs of the information provider. It is therefore important to consider ways of introducing user-tailoring into information presentation. In this paper we explore the use of a set of intelligent agents for information presentation on the World Wide Web. These agents use natural language processing (NLP) techniques to generate appropriate descriptions. We argue that a key benefit of this technology is the potential for data reuse under varying presentation constraints, and examine the particular case of multilingual presentation to suggest one way forward for this technology.

This work builds on existing work in tailored knowledge base description — the PEBA-II system (Milosavljevic, Tulloch and Dale 1996) , the ILEX system (Knott *et al* 1996), and joint work (Dale *et al* in press) — and describes how knowledge base entities can be dynamically described. We also discuss the ease with which we were able to port the PEBA-II system to a new domain in the POWER system, due to the use of a *phrasal lexicon* and a re-usable agent-based system architecture, and explore how these features of our approach might also be applied in a multilingual generation system.

2 Tailoring methodologies

Information presentation should be tailored to the needs of a user, and the needs of the information provider. For instance, newcomers to a domain will want to see information at a different level of detail than experts. In the context of a virtual museum, a museum curator may have a particular message he hopes to convey to visitors, and the visitors may have an agenda in visiting a particular exhibition. Different people may have different perspectives on individual objects; John might want to know the history of an object while Jane is more interested in its composition.

Such tailoring can be achieved in various different ways. We consider a few current methodologies below.

2.1 Adaptive Hypertext

Brusilovsky (1996) reviews adaptive hypermedia systems. These systems construct user models which drive the adaptation of a hypermedia page. This adaptation might involve, for example, limiting the browsing space a user has access to or prioritizing and/or annotating hyperlinks to guide a user through the available information.

These systems work with static textual documents, but vary the access which a user has to those documents. *Adaptive navigation support systems* support variations in which links are available to a user at a given point, or in the order in which the links are presented, and *adaptive presentation systems* include or suppress certain bits of text depending on the needs of a user.

Although the use of such approaches is an important step towards the tailoring of information to a user, they do not go as far as they might. For example, they do not consider issues of textual coherence — what is the effect of limiting access to texts or changing the order or their presentation? Furthermore, they require all the texts potentially available to a user to be pre-written, which demands a large amount of work when many levels of variability are introduced, and complete rewriting of texts when changes or additions are made in the information to be conveyed. This is an issue for multilingual systems in particular, since the same information must be expressed in several different ways (languages).

2.2 Natural Language Generation

Natural language generation (NLG) systems aim to produce natural language text from an underyling representation of knowledge. The basic components of any natural language generation system can be identified as follows:

- **Content selection:** The selection of information to be included in a text.
- **Text structuring:** The organisation and ordering of the selected information.
- **Surface realisation:** The mapping of the information to be conveyed to well-formed sentences and paragraphs on the basis of a grammar and lexicon for the target languages.

These systems involve goal-driven planning, in which the formulation of a text must satisfy a communicative goal. Content selection, for instance, will differ for different goals as different information will be relevant for each goal. Such goals might be instantiated from schemas (McKeown 1985), or a model of a user's interests/background can provide communicative goals which drive the generation of a text designed specifically for that user.

Each of the above components is independent and embodies knowledge specific to the task the component must achieve. They can therefore be construed as agents which reason on the basis of a current communicative goal about what information to express or how to express that information. Each agent is endowed with a detailed model of its task, and with reasoning procedures for achieving that task relative to the specific context. For example, the text structuring agent might utilize a finite state automata for determining what order to express selected information (chosen by the content selection agent) in. The agents act in sequence to produce the desired output texts. The system architecture of the PEBA-II and POWER systems (to be described in Section 4) is found in Figure 1 and shows the sequence in which agents act and the resources which they utilize. It does not include an explicit content selection agent, because in this system a command given by a user serves to delimit the portion of the knowledge base which will be expressed as a result of reasoning by the text planning and surface realisation agents.

This means that information only needs to be encoded once, and texts can be produced dynamically from that information in various output forms. A generated text might, for example, vary at the surface realisation level in vocabulary or structure for different age groups or in languages for different linguistic groups, or at the content selection level through presentation of different bits of information.

2.3 Dynamic Hypertext

Dale *et al* (in press) discusses the application of NLG techniques to the hypermedia context. The result is a system which *dynamically* creates a hypertext network and hypertext nodes (documents) in response to information about a user or particular communicative goals.

The use of language technology in the hypertext context allows more variation in what information is presented to a user, and in the way that information is presented. The system decides whether a particular string of text should be marked with a hypertext link, on the basis of whether there is more to say about a concept and what the system knows.

This approach is an amalgamation of the previous two approaches, which allows both the incorporation of adaptive hyperlinks and the development of a user model which includes information on what the user has previously explored in the system. Dynamic hypertext systems therefore have the advantages of both previous approaches in terms of flexibility in information presentation, and give the system control over the coherence of that presentation. These systems can be viewed as creating a conversation between the user and the system, in which the system responds adaptively to the high-level discourse planning the user performs by following hyperlinks.

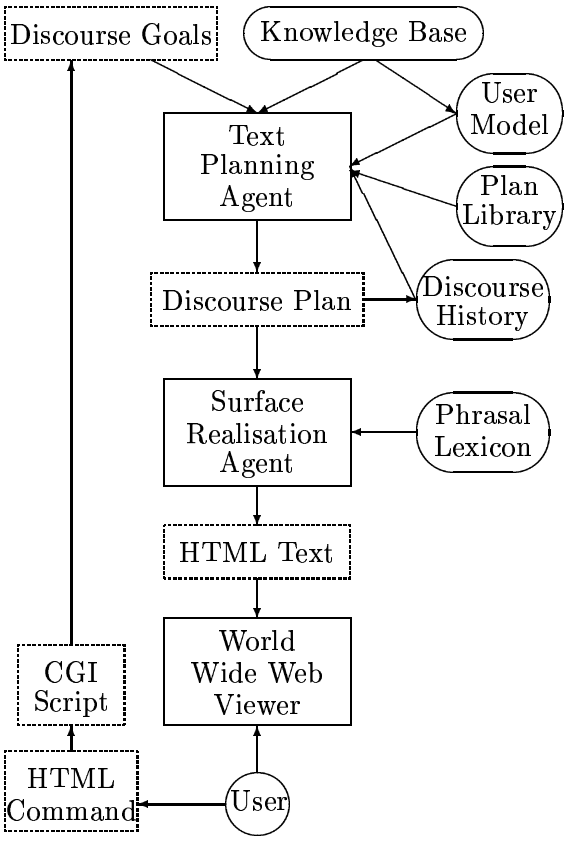

Figure 1: The PEBA–II System Architecture.

3 What are the benefits of this technology?

It is important to identify the value of tailoring technologies. We highlight a few benefits here:

- **Customized information delivery:** Whether perusing a virtual museum, doing on-line shopping, checking out the day's news, or searching for information on a particular topic, a user normally has a set of goals and/or interests which impact on what information he might consider relevant. Adaptive technologies allow those goals and interests to act as a filter on what information is actually presented to that user.

- **Reusability of knowledge sources:** The use of adaptive technologies enables the same information to be presented in the many different contexts where that information might be relevant. This avoids replication of information, as it need only be represented in one location, yet it can be accessed whenever it becomes relevant and inserted dynamically.

- **Avoiding information redundancy:** Information that has already been presented to a user can be suppressed, thus avoiding redundancy in the output. The system can keep track of what information has already been presented to a user (potentially including interactions with the user of previous occasions), and can either leave it out entirely or refer back to it in some way (e.g. via a hyperlink).

Natural language processing techniques in particular have certain advantages for the development of adaptive systems because the *content* of (hyper)texts can be changed, as well as the links between them. Furthermore, since an NLG system actually creates texts, it has much more control over the coherence of the text output than a system which only includes or hides

paragraphs (or another text unit) and has no conception of the impact of those decisions on overall coherence for the end user.

The PEBA-II system (Milosavljevic 1997), for example, can make explicit comparisons between animals to clarify a particular description, but which comparisons are made depend on what information the user has already seen. Dynamic generation of the comparison avoids the need for pre-definition of every possible comparison that might be relevant to any user — the comparison is *constructed* by the system on the basis of stored information rather than hand-created by a human and simply retrieved by the system. The construction of information has certain benefits, discussed by Dale *et al* (in press) and summarized here:

- Reduced text construction costs
- Variation according to purpose
- Variation according to user characteristics
- Variation according to resource boundedness

4 The state of the art: NLP & Object Description

To date there have been several systems implemented which make use of natural language processing techniques to incorporate dynamism in object description. Text in these systems is customised in terms of content and/or presentation.

The PEBA-II system (Milosavljevic, Tulloch and Dale 1996) interactively generates hypertext descriptions and comparisons of animals. The generated text varies according to the expertise level of the system user, the animals which the user can be assumed to know about, and the animals for which the user has already seen descriptions within the system. The output of this system is constructed on the basis of a taxonomic knowledge base, in conjunction with a phrasal lexicon (see Section 5.2), and shows how texts can be tailored to the needs and experiences of individual users. An example of PEBA-II's knowledge base appears in Figure 2.

The ILEX system (Knott *et al* 1996, Hitzeman *et al* 1997) aims to simulate the interaction between a museum tour guide and a visitor, by dynamically generating hypertext pages which vary on the basis of what information a visitor has already seen (the **discourse history**) and the goals which the system is attempting to achieve in the course of the particular interaction. ILEX utilizes a combination of abstract representations which are mapped via templates to a sentence or part of a text and pre-written texts which are dynamically inserted at appropriate points in the interaction. In this way, the quality of the output text is more similar to that which a human expert would produce than that of most NLG systems, yet the output text is not completely static. Transfer of this methodology to the multilingual context would, however, require all of the pre-written texts to be translated by humans to other languages, and so it is unclear what would be gained by using this particular technology.

We have implemented a new system, called POWER (PEBA-based Ontology With Enhanced Realizations) which builds on the PEBA-II system (Milosavljevic 1997), porting it from the animal domain to the domain of computers which might be on exhibition in a museum. Figure 3 shows a fragment of the knowledge base for the POWER system. Note that the structure of the entities precisely parallels that of the PEBA-II system in Figure 2. The changes which were made to the underlying PEBA-II system involved only the following:

- Construction of knowledge base entries for each of the objects in the new domain, and inclusion of these objects into an ontology
- Encoding of the linguistic realizations of the semantic concepts introduced into the knowledge base for the new objects (i.e. specification of the phrases which correspond to concepts such as "`general-computation`"; see Section 5.2))

```
(hasprop Echidna
  (linnaean-classification Family))
(distinguishing-characteristic Echidna Monotreme
  (body-covering sharp-spines))
(hasprop Echidna
  (geography found-Australia))
(hasprop Echidna
  (social-living-status lives-by-itself))
(hasprop Echidna
  (diet eats-ants-termites-earthworms))
(hasprop Echidna
  (length (quantity (lower-limit (unit cm) (number 35))
                    (upper-limit (unit cm) (number 60)))))
(hasprop Echidna
  (weight (quantity (lower-limit (unit kg) (number 2))
                    (upper-limit (unit kg) (number 7)))))
```

Figure 2: A portion of the PEBA-II knowledge base

- Addition of some domain-specific information about the properties in the computer domain which should be described

None of the agents in the PEBA-II generation system needed to be modified. The description and comparison schemas, surface realisation techniques, and content selection mechanisms transferred directly to the new domain.

In Figures 4-6 we can see the kind of descriptions generated by POWER from the new knowledge base entries, with the support of the PEBA-II underlying architecture. Figure 4 shows a description of the Analytical Engine, including a comparison to a potential confuser, the Difference Engine (also invented by Charles Babbage), which is an analogue computer. Figure 5 shows a description of a General Purpose Digital Computer. The generation at this point takes account of the fact that the user was just visiting the Analytical Engine node, by noting this in the text and constructing the list of subtypes of the General Purpose Digital Computer *relative* to the node just visited (by saying "Apart from the Analytical Engine, which you just saw, the General Purpose Digital Computer has the following subtypes:"). Figure 6 shows a comparison of two objects generated on the basis of the system's knowledge of properties of those two objects. In principle, a comparison can be generated for any combination of objects for which there is information in the system. For a system with many objects, the use of NLG techniques means that each comparison does not have to be pre-constructed by hand but can be reasoned about dynamically by the system and thereby avoids a resource-intensive manual task.

What we have learned from the ease with which we were able to port PEBA-II to a new domain in POWER is that the agents within a modular NLG system can be re-used in new domains, to the extent that the structure of the descriptions required in the new domains parallel that of the original domain.

```
(hasprop Analytical-Engine
  (inventor c-babbage))
(hasprop Analytical-Engine
  (date-of-invention (quantity (exact (unit year) (number 1830)))))
(hasprop Analytical-Engine
  (purpose general-computation))
(hasprop Analytical-Engine
  (height (quantity (upper-limit (unit cm) (number 210)))))
(hasprop Analytical-Engine
  (width (quantity (upper-limit (unit cm) (number 120)))))
(hasprop Analytical-Engine
  (length (quantity (upper-limit (unit cm) (number 300)))))
(hasprop Analytical-Engine
  (date-of-acquisition (quantity (exact (unit year)
                                        (number 1978)))))
(hasprop Analytical-Engine
  (source-of-acquisition d-swade))
```

Figure 3: A portion of the POWER knowledge base

Figure 4: A description of the Analytical Engine generated by POWER

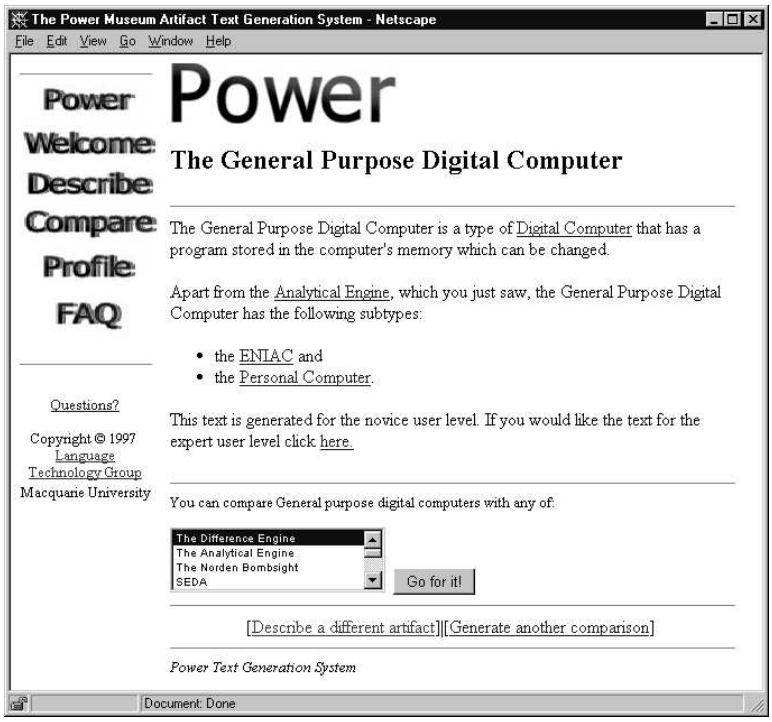

Figure 5: A description of the General Purpose Digital Computer generated by POWER, which takes into account that the Analytical Engine node was just visited by the user

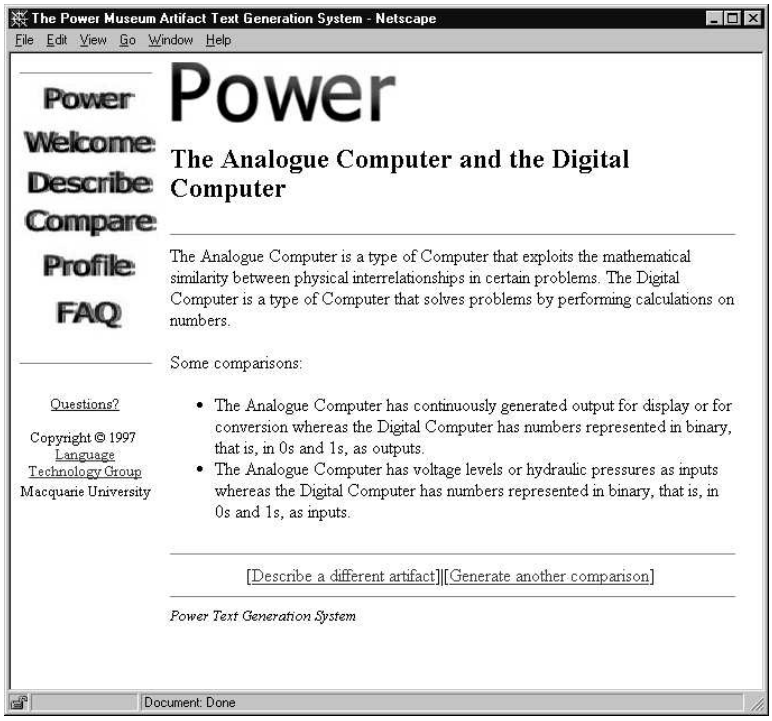

Figure 6: A comparison of analogue and digital computers generated by POWER

5 A proposal for Multilingual Generation

5.1 The components of a multilingual dynamic generation system

The kind of multilingual generation system we propose to implement will make use of current NLG techniques, as in Paris and Vander Linden (1996). Given an underlying language neutral knowledge source, text can be produced in several languages in parallel by providing the systems with the appropriate linguistic resources. In general, text for the same discourse goal but in different languages may vary in both in their discourse structures and in their surface realisations (vocabulary and syntax). As a result, the architecture for a multilingual generation system should allow for this variation, as in Paris and Vander Linden (1996). In some domains, however, it is possible, as a good approximation, to keep the *discourse structure* constant and vary only the *surface realisations*. This is what we intend to do in the proposed system. Thus the only agent within the NLG architecture which will need to change is the surface realisation agent. It will need to be expanded to include a linguistic model for each of the target languages.

The grammars of languages differ, as do their syntax and semantics for individual words. Furthermore, the mapping between concepts and words can vary between languages. Clearly realisations in different languages from the same underlying information can differ dramatically in lexical choice and syntactic structure. This imposes certain requirements on the underlying representation utilised by each of the agents in the NLG system, and on the surface realisation agent.

- **Language neutral representation:** The information in the system domain must be represented in a language-neutral way; that is, the concepts must be captured with minimal bias to their expression in one language or another.

- **Variability:** Any adaptations which the system makes to the information presented should be reflected at the level of the underlying representation, so that the variability is consistent across the target languages. This includes variations in vocabulary choice to the extent that each language makes relevant and consistent distinctions between vocabulary types (e.g. technical vs. general terminology).

- **Structured mapping:** The underlying representation of information must support the realisation of that representation in various target languages. A mapping to particular surface realisations from the representation must be defined within the surface realisation agent for each language, taking into consideration the words to which particular concepts correspond and the grammatical structure of each target language.

- **Language independence:** Variations in the realisations for different languages, whether in surface structure, lexical choice, or the breakdown of information into individual sentences, must be encapsulated in individual language modules within the surface realisation agent. The definitions of these processes for each target language should be isolated from the definitions for each other target language.

These requirements are largely met by the existing structure we find in PEBA-II and POWER, since the knowledge base in these systems is encoded in a representation language which is then mapped, via the phrasal lexicon (see below), to surface realizations. These systems are therefore a good starting point for the exploration of multilingual information presentation. Our hypothesis, based on our experience in porting PEBA-II to POWER, is that through addition of language-specific phrasal lexica into the existing PEBA-II/POWER structure, and with the incorporation of sentence planning procedures for target languages other than English, the systems should extend to the multilingual context and should allow the same range of output variation as the original systems. Thus, we should be able to *reuse* the underlying data in the original systems for multilingual generation. There are, however, a few issues deriving from the use of a phrasal lexicon to which we now turn.

5.2 Phrasal Lexica

A technique in NLG for handling the mapping from the knowledge base to linguistic realisation is the use of a *phrasal lexicon*, in which linguistic "chunks" can be recorded at a level higher than individual words. The surface realisation agent then controls the combination of these "chunks" to form sentences.

A lexical entry in this approach might consist of a full phrase which corresponds to a particular concept represented in the knowledge base. Milosavljevic, Tulloch and Dale (1996) argue that since the knowledge representation in a particular domain often uses complex elements, the linguistic elements to which they correspond should be equally complex. So, for example, they have a lexical item whose orthography is "is a carnivore and eats ants, termites and earthworms" which corresponds to the concept `eats-ants-termites-earthworms` in the knowledge base, rather than building up the structure of the verb phrase from the individual words "is", "a", "carnivore", "and", "eats", "ants", etc. They state,

> The use of phrasal lexical items of this kind has two specific advantages: **Reuse and Efficiency**. If we repeatedly realise a semantic element in the same way, it is better to remember this and avoid rebuilding the surface form each time.

Such phrasal units must exist in the lexicon regardless, due to the existence of idioms and other non-compositional linguistic material (Verspoor 1997).

The extension of the technique of utilizing phrasal lexica to the multilingual context means defining the lexica of the target languages in terms of phrases, and defining grammars which control the appropriate combination of these phrases for specific languages. This will demand a fair amount of linguistic sophistication of the lexicon, in order to account for the realisation of person, number, and gender agreement in the languages which require it, but the grammars utilised by the surface realisation agent can be considerably less complex than grammars which start at the level of individual words and must build up all syntactic structure.

An important issue which our research will address is the appropriate level of semantic granularity for the knowledge representation and correspondingly for the phrasal lexica for the target languages. The break-down of information into representational elements must meet two competing desiderata:

- **Coverage:** Representational elements must be fine-grained enough to capture the full range of concepts and predicates which are relevant to the domain and are likely to be realized in subtle variations in the target texts.
- **Reusability:** Representational elements must be broad enough to reflect generalisations which can be made about related concepts.

This granularity is furthermore affected by the interaction of the knowledge representation and the phrasal lexica, since the representational elements must be chosen in such a way that they consistently correspond to a particular realisation in a particular language. The appropriate level of granularity for the representation of an individual concept may therefore differ for different languages. As a result the representation, and the corresponding phrasal elements, may need to be adjusted for such concepts as additional target languages are included. Representational granularity in the multilingual context will likely differ from that in a monolingual context, and precisely in what ways is an area of investigation we are interested in pursuing.

5.3 Benefits of natural language generation techniques

Machine translation systems can be viewed as one approach to providing multilingual presentation of information. However, for the purposes of dynamic systems which incorporate user-tailoring and generate texts on the basis of underlying data, this approach is inappropriate. The

extension of these dynamic systems to the multilingual context can more effectively make use of natural language generation techniques, avoiding some of the difficulties inherent in machine translation systems and taking advantage of the underlying structured data (Hartley and Paris 1997).

In particular, machine translation systems work with unrestricted natural language as input, while an NLG system creates output texts on the basis of a constrained, pre-defined representation language. This means that the NLG system does not have to address issues of coverage or robustness. All of its inputs can be expected to fit within a particular framework, and it will be designed to generate a particular set of possible sentences from the inputs. Within a well-designed NLG system, no representation should be accessed or constructed by the content selection or text planning agents which cannot be mapped to an output form by the surface realisation agent, while in a machine translation system, it is virtually impossible to allow for the full range of variation and creativity which may exist in natural language inputs.

Machine translation systems need to be able to both do natural language interpretation, to establish a representation of the information embodied by the source text, and natural language generation, to produce the target texts. NLG systems eliminate the need for an interpretation step and can focus on the generation of high quality output.

Furthermore, the ambiguity of natural language is less of a problem in the generation context as the input to the NLG system can be assumed to be unambiguous due to the design of the representation language. A certain piece of information in the knowledge base should permit only one interpretation. The same is clearly not true of a piece of input information that is expressed in natural language.

An accurate machine translation system in a given domain depends on well-structured and easily interpretable documents in a given source language and any changes that need to be made to the information to be conveyed must be made in that source language in a way that will not "break" the translation system. In contrast, an NLG system works with a previously structured source which can be changed as needed (within the constraints established by the representational system). For instance, the TECHDOC generator (Rosner and Stede 1994, Stede 1996) has been used to perform multilingual generation in English, French, and German. The particular domain in this case were maintenance instructions from automobile manuals. The aim was to be able to generate these instructions in all three languages from the same underlying knowledge source. By having such a system, manufacturers would no longer need to manually write instructions, have them translated and then re-iterate when more changes were necessary. All that would be necessary is maintaining the knowledge base and regenerating the manuals when necessary.

The obvious real benefit of NLG for multilingual information presentation is that it allows variation in the language of the output text without requiring pre-written texts in each output language. Harnessing the systematicity of languages for the purpose of automatically mapping from a semantic representation to a linguistic realisation means that every text, and every potential variation of that text for individual users and contexts, does not need to be created in advance by a human text writer, but rather can be generated in the form required by the NLG system.

This approach does, however, demand the construction of an underlying knowledge representation and development of phrasal lexicons at an appropriate level of granularity for mapping to individual target languages. This work can be arduous and time-consuming; clearly there is a need for the development of techniques for automatic acquisition of knowledge bases and phrasal lexicons for the NLG task. Some of these techniques are under investigation, but still are the greatest hurdle for the use of NLG techniques. Once a knowledge base has been constructed, however, it should be reusable in any generation system, including but clearly not limited to a multilingual system, for which mappings to linguistic realizations for knowledge base elements have been defined.

6 Conclusions

A system composed of intelligent agents embodying techniques from natural language generation can provide dynamic flexibility in information presentation, to accommodate the differing needs of different users while allowing data re-use. In this paper we described the PEBA-II system and the porting of that system to a new domain in the POWER system. These are systems which dynamically generate descriptions of objects defined in a knowledge base, and in which distinct agentive modules take information about the user and his knowledge into account in order to construct a specific aspect of those descriptions. We saw that with a simple substitution of domain-specific information, the overall functionality of the original system was easily maintained. We speculated that the same structure can be utilized for multilingual information presentation, requiring only an extension of the surface realisation agent to include grammatical models of each target language, pending investigations into the implications for knowledge base structure of the use of a multilingual phrasal lexicon.

References

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

Dale, R., J. Oberlander, M. Milosavljevic, and A. Knott (in press). Integrating natural language generation and hypertext to produce dynamic documents. *Interacting with Computers.*

Hartley, A. and C. Paris (1997). Multilingual document production: From support for translating to support for authoring. *Machine Translation 12*(1-2), 109–129.

Hitzeman, J., C. Mellish, and J. Oberlander (1997). Dynamic generation of museum web pages: The intelligent labelling explorer. Presented at the International Conference on Museums and the Web, Los Angeles, March 1997.

Knott, A., C. Mellish, J. Oberlander, and M. O'Donnell (1996). Sources of flexibility in dynamic hypertext generation. In *Proceedings of the 8th International Workshop on Natural Language Generation*, Herstmonceux, Sussex, U.K., pp. 151–160.

McKeown, K. R. (1985). *Text Generation.* Cambridge, U.K.: Cambridge University Press.

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

Milosavljevic, M., A. Tulloch, and R. Dale (1996). Text generation in a dynamic hypertext environment. In *Proceedings of the 19th Australasian Computer Science Conference*, Melbourne, Australia.

Paris, C. and K. Vander Linden (1996, July). DRAFTER: An interactive support tool for writing multilingual instructions. *IEEE Computer*, 49–56.

Rosner, D. and M. Stede (1994). Generating multilingual documents from a knowledge base: The TECHDOC project. In *Proceedings of the International Conference on Computational Linguistics*, Kyoto. COLING.

Stede, M. (1996). *Lexical Semantics and Knowledge Representation in Multilingual Sentence Generation.* Ph. D. thesis, University of Toronto.

Verspoor, C. M. (1997). *Contextually-Dependent Lexical Semantics.* Ph. D. thesis, Centre for Cognitive Science, University of Edinburgh.