

# Strategies for Comparison in Encyclopædia Descriptions

Maria Milosavljevic and Robert Dale  
Microsoft Research Institute  
Macquarie University  
North Ryde NSW 2109 Australia  
{mariam,rdale}@mpce.mq.edu.au

## Abstract

Comparisons are typically employed to distinguish similar entities, or to illustrate a property of an entity by referring to another commonly known entity which shares that property. Based on an analysis of a corpus of encyclopædia texts, we define three types of comparisons and outline some strategies for applying these in the generation of entity descriptions. We describe how these comparison strategies are used within the PEBA-II hypertext generation system to generate descriptions of animals.

## 1 Introduction and Aims

In this paper, we outline some strategies for comparison which we use in PEBA-II, a hypertext generation system which produces encyclopædia descriptions of entities as World Wide Web (www) documents, based on an underlying taxonomic knowledge base. PEBA-II is part of a larger research programme built around the idea of an intelligent on-line encyclopædia, where the descriptions produced by the system vary for different users and at different times. Our work is grounded in the domain of animal descriptions, although similar issues arise in many other domains.

Comparisons are widespread within existing encyclopædia descriptions. In particular, when describing a new concept to a user, a

comparison may be made with reference to other known concepts or ideas, enabling the hearer to more easily process and understand the new material (see [Milosavljevic 1996]). So, for example, in the context of animal descriptions, if the user knows about the porcupine and requests a description of the echidna, then we might describe the echidna by highlighting both its similarities to and differences from the porcupine.

Clearly this requires us to make use of some notion of a user model, and in a way that is distinct from previous work in user-modelling in text generation. Our aim is to produce texts which introduce new concepts by reference to existing knowledge the user is assumed to have, thus employing the user model to greater advantage; past research (see, in particular, [Paris 1987]) has concentrated on avoiding the production of texts which repeat what the user already knows.

By making use of a discourse model, we can also generate comparisons that take account of entities that have been mentioned in the previous discourse. This is particularly important in the context of the dynamic construction of hypertext documents from an underlying representation: by employing text generation techniques, we can produce context-dependent descriptions which vary depending on the information which has already been presented to the user, thus overcoming some of the limitations of hypertext documents which have been constructed simply by breaking an existing linear text into

pieces. As has been noted by others (see, for example, [Reiter *et al* 1992] and [Moore 1989]), the dynamic generation of hypertext also permits the user to effectively drive the text generation system, alleviating from the system some of the responsibility of reasoning about what to present to the user.

In Section 2, we provide an overview of the PEBA-II system; in Section 3, we provide a definition of comparison and identify three types of comparison on the basis of a corpus analysis; and in Section 4 we describe how the corresponding discourse strategies are implemented in PEBA-II. Section 5 ends the paper by pointing to some future research directions.

## 2 An Overview of PEBA-II

The architecture of the PEBA-II system is shown in Figure 1; the components are as follows.

The knowledge base that currently underlies the system has been hand-constructed from an analysis of encyclopædia descriptions of animals and constitutes a taxonomy of the Linnaean animal classes with their associated properties. Particular properties may also be labeled as distinguishing for a specific class.

The plan library consists of discourse plans which are used by the text planning component. Currently, the system makes use of two high level discourse plans, which we name identify and compare-and-contrast. The identify discourse plan is used to describe an entity and the compare-and-contrast discourse plan is used to compare two entities. These discourse plans are similar in spirit but rather different in content to the similarly-named schemas used by McKeown [1985], with a number of the differences arising from the fact that we are generating hypertext pages.

A new discourse goal is generated by the user clicking on a hypertext link in the current document being viewed. Given this new goal, the text planning component selects any relevant information from the knowledge base

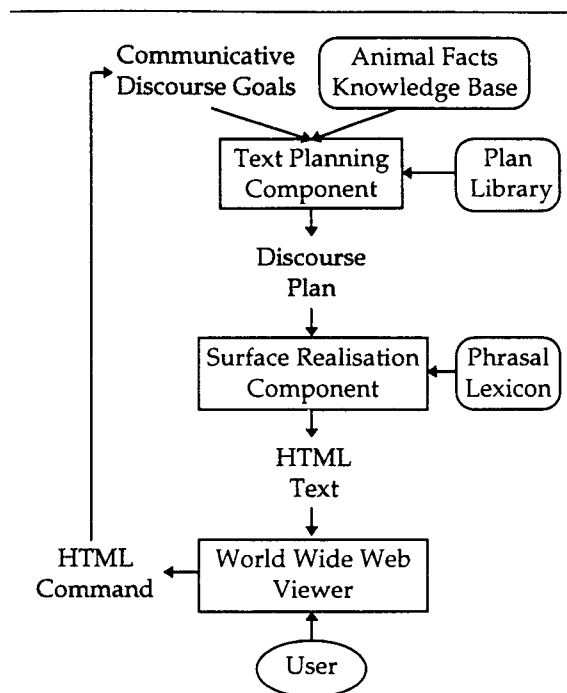


Figure 1: The architecture of the PEBA-II system

and organises the information according to the current discourse plan. The leaves of the instantiated discourse plan are then realised via a simple template mechanism.<sup>1</sup>

The output from the PEBA-II system is a document marked up using a subset of HTML commands. This document may be displayed using any WWW document renderer such as Mosaic or Netscape. The user poses new discourse goals to the system by clicking on any of the hypertext tags, and the cycle continues.<sup>2</sup>

The combination of text generation and hypertext has been explored by others, most notably in Moore's [1989, 1995] PEA and in Re-

<sup>1</sup>Although we have experimented with using Elhadad's [1992] FUF realisation engine, for the texts we currently generate a template-based mechanism is faster and seems quite adequate. Speed is important in the context of Web-based generation: see Tulloch and Dale [1995] for some ideas on addressing the problems here.

<sup>2</sup>A version of PEBA-II is available on the Web at URL: <http://www.mpce.mq.edu.au/msi/peba.html>.

iter *et al*'s [1992, 1995] IDAS. PEBA-II is closest in concept to the IDAS system; a more detailed description of PEBA-II can be found in [Milosavljevic, Tulloch and Dale 1996]. Knott *et al* [1996] discuss some further issues involved in combining hypertext with natural language generation.

## 3 Defining Comparisons

### 3.1 Data analysis

A corpus analysis has been conducted to identify how comparisons are used in encyclopædia articles, so that these techniques may be built into the PEBA-II system. In the first instance, we have concentrated on the domain of animal descriptions; we intend to widen the scope of this analysis to other domains in order to provide a more domain-independent theory of comparative forms. The two encyclopædias analysed were Microsoft Encarta [Microsoft 1995] and Groliers Multimedia Encyclopædia [Groliers 1992]; each encyclopædia yielded around 1200 animal entries, and from these we collated a subcorpus of sentences involving comparison. This subcorpus contains 1722 sentences from the Encarta corpus, and 1557 from the Groliers corpus.

The aim of the corpus analysis was to reverse-engineer the comparisons found in animal descriptions in order to answer the following questions:

- What entities are compared in descriptive texts and how do they relate to each other? What properties of these entities are used in comparisons?
- Why are particular entities compared? Why are some entities better comparators than others?
- What techniques do we need to build into a text generation system to be able to produce similar comparisons?

## 3.2 Some Definitions

### 3.2.1 Comparison

We will adopt the following definitions:

A *comparative proposition* is a proposition whose purpose is to draw the hearer's attention to a difference or a similarity that two entities have for the value of a shared attribute.<sup>3</sup>

A *comparison* is the linguistic realisation of a set of one or more comparative propositions, where the purpose of the set of propositions is to draw the hearer's attention to one or more differences or similarities between two entities.

We have identified three different types of comparative forms that appear in descriptive texts, which we refer to here as DIRECT COMPARISONS, CLARIFICATORY COMPARISONS, and ILLUSTRATIVE COMPARISONS. Of these three types, only the first has been explored to any great degree in the context of natural language generation: both McKeown [1985] and Maybury [1995] have looked at various aspects of direct comparisons.

### 3.2.2 Direct Comparisons

A DIRECT COMPARISON is a comparison whose purpose is to compare two entities where neither entity is more central to the discourse than the other. In the context of a language generation system like PEBA-II, direct comparisons arise when the user enters a request such as: *What is the difference between the Echidna and the African Porcupine?* PEBA-II generates the text shown in Figure 2 in response to such a query.

This text is essentially 'bi-focal': the echidna and the porcupine are equally impor-

---

<sup>3</sup>In the terminology we adopt here, a PROPERTY is a tuple consisting of an ATTRIBUTE and a VALUE; for example, (colour, red).

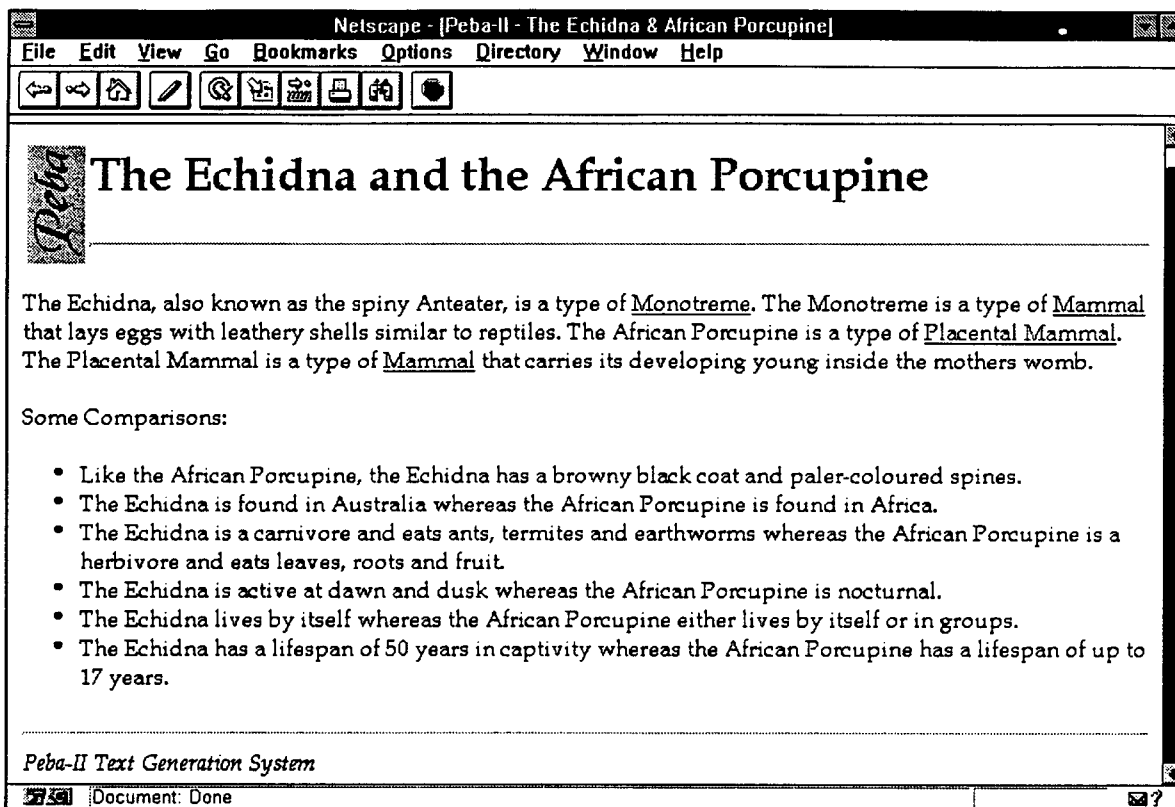


Figure 2: A direct comparison as generated by PEBA-II

tant, and the purpose of the text is to determine their similarities and differences based on both their relationship within a taxonomy of animals (their lowest common ancestor) and their attributes. This is, of course, the same notion of comparison that is used in McKeown's [1985] TEXT system.

The key point here is that direct comparisons are generally user-initiated. More interesting from the point of view of text generation are clarificatory and illustrative comparisons: here, the entity being described by the system is described in relation to some other entity chosen by the system.

### 3.2.3 Clarificatory Comparisons

A CLARIFICATORY COMPARISON is a comparison whose purpose is to describe an entity

by distinguishing it clearly from another entity with which it might be confused or with which it shares a number of salient properties. In such cases we will refer to the first entity as the FOCUSED ENTITY, and to the second entity as the POTENTIAL CONFUSOR.

The main difference between a clarificatory comparison and a direct comparison is that a clarificatory comparison is made within a text whose purpose is to describe one entity and not purely to provide a comparison between two entities. A clarificatory comparison serves to describe the focused entity; thus, it corresponds to the user entering a request such as *What is the echidna?* In such a case, instead of describing the echidna in isolation, the system may choose to describe it using a clarificatory comparison with the porcupine.

There are two reasons why a clarificatory comparison might be used:

- The focused entity might be extremely similar to another entity, and therefore often confused with that entity. In this case, it is important that, when describing the focused entity, it is sufficiently distinguished from the potential confuser.
- Alternatively, an entity sharing a number of salient features with the focused entity might already be known to the user; in such a case, a clarificatory comparison between these entities may aid the user's understanding of the focused entity.

For example, consider the following text extracted from the animal corpus:

Sheep, are hollow-horned ruminants belonging to the genus *Ovis*, suborder Ruminata, family Bovidae. Similar to goats, sheep differ in their stockier bodies, the presence of scent glands in face and hind feet, and the absence of beards in the males. Domesticated sheep are also more timid and prefer to flock and follow a leader. [Groliers 1992].

In this text, the focused entity (the sheep) is very similar and might often be confused with the comparator entity (the goat); this is particularly true of some wild sheep. A reader who is familiar with the comparator entity will also more easily form a mental picture of what the focused entity is like.

There are a number of interesting research issues here:

- How is a comparator entity selected? For example, a very appropriate comparator for the echidna is the porcupine, but the two entities are not closely related within the Linnaean taxonomy of animal classes. The reason for the choice of comparator entity here lies in the fact that both

animals possess sharp spines—this is the only salient property the animals share.

- How do we make clarificatory comparisons which do not cause the user to make incorrect inferences? For example, if a user who is not familiar with sheep requests a description of the sheep and the system describes the sheep by informing the user of its similarities with the goat and not their differences, then the user could be led to believe that the two animals are more similar than they are in reality. The text shown above very carefully describes both similarities and differences for only the most salient features which clearly distinguish the animals.

A user model is advantageous here since the importance of different attribute types will vary from person to person. For example, if external appearance is the most important attribute, then we would want to compare the echidna to the porcupine. If, on the other hand, reproduction is considered a more important feature, then we might compare the echidna to the platypus. The geographical location of the user can also play an important role: for example, in the texts that we have examined, squirrels are often used as comparators; but Australians are not necessarily familiar with the features of squirrels, and some North Americans might only know of the existence of black squirrels.

### 3.2.4 Illustrative Comparisons

An ILLUSTRATIVE COMPARISON is a comparison whose purpose is to describe one or more attributes of an entity by referring to the same attribute(s) of another entity with which the user is familiar. The difference between an illustrative comparison and a clarificatory comparison is that in an illustrative comparison, the comparator entity, although usually of a similar type (in this case, an animal), may only share one attribute with the focused entity, and is not necessarily similar in any other

way to the focused entity.

Here are some illustrative comparisons from our corpus:

- Powerful and aggressive animals *about the size of a large dog*, baboons have strong, elongated jaws, large cheek pouches in which they store food, and eyes close together. [Microsoft 1995]
- [Aye-aye] are *about the size of a large cat* and have long, bushy tails, a shaggy brown coat, and large ears. [Microsoft 1995]
- *About the size of a small fox*, [the Aye-aye] has a long, bushy tail, moderately large eyes, thick fur, and a pair of enlarged front teeth *resembling those of rodents*. [Groliers 1992]
- This echolocation system, *similar to that of the bat*, enables the dolphin to navigate among its companions and larger objects and to detect fish, squid, and even small shrimp. [Microsoft 1995]
- Slightly larger than chinchillas, the mountain viscachas have long, *rabbitlike ears* and a long *squirrel-like tail*. [Microsoft 1995]

In each of these sentences, an illustrative comparison is made so that the reader can more easily grasp the concept being described. Instead of describing the size and proportion of the viscacha's ears in absolute terms, a reference to the rabbit's ears makes it easier for the reader to understand what the ears really look like.

There is a great deal of scope for tailoring descriptions to a user's knowledge here: for example, illustrating the size of the aye-aye with the fox might be appropriate for a user who is familiar with the fox; however this illustration might not be appropriate for someone located in Australia, since the fox is not found in Australia. The features of a particular animal (the sheep, for example) might also vary geographically.

## 4 Implementing Comparison Strategies

Above, we identified three particular types of comparisons that are present in our corpus. In PEBA-II, each corresponds to a particular discourse strategy for generating a hypertext page. In this section, we describe how these strategies are implemented within PEBA-II.

### 4.1 Choosing Amongst the Strategies

We are faced with two interdependent questions: when do we decide to describe an entity by comparing it to another entity, and how do we decide which type of comparison to use?

Recall from earlier that PEBA-II can address two different discourse goals: requests to describe some specified entity, and requests to compare two specified entities. The latter discourse goal corresponds, of course, to the category of direct comparisons we identified above. As we noted earlier, direct comparisons are thus user-initiated. We are more interested here in how PEBA-II decides when it is appropriate to use either a clarificatory comparison or an illustrative comparison. Each becomes an option when PEBA-II has been asked to describe some specified entity. A clarificatory comparison is generated whenever the entity to be described is known to have a POTENTIAL CONFUSOR: our implementation of this strategy is currently very simple, and is described in Section 4.3. Illustrative comparisons are the focus of the current work, and we describe our approach to these in Section 4.4.

### 4.2 Direct Comparisons

As mentioned earlier, the PEBA-II system allows the user to request one of two actions: to describe a single entity or to compare two entities. A direct comparison is generated by PEBA-II whenever the user requests a comparison between two entities. Using a

corpus-derived property classification system, the discourse plan used here pairs up those attributes which are of a similar type (for example, measurements such as height and length) and compares their values. An example WWW page generated using this strategy is shown in Figure 2.

### 4.3 Clarificatory Comparisons

The purpose of a clarificatory comparison is to ensure that the reader does not confuse the entity being described with some other entity. Such confusions are possible when the entity being described is similar in relevant respects to some other entity.

We could try to generate such clarificatory comparisons from first principles: when we have to describe some entity *e*, we could search the knowledge base for entities which share properties with *e*, and then use some mechanism to determine whether there is any chance that the two entities might be confused. We could then phrase our description of *e* to make sure that we distinguish *e* from such potential confusors. For example, in describing the rabbit, it may be important to distinguish it from the very similar hare in order to avoid confusion.<sup>4</sup> There are problems with such an approach: searching the knowledge base in this way would be a very costly process: it assumes a rather more complete knowledge base than we may be able to rely on; and, most important of all, it assumes that we can determine likelihood of confusability on the basis of some metric—but it is not at all clear what such a metric might be.

Our current solution to these problems is to sidestep them entirely: for each entity that has a potential confusor—for example, sheep and goats—we specify this explicitly in the

<sup>4</sup>There are clearly ideas we might use here in McCoy's [1988] work on correcting a user's misconceptions; however, the real issue here lies in determining *whether* such a misconception might arise from a generated comparison (see Zukerman and McConachy [1993, 1995] for some work in this area).

knowledge base by means of a clause of the following form:

```
(hasprop sheep  
  (potential-confusor goat))
```

Then, whenever we have to describe the sheep, we know immediately that it has a potential confusor in the goat, and invoke a discourse strategy that makes an explicit comparison between the two entities. The resulting text includes a comparison with the goat but is aimed at describing the sheep and hence goes further than a direct comparison between the sheep and goat.

Hard-coding potential confusors might be considered an 'easy way out', although it is our view that this is one of many places in NLG where there is benefit in adopting solutions that make use of precomputed information in preference to working things out from first principles. For example, singling out potential comparator entities in this way is no different in principle to explicitly marking in the knowledge base those properties which are distinguishing characteristics, a tactic that both McKeown [1985] and we ourselves use.<sup>5</sup> We have adopted this philosophy for various design decisions made in the development of PEBA-II, so that, for example, we also make use of a phrasal lexicon as a repository of precomputed mappings from semantic units to multi-word lexico-syntactic resources (see [Becker 1979] for an early justification for this approach). Again, a similar philosophy underpins the use of precomputed lists of preferred attributes in the work on the generation of referring expressions reported in [Reiter and Dale, 1992]. Our position is that such methods can be a virtue rather than a vice, since they allow broad coverage systems to be built more quickly.

<sup>5</sup>Note that Maybury [1995], on the other hand, outlines an algorithm for determining the distinguishing characteristics for an entity from first principles.

## 4.4 Illustrative Comparisons

Currently, most of our attention is focused on the third category of comparisons, those we have termed illustrative comparisons. These are cases where one or more attributes of an entity being described are compared to those of a common object with which the reader is assumed to be familiar. For the present discussion, we will concentrate on the attributes of size and weight, and the mechanisms used to produce illustrative comparisons that indicate these attributes of the entity being described. These are probably two of the easiest properties to deal with; it remains to be seen to what extent the mechanisms we propose will generalise to other attributes.

For illustrative comparisons, there are two questions to be answered:

- How do we decide whether an illustrative comparator should be introduced?
- How do we decide which comparator to choose when there are multiple candidates?

We could perform these comparisons using a similar approach to that which we adopted for clarificatory comparisons: for each entity-attribute pair we could specify some entity that can be used as a comparator. Thus, we might have clauses in the knowledge base that look like the following:

```
(hasprop baboon  
  (illustrative-comparator size dog))
```

However, this would be unwieldy: part of the justification for taking this approach in the case of clarificatory comparisons is that we would expect a relatively small subset of the entities in the knowledge base to have potential confusors, and so the cost of explicitly encoding a representation of these potential confusors is not too great. However, virtually any entity-attribute pair might be described using an illustrative comparison, and so we need some way of generalising the processing here.

We do this by making use of the notion of a `COMMON COMPARATOR SET`. This is a set of entity types that can be compared against for illustrative purposes. For the moment, a common comparator set is defined for each attribute we might wish to describe; there may be some scope for interesting generalisations later. We focus here on the size and weight attributes: for both of these, our common comparator set is the set

```
(human, dog, cat)
```

Note that the common comparator set for any given attribute is

- domain specific: different comparator sets for size and weight will be appropriate in different domains;
- user specific: it is likely that different comparator sets will be appropriate for different users; and
- in principle extensible, both directly and indirectly: we can imagine the user explicitly being allowed to specify a set of comparator objects, or we could dynamically extend the set used on the basis of the ongoing discourse history.

There may be ways of building or precompiling a common comparator set automatically using the knowledge base and information from a user model, but for the moment we assume that it has been preconstructed. Given an entity  $e$  we want to describe and some attribute  $a$  of the entity we want to communicate, we use the algorithm in Figure 3.

The procedure used here for finding the best match is one that in our current experiments looks acceptable, although it is likely to be applicable only for a relatively narrow range of attributes. There are a number of obvious deficiencies, all of which we are currently exploring:

- Properties are not independent: for example, we have found that, when deal-



---

To describe attribute  $a$  of entity  $e$  (the focused entity):

- Identify comparator set  $S_a$  for attribute  $a$
- $Val$  = median value of  $a$  for  $e$
- For each  $e_i \in S_a$ , identify location of  $Val$  on the range of values that  $e_i$  has for  $a$
- Choose best match:
  - choose  $e_i$  whose median value for  $a$  is closest to  $Val$
  - if this doesn't select uniquely from amongst the comparator set then choose  $e_i$  whose range for  $a$  is closest to  $Val$

Figure 3: Choosing a comparator object

---

ing with size, we also need to take account of similarity of body-form in determining which entity makes the best comparator, and so our current mechanism distinguishes three different size measurements: height, length and shoulder-height.

- Similarity and difference are not completely distinct: the similarity of two values for a particular attribute should be viewed as a scale of similarity rather than as a binary distinction.
- The user's degree of familiarity with the potential comparators can help in making a choice.
- The degree of relatedness between the two entities can also play a role in choosing the best comparator.

So far, however, the results of the simple method we have outlined seem promising. For example, PEBA-II currently generates the following sentences:

- The platypus is about the same length as a domestic cat.
- The baboon has about the same shoulder height as a domestic dog.

Note that the use of a common comparator set in conjunction with the algorithm specified here means that we can separate the domain-specific aspects of the computation from the domain-independent aspects; in principle, the aim is that the comparator set specifies domain-specific information, but the algorithm itself is domain independent.

As always, our methodology is to pursue solutions that first assume a considerable amount of precompiled knowledge and then introduce generalisability and flexibility through subsequent parameterisation, rather than beginning with a very limited coverage solution that works from first principles. It is our view that this methodology is the only one that is likely to be successful for broad coverage, practical NLG systems.

## 5 Conclusions and Future Work

In this paper, we have:

- described the PEBA-II system, as an example of a system which integrates natural language generation and hypertext in the provision of user-tailored information;
- defined some notions relevant to the study of comparison; and
- looked at the concept of illustrative comparison in detail, with the aim of defining a mechanism for generating such comparisons that embodies a clear distinction between domain-dependent and domain-independent information.

For future work, we intend to elaborate upon and extend further the techniques described here. In particular, we intend to make use of these notions in the generation of comparisons which take account of the discourse history; some examples of this phenomenon are discussed in [Dale and Milosavljevic 1996].

## Acknowledgements

We would like to thank Mike Johnson of Macquarie University and the members of the Language Technology Group at Microsoft Research Institute for many discussions related to this work.

## References

- J. D. Becker [1979] The Phrasal Lexicon. In *Proceedings of the Conference on Theoretical Issues in Natural Language Processing*, Cambridge, MA, pp. 70–77.
- Robert Dale and Maria Milosavljevic [1996] *Authoring on Demand: Natural Language Generation in Hypertext Documents*. In *Proceedings of the First Australian Document Computing Conference*, Melbourne, Australia.
- Michael Elhadad [1992] *Using Argumentation to Control Lexical Choice: A Functional Unification Implementation*. PhD Thesis, Columbia University.
- Groliers [1992] *The New Grolier Multimedia Encyclopædia*. Copyright Grolier Incorporated, (c) 1987-1992 Online Computer Systems, Inc.
- Alistair Knott, Chris Mellish, Jon Oberlander and Mick O'Donnell [1996] Sources of Flexibility in Dynamic Hypertext Generation. In this volume.
- Kathleen McCoy [1988] Reasoning on a Highlighted User Model to Respond to Misconceptions. *Computational Linguistics*, 14(3):52–63.
- Kathleen McKeown [1985] *Text Generation - Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press.
- Mark T Maybury [1995] Using Similarity Metrics to Determine Content for Explanation Generation. *Expert Systems with Applications*, 8(4):513–525.
- Microsoft [1995] *Microsoft (R) Encarta'95 Encyclopædia*. Copyright (c) 1994 Microsoft Corporation. Copyright (c) 1994 Funk and Wagnall's Corporation.
- Maria Milosavljevic [1996] Introducing New Concepts Via Comparison: A New Look at User Modeling in Text Generation. In *Proceedings of the Fifth International Conference on User Modelling, Doctoral Consortium*, pp. 228–230.
- Maria Milosavljevic, Adrian Tulloch and Robert Dale [1996] Text Generation in a Dynamic Hypertext Environment. In *Proceedings of the Nineteenth Australasian Computer Science Conference*.
- Johanna Moore [1989] *A Reactive Approach to Explanation in Expert and Advice-Giving Systems*. PhD Thesis. UCLA.
- Johanna Moore [1995] *Participating in Explanatory Dialogues*. MIT Press, Cambridge, MA.
- Cecile Paris [1993] *User Modelling in Text Generation*. Pinter Publishers.
- Ehud Reiter and Robert Dale [1992] A Fast Algorithm for the Generation of Referring Expressions. In *Proceedings of Coling-92*, Nantes, France, August 1992.
- Ehud Reiter, Chris Mellish and John Levine [1992] Automatic Generation of On-line Documentation in the IDAS Project. In *Proceedings of the Third Conference on Applied Natural Language Processing*, Trento, Italy, March–April 1992.
- Ehud Reiter, Chris Mellish and John Levine [1995] Automatic Generation of Technical Documentation. *Applied Artificial Intelligence*, 9(2):259–287.
- Adrian Tulloch and Robert Dale [1995] Speeding up Linguistic Realisation by Caching Previous Results. In *Poster Proceedings of the Eighth Australian Joint Artificial Intelligence Conference*, Canberra, November 1995.
- Ingrid Zukerman and Richard McConachy [1993] Generating Concise Discourse that Addresses a User's Inferences. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, pp. 1202–1207. Morgan Kaufmann Publishers.
- Ingrid Zukerman and Richard McConachy [1995] WISHFUL: A Discourse Planning System that Considers a User's Inferences. *Technical Report*, Department of Computer Science, Monash University, Australia.