# Generating Indirect Anaphora

Claire Gardent & Kristina Striegnitz

Department of Computational Linguistics
Universität des Saarlandes, Saarbrücken, Germany
{claire,kris}@coli.uni-sb.de*

**Abstract**

Much information in natural language can be left implicit. From the generation perspective, this raises the problem of how to model the processes and in particular, the reasoning, underlying such implicitness. In this paper, we concentrate on the generation of one of the many natural language constructs supporting implicitness namely, indirect anaphora. We first summarize the inferences governing the use of indirect anaphors. We then show how indirect anaphors can be generated within a generation architecture which interleaves sentence realization with contextual reasoning.

**Keywords:** natural language generation, automated reasoning, indirect anaphora

## 1 Introduction

Indirect anaphora[1] are definite noun phrases that occur without an explicit antecedent in the surrounding text. Here is an example (we use bold face to highlight indirect anaphors and italics to indicate the noun phrase in the discourse context that the anaphor is related to) :

(1) The Young Woman scans *the restaurant* with this new information. She sees all **the patrons** eating, lost in conversations. **The tired waitress**, taking orders. **The busboys** going through the motions, collecting dishes. **The manager** complaining to **the cook** about something. A smile breaks out on the Young Woman's face[2].

The definite expression *the patrons* in this example is of course to be interpreted as referring to the patrons of the restaurant mentioned in the first

*This work was partially supported by the Project InDiGen in SPP–Sprachproduktion, grant by the Deutsche Forschungsgemeinschaft to the University of Saarbrücken.

[1]aka textual ellipsis, bridging anaphora, inferable, associative anaphora, contiguity anaphora, first-mention definites

[2]from: *Pulp Fiction*, film-script by Quentin Tarantino

sentence. The definite article signals that the intended referent is not any group of patrons, but a particular group of patrons. However, the information which makes this group identifiable in the current context, namely its relation to the earlier mentioned restaurant, is left implicit in this text. The same holds for *the tired waitress*, *the busboys*, *the manager*, and *the cook*.

Now, consider (2), a modified version of (1), where the missing information linking the definite descriptions into the context is made explicit.

(2) The Young Woman scans *the restaurant* with this new information. She sees all **the patrons of that restaurant** eating, lost in conversations. **The tired waitress of the restaurant**, taking orders. **The busboys of the restaurant** going through the motions, collecting dishes. **The manager of the restaurant** complaining to **the cook of the restaurant** about something. A smile breaks out on the Young Woman's face.

This text is awkward, and it is so not just because of the recurrence of the construction *of the restaurant*, but because the information conveyed by this expression is simply not necessary. Schwarz (2000) briefly reports on the results of acceptability tests that confirm this observation: the information left implicit by indirect anaphors is not noticed as missing, and in certain cases providing this information leads to less acceptable texts. Thus in order to produce natural sounding texts, a natural language generation (NLG) system must be able to leave some of the information that is inferrable by the hearer implicit. For instance, in the context given in 1, it must be able to generate the indirect anaphor **the patrons** rather than the fully explicit **the patrons of that restaurant**. In this paper, we look at the inferences underlying the use of indirect anaphors. We then show how indirect anaphors can be generated within an NLG architecture which interleaves sentence planning and realization with contextual reasoning.

The structure of the paper is as follows. In section 2, we motivate and introduce the NLG architecture assumed in this paper. In section 3, we investigate the contextual restrictions governing the use of indirect anaphors. In section 4, we show how the generation architecture described in section 2 supports the generation of indirect anaphors.

## 2 Generation architecture

Much work on generating definite descriptions has focused on determining a set of properties which uniquely identify the referent being talked about (cf. for instance, Dale and Reiter 1995; Dale and Haddock 1991). In such approaches the selection of this set of properties and the surface realization of the corresponding definite NP are two separate and independent processes. As the following examples illustrate however, not all identifying information needs to be explicitly realized in the definite NP. It can be explicitly given elsewhere in the sentence, as e.g. in Example (3a) which is appropriate in contexts where there are several hats, if only one of them contains a rabbit. Or it can be given implicitly through presuppositions of the sentence, as e.g. in Example (3b), which can be uttered in contexts with several rabbits, if there is only one that is in a hat.

(3) a. The rabbit in **the hat**

b. Remove **the rabbit** from the hat.

Thus an important shortcoming of this first type of approach is that it fails to address the question of when and how some of the identifying properties may be left implicit when realizing the definite NP.

To remedy this shortcoming, Stone and Webber (1998) argue that a generation architecture is needed which interleaves – rather than pipelines – contextual reasoning (e.g. the computation of a uniquely identifying set of properties) and sentence realization (e.g. the surface realization of a definite NP). Specifically, they show that the SPUD architecture presented in Stone and Doran's (1997) appropriately captures the type of textual economy illustrated by the above examples.

In this paper, we follow Stone and Webber (1998) and adopt a SPUD-like architecture to generate indirect anaphors. This architecture consists of three main components:

- A Lexicalized Tree Adjoining Grammar (LTAG, Joshi and Schabes 1997).
- A chart based generation algorithm (Kay 1996; Carroll et al. 1999) which uses the LTAG to build a linguistic structure fulfilling the initial communicative goal.
- An automated theorem prover which checks whether the semantic and pragmatic constraints of the linguistic structure being built are satisfied by the context.

## 2.1 The Grammar

For a precise definition of LTAGs we refer the reader to Joshi and Schabes (1997). In what follows, we sketch the properties of LTAG that are important for the purpose of this paper.

An LTAG consists of a set of *elementary* trees (the lexicon) and of two operations on trees called *adjunction* and *substitution*.

The lexicon contains two types of trees: *initial* trees (such as e.g. tree ($\alpha 1$) in Figure 2) which are used to encode the basic syntactic frame of syntactic functors and *auxiliary* trees (such as e.g. tree ($\beta 1$) in Figure 2) which encode modifiers e.g. adjectives, prepositional phrases (PP) or adverbs. Initial trees may contain so-called substitution nodes (marked with $\downarrow$) whereas auxiliary tree must have a unique foot node (marked with $\star$) i.e. a frontier node labeled with the same category as the root of the tree.

The two operations, substitution and adjunction, are then used to combine trees into bigger trees. Intuitively, substitution inserts a tree with root category $X$ at some substitution node $n$ with category $X$ in some other tree. Adjunction on the other hand, caters for recursion and permits inserting an auxiliary tree with root and foot node category $X$ at a node labeled with category $X$ in some other tree. Substitution and adjunction are illustrated in figure 1.

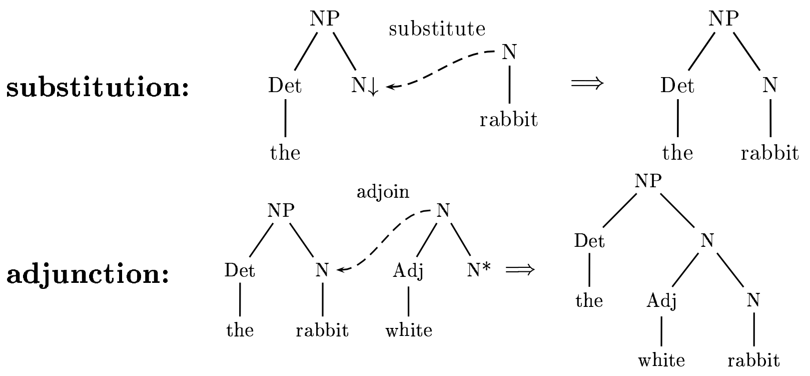


Figure 1: Examples of substitution and adjunction.

The lexicon we use for generation follows the guidelines set in Stone and Doran's (1997) SPUD generation system. In particular, it has the following properties:

- The trees are associated with semantic and pragmatic information.
- The semantic representation language is a flat semantics (Hobbs 1985; Copestake et al. 1999).

Example lexicon entries are given in Figure 2. Nodes in the trees are labeled with the usual syntactic categories and, additionally, variables over object level terms. The intuition behind this is that these terms refer to the entity described by the syntactic node. The entity associated with the 'S', 'VP', and 'V' nodes in tree ($\alpha 1$), for example, is co-referent with the first argument of *lazy*. It describes the state, while the entity associated with the 'NP' node describes the object which is lazy and is therefore co-referential with the second argument of *lazy*. For the purpose of this presentation, we will sometimes associate further semantic information with tree nodes: for instance, the N node in $\alpha_2$ also carries a variable ($N_{sem}$) over the properties described by the missing nominal. This will make a formulation of the uniqueness condition easier. As mentioned above, trees are associated with semantic information (S) which is collected as trees are put together to form bigger trees. The semantics of a complex tree is just the union of the semantic representations of the trees it consists of. The pragmatic information (P), encoding the contextual restrictions governing the use of the lexical items, is similarly collected up the tree. The constraint on the status of the described entity acts as a kind of licensing condition for prediction (selection) of certain lexical element. In particular, the tree for "the" is associated with a uniqueness constraint and the requirement that the entity described by the definite description be anaphoric.

## 2.2 The Generation Algorithm

The task of a NLG system is as follows. Starting from some communicative goal $G$, a generator must produce a text $T$ which achieves this goal in a given

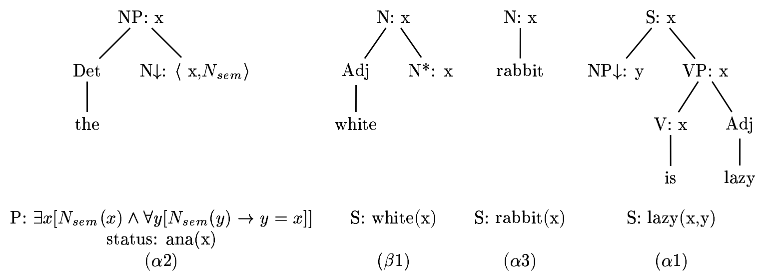


Figure 2: Some lexical items.

context $C$. Here we concentrate on small (i.e. one-sentence) texts and restrict ourselves to simple communicative goals of the form: *describeX*, where $X$ is an eventuality or an entity. Furthermore, we assume the context to consist of a set of formulas representing (i) the meaning of the previous discourse, (ii) the world knowledge that we assume is common to both hearer and speaker, and (iii) the relevant knowledge of the speaker. We take (i) and (ii) to constitute the common ground.

Given this, the basic generation algorithm can be summarized as follows:

Given the initial communicative goal *describe* $e$:

1. Build a *complete* and *minimal* syntactic structure $T$ describing $e$ and check its contextual conditions against the speaker context
2. If the context satisfies these conditions, i.e. if the tree is *contextually appropriate*, then halt; Else adjoin modifiers to $T$ till these conditions are satisfied.

Importantly, contextual restrictions are not checked recursively, i.e. each time a complete and minimal syntactic structure is built. Rather a minimal backbone tree is built for the whole sentence (or in the case of longer descriptions, for the whole discourse) before contextual restrictions are checked and, if necessary, modifiers adjoined. This two step procedure is similar to the chart generator for (semi)-lexicalist grammars presented in Carroll et al. (1999) and so avoids the exponential complexity arising from a more traditional treatment of intersective modifiers (Brew 1992; Carroll et al. 1999).

## 2.3 An Example: Generating "the white rabbit"

Assume the following situation:

(4) **Discourse Model:** $cat(a) \wedge rabbit(b) \wedge rabbit(c) \wedge white(b) \wedge black(c) \wedge a \neq b \wedge b \neq c \wedge a \neq c$[3]

[3]In future examples, we will not explicitly list all of these inequality conditions, although they will be there for each two different constants appearing in the discourse model and the speaker's knowledge.

**Speaker's Knowledge:** $lazy(s, b)$

**Goal:** describe (state) $s$

That is, the discourse model contains information about a cat and two rabbits, one of which is white and the other one black. The communicative goal for the next utterance is to describe the fact that entity $b$, i.e. the white rabbit which is already in the discourse model, is lazy. We will now use this example to illustrate how the generation algorithm sketched above works.

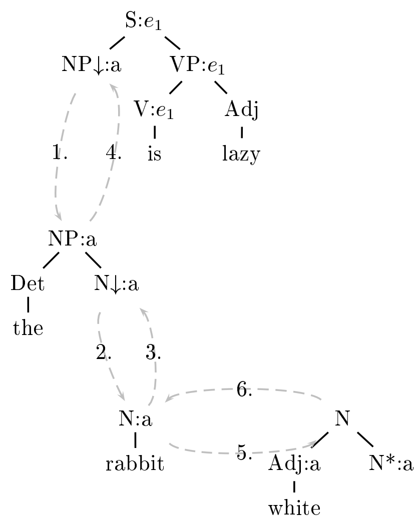


Figure 3: Generating *The white rabbit is lazy*

The system's initial goal will be to build a tree that (i) fulfills $G =$ *describe* $s$, (ii) is syntactically complete and (iii) is contextually appropriate. Figure 3 illustrates the process of generating *the white rabbit is lazy*, which is an utterance satisfying this goal.

First, the tree for *the rabbit is lazy* is built. This roughly works as follows. A prediction step selects elementary trees (in this case, $\alpha 1$) from the lexicon such that (i) their root is labeled with S and (ii) their semantic representation unifies with one or more of $s$'s properties. The resulting item $\alpha 1$ is not syntactically complete, hence generation proceeds where the next goals are determined by the labels of the open substitution nodes. In this case, the system must generate a noun phrase which describes entity $b$, where $b$ is discourse old, i.e. $b$ is an element of the discourse model. Therefore, ($\alpha 2$) is selected, which again has an open substitution site leading to the selection of ($\alpha 3$). This item is syntactically complete and can thus be combined with the previous item, which in turn can be combined with the first item that was put into the chart to form a syntactically complete parse tree rooted in 'S'.

At this stage, the generator has built a *complete* and *minimal* syntactic structure describing the initial goal $s$. Now, the automated theorem prover is called to verify that its contextual requirements are fulfilled. Let the common ground $\mathcal{C}$ be the discourse model specified in (4). Given the contextual requirements associated with the definite article, the task is to verify that $\mathcal{C}$ is consistent with $\exists x[rabbit(x) \wedge \forall y[rabbit(y) \rightarrow x = y]]$. Since this is not the case (there is more than one rabbit), generation proceeds to the adjunction phase, in which modifiers are adjoined to the initial structure. In this case, ($\beta$4) can be adjoined yielding *the white rabbit is lazy.* Again the theorem prover is called to check the contextual requirements of the new tree. This time, the task is to check whether $\exists x[rabbit(x) \wedge white(x) \wedge \forall y[rabbit(y) \wedge white(y) \rightarrow x = y]]$ is consistent with $\mathcal{C}$. The answer is "yes" and therefore generation halts.

## 3 Definite Descriptions Without Antecedents

In the previous section we have seen how definite descriptions could be generated which refer to an entity that is *discourse old* i.e. that has been explicitly mentioned in previous discourse.

However, definite descriptions are not restricted to such cases, as e.g. text (1) in the introduction shows. In fact, corpus studies, such as the one presented in Vieira and Poesio (1996), have shown that only 30% of all definite noun phrases in texts have an explicit antecedent. About 20% can be analyzed as *indirect anaphors*, that is as anaphors which refer to entities that have not been explicitly mentioned before. In such cases, there is an (implicit) link between the referent of the indirect anaphor and some object or event mentioned in the previous sentence and which we will call the *anchor.*

(5) a. There is *a great Italian restaurant* in Market Street. **The cook** is from Naples.

b. I walked into *the room.* **The chandelier** sparkled brightly.

c. John was *murdered* yesterday. **The knife** lay nearby.

Clark (1975) was one of the first to examine the types of relations that link an indirect anaphor to its anchor expression. The strongest link is, when the referent of the indirect anaphor is a necessary part of the anchor, as in (5a). Every restaurant has at least one cook, so that the hearer knows that there is a cook, once he knows that there is a restaurant. Things are different in (5b), since not every room contains a chandelier. However, chandeliers are usually in rooms and rooms usually have lamps, which makes the link between the room and the chandelier a plausible one. Similarly, in (5c): usually, some tool is involved in a murder and a knife is not an uncommon one. As example (6) shows, if no plausible link to an entity in the discourse context can be found, the discourse is marked.

(6) We went to *this nice Italian restaurant* in Market Street, yesterday. ? **The baby orang utan** was really cute.

So in general, the use of an indirect anaphor requires that a plausible relation can be inferred to hold between its referent and that of the anchor. In order to generate indirect anaphors however, several further questions need to be answered.

First: is it possible to restrict the set of relations that may hold between indirect anaphors and anchors? An answer to this question would require an extensive corpus study which we have not yet carried out. On a smaller scale, Schwarz (2000) gives an extensive description of the different ways in which an indirect anaphor can be anchored to the context and note that *part_of* relations between the referent of the indirect anaphor and the anchor seem to be the most common type. In what follows we will restrict ourselves to examples of indirect anaphors which involve the *part_of* relation.

Second: What determines the set of *distractors*, i.e. the set of entities from which the definite description must distinguish the intended referent? In the previous section, we just assumed that all entities in the discourse model were potential distractors to each other. However, the examples in (7) show that entities that are not in the discourse model but that are inferrable from it may also act as distractors. In Example (7a), for instance, the cook of the restaurant in Market street is a distractor to the cook the speaker met at the party as is illustrated by the markedness of the discourse. Similarly, in Example (7b) the cooks of the two restaurants are distractors to each other.

(7) a. ?? I met *a cook* at Peter's party last week. He told me about this very outstanding restaurant in Market Street which is beautifully decorated and serves excellent food. **The cook** is from Paris and has won several prizes.

b. ?? There is *a Thai restaurant* in Market Street and a Chinese restaurant in Canal Street. **The cook** always goes for a walk in the park in the mornings.

On the other hand, other factors seem to play a role which in turn, constrain the set of distractors. Thus Erkü and Gundel (1987) have shown that focus[4] has some influence as illustrated by the following examples.

(8) a. We stopped for drinks at *the New York Hilton* before going to the Thai restaurant. **The waitress** was from Bangkok.

b. We stopped for drinks at the New York Hilton before going to *the zoo.* ? **The baby orang utan** was really cute.

c. We stopped for drinks at the New York Hilton before going to *the zoo.* **The baby orang utan at the zoo** was really cute.

Under the assumption that entities mentioned in a subordinated clause are not in focus, neither the Thai restaurant nor the zoo is in focus, while the New York Hilton is in focus in all three examples. In (8a) the New York Hilton serves, almost unambiguously, as an anchor for the indirect anaphor

[4]They assume a notion of focus, or psychological focus as Gundel (1999) calls it, that is similar to the backward looking center of centering theory (Grosz et al. 1995).

*the waitress*, although encyclopedic knowledge would make linking the waitress from Bangkok to the Thai restaurant more probable. Similarly, the structure of (8b) seems to suggest the New York Hilton as an anchor of *the baby orang utan* (although the zoo would provide a more plausible anchor). By adding additional information which makes the way the orang utan is anchored in the context explicit, as in Example (8c), linking the orang utan to the zoo becomes possible, though.

Although focus seems to play an important role in restricting the set of possible distractors, further additional factors are probably influential. For instance, Asher and Lascarides (1998) have given some evidence that discourse structure is one of them. In general however, it is still unclear exactly which factors interact in determining the set of distractors of a given referent.

To summarize, we can formulate the following constraints on the use of indirect anaphors:

**Familiarity.** the referent of an indirect anaphor must be familiar in that one of the following clause must be true:

- The referent is an element of the discourse model. In the taxonomy of Prince (1981), it is said to be *discourse old*.
- The anchor is focussed and the existence of the referent is inferrable from the common ground (cf. Example (8a)): in Prince's terms, it is an *inferrable*.
- The existence of the referent is inferrable from the common ground and the description links it into the discourse context (cf. Example (8c)): it is a *containing inferrable*.

**Uniqueness.** The referent must be uniquely identifiable.

## 4 Generating Indirect Anaphora

In Section 2, we presented a grammar which supports the generation of definite descriptions referring to some discourse old entity. We now show how the grammar can be adapted to allow the generation of indirect anaphors, i.e. of definite NPs that refer to some inferrable entity. The proposal draws on the analyses of indirect anaphora proposed in the formal semantics literature, e.g. Chierchia (1995), and makes use of the contextual reasoning/surface realization interface supported by the SPUD-like architecture to implement the pragmatic and semantic constraints governing the use of indirect anaphors.

In Chierchia (1995), definite descriptions are assigned the following semantics:

$$\lambda Q.Q(\iota x(N(x) \wedge R(x,y)))$$

$N$ represents the semantics of the nominal introduced by the definite description, $y$ is some discourse old entity and $R$ is some relation that can be infer to hold between the referent $x$ of the definite description and the referent $y$ of the anchor. Thus the definition says that a definite description introduces some

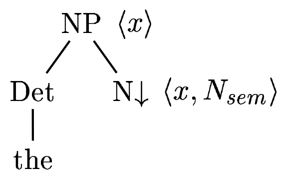


**pragmatic constraints:**

$$
\begin{array}{ll}
\textbf{familiarity} & d\_old(x) \\
& \vee \;\; \exists y\ [in\_focus(y)\ \&\ \exists z\ [N_{sem}(z)\ \&\ part\_of(z,y)]] \\
& \vee \;\; \exists y\ [d\_old(y)\ \&\ (N_{sem}(x) \rightarrow part\_of(x,y))]
\end{array}
$$

$$
\textbf{uniqueness} \quad \forall y[(N_{sem}(y)\ \&\ distractor(y,x)) \rightarrow x = y]
$$

Figure 4: Lexicon entry for *the*.

unique individual $x$ which is related by some relation $R$ to some discourse old entity $y$. Since $R$ can be resolved to the identity relation, the definition covers the direct as well as indirect anaphors.

We here assume a similar analysis but adapt it to further encode the various constraints on the use of indirect anaphors discussed in section 2. In particular:

- We fine-tune the familiarity condition encoded in Chierchia (1995)'s treatment to account for the focus-restriction discussed in the previous section and for containing inferrables.

- We restrict our analysis to those cases where $R$, the relation to be inferred between the anaphor and its anchor, is the *part_of* relation. This permits avoiding higher-order quantification and thus remaining with the use of a first-order theorem prover. In future research, we hope to identify through corpus studies, a set of relations prominently used in indirect anaphors, which would permit extending $R$ to a finite disjunction.

The lexicon entry for the definite article *the* is presented in Figure 4. The pragmatic constraints, which encode the conditions we formulated at the end of the previous section, will be checked during generation by posing an appropriate query to the theorem prover. We assume that as the discourse is processed, the discourse model is updated with the relevant information and in particular with predications of the type $d\_old(x)$ and $in\_focus(x)$ and the corresponding negative information ($\neg d\_old(x)$ and $\neg in\_focus(x)$[5]). The semantics is as follows: $d\_old(x)$ is true iff entity $x$ has been mentioned in previous discourse, and $in\_focus(x)$ is true iff entity $x$ is focussed in the sense of Erkü and Gundel (1987). Furthermore, the uniqueness constraint makes use of the predicate $distractor$. $distractor(x,y)$ is true iff $x$ is a distractor to $y$. As we have seen in

[5] Again, for clarity of the presentation, we will only list the positive information on focusing and discourse oldness in the following examples.

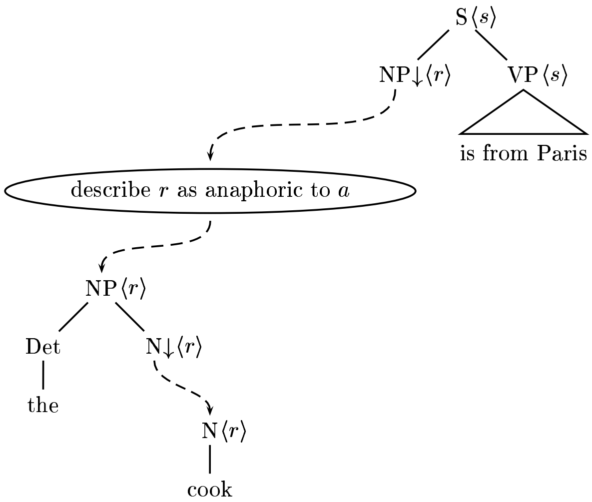


Figure 5: Generating an indirect anaphor (Example (9))

the previous section, the distractor set depends on the salience of the intended referent. We define $distractor(x, y)$ as follows:

- focussed items are potential distractors to any entity
  $\forall xy[in\_focus(x) \rightarrow distractor(x, y)]$
- entities inferable from focussed items are also distractors to any entity
  $\forall xy[\exists z[in\_focus(z)\ \&\ part\_of(x, z)] \rightarrow distractor(x, y)]$
- entities which are discourse old are distractors to all other discourse old entities which are not in focus
  $\forall xy[(d\_old(x)\ \&\ d\_old(y)\ \&\ \neg(in\_focus(y))) \rightarrow distractor(x, y)]$

We will now go through some of the examples that were given in Section 3 to illustrate the workings of the generator.

We start with a simple example characterized by the following setting:

(9) **Previous Discourse** *There is a great Italian restaurant in Market Street.*

**Discourse Model** ($DM$) $restaurant(r) \wedge in\_focus(r) \wedge \ldots$

**World Knowledge** ($WKL$)
$\forall x[restaurant(x) \rightarrow \exists y[cook(y)\ \&\ part\_of(y, x)]]$

**Speaker's Knowledge** ($SKL$) $cook(c) \wedge part\_of(c, r) \wedge from\_Naples(s, c)$

**Goal** describe $s$

The previous discourse has introduced entity $r$ into the discourse model and has also placed it into focus. The world knowledge contains the information that all restaurants have a cook, and the speaker also knows that the cook of that Italian restaurant introduced in the previous discourse is from Naples. The speaker's goal is to convey this fact to the hearer. Figure 5 outlines the generation process. After having selected the initial tree, expressing that some entity is from Naples, the generator has to deal with the open substitution node, requiring a noun phrase which describes entity $c$. $c$ is not in the discourse model, but since it is part of the speaker's knowledge that $c$ can be related by a *part_of* relation to entity $r$ which is an element of the discourse model, the generator will try to link $c$ to $r$ by an (indirect) anaphor. This subgoal (describe $c$ as anaphoric to $r$) licenses the selection of the definite article. After substituting "cook", the result can be combined with the first tree, giving us a syntactically complete parse tree yielding *The cook is from Naples.* Now, it has to be checked whether this tree is also contextually appropriate. The corresponding inference tasks are the following:

**familiarity**

$$\begin{aligned} DM \cup WKL \models\ & d\_old(c) \vee \\ & (\exists x\ [in\_focus(x)\ \&\ \exists y\ [cook(y)\ \&\ part\_of(c,x)\ ]]) \vee \\ & (\exists x\ [d\_old(x)\ \&\ (cook(c) \rightarrow part\_of(c,x))]) \end{aligned}$$

**uniqueness**

$\forall x[(cook(x)\ \&\ distractorOf(x,c)) \rightarrow x = c]$ is consistent with $DM \cup WKL \cup SKL$

They are both satisfied: (i) $r$ is in focus and world knowledge implies that $r$ has a cook, (ii) there are no cooks other than $c$ in the discourse model or inferable through world knowledge.

Now, let's have a look at a more interesting case. Here is what the discourse context will look like, after uttering the first sentence of Example (8a).

(10) **Previous Discourse** *We stopped for drinks at the New York Hilton before going to the Thai restaurant.*

**Discourse Model** $New\_York\_Hilton(a) \wedge in\_focus(a) \wedge restaurant(b) \wedge \ldots$

**World Knowledge** $\forall x[(restaurant(x) \vee bar(x)) \rightarrow \exists y[waitress(y)\ \&\ part\_of(y,x)]]\ \&\ \forall x[New\_York\_Hilton(x) \rightarrow bar(x)]$

**Speaker's Knowledge** $waitress(r) \wedge part\_of(r,a) \wedge from\_Bangkok(s,r)$

**Goal** describe $s$

The generator will produce a syntactically complete tree yielding *The waitress is from Bangkok.* Entity $a$, the anchor, is focussed, so that the familiarity constraint is satisfied. The uniqueness constraint requires that $r$ has no distractors. That is in fact the case here, because even though world knowledge tells us that restaurants may have waitresses as well, those waitresses are no distractors since entity $b$ is not focussed while $a$ is. If the intention had been to express that the waitress of the Thai restaurant was from Bangkok, neither of the constraints would have been satisfied. The generator would have added more information yielding e.g. *The waitress of the Thai restaurant is from Bangkok*, which satisfies all constraints.

As a final example, consider Example (6). The situation after uttering the first sentence will be as follows:

(11) **Previous Discourse** *There is a great Italian restaurant in Market Street.*

**Discourse Context** $restaurant(r) \wedge infocus(r) \wedge \ldots$

**World Knowledge** $\forall x[restaurant(x) \rightarrow \exists y[cook(y)\ \&\ part\_of(y,x)]]$

**Speaker's Knowledge** $baby\_orang\_utan(c) \wedge part\_of(c,r) \wedge very\_cute(s,c)$

**Goal** describe $s$

Here, the use of the indirect anaphor *the baby orang utan* is not licensed due to a failure of the familiarity constraint: it does not follow from world knowledge and/or the discourse model that the restaurant has a baby orang utan. A more informative definite description is therefore generated through adjunction e.g., *the baby orang utan from the restaurant.*

## 5 Conclusion

Although we are often not aware of it, natural sounding texts are texts which leave much of the communicated information implicit or in Stone and Webber's (1998) terms, texts that display *textual economy.* As Stone and Webber (1998) showed, textual economy imposes some strong constraints on the generation architecture. In particular, it requires that reasoning interact with surface realization so that information that is inferrable from the current discourse model can be omitted in the phonological realization of the message being communicated.

In this paper we've concentrated on one of the many natural language constructs that support textual economy – namely, definite descriptions that function as indirect anaphors – and we have shown how a SPUD like architecture could be used to generate them. At a more general level however, the observation is that presuppositions and anaphors are an important means to achieve textual economy and that therefore, their investigation from an NLG perspective is of prime importance if generators are to be developed which produce natural sounding texts. In future work, we intend to extend our investigation to further presuppositional items and to the modeling of their interaction within a generation scheme.

Another direction of research we intend to pursue in more detail concerns the inferences that arise when presuppositions/anaphors are interacting with world knowledge as in this paper. What reasoning tools are best suited for these inferences? Can, e.g., more efficient inference schemes, such as the ones developed in the field of description logics, be used to support the type of reasoning involved in supporting textual economy. Another important issue is also how the relevant world knowledge can be automatically selected (rather than entered by hand).

## References

Asher, N. and A. Lascarides (1998). Bridging. *Journal of Semantics 15*(1), 83–113.

Brew, C. (1992). Letting the Cat out of the Bag. In *Proceedings of COLING '92.*

Carroll, J., A. Copestake, D. Flickinger, and V. Poznzński (1999). An Efficient Chart Generator for (Semi-)Lexicalist Grammars. In *Proceedings of ACL*.

Chierchia, G. (1995). *Dynamics of Meaning: Anaphora, Presupposition and the Rhetory of Grammar*. Chicago: University of Chicago Press.

Clark, H. H. (1975). Bridging. In Schank and Nash-Webber (Eds.), *Theoretical Issues in Natural Language Processing*. MIT Press. Reprinted in: Johnson-Laird and Wason (eds.) *Thinking. Readings in Cognitive Science.* Camebridge University Press. 1977.

Copestake, A., D. Flickinger, and I. Sag (1999). Minimal Recursion Semantics: An Introduction. Draft. Available at `http://www-csli.stanford.edu/~aac/papers/newmrs.ps`.

Dale, R. and N. Haddock (1991). Content Determination in the Generation of Referring Expressions. *Computational Intelligence 7*(4).

Dale, R. and E. Reiter (1995). Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions. *Cognitive Science 18*, 233–263.

Erkü, F. and J. Gundel (1987). The Pragmatics of Indirect Anaphora. In J. Verschueren and M. Bertuccelli-Papi (Eds.), *The Pragmatic Perspective*, pp. 533–545. John Benjamins Publishing Company.

Grosz, B., A. Joshi, and S. Weinstein (1995). Centering: A Framework for Modelling the Local Coherence of Discourse. *Computational Linguistics 21*(2), 203–225.

Gundel, J. K. (1999). On Different Kinds of Focus. In P. Bosch and R. van der Sandt (Eds.), *Focus. Linguistic, Cognitive, and Computational Perspectives*, pp. 293–305. Camebridge University Press.

Hobbs, J. (1985). Ontological Promiscuity. In *Proceedings of ACL*, pp. 61–69.

Joshi, A. and Y. Schabes (1997). *Tree-Adjoining Grammars*, Chapter 2. Berlin: Springer-Verlag.

Kay, M. (1996). Chart Generation. In *34th Annual Meeting of the Association for Computational Linguistics*, Santa Cruz, California, pp. 200–204.

Prince, E. F. (1981). Toward a Taxonomy of Given-New Information. In P. Cole (Ed.), *Radical Pragmatics*, pp. 223–255. Academic Press.

Schwarz, M. (2000). *Indirekte Anaphern in Texten: Studien zur domänegebundenen Refernz und Kohärenz im Deutschen.* Tübingen: Niemeyer.

Stone, M. and C. Doran (1997). Sentence Planning as Description Using Tree Adjoining Grammar. In *Proceedings of ACL*, pp. 198–205.

Stone, M. and B. Webber (1998). Textual Economy Through Close Coupling of Syntax and Semantics. In *Proceedings of INLG 1998*, pp. 178–187.

Vieira, R. and M. Poesio (1996). Processing definite descriptions in corpora. In *Proc. of DAARC.*