# Integrating Referring and Informing in NP Planning

Michael O'Donnell, Hua Cheng,
Department of Artificial Intelligence,
University of Edinburgh,
80 South Bridge, Edinburgh, Scotland, EH3 9JF.

Janet Hitzeman
Human Communication Research Centre,
University of Edinburgh
2 Buccleuch Place, Edinburgh, Scotland, EH8 9LW.

## Abstract

Two of the functions of an NP are to refer (identify a particular entity) and to inform (provide new information about an entity). While many NPs may serve only one of these functions, some NPs conflate the functions, not only referring but also providing new information about the referent. For instance, *this delicious apple* indicates not only which apple the speaker is referring to, but also provides information as to the speaker's appreciation of the apple.

This paper describes an implemented NP-planning system which integrates informing into the referring expression generation process. The integration involves allowing informing to influence decisions at each stage of the formation of the referring form, including: the selection of the form of the NP; the choice of the head of a common NP; the choice of the Deictic in common NPs; the choice of restrictive modifiers, and the inclusion of non-referring modifiers. The system is domain-independent, and is presently functioning within a full text generation system.

## 1 Introduction

Two of the functions of an NP are to refer (identify a particular entity) and to inform (provide new information about an entity). In most cases, a given NP may serve only one of these functions. However, in some cases, the writer/speaker may choose to conflate the functions, providing an NP which not only refers but also provides new information about the referent. For instance, *this delicious apple* indicates not only which apple the speaker is referring to, but also provides information as to the speaker's appreciation of the apple.

Most of the work on NP planning has considered only the referring function of the NP (e.g., Dale 1988, 1989; Reiter 1990; Reiter & Dale 1992; Horacek 1995). Appelt (e.g., Ap-pelt 1985; Appelt & Kronfeld 1987) has considered the question of integrating referring and informing, although rather briefly, and without much detail. This paper will extend upon his discussion, and describe its role in ILEX, a text generation system which delivers descriptions of entities on-line from an underlying knowledge-base (see Mellish *et al.* 1998). ILEX is at present generating descriptions in the museum domain, in particular, that of 20th Century jewellery.

Our focus on this topic has grown out of the need to integrate two strands of research within ILEX. One strand involves the work on anaphora by Janet Hitzeman. She implemented a module to construct contextually appropriate referring expressions within ILEX, based on Centering Theory (Grosz *et al.* 1986). See Hitzeman *et al.* 1997.

The second strand involves the aggregation module (implemented by Hua Cheng, see Cheng *et al.* 1997). The task of this module is to repackage discrete informational units into single complex sentences. She is presently exploring the aggregation of information into the NP, for instance *this gold and silver ring, designed by King.*

These two functions, the referring and the informing, interfere with each other, to the extent that each wishes to control the construction of the NP form. These tasks thus need to pay regard to each other, and this paper, and the implementation it describes, are an attempt to answer this need.

Appelt's approach seems to be to build an NP

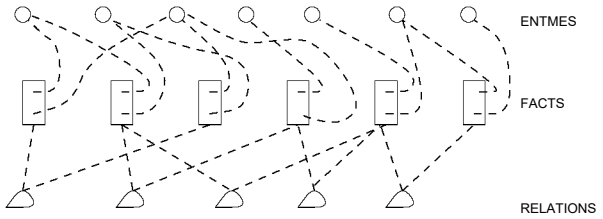


Figure 1: The Information Graph

for referring, then either modify the elements (e.g., substitution of the head noun) or fill unused structural slots with non-referring information. However, we have found that the two tasks of referring and informing can be more highly integrated, with each decision within the construction of the NP taking into account the needs of both tasks, rather than satisfying the referring function first, then looking to the informing function. In other words, we follow an integrated rather than pipeline approach.

Section 2 will describe how information is represented in ILEX. Section 3 describes the interface between the text-planner and the NP- planner, the input specification for the NP-planner. Section 4 discusses the syntactic structure of the NP, and which syntactic positions allow informing without interfering with referring. Section 5 details the referring expression generator which integrates referring and informing goals. An example of the generation process is given in section 6 and section 7 summarises.

## 2 Information Representation

To properly describe our NP-planning process, we need to describe how information is represented in ILEX. Domain knowledge is represented in terms of an *information graph,* which represents the *entities* of the domain, the interrelation between these entities (called *facts);* and the *relations* between these facts (e.g., a causal relation between two facts). Figure 1 shows an abstract representation of an information graph. At present, relations between facts are not used in the NP-planner, so will not be discussed further here.

Initially, the information graph representation was developed for text planning. However, following a suggestion from Alistair Knott, we have found it useful to use it for NP-planning as well.

### 2.1 Entities

Entities represent the objects of the domain. In the Museum domain, this includes not only the museum artifacts, but their designers, the materials they are made from, the styles they are made in, the periods they belong to, the locations of their manufacture, etc.

Entities are typically *specific* entities: real-world individuals. However, some of the entities will be *generic* entities, those representing classes of entities, such as *Art-Deco jewellery.* We also cater to *sets* of entities, which can be realised through either plural anaphora, or conjunctive NPs.

### 2.2 Facts

A fact in ILEX represent a relation between two entities. These relations may be *processual,* e.g., that X made Y: maker(J-999, King01); or *stative* (e.g., that X is a type of Y: isa(j-999, generic-brooch).

Each fact is represented as an attribute-value structure as below:

| | |
|---|---|
| -Pred: Argl: | |
| Arg2: | "maker" J-999 King01 |
| Polarity: | positive |
| Status: | indefeasible 0 6 8 |

Assimilation: Importance: _Interest:

Note that apart from the predicate and argument information, several other fields qualify the informational status of the fact, including the *polarity* (whether or not the relation holds), and *defeasibility* (distinguishing between hard facts about the entity, and those which are only tendencies, e.g., *Art-Deco jewellery tends to be made of enamel* (see Knott *et al.* 1997 for discussion of defeasibility in ILEX). The remaining fields, having a stronger affect on NP-planning, include:

- Assimilation: the degree to which the system considers the user to have understood the information. This is of particular importance to reference, since adequate reference usually requires the user to know the information used for reference (see later for exceptions).

- Importance/Interest: the degree to which the fact is considered important for the system to deliver to the user, and the system's estimate of the degree of interest to the user. These values are represented for each predicate type as a whole, and vary for different user models. These values are used when selecting the facts to use to produce a unique reference.

## 3 NP Specification

One of our goals in the design of the ILEX NP-planner was to provide a clean interface between text-planning and NP-planning, such that the text planner can specify what it wants from the NP without needing to know about syntax at all. To this end, we have developed a two-level specification of the NP, one at the semantic level, and one at the syntactic level. The text-planner specifies the NP only at the semantic level, leaving details of syntax totally to the NP-planner.

### 3.1 The NP Specification Interface

The interface between the text-planner and the NP-planner is in the form of an attribute-value matrix, the attributes of which are:

1. Cat: the function of the NP being produced. The NP-planner allows a wide range of NP functions, not only referring, shown in figure 2 and discussed below:

   *(a) Referring:* an NP which uniquely or non-uniquely refers to the referent. More delicate options can be specified, such as *refer-by-name, refer-by-type,* or *refer-by-pronoun;* and also whether the reference should be unique or not.

   *(b) Describing:* an indefinite NP giving an arbitrary number of the entity's attributes without attempting to be definitive.

   *(c) Classifying:* an indefinite NP which provides only the superclass of the item, e.g., *this is* a brooch

   *(d) Defining:* for generic entities, an NP which provides the entities defining characteristics, e.g., *a necklace is* an item of jewellery worn around the neck

   *( ) Eliciting:* a wh- NP for the referent. Eliciting can be selective, e.g., *which designer* or non-selective, e.g., *Who.*

   If *referring-np* is selected, various sub-types of reference can also be preselected through this slot (the specification of the *Cat* can be logically complex).

   If no preselection is made by the text-planner, the system will decide NP function on the basis of constraints and defaults. For instance, in the usual case, the text-planner will specify only *referring-np,* and leaves it up to the NP planner to choose the exact type of reference.

   In some cases, the system will override the preselection if it is incompatible with the referring environment. For instance, if the text-planner specifies *refer-by-name,* but no name information is provided for the entity, then a choice between *refer-by-type* and *refer-by-pronoun* will be made. The NP-specification can thus be seen to offer a flexible interface, allowing the text-planner to determine the amount of control it desires to assert.

2. Sem: the referent of the NP, an entity in the information graph (or a set of such entities if desired, realised as either a coordinated NP, or using plural anaphora).
3. Syn: the slot to be constructed by the NP-planner, a syntactic structure.
4. Orth: the slot to hold the eventual surface string for the NP. If the sentence planner provides a filler for this slot, then NP-planning is avoided and the string is used. ILEX thus allows canning of NPs when needed.
5. Agenda: a list of fact-ids which are to be incorporated into the NP if possible. The aggregation module uses this slot to state its requirements from the NP, which facts the NP is to express.
6. Restrictions: a list of fact-ids which should not be used in the NP. For instance, we might wish to avoid generating the sentence *Other jewels designed by Jessie M. King include a brooch designed by Jessie M. King.* To avoid such sentences, we place

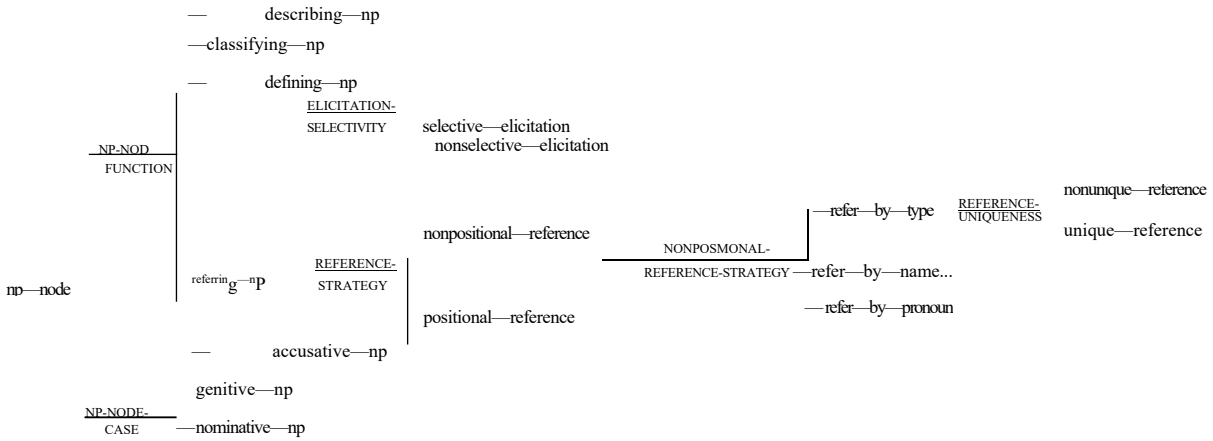


Figure 2: System Network for Nominal Function

the fact-id of the 'designer' fact into the Restrictions field for the mention of the brooch. The designer fact will not then be used as a referential restrictor.

A sample NP-specification is shown below:

| | |
|---|---|
| Cat: Sem: | unique-reference |
| Agenda: | J-999 |
| Restrictions | (FN-48 FN-56) |
| : | (FN-59) |

...where FN-48 and FN-56 are facts to include in the reference, and FN-59 is a fact to avoid. **Problems of Modularity:** One of the problems of a clean separation between NP-specification and NP-planning is that it might not be possible to incorporate all facts on the informing agenda into the NP. However, given that NPs syntactically allow any number of non-referring post-modifiers, our planner will handle any arbitrary number of facts on the agenda. However, in terms of intelligibility, too many post-modifiers will produce unintelligible NPs. We make the simplifying assumption that the text planner uses some metrics (simple or complex) to avoid overloading the NP. We are merely providing a mechanism to support NP-planning once the agenda is specified.

### 3.2 Focus Spaces

Apart from the text-planner's specification of the NP-form, the text-planner also maintains some variables concerning the current referring environment, mainly in terms of various focus spaces. These various spaces are:

I. **Mentioned Entities:** entities mentioned at some point within the discourse.

2. **Shared Entities:** entities which the system assumes the addressee to know about. These entities include world-knowledge (e.g., *Ronald Reagan),* but also entities mentioned previously in the discourse (mentioned-entities), and entities in the immediate context (focal-objects). Entities in this space are potential confusers for definite reference with *the.*

3. **Focal Objects:** the focal space includes a set of entities which may potentially be referred to as *this x.* Firstly, we have the *Prior-Cb* (backward looking centre, usually the subject of the prior sentence). Entities directly related to this may also be focal. This is also called the local focus in our system. Then there is the page-focus, the focus of the current object description in the ILEX system, e.g., *this brooch.* Other objects are also focal by being part of the immediate context of the reader/writer. In a web-browsing environment, this might include the current page *(this page),* or parts of the page *(this picture).*

In addition to the focal spaces, there are also variables holding individual focal objects, including the *Cb* and *Prior-Cb, Page-Focus* and *Discourse-Focus*. We allow pronominalisation only when the object being referred to is Prior-Cb, which seems to produce coherent reference.

## 4 NP Structure for Referring and Informing

The NP-Planner has distinct procedures for each of the NP functions, one for classifying, one for referring, one for eliciting, etc. Due to lack of space, we will focus from now only on NPs primarily serving a referring function. Other nominal functions will be covered in a later paper.

The issue of how to construct a referring NP is well explored. The issue remaining is how non-referring information can best be incorporated into referring NPs. This section will look at the locations in the NP which can express non-referring information, without interfering in the reference of the NP. The following section will describe the algorithm which allows referring and informing to be integrated.

### 4.1 Nucleus and Satelite Composition of the NP

We consider the basic structure of the NP to fall into two components: a *nucleus,* which performs the nominal function of the NP, and optional *satelites,* where additional information can be placed.' The nucleus of the NP consists of all slots before the head, and the defining post-modifying slots (e.g., defining relative clauses, or prepositional phrases). The satelite elements are typically realised by indefinite NPs, or non-defining relative clauses (or complexes of such), e.g., *[Nuc: this brooch], [Sat: designed by Jessie M. King],* or *[Nuc: King], [Sat: a Scottish designer]*.

In our model, all referring is performed by the nucleus — the satelite(s) are non-defining, i.e., perform only an informing role. However, as will be discussed below, the nucleus can also contain non-referring information.

*Nucleus* and *satelzte* are terms taken from RST (Rhetorical Structure Theory, e.g., Mann & Thompson, 1987), although usually applied to the relations between sentences.

### 4.2 The Structure of the Nucleus

The range of slots in a systemic analysis of the NP, in the order they typically appear (after Halliday 1994), appears below, and figure 3 shows a typical NP structure:

```
(Deictic) - (Numerator) - (Epithet*) -
  (Classifier*) - Thing - (Qualifier*)

Key: 0 - optionality
   * - any number of this slot may
         occur
```

### 4.3 Informing within the Nucleus

While primarily for referring, non-referring information can sometimes be included in the nucleus without interfering with the referring function. For instance, we can add information to an already uniquely-referring NP, making the reference more explicit: *The [granny smith] apple on the table; this [enjoyable] book*. The degree to which informing and referring can be so integrated varies from domain to domain.

The major constraint we seem to face is that there is a degree of expectation under conversational implicature that the speaker refers using information known to the addressee (see Dale & Reiter 1996). Thus, in a situation where only one apple is visible, if I say *pass me the Spanish apple,* the addressee might be confused by the inclusion of the superfluous information, and perhaps think there must be another apple somewhere.

However, in some registers this form of reference seems to offer no problems. Appelt (1985) mentions the case of the speaker pointing at some implement and saying *use the wheel-puller*. The addressee, not knowing the name, but having the item identified through pointing, accepts the naming. We thus have an NP whose head-noun is not serving a referring function, but rather an informing function, since the referring function was otherwise fulfilled.

The newspaper genre is particularly strong on this type of reference, as shown by the newspaper article below:

*Student* fights for life after flat fire:
*A young student* was today fighting for

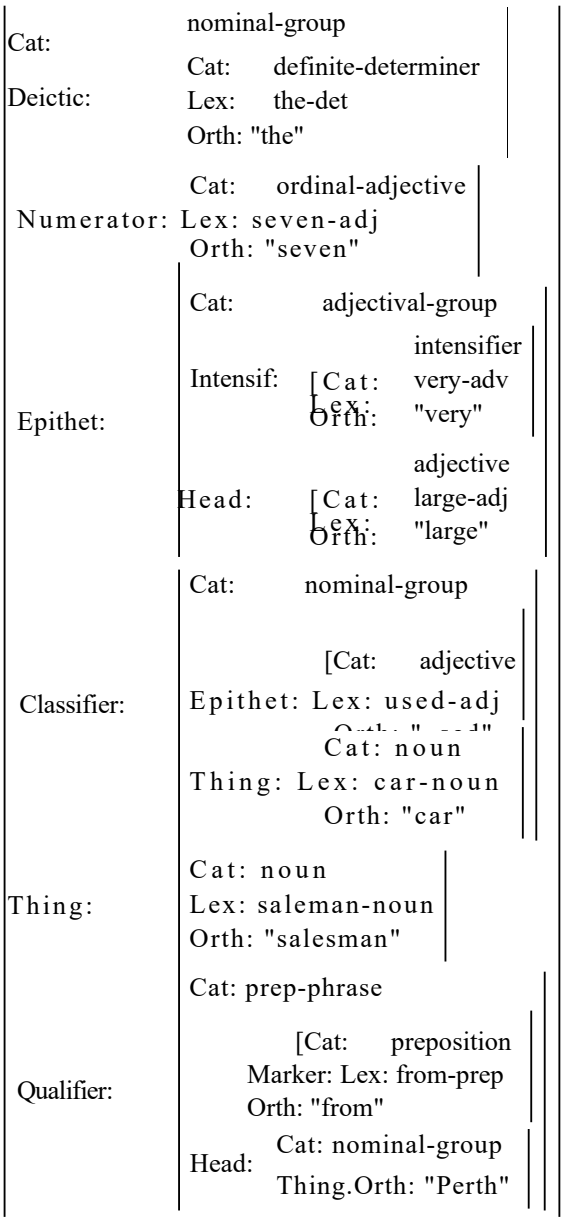


Figure 3: A Sample NP Structure

her life after fire ripped through her Edinburgh flat. *Nicola Graham* is in a "serious but stable" condition at the specialist burns unit in St John's Hospital, Livingston. Firefighters suspect the blaze may have been started by a dropped cigarette in *Miss Graham's* bedroom. The *19-year-old* was transferred from Edinburgh Royal Infirmary to St John's for emergency burns treatment. ..."

The sequence of references to the student successively add new information: *A young student:* Age and occupation; *Nicola Graham:* Name; *Miss Graham:* Marital status; *the 19-year-old:* Age. This writer is not depending on assimilated information to refer, but, depending on the lack of potential confusors, is successfully referring with new information. While this style is more typical of newspaper reporting, where compact information delivery is important, it is still an issue which needs to be addressed in any NP-planner.

In the register of museum object descriptions, it seems that the degree to which new information can be included in the nucleus is limited. New information seems not to be appropriate in the Deictic, Classifier, Thing or Qualifier slot, but is generally allowed in the Numerative and Epithet slots. This makes some degree of sense, since these slots are the least restrictive. The Numerative can be used restrictively when used contrastively, e.g., *the five cups (but not the set of three),* but this is rare. Epithets generally add qualitative information, and are thus less restrictive.[2]

Another approach is to examine the semantic types of pre-modifier elements, to see which, when inserted for informing reasons, seem to interfere with the referring function. We have found some of our fact-predicates interfere more, some less. As a result of this, we maintain a list of fact-predicates which are judged, for the current domain, to be suitable for pre-modifier slots without interfering with reference. This allows us to produce, for instance, *this [important] designer; the [gold and enamel] brooch designed by King; the [quite influential] Art-Deco style.*

## 5 The Planning Algorithm

The tension in planning definite reference derives from the need to serve both the referring and informing functions. The referring function is mainly concerned with the *Sem* slot of the NP-specification: the task is to point uniquely at the filler of this slot, distinguishing it from all other entities. The planner may need to use any or all of the syntactic slots of the nucleus

[2] A study of 20 randomly chosen museum descriptions , undertaken by Cheng, from four museums and galleries revealed that only 1/3 of Epithets act restrictively.

to do this.

On the other hand, the informing function is primarily concerned with the *Agenda* slot — the function is satisfied if all the facts in this slot are expressed somewhere in the NP (whether nucleus or satelite). While these facts can be placed in satelite position, it is often more coherent to place them within the nucleus.

As such, the two functions are in competition for the syntactic slots, and structural decisions good for one function may be sub-optimal for the other. The usual approach is to allow the referring component to go first, generating the desired referring form. Then, the needs of the informing component are fitted into this structure.

However, we have found it far more rewarding to allow all syntactic decisions to be mutually negotiated between the two functions. Below, we describe the definite description algorithm as used in ILEX.

### 5.1 Construction of NP

The steps of building the NP are as follows:

**Build Nucleus:** Since we wish to choose a referring expression which opportunistically serves some of the informing function, we will build the nucleus on the basis of i) the referential context, and ii) the agenda of things to say:

1. *Location of Assimilated Agenda:* the facts on the agenda need not all be new information (unassimilated) — the text-planner may place previously given information on the agenda, perhaps for some pragmatic reason, e.g., to ensure that the addressee is aware of some fact at this point of the discourse.

   Assimilated facts on the agenda play an important role in our algorithm, since we will use them preferentially for referring.

2. *Choose Referential form:* we need to choose between *refer-by-name, refer-by-pronoun* and *refer-by-type,* on basis the referential context. Where the referential context allows more than one choice, we refer to the assimilated-agenda to help. For instance, if a *Name* fact is the agenda (and assimilated), we might favour refer-by-name, if *Gender* is the sole fact on the agenda, pronominal might be favoured. With several (assimilated or unassimilated) facts on the agenda, a common-group is preferred since it offers more opportunities for inclusion of facts (although proper-np expression also supports non-referring post-modification).

3. *Choose Head Noun:* for common noun-phrases, the head noun will be chosen from the most specific assimilated *isa* fact about the object. However, an assimilated *isa* fact on the agenda is allowed to override the default.

4. *Choose Determiner:* for common noun-phrases, the determiner will be chosen on the basis of the objects focal status (e.g., *this* if focal, *the* otherwise). If a fact specifying the owner of the entity being expressed is on the agenda (and assimilated), then a genitive deictic will be used.

5. *Choose Restrictive Modifiers:* if the common-noun-phrase form was selected, then we need to determine which modifiers are to be included to produce a unique reference. For instance, if we have the *Style* fact of a jewel on the agenda, and it is assimilated, then that will be preferred as a restrictive modifier. See below for more detail.

6. *Fill in Unused Slots:* When we have a functioning referring form, then we can add information from the agenda into the unused slot, e.g., *this book* + *enjoyable* *this enjoyable book*. The system is provided with a list of fact-predicates which can be expressed in pre-modifier slots, e.g., in the Jewellery domain, Materials — *this gold and enamel brooch,* Fame — *the famous designer called Jessie M. King.*

**Add Satelites:** Any information which was not consumed in the nucleus can now be placed into non-defining satelites, e.g., *[Nuc: Jessie M. King], [Sat: a Scottish designer].*

### 5.2 The Confusor Set

The confusor set is the set of entities which a partially constructed NP unintentionally refers to. For instance, if we have only selected the head noun, *brooch,* then the confusor set is all

brooches known to the system apart from the intended referent.

The process of building a referring expression can be seen as successively reducing the confusor set until it is empty. Assuming a common-np, the steps in this reduction are:

1. *Set Initial Confusors* based on focus status:
   - *Sem is member of Focal-objects:* Deixis: proximal (this/these). Con-fusors: Focal-objects, without Sem.
   - *Entity has an assimilated Owner fact:* Deixis: Owner. Confusors: other entities owned by Owner.
   - *Entity has been mentioned already on this page:* Deixis: nonproximal (the). Confusors: Other entities mentioned on page.
   - *Default:* Deixis: nonproximal (the). Confusors: Shared-Entities.
2. *Restrict on class:* Choose a head noun for the item, and eliminate all confusors which do not take the class.
3. *Add Restrictive Modifiers:* Choose a subset of assimilated facts which eliminates all confusors (see next sub-section).
4. *Insert "One of" if needed:* (not yet implemented) If insufficient assimilated facts to eliminate all confusors, insert "one of' or "another of' into the Pre-Deictic slot. "another of' is used if the confusor is already mentioned on the page.

### 5.3 Choosing Restrictive Modifiers

There are a number of strategies used to select the optimal set of restrictive modifiers to produce unique reference. There seems to be two main approaches. One attempts to select the smallest subset of modifiers which uniquely refers (e.g., Reiter 1990; Dale 1989). A solution which offers better computational complexity is based on the premise that some fact-types are better suited as restrictive modifiers than others, and thus restrictive modifiers are chosen by incrementally taking the next modifier from the list (e.g., Reiter & Dale 1992).

In ILEX, we follow the incremental approach, adding restrictors in order from our (domain-dependent) list (but only if the restrictor eliminates some confusors). We have found that ordering restrictors in terms of goodness, the NPs we generate are of better quality.

| Information | On Agenda? | Assim.? |
|---|---|---|
| Class: apple | no yes | yes |
| Owner: John | no yes | no |
| Color: Red | yes no | yes |
| Variety: Granny Smith | | no |
| Position: on table | | yes |
| Taste: good | | no |

Figure 4: An Example Information Base

The need to integrate informing into the process changes the process slightly. As stated above, the text-planner is allowed to place assimilated, as well as unassimilated, information on the agenda. If this has not happened, then we use the standard incremental strategy. However, if the text-planner has placed assimilated information on the agenda, then our planner places these at the front of the preferred-restrictors list.

We note however, that there are cases where, while the text-planner may want the fact to be included, the fact is not a suitable restrictor. For instance, including the *place-of-wearing* fact on the agenda could result in an NP like *the gold necklace that is worn around the neck*. However, since the place-of-wearing does not actually discriminate (given all necklaces are worn around the neck), the fact was not used restrictively, and was later realised in a satelite of the NP, e.g., *the gold necklace, which is worn around the neck*. However, there may be facts which are partially restrictive, but nevertheless poor candidates for restriction. Our algorithm does not cater to these cases as yet.

## 6 An Example

Agenda: Assume we are talking about an apple, and have the information as in figure 4 to express. In short, the facts on the agenda are: Owner, Variety, and Position.

Referential Context: Assume also that we have several red apples, but only one on the table. The apple above has been mentioned, but not for a while, with other apples mentioned since.

Stage 1: Building the Nucleus:

1. *Choose Referential form:* Since the item is not the *Cb,* we cannot use a pronoun. Since it doesn't have a proper-name, proper-noun reference is also out. We are forced to use a common noun-phrase.

2. *Choose Restrictive Modifiers:* We have a set of potential referential restrictors of: (Class Owner Color Variety Position Taste). Of these, we can only refer using assimilated roles, so we can use: (Class Color Position). We also have the agenda role-list of: (Owner Variety Position). of which the assimilated items are: (Position). Since the Class fact is assimilated, we automatically take the class as the head of the referring NP, e.g., *the apple.* This is not however unique, so we need to add more restrictions. We use the first (and only) item in the assimilated agenda: Position: *the apple on the table.* This happens now to be unique, so we have a functional referring NP.

3. *Fill in Unused Slots:* This leaves two facts unexpressed: Owner and Variety. The Owner predicate can normally be expressed in one of two slots of the nucleus:

   - the *Deictic* slot e.g., *John's apple on the table;* or,
   - the *Qualifier* slot (after the Head noun, e.g., *the apple that John owns on the table.* (I assume here that nonrestrictive relative clauses are always satelites, discussed below).

   In both of these slots, the inclusion of unassimilated *Owner* information seems to mess up the reference, seemingly because it implies the reader should already know the ownership. We thus leave the Owner role for expression in a satelite position (realised as a non-restrictive relative clause, e.g., *the apple on the table, which john owns).*

   The *variety* fact can be realised best through the Classifier slot, e.g., *the Granny Smith apple on the table.* This does not seem to interfere with the referring function, so this fact-type would occur on our list of facts which can appear in a pre-modifier slot without interfering with the referring function.

This stage thus ends with the referring slot consisting of: *the Granny Smith apple on the table.* We have only one item left on the agenda, the *Owner* fact.

Stage 2: Adding Satelites The *Owner* fact can be incorporated into the NP as a satelite (as a non-referring relative clause), e.g., *the Granny Smith apple on the table, which John owns.*

## 7 Conclusions

We have improved on the integration of referring and informing within NP generation by allowing informing to influence decisions at each stage of the formation of the referring form. Previous np-generation systems only satisfy informing goals after the referring form has been determined.

The points of intervention in the referring process include: the selection of the form of the NP; the choice of the Deictic in common NPs; and choice of restrictive modifiers. Information remaining on the agenda at this point is expressed in non-referring slots of the NP, in particular, the Epithet slot, or non-referring post-modifier slots. The use of an Agenda slot in the NP-specification is the main addition, which allows the Aggregation component to interface with the referring expression generator.

## 8 Acknowledgements

The ILEX project is supported by EPSRC grant GR/K53321. Thanks to Alistair Knott for ideas contributing to the implementation, and to Renate Henschel for valued comments on this paper.

## 9 References

Appelt, D E. 1985. "Planning English Referring Expressions". Artificial Intelligence, 26, pp 133.

Appelt, D and Kronfeld A. 1987. "A Computational Model of Referring". In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence,* Milan, Italy, August 2328, 1987, pp 640-647.

Cheng, Hua, & Chris Mellish. 1997. "Aggregation in the Generation of Argumentative texts". *Proc. of PhD Workshop on Natural Language Generation, 9th European Summer School in Logic, Language and Information (ESSLLI97).* Aug. 1997, France.

Dale, Robert. 1988. *Generating Referring Expressions in a Domain of Objects and Processes.* Ph.D. Thesis, Centre for Cognitive Science, University of Edinburgh.

Dale, Robert. 1989. "Cooking up referring expressions". *Proceedings of ACL-89.* Vancouver, pp 68-75.

Dale R. and E. Reiter. 1996. "The Role of the Gricean Maxims in the Generation of Referring Expressions". *Working Notes for the AAAI Spring Symposium on Computational Implicature,* Stanford, 1996, pp 16-20.

Grosz, Barbara J., Aravind K. Joshi and Scott Weinstein. 1995. "Centering: A Framework for Modeling the Local Coherence of Discourse". *Computational Linguistics,* Volume 21, Number 2, June 1995, pp 203-225.

Halliday, M.A.K. 1994. *Introduction to Functional Grammar.* 2nd edition. London: Edward Arnold.

Hitzeman, Janet, Chris Mellish & Jon Oberlander. 1997. "Dynamic Generation of Museum Web Pages: The Intelligent Labelling Explorer". *Proceedings of the Museums and the Web Conference,* Los Angeles, March 1997.

Horacek, Helmut. 1995 "More on Generating Referring Expressions". *Proceedings of the 5th European Workshop on Natural Language Generation.* Leiden, The Netherlands.

Knott, Alistair, Michael O'Donnell, Jon Oberlander, Chris Mellish. 1997. "Defeasible Rules in Content Selection and Text Structuring". *Proceedings of the 6th European Workshop on Natural Language Generation.* March 24 - 26, 1997 Gerhard-Mercator University, Duisburg, Germany.

Mann, William & Sandra Thompson, 1987. "Rhetorical Structure Theory: A Theory of Text Organization". Technical Report ISI/RS-87-190.

Mellish, C., O'Donnell, M., Oberlander, J. and Knott, A. 1998 "An architecture for opportunistic text generation". *Proceedings of the 9th International Workshop on Natural Language Generation.* 5-7 August 1998. Prince of Wales Hotel, Niagara-on-the-Lake, Ontario, Canada.

Reiter, E. 1990 *Generating appropriate natural language object descriptions.* PhD Thesis, Harvard University.

Reiter, E. and Dale R. 1992 "A Fast Algorithm for the Generation of Referring Expressions". *Proceedings of COLING-92.* Nantes, 1992.