

Integrating Referring and Informing in NP Planning

Michael O'Donnell † and Alistair Knott ‡ and Janet Hitzeman ‡ and Hua Cheng †

† Department of Artificial Intelligence, 80 South Bridge

‡ Human Communication Research Centre, 2 Buccleuch Place
University of Edinburgh, Scotland

March 15, 1998

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.

Submission Type: Paper

Topic Areas: NP-planning, object description generation

Author of Record: Mick O'Donnell, micko@dai.ed.ac.uk

Under consideration for other conferences (specify)? no

Integrating Referring and Informing in NP Planning

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). However, Appelt (e.g., Appelt 1985; Appelt & Kronfeld 1987) has considered the question of integrating referring and informing. 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 Knott *et al.* 1996). 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 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, allowing the informing task to have influence at every choice-point in the construction of the referring process.




Figure 1: The Information Graph

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 inter-relation 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 involves 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 entit-

ies 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:	"maker"
Arg1:	J-999
Arg2:	King01
Polarity:	positive
Status:	indefeasible
Assimilation:	0
Importance:	6
Interest:	8

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

If the referring function is selected, various sub-types of reference can also be preselected through this slot. If no preselection is made by the text-planner, the system will decide NP function on the basis of constraints and defaults. 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 the entity has no name provided, a choice between *refer-by-type* and *refer-by-pronoun* will be made.

2. **Sem:** the referent of the NP, an entity in the information graph.
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. This is used by the aggregation module to specify which facts 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 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:	unique-reference
Sem:	J-999
Agenda:	(FN-48 FN-56)
Restrictions:	(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

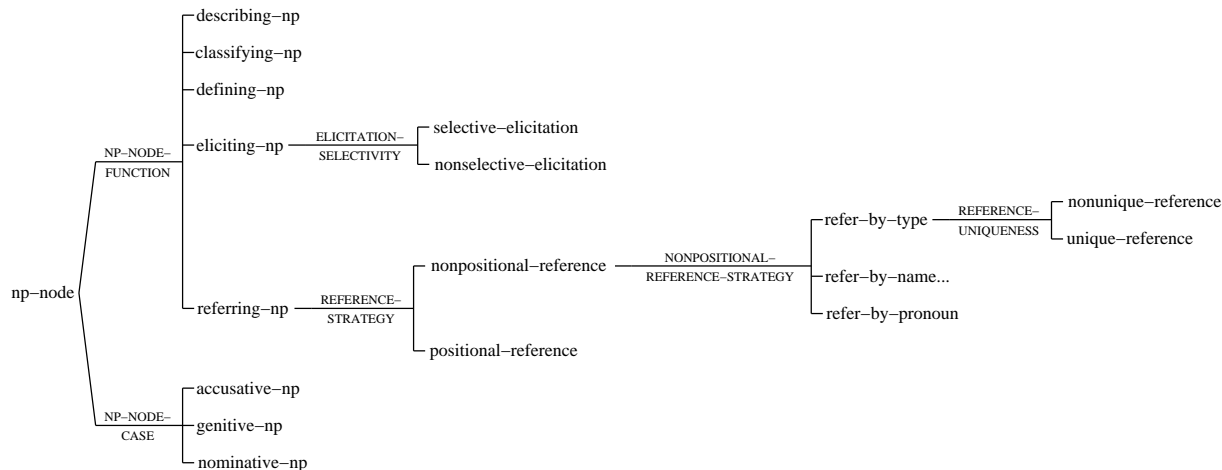


Figure 2: System Network for Nominal Function

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:

1. **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 Syntactic Placement of Non-Referring Information

In this section, we will consider where non-referring information can be placed within an NP. Before we begin this discussion, we will introduce the structure of the NP as used in ILEX.

4.1 Referring and informing components of an NP

We consider the basic structure of the NP to fall into two components: a *core*, which performs the nominal function of the NP, and an optional *periphery*, which is used to contain additional information. The peripheral element is typically realised by an indefinite NP, or a non-defining relative clause (or complexes of such), e.g., *this brooch, designed by Jessie M. King*, or *King, an Scottish designer*.

Given that we are focusing on NPs serving a referring function in this paper, we can relabel the nucleus as the *referring component* - where the referential function is satisfied; and the *informing component* - a component where only informing is done. Note however, that informing can also be performed within the performing component, as will be discussed below.

4.2 Incorporating information into the referring component

The range of slots in a systemic analysis of the NP, include, in the order they typically appear (after Halliday 1994):

(Deictic) ^ (Numerator) ^ (Epithet*) ^
(Classifier*) ^ Thing ^ (Qualifier*)

Key: () - optionality
* - any number of this slot may occur

Some of these slots support the addition of non-referring (non-assimilated) information without interfering with referring, while others are less supportive. 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, (see Dale & Reiter 1996), i.e., there is an expectation that the speaker refers using information known to the addressee. 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 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 referring component 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.¹

Another approach is to examine the semantic types of pre-modifier slot fillers, to see which, when inserted for informing reasons, seem to interfere with the referring function. We have found some of our fact-predicates are more suitable. 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

¹A study of 20 jewellery domain texts by Cheng revealed that only 1/3 of Epithets act restrictively.

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 NP-Planner has distinct procedures for each NP-function, one for classifying, one for referring, one for eliciting, etc. Due to lack of space, we will give details here only of the module for referring, and at that, only on definite reference.

The locus of the referring function is the Sem slot of the NP specification. This is the entity to be expressed. The Agenda slot is the locus through which the informing function works. The slot contains the list of facts which the aggregation module considered should be expressed in this NP. Note that the facts on the agenda need not all be new information (unasimilated) – 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.

5.1 Construction of NP

The steps of building the NP are as follows:

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

1. *Location of Assimilated Agenda:* referring will only succeed using information available to the listener, for instance, we cannot refer to *the arts and craft brooch* if the user hasn't yet been told that that brooch was in the Arts and Craft style. We thus need to select out those roles on the agenda which are assimilated. These will be preferred in building a referring expression.
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 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).
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 premodifier slots, e.g., in the Jewellery domain, Materials – *this gold and enamel brooch*, Fame – *the famous designer called Jessie M. King*.

Build Informing Component: Any information which was not consumed in the referring component can now be constructed into the informing component, e.g., [*RC: Jessie M. King*], [*IC: a Scottish designer*].

5.2 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 chooses restrictive modifiers 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.

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 does not happen, 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 the informing component 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.

6 An Example

Agenda: Assume we are talking about an apple, and have the information as in figure 3 to express. In short, the facts on the agenda are: Owner, Variety, and Position.

Referential Context: Assume also that we

Information	On Agenda	Assimilated
Class: apple	no	yes
Owner: John	yes	no
Color: Red	no	yes
Variety: Granny Smith	yes	no
Position: on table	yes	yes
Taste: good	no	no

Figure 3: An Example Information Base

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 Referring Component:

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 component.
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 referring component:
 - 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 non-restrictive relative clauses are always

in the informative component, discussed below).

In both of these slots, the inclusion of unasimulated *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 the informative component (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: Building the Informative Component The *Owner* fact can be incorporated into the Informative component 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 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. Of course, the algorithm for referring

has been modified substantially to allow agenderised facts to be incorporated into the form.

8 References

- Appelt, D E. 1985. "Planning English Referring Expressions". *Artificial Intelligence*, 26, 1-33.
- 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 23-28, 1987, pp640-647.
- Cheng, Hua, Chris Mellish & Mick O'Donnell. 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, Aix-en-Provence, 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. Pp68-75.
- 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, Chris Mellish, Jon Oberlander & Mick O'Donnell. 1996. "Sources of Flexibility in Dynamic Hypertext Generation". *Proceedings of the 8th International Workshop on Natural Language Generation*, Herstonceux Castle, UK, 13-15 June 1996.
- 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 .
- 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.