

TUNA: Towards a UNified Algorithm for the Generation of Referring Expressions

Case for Support

Part I: Previous Research and Track Record

About ITRI (General)

The **Information Technology Research Institute** at the University of Brighton, (ITRI, <http://www.itri.brighton.ac.uk/>) is a centre for research on Computational Linguistics, with 11 research staff and 7 research students working in this area. Our work covers document generation, knowledge editing, and lexicons. The ITRI is recognized to be among the leading groups internationally in the field of natural language processing, as was recently confirmed by the International Review of UK research in Computer Science (Schneider and Rodd 2001).

Language Generation at ITRI

Natural Language Generation (NLG) is one of the main foci of the ITRI, where it is studied from a theoretical angle and also in connection with such applications as document authoring, question-answering, and knowledge editing. The ITRI has hosted a number of projects that have NLG as their core:

- DRAFTER (EPSRC, 1993-97, Drafting assistant for technical writers; highly commended in DTI review of SALT initiative; nominated for a BCS innovation award)
- GNOME (EPSRC, 1997-2000, Generation of NOMinal Expressions)
- ICONOCLAST (EPSRC, 1997-2000, Integrating CONstraints On Content, Layout And STyle)
- RAGS (EPSRC, 1998-2001, Reference Architecture for Generation Systems)
- CLIME (EU, 1998-2001, Computerised Legal Information Management and Explanation)
- AGILE (EU, 1998-2001, Automatic Generation of Instructions in Languages of Eastern Europe)
- PILLS (EU, 2000-2001, Patient Information Leaflets)
- NECA (EU, 2001-2004, Net environment for Embodied Conversational Agents).

Most closely related to the present proposal is GNOME, which focused on deciding whether a full definite description ('the so and so') is called for instead of, for example, a personal pronoun (Kibble 1999, Poesio et al. 2000). Following up on GNOME, we have embarked on a strand of new work which addresses the structure of definite descriptions (e.g., Paraboni 2000, Piwek and Beun 2001, Power 1999, van Deemter and Kibble 2000, van Deemter 2000, van Deemter and Halldórsson 2001, van Deemter 2002, van Deemter and Kraemer (forthcoming)). Unlike the GNOME project, this work focuses primarily on situations where the referent is mentioned for the first time, so that (nondeictic) pronominal reference is not an option. In 1999, Van Deemter and Kibble organized a workshop 'Generation of Nominal Expressions' which contains some of the seeds for the present proposal. A book containing papers from this workshop will appear in the Autumn of 2002 (van Deemter and Kibble 2002).

Kees van Deemter (Principal Investigator; PhD University of Amsterdam 1991) joined the ITRI in 1997, where he holds a permanent position as Principal Research Fellow. He did a post-doctoral year at Stanford University (1992-3), working on logic and formal semantics of natural

language. He was a tenured research scientist at Philips Electronics in Eindhoven (1984-92 and 1993-97), working on concept-to-speech synthesis, dialogue systems, and language generation. At ITRI, he has continued his work in language and document generation. He codirected the ITRI's work on GNOME, was a consultant of the RAGS project, and has published widely about referring expressions. With Matthew Stone, Rutgers University, he will be teaching a course on the Generation of Referring Expressions at the 2002 ESSLI summer school in Trento.

Richard Power (Principal Investigator; PhD University of Edinburgh 1974) joined ITRI in 1993, where he is a (permanent) Principal Research Fellow. After his PhD, he was a research fellow at the University of Sussex (1975 - 1978). As a chief scientist of Artificial Intelligence Software in Italy he developed expert systems and coordinated an ESPRIT project (PECOS - PERSpectives on COoperative Systems) in Computer Supported Cooperative Work. At ITRI, he developed the WYSIWYM technique for knowledge editing, which was nominated for a 1998 British Computer Society Design Award. Richard has extensive experience in various aspects of Natural Language Generation. He has co-directed a number of projects, including PILLS and ICONOCLAST, which applies Constraint Logic Programming to document generation. With Donia Scott and Kees van Deemter, he has supervised numerous PhD projects.

Emiel Krahmer (Visiting Fellow; PhD Tilburg University 1995) is an assistant professor at Tilburg University. He is a computational linguist by training and for his PhD investigated formal approaches to presupposition and anaphora. An updated version of his dissertation was published by CSLI Publications in 1998. After obtaining his PhD he worked at the Eindhoven University of Technology, primarily on natural language generation for VODIS, a European research project investigating the role of language and speech technology in the car. He is co-supervisor of three PhD projects on language generation and machine learning. His interests include the generation of referring expressions, embodied conversational agents, functional aspects of prosody, and machine learning of dialogue strategies. Emiel has co-authored over 50 publications on these topics, some of which are coauthored with Kees van Deemter.

Some relevant publications by the investigators and their associates:

- Kibble. 2001. A Reformulation of Rule 2 of Centering Theory. *Computational Linguistics* **27**, 4, pp.579-587.
- Krahmer and Theune. (In press.) Efficient Context-Sensitive generation of Referring Expressions. To appear in K. van Deemter and R. Kibble (Eds.), "Information Sharing". CSLI Publications, Stanford.
- Krahmer, van Erk, and Verleg. (In press.) Graph-based Generation of Referring Expressions. To appear in *Computational Linguistics*.
- Power. 1999. Generating referring expressions with a unification grammar. Proceedings of the 9th Conference of the European Association for Computational Linguistics (EACL'99), Bergen, Norway, pp. 9-14.
- Krahmer and van Deemter. 1998. "On the Interpretation of Anaphoric Noun Phrases: Towards a Full Understanding of Partial Matches". *Journal of Semantics* **15** (2), pp. 355-392.
- Paraboni and van Deemter (in press). Generating Easy References: the case of Document Deixis. To appear in Procs. of INLG 2002.
- van Deemter. 2000. Generating Vague Descriptions. In Procs. of First Int. Conf. on Natural Language Generation (INLG-2000), Mitzpe Ramon.
- van Deemter and Kibble. 2000. "On Coreferring: Coreference in MUC and Related Annotation Schemes". *Computational Linguistics* **26**, 4, pp. 629-637 (Squib).
- van Deemter and Halldórsson. 2001. Logical Form Equivalence: the Case of Referring Expressions Generation. In Procs. of 8th European Ws. on Natural Language Generation (EWNLG-2001), Toulouse.
- van Deemter. 2002. Generating Referring Expressions: Extending the Incremental Algorithm. *Computational Linguistics* **28** (1), pp.37-52.
- van Deemter and Krahmer (forthc.) Towards a broad coverage algorithm for generating referring expressions.

Part II: Description of Proposed Research

A: Topic area and Background

The goal of the project is to design a new algorithm for the Generation of Referring Expressions (GRE) that generates appropriate descriptions in a far greater variety of situations than its predecessors. This greater *expressivity* will be achieved by allowing the generator to generate descriptions that are structurally more complex than those generated by previous algorithms. (For example, the algorithm will enable the generation of NPs like ‘the large dogs in front of the doghouse without a roof’.) We will explain why this is important, why it is difficult, and how it will be done.

The ability to verbally single out objects is an important aspect of the human communicative apparatus (e.g., Clark 1992), and it is this ability that is modeled computationally by GRE. Not surprisingly, GRE is a key task of nearly all Natural Language Generation (NLG) systems (e.g., Reiter and Dale 2000): even the simplest language generating systems (based on templates with gaps) tend to contain a GRE component, which *invents* ways of referring to objects in the domain, rather than looking them up in a list (Reiter 1995). We will assume that the choice between different syntactic types of expressions (pronouns, demonstratives, etc.) has been made by another component of the system and focus on the task of inventing appropriate descriptions, of the kind that can be realised as full noun phrases. The project will focus primarily on semantic questions involving the factual content of the descriptions. The associated PhD project will focus on matters relating to linguistic realization, asking how the factual content of a complex description is best put into words.

In the last 15 years, a number of GRE algorithms have been proposed, the most notable of which is the ‘Incremental’ algorithm (Dale and Reiter 1995, henceforth D&R), which generates referring expressions in a computationally efficient way. These early GRE algorithms were extremely limited, doing little else than conjoin atomic properties. For example, by conjoining the properties ‘dog’ and ‘brown’, an animal might be characterized as the intersection between the set of things that are brown and the set of things that are dogs (i.e., the brown dogs). Even if multimodal aspects like pointing (e.g., Kobsa et al. 1986, Landragin et al. 2001, van der Sluis 2001) are put aside, many kinds of descriptions could not be generated, because the following simplifying assumptions were made:

- a1. The target must always be a single object.
- a2. Salience is modeled as a dichotomy: some objects are salient, all others are not.
- a3. An object has (or does not have) a property regardless of context
- a4. A property never consists of a relation to another object
- a5. Referring expressions do not use negations or conjunctions

Basic GRE algorithms are often unable to generate a unique description, even though some very natural ones may exist. To remedy such limitations, a number of new algorithms were proposed, often in the shape of extensions of the Incremental Algorithm. Each of them abandons one or two of the above-mentioned assumptions, addressing such problems as the following, thereby taking a step in the direction of a **generalized** type of GRE:

- p1. Targets that are *sets* (e.g., ‘the brown dogs’, Stone 2000, van Deemter 2000)
- p2. *Degrees* of salience (Krahmer and Theune 1999; see below)
- p3. *Context-dependent* and vague properties (e.g., ‘the large dog’, van Deemter 2000)
- p4. *Relational* properties (e.g., ‘the dog in the shed’, Dale & Haddock, Horacek, Krahmer & Theune 1999)
- p5. *Boolean* properties (containing negation, conjunction) (van Deemter 2002, Gardent (in press); for examples see below)

Basic GRE algorithm

/ / / / | \ \ \ \
.. .. (Various Extensions of Basic algorithm)
\ \ \ \ | / / / /

New, unified GRE algorithm

The aim of the project is to design a generalized GRE algorithm that combines and surpasses all existing algorithms in this area. Unfortunately, it is not possible to simply merge the existing algorithms. This is true for three reasons. *Firstly*, and most obviously, the algorithms have been formalized in different ways. *Secondly* and more crucially, each of the algorithms is in need of significant finetuning, in particular because the ‘incremental’ approach to GRE tends to give bad results when applied to generalised GRE. The ‘Boolean’ algorithm presented in Van Deemter 2002, for example, will always prefer a description containing fewer disjunctions over one with more disjunctions, even if the latter is much simpler in other ways. (For example, it chooses ‘the shorthaired scruffy white (...) arthritic dogs’ instead of ‘the poodles and the chihuahuas’, if both descriptions denote the same set. See Van Deemter 2002 for discussion.) Similar observations can be made about the treatment of relations by Krahmer and Theune, whose algorithm is equally unable to take the complexity of an entire description into account. Algorithms of this kind have to be viewed as ‘proof of concept’ algorithms, demonstrating how a new class of descriptions can be generated in principle, without claiming to choose the best possible description always (or even most of the time). But, *thirdly*, the problem of choosing an appropriate description becomes even larger when different algorithms are combined. Suppose that all algorithms can be expressed in the same formalism (see **Methodology**), and that each of them has been properly finetuned. Then our situation might be compared to that of a mechanical engineer in 1870 who knows how to build wheels, brakes, combustion engines, etc., and who is now trying to build a *car*. Most components were originally designed for coaches, so they will have to be re-designed to cope with the speed engendered by the combustion engine. More dramatically, some completely new inventions are needed, such as the differential gear. Drastic innovations will also be needed to design a unified GRE algorithm.

Addressing all the different problems mentioned above in one algorithm gives rise to a wholly new set of questions, computational as well as linguistic. At the most general level, the problem may be described as an *embarrassment of riches*. GRE is all about choosing appropriate combinations of properties; the problem of choosing the best description becomes even harder now that several types of referential possibilities are taken into account. Consider, for example, what happens when relational and Boolean combinations are allowed: When trying to refer to a set of animals, the generator may now have to choose between, for example,

- ‘the dogs in the doghouse’ (relational),
- ‘the labradors *and* the poodles’ (boolean),
- ‘the labradors that are *not* in the house’ (relational and Boolean),

and so on. Owing to this new wealth of referential possibilities, the search space becomes huge, while the choice between two candidate descriptions may be difficult. In addition to the problems of search and choice, there are ‘combinatorics’ issues that arise when different referential mechanisms are combined. For example,

- *Combining Salience and Sets*. It is not clear how to generate references to sets in a context-sensitive manner (i.e., in a manner that takes the salience of objects properly into account). Krahmer and Theune, for example, modify the Incremental

Algorithm in such a way that ‘the labrador’ can be generated provided the target referent is *the most salient* labrador in the domain. But if this idea is generalized to *sets* of referents (Van Deemter 2002), ambiguities can be introduced. Suppose, for example, I’m telling you about five labradors who have taken refuge in my shed (becoming salient as a result), and more specifically about two of them that I want to keep (which makes these two even more salient). In this situation, ‘the labradors’ becomes ambiguous between the set of five and the set of two, upsetting one of the usual assumptions behind GRE: that referents are individuated uniquely.

- *Combining Saliency and Vagueness.* It is unclear how saliency interacts with other gradable properties, when these are expressed in the description. On Krahmer and Theune’s account of saliency, for example, ‘the labrador’ means ‘the most salient labrador’. But this leaves it unclear what ‘the *large* labrador’ means (where ‘large’ is gradable). It could, for example, mean ‘the largest one of the labradors that are sufficiently salient’, or ‘the most salient one of the labradors that are sufficiently large’, and so on (van Deemter 1999). As a result, it is unclear under what circumstances such descriptions can usefully be generated.

These open questions are partly, but not exclusively, empirical. **Computational complexity**, in particular, has always been a major issue in GRE (e.g., Dale and Reiter 1995, Bateman 1999) and this problem is exacerbated sharply by the new referential possibilities, and by the fact that referents can be sets (Van Deemter 2002). In the section **Methodology**, we will explain how the computational and the empirical challenges in this area will be tackled.

B: Programme and Methodology

Aims, objectives and timeliness

Aim: to achieve a substantial improvement in the ‘logical’ scope and the empirical quality of computational generators of referring expressions. The resulting ‘generalized’ GRE program will be able to produce referring expressions in a much wider variety of situations than its predecessors, and they will become more natural and varied. Our understanding of the notion of reference and the act of referring should be greatly advanced as a result.

Objective: to design and build an empirically based unified algorithm which combines and surpasses the expressive power of existing GRE algorithms. The expressions generated can be simple or complex (e.g., having Boolean and/or relational structure), and they may refer to sets as well as individuals; in addition, they will be able to incorporate context-dependent and vague properties, and pointing.

The project is **timely** because it answers questions triggered by recently developed algorithms (p1-p5 above). ITRI and Tilburg University are well placed to do this work, as a result of their previous work in GRE and other areas of NLG, and in formal/computational semantics.

Methodology

A test domain just rich enough to involve all the issues mentioned above is reference to *musicians* within an ensemble. The GRE system will, for example, be able to refer to those two people, in a given string quartet, who do *not* play the violin. Since orchestras can contain around 50 players (implying 2^{50} possible referents when plurals are included), computational feasibility is a relevant concern. To solve the problems noted in section **A**, a combination of

computational and empirical work will be pursued. To keep things manageable, we will look at GRE in isolation from other aspects of NLG (unlike, e.g., Stone and Webber 1998).

As a preliminary step towards addressing the computational issues, we took a formalism based on *labeled directed graphs*, originally designed for the problem of *basic* GRE (Krahmer et al. (in press)) and demonstrated that it can be applied to *generalized* GRE (van Deemter and Krahmer, forthcoming). The basic idea of graph-based GRE is to express both the domain and the referring expression as a labeled directed graph that contains a designated element. In this way, it becomes possible to check whether the candidate description refers uniquely to a referent, by *comparing* the two graphs. Speaking informally, this involves checking whether there exists, within the Domain Graph, any *other* subgraph to which the Description Graph is isomorphic, while still having the same designated element (i.e., the same referent). If no other such subgraphs exist then the description is a uniquely referring expression.

Consider the Domain Graph of Fig. 1, where s_2 is a designated element (and where size is modeled as if it were a non-gradable property, for simplicity). One example of a Description

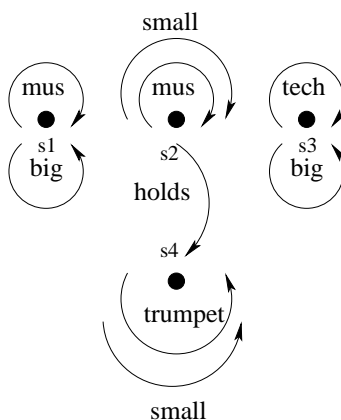


Figure 1: A scene graph involving four objects; mus = musician; tech = technician

Graph referring uniquely to s_2 is the subgraph consisting of s_2 and the edges labeled ‘mus’ (‘is a musician’) and ‘small’ (‘is small’); another is the subgraph consisting of s_2 and s_4 (with s_2 as the designated element) and the edge labeled ‘holds’. Subgraph isomorphisms can be applied to *sets* by replacing the referent in the Domain Graph by a *set* of referents. Writing (v, H) for a graph H which has v as its designated element, a description pair (v, H) refers to the scene pair (W, G) iff (v, H) is connected and, for each $w \in W$, there exists an isomorphism π from (v, H) to a subgraph (w, G') of (w, G) such that $\pi(v) = w$; a description pair (v, H) refers *uniquely* to (W, G) iff (v, H) refers to the scene pair (W, G) and there is no $w' \in W$ such that (v, H) refers to (w', G) . *Negations* of atomic properties can be accommodated if Domain Graph and Description Graph are both expanded: the property of *not holding* a given trumpet is added to the graph, for example; after expansions of this kind, the algorithm proceeds as normal. Other, similar extensions are addressed in van Deemter and Krahmer (forthcoming).

Other formalisms will be considered where appropriate, but Graph-based GRE promises to offer some important advantages:

- The framework can support a wide variety of algorithms: greedy, incremental, etc. (Krahmer et al. (in press)). Combinations with constraint satisfaction (Gardent (in press)) or related approaches are also possible (see Rudolph 2000).
- The framework allows the choice between different descriptions to be based on global properties of the description (e.g., the number of logical operators plus the number of relations used)

- Efficient implementations of key algorithms (e.g., for determining whether the relation of subgraph isomorphism holds between two graphs) are available (e.g., Eppstein 1999, Liebers 2001).

Empirical data are needed to ‘fill’ the framework with an appropriate algorithm. In the graph-based algorithm, this means choosing a *cost* function which computes the cost of an entire description graph based on, for example, the number of nodes and arcs that it contains. For this purpose, language corpora such as the British National Corpus will be used. Corpus-based methods, for example, will be used for determining reasonable upper limits on the complexity of a description (for example, measured in terms of the total number of properties expressed in one noun phrase). Such upper limits are crucial, because they allow the worst-case running time of GRE algorithms to remain polynomial as a function of all the relevant variables (e.g., the number of properties in the language and the number of elements in the domain, van Deemter 2002). Corpus-based methods are limited, however, because the semantics of the expressions in the corpora (e.g., the intended referent and the set of ‘distractors’) is often unknown. Moreover, a noun phrase that never occurs in a given situation might still have been perfectly acceptable in that situation. For these reasons, we plan to do two types of experiments, one of which investigates the behaviour of writers while the other makes use of judgements from readers.

- Experiments involving controlled elicited utterances. In these experiments, subjects will be asked to refer to a target referent. Where it is desirable to limit their options, a simple, menu based interface will be used which lets subjects choose the words of the referring expression in the style of NLMenu (Tennant 1986). This interface will be implemented as a trivial variant of WYSIWYM-style interfaces developed at the ITRI (Power and D. Scott 1998).
- Experiments that focus on ease of interpretation. In these experiments, subjects will be asked to evaluate different descriptions, saying which expression they prefer (cf. Paraboni and van Deemter (in press) for a similar experiment.) They may also be asked to determine the referents of a variety of referring expressions, in which case interpretive errors and/or reading times will reveal which expression counts as easiest to interpret.

These experiments will also shed light on the choice between the definite and the indefinite article and on the usefulness of logically redundant properties in a description. For example, we fully expect that these experiments will force us to start paying attention to the effort needed to *interpret* a referring expression and to find its referent. Experiments of the second type, which were done for a highly specialized class of referring expressions (Paraboni and van Deemter 2002), for example, suggest that readers have a preference for descriptions that are longer than those predicted by an Incremental algorithm, *if* the added properties make it easier to locate the referent.

Beyond factual content. The project will focus on semantics related tasks. This involves content selection, but also aggregation, for example, to reduce (a) to (b), which has the same factual content (being logically equivalent) but is likely to be easier to process (McCluskey 1965, van Deemter and Halldórsson 2001):

- a. The X that are A_1 and ... and the X that are A_n
- b. The X that are A_1 or ... or A_n

The appropriateness of a referring expressions is likely to depend on issues like lexical choice and linguistic realization (e.g., ‘the violin owned by Mary’ vs. ‘Mary’s violin’) as well as on semantic issues. These will be studied in an associated PhD project, which will build on existing approaches to realization (e.g., Stone and Webber 1998, Krahmer and Theune 1999, Malouf 2000), taking them into the realm of generalized GRE. This will, for example, involve the choice between ‘the old violas’, ‘all the old violas’, and ‘every old viola’.

Programme of work

The work packages of the project correspond with the main steps that we intend to take. They are presented here in chronological order (see diagram), but towards the end, order is merely indicative. Each package will deliver a new GRE system including at least a trivial realization component; each except the last one will also involve empirical work. Dissemination will take place in all WPs except the first.

One of the main difficulties confronting us lies in the *combination* of different GRE problems (e.g., reference to sets, relational properties, context dependency, salience, etc.) For this reason, we have to make sure that, when a new problem is addressed (typically in a new work package), solutions to all previously-tackled problems remain applicable.

WP (i) Preparatory technical and corpus work. Since the graph-based algorithms of Krahmer et al. (in press) have been installed at the ITRI in recent months, technical preparations will focus on the development of a tool to support controlled elicitation of referring expressions (see under **Methodology**). *Deliverables:* experimentation tool; report on complexity of referring expressions in BNC.

WP (ii) Individuals as targets. Exploration of various ‘global’ algorithms for calculating costs, also taking salience into account. The algorithm will be finetuned to make sure that relational properties are used where this makes references more natural. *Deliverable:* GRE system (documented).

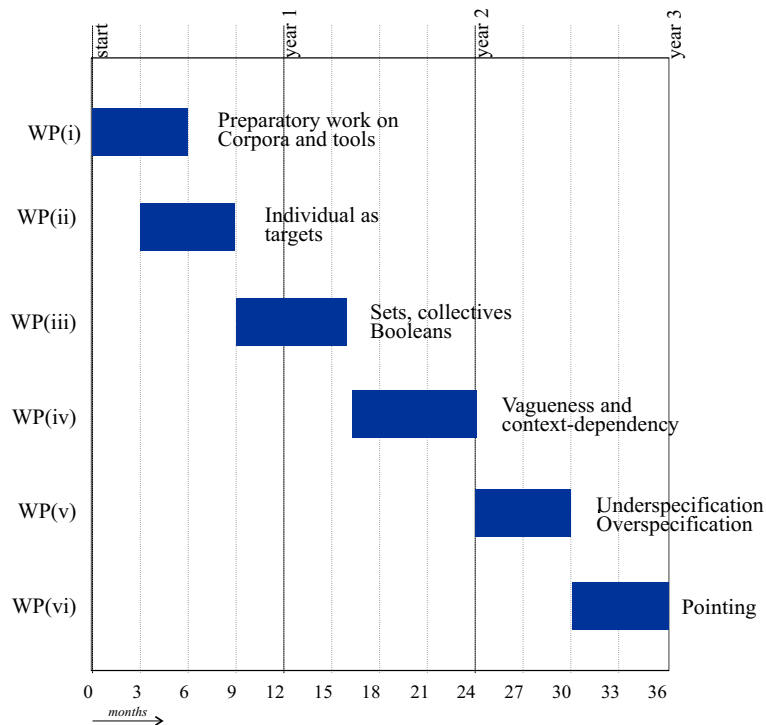
WP (iii) Sets as targets; collectives; Booleans. The next step is to generalize the procedure so as to allow reference to sets. Special care will be taken to allow a proper treatment of collective properties (e.g., Stone 1999). When this has been done, negation and disjunction will be added. One of the main challenges will be to find a proper role for the notion of salience in the new situation. *Deliverable:* GRE system (documented); conference and/or journal papers.

WP (iv) Context-dependent and vague properties. Next, we will study the effect of context-dependency and vagueness on GRE. Unlike the representation in Fig. 1, for example, this implies using different standards of size depending on the class of objects involved (e.g., ‘the small musician’ vs. ‘the small technician’ vs. ‘the small trumpet’). care will be taken to account for degrees of salience and to allow reference to sets at the same time. *Deliverable:* GRE system (documented); conference and/or journal papers.

WP (v) Under- and overspecification. Next, we will study under what circumstances it is proper for a referring expression to *under* or *overspecify* a target, taking various factors into account (e.g., Jordan and Walker 2001), including ease of interpretation (Paraboni and van Deemter 2002). *Deliverable:* GRE system (documented); conference and/or journal papers.

WP (vi) Pointing. Finally, we plan to explore how pointing affects reference. The main new challenges here are to integrate textual and spatial aspects of salience, and to account for pointing to sets, and for imprecise pointing, where the set of objects pointed to is not precisely delimited. (This WP will focus on conceptual and computational issues; for experimentation, see *Staffing*, in section E.) *Deliverable:* GRE system (documented); conference and/or journal papers.

The following diagram visualises the programme of work:



C: Relevance to Beneficiaries

1. GRE is a key component of most NLG systems, hence anyone working in NLG is a potential beneficiary.
2. Users of practical generation systems are beneficiaries. This includes commercial systems based on templates with gaps, such as report generators.
3. Since GRE models an important aspect of linguistic competence, theoretical linguists (especially semanticists and pragmaticists) will benefit from the insights gained in the project.

D: Dissemination and Exploitation

Journals: e.g., *Computational Linguistics*, *Computational Intelligence*, *Cognitive Science*, *Artificial Intelligence*, *Journal of Semantics*. **Conferences:** ACL, COLING, INLG, IJCAI. The student project will result in a **PhD thesis** at the University of Brighton

E: Justification of Resources: Staffing etc.

- Dr van Deemter (Brighton) will be managing the project, with substantial input from Dr Power (Brighton), particularly concerning implementation aspects.
- **Post Graduate RA:** The RA will carry out most of the work in the project. Affinity for formal issues in semantics, and experience in programming are essential. Because the existing implementation of graph-based GRE uses JAVA, experience using JAVA is highly desirable.
- **Ph.D. student:** The student project, which would ideally start in month 1 and cover the material of WPs i-vi in that order, will focus on linguistic realization and lexical choice. A good background in computer science or computational linguistics is required.
- Dr Krahmer (Tilburg, The Netherlands) will co-supervise the PhD student and provide assistance with the graph-based software developed by him and his colleagues

in Eindhoven/Tilburg. Extensive experimentation with multimodal aspects of GRE is also foreseen to take place in Tilburg under his supervision, and these latter activities will be coordinated with the present project (WP (vi)).

Equipment: Two desk-top computers. For compatibility with existing systems at the ITRI, state-of-the-art SUN work stations will be used.

Travel: In addition to conferences, the project will require travel between Brighton and Tilburg: approx. three times during the first year and twice yearly during the second and third year.

Literature (not by investigators)

- Bateman. 1999. Using Aggregation for Selecting Content when Generating Referring Expressions. In Procs. of ACL 1999, University of Maryland.
- Clark. 1992. *Arenas of Language Use*. University of Chicago Press & CSLI Publications, Stanford.
- Dale and Haddock. 1991. Generating Referring Expressions involving Relations. Procs. EACL, Berlin, pp. 161-166.
- Dale and Reiter. 1995. Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions. *Cognitive Science* **18**, pages 233-263.
- Eppstein. 1999. Subgraph isomorphism in planar graphs and related problems. *J. Graph Algorithms and Applications*, **3** (3):1-27.
- Gardent. (in press) Generating Minimal Definite Descriptions. To appear in Procs. of ACL-2002.
- Horacek. 1997. An Algorithm for Generating Referential Descriptions with Flexible Interfaces. In Procs. of 35th. ACL conference, 1997, Madrid.
- Jordan and Walker. 2001. Learning attribute selections for non-pronominal expressions. In Procs. of ACL-2000, Hong Kong.
- Kobsa et al. 1986. Combining Deictic Gestures and Natural Language for Referent Identification. In Procs. of Coling 1986, pp. 356-361. Bonn.
- Landragin et al. 2001. Visual Salience and Perceptual Grouping in Multimodal Interactivity. In Procs. of Int. Ws. on Information Presentation and Natural Multimodal Dialogue. Verona, Dec. 2001
- Liebers. 2001. Planarizing graphs. *J. Graph Algorithms and Applications*, **5** (1):1-74.
- Malouf. 2000. The order of pronominal adjectives in natural language generation. In Procs. of 38th ACL Annual Meeting of the Association for Computational Linguistics, Hong Kong.
- McCluskey. 1965. *Introduction to the Theory of Switching*. New York. McGraw-Hill.
- Paraboni. 2000. An algorithm for generating document-deictic references. Procs. of w/s Coherence in Generated Multimedia, assoc. with First Int. Conf. on Natural Language Generation (INLG-2000), Mitzpe Ramon.
- Piwek and Beun. 2001. Multimodal Referential Acts in a Dialogue Game: From Empirical Investigations to Algorithms. Proceedings of the International Workshop on Information Presentation and Natural Multimodal Dialogue (IPNMD-2001), Verona, Italy.
- Poesio, et al. 2000. Specifying the Parameters of Centering Theory: a Corpus-Based Evaluation using Text from Application-Oriented Domains", Proc. of the 38th ACL, Hong Kong, October.
- Power and Scott. 1998. Multilingual Authoring using Feedback Texts. In Proc. of ACL conference, Montreal.
- Reiter. 1995. NLG versus Templates, in Procs. of European workshop on NLG, Leiden.
- Reiter and Dale. 2000. *Building Natural language Generation Systems*. Cambridge University Press, Cambridge.
- Rudolph. 2000. Utilizing Constraint Satisfaction Techniques for Efficient Graph Pattern Matching. In Procs. of 6th International Workshop on Theory and Application of Graph Transformations. Springer, Berlin.
- Schneider and Rodd. 2001. International Review of UK research in Computer Science. Report commissioned by EPSRC, BCS and IEE.
- van der Sluis. 2001. An empirically motivated algorithm for the generation of multimodal referring expressions. In Procs. of ACL student session. Toulouse.
- Stone and Webber. 1998. Textual Economy through Close Coupling of Syntax and Semantics. In Procs. of INLG 1998, pages 178-187.
- Stone. 2000. On Identifying Sets. In Procs. of INLG-2000, Mitzpe Ramon.
- Tennant. 1986. The commercial application of natural language interfaces". In Procs. of Coling 1986, Bonn.