Embedding New Information into Referring Expressions

Hua Cheng

Department of Artificial Intelligence, University of Edinburgh
E17, 80 South Bridge, Edinburgh EH1 1HN, U.K.

Email: huac@dai.ed.ac.uk

Abstract

This paper focuses on generating referring expressions capable of serving multiple communicative goals. The components of a referring expression are divided into a referring part and a non-referring part. Two rules for the content determination and construction of the non-referring part are given, which are realised in an embedding algorithm. The significant aspect of our approach is that it intends to generate the non-referring part given the restrictions imposed by the referring part, whose realisation is, on the other hand, affected by the non-referring part.

1 Components of a Referring Expression

The referring expression is a very important and complex construction in languages. It can serve multiple communicative goals including referring to an object, providing new information about it, and expressing the speaker's emotional attitude towards it (Appelt, 1985). Although a formal model of referring built within the framework of a general theory of speech acts and rationality is given in (Appelt and Kronfeld, 1987), and this can be used to explain how referring acts achieve multiple goals, there is a gap between the general model and the planning of the linguistic content of a referring expression.

We divide the constituents in a referring expression¹ into two parts based on their communicative goals and the rules for their content determination and realisation. They are *a referring part*, which intends to refer to an object and *a non-referring part*, which intends to provide additional new information about the object. For example, in "**the actual writing style of Xuanzong, who was a well-known calligrapher**", the bold faced items belong to the referring part, and the underlined ones to the non-referring part.

The division is a pragmatic one and the two parts are closely related to each other. On the one hand, the referring part puts both syntactic and semantic

constraints on the presenting of the non-referring part. The syntactic constraint concerns mainly the available syntactic slots around the head. The semantic constraint will be introduced in section 3. On the other hand, the possibility of adding a non-referring part can make some realisations of a referent preferred over others. When generating referring expressions, multiple factors should be considered, which include Centering Theory (Grosz *et al.*, 1995) and stylistic preferences such as avoiding too many repetitions. If we are to satisfy all constraints to some extent, we may need to consider more than one possible realisation of a referent, choosing among those that do not significantly affect the coherence of the text. Then one of the realisations that is most suitable for adding new information can be selected.

A great amount of work has been done on generating various types of referring expressions, which addresses the referring part, while little has addressed the generation issues with respect to the other part, except that in (Scott and de Souza, 1990), the relation between embedding and rhetorical relations is discussed and several heuristics for combining sentences using embedding are given. But this is far from enough for generating an appropriate referring expression.

2 System Architecture

We design an algorithm to generate referring expressions consisting of both parts. The referring part is generated by the referring process (Dale, 1992), while the non-referring part is generated by a subtype of the aggregation process called embedding, which selects suitable facts and realises them as components within the structure of a referring expression. The algorithm fits into the text planner of ILEX (Oberlander *et al.*, 1998).

ILEX is an adaptive hypertext system generating museum object descriptions. In ILEX, pieces of domain knowledge that may be worth expressing in a text are represented as nodes and links in a graph called the Content Potential. Two kinds of nodes

¹Only singular referring expressions that are primarily for referring to physical objects are considered here.

useful for referring expression generation are entity nodes and fact nodes². A fact is represented as *Predicate(Arg1,Arg2)*. A revised version of Text Structure (TS) (Meteer, 1992) is used as an intermediate level of representation between the text planner and the sentence realiser, which provides syntactic constraints to the text planner while abstracting away from linguistic details. The Text Structure uses a unified representation for structures both above and below sentence level, so that abstract sentence planning can be done in text planning.

The text generation process follows roughly four steps: 1) The text planner selects a set of facts to be expressed and the best rhetorical relations between them³. 2) The text planner builds the TS for each fact in the set. For each entity in a chosen fact, the referring process produces a list of possible realisations that will unambiguously refer (the referring part). Based on the constraints imposed by the referring part, the embedding process finds from the set all the unexpressed facts whose *Arg1*s are that entity⁴, and makes embedding decisions including what to embed, what syntactic form the embedded parts should take and which realisation for the entity is preferred, according to the principles in the next section. This step iterates until the TS for all facts is built. 3) The aggregation process goes through the TS for parataxis possibilities. 4) The appropriately simplified TS is sent to the surface realiser, where the natural language text is generated.

We distinguish between two types of parataxis: semantic and textual. Semantic parataxis concerns facts that have two identical semantic constituents or a rhetorical relation between them, while textual parataxis deals with any adjacent facts from text planning, with no rhetorical connection between. In step 3), both types of parataxis are performed.

3 Generating the Non-Referring Part

A referring expression is primarily for referring to an entity. So the addition of a non-referring part should not interfere with this primary function. We summarise two principles that the non-referring part must obey, which have been realised in our embedding algorithm in a simple way.

²Each entity node corresponds to a domain object; each fact node represents a relation between two entities and can be expressed as a single sentence in language.

³Details of the text planning algorithm can be found in (Oberlander *et al.*, 1998).

⁴The chosen fact actually forms the nucleus of Elaboration, and the facts collected by embedding form the satellites.

1. **The non-referring part should not confuse the reader about the referent indicated by the referring part.** That is, if the referring part can uniquely identify the referent, the reader should not be confused over which object the referring expression is about because of the addition of the non-referring part. For example, in the description of a currently focal object which is a necklace, we might say "*The necklace is made from gold*". Suppose we also want to inform the readers that the necklace has floral motifs. We should use "*The necklace, which has floral motifs, is made from gold*" rather than "*The necklace with floral motifs is made from gold*" because the latter may make the readers think that the sentence is about a necklace which is not the focal object.

Based on both the properties of English and our analysis of real museum descriptions, we find that additional information is provided by evaluative adjectives, non-restrictive clauses, and almost all grammatical constituents in an indefinite and a demonstrative noun phrase. These characteristics are captured by embedding rules. For example, the definition of one rule that embeds a prepositional phrase is:

```
(def-embed-rule
 :name with-phrase ;the name of this rule
 :priority 4
 :type prep-phrase ;the type of embedding
 :constraints
 ((:type pred Generalized-Possession)
 (:type refer (:or demonstrative indefinite)))
 :RT ((:rel-parent Adjunct)
 (:textual-sem With-Prep-phrase)))
```

In the definition, *priority* is the order in which the rule should be tried, where those rules producing simpler syntactic forms always have higher priority (Scott and de Souza, 1990); *constraints* is the restrictions that must be satisfied by the predicate and arguments of the embedded fact and the realisation of the referring part. In the above example, the required semantic category of the predicate is specified, which is used to select suitable facts for embedding; *RT* is the resource tree for building the TS for the embedded component.

Assume we have two facts *F1=style(J1,Organic)* and *F2=hasqual(J1,Floral-motif)*. Without using embedding, we might generate "*The necklace is in the Organic style. It has floral motifs*". Suppose *F1* and *F2* are selected by the text planner and the embedding process respectively, and the referring form of the entity *J1* can be demonstrative, definite or pronoun. Applying the above embedding rule,

we would realise *F2* as a post-modifier of the *Arg1* of *F1*, and choose demonstrative, as “*This necklace with floral motifs is in the Organic style*”.

2. The non-referring part should not reduce the readability of the text. There are several restrictions concerning readability:

1) Complexity of a referring expression: the generated expressions should not be too complex to read. We use a fixed number of syntactic slots to restrict the maximum amount of information that can be expressed. But the actual complexity is decided by user models. At present we only distinguish between adults and children. According to observations in psycholinguistic research, embedded clauses in subjects are a major obstacle to comprehensibility (Coleman, 1962). So for children, the system generates fewer non-restrictive clauses than for adults and none at all in subjects.

2) Compatibility with other aggregation possibilities: only semantic paratactic and hypotactic relations between facts are considered here. Complex embedded components like non-restrictive clauses may interrupt the semantic connection between a set of sentences. For example, if we do not consider such connections while making embedding decisions, we would generate a sentence like: “*This jewel is made of gold, sapphire, a kind of precious stone and enamel, which is often used to produce a shiny surface*”. It is not good compared with: “*This jewel is made of gold, sapphire and enamel. Sapphire is a kind of precious stone, and enamel is often used to produce a shiny surface*”.

Adjectives would not have such negative effect in most cases, especially when the paratactic parts have syntactically symmetrical modifications, like “*The bracelet has a slightly flared band and a swelling midsection*.” Prepositional phrases fall between adjectives and relative clauses in their effect.

Also when one fact is to be embedded, it is necessary to check if there are facts semantically related to it, which should be embedded together. For instance, it is bad to say “*The necklace, which is made from gold, is in the Organic style. It is also made from enamel*”.

So before embedding a fact, our embedding algorithm considers the possibilities of other types of aggregation, and only embeds if the embedded properties can be realised as a syntactic form other than a non-restrictive clause in possible paratactic nuclei, and all of the semantically related facts can be embedded at the same time. This means that em-

bedding has a lower priority than parataxis and hypotaxis, which reflects the relationship between the weakest rhetorical relation, Elaboration, and other types of rhetorical relations.

4 Future Work

This paper discusses our ongoing work on how to embed new information into a referring expression. While the restrictions concerning the second principle are currently implemented in a procedural way, it is possible to formalise them as constraints within the embedding rules.

An interesting problem is the relation between embedding and entity-based coherence, which exists between spans of text in virtue of shared entities (Oberlander *et al.*, 1998). When a fact is embedded into another one, the entity inside it may become unavailable for an entity-based move, and the smooth transfer from this fact to its elaborating facts is cut off. The effect of embedding on local and global coherence is to be exploited more in future work, and a comprehensive evaluation is indispensable.

Acknowledgement This research is supported by a University of Edinburgh Studentship. The author appreciates the comments from Dr. Chris Mellish, Dr. Mick O'Donnell and the four anonymous reviewers.

References

- Appelt, D. 1985. Planning English Referring Expression. *Artificial Intelligence*, 26:1-33.
- Appelt, D and Kronfeld, A. 1987. A Computational Model of Referring. In *Proceedings of the Tenth IJCAI*, 640-647.
- Coleman, E. 1962. Improving Comprehensibility by Shortening Sentences. *Journal of Applied Psychology*, 46:131-134.
- Dale, R. 1992. *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes*. MIT Press.
- Grosz, B, *et al.* 1995. Centering: A Framework for Modelling the Local Coherence of Discourse. *Computational Linguistics*, 21:203-226.
- Meteer, M. 1992. *Expressibility and The Problem of Efficient Text Planning*. Pinter Publishers Ltd.
- Oberlander, J. *et al.* in press. Information Structure and Non-canonical Syntax in Descriptive Texts. *Text Representation: Linguistic and Psycholinguistic Aspects*. Benjamins Publisher.
- Scott, D. and de Souza, C. 1990. Getting the Message Across in RST-based Text Generation. *Current Research in NLG*, 47-73.