

Integrating Gricean and Attentional Constraints

Rebecca J. Passonneau
Bellcore
445 South Street
Mornstown NJ 07960
beck@bellcore.com

Abstract

This paper concerns how to generate and understand discourse anaphoric noun phrases. I present the results of an analysis of all discourse anaphoric noun phrases (N=1,233) in a corpus of ten narrative monologues, where the choice between a definite pronoun or phrasal NP conforms largely to Gricean constraints on informativeness. I discuss Dale & Reiter's [To appear] recent model and show how it can be augmented for understanding as well as generating the range of data presented here. I argue that integrating centering [Grosz *et al.*, 1983] [Kameyama, 1985] with this model can be applied uniformly to discourse anaphoric pronouns and phrasal NPs. I conclude with a hypothesis for addressing the interaction between local and global discourse processing.

1 Introduction

This paper concerns how to generate and understand discourse anaphoric noun phrases, or noun phrases (NPs) that evoke a discourse entity already in the discourse model (Webber [1978]). Dale [1989] [1992] implements Gricean constraints on informativeness for generating discourse anaphoric NPs. However, his model follows the tradition of assuming that distinct constraints govern pronouns versus phrasal NPs (cf. [Reichman 1985] [Grosz and Sidner, 1986]). Centering [1983] [1985], a model of local attentional state [1979], has been applied primarily to definite pronouns. I argue that Gricean constraints should be applied equally to discourse anaphoric pronouns and phrasal NPs, and that integrating centering and informational constraints covers a broader range of cases. In §2, I present an analysis of all discourse anaphoric NPs (N=1,233) in a corpus of ten narratives showing that semantic explicitness depends largely on informational constraints. Discourse anaphoric NPs almost never provide new information and are rarely more informative than necessary. In §3, I show how Dale & Reiter's [To appear] generation model can be augmented to apply uniformly to pronouns and phrasal NPs for both generation and understanding. While centering has been

"The work reported here was not supported by Bellcore

used to account for informationally under-specified pronouns, I argue that centering also accounts for certain over-specified phrasal NPs. In §4, I integrate centering with the augmented Gricean model and discuss the extended coverage. Results in §2 include a one-way correlation of overly informative discourse anaphoric NPs with shifts in global discourse structure. In the conclusion, I discuss directions for extending the integrated model in ways that might indirectly account for this correlation.

2 Analysis of a Coded Corpus

In this section I present the results of an analysis of all discourse anaphoric NPs in a corpus of spoken narratives directed at the question of how informative NPs are, relative to their contexts of occurrence. The first subsection describes the corpus and coding features. The next subsection presents results showing that discourse anaphoric NPs in the corpus, whether pronominal or phrasal, are rarely more informative than necessary, and if so, tend to occur at shifts in global discourse structure.

Fig. 1 identifies four possibilities regarding the semantic informativeness of an NP relative to its context. Three of them pertain to the following Gricean principles, referred to by Dale [1989] as informational adequacy and efficiency: the speaker should be sufficiently informative to unambiguously identify the intended referent (adequacy), and the speaker should be no more informative than necessary (efficiency). The boxed pronouns in (2a) of Fig. 1 are both adequate and efficient (well-specified): it is clear what the pronouns refer to; less informative forms (zero pronouns) would be ungrammatical. The phrasal NPs in (2b) are adequate but not efficient (over-specified). The pronominal NP in (2c) is inadequate (under-specified, efficiency does not apply to inadequate NPs): "it" could refer either to the ladder or the tree. A fourth possibility is that an NP may perform two functions, to identify the referent and to add information about it, as in (2d) (over-determined). In Fig. 1, the feature +/- increasing distinguishes between over-determined and over-specified NPs.

2.1 Data Coding

The corpus consists of ten narrations from Chafe's *Pearl Stories* [1980]. Chafe recorded and transcribed subjects who had been asked to view the same movie and describe

1	A man ₁ saw a ladder ₂ leaning against a pear tree ₃	
2a	Later he ₁ moved it ₂ to a different tree	+adequate +efficient well specified
2b	The man ₁ moved the ladder ₃ to a different tree	+adequate -efficient -increasing over specified
2c	It (?) was tall	-adequate under specified
2d	The contented pear picker ₃ was done for the day	+adequate +increasing over determined

Figure 1 Relative Informativeness

it to a second person. The movie contained seven sequential episodes about a man picking pears. It had a vivid sound track, but no language. As part of a long term study of the relationship between linguistic features and discourse structure [Passonneau and Litman, 1993] [Litman and Passonneau 1995b] [Litman and Passonneau, 1995a], discourse anaphoric NPs in the corpus had already been coded for coreference relations and location. Location of an NP is represented here in terms of the containing sentential utterance and discourse segment as described below. Fig 2 illustrates an excerpt

S	U _i	
6	28	And you think Wow this little boy is probably going to come and see the pears
	29	[ps] and [ps] he is going to take a pear or two and then go on his way
7	30	[ps] Um but the little boy comes
	31	[ps] and uh [I O] he, doesn't want just a pear
	32	he wants a whole basket
8	33	[ps] So he puts the [ps] bicycle down
	34	and he [ps] you wonder how he is going to take it with this

Figure 2 Narrative Excerpt Illustrating Informativeness

Cbafé [1980] identified three types of prosodic phrases from graphic displays of intonation contours. A period indicates a phrase terminated by a pitch fall, a question mark indicates final level or rising pitch, and a comma indicates phrase final—not sentence final—intonation. The transcriptions here show all repeated and incomplete words and phrases, non-lexical articulations such as "uh, urn, tsk", and vowel lengthening as indicated by - . Pause locations are shown as '[ps]'.

Sentential utterances are defined to be a non-overlapping sequence of units that completely covers the discourse. Briefly, a new sentential utterance begins with a functionally independent clause (FIC) if it is immediately adjacent to the preceding FIC. Otherwise it begins at the onset of the prosodic phrase where the next FIC begins. An FTC is a tensed clause that is not a verb argument, a restrictive relative clause, or one of a set of formulaic "interjection" clauses (e.g., "You know" with no clausal argument, for full details cf [Passonneau, 1994]). Material between clauses includes sentence or word fragments, and non-lexical articulations (e.g., "um"). Locations and sequence numbers of the seven sentential utterances in Fig 2 are shown in column 2.

The global context is structured into sequential segments, multi-utterance units whose utterances are presumed to be more related to one another semantically

and pragmatically than to other utterances. The segments numbered 6-8 (col 1 of Fig 2) were derived from an empirical study described in [Passonneau and Litman, 1993]. Each narrative was segmented by 7 new, untrained subjects. Subjects were instructed to place segment boundaries in transcripts whenever the narrator had finished one communicative task and begun a new one. They were restricted to placing boundaries between prosodic phrases. To focus their attention on the criterion, subjects were also instructed to label segments with a brief description of the speaker's intention.

The size and number of segments per subject per narrative varied widely, from a rate of 5.5% to 41.3% (Avg = 16%) with segment widths ranging from 1 to 49 phrases (Avg = 5.9). Despite this variation, the number of times 4 to 7 subjects assigned boundaries in the same place was extremely significant (using Cochran's Q [1950], cf [Passonneau and Litman, 1993]). We took agreement among at least 4 subjects as the threshold for empirically validated boundaries.

A surface constituent is considered to be a discourse anaphoric NP if it occurs in free variation with syntactically prototypical NPs and corefers with a preceding NP (cf [Passonneau, 1994]). One type of empty category is also included, namely zero pronoun subjects of FICs conjoined by ", ' and" etc. In Fig 2 the sequence of coreferential NPs used to refer to the little boy are coindexed. Segments 7 and 8 in Fig 2 both begin with an utterance containing an NP referring to the boy. At the onset of segment 7, a phrasal NP is used to refer to him (U₃₀) whereas at the onset of segment 8 (U₃₃) a definite pronoun is used. But a pronoun could have replaced the phrasal NP in U₃₀ with no loss of information. So the phrasal NP is over-specified but not over-determined, the attributes "boy" and "little" were already mentioned in U₂₈. The pronoun subject in U₃₃ is locally well-specified because the boy is the only animate entity mentioned in U₃₂; it is globally well-specified because the boy is the only entity in the discourse with a bicycle.

2.2 Analysis of Informational Constraints

The goal of the analysis is to determine whether relative informativeness of NPs correlates with global discourse structure (cf [Reichman, 1985] [Grosz and Sidner 1986]). Any phrasal NP that is discourse anaphoric is potentially over-specified, whereas a definite pronoun will only be over-specified if a zero pronoun could have been used. I first sorted the discourse anaphoric NPs in the corpus (N=1,233) into the three categories of phrasal NPs (PhrNPs, N=563), explicit pronouns (PROs: definite, indefinite, demonstrative, N=544), and zero pronominals (ZPs, N=126). Then I identified all pairs of coindexed NPs where NP₂ was more explicit than NP₁. This procedure identified 128 discourse anaphoric NPs in the corpus that were potentially over-specified or over-determined. The sole over-determined NP, illustrated in Fig 3, occurs relatively late in the narrative (U₈₅), it seems mainly to provide contrast (cf "that old man" vs "those little boys").

Potentially over-specified NPs were sorted into four mutually exclusive categories—well-specified, segment

U_r
 84 [ps] You just know that those little boys are going to go back
 [ps] to where the pear tree was
 85 and you just know that old man is going to see [ps] these little
 boys coming and say "Ha you're the ones who stole the pears"

Figure 3 Over-determined NP

onset attentional shift, and reiterative. A potentially over-specified NP is well-specified if a less explicit form would have been ambiguous or unclear. The containing utterance is included in the context since the proposition expressed in an utterance can disambiguate a referring expression. A potentially over-specified NP that is not well specified but which occurs, in the first utterance of a new segment, is classified as a segment onset. The segments in the coded Pear corpus arguably contain intra-segmental shifts of attention associated with changes in temporal aspect, or shifts in discourse reference time (for definitions assumed here of [karneyama *et al* 1993]). The third category, attentional shift consists of these cases. A fourth catch-all category includes e.g., repetitions, repairs, contrastive NPs and unexplained cases.

Table 1 indicates that most potentially over-specified NPs (N=127) were either well-specified (46%) or occur at an empirically verified segment onset (16%) or a hypothesized attentional shift (23%). Of the 69 NPs whose nearest antecedent was in a distinct segment, 29% occurred at a segment onset. Over a third (36%) of the NPs whose antecedent was in the same segment and 12% of those whose antecedent was in a distinct segment occurred at an intra-segmental attentional shift. In sum in the coded Pear corpus, NPs that re- evoke existing entities seem to be rarely over-specified (68/1233 or 5.5%) or over-determined (1/1233). Of the 68 over-specified cases (columns 2-5), 20 (30%) correlate with segment onsets independently identified by naive subjects, and 29 (42%) appear to correlate with intra-segmental attentional shifts. Thus, an over-specified NP is more likely than not to correlate with an attentional shift (72%). Note however, that the reverse implication does not hold, that is it is not the case that a segment shift is likely to be signalled by an over-specified NP.

2.3 Focused Attribute Sets

To account for the choice of modifiers in phrasal discourse anaphoric NPs, it is necessary to determine how attributes are selected from the information known about a discourse entity. According to Grice's [1975] maxim of quality, speakers should be relevant. With respect to discourse anaphoric NPs in the Pear stones

Antecedent Segment	Well Specified	Segment Onset	Atten Shift	Other	Total
Same	22		21	15	58
%	38%		36%	26%	100%
Prev	37	20	10	4	69
%	53%	29%	14%	6%	100%
Totals	59	20	29	19	127
%	46%	16%	23%	15%	100%

Table 1 Potentially Over-Specified NPs

NP modifiers are derived from what I refer to as focussed attribute sets, independent of whether the NP is over-specified. Focussed attribute sets comprise the following three categories of relevance. First, an attribute set can be in focus because it was mentioned in the most recent phrasal NP. For example, in Fig. 2, the boy is referred to in U₃₀ as 'the little boy,' repeating attributes mentioned in the last phrasal NP referring to the boy (in U₂₈).

Second, the focussed attribute set may specify the most recently mentioned location of an entity. The subject NP in U₁₇ of Fig. 5 (§3.2) refers to one man as 'the man up in the tree' to distinguish him from the second man who came by with a goat. The free is the last mutually known location of the former. Finally, an attribute set can be in focus because it pertains to a key narrative event that the entity has been an agent of. Intuitively, an event is more central to a narrative the more difficult it is to describe the narrative without mentioning that event. Operationally, key events occur more frequently than others both within and across narratives. For example, the main adult character is often described as 'the pear picker' or as 'the man who was picking pears' (see U₁₀₈ of Fig. 6 §4), and so on. The other main character is often described as 'the thief' 'the boy who stole the pears' 'the boy with the pears,' and so on.

How to order the focussed attribute sets for a given discourse entity is a topic for further investigation. Here I simply assume that the three types of attribute sets mentioned above—where applicable—are in focus. I also assume that the focussed attribute sets of an entity (FAV_e) are updated as the discourse progresses.

3 Modelling Informativeness of NPs

The data reported above indicates that in the Pear corpus definite pronouns and phrasal NPs are rarely over-specified or over-determined. In this section, I describe a processing model to account for this observation. In the next section I discuss how centering can be integrated with this model to account for under-specified pronouns, and certain over-specified phrasal NPs. First I briefly review Dale's [1989] [1992] model including his more recent work with Reiter [to appear]. Then I modify this model to apply to understanding as well as generation to include the current utterance in the context of evaluation to apply informational constraints uniformly to pronouns and phrasal NPs and to select modifiers on the basis of focused attribute-value pairs.

3.1 Distinguishing Descriptions

Dale [1989] generates anaphoric pronouns and phrasal NPs by distinct means. In EPICURE [1989], a system for generating recipes, a definite pronoun is always generated to refer to the discourse center which is analogous to the backward-looking center of [Grosz *et al* 1983] [Kameyama, 1965] but is domain specific. It is the entity that results from the next recipe operation. For example, rice₁ will be the center after an utterance of *Stir the rice₁*.

Dale [1989] requires phrasal NPs to be distinguishing descriptions. As in Webber [1978], Dale assumes that the discourse mode represents the discourse entities that

have already been evoked and the attribute-value pairs describing them. For any set of entities U , Dale [1989] defines a distinguishing description of an entity e in U to be a set of attribute-value pairs that are true of e and of no other members of U . This enforces adequacy. He defines a minimal distinguishing description to be one where the cardinality of the attribute-value pairs cannot be reduced. This addresses efficiency.¹

Dale [1989] defines the discriminatory power (\mathcal{F}) of an attribute-value pair $\langle A, V \rangle$ that is true of a discourse entity e in a universe of entities U in terms of the cardinality N of U , and the total number n of entities in U that $\langle A, V \rangle$ is true of:

$$\mathcal{F}(\langle A, V \rangle, U) = \frac{N-n}{N-1}$$

\mathcal{F} ranges in value from 0 to 1. If $\langle A, V \rangle$ is true of only one of the entities in the set U , then $\mathcal{F}_{\langle A, V \rangle}$ is 1, and $\langle A, V \rangle$ is a distinguishing description of the entity.

Dale's [1989] algorithm for constructing a distinguishing description of e in U given a set \mathcal{P} of attribute-value pairs that are true of e briefly works as follows. First compute \mathcal{F} for each member of \mathcal{P} . If all values of \mathcal{F} are 0, no unique description can be constructed. Otherwise, select the attribute-value pair with the highest value to add to the description, and reset U to be only those entities in the initial U that the selected attribute value pair is true of. Repeat this process terminating when an attribute-value pair with a discriminatory power of 1 has been selected. The selected attribute-value pairs constitute the input description for a surface NP.

In recent work, Dale & Reiter [To appear] enforce a range of Gricean constraints using an algorithm based on human behavior that is simpler and faster than their previous algorithms [Dale 1989] [Reiter, 1990]. It performs less length-oriented optimization thus balancing brevity against lexical preference. The output NPs are not guaranteed to be maximally short because humans occasionally use unnecessary modifiers. The 5.5% rate of over-specified discourse anaphoric NPs in the Pear data also supports the relaxation of brevity, but is partly conditioned by attentional factors (cf §§4-5).

3.2 C_describe

In this section I illustrate the role of *c_describe* in processing definite pronouns and phrasal NPs. *C_describe* is a 4-place relation among a discourse entity E , a surface NP, the current utterance context λU , and the discourse context C that requires $\lambda NP \lambda U$ to be a distinguishing description of E relative to C . For generation, NP is solved for given an instantiation of the remaining three arguments, whereas E is solved for during understanding (assuming Prolog's control structure).

A definite pronoun that is a distinguishing description is also a minimal distinguishing description because its length is 1. In generation, *C_describe* attempts first to find a definite pronoun to satisfy the uninstantiated NP argument, succeeding if the pronoun is a distinguishing

¹ Cf. Reiter [1990] for a discussion of problems in generating maximally efficient NPs using Dale's framework and Dale & Reiter [To appear] for an argument that maximal efficiency is psychologically implausible.

S_i	U_j	
2	4	[ps] A nd he (e_1) [ps] fills his- thing with pears
	5	and ZERO (e_1) comes down
	6	and there s a basket he (e_1) puts them in
3	7	[ps] A nd you see- [ps] passerbys on
		bicycles and stuff go by
	8	[ps] A nd [ps] then a boy (e_2) comes by
		[ps] of a bicycle
	9	the man (e_1) is in the tree
	10	[ps] and the boy (e_2) gets off the bicycle
	11	and ZERO (e_2) looks at the man (e_1)

Figure 4 Excerpt from Narrative 9

description. For generating the pronoun "he" in U_4 of Fig. 4, the arguments of *c_describe* are

- (1) *c_describe*(e_1 , NP,
 $\lambda \lambda$ fills' (λ his thing" "pears"), FS₁)

The utterance context is assumed to be a feature structure co-indexed with any relevant discourse entities other than the uninstantiated variable E .² By using the utterance as part of the input in solving for NP, given information that appears anywhere in the current utterance can filter entities from the discourse context following Dale's [1989] algorithm. New information about an entity in the utterance is not mutually known, and has no discriminatory power [Dale, 1989].

For present purposes, the last argument of *c_describe* is first instantiated to the most recent focus space, and in turn to other focus spaces until a solution is found. Dale [1989] takes the universe of discourse to be partitioned into focus spaces (cf [Grosz and Sidner 1986]) with the most recent focus space being the most accessible and making no assumptions regarding relative accessibility of earlier focus spaces. Similar assumptions are made here. I assume that segment boundaries in the Pear corpus correspond to focus spaces, and that some focus spaces may be composed of others. I assume the existence of an inference mechanism that constrains how focus spaces are signalled during generation, and how focus spaces are inferred during understanding. In recent work, for example, Litman and I report on algorithmic methods for identifying segment boundaries in the Pear corpus using features of prosody, cue words and referential NPs [Litman and Passonneau, 1995a]. Given such a mechanism, a new focus space would be added to the discourse model after a segment onset has been processed.

In (1), FS₁ appears as the initial context argument of *c_describe*. The only animate entity in FS₁ is e_1 , previously described as a man picking pears in a pear tree who looks like a farmer, is plump, has a mustache and is wearing a white apron (utterances 1-3, not shown here). The feature structures corresponding to all but one of the definite pronouns "he, she, it" or "they" will be rejected as a description of e_1 because e_1 is neither plural, non-animate or female.³ The pronoun "he", represented

² For simplicity, the utterance context represents certain semantic arguments as quoted strings.

³ For simplicity, I am ignoring the difference between grammatical gender and sex.

S _i	U _i	
	13	and it is just a monotonous kind of thing for him (e ₁)
4	14	[pa] And a man (e ₂) comes along with a goat
	15	[pe] and the goat obviously is interested in the pears
	16	But the man (e ₂) just walks by with the goat
5	17	And the man (e₁) up in the tree doesn't even notice

Figure 5 Phrasal NP to Avoid Ambiguity

as the attribute-value pairs (<type human> <gender male>, <cardinality 1>), not only describes e₁ it is also a minimal distinguishing description

An analogous process applies to understanding the same pronoun in U₄, with the entity variable E uninstantiated NP instantiated to *he* and the utterance and discourse context instantiated as above. Given a distinguishing description there is guaranteed to be exactly one solution to E. However, the search problem increases with the size of the context. Partitioning the search space into focus spaces controls the search through the discourse model to some degree. (Integrating centering with *c-describe* as described below guides the search even further.) For present purpose, *c-describe* returns E instantiated to c₁ after searching through the entities in FS_i. The remaining NPs exemplified here are understood in a similar fashion.

Given a context where there is no definite pronoun solution to NP *c-describe* will attempt to construct a phrasal NP, preferably with no modifiers. In FIG A a new, singular male human entity is added to the context at U₉—a boy who comes by on a bicycle (e₂). Subsequent references to the boy or the man must discriminate between them: the utterance context for the subject NP of U₉—A \ m' \, the tree)—does not identify e_j because U_B—*he comes down*—leads to the inference that the man is no longer in the tree. However e₁ is a male adult and e₂ is a male child, a distinction encoded by the common nouns *man* versus *boy*. Since *man* is what Dale & Reiter [To appear] refer to as a basic attribute, *man* will be selected as the head noun. The determiner will be definite because the entity is already in the context (but cf [Passonneau 1904]). The resulting NP *the man* is a minimal distinguishing description because no pronoun is a distinguishing description.

Fig 5 illustrates a context when a phrasal NP without modifiers could not both have a head noun that specifies a basic attribute, and be a distinguishing description. It also illustrates the problematic nature of relations among distinct focus spaces. In generating the subject NP in U₁₇, the last argument of *C-describe* is first instantiated to FS₄. The pears referred to in U₁₅ of segment 4 are some pears that e₁ picked, so in order to interpret U₁₇, e_j must be brought into focus. This side-effect of resolving the reference to the pears could be implemented by adding d to FS₄ or by resetting the current focus space to a more encompassing focus structure that includes FS₃ and FS₄. I believe further empirical work is needed to resolve such issues. In any case, I assume that the context for generating U₁₇ includes both e₁ and e₂. Because these two entities are the same type, a distinguishing description of e₁ must contain discriminatory

modifiers. Features for generating the modifiers are selected from FAV_{e₁}, which here contains only two sets of salient attributes. Since e_j's location is the most recently evoked it is used in generating the NP.

Above I noted that centering can add structure to the search space for understanding discourse anaphoric NPs. Fig 4 illustrates another reason to integrate centering with *c-describe*. In U₁₀ of Fig 4, the subject NP ("The boy") is not a pronoun even though the utterance context is a distinguishing description of e₂. The boy (e₂) is mutually known to have been on a bicycle at the time of the event mentioned in utterance UR. Temporal processing (cf [Kameyama *et al*, 1993]) would lead to the inference that the boy is still on the bicycle after U_q. Thus a definite pronoun is presumably well-specified and the model presented so far would generate *he*. However, a pronoun would produce a garden path effect in this context. If it would be interpreted as referring to the man until *bicycle* has been processed.

4 Centering and Informativeness

The *edesebnt* relation has three limitations that centering can compensate for. First *c-describe* constrains the semantic content of a discourse anaphoric NP, but not its grammatical role. Second as noted below centering predicts that a pronoun can be under-specified. Third, in explanation is needed for the over-specified NP *the boy* in U₁₀ of Fig 2. In this section I indicate how centering is interleaved with *c-describe*. Centering is a more local process so it applies first.

4.1 Centering

Centering is a model of local focus of attention that constrains the use of definite pronouns [Grosz *et al*, 1983] [Kameyama 1985]. One of the discourse entities [Webber, 1978] evoked by an NP in an utterance U_i may be the backward-looking center (CD) [Grosz *et al* 1983] of U_i, the current local focus of attention. Alternatively the CB of U_i, (CB_i) might not be explicitly mentioned (realized) in the utterance [Grosz *et al* 1983]. The discourse entities mentioned in U_i comprise the forward looking centers (CFs) ordered by increasing obliqueness of grammatical role [Kameyama, 1985] [Passonneau, 1989] to represent the likelihood that they will be mentioned in the subsequent utterance. The centering principle [Grosz *et al*, 1983] predicts that if CB_{L1} and CB_(f,) are the same entity, then the NP evoking CB_i will be a third person, definite pronoun.

- (2) a Carmellaj went to the bookstore
b Afterwards, shej gave Rachel* a new book
c Shej's a true bibliophile

Example (2) illustrates that where the s* mantics of the utterance and commonsense reasoning do not discriminate among possible referents for an ambiguous pronoun there is an independent effect of local attentional constraints. Centering predicts that the preferred interpretation of the pronoun in (2c) is Carmella. But in this context, neither the pronoun alone nor the utterance is a distinguishing description of anyone, so the pronoun is under-specified.

- (3) a Carmellaj went to the bookstore
 b On her way home she_j saw Rachel_t
 c She* looked pale

kameyama [1985] used examples like (3) to illustrate how commonsense reasoning and lexical semantics can override the default centering predictions for pronoun interpretation. Centering would predict that 'Carmella is the backward-looking centre of (3b), and that the default interpretation of the pronoun in (3c) would thus be Carmella. Instead, (3c) is interpreted as a continuation of the description of the perceptual event in (3b). Kameyama [1986] posits property sharing of features of adjacent utterances as a constraint on CB where the shared property can be subject (or non-subject) grammatical role (cf [Passonneu '18*]) as in (2), or what she refers to as empathy as in (3). Note that because Rachel is already known to be the object of the perceptual event in (3b), the utterance context in (3c) is a distinguishing description of 'Hdch'!

4.2 Integrated Model

Fig 6 shows all of one segment and part of another one where the subject pronouns of all the utterances are coreferential. On the one hand, the CB of this segment initial utterance U₁₀₅ is the same as the CB of U₁₀₇, conflicting with the idea expressed in [Grosz et al., 1986] (that centering transitions reflect global discourse coherence (cf [Passonneu '18*] 'fo appear]). On the other hand, integrating centering and *c-describe* can account for both NPs in UIOB and support inferences consistent with a global focus shift.

Earlier in the narrative excerpted in Fig 6 three boys helped the pear thief after he had fallen off of his bicycle and were rewarded with three pears. Segment 21 describes their adventures after the pear thief leaves. In generating utterance U₁₀₅ the input to the generator will be a representation of an event in which the boys eat their pears. The set of three boys is designated as the new CB. Because CB_(f₁₀₅) is the same as CB_(f₁₀₇), it should be realized as a pronoun [Grosz et al., 1983] and by property sharing [kameyama 1985] [Passonneu, 1989] it should be realized as the subject of the current utterance. Centering and Gnce constraints coincide here in that the definite pronoun "they" is also a minimal distinguishing description.

To generate the phrasal NP object in U₁₀₈ the process is analogous to that discussed above for generating "the man up in the tree" in Fig 5. The context argument of *c-describe* is first set to Cf_(f₁₀₇). Since neither Cf_(f₁₀₇) nor the most accessible focus space (FS₂₁) contains a representation of e_T the context argument will be reset until

§	U _j	
21	105	[ps] So they (e ₁) are walking along
	106	and they (e ₁) brush off their pears (e ₃)
	107	and they (e ₁) start eating it (e ₃)
22	108	Then they (e ₁) walk by [ps]
		the man who was picking the pears (e ₂)

Figure 6 Excerpt from Narrative 1

e[^] is in a focus space on the focus stack. Focused attribute sets are then used to generate the relative clause.

For understanding the subject NP in UIDB, recall that centering applies before *c-describe*. The subject pronoun will be assumed to realize the CB of the utterance, and will be assigned the default interpretation of e_i. Application of *c-describe* leads to the recognition that "they" is also a distinguishing description of e_i relative to CF_(f_{inT}). In this fashion, centering prunes the search space to the single entity satisfying the informational constraints imposed by *c-describe*. In understanding the object NP the context argument must be instantiated to a more inclusive focus space, since neither the previous utterance nor the previous segment contains any entities described by this NP.

The integrated model also accounts for the problematic phrasal NP in Fig 4, discussed above. We saw that for U₀ and U₁₀, repeated below, the phrasal subject of U₄ was well-specified but the phrasal subject of U₁₀ was over-specified, and a pronoun would be generated instead. But as noted above, a pronoun subject would have a garden path effect.

U₄ the man (e_j) is in the tree (e₃),

U₁₀ and the boy (e_j) gets off the bicycle (e[^])

kameyama's version of centering [1986] differs from [Grosz et al. 1983] in allowing an utterance to have a null CB. I'm would have a null CB because there is no definite pronoun constrained by property sharing that corefers with an NP in the previous utterance, in fact no NPs in U₁₀ refer to entities mentioned in U₄. A definite pronoun subject in U₁₀ would be assumed to be CB_(f₁₀) and would be inferred to refer to e_i. This accounts for the garden path effect. Consequently, a pronoun must be blocked. Because no entity in U_n is referred to in U₁₀ the input for generating U₁₀ will be annotated as having a NULL CB. This imposes output constraints requiring the subject and object NPs to be other than definite pronouns. As a consequence, *c-describe* will not try to find a pronoun solution to the uninstantiated NP argument. In the first phrasal NP solution the head would denote a basic category and the NP would have no modifiers, thus generating the existing phrase "the boy". In sum, centering relaxes the constraint otherwise imposed by *c-describe* that an NP cannot be over-specified.

5 Conclusion

I have presented an analysis of discourse anaphoric phrasal NPs in a corpus of narrative monologues showing that pronouns and phrasal NPs are rarely over-specified. Future research should indicate to what degree this generalization applies to other genres and modalities. Centering predicts conditions under which an under-specified pronoun can be used, but says little about the interpretation of phrasal NPs. I have outlined a processing model that integrates the attentional constraints of centering with aspects of Gnce's maxims of quantity and quality. For enforcing the maxim of quantity, I rely on Dale's algorithm for constructing distinguishing descriptions [1989] [1992], which I apply uniformly to pronouns and phrasal NPs for both generation and understanding.

For enforcing the maxim of quality I combine aspects of Dale & Reiter's [To appear] preferred attributes with the construct of focussed attribute sets derived from the corpus analysis. In contrast to Dale & Reiter [To appear], distinguishing descriptions are evaluated using the current utterance context as a filter and by instantiating the discourse context successively to the Cf list of the preceding utterance, then the current focus space, then other focus spaces, until a solution is found.

Centering provides one mechanism for relaxing the requirement that an NP (either pronominal or phrasal) should be a distinguishing description. Another mechanism would be needed to relax informational constraints at shifts in focus structure so as to account for the one-way implication of over-specified NPs with global shifts of attention (Table 1). However further investigation is needed to determine how to integrate local and global discourse processing. When neither the Cf list nor the current focus space is the appropriate context for understanding or generating a discourse anaphoric NP I have assumed that either an earlier focus space or a more inclusive one must be accessed. Some of the examples presented here suggest that the contextual dependencies captured by the use of focused attributes might constrain the relation of each new utterance to the global discourse model. For example the segment onset in Fig. 6 (U_{108}) contains two NPs, one of which is the same as the CB of the preceding utterance. Maintaining the same CB relates U_{108} and its focus space (FS_{22}) to the most recent focus space FS_{21} . But the object NP expresses attributes last mentioned in segment 17 thus relating U_{108} to the earlier focus space FS_{17} . If the global structure is a tree the relation of U_{108} to both segments 21 and 17 might indicate how high up in the tree to locate the new focus space. Alternatively, an investigation of such relations might provide evidence about the nature of global structure, such as whether it is a tree or a lattice.

Acknowledgements

This work was partly supported by NSF grant IRI 91 13064. Thanks to Robert Dale, Ehud Reiter and Megumi Kameyama for valuable comments on ideas presented here.

References

- [Chafe 1980] W. L. Chafe. *The Pear Stories*. Ablex Publishing Corporation, Norwood, NJ, 1980.
- [Cochran 1950] W. G. Cochran. The comparison of percentages in matched samples. *Biometrika* 37: 256-266, 1950.
- [Dale and Reiter, To appear] R. Dale and E. Reiter. Computational interpretations of the Gricean maxims in the generation of referring expressions. To appear in *Cognitive Science*.
- [Dale, 1989] R. Dale. Cooking up references. In *Proceedings of the 27th ACL*, 1989.
- [Dale 1992] R. Dale. *Generating Referring Expressions*. MIT Press, Cambridge, MA, 1992.
- [Grice, 1975] H. P. Grice. Logic and conversation. In *Syntax and Semantics, Vol. 3, Speech Acts*, pages 41-58. Academic Press, New York, 1975.
- [Grosz and Sidner 1986] B. J. Grosz and C. L. Sidner. Attention intentions and the structure of discourse. *Computational Linguistics* 12: 175-204, 1986.
- [Grosz et al 1983] B. J. Grosz, A. K. Joshi and S. Weinstein. Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st ACL*, 1983.
- [Grosz et al 1986] B. J. Grosz, A. K. Joshi and S. Weinstein. Towards a computational theory of discourse interpretation. 1986. Ms.
- [Kameyama et al 1993] M. Kameyama, R. Passonneau and M. Poesio. Temporal centering. In *Proceedings of the 31st ACL*, 1993.
- [Kameyama 1985] M. Kameyama. *Zero Anaphora: The Case of Japanese*. PhD thesis, Stanford University, 1985.
- [Kameyama 1986] M. Kameyama. A property-sharing constraint in centering. In *Proceedings of the 24th ACL*, New York, 1986.
- [Litman and Passonneau 1995a] D. Litman and R. Passonneau. Combining multiple knowledge sources for discourse segmentation. In *Proceedings of the 33rd ACL*, 1995.
- [Litman and Passonneau 1995b] D. Litman and R. Passonneau. Developing algorithms for discourse segmentation. In *44th Spring Symposium Series on Empirical Methods in Discourse Interpretation and Generation*, 1995.
- [Passonneau and Litman 1993] R. J. Passonneau and D. Litman. Intention based segmentation: Reliability and correlation with linguistic cues. In *Proceedings of the 31st ACL*, 1993.
- [Passonneau and Litman To appear] R. J. Passonneau and D. Litman. Empirical analysis of three dimensions of spoken discourse: Segmentation, coherence and linguistic devices. In E. Hovy and D. Scott, editors, *Burning Issues in Discourse*. Springer-Verlag, Berlin, To appear.
- [Passonneau 1989] R. J. Passonneau. Getting at discourse referents. In *Proceedings of the 27th ACL*, 1989.
- [Passonneau 1994] R. J. Passonneau. Protocol for coding discourse referential noun phrases and their antecedents. Technical report, Columbia University, 1994.
- [Passonneau 1994] R. J. Passonneau. Frame shifts: Perlocutionary meaning and discourse reference. In *Conference on Focus and Natural Language Processing*, Kassel, Germany, 1994.
- [Passonneau To appear] R. J. Passonneau. Interaction of the segmental structure of discourse with explicitness of discourse anaphora. In F. Prince, A. Joshi and M. Walker, editors, *Proceedings of the Workshop on Centering Theory in Naturally Occurring Discourse*. Oxford University Press, To appear.
- [Reichman 1985] R. Reichman. *Getting Computers to Talk Like You and Me*. MIT Press, Cambridge, Massachusetts, 1985.
- [Reiter 1990] E. Reiter. *Generating appropriate natural language object descriptions*. PhD thesis, Harvard University, 1990.
- [Sidner 1979] C. L. Sidner. Towards a computational theory of definite anaphora comprehension in English discourse. Technical report, MIT AI Lab, 1979.
- [Webber 1978] B. Webber. A formal approach to discourse anaphora. Technical Report 3761, Bolt Beranek and Newman Inc., 1978.