Strategies for Specifying Relations*

Henk Zeevat

ILLC, University of Amsterdam

henk.zeevat@uva.nl

Abstract

The generation perspective taken by optimality theoretic syntax has a lot in common with the insights gained in natural language generation. This paper explores how insights about NP generation (e.g. [Reiter and Dale(2000)], [van Deemter(2004)]) can be made fruitful for explaining the semantics and pragmatics of sentences with more than one plural NP by exploiting optimality theoretic pragmatics, as well as the list construction in discourse grammar. Though this is an exploration only (anaphora is completely neglected and no attempt is made at covering the lexicon), I would claim to show here:
1. the naturalness of the cumulative readings
2. how the different quantificational schemata arise
3. how exhaustivity implicatures arise
4. how differential implicatures arise
5. how to disambiguate double plural sentences

1 How to specify relations

A normal natural language generation task is to specify a relation, given as a set of sequences of of objects, as computed e.g. by a relational database[1]. One of the subtasks here is to construct singular referential expressions. It is customary in NL generation to have preferences for the kind of NP employed. This is an example statement of the preferences, meant to be correct for Dutch and extended from [Reiter and Dale(2000)].

(1)   first and second person pronouns > reflexives > 3rd person
         pronoun > deictic pronouns > anaphoric definites and short
         names > full demonstrative NPs > full descriptions and full
         names > indefinites > generics

---

*In partial fulfillment of a promise I made to Kjell-Johan to specify all my prejudices about topics.

[1]It is certainly no accident that the pioneering PHLIQA project in the 1970's, an NL interface with a relational database, was the context for the work reported in [Scha(1981)] and [Scha(1983)]

The interpretation of the hierarchy is that if the triggering condition for one kind of NP is met, it should be employed in preference over the NPs in the lower parts of the hierarchy. The triggering conditions are the conditions under which the NP can be used and the whole scheme can be captured by saying that a certain type of NP must be used if it can be used unless there is a more preferred type that can also be used. First and second person pronouns can be used if the intended referent is the speaker or the hearer, reflexives can be used if they have a c-commanding antecedent in the same clause, 3rd person pronouns if the referent is given and highly activated. Deictical NPs are made possible by being present in the visual field, short definites and short names by the referent being activated (or for short definites employing a relational noun, if the bear the expressed realtion to a highly activated referent). Long definites are allowed by the possibility of a definition using common ground knowledge and indefinites are the last resort if everything else fails.

Singular quantifiers are missing from the scheme (and do not belong there since they are not referring even in the extended sense given to that concept in discourse semantics). But they will be discussed later on in this paper.

The scheme can be explained in optimality theory (a fuller treatment is given in [Zeevat(2000)]) by assuming a set of expressive constraints that force the expression of certain features, like reflexive, person, identifiable etc. The constraints can be left unordered since it seems that apart from 1st/2nd person not implying reflexive, the higher triggering conditions all entail the lower ones[2] Such constraints also give rise to implicatures associated with the choice of elements from the lower part of the hierarchy: the hearer is given to understand that e.g. the choice of "a woman" implies that various features are missing like reflexive, first person, identifiability by a property or by a function, contextual salience or visibility in the perceptual field of the conversation partners. Notice that these features are definable in terms of the common ground between the speaker and the hearer.

Grice notes that A in saying (2)

(2)   I saw Smith in town with a woman.

implicates that the woman is not Smith's wife, but there is a far larger class of women she is implicated not to belong to (basically no woman in the common ground of speaker and hearer or functionally related to those) and the explanation for these implicatures —including Grice's implicature— are the expressive constraints the speaker is supposed to follow.

The hierarchy does not substantially change for plural reference.

(3)   first and second person pronouns > reflexives > reciprocals >
         3rd person pronouns > deictic pronouns > anaphoric definites
         and short names > full demonstrative NPs > full descriptions
         and full names > cardinal and estimating indefinites > bare
         plurals and covering plural definites

---

[2]This suggests that Panini's elsewhere principle also explains what is going on here. The most specific rule for referring needs to be followed and this is e.g. using the third person pronoun if the objects is activated as well as discourse old and identifiable.

These hierarchies can be used for specifying a relation R. For simplicity, I am here assuming that the relation is given by its extension, i.e. as a set of n-ary sequences of objects, with the objects coming from domains A_i with $i \leq n$. The projections $\pi_i(R)$ are defined as $\pi_i(R) = \{a_i :< a_1, \ldots, a_i, \ldots, a_n >\in R\}$. I also assume that there is a given NL sentential scheme $\zeta\alpha_1, \ldots \alpha_n$ that expresses that $x_1, \ldots, x_n$ stand in the relation R iff $\alpha_1 \ldots \alpha_n$ are replaced by names for $x_1, \ldots, x_n$ respectively.

Under these assumptions, one can define a simple and natural strategy for specifying the given relation. Using the hierarchy, a referring expression is selected for each projection π_i and filled in for α_i in $\zeta\alpha_1, \ldots \alpha_n$. This will be the *default strategy* in the rest of the paper. It applies both to the singular and to the plural.

For singular relations (where all projections have cardinality 1) and quite a substantial class of plural relations this is a successful strategy. E.g. for a single plural projection in an otherwise singular relation, this gives an optimal specification, i.e. one where the hearer is maximally informed given the possibilities provided by the common ground between speaker and hearer. In case definites can be used everywhere and can in fact be used by the hearer to determine the referent, the hearer can reconstruct the input relation.

A single plural projection does no harm to this property and the cases where the projection is a single collection (rather than the set of the elements collected in it) also do not lead to specifications that are less informative than is possible.

But if there are two properly plural projections, information may be lost under the simple strategy: the information about how the members of the projections are related. This does not mean that the strategy is not used. It still is and then gives rise to intended so-called cumulative readings, but their use seems to be either a final resort: the common ground does not allow full specification or a full specification is not the goal of the conversation. For the other cases, there are alternative strategies: the *distributive strategy* and the *list strategy*.

The simple strategy operates properly for homogeneous relations:

A relation $R \subseteq A_1 \times \ldots \times A_n$ is *homogeneous* iff $R = B_1 \times \ldots \times B_n$ with for $i < n$ $B_i \subseteq A_i$.

Purely singular sentences specify homogeneous relations and all monadic relations are homogeneous. Any polyadic relation with a single plural projection is still homogeneous. It is clear that homogeneous relations can be effectively specified by generating the singular or plural NP appropriate for each of their projections.

For non-homogeneous relations there is the distributive and the list strategy. In the distributive strategy, the relation is split into a set of relations $R_a = \{< a, a2, \ldots, a_n >: a \in \pi_1(R)\}$ and it works only if all the other projections of these relations are uniform in one of the senses described below. If this condition is not met, the list strategy is the only one that will lead to a proper specification.

The best case is *identity uniformity*.

(4) $\forall a, b \in \pi_1 \ \pi_i(R_a) = \pi_i(R_b)$
       ex. Every boy kissed the two girls.

The next best case is (definable) functional uniformity.

(5)   there is a functional relation F with a name α such that
        $\forall a \in \pi_1 \pi_i(R_a) = F(a)$
        ex. Every boy kissed his nieces.

The next best uniformity is kind and number/estimation uniformity.
there is a kind K and a number/estimate N with names α and θ such that $\forall a \in \pi_1 \; \forall i \leq n \; \pi_i(R_a)$ has N members of kind K.
ex. Every boy kissed three aunts.
Every boy kissed many girls.
The weakest form of uniformity is kind uniformity:

(6)   there is a kind or set K with name α such that $\forall a \in \pi_1 \; \forall i \leq n$
        all members of $\pi_i(R_a)$ are of kind K.
        ex. Every boy kissed girls.

In essence, in finding the NP for the non-distributing projections $\pi_i(R)$, the hierarchy of NPs still applies. One is looking for the highest NP in the hierarchy such that the NP meets the triggering condition for each object $\pi_i(R_a)$ for all a in the distributing projection. The nature of the uniformity determines how far this can go: identity uniformity gives no restriction, functional uniformity can be marked by 3rd person pronouns bound by a, definite and possessive markers, kind and number uniformity leads to NPs like "many chicken", "some books" and "three cakes", while kind uniformity is exclusively expressed by bare plurals and covering definites.
The distributing projection itself may follow the default strategy, though explicit quantifiers ("every boy") are also possible.
The distributive strategy presupposes uniformities in the relation: each projection needs to be uniform at the very least for kind. The hierarchy of NPs gives also here the correct implicatures: if a bare plural is used, one can infer that the uniformity in that projection does not go up to identity, functional relation or count/estimate.
The list strategy finally cuts up the first projection into a partition and divides the relation over that partition. The cells of the partition determine subrelations which should be presented in turn. All strategies are in principle possible for presenting the subrelations, but it is sensible to go for efficiency here, by selecting the partition in such a way that the default strategy or the distributional strategy applies. The following list illustrates the list strategy for a given relation relation $LIKE \subseteq GIRLS \times DANCES$ on which the first projection is split into two singletons and the partition results into two homogeneous relations which are specified by the names of the girls and the dances.

(7)   Clara likes waltzing.
        Maaike likes belly dancing.

The list strategy gives rise to the discourse relation "list" which allows recursive uses. It comes with a special implicature: that the partition is a partition, i.e. that the different cells are exclusive (the kind of differential implicatures that also arise with the discourse relation "contrast") and that their union gives the whole projection (a special exhaustivity implicature going with the list discourse relation). Unlike

4

the default and the distributive strategies, the list strategy always leads to the most informative specification and -where definites can be used all around- allows the full reconstruction of the relation by the hearer.

2 A Theory of Interpretation: OT Pragmatics

In other work ([Zeevat(2001)]) I defended the following theory of interpretation given by four OT constraints. It is called a pragmatics and not a semantics because while it is a theory of how to determine the change that a given utterance makes to the common ground (conversational implicature, context change potential or update semantics -in this case also allowing "downdates"-) and thereby also a natural account of the truth conditions of sentences, it is not determined by the syntactic structure of the utterances alone but also by contextual factors and pragmatic principles.
The theory consists of a system of four defeasible constraints given here:

(8)   **FAITH > CONSISTENT > *NEW > RELEVANCE**

FAITH is the principle that an interpretation is good if one could have expressed that interpretation in the same way oneself, if one were the speaker. Utterances that are imperfect do not remain uninterpreted: one goes for the interpretation for which there are fewest violations of this principle. If one would have a full OT syntax and phonology, the principle would be finding an interpretation for which the utterance is as optimal as possible i.e. the interpretation should be such that there is no alternative interpretation for which the utterance is more optimal given the OT syntax and phonology. Requiring full optimality is too strong: **FAITH** should not give up with pronunciation or syntactic mistakes. It should also not give up for utterances that make use of non-conventional means.
CONSISTENT discards interpretations which are inconsistent with the context or are implausible given the context, provided of course there are others available that are more plausible and consistent. If the language can mark inconsistency with the context (e.g. by adversative marking) or implausibility given the context (e.g. by mirative marking), **FAITH** as the stronger constraint would switch this process off, by instead preferring inconsistent or implausible interpretations.
*NEW is the general preference for the old, expected, familiar and activated and the prohibition to add anything to the interpretation without good cause. Good cause is exclusively given by the needs of **FAITH**, **CONSISTENT** or ***NEW** itself. It prefers presupposition resolution over presupposition accommodation, partial presupposition resolution over full accommodation, highly activated antecedents for pronouns over less activated ones, old referents over new referents if both are possible.
RELEVANCE is the principle that any question that is activated in the context and that seems to be addressed by the utterance is in fact answered by the utterance. It prefers interpretations where the questions are answered over those where they are not. There is an interaction here with the strategic considerations of section 1. Wherever decisions are conditioned on relevance, the **RELEVANCE** principle guarantees that the hearer will faithfully reconstruct those decisions.

3   The meaning of plurals

The following two examples of double plurals are taken from [Scha(1981)] as well as the readings attributed to them. I will show how these readings follow from the possible strategies and pragmatics.

(9)   a. The lines cross the circles.
          b. 200 firms bought 300 computers.

They give rise to many readings. The first example identifies contextually given sets LINE of lines and CIRCLE of circles. Assuming those, the readings are as in (10).

(10)   a. $cross(LINE, CIRCLE)$
          b. $\forall x \in LINE \; \forall y \in CIRCLE \; cross(x, y)$
          c. $\forall x \in LINE \; \exists y \in CIRCLE \; cross(x, y)$
          d. $\forall x \in LINE \; \exists y \in CIRCLE \; cross(x, y) \; \wedge \; \forall y \in CIRCLE \; \exists x \in LINE \; cross(x, y)$
          e. $\exists x \in LINE \; \exists y \in CIRCLE \; cross(x, y)$
          f. $\forall y \in CIRCLE \; \exists x \in LINE \; cross(x, y)$

Here (10a.) is the (implausible) collective reading (compare: The boys lifted the pianos).
The second example has 18 readings based on the following ambiguities:

(11)   collective or not per coordinate
          precise/at least per cardinal
          scope
          cumulative or not

The point of the pragmatic system is that it gives combinations of a reading with a pragmatic profile which can be matched to the discourse situation as perceived by the interpreter.
If the speaker is taken to specify the crossing relation (restricted to the lines and the circles) or the buying relation (restricted to firms and computers) by means of the default strategy, in addition to the cumulative reading (= the double existential reading if we assume both definites in (10a.) are covering definites) one gets an implicature that the relation is homogeneous (the condition under which the default strategy is a strategy for specification). This then gives the double universal reading based on covering definites (13) to (10a.) (and in principle also to (10b.)).

(12) $\exists X \subseteq LINE \; \exists Y \subseteq CIRCLE \; \forall x \in X \; \forall y \in Y \; cross(x, y)$

If the business of the speaker is specification, this also gives exhaustivity implicatures: no more firms bought computers, no more computers were bought by firms. This gives precisely 200 and precisely 300 as a meaning of the cardinals. The double universal reading does not make sense for the relation of buying restricted to firms and computers (firms buy their own computers), so the assumption of the default

strategy for specifying the relation in (10b.) leads to the cumulative reading with the exhaustivity implicatures, where the question which firm bought what computer is left open. For (10a.) specification also brings the implicature that all and not only some of the lines and circles are involved. Otherwise, the hearer would not be able to reconstruct the relation. So (11b.) is the result. The existential readings (11c.) (11e.) and (11e.) result from not assuming full specification.
The full result for (9b.) is (13).

(13) $\exists X \subseteq FIRM \; \exists Y \subseteq COMPUTER \; (\#X = 200 \wedge$
$\#Y = 300 \wedge \; \forall x \in X \; \exists y \in Y \; buy(x,y) \wedge \; \forall y \in Y \; \exists x \in$
$X \; buy(x,y) \wedge$
$\forall z \in FIRM \; \forall v \in COMPUTER \; (buy(z,v) \rightarrow z \in X \wedge \; v \in$
$Y))$

But with a different predicate homogeneity is quite plausible, as in (14)

(14) 3 boys saw 5 girls.
$\exists X \subseteq BOY \; \exists Y \; \subseteq GIRL \; (\#X = 3 \wedge \#Y = 5 \wedge$
$\forall x \in X \; \forall y \in Y \; see(x,y) \wedge$
$\forall z \in BOY \; \forall v \in GIRL \; (see(z,v) \rightarrow z \in X \; \wedge v \in Y))$

The other plausible reading is the distributive one with exhaustivity implicatures.

(15) $\exists X \subseteq FIRM \; (\#X = 200 \wedge \forall x \in X \; \exists Y \subseteq COMPUTER$
$(\#Y = 300 \wedge \; \forall y \in Y \; buy(x,y) \wedge \; \forall z \in COMPUTER$
$(buy(x,z) \rightarrow z \in Y) \wedge$
$\forall v \in FIRM \; \exists w \in COMPUTER \; (buy(v,w) \rightarrow v \in X)))$

The other readings can arise by assuming that full specification is not sought by the speaker. E.g. in the last case one may be interested in how many firms bought a substantial number of computers with substantial being pragmatically defined as minimally 300. This removes both exhaustivity implicatures. Or (10b.) can be part of a list where groups of firms are listed according to the number of computers they bought. The exhaustivity of 300 is then maintained but the exhaustivity of 200 is now with respect to *buying 300 computers* and not with respect to *buying computers*[3].
Notice that there is a considerable distance here from compositional semantics. The "logical forms" that were given describe the input relation under the assumption that the speaker is trying to do something specific with it (specify it, count it, etc.), given her knowledge of the input relation and the common ground with the hearer. The semantic/pragmatic contraints on the relations derive from the strategy that the speaker is apparently following in order to achieve her aim. It is in principle an accident, if there exists a compositional recipee for deriving the logical form, though it may be argued that the existence of such a recipee could be a factor in making communication easier.

---

[3]The exactly 200 reading would still be what one gets. In the case described by (15) one gets an even stricter reading for the "quantifier" 200.

The strongest argument for the picture of interpretation sketched in this and the previous section is the optionality of the grammatical marking involved. Distributivity can be forced by determiners like "every", "each" and "no" or by floating markers like "each", but it is not necessary to do so. Likewise collective interpretations can be marked by "together" and "all" rules out covering interpretations of definites and bare plurals (even of cardinal NPs). But often it is not necessary to use a marker. This points to a system of defaults that one can mark against. The optionality of the markers does not make proper sense under a compositional approach[4] since it seems counterproductive to leave the interpreter in the dark with respect to which of the many readings applies and it is equally hard to see why the markers would be recruited from other material if the interpreter is anyway free to insert covert operators as she sees fit during interpretation.

The system of defaults is given to some extent by the pragmatic system. **RELEVANCE** forces the construction of goals and questions the current sentence must contribute to, given the goals and questions already in the common ground and so is responsible for mostly assuming that the speaker is trying to specify a fully known relation in an effective way, thus giving homogeneous interpretations when the default strategy is already assumed. (Constraints ***DISTRIBUTIVE** and ***SPLIT** make the default strategy in effect the default strategy), but also weaker interpretations when the relation is not fully known or when full specification is not the best way to contribute tot he conversation. **CONSISTENT** is important to rule out implausible homogeneous readings (as in our second example) or implausible collective interpretations (as in our first example). ***NEW** does not seem to play any role.

How did a system of this kind get in place? The fossils seem to be still around, in the form of bare nouns and floating quantifiers. Starting from the assumption that it all starts with nouns and verbs, initially noun+verb and noun+verb+noun were highly ambiguous, with the nouns being everything: kind names, pronouns, existential and universal quantifiers, definites and indefinites. Recruitment of demonstratives in an adnominal position makes it possible to disambiguate definites. Adverbs can take on the role of distribution/collection/existential markers. Finally these can coalesce with nouns into NP determiners. A separate development of cardinals and estimators from adjectives makes it possible to disambiguate towards indefiniteness.

The new recruitments have pushed some of the original uses of bare nouns away towards the more unclear uses: where there is no identification of the referent or group of referents, no counting is possible and there is no distribution.

4   Exhaustivity Implicatures

The strategies I discuss in the first section also give a theory of scalar and other implicatures. Anybody who has been thinking about problems in natural language generation knows the effect: choices need to be made in generation, the fact that a certain choice is made indicates that the speaker assumes something.

---

[4]Discourse Representation Theory ([Kamp and Reyle(1993)]) seems to have no advantages in this respect.

Our strategy counts where counting is possible and classifies only where that is impossible or irrelevant. If counting takes place, it indicates that (functional) identification is impossible. Splitting the most prominent projection indicates that the most prominent elements of the split form a partition of the most prominent projection and that the remaining relations are uniform. These are exhaustivity implicatures. Counting gives the special exhaustivity implicatures that used to be called scalar implicatures. They are computed as a side effect of the hearer checking **FAITH** in her joint context with the speaker. For this she needs to put herself in the hearer's shoes. As a good Gricean she needs to discover the speaker's intentions on the basis of the utterance. Many other Gricean implicatures can be reduced to similar strategies. E.g. the (natural) strategy of telling a story in the order of the events is responsible for moving up reference times. The strategic obligation in complex assertions of marking whether a subordinate clause holds or does not hold according to the speaker (an obligation that cannot always be carried out as in: My husband believes I am cheating on him but he does not know it) forces the choice of "when" or "because" instead of "if" when the complement is true and a choice of the irrealis when it is false and forms the basis of clausal implicatures.

5   Relations are Topics are Questions

Many ([van Kuppevelt(1995)], [Umbach(2001)], [Krifka(1992)], [Zeevat(1994)]) have assumed that topics are questions. If topics are wh-questions $?x_1, \ldots, x_n\varphi$ then they are relations or at least closely related to relations. The view that they are relations is Scha's [Scha(1983)], but the views of [Hamblin(1973)] or [Karttunen(1977)] or [Groenendijk and Stokhof(1984)] are not far removed. The point here is that for a satisfactory treatment of the relation between questions and reduced answers, one needs to get hold of the relation somehow and all accounts should allow for that.

If one takes a conversational turn in order to settle a topic, one needs to specify a relation. This is what is done all the time, but often the specification problems are rather trivial. Sometimes the relation is on a high level, e.g. in giving the arguments against a theory or in listing the possible causes of John not being there. Good Griceans standing in the speaker's shoes see the task the speaker is trying to carry out: telling them what happened last night at the party, explaining why so little progress has been made with the paper, giving the list of the shopping, explaining where one should go to when visiting Düsseldorf. If they grasp the speaker's intention, they grasp which relation she is trying to specify and how she is trying to do it. I am here playing out [Grice(1957)] against the later Grice who believed in literal meaning as a basis for computing the conversational contribution. **FAITH** is all but it produces all possible conversational contributions and not a core from which the contribution has to be computed.

References

[Grice(1957)] Grice, H. (1957). Meaning. *The Philosophical Review*, **66**(3), 377–388.

[Groenendijk and Stokhof(1984)] Groenendijk, J. and Stokhof, M. (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.

[Hamblin(1973)] Hamblin, C. (1973). Questions in montague english. *Studies in Language*, pages 41–53.

[Kamp and Reyle(1993)] Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.

[Karttunen(1977)] Karttunen, L. (1977). Syntax and semantics of questions. *Linguistics and Philosophy*, **1**, 3–44.

[Krifka(1992)] Krifka, M. (1992). Presupposition and focus in dynamic interpretation. *Journal of Semantics*, **10**.

[Reiter and Dale(2000)] Reiter, E. and Dale, R. (2000). *Building Natural-Language Generation Systems*. Cambridge University Press, Cambridge.

[Scha(1981)] Scha, R. (1981). Distributive, collective and cumulative quantification. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*, pages 483–512. AMC, Amsterdam.

[Scha(1983)] Scha, R. (1983). *Foundations for Question Answering*. Ph.D. thesis, University of Groningen.

[Umbach(2001)] Umbach, C. (2001). Contrast and contrastive topic. In *Proceedings of the ESSLLI 2001 Workshop on Information Structure, Discourse Structure and Discourse Semantics*. Helsinki.

[van Deemter(2004)] van Deemter, K. (2004). Towards a probabilistic version of bidirectional ot syntax and semantics. *Journal of Semantics*, **21**(3), 407–421.

[van Kuppevelt(1995)] van Kuppevelt, J. (1995). Discourse structure, topicality and questioning. *Journal of linguistics*, **31**, 109–147.

[Zeevat(1994)] Zeevat, H. (1994). Applying an exhaustivity operator in update semantics. In H. Kamp, editor, *Ellipsis, Tense and Questions*. ILLC, Amsterdam.

[Zeevat(2000)] Zeevat, H. (2000). Semantics in optimality theory. In H. de Hoop and P. Hendriks, editors, *Optimality Theoretic Semantics*. OTS, Utrecht.

[Zeevat(2001)] Zeevat, H. (2001). The asymmetry of optimality theoretic syntax and semantics. *Journal of Semantics*, **17**(3), 243–262.