

Generating Vague Descriptions

Kees van Deemter

ITRI, University of Brighton

Lewes Road, Watts Building

Brighton BN2 4GJ, United Kingdom

Kees.van.Deemter@itri.brighton.ac.uk

Abstract

This paper deals with the generation of definite (i.e., uniquely referring) descriptions containing semantically vague expressions ('large', 'small', etc.). Firstly, the paper proposes a semantic analysis of vague descriptions that does justice to the context-dependent meaning of the vague expressions in them. Secondly, the paper shows how this semantic analysis can be implemented using a modification of the Dale and Reiter (1995) algorithm for the generation of referring expressions. A notable feature of the new algorithm is that, unlike Dale and Reiter (1995), it covers plural as well as singular NPs. This algorithm has been implemented in an experimental NLG program using PROFIT. The paper concludes by formulating some pragmatic constraints that could allow a generator to choose between different semantically correct descriptions.

1 Introduction: Vague properties and Gradable Adjectives

Some properties can apply to an object to a greater or lesser degree. Such continuous, or *vague* properties, which can be expressed by, among other possibilities, gradable adjectives (e.g., 'small', 'large', e.g. Quirk et al. 1972 sections 5.5 and 5.39), pose a difficult challenge to existing semantic theories, theoretical as well as computational. The problems are caused partly by the extreme *context-dependence* of the expressions involved, and partly by the resistance of vague properties to discrete mathematical modeling (e.g., Synthese 1975, Pinkal 1995). The weight of these problems is increased by fact that vague expressions are ubiquitous in many domains.

The present paper demonstrates how a Natural Language Generation (NLG) program can be enabled to generate uniquely referring descriptions containing one gradable adjective, *despite* the vagueness of the adjective. Having presented a semantic analysis for such vague descriptions, we describe the semantic core of an NLG algorithm that has numerical data as input and vague (uniquely referring) descriptions as output.

One property setting our treatment of vagueness apart from that in other NLG programs (e.g. Goldberg 1994) is that it uses vague properties for an exact task, namely the ruling out of distractors in referring expressions (Dale and Reiter 1995). Another distinctive property is that our account allows the 'meaning' of vague expressions to be determined by a combination of linguistic context (i.e., the Common Noun following the adjective) and nonlinguistic context (i.e., the properties of the elements in the domain).

2 The Meaning of Vague Descriptions

Several different analyses are possible of what it means to be, for example, 'large': larger than average, larger than most, etc. But there is not necessarily just one correct analysis. Consider a domain of four mice, sized 2,5,7, and 10cm.¹ In this case, for example, one can speak of

1. *The large mouse*
(= *the one whose size is 10cm*), and of
2. *The two large mice*
(= *the two whose sizes are 7 and 10cm*).

Clearly, what it takes to be large has not been written in stone: the speaker may decide that 7cm is enough (as in (2)), or she may set the standards higher (as in (1)). A numeral (explicit, or implicit as in (1)), allows the reader to make inferences about the standards employed by the speaker.² More precisely, it appears that in a definite description, the absolute form of the adjective is semantically equivalent with the superlative form:

The n large mice – The largest n mice
The large mice – The largest mice
The large mouse – The largest mouse.

¹For simplicity, the adjectives involved will be assumed to be one-dimensional. Note that the degree of precision reflected by the units of measurement affects the descriptions generated, and even the objects (or sets) that *can* be described, since it determines which objects count as having the same size.

²Thanks are due to Matthew Stone for this observation.

This claim, which has been underpinned by a small experiment with human subjects (see Appendix), means that if a sentence containing one element of a pair is true then so is the corresponding sentence containing the other. There are bound to be differences between the two forms, but these will be taken to be of a *pragmatic* nature, having to do with felicity rather than truth (see section 5.2).

An important qualification must be made with respect to the analysis that we propose: to simplify matters, we assume that the *entire* domain of relevant individuals is available and that it is this domain alone which is taken into account when the adjective is applied. In the case of the example above, this means that all mice are irrelevant except the four that are mentioned: no other knowledge about the size of mice is assumed to be available.³

2.1 A Formal Semantics for Vague Descriptions

Let us be more precise. In our presentation, we will focus on the adjective ‘large’, without intended loss of generality. For simplicity, ‘large’ will be treated as semantically one-dimensional.

i. ‘**The largest n mouse/mice**’. Imagine a set C of contextually relevant animals. Then the NP ‘The largest n mouse/mice’ ($n > 0$) presupposes that there is an $S \subseteq C$ that contains n elements, all of which are mice, and such that (1) $C - S \neq \emptyset$ and (2) every mouse in $C - S$ is smaller than every mouse in S . If such a set S exists then the NP denotes S . The case where $n = 1$, realized as ‘The [Adj]-est [CN_{sg}]’ (sg = singular), falls out automatically.

ii. ‘**The largest mice**’. This account can be extended to cover cases of the form ‘The [Adj]-est [CN_{pl}]’ (pl = plural), where the numeral n is suppressed: these will be taken to be ambiguous between all expressions of the form ‘The [Adj]-est n [CN]’ where $n > 1$. Thus, in a domain where there are five mice, of sizes 4,4,4,5,6 cm, the only possible value of n is 2, causing the NP to denote the two mice of 5 and 6 cm size.

iii. ‘**The n large mouse/mice**’. We analyse ‘The n [Adj] [CN]’ ($n > 0$) as semantically equivalent with the corresponding NP of the form ‘The [Adj]-est n [CN]’. ‘The two large mice’, for example, denotes a set of two mice, each of which is bigger than all other contextually relevant mice.

iv. ‘**The large mice**’. Expressions of this form can be analysed as being of the form ‘The n [Adj] [CN]’ for some value of n . In other words, we will take

³In other words, only *perceptual* context-dependence is taken into account, as opposed to *normative* or *functional* context-dependence Ebeling and Gelman (1994).

them to be ambiguous or unspecific – the difference will not matter for present purposes – between ‘The 2 large mice’, ‘The 3 large mice’, etc.

3 Generation of Crisp Descriptions

Generation of descriptions covers a number of tasks, one of which consists of finding a set L of properties which allows a reader to pick out a given unique individual or set of individuals. The state of the art is discussed in Dale and Reiter (1995), who present a computationally tractable algorithm for characterizing *individuals*. This algorithm (henceforth D&R), deals with vague properties, such as *size*, to some extent, but these are treated as if they were context-independent: always applying to the same sets of objects.

In many cases, generating vague descriptions involves generating a plural and no generally accepted account of the generation of *plural* descriptions has been advanced so far. In the following section, therefore, a generalization or D&R will be offered, called D&R_{Plur}, which focuses on *sets* of individuals. Characterization of an *individual* will fall out as a special case of the algorithm.

3.1 Plural Descriptions: Dale and Reiter generalized

The *properties* which form the basis of D&R_{Plur} are modeled as pairs of the form ⟨Attribute, Value⟩. In our presentation of the algorithm, we will focus on complete properties (i.e., ⟨Attribute, Value⟩ pairs) rather than attributes, as in Dale and Reiter (1995), since this facilitates the use of set-theoretic terminology. Suppose S is the ‘target’ set of individuals (i.e., the set of individuals to be characterized) and C (where $S \subseteq C$) is the set of individuals *from which* S is to be selected.⁴ Informally - and forgetting about the special treatment of head nouns - what happens is the following: The algorithm iterates through a list P in which the properties appear in order of ‘preference’; for each attribute, it checks whether specifying a value for that attribute would rule out at least one additional member of C ; if so, the attribute is added to L , with a suitable value. (The value can be optimized using some further constraints but these will be disregarded here.) Individuals that are ruled out by a property are removed from C . The process of expanding L and contracting C continues until $C = S$. The properties in L can be used by a linguistic realization module to produce NPs such as ‘The white mice’, ‘The white mice that are pregnant’, etc. Schematically, the algorithm goes as follows: (*Notation*: Given a property Q , the set of objects that have the property Q is denoted $[[Q]]$.)

⁴Note that C contains r , unlike Dale and Reiter’s ‘contrast set’ C , which consists of those elements of the domain from which r is set apart.

$L := \emptyset$ {# L is initialized to the empty set #}
 For each $P_i \in P$ do
 If $S \subseteq [[P_i]]$ & $C \not\subseteq [[P_i]]$ {# Adding P_i
 would remove distractors from C #}
 then do
 $L := L \cup \{P_i\}$ {# Property P_i is added
 to L #}
 $C := C \cap [[P_i]]$ {# All elements outside
 $[[P_i]]$ are removed from C #}
 If $C = S$ then Return L {# Success #}
 Return Failure {# All properties in P have been
 tested, yet $C \neq S$ #}

‘Success’ means that the properties in L are sufficient to characterize S . Thus, $\bigcap\{[[P_i]] : P_i \in L\} = S$. The case in which S is a singleton set amounts to the generation of a singular description: $D\&R_{Plur}$ becomes equivalent to $D\&R$ (describing the individual r) when S in $D\&R_{Plur}$ is replaced by $\{r\}$.

$D\&R_{Plur}$ uses hill climbing: an increasingly good approximation of S is achieved with every contraction of C . Provided the initial C is finite, $D\&R_{Plur}$ finds a suitable L if there exists one. Each property is considered at most once, in order of ‘preference’. As a consequence, L can contain semantically redundant properties – causing the descriptions to become more natural, cf. Dale and Reiter 1995 – and the algorithm is polynomial in the cardinality of P .

Caveats. $D\&R_{Plur}$ does not allow a generator to include collective properties in a description, as in ‘the two *neighbouring* houses’, for example. Furthermore, $D\&R_{Plur}$ cannot be employed to generate conjoined NPs: It generates NPs like ‘the large white mouse’ but not ‘the black cat *and* the large white mouse’. From a general viewpoint of generating descriptions, this is an important limitation which is, moreover, difficult to overcome in a computationally tractable account. In the present context, however, the limitation is inessential, since what is crucial here is the interaction between an Adjective and a (possibly complex) Common Noun following it: in more complex constructs of the form ‘NP and the Adj CN’, only CN affects the meaning of Adj.⁵ There is no need for us to solve the harder problem of finding an efficient algorithm for generating NPs uniquely describing arbitrary sets of objects, but only the easier problem of doing this whenever a (nonconjunctive) NP of the form ‘the Adj CN’ is possible.

4 Generation of Vague Descriptions

We now turn our attention to extensions of $D\&R_{Plur}$ that generate descriptions containing the expression

⁵In ‘The elephant and the big mouse’, for example, the mouse does not have to be bigger than *any* elephant.

of one vague property. Case i of section 2.1, ‘The largest n chihuahuas’ will be discussed in some detail. All the others are minor variations.

Superlative adjectives. First, ‘The largest chihuahua’. We will assume that *size* is stored (in the KB that forms the input to the generator) as an attribute with exact numerical values. We will take them to be of the form n *cm*, where n is a positive natural number. For example,

$type = dog, chihuahua$
 $colour = black, blue, yellow$
 $size = 1cm, 2cm, \dots, 10cm.$

With this KB as input, $D\&R$ allows us to generate NPs based on $L = \{yellow, chihuahua, 9cm\}$, for example, exploiting the number-valued attribute *size*. The result could be the NP ‘The 9cm yellow chihuahua’, for example. The challenge, however, is to generate superlatives like ‘The *largest* yellow chihuahua’ instead.

There are several ways in which this challenge may be answered. One possibility is to replace an exact value like *9cm*, in L , by a superlative value whenever all distractors happen to have a smaller size. The result would be a new list $L = \{yellow, chihuahua, largest_1\}$, where ‘largest₁’ is the property ‘being the unique largest element of C ’. This list can then be realized as a superlative NP. We will present a different approach that is more easily extended to plurals, given that a plural description like ‘the 2 large mice’ does not require the two mice to have the same size.

Suppose *size* is the only vague property in the KB. Vague properties are less ‘preferred’ (in the sense of section 3.1) than others (Krahmer and Theune 1999).⁶ As a result, when they are taken into consideration, all the other relevant properties are already in L . For instance, assume that this is the KB, and that the object to be described is c_4 :

$type(c_1, c_2, c_3, c_4) = chihuahua$
 $type(p_5) = poodle$
 $size(c_1) = 3cm$
 $size(c_2) = 5cm$
 $size(c_3) = 8cm$
 $size(c_4) = size(p_5) = 9cm$

At this point, *inequalities* of the form $size(x) > m$ *cm* are added to the KB. For every value of the form n *cm* occurring in the *old* KB, all inequalities of the form $size(x) > n$ *cm* are added whose truth *follows* from the old KB. Inequalities are more

⁶Note, by contrast, that vague properties tend to be *realized first* (Greenbaum et al. 1985, Shaw and Hatzivassiloglou 1999). Surface realization, however, is not the topic of this paper.

preferred than equalities, while logically stronger inequalities are more preferred than logically weaker ones.⁷ Thus, in order of preference,

$$\begin{aligned} \text{size}(c_4), \text{size}(p_5) &> 8cm \\ \text{size}(c_3), \text{size}(c_4), \text{size}(p_5) &> 5cm \\ \text{size}(c_2), \text{size}(c_3), \text{size}(c_4), \text{size}(p_5) &> 3cm. \end{aligned}$$

The first property that makes it into L is ‘chihuahua’, which removes p_5 but not c_4 from the context set. (Result: $C = \{c_1, \dots, c_4\}$.) Now size is taken into account, and the property $\text{size}(x) > 8cm$ singles out c_4 . The resulting list is $L = \{\text{chihuahua}, > 8cm\}$. This implies that c_4 is *the only* chihuahua in the KB that is greater than $8cm$ and consequently, the property $\text{size}(x) > 8cm$ can be replaced, in L , by the property of ‘being larger than all other elements of C ’. The result is a list that may be written as $L = \{\text{chihuahua}, \text{largest}_1\}$, which can be employed to generate the description ‘the largest chihuahua’.

Plurals can be treated along analogous lines. Suppose, for example, the facts in the KB are the same as above and the target set S is $\{c_3, c_4\}$. Its two elements share the property $\text{size}(x) > 5cm$. This property is exploited by $D\&R_{Plur}$ to construct the list $L = \{\text{chihuahua}, > 5cm\}$. Analogous to the singular case, the inequality can be replaced by the property ‘being a set all of whose elements are larger than all other elements of C ’ (largest_n , for short), leading to NPs such as ‘the largest chihuahuas’. Optionally, the numeral may be included in the NP (‘the *two* largest chihuahuas’).

– ‘**Absolute**’ adjectives. The step from the superlative descriptions of case **i** to the analogous ‘absolute’ descriptions is a small one. Let us first turn to case **iii**, ‘The n large mouse/mice’. Assuming the correctness of the semantic analysis in section 2, the NP ‘The n large mouse/mice’ is semantically equivalent to the one discussed under **i**. Consequently, an obvious variant of the algorithm that was just described can be used for generating it. (For pragmatic issues, see section 5.2)

Finally, case **iv**, ‘The large mice’. Semantically, this does not introduce any new problems, since it is to case **iii** what case **ii** is to case **i**. According to the semantic analysis of section 2.1, ‘The large mice’ should be analysed just like ‘The n large mouse/mice’, except that the numeral n is suppressed. This means that a simplified version (i.e., without a cardinality check) of the algorithm that takes care of case **iii** will be sufficient to generate descriptions of this kind.

⁷E.g., $\text{size}(x) > m$ is preferred over $\text{size}(x) > n$ iff $m > n$. The preference for *inequalities* causes the generator to avoid the mentioning of measurements unless they are needed for the identification of the target object.

5 Conclusions and loose ends

We have shown how vague descriptions can be generated that make use of one vague property. We believe our account to be an instructive model of how the ‘raw data’ in a standard knowledge base can be presented in English expressions that have a very different structure. The numerical data that are the input to our algorithm, for example, take a very different form in the descriptions generated, and yet there is, in an interesting sense, no loss of information: a description has the same reference, whether it uses ‘exact’ information (‘The $3cm$ mouse’) or ‘vague’ information (‘The large mouse’).⁸

5.1 Limitations of the semantic analysis

Our proposal covers the generation of vague descriptions ‘from absolute values’, which is argued in Dale and Reiter (1995, section 5.1.2) to be most practically useful. When vague input is available (e.g., in the generation component of a Machine Translation system, or in WYSIWYM-style generation (Power and Scott 1998)), simpler methods can be used. Our own account is limited to the generation of definite descriptions and no obvious generalization to indefinite or quantified NPs exists. Other limitations include

- Descriptions that contain properties for other than individuating reasons (as when someone asks you to clean ‘the dirty table cloth’ when only one table cloth is in sight). This limitation is inherited directly from the $D\&R$ algorithm that our own algorithm extends.
- Descriptions containing more than one vague property, such as ‘The fat tall bookcase’, whose meaning is more radically unclear than that of definite descriptions containing only one vague term. (The bookcase may be neither the fattest nor the tallest, and it is not clear how the two dimensions are weighed.)
- Descriptions that rely on the salience of contextually available objects. Krahmer and Theune (1998) have shown that a contextually more adequate version of $D\&R$ can be obtained when degrees of salience are taken into account. Their account can be summarized as analysing ‘the black dog’ as denoting the unique most salient object in the domain that is both black and a dog. (Generalizations of this idea to $D\&R_{Plur}$ are conceivable but nontrivial since not all elements of the set S have to be equally salient.) Our own extensions of $D\&R$ (and perhaps $D\&R_{Plur}$) could be ‘contextualized’ if the

⁸This may be contrasted with the vague expressions generated in (Goldberg et al. 1994), where there is a real – and intended – loss of information. (E.g., ‘Heavy rain fell on Tuesday’, based on the information that the rainfall on Tuesday equalled $45mm$.)

role of salience is changed slightly: focusing on the singular case, the algorithm can, for example, be adapted to legislate that ‘the large(est) mouse’ denotes the largest of all those mice that are salient (according to some standard of salience). Note that this analysis predicts *ambiguity* when the largest mouse that is salient according to one standard is smaller than the largest mouse that is salient according to a more relaxed standard. Suppose, for example,

Salient (strict):

m1 (2cm), *m2* (5cm)

Salient (relaxed):

m1 (2cm), *m2* (5cm), *m3* (7cm);

then ‘the large(est) mouse’ may designate either *m2* or *m3* depending on the standards of salience used. What this illustrates is that salience and size are both vague properties, and that – as we have seen under point b – combining vague properties is a tricky business.

5.2 Pragmatics

An experimental PROFIT (Erbach 1995) program has implemented the algorithms described so far, generating different descriptions, each of which would allow a reader/hearer to identify an object or a set of objects. But of course, an NLG program has to do more than determine under what circumstances the use of a description leads to a true statement: an additional problem is to choose the most appropriate description from those that are semantically correct. This makes NLG an ideal setting for exploring issues that have plagued semanticists and philosophers when they studied the meaning of vague expressions, such as whether it can be *true* for two objects *x* and *y* which are indistinguishable in size that *x* is large and *y* is not (e.g. Synthese 1975). The present setting allows us to say that a statement of this kind may be true yet infelicitous (because they conflict with certain pragmatic constraints), and consequently to be avoided by a generator.

As for the choice between the ‘absolute’/superlative forms of the gradable adjective, we conjecture that the following constraints apply:

- C1. **Distinguishability.** Expressions of the form ‘The (*n*) large [CN]’ are infelicitous when the *smallest* element of the designated set *S* (named *x*) and the *largest* CN smaller than all elements of *S* (named *y*) are perceptually indistinguishable.
- C2. **Natural Grouping.** Expressions of the form ‘The (*n*) large [CN]’ are better avoided when the difference in size between *x* and *y* is ‘comparatively’ small. One way of making this precise is by requiring that the difference between *x* and

y cannot be smaller than that between either *x* or *y* and one of their neighbouring elements. Consider, for example, a domain consisting of mice that are 1cm, 1cm, 2cm, 7cm, 9cm and 9cm large; then C2 predicts that the only felicitous use of ‘the large mice’ refers to the largest three of the group.

- C3. **Minimality.** Otherwise, preference is given to the absolute form. This implies that when objects of only two sizes are present, and the difference is perceptually distinguishable, the absolute form is preferred over the superlative form. (For example, in a domain where there are two sizes of pills, we are much more likely to speak of ‘the large pills’ than of ‘the largest pills’.) In languages in which the superlative form is morphologically more complex than the absolute form, constraint C3 can be argued to follow from general Gricean principles (Grice 1975)).

As for the presence/absence of the *numeral*, we conjecture that the disambiguating numeral (as in ‘the *n* large mice’ or ‘the *n* largest mice’) can be omitted under two types of circumstances: (1) when any ambiguity resulting from different values of *n* is likely to be inconsequential (see Van Deemter and Peters (1996) for various perspectives); (2) when the domain allows only one ‘natural grouping’ (in the sense of C2). Before and until a more accurate version of the notion of a natural grouping is available (perhaps using fuzzy logic as in Zimmermann 1985), generators could be forbidden to omit the numeral, except in the case of a definite description in the singular.

Appendix: A Supporting Experiment

Human subjects were asked to judge the correctness of an utterance in a variety of situations. The experiment was set up to make plausible that, in a situation in which only perceptual context-dependence (see section 1) is relevant, expressions of the form ‘the *n* large CN’ can be used whenever certain simple conditions are fulfilled. Note that this (\Rightarrow) direction of the hypothesis is most directly relevant to the design of a generator, since we expect a generator to avoid mistakes rather than always use an expression whenever it is legitimate.

Hypothesis (\Rightarrow): In a situation in which the domain *D* represents the set of perceptually relevant objects, an expression of the form ‘the *n* large CN’ (where $n \geq 1$), can be used to refer to a set *S* of cardinality *n* if all objects in $D - S$ are smaller than any of the *n*.

The experiment explores whether ‘the n large CN’ can refer to the n largest objects in the domain, whether or not this set of objects is held together by spatial position or other factors. Subjects were presented with 26 different situations, in each of which they had to say whether the sentence

The two high numbers appear in brackets

would constitute a *correct* utterance. The literal text of our question was:

Suppose you want to inform a hearer *which numbers in a given list appear in brackets*, where the hearer knows what the numbers are, but not which of them appear in brackets. For example, the hearer knows that the list is 1 2 1 7 7 1 1 3 1. You, as a speaker, know that only the two occurrences of the number 7 appear in brackets: 1 2 1 (7) (7) 1 1 3 1. Our question to you is: Would it be *correct* to convey this information by saying “The two high numbers appear in brackets”? (...).

All subjects were shown the 26 situations in the same, arbitrary, order. Each situation presented to the subjects contained a list of nine numbers. In 24 cases, the lists had the following form:

1 1 1 x y z 1 1 1,

where each of x, y, z equalled either 6 or 9, and where there were always two numbers among x, y, z that appear in brackets. In 16 out of 24 cases, the two bracketed positions are right next to each other, allowing us to test whether spatial contiguity plays any role. Subjects were presented with two additional situations, namely 1 1 1 (6) 1 (7) 1 1 1 and 1 1 1 (7) 1 (6) 1 1 1 in which, unlike the other 24 situations, the two largest numbers are not equally large, to make sure that the descriptions do not require the elements in their denotation to be similar in that respect. Our questions were presented via email to 30 third-year psychology/cognitive science students at the University of Durham, UK, all of whom were native speakers of English and ten of which responded.

Results: Eight subjects responded in exact conformance with the analysis of section 2.1, marking all and only those five sequences in which *the highest 2* numbers appeared in brackets. Only two subjects deviated slightly from this analysis: one of the two (subject 9) described all the expected situations as ‘correct’ plus the two cases in which two contiguous 6-es appeared in brackets; the other subject (subject 10) appears to have made a typing error, confusing

two subsequent situations in the experiment.⁹ All other responses of subjects 9 and 10 were as predicted. This means that all subjects except subject 10 were consistent with our ‘ \Rightarrow ’ hypothesis. The experiment suggests that the converse of the hypothesis might also be true, in which it is claimed that expressions of the form ‘the n large CN’ *cannot* be employed to refer to the set S *unless* S consists of the n largest objects in D :

Hypothesis (\Leftarrow): In a situation in which the domain D represents the set of perceptually relevant objects, an expression of the form ‘the n large CN’ (where $n \geq 1$), can *only* be used to refer to a set S of cardinality n if all objects in $D - S$ are smaller than any of the n .

Again disregarding subject 10, eight out of nine subjects act in accordance with Hypothesis \Leftarrow , while only one appears to follow a somewhat more liberal rule. Given these findings, it appears to be safe to build a generator that implements *both* hypotheses, since none of our subjects would be likely to disagree with any of the descriptions generated by it.

This experiment has evident limitations. In particular, it has no bearing on the pragmatic constraints suggested in section 5.2, which might be tested in a follow-up experiment.

Acknowledgements

Thanks are due to: Richard Power for discussions and implementation; Emiel Kraemer, Ehud Reiter and Matthew Stone for comments on an earlier draft; Hua Cheng for observations on linguistic realization; Rosemary Stevenson and Paul Piwek for their help with the experiment described in the Appendix.

6 References

- Dale and Reiter 1995. R. Dale and E. Reiter. Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions. *Cognitive Science* 18: 233-263.
- Ebeling and Gelman 1994. Ebeling, K.S., Gelman S.A. 1994. Children’s use of context in interpreting “big” and “little”. *Child Development* 65(4): 1178-1192.
- Erbach 1995. G. Erbach. Web page on the PROFIT

⁹The situations that we suspect to have been confused are 1 1 1 (9) (9) 9 1 1 1, which was marked as correct (although, remarkably, none of the other ‘three nines’ situations were marked as correct) and 1 1 1 (9) (9) 6 1 1 1.

- programming language, <http://coli.uni-sb.de/erbach/formal/profit/profit.html>.
- Goldberg et al. 1994. E. Goldberg, N. Driedger, and R. Kitteridge. Using Natural-Language Processing to Produce Weather Forecasts. *IEEE Expert* **9** no.2: 45-53.
 - Greenbaum et al. 1985. "A Comprehensive Grammar of the English Language". Longman, Harlow, Essex.
 - Grice 1975. P. Grice. Logic and Conversation. In P. Cole and J. Morgan (Eds.), "Syntax and Semantics: Vol 3, Speech Acts": 43-58. New York, Academic Press.
 - Krahmer and Theune 1999. E. Krahmer and M. Theune. Generating Descriptions in Context. In R. Kibble and K. van Deemter (Eds.), Procs. of workshop The Generation of Nominal Expressions, associated with the 11th European Summer School in Logic, Language, and Information (ESSLLI'99).
 - Pinkal 1995. M. Pinkal. "Logic and Lexicon". Oxford University Press.
 - Power and Scott 1998. R. Power and D. Scott. Multilingual Authoring using Feedback Texts. In Proc. COLING/ACL, Montreal.
 - Quirk et al. 1972. R. Quirk, S. Greenbaum, and G. Leech. "A Grammar of Contemporary English". Longman, Harlow, Essex.
 - Shaw and Hatzivassiloglou 1999. Ordering Among Premodifiers. In Procs. of ACL99, Univ. Maryland.
 - Synthese 1975. Special issue of the journal *Synthese* on semantic vagueness. *Synthese* **30**.
 - Van Deemter and Peters 1996. K. van Deemter and S. Peters (Eds.) "Semantic Ambiguity and Underspecification". CSLI Publications, Stanford.
 - Zimmermann 1985. H. J. Zimmermann. "Fuzzy Set Theory - and its Applications". Kluwer Academic Publishers, Boston/Dordrecht/Lancaster.