

Defeasible Rules in Content Selection and Text Structuring*

Alistair Knott Mick O'Donnell Jon Oberlander Chris Mellish

Human Communication Research Centre, Department of Artificial Intelligence
University of Edinburgh

Abstract

This paper outlines a number of ways in which defeasible rules can contribute to the content selection and discourse structuring components of a text generation system. We suggest that, for certain types of descriptive text, the characterisation of discourse structuring mechanisms as operations on or involving defeasible rules provides an attractive framework for addressing important issues in content selection/structuring. We describe an architecture which incorporates defeasible rules into a systemic model of generation, and illustrate its use in the description of objects in a museum gallery.

While defeasible rules are traditionally used in theorem-proving applications, to make predictions about the consequences of known facts, we are here concerned with three separate issues: (i) how such rules may need to be expressed by the NLG system in order to achieve its goals; (ii) how their interaction with facts about particular objects enables the use of valuable coherence relations; and (iii) how they are relevant in taking account of the system's user model.

1 Introduction: the interaction of general and particular facts in the museum domain

This paper describes aspects of ILEX-1, a text generator which operates in a museum gallery, producing descriptions of a series of objects encountered during a virtual visit. The user of the system is free to choose any objects in the gallery to look at, in any order; the goal of the system is to produce an informative description for each object chosen, in such a way as the sequence of object descriptions eventually produced forms a coherent whole.

One of the interesting issues which this domain allows us to explore is the way in which a finite collection of objects can be used to make statements about generic classes of objects—‘**generalisations**’. Artefacts in a museum are often significant wholly or partly in being *representative* of larger classes of objects in the world: the important lessons to be learned in a museum often relate to these generalisations, rather than to the actual objects seen. Human learning hinges on the ability to use facts about particular entities to illustrate, test, or even establish generalisations about classes of entities. The issue in the present paper is how best to take account of this ability in users of the system—indeed, how to take advantage of it—when structuring object descriptions. Our main suggestion is that a notion of **defeasible rules**, as used in a number of nonmonotonic logics, can play a useful role to this end.

*This work was supported by the UK Engineering and Physical Sciences Research Council, through grant GR/K53321. The support of the Economic and Social Research Council for HCRC is also gratefully acknowledged.

We will begin in Section 2 by illustrating the benefits of attention to generalisations in the description of objects in the museum domain. In Section 3, we introduce the notion of defeasible rules, which provides a means of representing generalisations with the right properties for our purposes. In Section 4, we describe three possible uses for these rules, for the generation of **quantifiers**, **coherence relations**, and **misconception corrections**. Section 5 describes a simple implementation of defeasible rules for these three purposes. Finally, in Section 6 we propose a number of extensions of the current implementation, which examine some interesting interactions between these three areas.

2 The use of generalisations: an illustration

Consider the following two texts, each of which describes a necklace in an exhibition of 20th century jewellery.

- (1) This necklace is in the arts-and-crafts style. It is made of silver, amethysts and pearls. It has very elaborate festoons. It has faceted stones.
- (2) This necklace is in the arts-and-crafts style. It is made of silver, amethysts and pearls. Arts-and-crafts jewels tend to be intricately worked; for instance, this piece has very elaborate festoons. However, unusually for arts-and-crafts jewellery, this piece has faceted stones. Most arts-and-crafts jewels (see for example the jewels in case 8) have cabochon stones.

Leaving aside any improvements that might result from aggregation, or from definitions of terms like *festoons* and *cabochon stones*, it should still be clear that the second example is far preferable to the first. This is true for several reasons. Firstly, Example 2 contains more information than Example 1: it informs about the class of arts-and-crafts jewels as well as about the particular jewel being described. This is likely to be the sort of material that the museum curator is really trying to get across. Secondly, this additional material allows a greater degree of *cohesion* in Example 2 than is possible in Example 1. In Example 1, the propositions describing the properties of the jewel being viewed do not stand in obvious relationships to each other; the most that we could do to improve the acceptability of the text is to aggregate, which would tend to obscure rather than accentuate the interesting relations. Finally, generalisations allow references to be made to other jewels that have or could be described, and are thus helpful in weaving together a sequence of object descriptions into a unified text. In the ILEX-1 domain, this is a particularly important function.

As should be clear from these examples, the use of generalisations, and of devices to signal their interaction with particular facts, combine to produce a much richer representation of the domain being described. In this paper we will be looking for a unified treatment of such phenomena, to allow a NLG system to produce similarly rich descriptions.

3 Defeasible rules

The starting point is to find a means of representing the semantics of generalisations. To represent a universal generalisation is unproblematic; the quantified expressions *all Xs are Y* and *no Xs are Y* can be used. However, many generalisations are vague, and admit of exceptions; as Example 2 makes clear, we also need a way of expressing these.

Many formalisms have been proposed for capturing the semantics of generic expressions; see for instance Carlson and Pelletier [3] for an overview. We will be adopting the terminology

of Asher and Morreau’s [1] commonsense entailment in what follows. In their notation, ϕ ’s *normally* ψ is expressed as $\forall x(\phi > \psi)$; the operator $>$ is a primitive in the logic, tailored for exactly this purpose, and the axioms of the logic are built around it. We will refer to structures such as $\phi > \psi$ as defeasible rules, and in the following we will refer to ϕ as the LHS and ψ as the RHS.¹ An example from the jewellery domain is given below:²

$$(3) \quad \textit{arts-and-crafts-jewel}(X) > \textit{intricately-worked}(X)$$

4 Three roles for defeasible rules

The communicative goals of the ILEX-1 system are of two sorts. Firstly, there are **content goals**, which are goals to communicate to the user certain key facts about the domain. Secondly, there are what we can call **formal goals**, which are goals to produce a text where these facts are well linked together. Defeasible rules can play a role in achieving both types of goals, as we will show in this section. We begin by considering the expression of defeasible rules; we then consider their use in definitions of certain coherence relations; and finally we consider their use in modelling reader misconceptions.

4.1 Defeasible rules and vague quantifiers

The most obvious application of defeasible rules is simply to represent facts to be expressed. According to the museum education professionals we have consulted, many of the important content goals relate to the communication of generalisations, and so we need a simple method for expressing them.

Perhaps the simplest way of expressing such facts is in terms of vague quantifiers: the expression *most Xs are Y* can be used to express $X > Y$ and the expression *few Xs are Y* to express $X > \neg Y$. Note that this suggestion is not trivial; in fact, it makes the powerful prediction that a reader presented with expressions such as these will exhibit patterns of reasoning consistent with the axioms of commonsense entailment. Such a hypothesis is rather simplistic, particularly as there are significant differences between *most Xs are Y* and the more standard generic *Xs are normally Y*; however, it leads to some productive suggestions, and will be assumed in what follows.

4.2 Defeasible rules and coherence relations

A second application for defeasible rules relates to the system’s formal goals, to link together the facts to be expressed in a coherent way. Coherence relations are often used to this end, and a number of these relations can be usefully defined in terms of defeasible rules. A number of coherence relations function to link generalisations to particular facts; in particular what we might call EXEMPLIFICATION, AMPLIFICATION and CONCESSION as illustrated below:³

$$(4) \quad \textit{Arts-and-crafts jewels tend to be intricately worked. For example, this piece has very elaborate festoons.}$$

¹The universal quantifier will be dropped henceforth, but should be assumed.

²More complex treatments of generic expressions, modelling the effect of focus on their interpretation (e.g. Carlson [2], Krifka [8], Cavedon and Glasbey [4]) might eventually be used; we are still exploring whether these will prove necessary in our domain.

³We will take the presence or suitability of the three connectives in these examples as definitional of their respective relations.

- (5) This piece has very elaborate festoons. *Indeed*, so do most arts-and-crafts jewels.
- (6) Arts-and-crafts jewels tend to be intricately worked, *but* this piece has clean, geometric lines.

Discourse theorists have often analysed EXEMPLIFICATION and AMPLIFICATION in terms of general rules: see for instance Mann and Thompson [10], and related discussion on implicature in Hirschberg [7]. It only remains to be noted that these rules can be defeasible as well as completely universal; the above examples each illustrate the former case. Recently it has also been argued that CONCESSION should be analysed in terms of defeasible rules; see for instance Grote *et al* [6], Overteegen [11]. According to these analyses, a defeasible rule can either be explicitly stated in the satellite of a CONCESSION (the first clause in the case above), or presupposed, as in the following case:

- (7) This piece is in the arts-and-crafts style, *but* it has clean, geometric lines.

Grote *et al* note several rhetorical uses for the CONCESSION relation, two of which are relevant in the present domain. Firstly, they can be used simply to prevent an incorrect implicature being drawn about a particular instance. If, for instance, we didn't add the second clause in Example 6, the reader might think that the object being described was intricately worked. Secondly, they can be used simply to inform about unusual circumstances in the world. While objects in a museum sometimes serve to represent a generic class of objects (as mentioned above), they can also be interesting precisely because they differ from the norm in some respect.

4.3 Defeasible rules and the user model

Finally, defeasible rules serve a useful purpose in relation to the user model. In describing an object or situation, it is necessary to take into account any misconceptions that a reader might have about it; the way a fact is expressed often depends on whether or not it presupposes such a misconception (see for instance the literature on negative polarity items). The point is that these misconceptions often derive from misconceived generalisations; addressing these is thus particularly important. Correcting such a generalisation can involve several types of quantifier and of coherence relation:

- (8) *Not all* arts-and-crafts jewellery is elaborate. *For example*, this piece is quite plain. *Indeed*, many arts-and-crafts jewels are plain.

Note that in an example such as this one both content goals and formal goals are being achieved: the quantified expressions in the first and third sentences convey important items of content, while the EXEMPLIFICATION and AMPLIFICATION relations each contribute towards the formal goal of coherent text.

A more complex issue to consider is how a generalisation is reached on the basis of individual instances. The reader's ability to induce generalisations on the basis of what he has seen must be taken into account when keeping track of the user model; for instance, if he has seen several elaborate arts-and-crafts jewels, the misconception addressed in Example 8 might be quite a natural one. Consider also the following sequence of descriptions.

- (9) **Description of J1:** This piece was designed by Jessie King, in 1905.
Description of J2: This piece was *also* designed by Jessie King, *but* around 1910.

How can we explain the *but* in the description of J2? Its function seems to be to alert the reader to a difference between J2 and J1, which needs to be made explicit in view of the similarity noted in respect of their designer. So far, we have thought of *but* as signalling a case where a defeasible rule is defeated. What might the defeasible rule be in this case? Presumably, something like the following:

(10) If Jessie King designed it, it was probably designed in 1905.

What is striking about this defeasible rule is that it is not one which needs to be stated explicitly in the text in order to be internalised. All that is needed is for the two facts ‘J1 was designed by Jessie King’ and ‘J1 was designed in 1905’ to be expressed. It seems to indicate that a defeasible rule can be internalised extremely quickly by a reader.

5 Generation using defeasible rules

In this section, we will describe the simple treatment of defeasible rules that is currently implemented in ILEX-1.

5.1 The Knowledge Base

The ILEX-1 Knowledge Base (KB) consists of two types of objects: *facts* and *rules*. Facts may involve either unary predicates (facts about *types* of objects, e.g. *elaborate(j1)*), or binary predicates (facts about relations between objects, e.g., *made-of(j1,gold)*).⁴ Rules can be either *indefeasible* or *defeasible* and the LHS can only be satisfied by facts (there is no chaining). Here we are only concerned with the defeasible rules.

5.2 Matching Defeasible Rules to the KB

One of the traditional problems with defeasible logics is their computational complexity; see e.g. Lifschitz [9]. However, complexity only becomes a problem when defeasible rules are used by a theorem-prover, to make inferences about what is not known. We are not in fact using them for this purpose, and we thus avoid the complexity problem. We are only using defeasible rules as a means of linking together existing facts about an object. When we find an object for which both the LHS and the RHS of a defeasible rule hold, then we can express not only the LHS and RHS independently, but we can also state that this combination of facts is meaningful and expected, as in the first part of Example 2.

When a defeasible rule is defeated, its LHS holds for a particular object, while its RHS does not. The object is thus seen as an exception to the rule. This also is worth stating, as in the second part of Example 2.

Two other results are possible: the LHS does not hold for the object, in which case the rule does not apply, and nothing is to be said; or secondly, the LHS does apply, but the KB has insufficient information to confirm or deny the RHS. In the latter case, we could possibly express a hedge; for instance, *being costume jewellery, we would expect this jewel to be cheaply made*.

All matching of defeasible rules to the objects they generalise over is done as a pre-compilation step, avoiding costly on-line calculation.

⁴This is a simplification: the ILEX-1 KB actually uses a fully-fledged typed unification system.




Figure 1: Discourse Structure for an EXEMPLIFICATION Relation

5.3 Expressing Defeasible Rules

When a jewel is selected, the content selection algorithm begins by finding all the simple facts that can be expressed about the jewel. Then, for all of these, a search is made of the defeasible/indefeasible rules which can be expressed in connection with these facts. For a rule to be expressible, the general class about which the rule holds must be introduced by a simple fact (as, for instance, in *this piece is an instance of arts-and-crafts jewellery*); otherwise, the expression of the rule will have the appearance of an inexplicable change of subject. The fact is then linked to the generalisation via the coherence relation DEFINITION. The generalisation is then in turn linked back to another simple fact about the jewel, by an appropriate relation: either EXEMPLIFICATION (if a fact is found which accords with the rule) or CONCESSION (if one is found which does not). An illustration of the discourse structure created for an EXEMPLIFICATION relation is given in Figure 1.

In the case of multiple facts which can be linked to a given rule, these are given as multiple satellites of the DEFINITION relation.

A similar procedure is used for misconception corrections, except that the rules searched are the mal-rules in the user model, rather than the rules in the domain KB. Only the EXEMPLIFICATION relation is currently used to provide counterexamples to misconceptions; but the following section suggests some interesting possible extensions.

6 Further interactions between defeasible rules and quantifiers

In the present section, the relationship between defeasible rules, coherence relations and quantifiers is explored in more detail, with a view to a more refined and theoretically interesting implementation of the ideas mentioned above.

6.1 A classification of quantifiers, including *many*

Until now, we have discussed cases where an object counts as an exception to a defeasible rule, where a CONCESSION is appropriate, and cases where an object serves to falsify an *indefeasible* rule which is mistakenly held to be true, where the quantifiers *not all* or *some* are appropriate. A third possibility should now be considered: if indefeasible rules can be shown to be untrue, then so, presumably, can defeasible ones. But how might this be achieved? Clearly, providing a single counterexample will not be enough; the whole point about defeasible rules is that they admit of exceptions. And yet it is clearly possible that a defeasible rule can be mistakenly held. So what are the criteria for deciding whether a defeasible rule is right or wrong in a particular case?

NL sentence	Representation	Description
<i>All Xs are Y</i>	$X \rightarrow Y$	indefeasible rule from X to Y asserted
<i>No Xs are Y</i>	$X \rightarrow \neg Y$	indefeasible rule from X to $\neg Y$ asserted
<i>Some Xs are Y</i>	$\neg(X \rightarrow \neg Y)$	indefeasible rule from X to $\neg Y$ denied
<i>Not all Xs are Y</i>	$\neg(X \rightarrow Y)$	indefeasible rule from X to Y denied
<i>Most Xs are Y</i>	$X > Y$	defeasible rule from X to Y asserted
<i>Few Xs are Y</i>	$X > \neg Y$	defeasible rule from X to $\neg Y$ asserted
<i>Many Xs are Y</i>	$\neg(X > \neg Y)$	defeasible rule from X to $\neg Y$ denied
<i>Many Xs are not Y</i>	$\neg(X > Y)$	defeasible rule from X to Y denied

Table 1: Rule-Based Definitions for ‘Vague’ and ‘Precise’ Quantifiers

Some terminology would be useful at this point. It is confusing to talk about a defeasible rule being ‘defeated’ by some facts; this could either mean that the facts count as exceptions to the rule, while the rule still stands, or it could mean that there are enough facts to suggest that the rule is actually wrong. In this latter case, we can speak of the rule as having been **overturned**; the term **defeated** will be reserved for the former case.

In order to overturn a defeasible rule *Most Xs are Y*, it is necessary to gather a certain number of instances of X s which are not Y . The hypothesis we propose here is that another vague quantifier, *many*, can be characterised as identifying just this number.

This is an attractive picture, because it provides a semantics for all of the principal quantifiers. *All* and *no* are used for asserting indefeasible rules. *Most* and *few* are used for asserting defeasible rules. *Some* and *Not all* are used for overthrowing indefeasible rules. *Many* and *Many...not* are used for overthrowing defeasible rules. Table 1 presents a summary of this picture of quantifiers.

The definition of *many Xs are Y* as denying the defeasible rule $X > \neg Y$ is central to this classification. Two arguments will be given to support this definition, which both hinge on the interaction between quantifiers and coherence relations.

Firstly, note an interesting interaction between the relation of CONCESSION and the quantifiers *some*, *many* and *most*. It is quite legitimate to link *most* and *some* using this relation:

(11) Most art-deco jewels are geometric. But some of them are asymmetrical.

But it seems a little strange to use the relation to oppose *most* and *many*:

(12) ? Most art-deco jewels are geometric. But many of them are asymmetrical.

The intuition that this latter example is unusual would need to be tested empirically, as it is not very strong. However, if the text does seem strange, then the proposed model can account for it: the second clause would be seen as doing more than indicating an exception to the rule expressed in the first clause; it would be overthrowing it. In which case, it is not appropriate to present the text as a concession. It would be more appropriate to present it as a straightforward disagreement:

(13) A: Most art-deco jewels are geometric.
B: That’s not true. Many of them are asymmetrical.

The fact that this seems to be a legitimate pattern of argument provides at least some support for our definition of *many*.

A second line of support is that the proposed definitions of quantifiers gives a rationale for the existence of another group of cue phrases, namely *indeed* and *in fact*. As has often been pointed out (e.g. Grice [5]), these phrases are useful as blockers of quantity implicatures, and are often used to replace one quantifier with another further along the monotonic scale. For instance:

(14) Many art-deco jewels are geometric. *Indeed*, most of them are.

The question is: why would a writer ever want to present the generalisation in two parts like this, rather than just coming straight out with the stronger quantifier? Such a policy becomes easier to understand if the former statement is seen as countering an existing supposition, for instance that *few art-deco jewels are geometric*.

This view is further supported by the suggestion that *indeed* does not operate transitively along the scale of monotonically increasing/decreasing quantifiers. There's something a little strange about the following two texts:

(15) ? Most art-deco jewels are geometric. *Indeed*, all of them are.

(16) ? Few art-deco jewels are geometric. *Indeed*, none of them are.⁵

Note that the problem is not simply due to the fact that the clauses introduced by *indeed* feature quantifiers at the extreme ends of the scale. There are no problems with the following texts, for instance:

(17) Many art-deco jewels are geometric. *Indeed*, all of them are.

(18) Not all art-deco jewels are geometric. *Indeed*, none of them are.

We now have a story which explains the difference in the behaviour of *indeed*. According to this story, the quantifiers *most* and *few* are not used to overthrow rules (defeasible or indefeasible), and therefore the implicatures which follow from them are hard to block with *indeed*; whereas the quantifiers *many* and *not all* can be used to overthrow rules (defeasible for *many*, and indefeasible for the other two), and therefore *indeed* can be used to block the implicatures which follow from them.

An objection to this account might be voiced on the basis of the complex quantifiers *most if not all* and *few if any*. These constructions are commonly seen as blocking implicatures following from *most* and *few* respectively. However, these cases are significantly different from those using *indeed* mentioned above. In the conditional constructions, it is only the *possibility* of the universally-quantified proposition which is asserted; while in the examples with *indeed*, the quantified expression is asserted categorically. Thus in the examples with *indeed* the implicatures are first drawn and then cancelled, while in those with the conditional construction they are prevented from being drawn in the first place. This difference points to very different uses for the two kinds of expression. The conditional expressions are suitable for use when the speaker simply has not encountered any exceptions to the generalisation in question, be it positive or negative, but has encountered enough instances of it to be able to affirm it. This is quite different from the use of *indeed* discussed above.

⁵Note that with *in fact* (often a good substitute for *indeed*) these texts are a lot better. But we would argue that this is because *in fact* can also be used to express a straightforward retraction of one claim, and the assertion of an amended one (in which case it is often accompanied by an apology, or a qualification such as *Well...*)

6.2 Choosing a quantifier for misconception corrections

The eight quantifiers listed in Table 1 seem to cover the eight possible attitudes to a rule quite neatly. However, it should be noted that many of the possible attitudes can be expressed in more than one way, using an alternative quantifier coupled with negation. For instance, note that instead of saying *some Xs are Y* to deny that $X \rightarrow \neg Y$, we could say *not all Xs are not Y*, which in fact makes the denial of the rule more explicit. How might we decide which version to produce when generating a text? This question will be addressed below.

6.2.1 Property-introducing and presupposition-denying quantifiers

Some Xs are Y should in fact be thought of as ambiguous. It can either be seen as countering the assertion that *no Xs are not Y* (as it has been analysed above), or simply as asserting a positive statement, in the absence of any presuppositions. What can we say about its use in this latter sense?

We suggest that when used in this way, without presuppositions, a quantifier is better understood in terms of a focusing account of text. If a writer is talking about one topic, and wants to move to a related topic, it is awkward simply to begin a new sentence with the new topic in subject position. Rather, the new topic should be *introduced* in a sentence about the current topic, so as to show how they are related; the following sentence can then make a smooth shift to the new topic. Now, imagine that the current topic is not a single entity, but a class of entities, and the writer wants to talk about a property which belongs to some of these. (Simply because it is an interesting thing to talk about.) The natural way to do so is in a quantified sentence. When it comes to choosing a quantifier, *all Xs are Y*, *some Xs are Y*, *most Xs are Y* and *many Xs are Y* can all be used simply to introduce an interesting property *Y*, in the absence of any presuppositions at all. But note that this is certainly not true of *no Xs are Y*, *not all Xs are Y*, *few Xs are Y* and *many Xs are not Y*.

We would predict that the scalar implicatures triggered by property-introducing quantifiers are different from those triggered by presupposition-denying quantifiers. If the writer's main concern is simply to introduce an interesting subset of objects, rather than to make a point about the size of this subset, we should not expect the choice of quantifier used to do so to carry as much significance. For instance, if a writer simply wants a way to introduce the property *Y* in the present general discussion of *Xs*, she might choose to say that *some X are Y* when in fact many of them are, or *many X are Y* when in fact most of them are; it simply wouldn't matter. As a consequence, we might expect readers to be less inclined to draw scalar implicatures from a quantifier if they recognised it as property-introducing.

6.2.2 Choosing between *not all* and *some...not*

An interesting consequence of the distinction between property-introducing and presupposition-denying quantifiers relates to the question of how to correct a misconception. We know that *some Xs are not Y* and *not all Xs are Y* are truth-conditionally the same, for instance: both forms can, according to the present model, be used to deny the rule that $X \rightarrow Y$. And yet, to express a misconception correction, the latter expression seems far preferable. Imagine that the reader has seen several art-deco jewels with faceted stones, but is now looking at one with cabochon stones. Example 20 seems by far preferable to Example 19:

(19) Some art-deco jewels have cabochon stones. This jewel is a case in point.

(20) Not all art-deco jewels have faceted stones. This jewel has cabochon stones.

The problem with Example 19 is simply that the quantified sentence is ambiguous between the ‘introducing interesting properties’ sense of *some* and the *denying general rule* sense. Naturally, the sentence still does deny the rule. But this fact is very implicit. In Example 20 there is no such ambiguity, and the intended effect of the sentence is much clearer.⁶

7 Conclusions

ILEX-1 is similar to many NLG systems in having to generate texts which satisfy both content and formal goals. Somewhat to our surprise, we have found that the representation of defeasible rules is essential for all of the following abilities:

- Achieving the educational goals of the system.
- Achieving text coherence through the use of varied and interesting coherence relations.
- Addressing the misconceptions that the user may have in advance or may develop during the session.

The assumption of a knowledge representation scheme based around defeasible rules is becoming vital to our research because this seems to be the most promising, if not the only, basis for formalising the appropriate use of certain coherence relations and quantifiers in descriptive text. The role of defeasible rules in ILEX-1 is, however, not a traditional theorem-proving one. Rather the rules serve as a way of binding together the isolated facts of the domain and allowing the user and the system to discuss their conceptualisations of the domain.

References

- [1] N Asher and M Morreau. Commonsense entailment: A modal theory of nonmonotonic reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia, 1991.
- [2] G Carlson. The semantic composition of english generic sentences. In G Chierchia, B Partee, and R Turner, editors, *Properties, Types and Meaning. Vol 2: Semantic Issues*, pages 167–191. Kluwer, Dordrecht, 1989.
- [3] G Carlson and F Pelletier. *The Generic Book*. University of Chicago Press, Chicago and London, 1995.
- [4] L Cavedon and S Glasbey. The role of context in the interpretation of generic sentences. In *Proceedings of the 10th Amsterdam Colloquium*, Amsterdam, 1996.
- [5] H P Grice. Logic and conversation. In P Cole and J L Morgan, editors, *Syntax and Semantics*, volume 3. Academic Press, New York, 1975.
- [6] B Grote, N Lenke, and M Stede. Ma(r)king concessions in English and German. In *Proceedings of the Fifth European Workshop on Natural Language Generation*, Leiden, the Netherlands, 1995.
- [7] J Hirschberg. *A theory of scalar implicature*. Garland Publishing, New York, 1991.
- [8] M Krifka. Focus and the interpretation of generic sentences. In G Carlson and F Pelletier, editors, *The Generic Book*, pages 239–264. University of Chicago Press, Chicago and London, 1995.

⁶Note that another possible way of making the misconception correction clear would be in the form: *Some art-deco jewels do not have faceted stones. For instance, this jewel has cabochon stones*. But even this is not completely clear—maybe because some ‘interesting properties’ happen to be negative.

- [9] V Lifschitz. On the satisfiability of circumscription. *Artificial Intelligence*, 28:17–27, 1986.
- [10] W C Mann and S A Thompson. Rhetorical structure theory: A theory of text organization. *Text*, 8(3):243–281, 1988. Also available as Tech Report RR-87-190, USC Information Sciences Institute, Marina del Rey, CA.
- [11] E Oversteegen. Causal and contrastive connectives. Technical Report R1.1.3b, The DANDELION Consortium, CLS, Tilburg/Nijmegen, 1995.