

An Empirical Investigation of Definiteness

Massimo Poesio

University of Essex, Computer Science and Cognitive Science

1 Motivation

The debate between those who believe that uniqueness is the defining property of definite noun phrases (Russell, 1905; Loebner, 1987; Kadmon, 1987; Abbott, 1999; Roberts, 2001) and those who instead argue that the determining factor is familiarity (Christophersen, 1939; Heim, 1982; Prince, 1981, 1992) is still ongoing (see also (Birner and Ward, 1994; Abbott, 2000)). We report the results of a study in which this question was investigated by means of corpus analysis. Perhaps the most novel feature of this work is that for our study we used a large number of annotators together with statistical tests of reliability to ensure that our results could be replicated, whereas most previous studies were carried out by a single person on the basis of her own intuitions.

2 Background

2.1 The Poesio and Vieira analysis of definite description use

The starting point of this work were the results of Poesio and Vieira (1998), who found that of the 1,400 definite descriptions in their corpus, only about 50% were subsequent mention or bridging references, whereas 50% were first mentions. Of the first mentions, about half (i.e., 25% of the total) were what Hawkins (1978) would call 'larger situation' definites, i.e., definite descriptions like *the pope* whose referent is supposed to be part of shared knowledge; whereas the other half includes what Loebner (1987) calls SEMANTICALLY FUNCTIONAL definites, i.e., definites like *the first man on the Moon*.

2.2 Loebner's theory of definiteness

Poesio and Vieira noted that the theory proposed by Loebner (1987) appears to be the one that accounts for the largest percentage of their data, but did not test this hy-

pothesis. Loebner’s main claim is that the paradigmatic case of definiteness are not anaphoric NPs, as suggested by familiarity theories such as Heim’s (1982), but semantically functional ones such as *the first person ever to row across the Pacific on his own*. Anaphoric definite descriptions are viewed by Loebner as a special case; according to him, a definite description with a sortal predicate, such as *the dog*, can nevertheless be licensed in a context if the predicate can be coerced to a function in that context.

3 Methods

3.1 The Data: the GNOME corpus

In order to test Loebner’s theory and compare it with one based on familiarity, we analyzed about 3,000 NPs from the GNOME corpus (Poesio, 2000). These 3000 NPs we annotated include 1,700 definites (proper names, ‘the’-nps, ‘this’-nps, ‘that’-nps, pronouns, and possessive NPs), 1,050 indefinites (‘a’-nps, bare NPs, and numerical NPs such as ‘three cars’) and 300 ‘other’ NPs (including quantified NPs, wh-NPs, and gerunds such as ‘smoking’ in ‘smoking is dangerous’).

3.2 Annotation

Naive (but trained) annotators were asked to mark both the ‘familiarity’ and the ‘functionality’ status of each NP according to our instructions. In order to understand the results and the limitations of the present work, it is crucial to understand that for both properties, in order to ensure replicability of results, we had to restrict the range of semantic judgments we asked our subjects to make.

Specifically, our scheme for marking the familiarity status of a discourse entity was affected by the results of Fraurud (1990) and Poesio and Vieira (1998), that showed that some of the distinctions proposed, e.g., in (Prince, 1981, 1992) could not be marked reliably. In particular, subjects generally can’t agree on whether a discourse entity is hearer-new or hearer-old, so that they tend not to agree on whether ‘larger situation’ references such as *the pope* are part of shared knowledge or not. As a result, we only considered as familiar those NPs which either mentioned an entity already introduced (i.e., they are familiar in the narrower sense of Heim, or discourse-old in the sense of Prince) or were semantically related to it (i.e., bridging references). Poesio and Vieira (1998) also found that it is virtually impossible to get subjects to agree on what counts as a ‘bridging reference’ if the notion is intended in the broad sense used by Clark (1977). As a result, in our study we restricted the range of bridging references we asked our subjects to mark to those which we could get our annotators to agree on, including set relations (*the family ... the children*) and generalized **part-of** relations (as in *the car ... the wheel*).

The annotation of the GNOME corpus followed a systematic manual, available from the GNOME project's home page at <http://www.hcrc.ed.ac.uk/~gnome>; here, we only briefly discuss anaphoric and functionality annotation.

3.3 Anaphoric information

Each NP in the GNOME corpus is marked with a $\langle ne \rangle$ tag and with a variety of attributes capturing syntactic and semantic properties. (Example attributes are *cat* (specifying the type of an NP), *gf* specifying its grammatical function, *deix* (whether the object is a visual deictic reference or not) and *generic* (whether the NP denotes generically or not).) A separate $\langle ante \rangle$ element is used to mark anaphoric relations; the $\langle ante \rangle$ element itself specifies the index of the anaphoric expression and the type of semantic relation (e.g., identity), whereas one or more embedded $\langle anchor \rangle$ elements indicate possible antecedents (the presence of more than one $\langle anchor \rangle$ element indicates that the anaphoric expression is ambiguous). We were able to annotate anaphoric references reliably after restricting the range of bridging references.

3.4 Annotating functionality

In addition to familiarity, our subjects also marked the NPs in our corpus according to their functionality status, using a scheme derived from Loebner's proposal. Our subjects were asked to classify NPs as semantically functional, discourse functional (i.e., functional in a given context), directly referring (e.g., proper names), relational, or sortal. We achieved good reliability on this ($\kappa = .82$).

4 Results

We then tested the correlation between definiteness of an NP and either its familiarity or functionality by means of a variety of statistical correlation tests, which tell us how well a given feature (in our case, DEFiniteness) is predicted by another one (e.g., FAMilarity or FUNCtionality).

Of the 1,700 definite NPs, about 900 were anaphoric or bridging references, and about 800 were unrelated to entities previously introduced by one of the semantic relations we were able to identify reliably. Of the indefinite NPs, 900 were totally unrelated, and 150 had some (quasi-) anaphoric relation with an existing discourse entity. This resulted in a significant correlation between familiarity and definiteness by the χ^2 test, the Φ test, and the Λ test, but the correlation wasn't very strong: e.g., the Λ test indicates that using information about familiarity can make the prediction of definiteness 14% more precise than it would be predicted by chance. By contrast, when looking at the correlation between functionality and definiteness, we found that only 60 definite

NPs had a sortal predicate, and 120 were relational; the rest were all functional. Among indefinite NPs, 930 were sortal, 20 relational, and 60 functional. In other words, the correlation between functionality and definiteness was found to be much stronger than the correlation between familiarity and definiteness: e.g., the results of the Λ test suggest that information about functionality can make the prediction about definiteness 77% more precise than expected by chance - about 5 times the error reduction obtained by using familiarity.

5 Discussion

It is not our intention to interpret these results as conclusively settling the argument. For one thing, the stronger correlation between functionality and definiteness was expected, given that Loebner's scheme incorporates information about familiarity by viewing familiar discourse entities as one type of functional entity. And supporters of familiarity theories could always argue that the notion of familiarity we were able to annotate is too narrow. Nevertheless, our results do indicate that at the present stage the predictions of theories based on functionality are easier to test than those of theories based on familiarity (because some types of familiarity information are difficult to identify), and support Poesio and Vieira's hypothesis that at the present state of knowledge such theories cover a wider range of definites than 'narrow' theories of familiarity. More in general, these results raise the question of the extent to which understanders share the same kind of semantic information.

In the talk we will go in more detail over issues we had to address in the annotation and the types of definite descriptions that are problematic for the two types of theories. As already noted, e.g., in (Poesio and Vieira, 1998; Abbott, 1999), the most difficult cases for familiarity theories are semantically functional discourse entities. Then there are annotation problems. One important reason why functionality theories fare better is that it is very hard to get annotators to agree on whether an NP refers to an entity which is part of shared knowledge but has not been mentioned before, whereas it's easier (although not trivial) to get them to agree on whether such NPs refer to a unique object or not. Likewise, it is very hard to get annotators to agree on whether a definite NP such as *the murderer* counts as a bridging reference, but it's easier to get them to agree that the head noun of that NP denotes a function.

The most important class of definites not covered by Loebner's proposal are generic definites such as *the tiger* in *the tiger is a fierce animal*; however, we hypothesize that these can be brought under the coverage of his theory by adopting Dayal's (1999) suggestion that these NPs are basically proper names of kinds. Loebner also has no real account of cases such as *the student of a linguist* or *the bank of a river* (Poesio, 1994; Abbott, 1999), that we will also discuss.

References

- Abbott, B. (1999). Support for a unique theory of definite descriptions. In T. Matthews and D. Strolovitch, eds., *Proc. of SALT IX*, pp. 1–15.
- Abbott, B. (2000). Definiteness and indefiniteness. In L. R. Horn and G. Ward, eds., *Handbook of Pragmatics*. Blackwell.
- Birner, B. and G. Ward (1994). Uniqueness, familiarity, and the definite article in English. In *Proc. of the Berkeley Linguistic Society*, pp. 93–102.
- Christophersen, P. (1939). *The Articles: A Study of Their Theory and Use in English*. Oxford University Press.
- Clark, H. H. (1977). Bridging. In P. N. Johnson-Laird and P. Wason, eds., *Thinking: Readings in Cognitive Science*. Cambridge University Press, London and New York.
- Fraurud, K. (1990). Definiteness and the processing of NPs in natural discourse. *Journal of Semantics*, 7:395–433.
- Hawkins, J. A. (1978). *Definiteness and Indefiniteness*. Croom Helm, London.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts at Amherst.
- Kadmon, N. (1987). *On Unique and Non-Unique Reference and Asymmetric Quantification*. Ph.D. thesis, University of Massachusetts at Amherst.
- Loebner, S. (1987). Definites. *Journal of Semantics*, 4:279–326.
- Poesio, M. (1994). Weak Definites. In *Proc. of SALT*, pp. 282–299.
- Poesio, M. (2000). Annotating a corpus to develop and evaluate discourse entity realization algorithms. In *Proc. of the 2nd LREC*, pp. 211–218. Athens.
- Poesio, M. and R. Vieira (1998). A corpus-based investigation of definite description use. *Computational Linguistics*, 24(2):183–216.
- Prince, E. F. (1981). Toward a taxonomy of given-new information. In P. Cole, ed., *Radical Pragmatics*, pp. 223–256. Academic Press, New York.
- Prince, E. F. (1992). The ZPG letter: subjects, definiteness, and information status. In S. Thompson and W. Mann, eds., *Discourse description: diverse analyses of a fund-raising text*, pp. 295–325. John Benjamins.
- Roberts, C. (2001). Uniqueness presuppositions in English definite Noun Phrases. Available from the Semantics Archive.
- Russell, B. (1905). On denoting. *Mind*, 14:479–493. Reprinted in *Logic and Knowledge*, ed. R. C. Marsh. London: George Allen and Unwin.