

# Formalising and specifying underquantification

Aurelie Herbelot  
University of Cambridge  
ah433@cam.ac.uk

Ann Copestake  
University of Cambridge  
aac@cl.cam.ac.uk

## Abstract

This paper argues that *all* subject noun phrases can be given a quantified formalisation in terms of the intersection between their denotation set and the denotation set of their verbal predicate. The majority of subject noun phrases, however, are only implicitly quantified and the task of retrieving the most plausible quantifier for a given NP is non-trivial. We propose a formalisation which captures the underspecification of the quantifier in subject NPs and we show that this formalisation is widely applicable, including in statements involving kinds. We then present a baseline for a quantification resolution system using syntactic features as basis for classification. Although the syntactic baseline provides a respectable 78% precision, our error analysis shows that obtaining true performance on the task requires information beyond syntax.

## 1 Quantification resolution

Most subject noun phrases in English are not explicitly quantified. Still, humans are able to give them quantificational interpretations in context:

1. Cats are mammals = *All* cats...
2. Cats have four legs = *Most* cats...
3. Cats were sleeping by the fire = *Some* cats...
4. The beans spilt out of the bag = *Most/All of the* beans...
5. Water was dripping through the ceiling = *Some* water...

We refer to this process as **quantification resolution**, that is, the process of giving an implicitly quantified NP a formalisation which expresses a *unique* set relation appropriate to the semantics of the utterance. For instance, the most plausible resolution of 1 can be expressed as:

6. All cats are mammals.

$|\phi \cap \psi| = |\phi|$  where  $\phi$  is the set of all cats and  $\psi$  the set of all mammals.

Resolving the quantification value of NPs is important for many NLP tasks, in particular for inference. We would like to be able to automatically perform the type of interpretations shown in 1 to 5. It will allow us to draw conclusions such as *If (all) cats are mammals and Tom is a cat, then Tom is a mammal* and *If (some) cats are in my garden, then (some) animals are in my garden*.<sup>1</sup>

The task of quantification resolution involves finding a semantic representation that goes beyond what is directly obtainable from a sentence's syntactic composition. We can write *the*( $x$ ,  $\text{cat}'(x)$ ,  $\text{sleep}'(x)$ ) as we would write *some*( $x$ ,  $\text{cat}'(x)$ ,  $\text{sleep}'(x)$ )<sup>2</sup>, but while the quantification semantics of *some* can be

<sup>1</sup>The type of entailment relying on word substitution is dependent on quantification: *(All) cats are mammals* doesn't imply that *(All) animals are mammals*.

<sup>2</sup>We use here a generalised quantifier notation where the first argument of the quantifier is the bound variable.

fully defined (given a singular NP, we are talking of one entity only), that of *the* cannot: in a singular NP introduced by *the*, the referent can either be a single entity or a plurality with various possible quantificational interpretations (cf *The cat is sleeping* vs *The cat is a mammal*).

This paper is an attempt to provide a formal semantics for implicitly quantified NPs which a) supports the type of inferences required by NLP, b) has good empirical coverage (beyond ‘standard’ linguistic examples), c) lends itself to evaluation by human annotation and d) can be derived automatically. We draw on work in formal linguistics, but by formulating the problem as quantification resolution, we obtain an account which is more tractable from an NLP perspective. We also present preliminary experiments that automate quantification resolution using a syntax-driven classifier.

## 2 Under(specified) quantification

The phenomenon of ambiguous quantification overlaps with **genericity**. Generic NPs have traditionally been described as referring to **kinds** (Krifka et al., 1995) and one of their most frequent syntactic expressions is the bare plural, although they occur in definite and indefinite singulars too, as well as bare singulars. There are many views on the semantics of generics (e.g. Carlson, 1995; Pelletier and Asher, 1997; Heyer, 1990; Leslie, 2008) but one of them is that they quantify (Cohen, 1996), although, puzzlingly enough, not always with the same quantifier:

7. Dogs are in my garden = *Some* dogs...
8. Frenchmen eat horsemeat = *Some/Relatively-many* Frenchmen... (For the *relatively many* reading, see Cohen, 2001.)
9. Cars have four wheels = *Most* cars...

This behaviour has so far prevented linguists from agreeing on a single formalisation for all generics. Note that relegating the various readings to a matter of pragmatics, formalising all bare plurals using an existential, is no solution as we are then unable to explain the semantic difference between, for instance, *Mosquitoes carry malaria* and *Some mosquitoes carry malaria*. The only accepted assumption is that an operator *GEN* exists, which acts as a silent quantifier over the restrictor (subject) and matrix (verbal predicate) of the generic statement.

In this paper, we take an approach which sidesteps some of the intractable problems associated with the literature on generics and which also extends to definite plurals, as discussed below. Instead of talking of ambiguous quantification, we will talk of **underspecified quantification**, or **underquantification**. By this, we mean that the bare plural, rather than exhibiting a silent, *GEN* quantifier, simply features a placeholder in the logical form which must be filled with the appropriate quantifier (e.g.,  $uq(x, \text{cat}'(x), \text{sleep}'(x))$ , where *uq* is the placeholder quantifier). This account caters for the facts that so-called generics can so easily be quantified via traditional quantifiers, that *GEN* is silent in all known languages, and it explains also why it is the bare form which has the highest productivity, and can denote a range of quantified entities, from existentials to universals. Using the underquantification hypothesis, we can paraphrase any generic of the form ‘X does Y’ as ‘there is a set of things X, *a certain number of which* do Y’ (note the partitive construction).

We now turn to definite plurals which have traditionally been thought to be outside of the genericity phenomenon and associated with universals (e.g., Lyons, 1999). Definite plurals do exhibit a range of quantificational behaviour and thus we argue that they should be studied as underquantified forms too. Consider the following, from Dowty (1987):

10. At the end of the press conference, the reporters asked the president questions.

Dowty remarks that it is not necessary that all reporters ask questions for the sentence to be true. In fact, it is only necessary that *some of them* did. Dowty says: “The question of how many members of the group referent of a definite NP must have the distributive property is in part lexically determined and in part determined by the context, and only rarely is every member required to have these properties.”

Following the existential reading, we can write:

11.  $some(x, reporter'(x), askQuestion'(x))$

The problem is that for Dowty, the NP refers to a ‘group’, i.e., to the reporters as a whole, and not to specific reporters. We don’t want to say ‘there is a small set of reporters, each of which asked a question’; we want to say ‘there is a large set of reporters – all those present at the press conference – and some of them asked a question’, i.e., we want to use a partitive construction. We follow Brogaard’s (2007) account of definite plurals as partitive constructions, where she examines the following:

12. The students asked questions.

Brogaard argues that, given  $X$ , the denotation of *the students*, a subset  $Y$  of  $X$  is selected via the quantifier *some* and that the verbal predicate applies (distributively) to  $Y$ . A similar account can be given of (10): there is a set of reporters, and a certain number of elements in that set (some reporters) asked questions — which is our desired reading. Note that all definite plurals can have this interpretation (e.g., possessives and demonstratives also).

We will next argue that the partitive construct observed in definite plurals can be generally applied to subject NPs and we will propose a single formalisation for all underquantified statements.

### 3 Formalisation

#### 3.1 Link’s notation (1983)

In what follows, we briefly define each item of notation used in this work, as taken from Link (1983). We illustrate the main points via examples over a closed world  $W$  containing three cats (Kitty, Sylvester and Bagpuss).

The background assumption for our formalisation is that, following Link, plurals can be represented as lattices. The star sign  $*$  generates all individual sums of members of the extension of predicate  $P$ . So if  $P$  is *cat*’, the extension of  $*P$  is a join-semilattice representing all possible sums of cats in the world under consideration. The join-semilattice of cats in world  $W$  is shown in Fig 1.




Figure 1: Join-semilattice of all cats in world  $W$

The sign  $\sigma$  is the sum operator.  $\sigma xPx$  represents the sum, or supremum, of all objects that are  $*P$ .  $\sigma^*xPx$  represents the **proper sum** of  $P$ s, that is, the supremum of all objects that are proper plural predicates of  $P$ . The sum includes (non-plural) individuals such as  $K$  or  $S$  while the proper sum doesn’t. In worlds where there is more than one object in the extension of  $*P$ ,  $\sigma xPx = \sigma^*xPx$ : e.g., in Fig 1, the sum of all cats is the same as the proper sum of all cats, i.e., the set  $\{K,S,B\}$ . (Compare this with a world where there is only one cat, say Kitty: then  $\sigma xPx = \{K\}$  while  $\sigma^*xPx = \emptyset$ ).

The product sign  $\amalg$  expresses an **individual-part** relation. The  $\cdot$  sign in combination with  $\amalg$  indicates **atomic part**. Following Chierchia (1998), we assume the same underlying lattice for both mass terms and count nouns, so we use the  $\amalg$  and  $\cdot$  operators for formalising quantification over mass entities.

#### 3.2 Collective and distributive predicates

Some predicates are collective: they refer to a group as a whole and not to its instances (13). Other predicates are always distributive (14):

13. Antelopes gather near water holes (\*Andy the antelope gathers near water holes.)
14. Three soldiers were asleep (Tom was asleep, Bill was asleep, Cornelia was asleep.)

Most verbal phrases, though, are ‘mixed predicates’ that accept both readings:

15. Three soldiers stole wine from the canteen.

(Tom, Bill and Cornelia went together to the canteen to steal wine or Tom, Bill and Cornelia each stole wine from the canteen.)

Collective predicates can be a source of confusion when trying to directly apply quantification to an ambiguously quantified NP:

16. (\*Some/Most/All) Americans elect a new president every five years.

Quantifying 16 seems initially impossible in shallow form: we cannot write  $all(x, american'(x), electPres'(x))$  as it seems to imply distributivity. However, we refer to the reporter example (10) and the latent partitive construct that we suggested existed in that (distributive) sentence. By similarity, we can say that there is a set  $X$  of Americans able to vote, and a subset  $Y$  of those — which in this case is selected by the quantifier  $all$  and is therefore equal to  $X$  — collectively elects the president.

### 3.3 Formalising the partitive construct

Following Link (1998) for the formalisation of collective and distributive predicates, we can write, for 10 and 16:

17.  $X = \sigma^*x \text{reporterAtPressConference}'(x) \wedge \exists Y[Y \sqsubseteq X \wedge \forall z[z \cdot \sqsubseteq Y \rightarrow \text{askques}'(z)]]$
18.  $X = \sigma^*x \text{votingAmerican}'(x) \wedge \exists Y[Y \sqsubseteq X \wedge \text{electPresident}'(Y)]^3$

For the collective case, we just apply the verbal predicate collectively.

We can then add the quantifier resolution. We assume a three-fold partitioning of the quantificational space, corresponding to the natural language quantifiers *some*, *most* and *all* (in addition to *one*, for the description of singular, unique entities). The corresponding set relations are:

19. if *some*( $\phi, \psi$ ) then  $0 < |\phi \cap \psi|$
20. if *most*( $\phi, \psi$ ) then  $|\phi - \psi| \leq |\phi \cap \psi|$
21. if *all*( $\phi, \psi$ ) then  $|\phi - \psi| = 0$

These set relations can be expressed in terms of the sets involved in the partitive construction: in 16, if  $X$  is the set of all Americans able to vote,  $Y$  the subset of  $X$  selected by the quantifier, and  $Z$  the set of all things that elect the president, then  $Y$  actually represents the intersection  $X \cap Z$ . We can thus write:

22.  $X = \sigma^*x \text{reporterAtPressConference}'(x) \wedge \exists Y[Y \sqsubseteq X \wedge \forall z[z \cdot \sqsubseteq Y \rightarrow \text{askques}'(z)] \wedge (0 < |Y|)]$
23.  $X = \sigma^*x \text{votingAmerican}'(x) \wedge \exists Y[Y \sqsubseteq X \wedge \text{electPresident}'(Y) \wedge (|X - Y| = 0)]$

The same principle applies to mass nouns. We show below a distributive example.

24. Water was dripping through the ceiling.

$$X = \sigma^*x \text{water}'(x) \wedge \exists Y[Y \sqsubseteq X \wedge \forall z[z \cdot \sqsubseteq Y \rightarrow \text{dripThroughCeiling}'(z)] \wedge (0 < |Y|)]$$

We thus write the underspecified quantifier as:

25.  $X = \sigma^*x P'(x) \wedge \exists Y[Y \sqsubseteq X \wedge Q(Y)] \wedge \text{quantConstraint}(X, Y)]$

where the  $\text{quantConstraint}$  ensures the correct cardinality of  $Y$  for various quantifiers and the predicate  $Q$  applies distributively or collectively depending on the semantics of the sentence.  $X$  and  $Y$  respectively denote the Nbar and NP referents in the quantified paraphrase of the statement.

<sup>3</sup>Note that in the two examples, we have restricted  $X$  to the relevant set of entities. We will not investigate here how this particular reference resolution takes place.

## 4 Kinds

In order to argue that our formalisation is applicable to all subject noun phrases, we must briefly come back to the case of generics which, in some linguistic accounts, are not seen as quantified (Carlson, 1977).<sup>4</sup> According to those accounts, the subject NP in sentences such as *The cat is a mammal* (the **kind**) can be regarded as an entity similar to proper nouns. The generic reading of the sentence then takes a straightforward subject/predicate formalisation of the type mammal'(cat'). The main argument in favour of such a representation is the existence of sentences where the verbal predicate seems to only be applicable to a species rather than to its instances:

26. The dodo is extinct.

Such cases, we claim, do not preclude quantification. We use the accounts of Chierchia (1998) and Krifka (2004), where a kind is defined as a function that returns the greatest element of the extension of the property relevant to that kind:  $Kind(X) = \sigma^*x X'(x)$ . This gives us the following for 26:

27.  $X = \sigma^*x dodo'(x) \wedge \exists Y[Y \sqcap X \wedge extinct'(Y) \wedge (|Y - X| = 0)]$

We stress however that we do not deny the validity of representations that involve a simple subject/predicate structure. It should be clear that the sentence *The cat is a mammal* has an interpretation where the species 'cat' is attributed the property of being a mammal. What we argue is simply that the meaning of the sentence also includes a quantificational aspect. We want, after all, to be able to make natural inferences about individual cats: if the cat is a mammal then Tom the cat is a mammal. We believe that both quantification and a subject/predicate formalisation are necessary to fully render the semantics of such sentences. We will also argue in Section 7 that for the purposes of computational linguistics, it is actually desirable to formalise the quantificational aspect separately, as part of the full semantics.

We should also note that the genericity phenomenon is usually seen as encompassing habitual constructions (Krifka et al., 1995). Our quantificational account of kinds will not necessarily be applicable to quantification of events and we do not wish to make any claims with regard to habituality in this paper. For completeness, we will however point out that, following Chierchia (1995) on indefinites, we see quantification adverbs as able to bind, and therefore quantify over individuals: according to this view, the most felicitous reading of *Mosquitoes sometimes carry malaria* is *Some mosquitoes carry malaria*, formalisable with 25.

## 5 Automatic quantification: first attempts

To our knowledge, no attempt at the automatic specification of quantification has been made before. In consequence, we start our investigation with the simplest possible type of machine learning algorithm, using as determining features the direct syntactic context of the statement to be quantified. The general idea of such a system is that grammatical information such as the number of a subject noun phrase and the tense of its verbal predicate may be statistically related to its classification.

### 5.1 Gold standard

We built a gold standard by re-using and expanding the quantification annotations we produced in Herbelot and Copestake (2010). This small corpus, which contains randomly extracted Wikipedia<sup>5</sup> sentences, provides 300 instances of triply annotated subject noun phrases. The categories used for annotation are the natural language quantifiers ONE, SOME, MOST, ALL and the label QUANT (for noun phrases of the type *some cats*, *most turtles* or *more than 37 unicorns* which, being explicitly quantified, do not enter our underquantification account and must be marked with a separate label). In order to convert the multiple

<sup>4</sup>A more comprehensive discussion can be found in Herbelot (2010).

<sup>5</sup><http://www.wikipedia.org/>

annotations to a gold standard, we used majority opinion when it was available and negotiation in cases of complete disagreement. There were only 14 cases where a majority opinion cannot be obtained.

The main issue with the resulting gold standard is its relatively small size. The 300 data points it provides are clearly insufficient for machine learning, but the annotation process is time-consuming and we do not have the resources to set up a large-scale annotation effort. As a trade-off, the first author of this paper annotated a further 300 noun phrases, thus doubling the size of the gold standard. As a precaution, we ran the classifier presented later in this section over the original gold standard and over the new annotations; no substantial difference in performance between the two runs was found.

Table 1 shows the class distribution of our five quantification labels over the 600 instances of the extended gold standard.

Class	Number of instances	Percentage of corpus
ONE	367	61%
SOME	53	9%
MOST	34	6%
ALL	102	17%
QUANT	44	7%

Table 1: Class distribution over 600 instances

We note, first, that the number of explicitly quantified noun phrases amounts to only 7% of the annotation set. This shows that the resolution of underquantification has potentially high value for NLP systems. Next, we remark that 61% of all instances simply denote a single entity, leaving 32% to underquantified plurals — 189 instances. This imbalance is problematic for the machine learning task that we set out to achieve. First, it means that the training data available for SOME, MOST and ALL annotations is comparably sparse. Secondly, it implies that the baseline for our future classifier is relatively high: assuming a most frequent class baseline, we must beat 61% precision.

## 5.2 Quantifying with syntax

Most of the remarks that can be found in the literature on the relation between syntax and quantification have been written with respect to the generic versus non-generic distinction. Although we have moved away from the terminology on genericity, the two following examples show the potential promises — and hurdles — of using syntax to induce quantification annotations.

- Noun phrases which act as subjects of simple past tense verbs are usually non-generic: *A cow says ‘moo’ / A cow said ‘moo’* (Gelman, 2004). However, the so-called ‘historic past’ is an exception to this rule: *The woolly mammoth roamed the earth many years ago.*
- The combination of a bare plural and present tense is a prototypical indication of genericity: *Tigers are massive* (Cimpian and Markman, 2008). But news headlines behave differently: *Cambridge students steal cow.*

We informally investigate the distribution of various grammatical constructions with respect to quantification, as obtained from our gold standard. Although some constructions give a clear majority to one or another label, that majority is not always overwhelming. For instance, consistently annotating bare plurals followed by a past tense as SOME would result in a precision of only 54%. It is therefore unclear how accurate a classifier based only on syntax can be. (Note that the quantification phenomenon is understood to be semantically complex and that syntax is only one of many features used in the annotation guidelines produced in Herbelot and Copestake, 2010.)

### 5.3 Features

We give the system article and number information for the noun phrase to be quantified, as well as the tense of the verbal predicate following it. In order to cater for proper nouns, we also indicate whether the head of the noun phrase is capitalised or not. Article, number and capitalisation information is similarly provided for the object of the verb. All features are automatically extracted from the Robust Minimal Recursion Semantics (RMRS, Copestake, 2004) representation of the sentence in which the noun phrase appears (obtained via a RASP parse, Briscoe et al., 2006). The following shows an example of a feature line for a particular noun phrase (the sentence in which the noun phrase appears is also given):

```
ORIGINAL: [His early blues influences] included artists such as Robert  
          Johnson, Bukka White, Skip James and Sleepy John Estes.  
FEATURES: past, possessive, plural, nocap, bare, plural, nocap
```

Note that articles belonging to the same class are labelled according to that class: all possessive articles, for instance, are simply marked as ‘possessive’. This is the same for demonstrative articles.

### 5.4 Experiments and results

The aim of this work is not only to produce an automatic quantification system, but also, if possible, to learn about the linguistic phenomena surrounding the underspecification of quantification. Because of this, we choose a tree-based classifier which has the advantage of letting us see the rules that are created by the system and thereby may allow us to make some linguistic observations with regard to the cooccurrence of certain quantification classes with certain grammatical constructions. We use an off-the-shelf implementation of the C4.5 classifier (Quinlan, 1993) included in the Weka data mining software.<sup>6</sup> We perform a 6-fold cross-validation on the gold standard and report class precision, recall and F-score.

Class	Precision	Recall	F-score
ONE	86% (362/422)	99% (362/367)	92%
SOME	60% (25/42)	47% (25/53)	53%
MOST	33% (2/6)	6% (2/34)	10%
ALL	53% (57/108)	56% (57/102)	54%
QUANT	100% (22/22)	50% (22/44)	67%

Table 2: Class precision and recall for the quantification task

The C4.5 classifier gives 78% overall precision to the quantification task. Table 2 shows per class results for the three tasks. The figures in brackets indicate the number of true positives for a particular class, followed by the total number of instances annotated by the system as instances of that class. The classifier performs extremely well with the ONE class, reaching 92% F-score. Already quantified noun phrases yield perfect precision and mediocre recall, as might be expected since we do not provide the system with a list of quantifiers. The system performs less well with the labels SOME, MOST and ALL.

In order to understand the distribution of errors, we perform a detailed analysis on the first fold of our data. Out of 100 instances, the classifier assigns 25 to an incorrect class. The majority of those errors (44%) are due to the fact that the classifier labels all singulars as ONE, missing out on generic interpretations and in particular on the plural reading of mass terms: out of 11 errors, 5 are linked to a bare singular). The next most frequent type of error, covering another 16% of incorrectly classified instances, comes from already quantified noun phrases being labelled as another class. These errors affect the recall of the QUANT class and the precision of the SOME, MOST and ALL labels in particular (most of those errors occur in plural noun phrases). The coarseness of the rules is again to blame for the remaining errors: looking at the decision tree produced by the classifier, we observe that all bare

<sup>6</sup><http://www.cs.waikato.ac.nz/ml/weka/>

plurals followed by a present tense, as well as all definite plurals, are labelled as *universals*, while all bare plurals followed by a past tense are labelled as *SOME*. This accounts for a further 7 errors. The last three incorrect assignments are due to a dubious capitalisation rule.

## 5.5 Correspondence with linguistics

We observe that most definite plurals (including demonstratives and possessives) are classified as either *MOST* or *ALL*. This fits the linguistic notion of a definite as being essentially universal (Lyons, 1999) but also misses out on the correct quantification of statements such as 10.

We note also that non-capitalised bare plurals followed by a present tense are similarly classed as *ALL*. This echoes the observation that the combination of bare plural and present is a typical manifestation of genericity (if one understands genericity as a quantification phenomenon close to universality). When followed by past or perfect tenses, an existential quantification with *SOME* is however preferred.

One of the puzzles opened by the classifier's decision trees is the use of the direct object feature to distinguish between *MOST* and *ALL* in the case of some definite plurals. Given Sentences 28 and 29, our classifier would label the first one as *ALL* and the second one as *MOST*.

28. *My cats like the armchair.* *ALL*

29. *My cats like the armchairs.* *MOST*

At first glance, the rule seems to be a mere statistical effect of our data. We will however remark that statements like 29 are reserved a special section in Link (1998), where they are introduced as 'relational plural sentences'. One of Link's claims is that those sentences warrant four collective/distributive combinations — as opposed to two only in the case where the object is an individual. So we can say in Sentence 29 that a collective of cats likes a collective of armchairs, or that this collective of cats likes each armchair individually, etc. This proliferation of interpretations makes uncertainties more likely with regard to who likes what, and to the quantification of the subject and object.

For now, we will simply conclude that, although a simple syntax-based classifier is able to classify certain constructs with high precision, other constructs are beyond its capabilities. Further, it is difficult to see how improvements can be made to the current classification without venturing outside of the grammatical context. For instance, it seems practically impossible to improve on the high-precision rule specifying that every singular noun phrase should be classified as *ONE*. Due to space constraints, we will not report any further experiments in this paper. However, preliminary investigations into the use of lexical similarity to resolve quantification ambiguity can be found in Herbelot (2010).

## 6 Previous work

The general framework of this proposal is an underspecification account close to that described in Pinkal (1996) or Egg (2010). Computational approaches to underspecified quantification have so far focused on the genericity phenomenon. Leaving aside the question of annotation, which is treated in Herbelot and Copestake (2010), research on genericity can be classified within two strands: theoretical research on defeasible reasoning and extraction of common sense knowledge. Attempts to model defeasible reasoning were made in the 1980s with, for instance, the developments of default logic (Reiter, 1980) and non-monotonic logic (McDermott and Doyle, 1982). With information extraction as aim, Suh et al. (2006) attempt to retrieve 'common sense' statements from Wikipedia. They posit that common sense is contained in generic sentences. Their system, however, makes simplifying assumptions with regard to syntax: in particular, all bare plurals (and bare plurals only) are considered generic. In general, common sense extraction systems tend to restrict the data they mine to avoid the problem of identifying genericity (e.g., Voelker et al., 2007).



## 7 Conclusion, with some remarks on semantics

We have shown in this paper that subject noun phrases that are not explicitly quantified could be represented in an underspecified form. We have also argued that this formalisation is applicable to all constructs, including so-called generics. We have introduced a syntax-based classifier for quantification resolution and discussed the limits of an approach relying on compositional information only.

We acknowledge that our quantificational account of noun phrases, and especially of generics, does not satisfy the common requirement that a formalisation be a full description of the semantic particularities of a linguistic phenomenon. We think, however, that this requirement has led to over-restrictive approaches. One of the debates surrounding generics, for instance, relates to whether they should be given a ‘rules and regulations’ or an inductivist truth condition (Carlson, 1995). Our view is that it would be a mistake to exclude either interpretation. Burton-Roberts’ (1977) *A gentleman opens doors for ladies* clearly has normative force and without doubt, also allows the hearer to make their own conclusions with regard to the intersection between the set of all gentlemen and the set of people opening doors for ladies.

Our view of semantics is that it is a layered system and that specifying the quantification semantics of a noun phrase does not mean providing the full semantics of that noun phrase. It may be argued that the ideal semantics of generics should be unified and integrate all possible aspects of meaning. But such a theory is yet to be developed for genericity and, from a computational point of view, may not even be desirable: a modular representation of meaning allows us to only formalise the aspects that we are interested in for a particular task, leaving the rest out.

The approach presented here can be said to implement the idea of ‘slacker’ semantics (Copestake, 2009) in that a) our experiments try to derive a specification from compositional information only and b) we only attempt to specify one aspect of the meaning of noun phrases (quantification), leaving other aspects unspecified. In the future, we would like to take away some of the slack in a) by using lexical semantics in the specification of quantification. In order to do this, a much larger corpus should be created for the training and testing of the system, and this will be our next task.

## References

- Briscoe, T., J. Carroll, and R. Watson (2006). The second release of the RASP system. In *Proceedings of the COLING/ACL on Interactive presentation sessions*, Morristown, NJ, USA, pp. 77–80.
- Brogaard, B. (2007). The But Not All: A Partitive Account of Plural Definite Descriptions. *Mind & Language* 22(4), 402–426.
- Burton-Roberts, N. (1977). Generic sentences and analyticity. *Studies in Language* 1, 155–196.
- Carlson, G. N. (1977). *Reference to Kinds in English*. Ph. D. thesis, University of Massachusetts at Amherst.
- Carlson, G. N. (1995). Truth-Conditions of Generic Sentences: Two Contrasting Views. In G. N. Carlson and F. J. Pelletier (Eds.), *The Generic Book*, pp. 224–237. Chicago: University of Chicago Press.
- Chierchia, G. (1995). Individual-level predicates as inherent generics. In G. N. Carlson and F. J. Pelletier (Eds.), *The Generic Book*, pp. 176–223. Chicago: University of Chicago Press.
- Chierchia, G. (1998). Reference to kinds across languages. *Natural Language Semantics* 6, 339–405.
- Cimpian, A. and E. M. Markman (2008). Preschool children’s use of cues to generic meaning. *Cognition* 107(1), 19–53.
- Cohen, A. (1996). *Think Generic: The Meaning and Use of Generic Sentences*. Ph. D. thesis, Carnegie-Mellon University at Pittsburgh.

- Copestake, A. (2004). Robust Minimal Recursion Semantics. <http://www.cl.cam.ac.uk/~aac10/papers/rmrsdraft.pdf>.
- Copestake, A. (2009). Slacker semantics : why superficiality , dependency and avoidance of commitment can be the right way to go. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, Athens, Greece, pp. 1–9.
- Dowty, D. (1987). Collective predicates, distributive predicates and all. In F. Marshall, A. Miller, and Z.-s. Zhang (Eds.), *The Third Eastern States Conference on Linguistics*, Columbus, pp. 97–115. The Ohio State University, Department of Linguistics.
- Egg, M. (2010). Semantic Underspecification. *Language and Linguistics Compass* 4(3), 166–181.
- Gelman, S. A. (2004). Learning words for kinds: Generic noun phrases in acquisition. In D. Hall and S. Waxman (Eds.), *Weaving a lexicon*. Cambridge, MA: MIT Press.
- Herbelot, A. (2010). *Underspecified quantification*. Ph. D. thesis, University of Cambridge.
- Herbelot, A. and A. Copestake (2010). Annotating underquantification. In *Proceedings of the Fourth Linguistic Annotation Workshop*, Uppsala, Sweden, pp. 73–81.
- Heyer, G. (1990). Semantics and Knowledge Representation in the Analysis of Generic Descriptions. *Journal of Semantics* 7(1), 93–110.
- Krifka, M. (2004). Bare NPs: Kind-referring, Indefinites, Both, or Neither? In O. Bonami and P. Cabredo Hofherr (Eds.), *Empirical Issues in Formal Syntax and Semantics*, pp. 111–132.
- Krifka, M., F. J. Pelletier, G. N. Carlson, A. ter Meulen, G. Chierchia, and G. Link (1995). Genericity: An Introduction. In G. N. Carlson and F. J. Pelletier (Eds.), *The Generic Book*, pp. 1–125. Chicago: Chicago University Press.
- Leslie, S.-J. (2008). Generics: Cognition and Acquisition. *Philosophical Review* 117(1), 1–47.
- Link, G. (1983). The Logical Analysis of Plurals and Mass Terms: a Lattice-Theoretical Approach. In R. Bauerle, C. Schwarze, and A. von Stechow (Eds.), *Meaning, Use, and Interpretation of Language*, pp. 302–323. Berlin: de Gruyter.
- Link, G. (1998). Plural. In *Algebraic Semantics in Language and Philosophy*. Stanford: CSLI Publications.
- Lyons, C. (1999). *Definiteness*. Cambridge: Cambridge University Press.
- McDermott, D. and J. Doyle (1982). Non-monotonic Logic I. *Artificial Intelligence* 13, 41–72.
- Pelletier, F. J. and N. Asher (1997). Generics and Defaults. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 1125–1177. Amsterdam: Elsevier.
- Pinkal, M. (1996). Radical Underspecification. In P. Dekker and M. Stokhof (Eds.), *Proceedings of the 10th Amsterdam Colloquium*, Amsterdam, pp. 479–498. de Gruyter.
- Quinlan, J. (1993). *Programs for Machine Learning*. San Francisco, CA: Morgan Kaufmann.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence* 13(1-2), 81–132.
- Suh, S., H. Halpin, and E. Klein (2006). Extracting Common Sense Knowledge from Wikipedia. In *Proceedings of the International Semantic Web Conference (ISWC-06). Workshop on Web Content Mining with Human Language Technology*, Athens, GA.
- Voelker, J., P. Hitzler, and P. Cimiano (2007). Acquisition of OWL DL Axioms from Lexical Resources. In *Proceedings of the Fourth European conference on The Semantic Web: Research and Applications*, Innsbruck, Austria, pp. 670–685. Springer Verlag.