COMPUTATIONAL SEMANTICS OF MASS TERMS

Jan Tore Lønning
Department of Mathematics
University of Oslo
P.O. Box 1053, Blindern
0316 Oslo 3, Norway

ABSTRACT

Although the formalisms normally used for describing the semantics of natural languages are far from computationally tractable, it is possible to isolate particular semantic phenomena and interpret them within simpler formal systems. Quantified mass noun phrases is one such part. We describe a simple formal system suitable for the interpretation of quantified mass noun phrases. The main issue of this paper is to develop an algorithm for deciding the validity of sentences in the formal system and hence for deciding the validity of natural language inferences where all the involved noun phrases are quantified mass noun phrases. The decision procedure is based on a tableau calculus.

INTRODUCTION

A formal semantics for a part of a natural language attempts to describe the truth conditions for sentences, or propositions expressed by sentences, in model theoretic terms, and thereby also the relation of valid inferences between sentences. From the point of view of computational linguistics and natural language understanding, it is important whether this relation of entailment can be made computational. In general, the question must be answered in the negative. All proposed formal semantics of, say, English are at least as complex as first order logic and hence at best semi-decidable, which means that if a sentence β is a logical consequence of a set of sentences Σ , then there exists a proof for β from Σ , but no effective way to find such a proof. Several proposals use even more complex logics, like the higher order intensional logic used in Montague grammar, which has no deductive theory at all.

We will not oppose to the view that English incorporates at least the power of first order logic and that even more complex formalisms may be needed to represent the meaning of all aspects of English. But we believe there are two different possible strategies when one is to study one particular semantic phenomenon in natural languages. The first one is to try to interpret the particular phenomenon into a system that attempts to capture all semantic aspects of the natural language. The other strategy is to try to isolate the particular semantic phenomenon one wants to study and to build a semantic

interpretation suited for this particular phenomenon. By following the latter strategy it might be possible to find systems simpler than even first order logic that reflect interesting semantic phenomena, and in particular to come up with systems that are computationally tractable.

Quantified mass noun phrases is one such phenomenon that can be easily isolated. The properties particular for the semantics of quantified mass terms have been difficult to capture in extensions of systems already developed for count terms, like first order logic. However, if one isolate the mass terms and tries to interpret only them, it is possible to build a model where their typical properties fall out naturally. We have earlier proposed such a system and shown it to have a decidable logic (Lønning, 1987). We repeat the key points in the two following sections. The main point of this paper is a description of an algorithm for deciding validity of sentences and inferences involving quantified mass terms.

The strategy of isolating parts of a natural language and giving it a semantics that can be computational is of course the strategy underlying all computational uses of semantics. For example, in queries towards data bases one disregards all truly quantified sentences, and use only atomic sentences and pseudo quantified sentences where e.g. *for all* means for all individuals represented in the data base. The system we present here contains genuine quantifiers like *all water* and *much water*, but contain other restrictions compared to full first order logic. In particular the mass quantifiers behave simpler with respect to scope phenomena than count quantifiers.

REPRESENTING QUANTIFIED MASS NOUNS

We will make use of a very simple formal language for representing sentences containing quantified mass nouns, called LM (Lønning, 1987). We refer to the original paper for motivation of the particular chosen format and more examples and repeat only the key points here.

1. A particular LM language consists of a (non-empty) set of basic terms, say: *water*, *gold*,

blue, hot, boil, disappear ..., and a (possibly empty) set of non-logical determiners, say: *much, little, less_than_two_kilos_of...*

2. Common to all LM languages are the unary operator: $-$, the binary operator: $*$, the logical determiners: *all, some*, and the propositional connectives: $\neg, \wedge, \vee, \rightarrow$.
3. A term is either basic or of one of the two forms $(t*s)$ and $(-t)$ if t and s are terms.
4. An atomic formula has the form $D(t)(s)$ where D is a determiner and t and s are terms.
5. More complex formulas are built by the propositional connectives in the usual way.

A model for the particular language is a pair consisting of a Boolean algebra $A = \langle A, +, *, -, 0, 1 \rangle$ and an interpretation function $[]$, such that

1. $[t]$ is a member of A for all basic terms t ,
2. $[D]$ is a set of pairs of members of A for all determiners D .

The interpretation of more complex expressions is then defined as an extension of $[]$:

1. $[-t] = -[t]$, the Boolean complement of $[t]$, $[t*s] = [t]*[s]$, the Boolean product (or meet) of $[t]$ and $[s]$.
2. $[D(t)(s)]$ is true provided $([t], [s]) \in [D]$, in particular $[All(t)(s)]$ is true provided $[t] \leq [s]$, and $[Some(t)(s)]$ is true provided $[t]*[s] \neq 0$.
3. The propositional part is classical.

To get an intuition of how the semantics work one can think of $[water]$ as the totality of water in the world or in the more limited situation one considers, of $[blue]$ as the totality of stuff that is blue and of $[disappear]$ as the totality of stuff that disappeared. However, one shall not take this picture too literally since the model works as well for abstract as for concrete mass nouns.

In the formalism, a sentence like (1a) is represented as (1b) and (2a) is represented as (2b) if the negation is read with narrow scope and as (2c) if the negation is read with wide scope.

- (1) (a) All hot water is water.
(b) $All(hot*water)(water)$
- (2) (a) Much of the water that disappeared was not polluted.
(b) $Much(water*disappeared)(-polluted)$
(c) $\neg Much(water*disappeared)(polluted)$

Formula (1b) is a valid LM formula. In general, the valid English inferences that become valid if a mass term like *water* is interpreted as *quantities of water*

and *all water* is read as *all quantities of water* are also valid under the LM interpretation.

In addition, this approach can explain several phenomena that are not easily explained on other approaches. Roeper (1983) pointed out that paraphrasing *water* as *quantities of water* was problematic when propositional connectives were considered. If some water disappeared and some did not disappear, there will be many quantities that partly disappeared and partly did not disappear. If *disappear* denotes the set of all quantities that wholly disappeared and *did not disappear* denotes the complement set of this set, then all quantities that partly disappeared will be members of the denotation of *did not disappear*. The sum of quantities that are members of the denotation of *did not disappear* will equal all the water there is. Roeper solved this problem by letting the quantities be members of a Boolean algebra and used a non-standard interpretation of the negation. In LM it naturally comes out by the Boolean complement as in (2b) and *water that did not disappear* is represented by $(water*(-disappear))$.

A main feature of the current proposal is that it introduces non-logical quantifiers co-occurring with mass nouns, like *much, little, most,...* in a straightforward way. A sentence like *much water was blue* does not say anything about the number of quantities of blue water, but says something about the totality of blue water, which is the way it is interpreted in LM.

It might seem a little like cheating that the system only introduces interpretations of mass quantifiers with minimal scope with respect to other quantifiers, that it is not possible to interpret one quantifier with scope over another quantifier. In particular, since this is a main reason to why the logic in the sequel becomes simple. However, it is characteristic for mass quantifiers that they get narrow scope.

- (3) (a) A boy ate many pizzas.
(b) A girl drank much water.

While it might be possible to get a reading of (3a) which involves more than one boy, i.e., one boy for each pizza, it is not possible to get a reading of (3b) involving more than one girl.

The only determiners that get a fixed, or logical interpretation in LM are: *all* and *some*. For the other determiners we can add various sorts of constraints on the interpretations that will make more of the inferences we make in English valid in the logic. For example, it has been claimed that all natural language determiners are *conservative* (or "live on" their first argument), i.e: $(a,b) \in [D]$ if and only if $(a,b*a) \in [D]$ (Barwise and Cooper, 1981).

Several determiners are *monotone*, either *increasing* or *decreasing*, in one or both arguments, e.g., *much* is monotone increasing in its

second argument: if $(a,b) \in [much]$ and $b \leq c$ then $(a,c) \in [much]$, and *less_than_two_kilos_of* is monotone decreasing in its second argument. Whether an inference like

Much water evaporated.
All that evaporated disappeared.
 $\therefore$ *Hence much water disappeared,*

becomes valid in the logic, will depend on whether the denotation of *much* is constrained to be monotone increasing in its second argument or not.

LOGICAL PROPERTIES

We will repeat shortly several of the properties of the logic LM shown in (Lønning, 1987) as a background for the decision algorithm for validity.

A Hilbert style axiomatization was given and it was shown that any set of LM sentences consistent in the logic has a model: the logic is complete and compact.

It was implicitly shown, but not stated, that any model for LM must be a Boolean algebra: let a model be any set A with one unary operation $[-]$, one binary operation $[*]$, and a binary relation $[All]$ then the model is a Boolean algebra with $[-]$ the Boolean complement, $[*]$ the Boolean product (meet) operation and $[All]$ the ordering $\leq$ on A .

It was also shown that the logic was complete and compact with respect to the smaller model class: the atomic Boolean algebras, i.e. any consistent set of sentences has a model which is an atomic algebra, and in fact a finite such.

From this, it was shown that LM with no non-logical determiners, let us call it LA, is equivalent to a subset of monadic first-order logic, hence it is decidable. It was also shown that the full LM is decidable. The argument is based on the fact that the number of possible models for a sentence is finite and decidable. This number grows rapidly, however. Already a sentence in LA with n different basic terms have $2^{(2^n)}$ different models, so the argument does not establish a good procedure for checking validity in practice. We will establish procedures that are better (in most cases) in the next section.

Several natural restrictions on determiners, like conservativity and monotonicity can be expressed completely in LM. It is not surprising that this can be done for a fixed LM language given the finiteness of its nature. The more important point is that the properties can be expressed in a uniform way independently of the atomic terms of the language, which in next section will give rise to uniform inference rules.

A DECISION PROCEDURE

We shall establish a procedure for deciding the validity of LM sentences. The procedure is a combination of a normal form procedure and a tableau procedure (see e.g. Smullyan, 1968).

We start with an LM formula ϕ for which we want to decide whether it is valid, $\models \phi$. This is equivalent to deciding whether $\neg\phi$ is satisfiable. One can think of the process as an attempt on building a model for $\neg\phi$. If we succeed, we have shown that ϕ is not valid and we have found an explicit counterexample; if we fail, we have shown that ϕ is valid. We assume all propositional connectives in ϕ to be basic: $\neg, \wedge, \vee$.

1. First we introduce a new unary quantifier *Null* such that *Null*(t) is a formula when t is a term. The meaning is that $[Null(t)]$ is true if and only if $[t]=0$. Then substitute

Null($t \cdot (-s)$) for *All*(t)(s)

Null($t \cdot s$) for $\neg$ *Some*(t)(s)

$\neg$ *Null*($t \cdot s$) for *Some*(t)(s)

This step is not necessary, but it gives a more convenient notation to work with and fewer instances to consider in the sequel. The substitutions correspond to the fact that *Some* and *Every* can be taken to be unary and the one can be defined from the other, as in the case with count nouns.

2. Then transform ϕ to conjunctive normal form, that is, a conjunction of disjunctions of literals, where a literal is an atomic or negated atomic formula.

3. Observe that (i) and (ii) are equivalent.

(i) $\models \psi_1 \wedge \psi_2 \wedge \dots \wedge \psi_n$

(ii) $\models \psi_1, \models \psi_2 \dots$ and $\models \psi_n$,

Hence we can proceed with each conjunct separately.

4. We shall now check $\models \psi_1 \vee \psi_2 \vee \dots \vee \psi_n$, where each ψ_i has the form *Null*(t), $\neg$ *Null*(t), *D*(t)(s) or $\neg$ *D*(t)(s). Observe that the following two formulas are equivalent, where $t+s$ is shorthand for $\neg((\neg t) \cdot (\neg s))$.

(i) $\neg$ *Null*(t) $\vee$ $\neg$ *Null*(s)

(ii) $\neg$ *Null*($t+s$)

(This corresponds to the equivalence between $\exists x \phi \vee \exists x \psi$ and $\exists x (\phi \vee \psi)$ in first order logic.) Hence contract all the literals of the form $\neg$ *Null*(t) to one.

5. We are then left with a formula of the form

$\neg$ *Null*(t) $\vee$ *Null*(s_1) $\vee \dots \vee$ *Null*(s_n) $\vee \psi_1 \vee \dots \vee \psi_m$,

where each ψ_i has the form *D*(u)(v) or $\neg$ *D*(u)(v), for some non-logical *D*. First assume that there are no ψ_i 's. Then observe that (i) and (ii) are equivalent.

- (i) $\vdash \neg \text{Null}(t) \vee \text{Null}(s_1) \vee \dots \vee \text{Null}(s_n)$
(ii) $\vdash \neg \text{Null}(t) \vee \text{Null}(s_1) \text{ or } \dots \text{ or } \vdash \neg \text{Null}(t) \vee \text{Null}(s_n).$

(If there are no s_i 's proceed with $\vdash \neg \text{Null}(t)$.) This equivalence might need some argument. That (ii) entails (i) is propositional logic. For the other way around, there are two possibilities. Either $\vdash \neg \text{Null}(t)$, which yields the equivalence. Otherwise, there exists a model **A** for the language of ϕ where $[t]=0$ and for all other terms s in the language: $[s]=0$ if and only if $\text{All}(s)(t)$ is a valid formula. Let ψ be the formula $\neg \text{Null}(t) \vee \text{Null}(s_1) \vee \dots \vee \text{Null}(s_n)$. Then ψ is valid if and only if ψ is true in **A**. To show this, it is sufficient to show that if there is a model **B** in which ψ is not true then ψ is not true in **A**. If **B** is a model in which ψ is not true, then $[t]=0$ and each $[s_i] \neq 0$ in **B**. Hence $\text{All}(s_i)(t)$ cannot be valid and $[s_i] \neq 0$ in **A** for each s_i . Since $[t]=0$ in **A**, ψ cannot be true in **A**. The same type of argument yields that $\neg \text{Null}(t) \vee \text{Null}(s_i)$ is valid if and only if it is true in **A**. If we write $\mathbf{A} \models \eta$ for η is true in **A**, the following equivalence is propositional logic and yields the equivalence above.

- (i) $\mathbf{A} \models \neg \text{Null}(t) \vee \text{Null}(s_1) \vee \dots \vee \text{Null}(s_n)$
(ii) $\mathbf{A} \models \neg \text{Null}(t) \vee \text{Null}(s_1) \text{ or } \dots \text{ or } \mathbf{A} \models \neg \text{Null}(t) \vee \text{Null}(s_n).$

6. a. We shall describe two different ways for checking $\vdash \neg \text{Null}(t) \vee \text{Null}(s_i)$. The first one proceeds by a transformation to normal form and may be the easiest one to understand if one is not accustomed to tableau calculus. The second one which uses a tableau approach is more efficient. First observe that (i) and (ii) are equivalent.

- (i) $\vdash \neg \text{Null}(t) \vee \text{Null}(s_i)$
(ii) $\vdash \text{Null}((-t) \cdot s_i)$, (i.e. $\vdash \text{All}(s_i)(t)$).

The last claim entails the first one since

$$\text{Null}((-t) \cdot s_i) \rightarrow \neg \text{Null}(t) \vee \text{Null}(s_i)$$

is a valid LM-formula. To see that (ii) entails (i) observe that if $\neg \text{Null}(t) \vee \text{Null}(s_i)$ is valid, it will in particular be true in the model **A** described in step 5, hence $\vdash \text{All}(s_i)(t)$.

To check $\vdash \text{Null}((-t) \cdot s_i)$ rewrite the term $(-t) \cdot s_i$ in disjunctive normal form: allow the symbol $+$ and write the term on the form $s_1 + \dots + s_m$ where each s_j has the form $u_1 \cdot \dots \cdot u_k$ and each u_j is either an atomic term or on the form $\neg v$ for an atomic term v . Then $\vdash \text{Null}(s_1 + \dots + s_m)$ if and only if $\vdash \text{Null}(s_1)$ and $\dots$ and $\vdash \text{Null}(s_m)$, and $\vdash \text{Null}(u_1 \cdot \dots \cdot u_k)$ if and only if there is a v such that one u_j equals v and another u_j equals $\neg v$.

b. The checking of $\vdash \neg \text{Null}(t) \vee \text{Null}(s_i)$ will be faster using a tableau procedure instead of rewriting to normal form. Note that the following are equivalent:

- $\vdash \neg \text{Null}(t) \vee \text{Null}(s_i)$
 $\vdash \text{Null}((-t) \cdot s_i)$
 $\vdash \neg \text{Null}(t + (-s_i))$

There is a close connection between propositional logic and Boolean algebras. To each term in LM, t , there corresponds a formula p_t in pure propositional logic such that $\neg \text{Null}(t)$ is valid in LM if and only if the corresponding formula p_t is a tautology: shift each basic LM term t with a corresponding propositional constant p_t , and exchange $\neg$ with $\neg$, $\cdot$ with $\wedge$, and $+$ with $\vee$. In particular, the following are equivalent:

- $\vdash \neg \text{Null}(t + (-s_i))$ (in LM)
 $\vdash (p_t \vee (\neg p_{s_i}))$ (in propositional logic)

(By the earlier mentioned connection between LM and first order logic this corresponds to the fact that a first order formula $\exists x \phi$ is valid if and only if ϕ is a tautology whenever ϕ is quantifier free.)

Step (6a) above is equivalent to checking this latter formula for validity by transformation to a normal form. Instead it can be checked by a standard tableau procedure (see e.g. Smullyan 1968).

We give a short description of the tableau approach to propositional logic. In order to verify a formula ψ , we try to build a model that falsifies it. To ease the description we assume that ψ is on negation normal form, that is, built up from literals by $\wedge, \vee$. The attempt to build a model is done by building a tree for ψ . We start with a root node depicted by ψ and use the following two rules:

1. For each node a depicted by a formula of the form $\gamma \vee \eta$ attach to each leaf below a in the tree constructed so far one successor node b depicted by γ and one successor node c to b depicted by η .
2. For each node a depicted by a formula of the form $\gamma \wedge \eta$ attach to each leaf below a in the tree constructed so far two new leaf nodes one depicted by γ and another one depicted by η .

The tree is complete when all formulas of the forms $\gamma \vee \eta$ and $\gamma \wedge \eta$ are reduced according to the rules above. A branch in a tree is called closed if there is a formula γ such that one node along the branch is depicted by γ and another node along the branch is depicted by $\neg \gamma$. A branch in a complete tree for ψ which is not closed describes a valuation that falsifies ψ . Conversely, if all branches in a complete tree for ψ are closed, ψ is valid. We illustrate with an example:

Just like there is a meaning preserving translation from LA into first order logic, LM can be translated into first order logic extended with generalized quantifiers. A proof procedure for first order logic, like a tableau or a sequent calculus, can be extended with rules for generalized quantifiers similar to the rules introduced here. If Q is a binary quantifier with no additional constraints on its interpretation then the following are equivalent.

- (i) $\vdash Qx(A(x), B(x)) \vee \neg Qx(C(x), D(x))$
- (ii) $\vdash \forall x(A(x) \leftrightarrow C(x)) \wedge \forall x(B(x) \leftrightarrow D(x))$

So to show that (i) is valid one has to show that (ii) is valid. This can be incorporated into a tableau or sequent calculus for first order logic. If the first order logic is monadic, as the logic we get after translating LM into first order logic is, one can use a similar procedure as the one described here. If the extended first order logic is not monadic, the procedure one gets when rules corresponding to the reduction from (i) to (ii) are included, becomes more complex.

EFFICIENCY

We chose to transform the formula being tested to normal form early in the procedure (step 2). Alternatively to the described algorithm one could think of using a tableau procedure all the way, and not first transform to conjunctive normal form in step 2. In general, transformation to normal form is slower than using a tableau procedure (cf. step 6 above). The reason we made the transformation to normal form was that this was necessary to split the formula in step 5 and step 7. In the procedure one gets by translating LM into first order logic (with generalized quantifiers) and using a tableau procedure from the start, it is not possible to split the tree similarly. If we for simplicity considers a formula with no generalized quantifiers, the pure tableau calculus will not lead to a separate tree for each s_i together with t but to one big tree containing all the s_i 's and roughly one copy of t for each s_i . This corresponds to the quantifier rules in a tableau calculus for first order logic: (i) for each formula of the form $\forall x\phi$ introduce one new formula $\phi(a)$ where a is some new term, (ii) for each formula $\exists x\psi$ introduce one new formula $\psi(a)$ for each term a introduced in the tree at a branch to which $\exists x\psi$ belongs. The successful separation in the described algorithm here will also be possible in a proof procedure for monadic first order logic.

The two different procedures will be of the same time complexity in worst cases. In the practical applications we have in mind, the procedure described here will be faster. Typically we want to check whether a formula β follows from $\alpha_1, \dots, \alpha_n$. This is the same as deciding whether $\neg\alpha_1 \vee \dots \vee \neg\alpha_n \vee \beta$ is

valid or not. The transformation to normal form will produce one additional copy for each $\vee$ within an α_i and each $\wedge$ within β . If each α_i and β are LM formulas that represent English sentences, they can each be expected to be relatively short and in particular not contain many $\vee$'s, so the number of copies made will be relatively small. On the other hand, the number of α_i 's may be large if they represent the dialogue so far or the agent's knowledge. It is therefore important that each disjunct can be split up as much as possible.

IMPLEMENTATION

The inference algorithm has been implemented in PROLOG. To test it out we have built a small (toy) natural language question-answering-system around it. The program reads in simple sentences and questions from the terminal and answers the questions. It can handle simple statements, like *If some of the hot coffee that did not disappear was black then much gold is valuable* (the fragment in Lønning 1987) and yes/no questions like *Did much water evaporate?* and *Was the old gold that disappeared valuable?* We have written the grammar and translation to LM in the built in DCG rules (Pereira and Warren, 1980).

Statements typed on the terminal are interpreted as facts about the world and stored as simple sentences $\phi_1, \dots, \phi_n$. When a question like *Did much water evaporate?* is asked, it is parsed and turned into a formula like $\psi: \text{Much}(\text{water})(\text{evaporate})$. Then the program proceeds to check the validity of $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi$. If it is valid, the program answers *yes*, otherwise it checks $(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \neg\psi$. If this is valid, the answer is *no*, otherwise the program answers that it *does not know*. When a statement is made the program checks whether it is consistent with what the program already knows before it is added to the knowledgebase.

The system is mainly made to test the inference algorithm and is not meant as an application by itself. But it illustrates some general points. It is a system where natural language inferences are made from natural language sentences and not from a fixed database. The system contains a complete treatment of propositional logic and illustrates a sound treatment of negation where failure is treated as *does not know* instead of negation. On the other hand, there is also a price to pay for incorporating full propositional logic. The system can only handle examples of a limited size in reasonable time.

CONCLUSION

We have here presented a computational approach to the semantics of quantified mass noun phrases. We think the semantics ascribed to quantified mass

nouns through a translation into LM is the one that most adequately reflects their particular semantic properties. In addition this semantics can be made computational in a way not possible for other approaches to the semantics of mass terms, like Bunt's (1985) which extends axiomatic set theory, or Link's approach (1983) based on Montague's higher-order intensional logic.

We have modified and adapted a tableau calculus to be used with mass terms and extended it with generalized quantifiers. Although the implementation we have made is of limited applicability, we hope that the algorithm can be used to incorporate quantified mass noun phrases into larger systems treating count terms. In particular, it should be possible to combine the algorithm with other approaches based on a tableau calculus, like the one described by Guenther, Lehmann and Schönfeld (1986).

REFERENCES

- Barwise, J. and R. Cooper: 1981, 'Generalized Quantifiers and Natural Language', *Linguistics and Philosophy*, 4, 159–219.
- Bunt, H.: 1985, *Mass terms and model-theoretic semantics*, Cambridge University Press, Cambridge.
- Guenther, F., H. Lehmann and W. Schönfeld: 1986, 'A theory for the representation of knowledge', *IBM J. Res. Develop.* 30, 39–56.
- Link, G.: 1983, 'The Logical Analysis of Plurals and Mass terms: A Lattice-Theoretical Approach', in Bäuerle et al. (eds.), *Meaning, Use, and Interpretation of Language*, Walter de Gruyter, Berlin.
- Lønning, J.T.: 1987, 'Mass Terms and Quantification', *Linguistics and Philosophy*, 10, 1–52.
- Pereira, F.C.N. and D.H.D. Warren: 1980, 'Definite Clause Grammars for Language Analysis — A Survey of the Formalism and a Comparison with Augmented Transition Networks', *Artificial Intelligence* 13, 231–278.
- Roeper, R.: 1983, 'Semantics for Mass terms with Quantifiers', *Nous*, 17, 251–265.
- Smullyan, R.M.: 1968, *First-Order Logic*, Springer, New York.