A Computational Model of Quantification in Natural Language

Proposal for MEng Thesis Project

Avril Kenney

12 May 2011

1 Introduction

In natural languages, there are words that express different quantities or amounts – for example, English has words such as "some" and "every". However, across languages there is a large amount of variation in which and how many such words exist, and there are many conceivable such words that are not known to exist in any languages. The purpose of this project is to examine the quantificational words that exist in natural languages, by making a computational model based on factors known to be relevant to other properties of language. Specifically, the meanings of words need to be representative of the real-world meanings that people need to communicate, learnable from linguistic input together with non-linguistic experience, and easily computable given human perceptual and reasoning abilities.

Quantificational words are a useful class of words to study because, unlike others such as nouns or adjectives, they do not refer to specific objects, actions, or properties in the real world; they are abstract linguistic constructs. The overall frequency of their use is high and is largely independent of non-linguistic circumstances. Furthermore, they are a "closed class" of words: each language has essentially an unchanging set of them, which makes them easier to enumerate and compare across languages.

By investigating patterns in the words expressing quantification in natural language, this project will provide insight into how constraints on human cognition and communication influence language.

2 Previous Work

Quantification has long been a topic of interest in linguistics (e.g., Barwise & Cooper, 1981; Keenan & Stavi, 1986); however, most of the research in this area has been focused on formally characterizing the syntactic and semantic details, rather than the relationship to communication. In English, "quantifiers" are often equated with "determiners" (which are, roughly, words that act as specifiers for nouns), but there is not currently a consensus on what words actually are "quantifiers." This project will therefore sidestep the issue of these formal categorizations and deal with "quantificational words," which will be defined loosely as single words used to describe relations between sets, where the relations have to do with quantity. The set of these words in English corresponds to what are sometimes called "set-relational quantifiers."

Languages vary widely in terms of how they express quantification. In many languages, quantificational words do not occupy the syntactic position of determiner as they (arguably) do in English. They are often expressed as adverbs or modifiers (similar to the English "dogs are always animals" as opposed to "all dogs are animals") (Matthewson, 1996; Bach, Jelinek, Kratzer, & Partee, 1995). Other languages use nouns with meanings like "part" or "small amount" (Everett, 2005). The point of this project is not to make universal claims or to characterize all of the constructions by which languages can express quantification; rather, the point is to see whether some of the patterns that do exist can be accounted for by a few general principles about language and communication.

One property that has been proposed to be true of all natural-language determiners is "conservativity" (Keenan & Stavi, 1986). A function f is defined to be conservative if for any sets S and T, $f(S,T) = f(S, S \cap T)$. For example, the determiner "every" is conservative ("every person is friendly" means the same thing as "every person is a person who is friendly"), whereas the adverb "only" is not conservative ("only people are friendly" does not mean the same thing as "only people are people who are friendly"). Some experiments have shown that people are better at learning conservative than non-conservative meanings for novel determiners (Graff, Romoli, Moro, & Snedeker, 2009; Hunter & Lidz, submitted), suggesting that cognition might in some way be better equipped to understand conservative meanings. Conservativity is just one example of a formal property that could be indicative of the types of meanings people can easily represent.

Recent work by Piantadosi, Tenenbaum, and Goodman (2010) has specifically investigated the types of "representation language" that people might be using to learn and reason about set-theoretic concepts. In general, people are better at learning simple concepts, i.e., the ones that can be described using fewer primitive symbols in the representation language. Certain types of concepts are also more easily learnable than others: for example, conjunctions (meanings that apply if all of their subparts apply, such as "blue circles") are easier to learn than disjunctions (meanings that apply if any of their subparts apply, such as "blue things or circles").

It is well established that certain quantificational words tend to be learned earlier than others by children acquiring language, and this order is related to the complexity of the meanings (Hanlon, 1978, 1981). In particular, Hanlon proposes three dimensions on which meanings can differ. A meaning can be generic or specific (with respect to the set that it quantifies over), or neither ("non-specific"), and in the order of acquisition, non-specific meanings precede specific which precede generic. A meaning can be collective or distributive (the difference between "all" and "each"), and collective meanings are acquired before distributive ones. And different meanings can have different presuppositional sets (for example, "both" presupposes that there are exactly two), and meanings for which the presuppositional set is the reference set precede those for which it is not (for example, "both" precedes "either").

Even among words with fairly straightforward literal meanings, there is a surprising amount of subtlety in how meanings are mentally computed, and phrases with the same meaning are sometimes computed by different methods. The word "most," which could be defined as "more than half," is probably actually represented as a superlative rather than a comparative – that is, a statement like "most people are friendly" would be computed as "more people are friendly than not" rather than "more than half of people are friendly" (Hackl, 2009). In general, phrases expressed as "at least" or "at most" require more processing than the corresponding "more than" or "fewer than" phrases that express the same meanings (Geurts, Katsos, Cummins, Moons, & Noordman, 2009).

There are also many meanings that people are capable of learning and computing but that do not exist as words. The obvious explanation for this fact would be that the existing words are already adequate for expressing the necessary meanings. In other categories of words, it has been shown that natural languages are indeed efficient at allocating a large space of meanings to a smaller set of words: cross-linguistics studies of words for colors (Regier, Kay, & Khetarpal, 2007) and for spatial relationships (Khetarpal, Majid, & Regier, 2009) found that the set of words in a given language tends to maximize informativeness. For quantification, another factor that needs to be taken into account is pragmatic inference: a word can convey more information than its literal meaning because certain inferences are licensed by the non-linguistic context. For example, in English the word "some" is often used to imply "some but not all," with the reasoning that if "all" were true then the speaker would have used that word instead.

All these results suggest that the quantificational words existing in natural language should be affected by specific features of human cognition.

3 Proposed Model

The model will be based on three factors: communicative efficiency, learnability, and computability. It will represent quantificational words as relationships between sets, possibly with additional information such as presuppositions, and for each of the three factors, it will be able to compute a score for a given word or set of words. Then it can be used to compute, given that a language has a certain number of quantificational words, what sets of words are predicted to be most likely to occur.

Making the model will involve several steps. First, it must be determined what the primitive building-blocks of meanings will be and how they can be combined to create allowable words. Then, methods must be formulated for how to evaluate words with respect to each of the three factors, and experiments must be run on people to confirm that the model is modeling what it is intended to. Finally, the factors must be put together to make predictions, and these predictions must be compared to existing languages.

3.1 Communicative Efficiency

The idea of communicative efficiency is that the set of words should cover the set of meanings that could need to be conveyed, in such a way as to maximize the probability that when a listener hears a word, he can correctly understand what the speaker was trying to communicate.

The case of quantification is more complicated than, for example, color, because it is not simply a matter of partitioning a space into segments and assigning words to them. The speaker might have varying degrees of certainty, or might want to specify his meaning to varying degrees of precision. Meanings can have significant overlap, and pragmatic implicatures can influence interpretations. Intended meanings should therefore be represented as distributions over possible reference sets.

Another important consideration is that not all meanings are equally likely to need to be communicated. Presumably there is some prior distribution over desired meanings, and it is not clear how this distribution could be uncovered, because an appoximation based on the distribution of words used would be biased by what words exist. To begin with, then, the model will use simple prior distributions, such as a uniform distribution over a finite number of proportions of a set.

To test the communicative efficiency part of the model, people can be taught specific sets of quantification-word meanings and be asked to use them to describe scenarios, and the communicative efficiency of a set of words can be calculated as how often other people are able to correctly identify what scenario is being described. If the model is accurate, it should produce values similar to the humans' results.

3.2 Learnability

The words that exist should be learnable from the type of input that a child learning a language would normally have available, and words that are easier to learn are predicted to exist more often.

The learnability of individual meanings depends on the complexity of their mental representations. As discussed in Section 2, this is related to what primitives the meaning is built from as well as how long of a combination it is.

Learnability is a property not only of individual meanings but also of sets of meanings. If the meanings of the words in a given language are very distinct and have little overlap among them, then it will probably be easier to learn that set of words than a set in which the individual words have equivalent individual learnability but are more similar to each other. It should also be more difficult to learn a large number of words than to learn just a few words, both because of memory constraints and because of the amount of available evidence.

To test the learnability part of the model, people can be shown example scenarios paired with nonsense-words representing the applicable quantificational meanings, and after being trained on these pairings, they can be asked whether each word applies to new scenarios. The more easily learnable a meaning is, the fewer examples people should need to see before they can correctly determine whether the word applies in a given situation.

3.3 Computability

In order for words to be used in normal conversation, people have to be able to compute their meanings, both to determine what word to use in a given situation and to determine the information trying to be communicated by a given word.

Computing whether a word applies to a scenario involves multiple steps. If the word has presuppositions, the person must check whether those presuppositions are met. Then it must be determined what sets are relevant to the meaning, and those sets must be checked to see whether the relation expressed by the meaning is true of them. Some properties, such as exact numbers greater than about three, cannot generally be computed without external mechanisms.

To test the computability part of the model, people can be taught quantificational meanings for nonsense-words, and then be shown scenarios and tested on how quickly they can determine whether a given meaning applies to each scenario. One complicating issue is that people will probably be faster on meanings that have words in their own language; to account for this, people could be tested only on meanings that they do not already have words for, and the computability of the remaining meanings could be interpolated from this data.

3.4 Full Model

Once the three subparts are established, they can be combined and used to assess the overall goodness of particular sets of quantificational words. Corpus data for various languages can be used to determine some of the sets of words that exist in real languages, and these can be compared to alternative sets that do not necessarily exist.

References

Bach, E., Jelinek, E., Kratzer, A., & Partee, B. H. (Eds.). (1995). *Quantification in natural languages*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Barwise, J., & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, *4*(2), 159-219.

Everett, D. L. (2005). Cultural constraints on grammar and cognition in Pirahã. *Current Anthropology*, *46*, 621-646.

Geurts, B., Katsos, N., Cummins, C., Moons, J., & Noordman, L. (2009). Scalar quantifiers: Logic, acquisition, and processing. *Language and Cognitive Processes*, *25*, 130-148.

Graff, P., Romoli, J., Moro, A., & Snedeker, J. (2009). *Exploring a learning bias against non-conservative determiners.* Presented at the 22nd Annual Meeting of the CUNY Conference on Human Sentence Processing.

Hackl, M. (2009). On the grammar and processing of quantifiers: *most* versus *more than half*. *Natural Language Semantics*, *17*, 63-98.

Hanlon, C. (1978). The emergence of set-relational quantifiers in early childhood. In P. S. Dale & D. Ingram (Eds.), *Child language – An international perspective* (p. 115-130). Baltimore: University Park Press.

Hanlon, C. (1981). Frequency of usage, semantic complexity, and the acquisition of set-relational quantifiers in early childhood. In C. E. Johnson & C. L. Thew (Eds.), *Proceedings of the second international congress for the study of child language* (p. 245-260). Washington, D.C.: University Press of America.

Hunter, T., & Lidz, J. (submitted). *Some (but not all) unattested determiners are unlearnable.*

Keenan, E. L., & Stavi, J. (1986). A semantic characterization of natural language determiners. *Linguistics and Philosophy*, *9*, 253-326.

Khetarpal, N., Majid, A., & Regier, T. (2009). Spatial terms reflect near-optimal spatial categories. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.

Matthewson, L. (1996). *Determiner systems and quantificational strategies: Evidence from Salish*. Unpublished doctoral dissertation, University of British Columbia.

Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2010). Beyond Boolean logic: exploring representation languages for learning complex concepts. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*.

Regier, T., Kay, P., & Khetarpal, N. (2007). Color naming reflects optimal partitions of color space. *Proceedings of the National Academy of Sciences*, *104*(4), 1436-1441.