

The constructive nature of natural language*

Allan Ramsay

School of Computer Science, University of Manchester

Allan.Ramsay@manchester.ac.uk

Ghadah Binhadba

School of Computer Science, University of Manchester

Ghadah.Binhadba@postgrad.manchester.ac.uk

Abstract


Natural language quantifiers are often described in truth conditional terms – ‘*A man is sleeping*’ is true if and only there is a man who is sleeping, ‘*Few men are immortal*’ is true if and only the proportion of men who will live forever is small, and so on. We will survey a number of examples that suggest that it is better to think of them as instructions to the hearer about how to update their model of the minutes of the conversation (and in the presentation of this work we will suggest how this should be done, but there is no space for that here).

1 Introduction

Consider (1):

- (1) A man is sleeping.

This contains two descriptive terms – ‘*man*’ and ‘*sleeping*’. Anyone reading this sentence will immediately form a picture in their minds of a situation where an entity that fits the first description is carrying out an activity that fits the second.



If you read (2), on the other hand, your immediate reaction is to say ‘*What man?*’ and ‘*When was this happening?*’.

- (2) The man had been sleeping.

*Presented at the ESSLLI 2017 workshop *QUAD: Quantifiers and Determiners* (Retoré, Steedman)

(1) and (2) contain the same descriptions, but different ‘instructions’. (2) tells the hearer to try to find things that match the descriptions – ‘*the*’ says ‘find a man’, the tensed form of the auxiliary says ‘find a salient past time’ – whereas (1) tells them to note that there are things that fit the descriptions.

Where should you find them? Not in the external world. It is evident that a term like ‘*the man*’ cannot be an instruction to find the one and only man in the universe; you have to find the one and only man in the current context. This becomes even clearer with pronouns and referential tenses of the kind in (2), where the context is ‘things that we have recently been talking about’.

It is, indeed, possible to refer to entities without accepting that they exist in the real world:

- (3) A: My boyfriend bought me a Ferrarri for my birthday
 B: I don’t believe you
 A: He did, he did
 B: Go on then, what colour is it?

B here can refer to the Ferrari without admitting that it exists, and indeed without admitting that she believes that A thinks it exists.

And it is also possible to use descriptions to pick out entities without believing the intended entity is the only thing in the universe that fits the description:

- (4) I saw a man and a woman yesterday. The man was wearing a blue dress.

Clearly, the man in the second sentence is the one introduced in the first. The uniqueness implicit in the use of ‘*the*’ is specified with respect to the context, not with respect to the entire universe.

This can be neatly captured by saying that ‘*the man*’ means ‘find the unique thing that can be proved to be a man either by using the information available in the minutes of this conversation¹ or by using the general knowledge that we share’. This captures the meaning of ‘*it*’ in (3) and ‘*the man*’ in (4), as well as more complex examples such as ‘*the first snowdrops of spring*’ in ‘*It always cheers me up when I see the first snowdrops of Spring*’² and ‘*John believes he will marry the richest debutane in Dubuque*’, but space precludes much more discussion of reference. The point to take away is that thinking about in terms of what can be **proved** given the information that the speaker and hearer believe they share provides, at least, a workable treatment of a range of awkward problems.

If referential and indefinite articles can be usefully treated as instructions to the hearer to manipulate their view of the minutes of the conversation, what about other quantifiers?

1.1 ‘*every*’ and ‘*most*’

What is a hearer supposed to do when they are told (5)?

- (5) Every existing theory of quantifiers is misguided.

¹All uses of language, barring perhaps diary entries, are conversations, even if, as with books, the participants are not known to each other and one of them never gets a turn: no-one would write a book unless they were expecting someone else to read it.

²Still waiting: lashing down with sleet outside my window right now

If you read this and decide to believe it, then you have a handy rule to use next time someone presents you with a theory of quantifiers. A handy rule to **use**. The point of making universal statements is to enable your hearer to use them for inference. They are not just supposed to sit there passively in the hearer’s head – interesting observations about what the world is like, perhaps, but not actually useful. They are rules which you hope will enable your hearer to do things (i.e, carry out inferences) that they would not otherwise have been able to do.

On this reading, universally quantified statements can also best be seen from a constructive point of view. They are rules, to be used, not statements, to be passively stored.

This becomes clearer if we consider ‘*most*’ rather than ‘*every*’.

- (6) Most formal dress shirts are white. Most white clothes show the dirt really badly.

Given (6), most people would infer that if you are going to wear a formal dress shirt then you should be more careful than usual about not getting gravy on your shirt.

This inference is not supported by interpretations of ‘*most*’ that say ‘*Most formal dress shirts are white*’ means $|white \cap formal_dress_shirt|/|formal_dress_shirt| > 0.5$. Quite apart from the fact that the set of dress shirts is essentially unbounded, so that $|formal_dress_shirt|$ is not really well defined, you just cannot go from $|A \cap B|/|A| > 0.5$ and $|B \cap C|/|B| > 0.5$ to $|A \cap C|/|A| > 0.5$.

If, however, we return to Reiter (1980)’s essentially proof theoretic notion of default logic³ then we can see exactly how (6) supports the conclusion that you should eat particularly daintily when wearing a dress shirt: to recap, Reiter’s suggestion is that you should use a rule involving ‘*most*’ as though it were a rule involving ‘*every*’ unless you have evidence that suggests that this is going to lead to a contradiction⁴.

1.2 ‘*few*’ and ‘*a few*’

- (7) A few students passed the semantics exam.
 (8) Few students passed the semantics exam.

What is the difference between (7) and (8)? They both say that a small number of students passed this exam; but (8) further says that the size of this group was disappointing, or surprising, or something like that. It says something about how big the group was, though what it says is rather vague, and it expresses an attitude to that size – a relation between the proposition that the number was small and the speaker’s hopes or expectations. Relationships of this kind between sets of propositions are essentially proof-theoretic. They allow the hearer to make inferences about the speaker’s internal state – that I am disappointed, for instance, because it suggests that I didn’t teach the semantics course very well. As such, they are better suited to a proof-theoretic account than to a model-theoretic one.

³Reiter does provide a model theory for this logic, but it is very artificial and is, indeed, little more than an unwrapping of the proof theory.

⁴e.g. if the shirt you are planning to wear is like Alec Guinness’ shirt in ‘*The Man in the White Suit*’.

1.3 and the others ...

Two pages isn't very much. There is much more to be said, but no space to say it. What I hope is that what I have managed to say is enough to suggest that the idea that quantifiers are instructions about what to do is worth exploring further; and then in a longer version of the paper, maybe it can be explored further.

References

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, **13**(1), 81–132.