The Impact of Syntax on the Interpretation and Graphical Depiction of Underspecified Propositions

Aaron Kalb, Dave Barker-Plummer[1],
Deonne Castaneda, Christopher Potts[2]
Stanford University, Stanford, CA 94305 USA
{kalb,dbp,deonne,cgpotts}@stanford.edu

Richard Cox
School of Informatics, University of Sussex
Falmer, E. Sussex, BN1 9QJ, UK
richc@sussex.ac.uk

Robert Dale
Center for Language Technology, Macquarie University
Sydney, NSW 2109, Australia
rdale@science.mq.edu.au

Abstract

Different representational systems permit differing degrees and forms of ambiguity and underspecification in the content they represent. Independently of this observation, a notable feature of natural language as a representational system is that it allows the same content to be expressed in different ways. In this paper, we examine the interaction of these two observations; in particular, we explore a number of linguistic forms involving underspecified content, and look at how subjects express the content of these linguistic forms both in logic and in diagrams. Our analysis demonstrates that variations in syntactic realization of the same semantic content lead to different interpretations of that content.

Keywords: logic; natural language; syntactic structure; diagrams; representations; negation.

Introduction

This paper takes as its starting point two widely-made observations. First, different representational systems permit different abstractions, and consequently, they permit underspecification on different dimensions. In particular, natural language (NL) and first-order logic (FOL) are two representational systems that permit underspecification of aspects of meaning that must be made explicit in diagrammatic representations. Direction is a case in point: consider the natural language statement *The house is adjacent to the park*; neither this sentence, nor a typical FOL rendering such as AdjacentTo(house, park), specifies the direction of adjacency, but a picture or diagrammatic rendering of the sentence must make this explicit. We will say that the NL representation of the state of affairs is **underspecified with respect to** the diagrammatic representation. This makes it clear that underspecification as defined here is a relational notion; however, for convenience in the remainder of this paper we will simply refer to representations as being **underspecified** when the relatum is obvious from the context.

The second observation we take as a starting point is that natural language affords multiple ways of realizing the same semantic content. This is often exemplified by reference to the fact that active and passive sentences, such as *Fred wrote the book* and *The book was written by Fred*, describe the same state of affairs. There may be contextual, or pragmatic, reasons for choosing one realization over the other, as commonly discussed under the heading of **information packaging** (Vallduvi, 1992); but the common view is that the semantics of the two sentences, *in terms of propositional content*, is the same.

We are interested in how these two phenomena interact. Our interest is motivated by an effect found in an earlier study (Cox, Dale, Etchemendy, & Barker-Plummer, 2008), in which the specificity of participants' responses to NL sentences containing negation differed markedly between their FOL translations and their diagrammatic interpretations. Specifically, it was found that in their FOL translations of the sentence ***d** is not a small dodecahedron*, participants overwhelmingly treated the predicates *small* and *dodecahedron* symmetrically, whereas their diagrams of the sentence tended to make **d** a dodecahedron that isn't small, rather than a small shape other than a dodecahedron. However, contextual confounds made the source of this effect hard to establish.

In this paper, we report on an experiment which sought to elucidate the effects of different possible factors on this phenomenon. In particular, we ask: if we have a number of natural language forms that express the same *underspecified* semantic content, what happens when subjects are asked to draw diagrams that require them to be more explicit? If the NL sentences truly express the same meaning, then we might expect to see similar distributions of the possible diagrammatic renderings, regardless of the NL surface form used. Alternatively, syntax or semantics may make some diagrammatic renderings more salient or available than others, This paper sets out to determine which of these alternatives hold.

---

[1]Center for the Study of Language and Information.
[2]Department of Linguistics.

Hypotheses

We explore two hypotheses in particular.

Hypothesis 1: When asked to translate from one representation into another that permits underspecification to be maintained, then in the absence of any contextual factors that encourage a more specific reading in the target representation, subjects will maintain the underspecification.

Hypothesis 2: When asked to translate from one representation into another that requires underspecification to be made specific, and there are a limited number of ways of doing this, we expect to see similar distributions across these solutions irrespective of superficial variations in the way the content is expressed in the source representation.

To test Hypothesis 1, we ask subjects to translate from NL to FOL. We make use of syntactic variations that represent the same semantic content; for example, the three sentences below all are expressions of the FOL statement $\neg(\text{Striped}(q) \wedge \text{Circ}(q))$:

(1)	**q** is not a striped circle	PREMOD
(2)	**q** is not a circle with stripes	POSTMOD
(3)	**q** is not striped and circular	COORD

The first two sentences are syntactically and semantically unambiguous. It is possible, with appropriate contextual cues, to encourage a more specific reading than the **wide-scoped** FOL statement above. For example, in spoken form, emphasis on either *striped* or *circle*, as in **q** *is not a striped circle* or **q** *is not a striped circle*, may encourage a **narrow-scoped** reading, corresponding to $\neg\text{Striped}(q) \wedge \text{Circ}(q)$ and $\text{Striped}(q) \wedge \neg\text{Circ}(q)$, respectively. However, in the absence of any such cues, Hypothesis 1 predicts that subjects will provide the wide-scoped reading.[3]

The third sentence is syntactically ambiguous (see Figure 1). Each parse corresponds to a different semantics, one of these being the wide-scoped reading, and the other the narrow-scoped-left reading. We would expect to find a distribution across these two readings in the FOL renderings.

To test Hypothesis 2, we ask subjects to translate from NL into diagrammatic realizations. We focus our analysis on those subjects who maintained underspecification in our test of Hypothesis 1, i.e., we leave aside any subjects who produce a narrow-scope reading for COORD sentences. We set up the diagram task conventions in such a way that there are only a limited number of possible ways of making the under-specified content specific. In particular, the wide-scope FOL above can be realized by three classes of diagrams:

---

[3]There is an extensive literature on scope ambiguity and its effects on human sentence processing (see, for example, (Kurtzman & MacDonald, 1993)) and on how discourse factors and lexical frequency impact on the processing of syntactic ambiguities (see, for example, (Trueswell, 1996; Spivey & Tanenhaus, 1998)); however, the focus of the former tends to be on quantifier scoping, and the latter is not obviously relevant to the kind of data we explore here. We are not aware of any existing work that looks at the processing of negated conjunctions of verbal complements, as explored here.
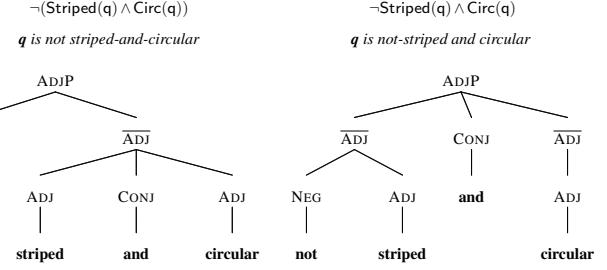

Figure 1: Two parse trees for **q** *is not striped and circular*

1. diagrams in which **q** is a circle that isn't striped (i.e. only the predicate Striped is realized-as-negated);

2. diagrams in which **q** is a striped object other than a circle (i.e. the predicate Circ alone is realized-as-negated); and

3. diagrams in which **q** is neither circular nor striped (i.e. both predicates are realized-as-negated).

Each of these realizations commits to some information left unspecified in the FOL sentence.

Our hypothesis predicts that the distribution of different diagrammatic realizations should be roughly similar irrespective of which surface NL form is being translated.

Methodology

The Subjects

Forty-one students enrolled in an introductory logic class at Stanford University took part. The experiment was conducted in the final weeks of the term. All of the background material necessary to complete the experimental task was presented within the first two weeks of the term. A key aim of the course is to teach the ability to distinguish the propositional content of sentences from their implicatures. The subject pool had therefore been primed to consider different possible interpretations of sentences and whether those interpretations depend on factors external to the sentence, such as common knowledge. Further, the students knowledge of FOL allows us to test for an unambiguous reading of sentences with multiple interpretations.

The class used (Barwise, Etchemendy, Allwein, Barker-Plummer, & Liu, 1999) as the textbook, and used the Tarski's World computer program for teaching the semantics of FOL. Tarski's World presents a system similar to the diagrammatic representation used in this experiment. In general, the materials were designed to parallel the structure of materials in the course, both in terms of the diagrammatic representations that were used, and the names used to refer to the distinct activities within the experiment.

Materials Administered

Subjects were given workbooks consisting of: (a) a page of study information; (b) a sheet consisting of 18 declarative natural language sentences to be translated; (c) a page of instructions; and (d) a half-page description with illustrative ex-

Table 1: Experimental Sentences in PREMOD Formulation

d is not a large dotted object	**k** is not a small circle
h is not a small dotted object	**l** is not a striped triangle
p is not a small striped object	**n** is not a dotted triangle
b is not a large triangle	**q** is not a striped circle
c is not a large circle	

amples of the diagrammatic representation and the first-order language to be used in the task. Following these were pages describing four activities to be completed. Three of these were presented as 'You Try It's (YTIs), and are described further below; these would be familiar to the subjects from (Barwise et al., 1999) as activities for becoming familiar with a concept. The final page of the workbook contained a more complex exercise in translation and realization, the contents of which are not discussed here.

Subjects were asked to complete all four activities with no time-limit and with no supervision. Only data from the second and third YTIs—translating the sentences into FOL and drawing realizations, respectively—are analyzed as part of this experiment; the remaining activities were included in order to format and embed the experimental tasks within an exercise form that the participants were familiar with from (Barwise et al., 1999). More importantly, they were designed to encourage subjects to submit spontaneous, naturalistic realizations.

The Tasks

As noted, subjects were given 18 natural language sentences in English. Nine of these were **negated logical conjunctions**, expressed in a form determined by the different conditions described below. The remainder of the sentences were fillers, also varying by condition, such as ***m** is a triangle that's not dotted*. The filler sentences all use the same vocabulary as the experimental sentence; some involve negation and others do not. All have unique readings.

The complete set of sentences (shown for condition PREMOD in Table 1) was counterbalanced such that each predicate is mentioned an equal number of times. In particular, of the nine experimental sentences, three mention size and pattern, three mention pattern and shape, and three mention size and shape; and each of the words *large*, *small*, *circle/circular*, *triangle/triangular*, *striped/stripes*, and *dotted/dots* are mentioned three times.

Subjects were asked to perform two tasks. The first task was to translate each sentence into a formal language of FOL, as discussed above. For the negated logical conjunctions, three FOL readings were possible, which we refer to as wide-scope, narrow-scope-left, and narrow-scope-right; for the example sentences introduced at the outset of the paper, these FOL readings are $\neg(\text{Striped}(q) \wedge \text{Circ}(q))$, $\neg\text{Striped}(q) \wedge \text{Circ}(q)$, and $\text{Striped}(q) \wedge \neg\text{Circ}(q)$ respectively.

The second task was to draw, for each sentence, a dia-

gram of a situation making the sentence true. We call these **diagrammatic realizations** of the sentences. We devised a highly constrained diagrammatic representation system in which objects have exactly three properties: shape, size, and pattern, with each of these properties having only two possible values (circle/triangle, small/large, striped/dotted). The students were asked to draw such objects in prepared spaces.[4]

Since the sentences have different readings, they may be realized in different ways, but the wide-scoped reading itself has multiple equally valid realizations. In the example above, **q** can be either a dotted circle, a striped triangle, or a dotted triangle. For such **multiply realizable** readings, each of the three possible realizations are equally valid, but the response requirement of a single diagram forces subjects to choose one.

The Conditions

The sentences of interest share the common property that they can be read as expressing the negation of a conjunction. Sentences (1)–(3), introduced earlier, are examples of such sentences. Each corresponds to one of three different conditions.

In a between-groups design, subjects were randomly allocated to one of these three conditions, named PREMOD (prenominal modifier, $N = 14$), POSTMOD (post-nominal modifier, $N = 11$) and COORD (coordination, $N = 16$). Subjects in each condition were presented with negated conjunctions expressed in one of these three forms.[5] Within each condition, subjects were randomly assigned to one of three random sentence orderings in order to control for possible priming effects within the stimulus sentences.

Data Collection and Encoding

Each worksheet was encoded independently by two coders. The FOL sentences and features of the diagrams[6] were recorded for each subject along with the condition that they were in. We also encoded which of the random sentence-orderings the subjects saw, but this information was not used for this study, as no systematic ordering effects were observed. Where they differed, the independent codings were

---

[4]These included guide lines for distinguishing large objects from small ones.

[5]The POSTMOD condition included an even mix of sentences with pattern expressed as a prepositional phrase (as above) and as a relative clause, as in ***q** is not a circle that's striped* In the COORD condition, the order of the predicates was varied, with some sentences of the form ***q** is not striped and circular* and others of the form ***q** is not circular and striped*. For the three sentences in the PREMOD and POSTMOD conditions which mention size and pattern but not shape, the word *object* is used as the noun (see the sentences describing **d**, **h**, and **p** in Table 1). Finally, in the POSTMOD condition, pattern is expressed post-nominally, but size is expressed as a pre-nominal adjective, as in the PREMOD condition, because the formulation ***k** is not a circle with small* is ungrammatical.

[6]These were encoded as large/small, striped/dotted, circle/triangle or as 'unclear' (if, for instance, a medium sized object were drawn), 'unspecified' (if, for instance, a shape were drawn with neither stripes nor dots), or 'other' (if, for instance, a square were drawn instead of a triangle or a circle).

Table 2: Readings: FOL scope by negation sentence condition

CONDITION	SCOPE		
	Wide	Left	Right
PREMOD ($N = 126$)	100%	0%	0%
POSTMOD ($N = 99$)	99%	1%	0%
COORD ($N = 132$)	42%	50%	0%

Table 3: Realizations of sentences with 'negatable heads'

CONDITION	REALIZED AS NEGATED		
	Both	Head Only	Mod. Only
PREMOD ($N = 68$)	34%	25%	41%
POSTMOD ($N = 38$)	21%	21%	58%

Table 4: Realizations of modifier-only (headless) sentences

CONDITION	REALIZED AS NEGATED		
	Both	First Only	Second Only
PREMOD ($N = 32$)	47%	18%	35%
POSTMOD ($N = 18$)	50%	17%	33%
COORD ($N = 52$)	71%	14%	15%

arbitrated by a third coder, who resolved disagreements.[7]

Results

Translations into FOL

We can measure the accuracy with which subjects completed the task of translating into FOL by considering their success in expressing an expected reading of each sentence. In the case of filler sentences, there is a unique expected FOL sentence, while for experimental sentences there are three possible readings for each. 78.6% of translations were expected. 92% of unexpected sentences were produced by four of the participants. Table 2 shows the proportions of each reading obtained for the experimental sentences.

Participants in the PREMOD and POSTMOD conditions almost universally wrote wide-scoped readings: only one sentence out of 225 was translated with a narrow-scoped reading. Subjects in the COORD condition displayed markedly different behavior. Table 2 gives the breakdown by condition, but the results are interesting when broken down by subject. 43.7% of subjects produced a wide-scoped reading for all (25%) or all but one (18.7%) of the nine sentences. 50.0% of these subjects always produced a narrow-scoped reading, and all of these were narrow-scoped-left.[8] The subjects who produced wide-scoped reading for all but one of the sentences were the only subjects to produced a mix of wide-

and narrow-scoped readings. In short, participants were systematic, with 50.0% always translating with narrow-scope-left, and 43.7% (almost) always translating with wide scope.

Thus, the results are consistent with Hypothesis 1: When subjects are asked to translate from one representation (NL) into another (FOL) that permits underspecification to be maintained, then the underspecification is indeed maintained. This is almost universally the case in the PREMOD and POSTMOD conditions, and also the case in around half of the COORD condition sentences, consistent with the fact that the latter are syntactically ambiguous, and one of the two parses is consistent with an underspecified reading.

Diagrammatic Realizations

Tables 3 and 4 show the results of encoding the realizations that students produced in the diagramming task. We recorded the predicates in the sentence that were **realized as negated** in each diagram. If the sentence is *q is not a striped circle* and the drawing was of a dotted circle, the predicate Striped is realized as negated and Circle is not.

Table 3 shows the results for sentences which are expressed syntactically with a head and modifier.[9] This pattern only arises in conditions PREMOD (**q** *is a striped circle*), and POSTMOD (**q** *is a circle with stripes*) in which the shape predicate is the head and the other is the modifier. The columns record whether both predicates, just the head predicate, or just the modifier predicate are realized as negated.

In the PREMOD and POSTMOD conditions, the sentences which do not mention shape (such as **d** *is not a large, dotted object*) contain a head predicate (*object*) which cannot be realized as negated. In COORD, all of the sentences lack a 'negatable' head (**q** *is striped and circular*). We will call such sentences **headless**, although this is not literally true in the PREMOD and POSTMOD conditions. Table 4 give the results for these sentences, with the columns indicating which modifier appears lexically first in the sentence.

Correspondence between Readings and Realizations Recall that approximately half of the subjects in the COORD condition wrote FOL sentences corresponding to the narrow-scoped-left reading of the sentences. These subjects universally drew diagrams consistent with this reading.

This suggests a strong alignment of readings with realizations, perhaps because the subjects referred to their FOL while producing the realizations, or because they arrived at the same mental representation on reading the sentence in preparation for translation into FOL and again in preparation for drawing their realization.

Similarly, subjects with wide-scoped readings in all three conditions drew diagrams consistent with this reading, although there are fewer possible incorrect realizations for these readings.[10] While variation in the narrow-scoped case

---

[7]The workbooks for all three conditions, an exemplar subject response, and the complete encodings can be downloaded from `http://openproof.stanford.edu/readingsandrealizations`.

[8]One subject (6.3%) wrote on the packet that the sentences were ambiguous, and submitted both a narrow-scoped left translation and a wide-scoped translation for each. Data for this subject was discarded.

[9]For convenience we talk from here on of 'sentences with heads and modifiers', although of course this refers to the heads and modifiers used in the descriptions of the objects.

[10]Subjects could only incorrectly realize **q** *is not a striped circle*
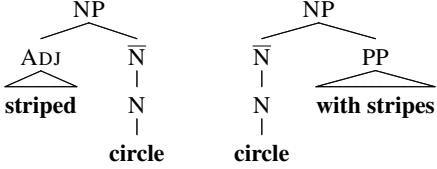
Figure 2: Asymmetric Parse Trees for Head-Modifier Constructions

would represent error (or at least inconsistency), variation in the wide-scoped case is expected.

Realizations of Wide-Scoped Readings

Subjects who obtained wide-scoped readings of the sentences have a choice of the realization that they can draw while remaining consistent with their reading. We focus on these subjects and that variability for the remainder of the analysis. Hence, we consider only the realizations for sentences with wide-scoped readings ($N = 100$ for PREMOD, $N = 63$ for POSTMOD, and $N = 60$ for COORD).

Heads vs. Modifiers

We first discuss the sentences that have heads that could have been realized as negated. These are just those sentences in conditions PREMOD and POSTMOD that mention the shape properties *circle* and *triangle*. For these sentences, the predicate which is expressed by a modifier is significantly more likely to be realized-as-negated than the predicate expressed by the head: In the PREMOD condition, the modifier is realized as negated 75% of the time while the head is realized as negated 59% of the time ($\chi^2 = 6.752$, $p < .01$). In the POSTMOD condition the modifier is realized as negated 79% of the time while the head is realized as negated 42% of the time ($\chi^2 = 10.794$, $p < .01$). Note that in some realizations both predicates are realized as negated.

This result mirrors that reported in (Cox et al., 2008). Our sentences analogous to *d is not a small dodecahedron* are those in the PREMOD and POSTMOD conditions which do not mention pattern. In our subjects' diagrams for these sentences, size was realized as negated significantly more often than shape. In 53.4% of the realizations of the 53 readings of the three sentences of the form *b is not a large triangle,* the size alone takes the negation. By contrast, participants negated just the shape or negated both predicates in only 22.3% and 24.3% of the realizations, respectively.

Modifier Choice

We now turn our attention to those sentences that only express properties via modifiers. All sentences in the COORD condition belong in this category, as do the sentences from the other two conditions which do not mention shape (e.g. *d is not a large object with dots*).

In 57% of the 105 realizations of these sentences, both predicates are realized as negated. In the 156 realizations of

by drawing a striped circle.

the other sentences (those with heads), both predicates are realized as negated only 35% of the time. This is a highly significant difference ($\chi^2 = 14.656$, $p < .001$).

It seems that when both predicates are expressed as modifiers, subjects are likely to realize them both as negated (perhaps because they must negate at least one and there is no obvious means of deciding which), while if one is expressed as a head, its identity is likely to be preserved.

It is worth noting, as well, that the tendency to realize both predicates as negated is most pronounced for sentences in the COORD formulation: both predicates are realized as negated in 71% of the realizations in this condition ($N = 52$), compared with 47% and 50% of the realizations of headless PREMOD and POSTMOD sentences, respectively. This may be because, in wide-scoped parses of a COORD formulation, the conjunction attaches to both arguments symmetrically (see Figure 1, left), so there are no structural differences whatsoever between the expressions of the two predicates.

Discussion

The results just discussed suggest that Hypothesis 2 does not hold. When subjects are asked to translate from one representation (NL) into another (a diagram) that requires underspecification to be made specific, the way in which this is done depends on the syntactic form used in the source representation. In particular, if a property is expressed via a syntactic nominal head, it is less likely to be realized as negated than when it is expressed via as a modifier.

There are other possible explanations for the observed behaviour, which we consider briefly below.

Ontological Primacy: *Perhaps shape as a concept is less readily negate-able than the other predicates.*

Since the only heads occurring in our sentences are the shape nouns *circle* and *triangle*, perhaps the phenomenon is due to some ontological primacy accorded to shape, but not to the other predicates. In our materials, shape is primarily seen as a *type* of object, whereas the other predicates are *attributes* of objects. If shape were protected because of its ontological status, rather than because of the way it is expressed, we would see these same results in conditions PREMOD and POSTMOD, since the only heads appearing in our sentences are the shape predicates. However, if it were the ontological status of shape that were protected, we would expect it to be protected in the COORD condition as well, even though in that condition shape is expressed as a modifier.

Among sentences in the COORD condition, however, shape is realized as negated 77.0% of the time ($N = 39$) (Figure 3)— just as much (more, in fact) than the other predicates. This strongly suggests that shape, as a concept, is not protected.

Surface Proximity: *Perhaps participants simply tend to negate the predicate closest to the word* not.

In sentences such as *q is not a striped circle*, *striped* is closer to *not* and perhaps this accounts for the preference for realizing this predicate as negated.
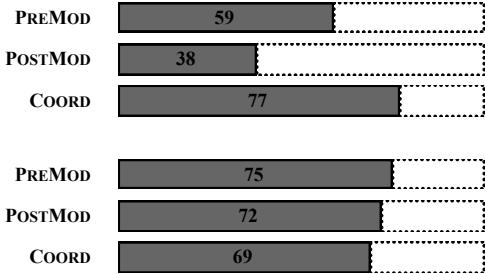
Figure 3: Among realizations of sentences which mention shape, % which negate shape (TOP) vs. % which negate the other predicate—size or pattern (BOTTOM)


only the first only the second both
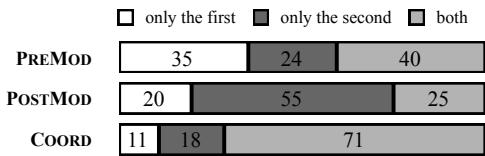
Figure 4: % of realizations of sentences which mention pattern and shape in which — predicate is negated.

PREMOD	*q is not a striped circle*	$N = 33$
POSTMOD	*q is not a circle with stripes*	$N = 24$
COORD	*q is not striped and circular*	$N = 19$

Looking at the readings of the three sentences which mention the predicates pattern and shape (Figure 4), we see that students are somewhat (though not significantly) more likely to negate just pattern than just shape in the PREMOD condition, where the pattern predicate occurs closest to the word *not*. However, when phrased so that the pattern predicate occurs farthest (in the POSTMOD condition with sentences like *q is not a circle with stripes*), we find that pattern continues to take the negation—this time, 2.8 times as often as just shape (more, in fact, than when it occurs in closer proximity to *not*). The difference in likelihood to realize-as-negated just the first vs. just the second predicate across the PREMOD and POST-MOD conditions is significant ($\chi^2 = 3.979, p < .05$). Moreover, we see no tendency whatsoever toward negating the closer predicate in the COORD condition, for any sentence.[11]

Conclusion and Future work

We set out to test two hypotheses, one of which suggested that subjects would maintain underspecification in their representations if this was possible, and a second which suggested that, if subjects had to translate into a representation that required more specificity than the source representation, then the results would be the same for semantically-equivalent source representations.

The evidence from our experiment supports the first hypothesis. This allowed us to go on to test our second hypothesis, where the results turned out to be surprising: we demonstrated that the same semantic content, expressed in natural language in different ways, leads to different interpretations when subjects are asked to express that information in diagrams which require them to choose a more specific representation.

This is unexpected. Of course, it is not suprising that the particular form of an utterance has an impact on how that utterance is interpreted; but such variations are usually considered to be in the realm of pragmatics, and more concerned with connotation than with denotation. The results here, however, indicate that how something is expressed has an impact not only in terms of the pragmatic aspects of interpretation, but also in terms of the state of affairs in the world the utterance is taken to describe.

If we characterize shape via a noun, then it is less likely to be negated than if it is expressed via an adjective or other modifier. It would appear that it is how things are described, or how, in Langacker's terms (Langacker, 1991), they are **construed**, that governs our interpretation; not what they are.

Acknowledgements

References

Barwise, J., Etchemendy, J., Allwein, G., Barker-Plummer, D., & Liu, A. (1999). *Language, proof and logic*. CSLI Publications and University of Chicago Press.

Cox, R., Dale, R., Etchemendy, J., & Barker-Plummer, D. (2008). Graphical revelations: Comparing students' translation errors in graphics and logic. In *Diagrams 2008, fifth international conference on the theory and application of diagrams.* (Herrsching, Germany, September 2008)

Kroch, A. (1974). *The semantics of scope in English*. Unpublished doctoral dissertation, Massachusetts Institute of Technology.

Kurtzman, H., & MacDonald, M. (1993). Resolution of quantifier scope ambiguities. *Cognition*, *48*, 243-279.

Langacker, R. (1991). *The foundations of cognitive grammar*. Stanford University Press.

Spivey, M. J., & Tanenhaus, M. K. (1998). Syntactic ambiguity resolution in discourse: Modeling the effects of referential context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1521–1543.

Trueswell, J. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of Memory and Language*, *35*, 566–585.

Vallduvi, E. (1992). *The informational component*. New York: Garland.

---

[11](Kroch, 1974) proposes 'a general surface ordering principle that fixes the initial scope order of the operator words in an English sentence according to their surface order'; however, in line with our findings, this claim is refuted by (Kurtzman & MacDonald, 1993).