

ASSIGNING A SEMANTIC SCOPE TO OPERATORS

Massimo Poesio

University of Rochester, Department of Computer Science
Rochester, NY 14627-0226, USA
poesio@cs.rochester.edu

Abstract

I propose that the characteristics of the scope disambiguation process observed in the literature can be explained in terms of the way in which the model of the situation described by a sentence is built. The model construction procedure I present builds an *event structure* by identifying the situations associated with the operators in the sentence and their mutual dependency relations, as well as the relations between these situations and other situations in the context. The procedure takes into account lexical semantics and the result of various discourse interpretation procedures such as definite description interpretation, and does not require a complete disambiguation to take place.

THE PROBLEM

Because new ways of obtaining semantically distinct interpretations for sentences are continuously discovered, coming to grips with ambiguity is becoming more and more of a necessity for developers of natural language processing systems, linguists and psychologists alike [9, 31, 7, 2]. In this paper, I am concerned with the *scopal* ambiguity of operators¹ [31, 33].

The attention of both psycholinguists and computational linguists interested in ambiguity has concentrated on the problem of combinatorial explosion. If the number of readings of an utterance were to actually grow with the factorial of the number of operators, even a simple sentence like (1), with 4 operators (the modal ‘should’, tense, an indefinite and a definite), would have $4! = 24$ scopally different readings. Two distinct questions thus must be answered: how can listeners (and how should machines) deal with the combinatorial explosion of readings? Do we really use the brute-force strategy of considering all of the available readings, and then choose among them? And, if we do choose among several readings, how is that done?

(1) We should hook up an engine to the boxcar.

To my knowledge, three positions on the problem of combinatorial explosion have been taken in the literature. Some have argued that there is no problem: our brains contain

more than enough machinery to process in parallel $4!$ interpretations. It’s unclear, however, whether this strategy is feasible when larger numbers of readings are concerned. A classical demonstration of the number of readings one may have to consider is (2), which has $11!$ interpretations if the standard treatment of quantification and modality is assumed.

(2) You can fool most people on most of the issues most of the time, but you can’t fool everybody on every single issue all of the time. [15]

Another position is that sentences like (1) are not semantically ambiguous, but *vague*. Consider for example (3):

(3) Every kid climbed a tree.

Here, one of the readings (the one in which the indefinite takes narrow scope) is entailed by the other (in which the indefinite takes wide scope). The claim is that (3) is interpreted in the vaguest possible way, and the strongest reading, if at all, is derived by pragmatic ‘strengthening’ [25]. A difficulty with this approach is that a vaguest reading doesn’t always exist. The two readings of (4), for example, are distinct.

(4) Few people speak many languages. [27]

Finally, it has been proposed that the reason why listeners do not seem to have problems in processing utterances like (1) is because they do not disambiguate. They build a non-disambiguated representation of the sentence and leave the interpretation open. This strategy might be advantageous for some kinds of applications² and it has been argued that a complete disambiguation never takes place [7].

No matter what processing strategy is chosen, the question of how listeners choose one particular interpretation cannot be ignored. All experimental work done on the subject of scopal ambiguity [20, 35, 26] indicates that subjects do have preferred interpretations when confronted with tasks which require understanding. In addition, sentences like (1), (5) and (6) clearly have preferred interpretations. However, the only answers to to this question that I have seen are based on heuristics.³

¹I use here the term *operator* as it is used by Heim [13], i.e., to mean either quantifier or modal/tense operator.

²E.g., machine translation [2].

³See [17] for an example of state-of-the-art techniques

- (5) A girl took every chemistry course. [20]
 (6) Each daughter of an admiral married a captain.

I present in this paper an hypothesis about interpretation that accounts for facts about scope disambiguation that were previously explained in the literature by stipulating a number of unmotivated principles. The proposal developed here is being applied to develop the module of the TRAINS-93 system [1] that handles scope disambiguation and reference interpretation. The goal of the TRAINS project is to develop a conversationally proficient planning assistant. More details about the project and the work presented here can be found in [29].

SCOPE DISAMBIGUATION FACTORS

Most proposals on scope disambiguation were developed to account for the general preference of the leftmost quantified phrase from taking wide scope in simple active sentences like (7):

- (7) Every kid climbed a tree.

Lakoff [27] proposed that this preference is due to the fact that sentences are parsed from left to right; “every kid” takes scope over “a tree” because it is processed first. (Kurtzman and MacDonald called this the *Left to Right principle*.)

Ioup [20] argued instead that “...in natural language, order has little to do with the determination of quantifier scope.” ([20], p.37). The preferred reading of (8), for example, is the one in which the NP “each child” takes wide scope.

- (8) I saw a picture of each child. [20]

According to Ioup, the relative scope of quantifiers is determined by the interaction of two factors. First of all, quantifiers such as “each” or “the” have the inherent property of taking wide scope over indefinites, which, in turn are lexically marked to take scope over plural quantifiers like “all.” This hypothesis is motivated by contrasts such as those in (9), and accounts for cases such as (8).⁴

- (9) a. I saw a picture of each child.
 b. I saw a picture of all the children.

Secondly, Ioup proposed that a hierarchy exists among grammatical functions, such that listeners tend to attribute to NPs in subject position wide scope over NPs in indirect object position, which in turn tend to take wide scope over NPs in object position. The hierarchy between grammatical functions accounts for the preferred reading of (7).

Ioup also observed that NPs in topic position tend to take wide scope. This is especially obvious in languages that have a specific grammatical category for topic, like Japanese or Korean. The Japanese sentence (10b) is ambiguous, but the reading in which the NP in subject position, “most students” takes scope over the NP in object position, “every language,” is preferred. This preference is maintained if the

⁴Van Lehn [35] and Hendrix [14] also studied the effect of lexical preferences, or ‘strengths’ as they are also called.

NP in object position is scrambled in sentence-initial position, as in (10c) (another counterexample to Lakoff’s left-to-right principle). If, however, the NP is marked with the topic-marking suffix “wa,” as in (10d), suddenly the preferred reading of the sentence becomes the one in which “every language” takes wide scope.⁵

- (10) a. Most students speak every language.
 b. Hotondo-no gakusei-ga subete-no gengo-o hanasu
 most-gen student-nom every language-acc speak
 c. Subete-no gengo-o hotondo-no gakusei-ga hanasu
 every language-acc most-gen student-nom speak
 d. Subete-no gengo-wa hotondo-no gakusei-ga hanasu
 every language-TOP most-gen student-nom speak

Several proposals attribute an important role to structural factors in assigning a scope to operators. Jackendoff [21] and Reinhart ([32], ch. 3 and 9) propose to account for the preferred reading of (7) by means of a *C-command principle* according to which a quantified expression is allowed to take scope over another quantified expression only if the latter is c-commanded by the former at surface structure.

Structural explanations (in the form of constraints on syntactic movement) have also been proposed to explain the constraint that prevents a quantifier to take scope outside the clause in which it appears, first observed by May [28] and called *Scope Constraint* by Heim [13]. This constraint is exemplified by the contrast in (11): whereas (11a) has a reading in which “every department” is allowed to take wide scope over “a student,” this reading is not available for (11b).

- (11) a. A student from every department was at the party.
 b. A student who was from every department was at the party.

Lexical semantics and commonsense knowledge also play an important role in determining the scope of operators. The contrast between the preferred readings of (12a) and (12b) can only be explained in terms of lexical semantics:

- (12) a. A workstation serves many users.
 b. A workstation can be found in many offices.

Kurtzman and MacDonald [26] set out to verify the empirical validity of several of these principles. The most crucial result is that none of the principles they set to verify can account for all the observed effects, and actually counterexamples to all of them—including the quantifier hierarchy—can be found. No evidence for a Left-to-Right processing principle was found. Kurtzman and MacDonald hypothesize that “...processes that are not strictly dedicated to the interpretation of scope relations may nonetheless influence the interpretation of quantifier scope ambiguities.” ([26], p.22). They conclude that “...the results leave open the question of whether the building and selection of representations of scope are mandatory processes” ([26], p.45).⁶

⁵Arguably, the closest thing to an explicit topic marker in English are certain uses of definite descriptions and the topicalization construction; in both cases, the topically marked NP tends to take wide scope.

⁶Their experiments are discussed in more detail in [29].

OVERVIEW OF THE PROPOSAL

Scope Disambiguation as Construction of an Event Structure

It is commonly assumed in the psycholinguistic literature on sentence interpretation that hearers interpret sentences by constructing a *model* of the situation described by the sentence [10, 22]. I propose that the scope assigned to the operators contained in a sentence is determined by the characteristics of the model construction procedure. The model being constructed, which I call *event structure*, consists of a set of *situation descriptions*, one for each operator, together with dependency relations between them. The task of the model construction procedure is to identify these situations and to establish dependency relations. The scope assigned by a hearer to an operator depends on the position of the situation associated with that operator in the event structure.

For example, I propose that the scope assigned to quantifiers depends on how their *resource situation* [3, 8] is identified. It is well-known that a sentence like (13):

(13) Everybody is asleep.

is not interpreted as meaning that every single human being is asleep, but only that a certain contextually relevant subset is. The process of identifying the set of individuals over which an operator quantifies is usually called *domain restriction*. In the case of, say, (7) whether “every kid” or “a tree” takes wide scope depends on how the listener builds a model of the sentence. If she starts by first identifying a situation containing the group of kids that “every” is quantifying over, and then proceeds to ‘build’ for each of these kids a situation which contains a tree the kid is climbing, then “every kid” will take wide scope. In other words, I propose that a listener has a preferred reading for a sentence if she’s able to identify the resource situation of one or more of the operators in that sentence (‘to picture some objects in her mind’), and to hypothesize dependency relations between these situations. If this process cannot take place, the sentence is perceived as ‘ambiguous’ or ‘hard to understand.’

The less context is available, the more the establishment of dependency relations between situations depends on the order in which the model is built, i.e., on the order in which the situations associated with the different operators and events are identified. This order depends in part on which NPs are perceived to be ‘in topic,’ and in part on general principles for building the conceptual representation of events (see below). In addition, some operators (e.g., definite descriptions) impose constraints on their resource situation.

A Model Construction Procedure: The DRT Algorithm

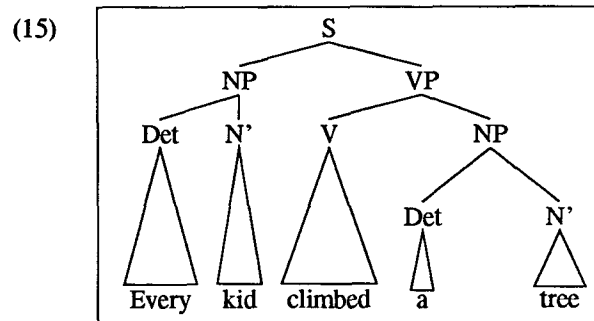
In order to make the intuition more concrete we need the details of the model construction procedure. Ideally, one would want to adopt an existing procedure and show that the desired results fall out automatically. Unfortunately, the model construction procedures presented in the psycholinguistic literature are not very detailed; often it’s not even clear what these researchers intend as a model. There is,

however, a discourse interpretation procedure that is specified in detail and has some of the characteristics of the model construction procedure I have in mind; I’m thinking of the *DRS construction algorithm* [23, 24].

The DRS construction algorithm consists of a set of rules that map discourses belonging to the language into certain “interpretive structures”. The output structures are called “Discourse Representation Structures” or “DRSs.” A DRS is a pair consisting of a set of *discourse referents* and a set of *conditions* (= predicates on the discourse referents). The construction algorithm works by first adding the syntactic structure of the sentence to the ‘root’ DRS representing the discourse up to that point, then applying the rules to the syntactic structure, thus adding discourse referents and conditions to the DRS. Consider how the algorithm is applied to obtain an interpretation for (7):

(14) Every kid climbed the tree.

The initial interpretation of (14) is the tree shown in (15).



The DRS construction rule for definites and universal quantification are as follows:

(Definite Descriptions) When a syntactic configuration containing a definite NP is met in a DRS K ,

1. Add a new discourse referent x to the root DRS,
2. Add a new condition to the root DRS representing the restriction on the indefinite NP,
3. Replace the NP with x in the syntactic configuration.

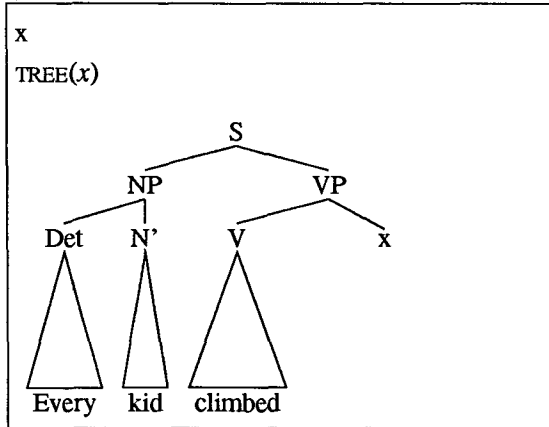
(Universal Quantification) When a syntactic configuration containing an NP with determiner “every” is met in a DRS K ,

1. Add a *complex condition* $K_1 \rightarrow K_2$ to K ,
2. Add a new discourse referent x to K_1 ,
3. Add a new condition to K_1 representing the restriction on the indefinite NP,
4. Replace the NP with the discourse referent in the syntactic configuration,
5. Move the syntactic configuration insides K_2 .

Both the rule for definites and the rule for universal quantification are triggered by (15). Two hypotheses are obtained; that obtained by applying first the rule for definite descriptions is shown in (16). Both of these hypothesis contain operators whose DRS construction rules haven’t been applied yet: this algorithm comes with a built-in notion of *partial*


hypothesis—a partial hypothesis is a DRS some of whose operators still have to ‘interpreted’ in the sense just mentioned.

(16)



The two partial hypotheses are made into complete hypotheses by applying the remaining rules; the complete hypothesis with the definite taking wide scope is shown in (17).

(17)



Modifying the DRS Construction Algorithm

Because the DRS construction rules depend on syntactic patterns, the role of structural factors in disambiguation can be taken into account—and a lot of data about disambiguation preferences can be explained without any further machinery. The Scope Constraint, for example, is embedded in the very semantics of DRT; and one can ‘build in’ the construction rules principles such as the c-command principle. (Kamp and Reyle do just that in [24].) The limitations of this approach are shown by examples in which the choice of an interpretation does not depend on the structure, like (12). Also, the rule for definites as just formulated is too restrictive: in cases like (18), for example, predicts the correct reading for the definite NP “the meeting,” but the wrong one for “the principal,” that, intuitively, takes narrow scope with respect to “every school:”

(18) Every school sent the principal to the meeting.

I propose that the role of lexical semantics, as well as the data accounted for in the literature by introducing principles such as the grammatical function hierarchy, the topic principle, and the quantifier hierarchy, can be accounted for by making the activation of the DRS construction rules depend on factors other than the syntactic structure of the sentence. The factors I propose to incorporate are (i) the semantics of lexical items, (ii) the results of the interpretation of opera-

tors in context, and (iii) the way the representation of events is built in memory.

In order to achieve this goal, I propose two main modifications to the standard DRS construction algorithm. First of all, I propose that the input to the algorithm is a *logical form*—a structure isomorphic to the s-structure, that carries however information about the semantic interpretation of lexical items. In this way, the role of semantic factors in interpretation can be taken into account; in addition, a semantic value can be assigned to a representation containing unresolved conditions or partial hypotheses. Secondly, I propose to make the application of the DRS construction rules depend on the identification of certain contextually dependent elements of the interpretation. The ingredients of the account thus include: a proposal about the input to the model construction procedure; a notion of what an event structure is; and an account of discourse interpretation. I discuss these issues in turn in the next sections.

THE LOGICAL FORM

As said above, the first difference between the interpretation procedure proposed here and the DRS construction algorithm illustrated above is that the rules I propose rely on semantical and contextual factors. I propose to do this by adding to standard DRT a new class of conditions, that I call ‘logical forms.’ Logical forms include semantic information about the lexical items occurring in the sentence. The logical form representation is the interface between the parser and the model construction algorithm, and can be compositionally obtained by a GPSG parser [11, 18] that couples a context-free grammar with rules of semantic interpretation. I first describe the language used to characterize the semantics of lexical items, *SEL* (for Simple Episodic Logic), then the syntax and interpretation of logical forms.

Lexical Semantics in Simple Episodic Logic

I introduce *SEL* by presenting the truth conditions I propose to assign to (18), repeated here for convenience:

(18) Every school sent the principal to the meeting.

The truth conditions usually assigned to (18) in a language with restricted quantification, and ignoring tense, are shown in (19); I propose instead to assign to (18) the interpretation specified by (20).

(19) (the m MEETING(m)
(\forall s SCHOOL(s)
(the p PRINCIPAL(p,s)
SENT(s,p,m))))

(20) (the m [$s_1 \models$ MEETING(m)] \wedge SHARED(spkr,hearer, s_1)
(\forall s [$s_2 \models$ SCHOOL(s)]
(the p [$s_3 \models$ PRINCIPAL(p,s)]
 \wedge SHARED(spkr,hearer, s_3)
SENT(s,p,m))))

(20) reads: there exists a unique m that is a meeting in a contextually specified *resource situation* s_1 , and for all s’s that are schools in a contextually specified resource situation s_2 the unique p such that p is the principal of s participates to m.

The intent of the expression used for the quantifier restrictions in (20) is to make it explicit that the situations from which the quantified elements are ‘picked up’ need not be the complete set of objects and relations at which the truth of (20) is evaluated. This is accomplished by introducing into the language an explicit relation \models (‘supports’) to represent ‘truth at a situation’ [8]. A statement of the form

$$[s_1 \models \text{MEETING}(x)]$$

evaluates to true in a situation s if the object—say, m —assigned to the variable x is a meeting in the situation s_1 . A *situation* is a set of objects and facts about these objects [8, 18]. I assume a language which allows us to make statements about situations, and an ontology in which situations are objects in the universe. Episodic Logic provides such a language and such an ontology [19, 18]; where not otherwise noted, the reader should assume that an expression of *SEL* has the semantics of the identical expression in Episodic Logic.

The restriction of the existential quantifier in (20) contains a *parameter* s_1 . Parameters are used in *SEL* to translate anaphoric expressions of English. A parameter behaves semantically as an open variable, a value for which has to be provided by context.⁷

I have assumed the following translations for the lexical items “every,” “meeting,” and “sent” (I have again ignored tense):

$$\begin{aligned} \text{“every”} &\rightsquigarrow \lambda P \lambda Q (\forall x [s_1 \models P(x)] Q(x)) \\ \text{“meeting”} &\rightsquigarrow \text{MEETING} \\ \text{“sent”} &\rightsquigarrow \text{SENT} \end{aligned}$$

The semantics assigned to definite descriptions needs a bit of an explanation. According to the *location theory* [12, 4] the major uses of definite NP’s, as well as the contrast between definites, indefinites, and demonstratives, can be accounted for by stipulating that a speaker, when using a definite article,

1. instructs the hearer to locate the referent in some shared set of objects, and
2. refers to the *totality* of the objects/mass within this set that satisfy the restriction.

I formalize this idea in [30] by associating to definite descriptions the translation below. A situation is ‘shared’ between x and y if every fact Ψ supported by that situation is mutually believed by x and y (see [30] for details).

$$\text{“the meeting”} \rightsquigarrow \lambda P (\text{the } x: ([\dot{S} \models \text{MEETING}(x)] \wedge \text{SHARED}(\text{spkr}, \text{hearer}, \dot{S})) P(x))$$

Syntax and Interpretation of the Logical Form

The translations seen above, together with the obvious context-free rules, result in the following LF for (18) (I have

⁷See [29] for details. The idea is to add to the parameters of evaluation an *anchoring function* a that provides the values for parameters, thus plays the role of ‘context’ in Heim’s proposal. The reader should be aware that while the notation and terminology I have adopted is borrowed from Situation Theory, parameters have a different semantic interpretation there [8].

used here, and elsewhere in the paper, a linear notation to save space):

$$(21) \text{ [CP [IP [NP } \lambda Q (\forall s [s_2 \models \text{SCHOOL}(s)] Q(s))] \text{ [VP [VP [V' 'SENT [NP } \lambda Q (\text{the } p [s_3 \models \text{PRINCIPAL}(p, \dot{k}) \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_3) Q(p))]]]] \text{ [PP 'TO [NP } \lambda Q (\text{the } m [s_1 \models \text{MEETING}(m) \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_1) Q(m))]]]]]]]$$

I propose that expressions like (21) can appear as conditions of DRSS. The syntax of LFs is as follows. Each internal node of (21) is labeled with a phrase category; the leaves are labeled with expressions of the form ‘ α ’, where α is an expression of *SEL* (and has therefore a ‘standard’ model theoretic denotation). I use the phrase structure system largely adopted in the Government and Binding literature, according to which the sentence is the maximal projection of an Infl node and is therefore labeled IP [34]. I also assume the existence of a maximal projection of complementizer CP above IP. Because I don’t discuss relatives here, I use the following simplified notation for NPs with determiners, such as “every school”:

$$[\text{NP } \lambda Q (\forall x [s_1 \models \text{SCHOOL}(x)] Q(x))]$$

LFs like (21) are usually treated in the natural language processing literature as uninterpreted data structures from which to ‘extract’ the readings [16, 17]. However, it has been recently proposed [31, 2, 33] that it is possible (and indeed desirable) to assign a denotation to expressions like (21). The reason is that in this way one can define a notion of sound inference—that is, one can specify what can and cannot properly be inferred from an expression like (21) prior to disambiguation; and therefore, a notion of ‘monotone disambiguation.’ I do not assume disambiguation to work monotonically, but I want to be able to treat expressions like (21) as full-fledged conditions so that a DRS containing a condition of this kind can be interpreted, and I need to be able to characterize a disambiguation step as *compatible* in the sense that it does not introduce any new readings. To do this I need LFs to have an interpretation.

Were it not for the problem that more than one interpretation can be associated to a single LF, one could easily define a recursive mapping EXT from logical forms to truth-theoretical denotations (functions from situations to truth values) in terms of the usual $\|\ \|\$ function, as follows:

$$\begin{aligned} \text{EXT}(\alpha) &= \|\alpha\| \\ \text{EXT}([\text{V}' \alpha]) &= \text{EXT}(\alpha) \\ \text{EXT}([\text{VP } \alpha]) &= \text{EXT}(\alpha) \\ \text{EXT}([\text{N}' \alpha]) &= \text{EXT}(\alpha) \\ \text{EXT}([\text{NP } \alpha \beta]) &= \text{EXT}(\alpha)(\text{EXT}(\beta)) \\ \text{EXT}([\text{IP } \alpha \beta]) &= \text{EXT}(\alpha)(\text{EXT}(\beta)) \\ &\text{if TYPE}(\text{EXT}(\alpha)) = \langle t_1, t_2 \rangle \\ &\text{and TYPE}(\text{EXT}(\beta)) = t_1; \\ &\text{EXT}(\beta)(\text{EXT}(\alpha)) \text{ otherwise.} \end{aligned}$$

Once this is done, one can reformulate the semantics of DRS in terms of situations and situations extensions instead of embeddings and embedding extensions, and interpret all conditions as functions from situations to truth values. (See [29] for details.)

Matters get more complicated when expressions with more than one reading like (21) are considered. Different ways for assigning a denotation to expressions with more than one interpretation have been proposed [2, 31]; my proposal derives from [31]. I use a Cooper storage mechanism [5] to define EXT in such a way as to allow for an LF to have more than one ‘indirect interpretation.’ Briefly, Cooper’s idea is to have a syntactic tree denote a set of *sequences*, each sequence representing a distinct ‘order of application’ in computing the interpretation of the sentence. For example, because in interpreting (22) one can either apply the translation of tense immediately or wait, EXT maps (22) in a set of two sequences, shown in (23).

$$(22) \quad [V' \text{ P } [NP \text{ } \lambda Q (\det x R(x)) Q(x)]]$$

$$(23) \quad \text{EXT}((22)) = \{ \{ \lambda x (\det x R(x)) P(x) \}, \\ \{ P, \lambda Q (\det x R(x)) Q(x) \} \}$$

I omit here the definition of the EXT function implementing Cooper storage, that is rather complex. For the current purposes, it is enough to understand that EXT associates to (21) a set of functions from situations to truth values, as in (24).

$$(24) \quad \text{EXT}((21)) =$$

{ the function denoted by

$$\| (\text{the } m [s_1 \models \text{MEETING}(m)] \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_1) \\ (\forall s [s_2 \models \text{SCHOOL}(s)] \\ (\text{the } p [s_3 \models \text{PRINCIPAL}(p, s)] \\ \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_3) \\ \text{SENT}(s, p, x)))) \|,$$

the function denoted by

$$\| (\forall s [s_2 \models \text{SCHOOL}(s)] \\ (\text{the } m [s_1 \models \text{MEETING}(m)] \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_1) \\ (\text{the } p [s_3 \models \text{PRINCIPAL}(p, s)] \\ \wedge \text{SHARED}(\text{spkr}, \text{hearer}, s_3) \\ \text{SENT}(s, p, x)))) \|, \text{ etc. } \}$$

Having done this, we can say that a DRS condition like (21) is verified the current situation s if one of the functions denoted by (21) maps s into 1.

BUILDING EVENT STRUCTURES

Not all assertions in a narrative or conversation are going to be about the same situation. In the conversations with the TRAINS system, for example, the participants can discuss both the state of the world and the state of the plan being developed. Maintaining this separation is crucial for the proper interpretation of definite descriptions, for example. The separation between the situations that are the topic of different sentences is achieved by translating sentences as *situation descriptions*. A situation description is a condition of the form:

$$(25) \quad s: \begin{array}{|c|} \hline x \\ \hline \Phi \\ \hline \end{array}$$


whose intuitive interpretation is that Φ provides a partial characterization of the situation s . The semantics of situation descriptions is defined as follows, using a semantics of DRSs in terms of *situation* extensions, as discussed in the previous section, and interpreting discourse markers as constituents of situations:

The condition $s:K$ is satisfied wrt the situation s' iff K is satisfied wrt the value assigned to s in s' .


I also propose the following constraint on the model construction rules:

Constraint on Interpretation : with the exception of the discourse markers interpreted over situations and of the situation descriptions, every discourse marker and condition has to be part of a situation descriptions.

Situation descriptions are added to the model by rules triggered by an LF whose root is a CP node. The rules (now shown for lack of space) delete the complementizer and its whole projection, and introduce a situation structure. The result is shown in (26).



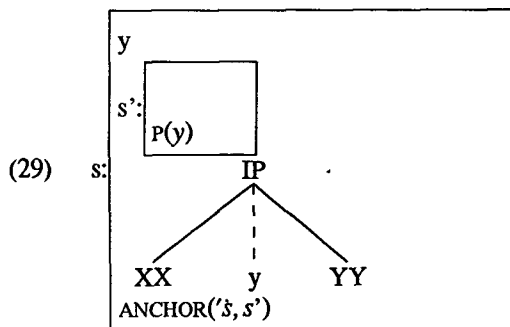

The constraint on discourse interpretation proposed above is implemented by forcing the rules that build situation structures to be triggered before any other rule; this is done by having every other rule being triggered by LFs whose root node is an IP. The result of this constraint is that a discourse model consists of a set of situation descriptions:



The DRSs produced by the standard DRT algorithm are semantically equivalent to the special case of a set of situation descriptions all describing the same situation s .

Models like the one in (27) enable the formalization of processes of resource situation identification like that described in [30]. I illustrate how my rules for interpreting operators differ from those of standard DRT, and how the interaction between model construction rules and discourse interpretation works, by means of the model construction rule for definites. The rule MCR-DD is triggered by the configuration in (28), and results in the configuration in (29). The notation used for the pattern indicates that this rule applies to a definite NP in any position within a syntactic tree whose

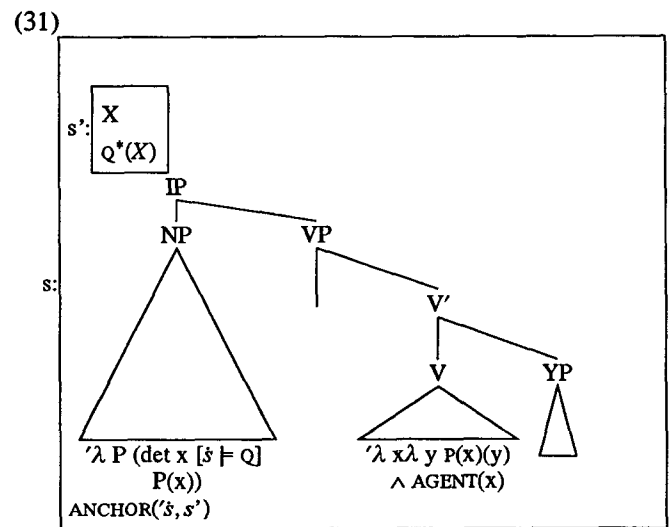
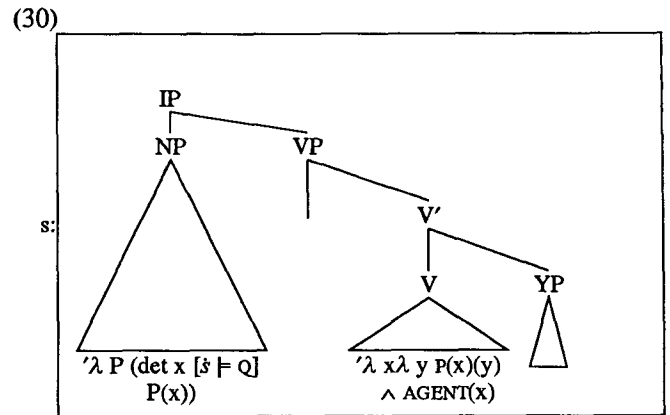
maximal projection is an IP node, without any intervening IP node.



The key observation is that the application of this rule, as well as of any other NP rule, depends on the hearer's previous identification of a resource situation for the definite description. The statement $\text{ANCHOR}('s, s')$ constraining the interpretation of s is added to the situation structure by the processes that identify the referent of the definite description; I describe these processes in detail in [30].⁸

Finally, I propose that, when context is missing, a default model construction procedure operates. It has been suggested [6] that the conceptualization of events follows an order reflected in the *thematic hierarchy* AGENT < LOCATION, SOURCE, GOAL < THEME proposed to account for phenomena like passivization [21]. Briefly, the idea is that 'the normal procedure for building an event description' is to follow the order in the hierarchy: first identify the agent, then the location, then the theme. This proposal can be formalized in the current framework by having rules that operate in case no other rule has, and that modify the model by introducing a resource situation for an operator and establishing anchoring connections. These rules depend both on the semantics of the verb and on the syntactic configuration. The rule that identifies the AGENT, for example, is triggered by the configuration in (30), and results in the configuration in (31), that allows for the rule for the NP to operate in that the resource situation of the operator has been anchored:

⁸A more conventional situation-theoretic framework is used there, but the analysis carries over to the framework in this paper.



These rules can of course originate conflicts with the results of other discourse interpretation processes. I assume the following conflict resolution rule: when two rules produce conflicting hypothesis, assume the result of the more specific rule. In general, the discourse interpretation rules are more specific than the default rules for constructing events representations, so they will be preferred.

Although lack of space prevents me from giving examples, rules relating the construction of the model to lexical semantics, such as those accounting for data like (12), can also be formulated.

AN EXAMPLE


We can now discuss in more detail the process of disambiguation of (18). I have presented the logical form for (18) above, as (21).

(18) Every school sent the principal to the meeting.

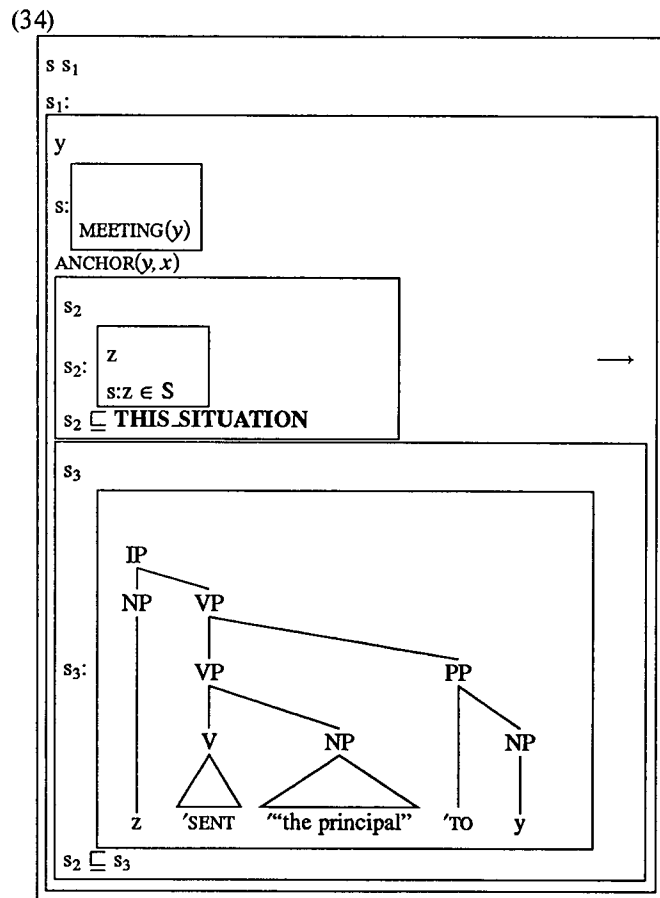
After identifying the situation descriptions, various interpretation processes take place, like those performing definite description interpretation described in [30]. These processes generate hypotheses about the anchoring of resource situations. Without entering into details, I assume that the context for (18) is provided by (32), that introduces into the

model the situation description in (33), containing a group of schools and a meeting.

(32) There was a meeting of the schools in the district.



Given this context, the discourse interpretation processes identify s as the resource situation for the NPs “every school” and “the meeting.” However, no unique principal can be identified in s . The activation of the model construction rules for universal quantification and definite descriptions results in the partial model in (34), in which s_1 and s_2 have been identified:



The model construction rule applied to the universal “every school” introduces a complex condition $K_1 \rightarrow K_2$ as usual, but both the restriction and the nuclear scope include situation descriptions. The situation description in the restriction, s_2 , is a subsituation of the situation at which the restriction is evaluated (denoted by the indexical constant **THIS_SITUATION**). The situation description in the nuclear scope, s_3 , is an extension of s_2 .

Now that a situation description for the resource situation of the universal and a discourse marker for the school have been introduced (s_2 and z , respectively), the rules for resolving the parametric component x of the interpretation of “the principal” can apply. The result is that z is chosen as antecedent of x , and s_2 is chosen as the resource situation for “the principal.” The model construction rule updates s_3 accordingly; the resulting event structure is equivalent to the interpretation of (21) specified by (20).

ACCOUNTING FOR THE DISAMBIGUATION DATA

I briefly return here the disambiguation principles, to show how the proposal just presented accounts for them. First of all, I’ll note that, under simple assumptions about the mapping between grammatical functions and theta-roles, there is a striking resemblance between the grammatical function hierarchy proposed by Ioup and the thematic hierarchy proposed by Jackendoff to account for facts about passives and reflexives. The facts accounted for by the grammatical function hierarchy principle can also be explained if we assume that the description of an event is constructed by identifying the filler of each thematic role in the order specified by Jackendoff’s thematic hierarchy.

Consider now the case of the other disambiguation factor proposed by Ioup, the lexically encoded preference for certain operators to take wide scope. Definite descriptions are the paradigmatic case of an operator that tends to take wide scope. This preference can be explained in terms of the model construction hypothesis as follows. The choice of a resource situation for definite descriptions is restricted by the constraint that this resource situation be either shared among the conversational participants, or related to shared knowledge by shared relations [12, 4]. In our dialogues, for example, definite descriptions are usually interpreted with respect to the ‘situation’ corresponding to the current visual scene, which is independent from other situations. It follows that a definite description will be assigned narrow scope relative to another operator only if (i) the resource situation of the definite is perceived to depend on this other resource situation, and (ii) this dependency relation is known to be shared.

As for the tendency for NPs in topic to take wide scope, an element of a sentence is said to be in topic if it is considered to be part of the background information on which the new information in the sentence depends. As the interpretation of the ‘new’ information in the sentence depends on the background information, it is plausible to assume that, in constructing a model for the sentence, the listener begins by applying the model construction rules for the operators perceived to be in topic (or explicitly marked as being in topic, in the case of Japanese). The interpretation of the operators not in topic, when determined at all, will depend on the interpretation of the operators in topic, resulting in the dependency relations between the related situations that I have assumed to be the way scope is represented.

Finally, I’ll note that, in the absence of contextual clues, whether a completely disambiguated event structure is actu-

ally constructed depends on how strong the model construction rules are supposed to be; it's perfectly possible that the activation of these rules is controlled by additional factors, such as the specific needs of a task to be performed.

ACKNOWLEDGMENTS

I wish to thank my advisor Len Schubert and James Allen, Howard Kurtzman, Peter Lasersohn, and Uwe Reyle for several suggestions, technical help, and constructive criticism. This work was supported by the US Air Force - Rome Laboratory Research Contract no. F30602-91-C-0010.

References

- [1] J. F. Allen and L.K. Schubert. The TRAINS project. TRAINS Technical Note 91-1, University of Rochester, Department of Computer Science, 1991.
- [2] H. Alshawi and R. Crouch. Monotonic semantic interpretation. In *Proc. 30th. ACL*, pages 32–39, University of Delaware, 1992.
- [3] J. Barwise and J. Perry. *Situations and Attitudes*. The MIT Press, 1983.
- [4] H. H. Clark and C. R. Marshall. Definite reference and mutual knowledge. In *Elements of Discourse Understanding*. Cambridge University Press, 1981.
- [5] Robin Cooper. *Quantification and Syntactic Theory*. D. Reidel, 1983.
- [6] W. Croft. *Syntactic Categories and Grammatical Relations: The cognitive organization of information*. University of Chicago Press, 1991.
- [7] Kees van Deemter. *On the Composition of Meaning*. PhD thesis, University of Amsterdam, 1991.
- [8] K. Devlin. *Logic and Information*. Cambridge University Press, 1991.
- [9] J.E. Fenstad, P.K. Halvorsen, T. Langholm, and J. van Benthem. *Situations, Language and Logic*. D.Reidel, 1987.
- [10] A. Garnham. On-line construction of representations of the content of texts. Reproduced by Indiana University Linguistics Club, 1982.
- [11] G. Gazdar, E. Klein, G. Pullum, and I. Sag. *Generalized Phrase Structure Grammar*. Blackwell, 1985.
- [12] J. A. Hawkins. *Definiteness and Indefiniteness*. Croom Helm, 1978.
- [13] I. Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts at Amherst, 1982.
- [14] G. G. Hendrix. Semantic aspects of translation. In D. Walker, editor, *Understanding Spoken Language*, pages 193–226. Elsevier, 1978.
- [15] J. Hobbs. An improper treatment of quantification in ordinary English. In *Proc. ACL-83*, pages 57–63, Cambridge, MA, June 1983.
- [16] J. R. Hobbs and S. M. Shieber. An algorithm for generating quantifier scopings. *Computational Linguistics*, 13(1-2):47–63, January-June 1987.
- [17] Sven Hurum. Handling scope ambiguities using domain-independent heuristics. Technical Report TR 88-12, University of Alberta, June 1988.
- [18] C. H. Hwang. *A Logical Approach to Narrative Understanding*. PhD thesis, University of Alberta, 1992.
- [19] C. H. Hwang and L. K. Schubert. Episodic logic: A situational logic for natural language processing. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters, editors, *Situation Theory and its Applications*, v.3. CSLI, 1993. To appear.
- [20] Georgette Ioup. Some universals for quantifier scope. In J. Kimball, editor, *Syntax and Semantics 4*, pages 37–58. Academic Press, New York, 1975.
- [21] R. Jackendoff. *Semantic Interpretation in Generative Grammar*. MIT Press, 1972.
- [22] P. Johnson-Laird. *Mental Models*. Harvard University Press, 1983.
- [23] H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam, 1981.
- [24] H. Kamp and U. Reyle. From discourse to logic. To appear, 1993.
- [25] R. Kempson and A. Cormack. Ambiguity and quantification. *Linguistics and Philosophy*, 4(2):259–310, 1981.
- [26] H. S. Kurtzman and M. C. MacDonald. Resolution of quantifier scope ambiguities. To appear., April 1992.
- [27] G. Lakoff. Semantic interpretation in generative grammar. In D. A. Steinberg and L. A. Jakobovits, editors, *Semantics: An interdisciplinary reader in philosophy, linguistics, anthropology, and psychology*. Cambridge University Press, 1971.
- [28] R. May. *The Grammar of Quantification*. PhD thesis, MIT, 1977.
- [29] M. Poesio. *Assigning a Scope to Operators in Dialogues*. PhD thesis, University of Rochester, Department of Computer Science, 1993.
- [30] M. Poesio. A situation-theoretic formalization of definite description interpretation in plan elaboration dialogues. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters, editors, *Situations Theory and its Applications*, vol.3, chapter 12, pages 343–378. CSLI, 1993. To appear.
- [31] Massimo Poesio. Relational semantics and scope ambiguity. In J. Barwise, J. M. Gawron, G. Plotkin, and S. Tutiya, editors, *Situation Semantics and its Applications*, vol.2, chapter 20, pages 469–497. CSLI, 1991.
- [32] T. Reinhart. *Anaphora and semantic interpretation*. Croom Helm, 1983.
- [33] U. Reyle. Dealing with ambiguities by underspecification: Construction, representation and deduction. *Journal of Semantics*, 3, 1993.
- [34] T. Stowell. *Origins of Phrase Structure*. PhD thesis, MIT, 1981.
- [35] Kurt A. VanLehn. Determining the scope of English quantifiers. Technical Report AI-TR-483, Artificial Intelligence Laboratory, MIT, Cambridge, MA, 1978.