Semantic Ambiguity and Perceived Ambiguity[*]

Massimo Poesio
University of Edinburgh
Centre for Cognitive Science
poesio@cogsci.ed.ac.uk

Abstract

I explore some of the issues that arise when trying to establish a connection between the underspecification hypothesis pursued in the NLP literature and work on ambiguity in semantics and in the psychological literature. A theory of underspecification is developed 'from the first principles', i.e., starting from a definition of what it means for a sentence to be semantically ambiguous and from what we know about the way humans deal with ambiguity. An underspecified language is specified as the translation language of a grammar covering sentences that display three classes of semantic ambiguity: lexical ambiguity, scopal ambiguity, and referential ambiguity. The expressions of this language denote sets of senses. A formalization of defeasible reasoning with underspecified representations is presented, based on Default Logic. Some issues to be confronted by such a formalization are discussed.

1 The Combinatorial Explosion Puzzle

The alternative syntactic readings of a sentence such as (1) probably number in the hundreds, whereas sentences such as (2) would have hundreds of thousands scopally distinct readings if all permutations of scope-taking sentence constituents were considered admissible readings. Yet, human beings appear able to deal with these sentences effortlessly.

(1) We should move the engine at Avon, engine E1, to Dansville to pick up the boxcar there, then move it from Dansville to Corning, load some oranges, and then move it on to Bath.

(2) A politician can fool most voters on most issues most of the time, but no politician can fool all voters on every single issue all of the time.

[*] This paper will appear in K. van Deemter and S. Peters (eds), *Semantic Ambiguity and Underspecification*, CSLI.

This COMBINATORIAL EXPLOSION PUZZLE is one of the most fundamental questions to be addressed by a theory of language processing, and a substantial problem for developers of Natural Language Processing (NLP) systems. NLP systems which have to perform non-linguistic actions like booking a flight in response to an user's utterance must arrive at the preferred interpretation of their input in the context of the conversation, if one exists; otherwise, they must realize that their input is ambiguous and request a clarification. Examples such as (1) and (2) indicate that such systems cannot adopt the sentence processing strategy of generating all the readings of an ambiguous sentence and choosing one of them, because there are too many such readings.[1] In order to develop such systems, a theory of ambiguity processing is needed that is consistent both with linguistic facts and with what is known about the way humans disambiguate.

Work on UNDERSPECIFIED REPRESENTATIONS such as (Alshawi , 1992; Poesio , 1991; Reyle , 1993)[2] differs from other work on discourse interpretation (Charniak and Goldman , 1988; Hobbs et al. , 1990; Pereira and Pollack , 1991; Dalrymple et al. , 1991; Hwang and Schubert , 1993; Kamp and Reyle , 1993) because it is explicitly motivated by the Combinatorial Explosion Puzzle, and aims at a unified account of all interpretation processes, including those that occur before the scope of all operators has been determined. The work on underspecified representations holds the promise of yielding a better account of the way interpretive processes such as scope disambiguation and reference resolution affect each other.

The existing theories of underspecification, however, have been motivated almost exclusively by computational considerations. For example, the semantics assigned to underspecified representations is designed so as to support those inferences that are deemed useful for an economical approach to disambiguation, rather than being motivated by an analysis of the phenomenon of ambiguity. In this paper I explore some of the issues that arise when trying to establish a connection between work on underspecification and, on the one side, work on ambiguity in semantics; on the other side, work on ambiguity in the psychological literature. A theory of underspecification is developed 'from the first principles', i.e., starting from a definition of what it means for a sentence to be semantically ambiguous and from what we know about the way humans deal with ambiguity. The goal is to arrive at a linguistically and cognitively plausible theory of ambiguity and underspecification that, in addition to computational gains, may provide a better understanding of how humans process language.[3]

Many of the issues discussed in this paper arose from work on the TRAINS project at the University of Rochester, in which the issues of language compre-

[1] With current technology, the real problem is not so much the size of the search space, but how to choose among these interpretations, most of which are plausible.

[2] An early attempt at a model of discourse interpretation of this kind was made by Hobbs, e.g., (Hobbs , 1983).

[3] A recent example of work also attempting to exploit the properties of underspecified representations to address linguistic questions is (Crouch , 1995).

hension, planning, and reasoning encountered in task-oriented natural language conversations are studied (Allen et al. , 1995). The theory of ambiguity proposed in this paper is the basis for the implemented surface discourse interpretation system SAD-93, used in the TRAINS-93 demo system. SAD-93 is described in (Poesio , 1994).

2 Ambiguity in Natural Language

2.1 Ambiguity and Grammar

2.1.1 Characterizing Ambiguity

The dictionary definitions of the terms AMBIGUITY and AMBIGUOUS try to capture the intuition that an expression is ambiguous if 'it has multiple meanings'. An example are the following entries, from Webster's:

- am.big.u.ous |am-'big-y*-w*s| aj [L ambiguus, fr. ambigere to wander about, fr. ambi- + agere to drive – more at agent]

  ...

  2: capable of being understood in two or more possible senses : equivocal - am.big.u.ous.ly av

- am.bi.gu.ity |am-b*-'gyu:-*t-e-| n : the quality or state of being ambiguous in meaning; also : an ambiguous word or expression

A more precise characterization of the notion of ambiguity is required to differentiate ambiguity from VAGUENESS or INDETERMINACY, for example, or to clarify notions such as HOMONIMY and POLYSEMY (see below).[4]

An early attempt at making the notion of ambiguity more precise was presented in (Lakoff , 1970). Lakoff proposed linguistic tests that could be used to tell whether a sentence was ambiguous or not.[5] (Lakoff's tests were meant to provide a way for distinguishing ambiguous sentences from indeterminate ones.)

[4] A sentence is INDETERMINATE, or UNSPECIFIED, if it is definitely true or false, but it could be made more specific. Zwicky and Sadock (1975) bring the example of the sentence *My sister is the Ruritanian secretary of state*, which is indeterminate as to whether "...my sister is older or younger than I am, whether she acceeded to her post recently or some time ago, whether the post is hers by birth or by merit," and so forth. The point is that these additional facts do not affect the truth value of the sentence. It hardly needs to be pointed out that just about every sentence is indeterminate / unspecific in some respects. I will use the term indeterminate for these sentences, and reserve the term UNDERSPECIFIED for sentences which may have different truth values depending on the way the facts are 'filled in' (see below).

[5] An example of these tests are the IDENTITY TESTS, one of which is the CONJUNCTION test. The (presumed) ambiguity of sentences such as *They say her duck* derives from the fact that the phrase *her duck* can either be a NP or a bare infinitival complement. The sentence *They saw her swallow* should have a similar ambiguity. If these two sentences were really ambiguous as claimed, a sentence such as *They saw her swallow and her duck* should only have two readings instead of four since conjunction requires its two arguments to be of the

Zwicky and Sadock (1975) showed, however, that such tests do not result in an unambiguous classification of sentences, and that a formal characterization of the concepts of ambiguity and indeterminacy was required even to understand what these 'ambiguity tests' really test.[6]

2.1.2 Meaning, Sense, and Ambiguity

One problem to be tackled in attempting to make precise the definition of ambiguity is to say what 'meanings' and 'senses' are. In modern semantic theory, the meaning assigned to an expression by a grammar is a function from contexts (or DISCOURSE SITUATIONS) to SENSES. Roughly speaking, the discourse situation provides a value for all context-dependent aspects of the sentence; the sense of a sentence (what we get once we resolve its context-dependent aspects) tells us under which circumstances in the world the sentence is true or false.[7]

Not all notions of 'sense' employed in the literature can serve as the basis for a definition of ambiguity. For example, of the various notions of PROPOSITION (the sense of sentences), the simplest is the one according to which propositions are truth values. But if we were to use this notion of sense, the sentence *Kermit croaked*, ambiguous between a reading in which Kermit utters a frog-like sound and a reading in which he dies, would be classified as unambiguous with respect to all models in which Kermit has both the property of dying and the property of producing a frog-like sound, or he (it) has neither property. In other words, in providing a definition of ambiguity we find the same need for a fine-grained notion of sense that has been observed in connection with the semantics of attitude reports.[8] A model-theoretic definition of ambiguity requires a finer-grained notion of proposition than simply truth values. In most recent semantic theories, senses are intensional objects; the simplest way of achieving intensionality is to use functions from possible worlds or situations to referents as one finds in Montague Grammar, where, for example, propositions are functions from possible worlds to truth values. This simple form of intensionality will be sufficient for the purposes of the present paper.[9]

same type, and therefore the 'crossed' readings should not be available. This is indeed the case. On the other hand, an indeterminate sentence such as *My sister is the Ruritanian secretary of state* maintains all of its indeterminateness once conjoined, as in *My sister is the Ruritanian secretary of state and a prominent composer.*

[6] For example, Zwicky and Sadock observed that such tests can only identify POLAR ambiguities (such as the ambiguity of GAME between two entirely independent readings), but not PRIVATIVE ambiguities, like those displayed by a term like DOG which can be used both to indicate generic individual of the Canis species and a male element of it.

[7] For a discussion of these assumptions, see (Kaplan , 1977), (Barwise and Perry , 1983), or chapter 2 of (Pinkal , 1995).

[8] In the case of attitude reports, the problem is to make sure that if *John is tall* and *John is stupid* are both true in a model, *Bill believes that John is tall* does not entail *Bill believes that John is stupid* in that model, assuming that propositions are the semantic correlate of sentential complements.

[9] More complex notions of propositions have been introduced in the literature on proposi-

2.1.3 A Semantic Theory of Ambiguity

The notions of 'meaning' and 'sense' just discussed are the starting point for the semantic account of the notion of ambiguity and its relation with vagueness developed by Pinkal ((1985), translated as (Pinkal , 1995)). Pinkal introduces the notion of INDEFINITENESS to subsume both ambiguity and vagueness. He defines indefiniteness as follows:

Definition 2.1 *A sentence is* SEMANTICALLY INDEFINITE *if and only if in certain situations, despite sufficient knowledge of the relevant facts, neither "true" nor "false" can be clearly assigned as its truth value.*

Pinkal formalizes the notion of indefiniteness in terms of PRECISIFICATION.[10] According to Pinkal, a linguistic expression is semantically indefinite if it has the potential for being made *precise* in distinct ways. For example, the sentence *The Santa Maria is a fast ship* containing the degree adjective *fast* can be 'made precise' (and assigned a definite truth value) either with respect to a context in which 'fast' is interpreted as 'fast for a modern ship', in which case the sentence is false; or with respect to a context in which 'fast' is interpreted as 'fast for a ship of her age', in which case the sentence can be true or false, depending on the class of comparison. Let p and q be two propositions. Proposition p is MORE PRECISE THAN q iff (i) p is true (false) under all states of the world under which q is true (false), and (ii) p is true or false under certain circumstances under which q is indefinite. The idea of PRECISIFICATION is defined as follows:

Definition 2.2 *Expression α in context* c CAN BE PRECISIFIED *to* s *if and only if (i)* s *is a sense that α can assume according to its meaning; and (ii)* s *is more precise than the sense of α in* c.

The connection between indefiniteness and precisification is provided by the following PRECISIFICATION PRINCIPLE:

Precisification Principle : A sentence is of indefinite truth value in a context if and only if it can be precisified alternatively to "true" or to "false".

which Pinkal also reformulates as follows:

Extended Precisification Principle An expression is semantically indefinite in a context iff it can assume different senses in that context.

Pinkal does not equate ambiguity with vagueness. His theory includes, in addition to the notion of precisification, additional criteria to differentiate different

tional attitudes, such as those used in Situation Semantics (Barwise and Cooper , 1993) or Property Theory (Turner , 1992).

[10] A treatment of ambiguity and vagueness in terms of precisifications was proposed early on in (Fine , 1975).

forms of ambiguity, as well as differentiating 'pure' ambiguity from 'pure' vagueness. The intuition he is trying to capture is that "...whether an expression is ambiguous or only vague is a question that cannot be cleared once and for all. Indefiniteness is perceived as ambiguity when alternative precisifications are predominant, as vagueness when an unstructured continuum presents itself:"

Ambiguity (Pinkal) : If the precisification spectrum of an expression is perceived as discrete, we may call it ambiguous; if it is perceived as continuous, we may call it vague.

Pinkal identifies two fundamental types of ambiguity, according to whether an expression has, or does not have, a 'wider' sense that could be taken as most 'basic'. For example, BALL does not have a wide sense of 'round object and dancing party', whereas AMERICAN may either mean 'person from the US' or 'person from the American continent'. He classifies expressions like AMERICAN which have a wider sense as having a MULTIPLICITY OF USE, whereas expressions such as BALL or GREEN which do require precisification are called NARROWLY AMBIGUOUS. The cases of ambiguity in the narrow sense are further distinguished in two classes, depending on whether they are subject to the PRECISIFICATION IMPERATIVE. Although the two interpretations of GREEN are distinct, it is possible of an object to be both green in the 'ripe' sense and green in the 'color' sense: for example, a green apricot. An object cannot, however, be a 'band' both in the musical group sense and in the piece of tape sense. Pinkal proposes that polysemous expressions behave like GREEN, and calls all of these expressions P-TYPE AMBIGUOUS; expressions like BAND, however, are true homonyms, and therefore he calls them H-TYPE AMBIGUOUS. These latter are defined as follows:

Precisification Imperative : An expression is H-type ambiguous iff its base level is inadmissible, i.e., if it requires precisification.

For my purposes, it's not particularly important whether the difference between homonimy and polysemy is completely captured by the Precisification Imperative; what is important is the claim that H-type ambiguous expressions need precisification, and furthermore, that the Precisification Imperative "is a second order phenomenon ...that lies beyond the scope of a strictly truth-conditional approach." ((Pinkal , 1995), p. 86–87). I will provide below independent reasons for including a formalization of reasoning in context in a treatment of ambiguity, and I will argue that such formalization provides the necessary tools to express the Precisification Imperative.

To summarize, a sentence is H-type ambiguous iff the grammar assigns to it distinct precisifications (senses) in a given discourse situation, and if the 'base level' of the expression requires precisification. Thus, the sentence *Kermit croaked* is considered ambiguous since in the 'empty context' that provides all the senses of the expression according to the grammar G for English, that sentence has two senses: the proposition that attributes to Kermit the property of producing the sound that frogs produce, and the proposition that attributes to Kermit the property of dying. (I am assuming here that terms like *Kermit* refer unambiguously.) On the other hand, the sentence *Kermit kissed Miss Piggy* would be considered unambiguous with respect to the same context.

Although Pinkal is only concerned with lexical ambiguity, the precisification approach can also be used to classify as ambiguous sentences which have more than one structural analysis (like the sentence *They saw her duck*) or are scopally ambiguous (cfr. the sentence *Everybody didn't leave*) whenever the grammar assigns to them more than one sense. I will discuss below how Pinkal's system can be extended to scopal and referential ambiguity.[11]

2.1.4 The Disjunction Fallacy

It is important to realize that saying that a sentence is ambiguous in a context if it has distinct precisifications is not the same as saying that an ambiguous sentence is equivalent to the disjunction of its distinct precisifications. Intuitively, in uttering S, whose two precisifications are the propositions P and Q, a speaker may have meant P or she may have meant Q, but the following does not hold:

$$[[\text{A } \mathbf{means} \text{ that P}] \vee [\text{A } \mathbf{means} \text{ that Q}]] \equiv [\text{A } \mathbf{means} \text{ that } [\text{P} \vee \text{Q}]]$$

To treat an ambiguous sentence in such a way would be tantamount to propose that an ambiguous sentence has a single sense in any given discourse situation, namely, the proposition that is true at a situation if either of the distinct

[11] Another way of providing a precise definition of ambiguity has been explored in the literature, that we might call *structural* or *syntactic*. An example of structural definition is the following, from Gillon ((Gillon , 1990), p. 400):

Ambiguity (Gillon) : *An expression is ambiguous iff the expression can accomodate more than one structural analysis.*

This definition relies on the assumption made in transformational theories of grammar such as Government and Binding theory (Chomsky , 1981; Haegeman , 1991), that each interpretation of a sentence is a quadruple ⟨PF,DS,SS,LF⟩, each of whose elements is a structured object: PF characterizes the phonetic interpretation, DS the predicate/argument composition of the sentence, SS its surface syntactic analysis, and LF its logical analysis. A sentence (string) is ambiguous iff it can be characterized by distinct quadruples, and this may happen not only for phonetical or syntactic reasons, but also for semantical reasons, since one of the components of a structural characterization of a sentence, the 'LF,' encodes the semantic interpretation of the sentence.

The problem with this definition is that a purely syntactic analysis of meaning introduces spurious distinctions: for example, unless something is said about invariance under renaming of variables, one would predict from a structural definition that even a sentence with a single quantifier such as *Every man left* is infinitely ambiguous, because all structures of the form $[_S\ [_{NP}$ every x man$][_x$ left$]]$, for any choice of the variable, are appropriate (and distinct) LF constituents of a sentence's interpretation. Another case of spurious ambiguity is discussed below. A semantic characterization of ambiguity avoids this problem.

interpretations of the sentence is true at that situation; but according to the definition above, an ambiguous sentence is one which has more than one sense at a discourse situation. For example, according to the definition of ambiguity discussed above, the listener of an utterance of *They saw her duck* could either interpret the speaker as saying that the contextually determined set of individuals denoted by the pronoun *they* saw a contextually specified female person lowering herself, or as saying that that set of individuals saw the pet waterfowl of that female person. According to the disjunction theory, instead, the listener would attribute to the speaker of that sentence a single meaning, albeit a disjunctive one; namely, that it was either the case that *they* saw a contextually specified female person lowering herself, or it was the case that *they* saw the pet waterfowl of that female person.[12]

I will refer to the idea that a semantically ambiguous sentence denotes the disjunction of its alternative interpretations as the DISJUNCTION FALLACY. The disjunction fallacy can be found in the literature in two forms. Its 'purest' form is the hypothesis that the interpretation process literally involves generating all of the senses of an expression and putting them together in a disjunction. In this form, the disjunction 'theory' is not simply counterintuitive; it doesn't explain the combinatorial explosion puzzle at all. As far as I know, this 'explicit' form of the theory has only been discussed jokingly. One can find in the literature, however, an 'implicit' form of the disjunction theory, in theories of underspecification that assign to underspecified expressions a semantics that makes them equivalent to the disjunction of their readings. One such proposal is (Poesio , 1991); the semantics of UDRSs is also disjunctive (Reyle , 1993).

2.2 The Role of Syntactic and Semantic Constraints

Although the number of logical form permutations that one can obtain for a particular sentence by, e.g., considering all the permutations of its operators may be rather large, constraints of a syntactic and/or semantic nature drastically reduce this number.

In the case of scopal ambiguity, for example, permutations may not correspond to actual readings for at least three reasons. First of all, some of these permutations result in logical expressions that are either ill-formed or contradictory,as noted by, e.g., Hobbs an Shieber (1987). For example, (3a), in the interpretation in which the pronoun *he* is anaphoric on the NP *every man,* does not have a reading in which the NP *the woman he$_i$ married* outscopes the NP *every man$_i$*. There is no well-formed logical expression that may represent this reading. Hobbs and Shieber point out that this constraint also prevents a quantifier to scope between a noun and its complement: for example, *a meeting* may

[12] A cute example of the problems with the theory is presented by (Stallard , 1987). If ambiguous sentences were to denote the disjunction of their readings, then the answer to the question *Does the butcher have kidneys?* should always be 'yes'.

not scope inside *most* but outside *each* in (3b).[13]

(3) a. Every man$_i$ loves the woman he$_i$ married.

b. Most people on each committee attended a meeting.

Another reason why the number of actual readings of a sentence is much smaller than the number of permutations of its operators is that two distinct permutations may correspond to semantically equivalent readings. For example, (4) has only one reading, even though (at least) two equivalent logical expressions can be obtained as the translation of the sentence.[14]

(4) A student saw a dog.

Finally, the readings corresponding to certain permutations may be unavailable because of syntactic constraints. Much work on uncovering readings that are absent due to constraints on syntactic trasformations and/or conditions on syntactic levels of representation has been done in the generative tradition (May , 1985).

Some of the constraints proposed in this literature have been proved to yield quite robust predictions. Perhaps the best known example of syntactic constraint is the observation that a quantifier cannot take scope outside the clause in which it appears. The observation that clauses serve as 'scope islands' goes back at least to Rodman (1976), but was discussed most extensively by May (1985); the constraint was called SCOPE CONSTRAINT by Heim (1982). The Scope Constraint is exemplified by the contrast in (5): whereas (5a) has a reading in which *every department* is allowed to take wide scope over *a student,* this reading is not available for (5b), even though arguably *from every department* and *who was from every department* have the same denotation.

(5) a. A student from every department was at the party.

b. A student who was from every department was at the party.

Although syntactic and semantic constraints do not rule out all possible readings—for example, a sentence like *They saw her duck* still has more than one interpretation under all of these theories—a theory of disambiguation must be such that these constraints can play a role.

2.3 Perceived Ambiguity

As noted by Hirst (1987), the discussions of ambiguity processing in the NLP literature tend to ignore the fact that humans are aware that sentences can be ambiguous, and that they can exploit the ambiguity of sentences for rhetorical effect. Raskin, for example, claims (1985) that humor crucially relies on ambiguity. He discusses examples such as the following (p. 25-26):

[13] Pereira (1990) argues that this constraint is best formulated as a condition on semantic derivations rather than as a condition on the syntax of logical expressions.

[14] This is one of the reasons for preferring a semantic account of ambiguity to a syntactic account which makes ambiguity depend on the existence of two distinct logical forms.

(6) The first thing that strikes a stranger in New York is a big car.

The joke relies on two assumptions about human processing: first, that the clause *the first thing that strikes a stranger in New York* gets interpreted before the end of the sentence, with *strikes* receiving the 'surprise' interpretation; and second, that the reader, upon reading *is a big car*, will go back, produce a second interpretation, and entertain both interpretations simultaneously. The joke could not be understood unless the hearer were able to entertain the two interpretations of the sentence simultaneously. These jokes can exploit other forms of ambiguity, e.g., scopal ambiguity, as in *Statistics show that every 11 seconds a man is mugged here in New York City. We are here today to interview him.*

The reader's ability to entertain more than one interpretation simultaneously is exploited in poetry, as well (Su , 1994). The linguistic articles discussing ambiguity are another literary form that exploits this possibility. Examples such as *They saw her duck* are a clear case of deliberate ambiguity; the whole point of these examples is to show that a sentence can have more than one interpretation. The writer relies on the reader being able to entertain more than one interpretation at once.

The opposite is true, as well: when clarity is a goal, writers and speakers tend to construct their sentences in such a way as to avoid ambiguity. Thus, most sentences one runs across in scientific texts or in transcripts of task-oriented conversations have a clearly preferred interpretation. This interpretation is sometimes suggested by the context, sometimes by means of DISAMBIGUATION MARKERS—expressions such as *each*, *a different*, or *the same* that suggest which interpretation is preferred. Thus, a writer will use sentences such as *Every kid climbed the same tree*, rather than *Every kid climbed a tree*, when we/she wants to make sure that the reader arrives at the interpretation in which there is a single tree.[15]

I will call the situation in which a listener arrives at more than one interpretation for an utterance PERCEIVED AMBIGUITY. A situation in which B perceives an utterance as ambiguous may result in B's appreciating the joke, the poetic phrase, or the point of the linguistic example; if the ambiguity is not perceived as intended, B may say saying something like *This is not very clear,* or perhaps *This sentence is ambiguous.* This situation can be informally characterized as follows:

Definition 2.3 *An utterance U by conversational participant A addressing conversational participant B in a discourse situation D is* PERCEIVED AS AMBIGUOUS IN D *by B if B's processing of U in D results in B obtaining distinct interpretations for U.*

[15] Such sentences are used, for example, to get a 'baseline' interpretation in psychological work on ambiguity.

The phenomenon of deliberate ambiguity suggests that the solution to the Combinatorial Explosion Puzzle cannot be that humans either generate only one interpretation at a time by using some clever heuristics, or do not generate any interpretation at all. Humans entertain more than one interpretation at a time, and they may not be able to choose one among them. This conclusion is also supported by psychological results. There is evidence, for example, that during both lexical processing and syntactic processing several hypotheses are generated in parallel, and only later filtered on the basis of contextual information (Swinney , 1979; Crain and Steedman , 1985; Kurtzman , 1985; Schubert , 1986; Gibson , 1991).[16] Kurtzman and MacDonald (1993) suggest a similar model for scope disambiguation. As far as reference interpretation is concerned, there is some evidence that all pragmatically available referents become active before a referent is identified (see, e.g., (Spivey-Knowlton et al. , 1994)).

These facts are consistent with the view of discourse interpretation taken in Artificial Intelligence, in which processes such as reference resolution or lexical disambiguation are modeled in terms of defeasible inference, which may result in alternative hypotheses. Examples include the theories of the effects of semantic priming on lexical disambiguation, as formalized, e.g., in Hirst's ABSITY system (Hirst , 1987) or, more recently, in statistically based terms; the theories about the effects of local focusing on the choice of pronoun antecedents such as (Grosz et al. , 1983); the work on temporal interpretation by Asher, Lascarides, and Oberlander (see, e.g., (Lascarides et al. , 1992)); and work on scopal disambiguation such as (Kurtzman and MacDonald , 1993; Poesio , 1994).

2.4 Semantic Ambiguity versus Perceived Ambiguity

A preliminary and, I hope, uncontroversial conclusion I intend to draw from the discussion on deliberate ambiguity and ambiguity processing is that a theory of ambiguity that aims at explaining the Combinatorial Explosion Puzzle needs to be concerned both with the interpretation that the grammar assigns to a sentence—i.e., what it means for a sentence to be semantically ambiguous—and with the process by which interpretations are generated, i.e., with what it means for an utterance to be perceived as ambiguous. On the one hand, the theory must explain why the disambiguation process will not generate all semantically available interpretations; on the other hand, it must predict that more than one interpretation will be generated. This conclusion is the central idea of this paper, indeed, what gives the paper its title. The inclusion of a theory of disambiguation will also remedy one of the omissions in Pinkal's theory, namely, how to formalize the Precisification Imperative.

The discussion of perceived ambiguity supports a stronger claim, namely, that semantic ambiguity and perceived ambiguity are distinct notions, in the

[16] The results about lexical disambiguation are fairly well established, but there is some controversy about syntactic processing. A constrasting view on syntactic disambiguation is discussed in (Frazier and Fodor , 1978).

sense that whereas a model of semantic ambiguity has to express the truth-conditional properties of an expression, the reasoning processes involved in disambiguation, and that may lead to a perceived ambiguity, consist of defeasible inferences that are not supported by the semantics of ambiguous expressions.

The distinction I intend to draw, then, is as follows. Semantic ambiguity is part of the specification of the grammar of a language; most, if not all, sentences are semantically ambiguous, but their ambiguity need not be noticed by listeners, and in fact it is typically discovered only by linguistic research. Perceived ambiguity, on the other hand, is a result of the interpretation process, that is defeasible in nature, and may therefore result in more than one interpretation in cases of miscommunication or when the speaker constructs the context appropriately to serve a rhetorical purpose, as in the puns presented above.

Some readers may wonder why the developer of a NLP system should be concerned with perceived ambiguity, i.e., with generating all of the contextually available interpretations of a sentence. The answer is that certain applications need this information. Consider the following example, again from the TRAINS domain. Say that the user utters *move the engine to Avon*, and say that two different engines have been discussed during the elaboration of the current part of the plan. Clearly, we do not want the system to just come out with a plausible guess about which engine was meant: instead, we want it to recognize the ambiguity and ask for clarification. In general, all systems that engage in conversations with their users need to be able to recognize an ambiguity, to ask for clarifications when necessary rather than guess one possible interpretation, and to make their own output unambiguous. (Of course, the theory of contextual disambiguation must be such that no spurious ambiguities are obtained.)

3 The Underspecification Hypothesis

All theories of semantic interpretation based on Montague's general program as exposed in *Universal Grammar* (Montague , 1970) assume that the grammar of a language $\mathcal{L}$ specifies two homomorphisms: one between syntactic trees and a DISAMBIGUATED LANGUAGE $\mathcal{DL}$, and a second one between the disambiguated language and objects of the model M (the senses). These two homomorphisms can be composed, thus making the intermediate level of the disambiguated language dispensable. The grammar assigns to an ambiguous expression of $\mathcal{L}$ distinct expressions of $\mathcal{DL}$, each of which has a unique interpretation.

A direct implementation of this strategy in an NLP system would require generating all senses of an ambiguous sentence-string, which would be clearly problematic. Many NLP systems, instead, make use of heuristic methods that generate only one interpretation and ignore the alternatives. These heuristics work fairly well fairly often; such systems, however, won't be able to perceive an ambiguity even when it would be helpful to do so. Other systems therefore split the semantic problem of computing all the interpretations of a sentence from the processing problem of generating these interpretations in context, by making use of an intermediate, UNDERSPECIFIED level of representation. One of the earliest examples of underspecified representations is the 'Logical Form' of Schubert and Pelletier (1982). The representation for (7) proposed by Schubert and Pelletier, shown in (8), is a typical example of these underspecified representations: quantifiers are left in place and the referent for the definite description *the tree* is not specified.

(7) Every kid climbed the tree.

(8) [<every kid> climbed <the tree >]

In more recent years, underspecified representations similar to Schubert and Pelletier's have been used by (Hobbs and Shieber , 1987), Fenstad *et al.* (1987), in Allen's textbook (1987) and, most recently, in the Core Language Engine (Alshawi , 1992); the 'uninterpreted conditions' produced during the intermediate steps of the DRT construction algorithm in (Kamp and Reyle , 1993) can be considered underspecified representations as well.

Underspecified representations were originally conceived as a way to solve a problem in system implementation, namely, separating 'context-independent' from 'context dependent' aspects of the interpretation, thus making either part reusable for different applications. Since the motivation was strictly computational, the underspecified representations used in most NLP systems are little more than data structures, in the sense that they do not have a interpretation other than the one provided by the procedures that interpret them. These representations 'encode' the ambiguity of a sentence in the sense that that sentence has the reading r iff that reading can be generated by repeatedly applying 'construction rules' to the underspecified representation.

In recent years, there has been growing interest for the hypothesis that the ability to encode multiple interpretations in an underspecified language may be (part of) the explanation of the Combinatorial Explosion Puzzle. The idea is that humans, as well, make use of an underspecified language that can encode distinct meanings implicitly, and therefore do not need to generate all of these meanings. A semantically ambiguous sentence, therefore, need not cause problems for a human to process, because it is not necessarily *perceived* as ambiguous in the sense discussed in the previous section. I will call this assumption the UNDERSPECIFICATION HYPOTHESIS:

Underspecification Hypothesis : Human beings represent semantic ambiguity implicitly by means of UNDERSPECIFIED REPRESENTATIONS that leave some aspects of interpretation unresolved.

My goal in the rest of the paper is to spell out the Underspecification Hypothesis both as a theory of grammar and as a theory of discourse interpretation. I assume, that is, that the hypothesis is correct, and try to answer questions such as: what kind of language are underspecified representations? what is their semantics? and, what kind of inferences are done with them?

The novel aspect of this work is that the answers I give are based on the discussion of semantic ambiguity and perceived ambiguity in the previous section. I hypothesize that underspecified representations are used by humans as the translation of expressions that are indefinite in the sense of Pinkal, and assign them a semantics that reflects this hypothesis. I assume that the disambiguation process is consists of defeasible inferences, and examine the characteristics of defeasible reasoning with underspecified representations. Although the same position towards disambiguation and defeasible has been adopted in the Core Language Engine, most of the issues I discuss have not been mentioned so far in the discussion on underspecified representations.

In the literature on underspecification, one often finds the argument that providing a semantics to underspecified representations is necessary because disambiguation requires inference, and therefore a 'logic of underspecification' is needed (Poesio , 1991; Reyle , 1993; van Deemter , 1991). However, it is not at all clear whether the process of disambiguation involves much semantically justified reasoning; disambiguation seems to consist mostly of defeasible inferences. It is fair to say that the debate on this issue is very open at the moment, as certified by a number of recent panels on the subject. But whatever the final conclusion on this topic will be, it is clear that under the perspective that the grammar of a language $\mathcal{L}$ is a mapping from elements of $\mathcal{L}$ to underspecified representations, the semantics of these underspecified representations becomes a central aspect of the specification of the grammar. Furthermore, it also becomes clear that the semantics of underspecified representations must be based on an analysis of semantic ambiguity, otherwise we wouldn't even know whether the form of underspecified representation we develop does the job it is supposed to do.

4 An Underspecified Theory of Ambiguity, Part I: Lexical Ambiguity

The simplest way to illustrate my implementation of the Underspecification Hypothesis is to start with lexical ambiguity. I present in this section a theory of grammar which makes use of an underspecified language to encode the 'ambiguity potential' of lexically ambiguous expressions, as well as a simple formalization of lexical disambiguation as defeasible inference over underspecified representations. In the next section I will show how to extend the approach presented here to deal with expressions that exhibit other forms of semantic ambiguity.

I want to emphasize that I start with lexical disambiguation for expository purposes only. Lexical ambiguity is the one case of ambiguity for which a 'generate and test' strategy may well be compatible with the psychological results, therefore the one for which the need for underspecified representations is less clear. Furthermore, I will only discuss cases of lexical ambiguity in the narrow sense, which is perhaps the least interesting case of lexical indefiniteness. Discussing lexical disambiguation, however, is the simplest way to explain how underspecified representations can be given a semantics related to Pinkal's proposals about ambiguity, and how to defeasible reasoning with underspecified representations. In the next sections I will generalize the approach introduced here to cases of ambiguity for which the underspecified approach is much more plausible. Furthermore, at least one theory of lexical disambiguation, Hirst's proposal (1987), makes use of 'Polaroid words' which are essentially underspecified interpretations of lexical items.

4.1 A Lexically Underspecified Grammar

The presentation of a lexically underspecified grammar below is centered on the example of (H-type) lexical ambiguity discussed above, the verb *croak*, which can take two precisifications. Let $\mathcal{L}$ be the language which consists of the single sentence *Kermit croaked.* This sentence is H-type ambiguous because it admits of two precisifications and it is subject to the precisification imperative. A 'Montagovian' grammar MG would map (syntactic analyses of) the sentence into distinct expressions of a 'disambiguated language' $\mathcal{DL}$, each of which denotes a function from discourse situations into intensional objects of the appropriate type (in this case, propositions). A grammar UHG that subscribes to the Underspecification Hypothesis, on the other hand, maps syntactic analyses of expressions of $\mathcal{L}$ into a single expression of a 'lexically underspecified language' $\mathcal{LXUL}$. The semantics of $\mathcal{LXUL}$ is based on the Precisification Principle: expressions of $\mathcal{LXUL}$ denote at each discourse situation a *set* of senses of the type they would be assigned by a Montagovian grammar.[17]

The lexically underspecified language $\mathcal{LXUL}$ has the following ingredients:

Terms : the constant k.

Predicates: $\mathbf{croak}_U$, $\mathbf{croak}_1$, and $\mathbf{croak}_2$.

Atomic Formulas: If t is a term and P is a predicate, then P(t) is a formula.

Formulas: If Φ and Ψ are formulas, then $\neg\Phi$ and $\Phi \wedge \Psi$ are formulas.

Note that in addition to two predicates $\mathbf{croak}_1$ and $\mathbf{croak}_2$, corresponding to the disambiguated senses of *croak*, the language includes an 'underspecified' predicate $\mathbf{croak}_U$. The interpretation function for $\mathcal{LXUL}$, $[\![\alpha]\!]^{\mathrm{M,d}}$, is defined as follows. Let M = $\langle U,F \rangle$ be a model just like the one that would be used for a disambiguated language $\mathcal{DL}$. The interpretation function $[\![\alpha]\!]^{\mathrm{M,d}}$ assigns to an expression α of $\mathcal{LXUL}$ a value with respect to M and a discourse situation d.

[17] The technique of assigning sets of senses as the denotation of sentences dates back at least to Hamblin (1973), who used it to extend Montague's fragment to questions.

- $[\![k]\!]^{\mathrm{M,d}}$ = {f_k}, where f_k : s → the denotation of 'Kermit' in s
- $[\![\mathbf{croak}_1]\!]^{\mathrm{M,d}}$ = {f_1}, where f_1 is a function such that f_1(s) → the set of croaking objects in s
- $[\![\mathbf{croak}_2]\!]^{\mathrm{M,d}}$ = {f_2}, where f_2 is a function such that f_2(s) → the set of dying objects in s}
- $[\![\mathbf{croak}_U]\!]^{\mathrm{M,d}}$ = {f_1, f_2}
- $[\![\mathrm{P(t)}]\!]^{\mathrm{M,d}}$ = {f | for a g ∈ $[\![\mathrm{P}]\!]^{\mathrm{M,d}}$ and an h ∈ $[\![\mathrm{t}]\!]^{\mathrm{M,d}}$ f(s) = g(s)[h(s)] }
- $[\![\Phi \wedge \Psi]\!]^{\mathrm{M,d}}$ = {f | for a g ∈ $[\![\mathrm{P}]\!]^{\mathrm{M,d}}$ and an h ∈ $[\![\mathrm{t}]\!]^{\mathrm{M,d}}$, f(s) = 0 if g(s) = 0 or h(s) = 0; f(s) = 1 if both g(s) = 1 and h(s) = 1 }
- $[\![\neg\Phi]\!]^{\mathrm{M,d}}$ = {f | for g ∈ $[\![\Phi]\!]^{\mathrm{M,d}}$, f(s) = 1 if g(s) = 0; f(s) = 0 if g(s) = 1 }

The language $\mathcal{LXUL}$ has been deliberately kept simple to make it clear that the underspecified languages I propose have two basic properties: (i) the value of an expression at a discourse situation is a set of senses of the type that a sense of that expression would have in a disambiguated language $\mathcal{DL}$; and (ii) expressions can be divided into expressions whose denotation at a discourse situation is a singleton set, such as k or $\mathbf{croak}_1$, and expressions such as $\mathbf{croak}_U$ that denote a non-singleton set. The latter expressions provide the interpretation for ambiguous expressions of $\mathcal{L}$.[18]

The clauses for application and the connectives show how ambiguity 'percolates up' from lexical items. The value of an expression like $\alpha(\beta)$ is obtained by taking the cross-product of the values of α and β, and it includes one function f per distinct pair of functions $\langle\alpha_1,\beta_1\rangle$ in the denotations of α and β. The value assigned by the function f to the situation $\underline{\mathrm{s}}$ is defined by applying a certain operation (in this case, application) to the values assigned to $\underline{\mathrm{s}}$ by the functions α_1 and β_1. Thus, if both the denotation of α and the denotation of β are singleton sets, the denotation of $\alpha(\beta)$ is also a singleton set; otherwise, ambiguity 'multiplies,' as it where.[19] The same 'multiplication' technique is also used to

[18] A lexical item can also be ambiguous in that it may be associated with lexical entries of different syntactic categories: for example, the word *duck* can either be interpreted as a noun or as an verb, as shown by the example *They saw her duck*. I assume a syntactic ambiguity in these cases, i.e., I assume that the grammar would assign two syntactic analyses to the word *duck*, each of which would then get an interpretation in $\mathcal{LXUL}$.

[19] It is perhaps worth emphasizing a difference between the semantics just sketched and virtually all other approaches to underspecification I am aware of. In this proposal, the underspecified language $\mathcal{LXUL}$ does not serve as a 'meta-language' to be given a semantics in terms of the values assigned to the expressions of a 'disambiguated language'; instead, it has a semantics of its own, defined bottom-up much in the way the semantics of $\mathcal{DL}$ would be defined. In other words, the approach just sketched does not rely on the assumption that a 'disambiguated language' can be defined, which, at the light of Pinkal's treatment of indefiniteness, appears to be questionable. For example, for Pinkal an expression is 'purely vague' is no natural precisification exists.

define the denotation of connectives.[20]

The following grammar generates an underspecified representation of *Kermit croaked* by mapping the semantically ambiguous predicate *croaked* into an 'ambiguous' predicate of $\mathcal{LXUL}$ as follows:

- S → NP VP; VP′(NP′)
- NP → Kermit; k
- VP → croaked; $\mathbf{croak}_U$

The underspecified translation of *Kermit croaked* in $\mathcal{LXUL}$, $\mathbf{croak}_U(k)$, denotes a set of two propositions at a situation d: the function that assigns 1 to a situation iff Kermit produced a frog-like sound in that situation, and the function that assigns 1 to a situation iff Kermit died in that situation. This makes the sentence indefinite in Pinkal's sense.[21] By contrast, an indeterminate sentence such as *Kermit is the Ruritanian secretary of state* would have a single sense at a given discourse situation.

Pinkal's Precisification Imperative is an attempt at making more precise the observation that human beings don't seem to have good intuitions concerning what follows from a H-type ambiguous sentence. Even when subjects are able to pass judgments about what follows from an ambiguous sentence, it's arguable that they do not give judgments concerning what follows from the underspecified representation: rather, they first generate one interpretation, then decide what follows from that. The conclusion that I would be inclined to draw is that a relation of semantic entailment capturing human intuitions can only be defined, if at all, between expressions whose interpretation is not subject to the Precisification Imperative. So, although it would be possible to define, for example, a 'strong' notion of entailment as what follows from all senses, this definition would be rather artificial. For this reason I will not attempt to define a notion of entailment between expressions of $\mathcal{LXUL}$; the readers interested in the issue are referred to Pinkal's book and to the discussion in van Deemter's dissertation (1991).

[20] I will ignore in what follows the issue of partiality, e.g., what happens when a conjunct has a value other than 0 or 1.

[21] The denotation assigned to an indefinite sentence by the grammar above is a simplification of the denotation that Pinkal would assign to such a sentence in (Pinkal , 1995), which, in addition the set of senses associated with a natural language expression, would also include a partial order relation of precisification between them. Such an order relation plays an important role in the meaning of vague sentences such as *Kermit is tall* is considered, which has distinct senses depending on the degree of precision with which the discourse situation is specified, but a less important one in the cases of 'narrow sense' ambiguity with which I am concerned here.

4.2 Discourse Interpretation and Perceived Ambiguity

4.2.1 Discourse Interpretation and Defeasible Reasoning

A theory of ambiguity processing solves the Combinatorial Explosion Puzzle if it does not require that all distinct interpretations of a semantically ambiguous sentence are actually generated. A grammar consistent with the Underspecification Hypothesis such as the one just discussed moves us one step towards that goal, since it only imposes the constraint that a single underspecified interpretation be generated.

On the other hand, we can conclude from the discussion of deliberate ambiguity and of the psychological work on ambiguity that a psychologically plausible theory of ambiguity must also predict that more than one interpretation may become available in a given context, although the number of such interpretations will in general be much smaller than the number of possible semantic interpretations.

As discussed above, the view of discourse interpretation that I am going to take is the one typically found in the AI literature, according to which disambiguation involves the generation of (possibly distinct) hypotheses in parallel by means of defeasible inference. This perspective is found, for example, in the work on abductive discourse interpretation by Hobbs and colleagues (Hobbs et al. , 1990), in the work on Bayesian disambiguation by, e.g., Charniak and his students (Charniak and Goldman , 1988) and in the work on DICE and discourse interpretation by Asher, Lascarides, and Oberlander (Lascarides et al. , 1992). Some of the formal models of defeasible reasoning that can be used to formalize the situation in which conflicting hypotheses are generated include Reiter's DEFAULT LOGIC (Reiter , 1980), the abductive model (Hobbs et al. , 1990), Bayesian Nets, and DICE (Lascarides et al. , 1992; Asher and Morreau , 1991).

4.2.2 Lexical Disambiguation Using Defaults

As a model of defeasible reasoning, I adopt Reiter's Default Logic. In default logic, the process that generates defeasible hypotheses is seen as the computation of the EXTENSIONS of a DEFAULT THEORY (D,W) where D is a set of default inference rules and W is a set of formulas. I will formalize discourse interpretation as the process of generating the extensions of the theory (DI,UF), where DI—the DISCOURSE INTERPRETATION PRINCIPLES—are default inference rules, and UF is a set of expressions of an underspecified language like $\mathcal{LXUL}$. Let us ignore for the moment the fact that the formulas in UF are underspecified representations. The Discourse Interpretation Principles formalize the defeasible inferences that take place in discourse interpretation, such as disambiguating inferences. These rules are operations that map a set of wffs that allow of a certain number of interpretations into a new set of wffs with a more restricted number of interpretations. An example of Discourse Interpretation Principle is the following:

$$\textbf{CROAK1-IF-FROG:}\quad \frac{\mathbf{croak}_U(x) \wedge \mathbf{frog}(x) : \mathbf{croak}_1(x)}{\mathbf{croak}_1(x)}$$

This inference rule[22] reads: if the set of wffs UF includes the fact that the object x has the property $\mathbf{croak}_U$ and the property **frog**, and if it is consistent to assume that the interpretation $\mathbf{croak}_1$ of $\mathbf{croak}_U$ was intended, then the inference rule CROAK1-IF-FROG produces a new set of wffs that includes the fact $\mathbf{croak}_1(x)$. The application of CROAK1-IF-HUMAN-LIKE would be blocked by the presence in UF of the wff $\neg\mathbf{croak}_1(k)$. Using Reiter's definition of extension in an 'intuitive' fashion, we can see that the default theory

$$(\{\text{CROAK1-IF-FROG}\}, \{\mathbf{croak}_U(k) \wedge \mathbf{frog}(k)\})$$

has the following (unique) extension:

$$\{\mathbf{croak}_U(k) \wedge \mathbf{frog}(k),\, \mathbf{croak}_1(k)\}$$

The denotation of a set of wffs $\{\Phi_1 \ldots \Phi_n\}$ will be defined as the denotation of the conjunction $\Phi_1 \wedge \ldots \wedge \Phi_n$ of these wffs. I also assume that the empty set of wffs denotes the function TRUE that is true at every situation. With this definition, and under the assumption that each 'unambiguous' interpretation of the word *croak* is incompatible with the others (i.e., under the assumption that $\mathbf{croak}_1(x) \vdash \neg\mathbf{croak}_2(x)$), the extension of (DF,UI) admits of only one denotation, the one under which the denotation of k produced a sound like the one frogs produce.[23]

A default theory always has an extension as long as all defaults are normal,[24] but it may have more than one extension if the set of Discourse Interpretation Principles contains two inference rules that both apply but generate a conflict. Consider, for example, the default theory consisting of a set of discourse interpretation principles DI$'$ that includes, in addition to CROAK1-IF-FROG, a second discourse interpretation principle (let's call it CROAK2-IF-HUMAN-LIKE) stating that the $\mathbf{croak}_2$ interpretation is plausible for human-like beings; and of a set of wffs UF$'$ including the fact that Kermit is a human-like being.

$$\begin{gathered}(\{\text{CROAK1-IF-FROG,CROAK2-IF-HUMAN-LIKE}\},\\ \{\mathbf{croak}_U(k) \wedge \mathbf{frog}(k),\, \mathbf{human\text{-}like}(k)\ \})\end{gathered}$$

this theory would have two extensions:

[22] CROAK1-IF-FROG is an *open* inference rule. Such rules act like inference rule schemas.

[23] See the discussion below.

[24] I.e., of the form $\alpha{:}\beta/\ \beta$.

1. $\{\mathbf{croak}_U(k) \wedge \mathbf{frog}(k), \mathbf{human\text{-}like}(k), \mathbf{croak}_1(k)\}$
2. $\{\mathbf{croak}_U(k) \wedge \mathbf{frog}(k), \mathbf{human\text{-}like}(k), \mathbf{croak}_2(k)\}$

Perceived ambiguity can now be redefined more precisely as the state that obtains when the default theory 'encoding' the listener's discourse interpretation processes has more than one extension; and the cases of deliberate ambiguity discussed in section §2 can be formalized as cases in which the speaker has 'reasoned about the other agent's reasoning,' as it were.

4.2.3 Constraints on Discourse Interpretation and the Anti-Random Hypothesis

Once we start allowing discourse interpretation processes like those just discussed, the Underspecification Hypothesis is not sufficient to explain the Combinatorial Explosion Puzzle anymore. The UH does not rule out a theory of discourse interpretation in which after an underspecified interpretation has been obtained, all possible senses of a sentence are generated. In fact, a lot of NLP systems work this way, as well as interpretation procedures such as Hobbs and Shieber's scoping algorithm (Hobbs and Shieber , 1987). In the framework for discourse interpretation just presented, theories of this kind could be formalized by including discourse interpretation principles that generate all the semantically justified interpretations at random. For the case of lexical disambiguation, for example, we could have a theory that includes the two following inference rules:

CROAK1-AT-RANDOM:

$$\frac{\mathbf{croak}_U(x)\colon\ \mathbf{croak}_1(x)}{\mathbf{croak}_1(x)}$$

CROAK2-AT-RANDOM:

$$\frac{\mathbf{croak}_U(x)\colon\ \mathbf{croak}_2(x)}{\mathbf{croak}_2(x)}$$

A theory of lexical disambiguation of this kind would simply produce all semantically justified interpretations of a sentence, and the Combinatorial Explosion Puzzle would remain a puzzle. To solve the puzzle, a theory of disambiguation must therefore supplement the Underspecification Hypothesis with constraints on discourse interpretation that ensure that only a few extensions are generated.

The constraints need not be the same for all classes of ambiguity. For certain classes of ambiguity, including perhaps lexical ambiguity,[25] the explanation

[25] See however (Hirst , 1987).

may simply be that the disambiguation process is *incremental*, i.e., it takes place as the text is processed word by word or constituent by constituent, and each ambiguity is resolved *locally*; in this way, only a small number of alternative hypotheses have to be considered every time. For other classes of ambiguity, however, such as scopal ambiguity and referential ambiguity, incremental processing does not seem to be the solution,[26] and different constraints must apply. In (Poesio , 1994), the following constraint was proposed:

Anti-Random Hypothesis (Informal) Humans do not randomly generate alternative interpretations of an ambiguous sentence; only those few interpretations are obtained that (i) are consistent with syntactic and semantic constraints and (ii) are suggested by the context.

The Anti-Random Hypothesis should be thought of as a 'meta-constraint' on theories of interpretation: if we intend to account for the Combinatorial Explosion Puzzle, we have to develop theories of interpretation (e.g., theories of parsing, or theories of definite description interpretation) that satisfy this constraint, i.e., in which discourse interpretation principles like CROAK1-AT-RANDOM and CROAK2-AT-RANDOM are not allowed.

In order to illustrate more concretely the difference between theories of discourse interpretation that satisfy the Anti-Random Hypothesis, and theories that do not, let us consider how one could formalize a theory of pronominal interpretation. A 'random' theory of pronoun interpretation would go as follows: first, compute all possible antecedents of the pronoun in the discourse. Then, generate an hypothesis for each of them, stating that the pronoun refers to that antecedent. Finally, rank these hypotheses according to their plausibility. A random hypothesis generation process usually leaves the task of choosing one hypothesis to plan recognition; the problem is that most often, the alternatives are equally plausible.

In contrast, CENTERING THEORY (Grosz et al. , 1983) is an example of non-random pronoun interpretation theory. According to centering theory, each utterance establishes a 'backward looking center' (Cb), and a pronoun is by default interpreted to refer to the Cb. (I am glossing over a number of complexities

[26] In a sentence such as *John gave a present to each child*, for example, the indefinite *a present* takes narrow scope with respect to the quantifier *each child*. The interpretation of the sentence must therefore either remain partially underspecified until the quantifier is processed, or be revised when the quantifier is encountered. Similarly, when processing a sentence such as *John always invites MARY to the movies*, whose preferred interpretation is that whenever John goes to the movies, he invites Mary, the restriction of the adverb of quantification *always* is not encountered until the PP *to the movies* is encountered.

There is little doubt that part of the solution to the Combinatorial Explosion Puzzle is that some forms of ambiguity, at least, are solved locally and incrementally. Garden-path phenomena, for example, are commonly interpreted as providing evidence for this hypothesis (Frazier and Fodor , 1978; Crain and Steedman , 1985; Altmann , 1989), and, as discussed above, similar effects can exploit forms of ambiguity other than syntactic ambiguity: e.g., scopal ambiguity. The examples just discussed suggest however that an incremental account of discourse interpretation, as well, must be supplemented with a theory of underspecification.

here.) Such a theory would generate a single (or a few) hypothesis concerning the antecedent of a pronoun; the other possibilities, although semantically possible, would simply never come up. Examples of theories of definite description interpretation, tense interpretation, the interpretation of modals in discourse, and scope disambiguation that satisfy the Anti-Random Hypothesis are discussed in (Poesio , 1994).

The Anti-Random Hypothesis can be made more formal in the framework for discourse interpretation adopted here by introducing a slightly different syntax for default inference rules, one in which the underspecified condition is syntactically separated from additional contextual requirements such as the requirement in CROAK1-IF-FROG that the object in question be a frog:

CROAK1-IF-FROG

$$\frac{\mathbf{croak}_U(x) : \ \mathbf{frog}(x) : \ \mathbf{croak}_1(x)}{\mathbf{croak}_1(x)}$$

Except for the fact that one of the prerequisite wffs is 'singled out', an inference rule thus rewritten has the same interpretation as one of Reiter's default rules. We can then require the contextual requirements to be non-trivial (i.e., not satisfied in every situation) as follows:

Anti-Random Hypothesis A discourse interpretation theory (DI,UF) is *Anti-Random* iff for all discourse interpretation principles α:β:δ/ γ in DI, β is not satisfied in every situation.

4.2.4 The Condition on Discourse Interpretation

The framework just introduced can also be used to formalize the 'second order' aspects of Pinkal's theory, such as the Precisification Imperative. The Precisification Imperative can be seen as imposing a constraint on the extensions of a discourse interpretation theory, namely, as the requirement that extensions include a 'disambiguating wff' like $\mathbf{croak}_1(k)$ for each H-type ambiguous constituent of the set UF such as $\mathbf{croak}_U(k)$. I will call this constraint CONDITION ON DISCOURSE INTERPRETATION. In first instance, the Condition on Discourse Interpretation might be formulated as follows, for the case of lexical ambiguity:

Condition on Discourse Interpretation (Preliminary): Each extension E of a discourse interpretation theory (DI,UF) must include, for each literal L in UF whose predicate is H-type ambiguous, a distinct DISAMBIGUATING LITERAL, i.e., a literal whose denotation is a single function among those in the denotation of L.

The definition of the Condition on Discourse Interpretation just given is not very general: it depends on the assumption that all cases of H-type ambiguity are originated by predicates. A simpler, and more general, formulation of the Condition on Discourse Interpretation can be obtained by generalizing the format for the discourse interpretation principles once more.

Default inference rules are typically used to augment a set of wffs with additional facts inferred by default: the fact that a particular bird flies, for example. But the purpose of discourse interpretation rules used for disambiguation, like CROAK1-IF-FROG, is to restrict the interpretation by eliminating certain readings. In this perspective, leaving the underspecified wffs around doesn't make much sense. I propose therefore to allow discourse interpretation principles to *rewrite* their 'triggering wff' whenever this wff encodes an H-type ambiguity, in addition to adding new wffs to a set. The more general format for discourse interpretation principles is as follows:

$$\frac{\alpha : \beta : \delta}{\sigma : \tau}$$

A rule of this form is an operation from sets of wffs into sets of wffs that given a set W of wffs containing α and β and not containing $\neg\delta$,[27] produces a set W$'$ of wffs containing τ, and in which α has been replaced by σ. I will call α the TRIGGERING CONDITION. For example, a version of CROAK1-IF-FROG in which the triggering condition $\mathbf{croak}_U(x)$ is rewritten by the consequent $\mathbf{croak}_1(x)$ is as follows:

CROAK1-IF-FROG

$$\frac{\mathbf{croak}_U(x) : \ \mathbf{frog}(x) \ : \mathbf{croak}_1(x)}{\mathbf{croak}_1(x) \ :}$$

If all disambiguation rules are rewritten in this format, a completely disambiguated extension can simply be characterized as one which doesn't contain any H-TYPE AMBIGUOUS WFFS. The notion of H-type ambiguous wff can be characterized either syntactically (by identifying certain syntactic constituents as specifying H-type ambiguity, and by classifying as H-type ambiguous a wff that contains one of these constituents)[28] or model-theoretically, e.g., by means of a function ι such that if X is a set of senses, ι(X) is 1 if the set of senses is admissible, 0 if it is inadmissible in Pinkal's sense. Whatever way we choose to define a H-type ambiguous wff, the Condition on Discourse Interpretation can now be formulated as follows:

Condition on Discourse Interpretation : An extension E of a discourse interpretation theory (DI,UF) cannot contain an H-type ambiguous wff.

[27] Strictly speaking, one should check that $\neg\delta$ does not occur in the extension itself, not in the intermediate sets of wffs; this form makes sense however once we adopt a 'syntactic' definition of extension (see below).

[28] Both the treatment of scopal ambiguity and the treatment of referential ambiguity proposed below are such that the constituents that introduce H-type ambiguity can be identified syntactically.

Notice that the statement of the Condition on Discourse Interpretation as a condition on pragmatic reasoning gives it the status of a felicity condition rather than of a hard constraint on interpretation.

4.2.5 Extensions, closure and consistency checking in an underspecified default theory

So far, I've been using the terminology from default logic as if the shift to an underspecified representation had no side effects, but this is not the case. Consider the way in which Reiter defines the notion of extension of a (closed)[29] default theory, for example:

Definition 4.1 *Let Δ = (D, W) be a closed default theory, so that every default of D has the form (α:β_1,..., β_m/w) where α,β_1,...,β_m,w are all closed wffs of L. For any set of closed wffs S $\subseteq$ L let ? (S) be the smallest set satisfying the following three properties:*

- *D1 W $\subseteq$? (S)*
- *D2 Th(? (S)) = ? (S)*
- *D3 If (α:β_1,..., β_m/w) $\in$ D and $\alpha \in$ S and $\neg\beta_1$, ..., $\neg\beta_m$ $\neg \in$ S then w $\in$? (S).*

A set of closed wffs E $\subseteq$ L is an EXTENSION *for Δ iff ? (E) = E, i.e., iff E is a fixed point of the operator ?.*

This definition crucially relies on the notion of deductive closure Th(S), defined as the set of wffs {w | S $\vdash$ w}; but what we said with semantic entailment holds for provability, as well: no clear notion exists of what it means for an expression of an underspecified language to follow from a set of wffs of the same language. Two routes are open to us. One is to define a notion of 'underspecified provability' $\vdash_U$, and to use $\vdash_U$ to define an 'underspecified' notion of closure Th$_U$(S). For example, we could say that w $\vdash_U$ w$'$ iff for each expression w$''$ that denotes a single one of the interpretations of w, w$''$ $\vdash_U$ w$'$. This route is not very appealing, however, if for no other reason that it's not clear that any way of defining an underspecified notion of provability will do.

The alternative is to adopt a new notion of extension that does not rely on deductive closure, i.e., one in which an extension is a fixed point of the operator ?$'$, which does not include condition D2 of the definition of ?:

- Let ?$'$(S) be the smallest set satisfying the following two properties:

[29] A closed default theory is one in which no default contains open variables. All really interesting cases of default inference rules do include such variables; but Reiter derives the definition of the extension of an 'open' default theory from the definition of extension for closed theories.

D1 W $\subseteq$?$'$(S)

D3 If (α:β_1,..., β_m/w) $\in$ D and $\alpha\in$ S and $\neg\beta_1$, ..., $\neg\beta_m$ $\neg$ $\in$ S then w $\in$?$'$(S).

Replacing ? with ?$'$ in the definition of extension has several consequences. First and foremost, dropping the requirement of deductive closure makes the test of whether it is consistent to assume β_1, ..., β_m essentially syntactic: i.e., it is possible for β_j not to be included in ?$'$(S) even though it is derivable from ?$'$(S). (In general, this definition of extension is a much closer description of the behavior of actual implementations of non-monotonic reasoning than the original definition.) And therefore, a logic defined in this way does not have the property of Reiter's logic that a (closed) default theory (D,W) has an inconsistent extension iff W is inconsistent.

In fact, each extension of a discourse interpretation theory under this new definition will, in general, be H-type ambiguous, some of the interpretations being inconsistent. However, if we adopt the 'rewriting' version of disambiguation discussed above, and impose the Condition on Discourse Interpretation, each extension will have a single interpretation, and therefore its consistency can be checked. I propose therefore to define the notion of extension of a discourse interpretation theory as follows:

Extension: A set of closed wffs E $\subseteq$ L is an EXTENSION for the discourse interpretation theory Δ iff E is a fixed point of the operator ?$'$ and satisfies the Condition on Discourse Interpretation.

5 Other Forms of Ambiguity

The theory of ambiguity introduced in the previous sections can be straightforwardly extended to obtain a treatment of two other classes of semantic ambiguity: scopal ambiguity and referential ambiguity. These extensions preserve the basic ideas of the theory, semantic ambiguity as multiplicity of meanings, and perceived ambiguity as multiple extensions of a default theory; what changes is that on the one hand, a more complex underspecified language is introduced, capable of encoding other forms of ambiguity; on the other hand, more complex inference rules are used.

5.1 Scopal Ambiguity

I will call the sentence constituents that modify the parameters of evaluation, and therefore affect the interpretation of other sentence constituents 'in their scope', OPERATORS. Examples of operators are quantifiers (that affect the choice of the variable assignment used to evaluate expressions in their scope) and modals (that affect the choice of the world / situation at which expressions in their scope are evaluated). As it is well-known, one cause of semantic ambiguity is that sentences may contain more than one operator, and their relative scope is not completely determined by the sentence's syntactic structure. Sentences that have more than one meaning due to the interaction between operators are called SCOPALLY AMBIGUOUS.[30]

Historically, most underspecified representations have been introduced to deal with scopal ambiguity. Typically, an intermediate step of processing is assumed in which operators are left 'in place,' as well as a subsequent step of processing in which their relative scope is determined by contextual processing. Schubert and Pelletier's underspecified representation of *Every kid climbed a tree* in (8) is an example of underspecified representation in which the operators are left 'in situ'.

These representations are typically justified in terms of ease of processing, and their ability to represent 'intermediate' readings. It is clear however that for the purposes of developing a 'principled' theory of ambiguity processing, it

[30] 'Ambiguity elimination' solutions to the combinatorial explosion puzzle, such as Kempson and Cormack's (1981) or Verkuyl's (1992) have had some success in showing that certain cases of 'ambiguity'—especially 'ambiguities' associated with plural noun phrases or certain classes of scopal ambiguities—are in fact cases of indeterminacy. Zwicky and Sadock noted that the identity tests do not classify a sentence as ambiguous if the propositions expressed by the sentence are such that one entails the other. This is the case, for instance, with sentences such as (9) (Lakoff , 1970; Zwicky and Sadock , 1975; Kempson and Cormack , 1981).

(9) Every kid climbed a tree.

(Kempson and Cormack , 1981) claim that sentences like (10) are not ambiguous, but indeterminate: according to them, such sentences semantically denote the weaker reading (the one in which the universal quantifier takes scope over the existential). The stronger reading is the result of pragmatic reasoning.

(10) Every linguistics student has read a book by Chomsky.

Kempson and Cormack propose in fact that all quantified sentences denote a single proposition; in this way, the combinatorial explosion puzzle disappears, at least as far as scopal ambiguity is concerned. However, it is not true in general that a sentence with two quantifiers has two interpretations, one of which entails the other. (11) does not have an interpretation weak enough to be entailed by all others, yet able to capture the truth conditions correctly.

(11) Few students know many languages.

A second problem with the proposal of Kempson and Cormack is that if one wants to claim that the meaning of a sentences such as (12a) is something like (12b), as Kempson and Cormack do, then one ends up predicting that the meaning of (12c) should be something like (12d), the strongest interpretation of the sentence. In other words, one either has to give up compositionality for sentences like (12b), or to abandon the strategy of letting sentences denote their weakest interpretation (Chierchia and McConnell-Ginet , 1990).

(12) a. Every kid climbed a tree.

b. $(\forall \text{ x } \mathbf{kid}(\text{x}) \supset (\exists \text{ y } \mathbf{tree}(\text{y}) \wedge \mathbf{climb}(\text{x,y})))$

c. It is not the case that every kid climbed a tree.

d. $\neg(\forall \text{ x } \mathbf{kid}(\text{x}) \supset (\exists \text{ y } \mathbf{tree}(\text{y}) \wedge \mathbf{climb}(\text{x,y})))$

It should also be clear that whatever the case for scopal ambiguity, other kinds of ambiguity, such as structural and H-type lexical ambiguity, cannot be reduced to indeterminacy.

would be much better to stick to as few new 'levels of representation' as possible.

In fact, there is no need to introduce a new level of representation. The two requirements on a scopally underspecified representation—that it allow representing the structural information provided by a sentence, and representing the intermediate steps of disambiguation—can be satisfied by using as an underspecified representation the syntactic structure of the sentence, augmented with information about the semantic interpretation of word-forms. In this way we can also maintain semantic translation of lexical items used in Montague grammar, that determine how they combine with other sentence constituents to determine a sentence's meaning.

The 'lexically and scopally underspecified language' $\mathcal{LSUL}$ I introduce to encode scopal ambiguity generalizes the language $\mathcal{LXUL}$ introduced in the previous section by allowing for arbitrary functional types. In this way, the lexical item *every* can be given its usual $\langle\langle e,t\rangle,\langle\langle e,t\rangle,t\rangle\rangle$ translation:

$$\lambda P.\ \lambda Q.\ (\forall\ \ x{:}\ P(x)\ (Q(x))).$$

The second augmentation to $\mathcal{LXUL}$ is the inclusion of tree-like expressions used to translate syntactic phrases. For example, the NP $[_{\mathrm{NP}}\ [_{\mathrm{Det}}$ every] $[_{\mathrm{N}}$ dog]] translates into the expression:

$$[_{\mathrm{NP}}\ \ [_{\mathrm{Det}}\ \ \lambda P.\ \lambda Q.\ \ (\forall\ \ x{:}\ P(x)\ (Q(x)))]\ [_{\mathrm{N}}\ \mathbf{dog}]]$$

The expression in (13) is the underspecified translation of the sentence *Every dog saw a frog*:

(13)

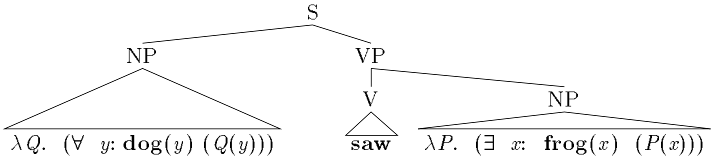

Besides reducing the number of representations floating around, this was of talking about scopal underspecification has two additional advantages over underspecified representations in which all syntactic information except for the position of operators is lost, such as Schubert and Pelletier's underspecified logical forms, Reyle's underspecified DRSs or the Core Language Engine's QLFs. First of all, the semantics of expressions such as (13) —that, for historical reasons, I call LOGICAL FORMS— can be computed in a completely classical fashion using the storage mechanism (Cooper , 1983), with the result that syntactic constraints on the available readings, such as the Scope Constraint (May , 1985; Heim , 1982), can play a role in determining the semantics of these objects, without the need for additional constraints such as the label ordering constraints used in UDRS (Reyle , 1993). Secondly, all structural information is preserved, not just information about the relative position of operators. Some of this syntactic information is used as a clue during disambiguation, for example, for interpreting pronouns, but also in certain theories of scopal disambiguation. (See, e.g., (Kurtzman and MacDonald , 1993) and (Poesio , 1994) for an account of scope disambiguation which esploits the syntactic information encoded by underspecified expressions such as (13).)

5.1.1 A Lexically and Scopally Underspecified Language

The semantics of the underspecified language $\mathcal{LSUL}$ is classically based on a set T of semantic types, the smallest set such that (i) e and t are types; and (ii) If τ and τ' are types, $\langle\tau,\tau'\rangle$ is a type. The set of meaningful expressions of type α is indicated by ME_α. The set of non-logical constant expressions of type α is indicated as $\mathrm{CE}_\alpha \subseteq \mathrm{ME}_\alpha$.

The semantics of $\mathcal{LSUL}$ is based on the same idea as the semantics of $\mathcal{LXUL}$. Natural language expressions are assigned objects of the same type that they would receive in (Dowty et al. , 1981) (as revised by Partee and Rooth (1983)), with the difference that, when I talk about 'meaningful expressions of type τ' below, therefore, I am really talking about expressions that denote sets of functions from the set of situations $\mathcal{S}$ to elements of $\mathrm{D}_\mathcal{T}$ (the domain of type τ). Thus for example, sentences are of type t both in Dowty, Wall and Peters' system, and in the current proposal; but a meaningful expression of type t in $\mathcal{LSUL}$ denotes a (function from a discourse situation to a) set of functions from situations to truth values.

The sets of meaningful expressions of $\mathcal{LSUL}$ include all the expressions in $\mathcal{LXUL}$:

- $\mathrm{CE}_e = \{k\}$.
- $\mathrm{CE}_{\langle e,t\rangle} = \{\mathbf{dog},\mathbf{frog},\mathbf{croak}_U,\ \mathbf{croak}_1,\ \mathbf{croak}_2\ \}$.
  (The language includes a single underspecified lexical interpretation, $\mathbf{croak}_U$.)
- $\mathrm{CE}_{\langle e,\langle e,t\rangle\rangle} = \{\mathbf{saw}\}$.
- If $\alpha \in \mathrm{ME}_{\langle\mathcal{T},\mathcal{T}'\rangle}$ and $\beta \in \mathrm{ME}_\mathcal{T}$, $\alpha(\beta) \in \mathrm{ME}_{\mathcal{T}'}$.
- If α and β are of type t, then $\alpha \wedge \beta$ and $\neg\alpha$ are in ME_t.

The set $\mathrm{ME}_\mathcal{T}$, for any type τ, includes a denumerably infinite set of VARIABLES of type τ. $\mathcal{LSUL}$ also includes lambda-abstracts and quantified expressions, defined below. The language also includes the new syntactic category of LOGICAL FORMS. The sets of logical forms of syntactic category XP, LF_{XP}, are defined as follows:

- $\mathrm{LF}_{\mathrm{Det}} = \{[_{\mathrm{Det}}\ \alpha] \mid \alpha \in \mathrm{ME}_{\langle\langle e,t\rangle,\langle\langle e,t\rangle,t\rangle\rangle}\}$;
- $\mathrm{LF}_{\mathrm{N}} = \{[_{\mathrm{N}}\ \alpha] \mid \alpha \in \mathrm{ME}_{\langle e,t\rangle}\}$;

- $\text{LF}_{\text{PN}} = \{[_{\text{PN}}\ \alpha] \mid \alpha \in \text{CE}_e\}$
- $\text{LF}_{\text{IV}} = \{\ [_{\text{IV}}\ \alpha] \mid \alpha \in \text{ME}_{\langle e,t\rangle}\}$;
- $\text{LF}_{\text{TV}} = \{[_{\text{TV}}\ \alpha] \mid \alpha \in \text{ME}_{\langle e,\langle e,t\rangle\rangle}\}$;
- $\text{LF}_{\text{NP}} = \{[_{\text{NP}}\ \alpha\ \beta] \mid \alpha \in \text{LF}_{\text{Det}}$ and $\beta \in \text{LF}_{\text{N}}\ \}$
  $\cup\ \{[_{\text{NP}}\ \alpha] \mid \alpha \in \text{LF}_{\text{PN}}\ \} \cup \{[_{\text{NP}}\ \alpha] \mid \alpha$ a variable of type e $\}$
- $\text{LF}_{\text{VP}} = \{[_{\text{VP}}\ \alpha\ \beta] \mid \alpha \in \text{LF}_{\text{TV}}$ and $\beta \in \text{LF}_{\text{NP}}\ \}$
  $\cup\ \{[_{\text{VP}}\ \alpha] \mid \alpha \in \text{LF}_{\text{IV}}\ \}$
- $\text{LF}_{\text{S}} = \{[_{\text{S}}\ \alpha\ \beta] \mid \alpha \in \text{LF}_{\text{NP}}$ and $\beta \in \text{LF}_{\text{VP}}\ \}$

Meaningful expressions are assigned a value with respect to a universe $\mathcal{U}$. I use below the notation $\underline{s}$ to indicate that $\underline{a}$ 'stands for' an object in $\mathcal{U}$, i.e., it is part of the metalanguage, as opposed to being a meaningful expression of the object language. The models with respect to which a CRT expression is evaluated include a set $\mathcal{S}$ of situations. The only fact about situations I use here is that they have CONSTITUENTS.

The interpretation of types with respect to $\mathcal{U}$ is defined as usual: $\text{D}_{e,U} = \mathcal{U}$; $\text{D}_{t,U} = \{0{,}1\}$; $\text{D}_{\langle a,b\rangle,U} = \text{D}_b^{D_a}$. In the rest of this paper, I generally drop the indication of the universe (e.g., I write D_e instead of $\text{D}_{e,U}$). The model of interpretation for CRT expressions is the triple $\langle \mathcal{U}, \mathcal{S}, \text{I}\rangle$. The interpretation function 'I' assigns an interpretation to constants of type τ.

The value of meaningful expressions is specified by a function $[\![.]\!]^{\text{M,g,d}}$ that includes an assignment function among its parameters, since the terms of $\mathcal{LSUL}$ include variables. The interpretation of variables is specified by the following clause:

- For α a variable of type τ, $[\![\alpha]\!]^{\text{M,g,d}}$ = the singleton set $\{\text{g}(\alpha)\}$, where $\text{g}(\alpha)$ is a constant function $\text{f} : \mathcal{S} \rightarrow \text{D}_{\mathcal{T}}$.

The interpretation of constants, connectives and application is as in $\mathcal{LXUL}$. The denotation of the other expressions is discussed below.

5.1.2 An Example of a Scopally and Lexically Underspecified Grammar

The following grammar extends the grammar discussed in section §4 by adding determiners and relations as new lexical items:

- PN → Kermit; k
- Det → every; $\lambda P.\ \ \lambda Q.\ \ (\forall\ \ x{:}\ P(x)\ \ (Q(x)))$
- Det → a; $\lambda P.\ \ \lambda Q.\ \ (\exists\ \ x{:}\ P(x)\ \ (Q(x)))$

- N → dog; **dog**
- N → frog; **frog**
- IV → croaked; $\mathbf{croak}_U$
- TV → saw; **saw**

and by adding phrase structure rules for NPs and transitive verbs:

- S → NP VP; $[_{\mathrm{S}}\ \mathrm{NP}'\ \mathrm{VP}']$
- NP → PN; $[_{\mathrm{NP}}\ [_{\mathrm{PN}}\ \mathrm{PN}']]$
- NP → Det N; $[_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathrm{Det}']\ [_{\mathrm{N}}\ \mathrm{N}']]$
- VP → IV; $[_{\mathrm{VP}}\ [_{\mathrm{IV}}\ \mathrm{V}']]$
- VP → TV NP; $[_{\mathrm{VP}}\ [_{\mathrm{TV}}\ \mathrm{V}']\ \mathrm{NP}']$

This grammar generates, in addition to lexically ambiguous sentences such as *Kermit croaked*, scopally ambiguous sentences such as *Every dog saw a frog.*

5.1.3 The Denotation of Logical Forms

The denotation of logical forms is specified using the *storage* method, developed by Robin Cooper as a way around a problem with Montague's QUANTIFYING IN technique, namely, the fact that in order to get all the readings of a scopally ambiguous sentence, one has to stipulate that the sentence is syntactically ambiguous (see (Dowty et al. , 1981)).

Cooper proposed that the value of a syntactic tree is a set of SEQUENCES, each sequence representing a distinct 'order of application' of the operators that may result in a admissible interpretation of a sentence. For example, the quantifier *a frog* can 'enter' the derivation of the VP *saw a frog* in two different ways. The narrow scope reading is obtained by immediately applying the interpretation of the quantifier to the translation of *saw*; but it is also possible to apply the predicate to the variable quantified over, and 'wait' before applying the quantifier, in which case the wide scope reading is obtained. The value of the NP *a frog*, then, is the set of two sequences shown in (14). One sequence consists of a single element, the 'traditional' Montague-style translation of *every frog.* The second sequence consists of two elements: the variable y, and the semantic translation of the quantified NP, put 'in storage'.

(14) $\{\langle \lambda P.\ (\exists\ y\text{: }\mathbf{frog}(y)\ (P(y)))\rangle,$
$\langle y,\ \lambda P.\ (\exists\ y\text{: }\mathbf{frog}(y)\ (P(y)))\rangle\ \}$

Ambiguity 'propagates up' as follows. The value of the VP *saw a frog* in (15) also consists of two sequences, one obtained by applying the first element of the first sequence in the denotation of *every frog* to the predicate **saw**, the other obtained by applying the predicate **saw** to the first element of the second sequence (the variable y). The result is as in (16).

(15) $[_{\mathrm{VP}}\ [_{\mathrm{V}}\ \mathbf{saw}]\ [_{\mathrm{NP}}\ \lambda P.\ (\exists\ y\text{: }\mathbf{frog}(y)\ \ (P(y)))]]$

(16) $\{\langle \lambda x.\ \ (\exists\ \ y\text{: }\mathbf{frog}(y)\ (\mathbf{saw}(y)(x)))\ \rangle,$
$\langle \mathbf{saw}(y),\ \lambda P.\ \ (\exists\ \ y\text{: }\mathbf{frog}(y)\ (P(y)))\ \rangle\}$

Finally, the value of a sentence is obtained by combining the value of the VP with the value of the NP in the usual fashion: the value of $[_{\mathrm{S}}$ *Every dog saw a frog*] is a set of two sequences, each representing a distinct reading of the sentence.

It's easy to see that Cooper's technique can be used to assign to underspecified representations like (13) a 'multiple sense' denotation like those assigned to lexically ambiguous expressions in the previous section. All that is needed is a function CV that assigns to each expression of the form $[_{\mathrm{XP}}\ \alpha]$ its 'Cooper Value'; the denotation of sentence translations like $[_{\mathrm{S}}\ \beta]$ can then be defined in terms of CV as follows:

- If $[_{\mathrm{S}}\ \alpha]$ is a logical form in LE_{S} then $[_{\mathrm{S}}\ \alpha]$ is a meaningful expression of type t. Let CV(α)(M,g,d) be the set of single-element sequences $\{\langle\{\sigma_{1_1},\ldots,\sigma_{1_j}\}\rangle, \ldots, \langle\{\sigma_{n_1},\ldots,\sigma_{n_k}\}\rangle\}$. Then
  $[\![[_{\mathrm{S}}\ \alpha]]\!]^{\mathrm{M,g,d}} = \{\sigma_{1_1},\ldots,\sigma_{1_j},\ldots\sigma_{n_1},\ldots,\sigma_{n_k}\}$

(I have taken into account the fact that an expression of our underspecified language denotes a set of objects, therefore each scopally disambiguated translation of a sentence will still denote a set of propositions.)

Cooper discusses in detail in (Cooper , 1983) how semantic and syntactic constraints on scope can be implemented as requirements that the storage be 'discharged' at certain positions—i.e., that no element in storage be 'carried across' syntactic constructions that produce scope islands, such as $\overline{\mathrm{S}}$. In this way, no operator in a clause may take scope over operators in an higher clause, or in a sister clause, thus enforcing the Scope Constraint discussed in §2.2.[31]

The CV function used to define the interpretation of logical forms is based on an implementation of the storage idea less general than Cooper's, but simpler. In order to arrive at a uniform specification of the Cooper Value of all constructs, it is useful to define construct-specific versions of application in which to 'bury' the differences in storage manipulation. These operations are defined as follows:

[31] The definition of storage generates spurious readings in the case of embedded NPs such as *a representative of every company*, that have to be eliminated via a separated filter (Keller , 1988). Keller introduced a 'nesting' technique that obviates the problem (1988). More recently, Pereira (Pereira , 1990) argued that the right scoping properties can be obtained without additional stipulations from the natural deduction approach to parsing. I only consider here 'basic' NPs that do not create problems for the simplest version of Cooper's technique.

- $\alpha\{\beta\}_{\mathrm{S}} = \begin{cases} \alpha(\beta)\text{, if Type}(\alpha) = \langle\langle e,t\rangle,t\rangle \text{ and Type}(\beta) = \langle e,t\rangle, \\ \quad \text{or if Type}(\alpha) = \langle\langle e,t\rangle,\langle e,t\rangle\rangle \text{ and Type}(\beta) = \\ \quad \langle e,t\rangle. \\ \lambda y.\ \ \lambda x.\ \ \beta(x)(y)\{\alpha\}_{\mathrm{VP}} \text{ if Type}(\alpha) = \langle\langle e,t\rangle,t\rangle \text{ and} \\ \quad \text{Type}(\beta) = \langle e,\langle e,t\rangle\rangle\text{, or if Type}(\alpha) = \\ \quad \langle\langle e,t\rangle,\langle e,t\rangle\rangle \text{ and Type}(\beta) = \langle e,\langle e,t\rangle\rangle. \\ \beta(\alpha)\text{, if Type}(\beta) = \langle e,t\rangle \text{ and Type}(\alpha) = e. \\ \text{Undefined otherwise.} \end{cases}$

- $\alpha\{\beta\}_{\mathrm{VP}} = \begin{cases} \alpha(\beta)\text{, if Type}(\alpha) = \langle e,t\rangle \text{ and Type}(\beta) = e. \\ \lambda w_e.\ \ \beta(\alpha(w_e))\text{, if Type}(\beta) = \langle\langle e,t\rangle,\langle e,t\rangle\rangle \\ \quad \text{and Type}(\alpha) = \langle e,\langle e,t\rangle\rangle. \\ \lambda w_e.\ \ \beta(\lambda z.\ \ \alpha(z)(w_e))\text{, if Type}(\beta) = \langle\langle e,t\rangle,t\rangle \\ \quad \text{and Type}(\alpha) = \langle e,\langle e,t\rangle\rangle. \\ \text{Undefined otherwise.} \end{cases}$

- $\alpha\{\beta\}_{\mathrm{XP}}$ (where XP is a category other than S or VP) =

  $$\begin{cases} \alpha(\beta)\text{, if Type}(\alpha) = \langle\tau,\tau'\rangle \text{ and Type}(\beta) = \tau. \\ \beta(\alpha)\text{, if Type}(\beta) = \langle\tau,\tau'\rangle \text{ and Type}(\alpha) = \tau. \\ \text{Undefined otherwise.} \end{cases}$$

Next, we need an operation that combines two sets of sequences into one. The result of applying this operation to two sets of sequences X and Y is the set of sequences obtained by (typed) applying the first element of a sequence in X to the first element of a sequence in Y and then merging the rest of the sequences, as follows:

- $\mathrm{X}\otimes_{\mathrm{YP}}\mathrm{Y}$, where X and Y are the sets
  $\mathrm{X} = \{\langle \mathrm{x}_{11},\mathrm{x}_{12}, \ldots \mathrm{x}_{1m}\rangle, \ldots \langle \mathrm{x}_{n1},\mathrm{x}_{n2}, \ldots \mathrm{x}_{nl}\rangle\}$ and
  $\mathrm{Y} = \{\langle \mathrm{y}_{11},\mathrm{y}_{12}, \ldots \mathrm{y}_{1p}\rangle, \ldots \langle \mathrm{y}_{q1},\mathrm{y}_{q2}, \ldots \mathrm{y}_{qr}\rangle\}$,
  and YP is any phrase category, is the set:
  $\{\ \langle \mathrm{x}_{i1}\{\mathrm{y}_{j1}\}_{\mathrm{YP}}, \mathrm{y}_{j2}, \ldots \mathrm{y}_{jp}\rangle, \mathrm{x}_{i2}, \ldots \mathrm{x}_{im} \mid \langle \mathrm{x}_{i1},\mathrm{x}_{i2}, \ldots \mathrm{x}_{im}\rangle \in \mathrm{X} \text{ and } \langle \mathrm{y}_{j1},\mathrm{y}_{j2}, \ldots \mathrm{y}_{jp}\rangle \in \mathrm{Y}\ \}$

We also need an operation to put operators into store, and one to 'discharge' them. The **storeaway** operation takes a set consisting a single single-element sequence and a RESULT, and returns a set that consists of two sequences: the original sequence, and a new sequence consisting of the result and the operator in store.

- **storeaway**(α,X), where X = $\{\langle\beta\rangle\}$, is $\{\langle\beta\rangle,\langle\alpha,\beta\rangle\}$

The **discharge** operation takes a sequence and applies all operators back to obtain a set of sequences with a single element and an empty store. For simplicity, we will assume that all operators are generalised quantifiers, i.e., of type $\langle\langle e,t\rangle,t\rangle$. (No other operators are specified in the grammar above.) **discharge** is defined as follows:

- **discharge**($\langle \mathrm{x}_1,\mathrm{x}_2, \ldots \mathrm{x}_m\rangle$) = $\{\mathrm{x}_m\{\ \mathrm{x}_{m-1}\{\ldots \mathrm{x}_2\{\mathrm{x}_1\}_{\mathrm{S}}\}_{\mathrm{S}}\ \}_{\mathrm{S}}\ \}$

discharge*(X), where X is a set of sequences, is the union $\bigcup_{x\in X}$ **discharge**(x). We can now specify the Cooper value of logical forms with respect to model M, variable assignment g, and discourse situation d as follows:[32]

- CV($[_{\mathrm{S}}\ [_{\mathrm{NP}}\ \alpha]\ [_{\mathrm{VP}}\ \beta]]$) = **discharge***(CV($[_{\mathrm{NP}}\ \alpha]$) $\otimes_{\mathrm{S}}$CV($[_{\mathrm{VP}}\ \beta]$))
- CV($[_{\mathrm{NP}}\ [_{\mathrm{PN}}\ \alpha]]$) = CV($[_{\mathrm{PN}}\ \alpha]$)
- CV($[_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \alpha]\ [_{\mathrm{N}}\ \beta]]$) = **storeaway**($\lambda Q.\ \lambda z.\ Q(z)$, CV($[_{\mathrm{Det}}\ \alpha]$) $\otimes_{\mathrm{NP}}$CV($[_{\mathrm{N}}\ \beta]$))
- CV($[_{\mathrm{NP}}\ \alpha]$) = $\{\langle[\![\alpha]\!]^{\mathrm{M,g,d}}\rangle\}$, where α is a variable of type e.
- CV($[_{\mathrm{VP}}\ [_{\mathrm{IV}}\ \alpha]]$) = CV($[_{\mathrm{IV}}\ \alpha]$)
- CV($[_{\mathrm{VP}}\ [_{\mathrm{TV}}\ \alpha]\ [_{\mathrm{NP}}\ \beta]]$) = CV($[_{\mathrm{TV}}\ \alpha]$) $\otimes_{\mathrm{VP}}$CV($[_{\mathrm{NP}}\ \beta]$)
- CV($[_{\mathrm{PN}}\ \mathrm{PN}']$) = $\{\langle[\![\mathrm{PN}']\!]^{\mathrm{M,g,d}}\rangle\}$
- CV($[_{\mathrm{Det}}\ \mathrm{Det}']$) = $\{\langle[\![\mathrm{Det}']\!]^{\mathrm{M,g,d}}\rangle\}$
- CV($[_{\mathrm{IV}}\ \mathrm{IV}']$) = $\{\langle[\![\mathrm{IV}']\!]^{\mathrm{M,g,d}}\rangle\}$
- CV($[_{\mathrm{TV}}\ \mathrm{TV}']$) = $\{\langle[\![\mathrm{TV}']\!]^{\mathrm{M,g,d}}\rangle\}$
- CV($[_{\mathrm{N}}\ \mathrm{N}']$) = $\{\langle[\![\mathrm{N}']\!]^{\mathrm{M,g,d}}\rangle\}$

There are three tricky aspects to the definition of CV: the discharge operation in the definition of the Cooper Value of a sentence translation, the definition of CV($[_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \alpha]\ [_{\mathrm{N}}\ \beta]]$) in which an operator is put in store, and the definition of CV($[_{\mathrm{VP}}\ [_{\mathrm{V}}\ \alpha]\ [_{\mathrm{NP}}\ \beta]]$ in which two stores are combined, and that has different results depending on whether the NP is of type e or is a quantifier. I'll illustrate these cases by looking at the main steps of the computation of the CV of (13):

[32] Strictly speaking, the form CV(α)(M,g,d) should be used. I omit the indices below.

1. $\mathrm{CV}([_{\mathrm{Det}}\ \lambda P.\ \lambda Q.\ (\exists\ y{:}\ P(y)\ (Q(y)))]) =$
   $\{\langle \lambda P.\ \lambda Q.\ (\exists\ y{:}\ P(y)\ (Q(y)))\rangle\}$
   (In what follows, I'll use **a** to indicate the lexical translation of the determiner *a*.)

2. $\mathrm{CV}([_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{a}]\ [_{\mathrm{N}}\ \mathbf{frog}]]) =$
   $\mathbf{storeaway}(\lambda Q.\ \lambda z.\ Q(z),\ \mathrm{CV}([_{\mathrm{Det}}\ \mathbf{a}]) \otimes_{\mathrm{NP}} \mathrm{CV}([_{\mathrm{N}}\ \mathrm{frog}])) =$
   $\{\langle \lambda Q.\ (\exists\ y{:}\ \mathbf{frog}(y)\ (Q(y)))\rangle,$
   $\langle \lambda Q.\ \lambda z.\ Q(z),\ \lambda Q.\ (\exists\ y{:}\ \mathbf{frog}(y)\ (Q(y)))\rangle\ \}$
   (I'll use [**a frog**] to stand for $'\lambda Q.\ (\exists\ y{:}\ \mathbf{frog}(y)\ (Q(y)))'$ below.)

3. $\mathrm{CV}([_{\mathrm{VP}}\ [_{\mathrm{TV}}\ \mathbf{saw}]\ [_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{a}]\ [_{\mathrm{N}}\ \mathbf{frog}]]]) =$
   $\mathrm{CV}([_{\mathrm{TV}}\ \mathbf{saw}]) \otimes_{\mathrm{VP}} \mathrm{CV}([_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{a}]\ [_{\mathrm{N}}\ \mathbf{frog}]]) =$
   $\{\langle \mathbf{saw}\{[\mathbf{a\ frog}]\}_{\mathrm{VP}}\rangle,\ \langle \mathbf{saw}\{\ \lambda Q.\ \lambda z.\ Q(z)\ \}_{\mathrm{VP}},\ [\mathbf{a\ frog}]\rangle\ \} =$
   $\{\langle \lambda w.\ \ (\exists\ y{:}\ \mathbf{frog}(y)\ \ (\mathbf{saw}(y)(w)))\rangle,$
   $\langle \lambda w.\ \lambda x.\ \mathbf{saw}(w)(x),\ [\mathbf{a\ frog}]\rangle\ \}$
   (Notice that the first element of the second sequence, in which [**a frog**] is kept in store, is still of type $\langle e,\langle e,t\rangle\rangle$.) I will use the abbreviations σ and σ' below for the first element of the first sequence and the first element of the second sequence, respectively.

4. $\mathrm{CV}([_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{every}]\ [_{\mathrm{N}}\ \mathbf{dog}]]) =$
   $\{\langle \lambda Q.\ (\forall\ x{:}\ \mathbf{dog}(x)\ (Q(x)))\rangle,$
   $\langle \lambda Q.\ \lambda z.\ Q(z),\ \lambda Q.\ (\forall\ x{:}\ \mathbf{dog}(x)\ (Q(x)))\rangle\ \}$
   (I'll use [**every dog**] to stand for $'\lambda Q.\ (\forall\ x{:}\ \mathbf{dog}(x)\ (Q(x)))'$ below.)

5. $\mathrm{CV}([_{\mathrm{S}}\ [_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{every}]\ [_{\mathrm{N}}\ \mathbf{dog}]]\ [_{\mathrm{VP}}\ [_{\mathrm{V}}\ \mathbf{saw}]\ [_{\mathrm{NP}}\ [_{\mathrm{Det}}\ \mathbf{a}]\ [_{\mathrm{N}}\ \mathbf{frog}]]]]) =$
   $\mathbf{discharge}^{*}(\mathrm{CV}([_{\mathrm{NP}}\ \alpha]) \otimes_{\mathrm{S}} \mathrm{CV}([_{\mathrm{VP}}\ \beta])) =$

$$\mathbf{discharge}^{*}\left(\begin{array}{l}\{\langle [\mathbf{every\ dog}]\{\sigma\}_{\mathrm{S}}\rangle,\\ \langle [\mathbf{every\ dog}]\{\sigma'\}_{\mathrm{S}},\ [\mathbf{a\ frog}]\rangle,\\ \langle\ \lambda Q.\ \lambda z.\ Q(z)\{\sigma\}_{\mathrm{S}},[\mathbf{every\ dog}]\rangle,\\ \langle\ \lambda Q.\ \lambda z.\ Q(z)\{\sigma'\}_{\mathrm{S}},\ [\mathbf{every\ dog}],\ [\mathbf{a\ frog}]\rangle\}\end{array}\right) =$$

$$\mathbf{discharge}^{*}\left(\begin{array}{l}\{\langle\ (\forall\ x{:}\ \mathbf{dog}(x)\ (\ (\exists\ y{:}\ \mathbf{frog}(y)\ (\ \mathbf{saw}(y)(x))))))\rangle,\\ \langle \lambda w.\ \ (\forall\ x{:}\ \mathbf{dog}(x)\ \ (\mathbf{saw}(w)(x))),[\mathbf{a\ frog}]\rangle,\\ \langle \lambda z.\ \ (\exists\ y{:}\ \mathbf{frog}(y)\ (\ \mathbf{saw}(y)(x))),[\mathbf{every\ dog}]\rangle,\\ \langle \lambda w.\ \ \lambda z.\ \mathbf{saw}(z)(w),\ [\mathbf{a\ frog}],[\mathbf{every\ dog}]\rangle\ \}\end{array}\right) =$$

   $\{\langle\ (\forall\ x{:}\ \mathbf{dog}(x)\ (\ (\exists\ y{:}\ \mathbf{frog}(y)\ (\ \mathbf{saw}(y)(x))))))\rangle,$
   $\langle\ (\exists\ y{:}\ \mathbf{frog}(y)\ (\ (\forall\ x{:}\ \mathbf{dog}(x)\ (\ \mathbf{saw}(y)(x))))))\rangle\}$

The Scope Constraint is enforced by requiring a complete discharge at the sentential level, which means no operators can 'move up' outside the sentence in which it occurs, although of course this couldn't occur in this grammar since it doesn't cover relative clauses, sentential complements or coordination. I have assumed that discharge only takes place at sentential level, i.e., there are no operators taking scope over VPs; doing this would complicate matters a bit in that a 'partial' discharge operation should be defined.

5.1.4 Lambda Abstracts

Some care is required in the system developed here to get a semantics for lambda-abstraction that preserves properties such as β- and η-reduction. The clause specifying the denotation of lambda-abstraction in Dowty, Wall and Peters's book is the following:

- If α is a variable of type τ and β a meaningful expression of type τ', then $\lambda\alpha.\beta$ is an expression of type $\langle\tau,\tau'\rangle$, and $\llbracket\lambda\alpha.\beta\rrbracket^{\mathrm{M,g,w,t}}$ is that function h: $\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'}$ such that for all objects $\underline{\mathrm{a}}$ in $\mathrm{D}_{\mathcal{T}}$, h($\underline{\mathrm{a}}$) is equal to $\llbracket\beta\rrbracket^{\mathrm{M,g}\{\alpha/\underline{\mathrm{a}}\}\mathrm{,w,t}}$.

If we generalize this clause in the 'obvious' way we get:

- If α is a variable of type τ and β a meaningful expression of type τ', then $\lambda\alpha.\beta$ is an expression of type $\langle\tau,\tau'\rangle$, and $\llbracket\lambda\alpha.\beta\rrbracket^{\mathrm{M,g,d}}$ is the set {f | where f: $\mathcal{S} \rightarrow (\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'})$ and for all situations $\underline{\mathrm{s}}$, f($\underline{\mathrm{s}}$) = h: $\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'}$, such that, for all objects $\underline{\mathrm{a}}$ in $\mathrm{D}_{\mathcal{T}}$, h($\underline{\mathrm{a}}$) is equal to h′($\underline{\mathrm{s}}$), for some h′ $\in \llbracket\beta\rrbracket^{\mathrm{M,g}\{\alpha/\underline{\mathrm{a}}\}\mathrm{,d}}$ }.

Lambda-abstraction defined in this way does not have the required properties. To show that it does not preserve η-reduction,[33] it is sufficient to consider the following example: let $\mathcal{S} = \{\underline{\mathrm{s}_1}, \underline{\mathrm{s}_2}\}$, $\mathrm{D}_{\mathcal{T}} = \{\underline{\mathrm{a}},\underline{\mathrm{b}}\}$, and let the expression β of type $\langle\tau,\tau'\rangle$ have the following denotation:

$$\left\{ \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \begin{bmatrix} \underline{\mathrm{a}} \rightarrow \underline{\alpha_1} \\ \underline{\mathrm{b}} \rightarrow \underline{\beta_1} \end{bmatrix} \\ \underline{\mathrm{s}_2} \rightarrow \begin{bmatrix} \underline{\mathrm{a}} \rightarrow \underline{\alpha_2} \\ \underline{\mathrm{b}} \rightarrow \underline{\beta_2} \end{bmatrix} \end{bmatrix}, \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \begin{bmatrix} \underline{\mathrm{a}} \rightarrow \underline{\alpha_3} \\ \underline{\mathrm{b}} \rightarrow \underline{\beta_3} \end{bmatrix} \\ \underline{\mathrm{s}_2} \rightarrow \begin{bmatrix} \underline{\mathrm{a}} \rightarrow \underline{\alpha_4} \\ \underline{\mathrm{b}} \rightarrow \underline{\beta_4} \end{bmatrix} \end{bmatrix} \right\}$$

Then $\llbracket\beta(\alpha)\rrbracket^{\mathrm{M,g}\{\alpha/\underline{\mathrm{a}}\}\mathrm{,d}}$ is as follows:

$$\left\{ \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \underline{\alpha_1} \\ \underline{\mathrm{s}_2} \rightarrow \underline{\alpha_2} \end{bmatrix}, \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \underline{\alpha_3} \\ \underline{\mathrm{s}_2} \rightarrow \underline{\alpha_4} \end{bmatrix} \right\}$$

and $\llbracket\beta(\alpha)\rrbracket^{\mathrm{M,g}\{\alpha/\underline{\mathrm{b}}\}\mathrm{,d}}$is as follows:

$$\left\{ \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \underline{\beta_1} \\ \underline{\mathrm{s}_2} \rightarrow \underline{\beta_2} \end{bmatrix}, \begin{bmatrix} \underline{\mathrm{s}_1} \rightarrow \underline{\beta_3} \\ \underline{\mathrm{s}_2} \rightarrow \underline{\beta_4} \end{bmatrix} \right\}$$

Then, under the definition above, $\llbracket\lambda\alpha.\ \beta(\alpha)\rrbracket^{\mathrm{M,g,d}}$ will contain the following function, that is not part of the denotation of β (hence, η-reduction is not a sound inference rule):

[33] I.e., that $\lambda\alpha.\ \beta(\alpha) \neq \beta$.

$$\left[\begin{array}{l} \underline{s_1} \rightarrow \left[\begin{array}{l} \underline{a} \rightarrow \underline{\alpha_1} \\ \underline{b} \rightarrow \underline{\beta_3} \end{array}\right] \\ \underline{s_2} \rightarrow \left[\begin{array}{l} \underline{a} \rightarrow \underline{\alpha_4} \\ \underline{b} \rightarrow \underline{\beta_2} \end{array}\right] \end{array}\right]$$

Intuitively, the problem with the definition above is that it does not 'preserve' the functions in the denotation of β. A definition of lambda-abstraction that does preserve these functions, and therefore preserves the soundness of β- and η-reduction, can be obtained as follows.[34]

The denotation function $[\![\alpha]\!]^{\mathrm{M,g,d}}$ used so far assigns a value to expression $\alpha \in \mathrm{ME}_{\mathcal{T}}$ in model M with respect to the parameters of evaluation g and d, $[\![\alpha]\!]^{\mathrm{M,g,d}} \subseteq \mathcal{S} \rightarrow \mathrm{D}_{\mathcal{T}}$. Another way of specifying the value of expressions is to define a function $[\![\alpha]\!]^{\mathrm{M,d}}$ that assigns as value to α at discourse situation d a set of functions of type (Ass $\rightarrow$ ($\mathcal{S}\rightarrow$ $\mathrm{D}_{\mathcal{T}}$)), from assignments to functions in ($\mathcal{S}\rightarrow$ $\mathrm{D}_{\mathcal{T}}$). For example, Dowty, Wall and Peters' clause for lambda abstraction could be rewritten as follows:

- If α is a variable of type τ and β a meaningful expression of type τ', then $\lambda\alpha.\beta$ is an expression of type $\langle\tau,\tau'\rangle$, and $[\![\lambda\alpha.\ \beta]\!] = \Lambda(\alpha,[\![\beta]\!])$, where $\Lambda(\alpha,\mathrm{Y})$ is that function f from Ass $\rightarrow$ ($\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'}$) such that for all g $\in$ Ass, $\underline{\mathrm{s}} \in \mathcal{S}$, $\underline{\mathrm{a}}$ in $\mathrm{D}_{\mathcal{T}}$, $\mathrm{f(g)}(\underline{\mathrm{s}})(\underline{\mathrm{a}}) = \mathrm{Y}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}})$.

This definition can then be generalized as follows:

- If α is a variable of type τ and β a meaningful expression of type τ', then $\lambda\alpha.\ \beta$ is an expression of type $\langle\tau,\tau'\rangle$, and $[\![\lambda\alpha.\ \beta]\!]^{\mathrm{M,d}} = \Lambda^{+}(\alpha,[\![\beta]\!])$, where $\Lambda^{+}(\alpha,\mathrm{Y}) = \{\Lambda(\alpha,\mathrm{m})$ for some $\mathrm{m} \in \mathrm{Y}\} = \{\mathrm{f} \in$ (Ass $\rightarrow$ ($\mathcal{S}\rightarrow$ ($\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'}$))) such that for all g $\in$ Ass, $\underline{\mathrm{s}} \in \mathcal{S}$, $\underline{\mathrm{a}}$ in $\mathrm{D}_{\mathcal{T}}$, and for some $\mathrm{m} \in \mathrm{Y}$, $\mathrm{f(g)}(\underline{\mathrm{s}})(\underline{\mathrm{a}}) = \mathrm{m}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}})$. }

Lambda-abstraction defined this way does support η-reduction.[35] Since this more general way of assigning a value is not needed to provide a semantics for the other constructs of $\mathcal{LSUL}$, I will continue using a function $[\![.]\!]^{\mathrm{M,g,d}}$, but the reader should keep in mind that a denotation function of this form is needed to deal with lambda abstraction, hence, with quantification. (And for referential ambiguity, as we will see below.)

[34] I wish to thank an anonymous reviewer for suggesting this solution to the problem just discussed, much simpler than the solution proposed in (Poesio , 1994).

[35] The proof is as follows. $[\![\lambda\alpha.\ \beta(\alpha)]\!]^{\mathrm{M,d}} = \{\mathrm{f} \in$ (Ass $\rightarrow$ ($\mathcal{S}\rightarrow$ ($\mathrm{D}_{\mathcal{T}} \rightarrow \mathrm{D}_{\mathcal{T}'}$))) such that for all g $\in$ Ass, $\underline{\mathrm{s}} \in \mathcal{S}$, $\underline{\mathrm{a}}$ in $\mathrm{D}_{\mathcal{T}}$, and for some $\mathrm{m} \in [\![\beta(\alpha)]\!]^{\mathrm{M}}$, $\mathrm{f(g)}(\underline{\mathrm{s}})(\underline{\mathrm{a}}) = \mathrm{m}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}})$. } Because of the definition of $[\![\beta(\alpha)]\!]^{\mathrm{M,d}}$, this is the set of functions f such that $\mathrm{f(g)}(\underline{\mathrm{s}})(\underline{\mathrm{a}}) = \mathrm{p}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}})[\mathrm{q}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}})]$, for some $\mathrm{p} \in [\![\beta]\!]^{\mathrm{M}}$ and some q in $[\![\alpha]\!]^{\mathrm{M}}$, i.e., of the functions which occur in $[\![\beta]\!]^{\mathrm{M}}$ since $\mathrm{q}(\mathrm{g}\{\alpha/\underline{\mathrm{a}}\})(\underline{\mathrm{s}}) = \underline{\mathrm{a}}$.

5.1.5 Quantification

The treatment of quantifiers in $\mathcal{LSUL}$ is based on Generalized Quantifiers Theory (Barwise and Cooper , 1981), i.e., the idea that determiners denote relations between two sets. The 'restricted quantification' notation used in the examples above is defined in terms of the two determiners **every** and **a**, as follows:

- $(\forall\ x{:}\ \Phi\ (\Psi)) \equiv_{\rm def}$ **every**$(\lambda x.\ \Phi, \lambda x.\ \Psi)$
- $(\exists\ x{:}\ \Phi\ (\Psi)) \equiv_{\rm def}$ **a**$(\lambda x.\ \Phi, \lambda x.\ \Psi)$

A 'single-valued' semantics for **every**$(\lambda\alpha.\ \Phi, \lambda\alpha.\ \Psi)$ could be defined, in first approximation, as in the following clause:

- Let F,G be meaningful expressions of type $\langle e,t\rangle$. Then **every**(F,G) is a meaningful expression of type t, and $[\![\textbf{every}(\mathrm{F,G})]\!]^{\mathrm{M,g,d}}$ is the function f s.t.

$$\mathrm{f}(\underline{\mathrm{s}}) = \begin{cases} \textbf{undefined} \text{ iff either } [\![\mathrm{F}]\!]^{\mathrm{M,g,d}} \text{ or } [\![\mathrm{G}]\!]^{\mathrm{M,g,d}} \text{ are} \\ \qquad \text{undefined;} \\ \textbf{1} \text{ iff } \{\underline{\mathrm{a}} \in \mathcal{U} \text{ s.t. } [[\![\mathrm{F}]\!]^{\mathrm{M,g,d}}(\underline{\mathrm{s}})](\underline{\mathrm{a}}) = 1\ \} \subseteq \\ \qquad \{\underline{\mathrm{b}} \in \mathcal{U} \text{ s.t. } [[\![\mathrm{G}]\!]^{\mathrm{M,g,d}}(\underline{\mathrm{s}})](\underline{\mathrm{b}}) = 1\ \}. \\ \textbf{0} \text{ otherwise.} \end{cases}$$

This definition can be generalized as follows into one that works in the case in which $[\![.]\!]^{\mathrm{M,g,d}}$ is a set:

- $[\![\textbf{every}(\mathrm{F,G}]\!]^{\mathrm{M,g,d}} =$

  $\{\mathrm{f} \mid$ there is an $\mathrm{h} \in [\![\mathrm{F}]\!]^{\mathrm{M,g,d}}$ and an $\mathrm{h}' \in [\![\mathrm{G}]\!]^{\mathrm{M,g,d}}$ such that

$$\mathrm{f}(\underline{\mathrm{s}}) = \begin{cases} \textbf{undefined} \text{ iff either } \mathrm{h}(\underline{\mathrm{s}}) \text{ is undefined or } \mathrm{h}'(\underline{\mathrm{s}}) \text{ is undefined;} \\ \textbf{1} \text{ iff for every } \underline{\mathrm{a}} \text{ s.t. } [\mathrm{h}(\underline{\mathrm{s}})](\underline{\mathrm{a}}) = 1,\ [\mathrm{h}'(\underline{\mathrm{s}})](\underline{\mathrm{a}}) = 1. \\ \textbf{0} \text{ otherwise.} \end{cases}$$

  $\}$

The interpretation of expressions of the form **a**$(\lambda\alpha.\ \Phi, \lambda\alpha.\ \Psi)$ is defined in a similar fashion, with the obvious semantics.

5.1.6 Scope Disambiguation by Defeasible Inference

Having extended the language into one that can be used to describe scopal underspecification, the framework for discourse interpretation developed in section §4.2 can also be used to formalize the inferences involved in scope disambiguation. Partially disambiguated interpretations can be represented by expressions which mix logical forms with 'traditional' expressions, as done in DRT. For example, one could formalize Ioup's (1975) Grammatical Function Principle, stating that an NP in subject position by default takes scope over NPs in other position, as follows:

GRAMMATICAL-FUNCTION-PRINCIPLE

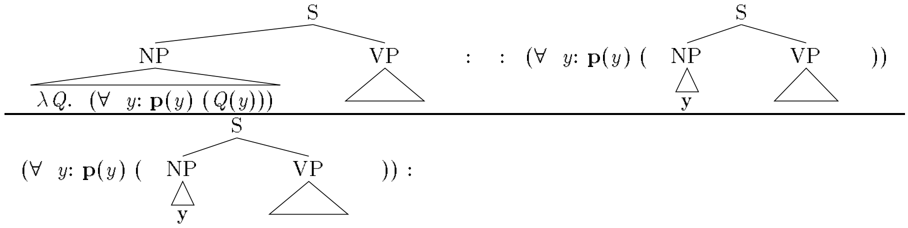

Logical forms in LF$_S$ are sentential expressions, and can therefore serve as triggering condition of discourse interpretation principles. They can also occur embedded in other expressions of $\mathcal{LSUL}$. During the scope disambiguation process, 'less ambiguous' expressions are inferred by deriving expressions such as (∀ y: $\mathbf{p}(y)$ ([$_S$ [$_{NP}$ y] [$_{VP}$ α]])) in which some quantifiers have been extracted, by a process very similar to the one used in the top-down version of the DRT construction algorithm (Kamp and Reyle , 1993). 'Partial' scopal disambiguation is thus represented by $\mathcal{LSUL}$ expressions which still contain logical forms.

As some readers will have already observed, the rule GRAMMATICAL-FUNCTION-PRINCIPLE does not satisfy the Anti-Random restriction proposed in §4.2: the rule does not contain a non-trivial restriction on the contexts in which it can operate. The already mentioned proposal in (Poesio , 1994) overcomes this problem by making the activation of scope disambiguation rules depend on whether the appropriate domain for the quantifier (its RESOURCE SITUATION) has been identified; a presentation of that proposal would however require introducing too much additional material.

5.2 Referential Ambiguity

5.2.1 Referential Expressions as Cases of Semantic Ambiguity

Yet another way in which the semantics of sentences is 'underspecified' by their syntax is in the interpretation of anaphoric expressions and other expressions whose interpretation has to be fixed in context. In semantics, referential expressions are traditionally translated as free variables whose interpretation depends on the choice of an assignment function (for the cases of deictic anaphora) or by assigning them the same variable bound by the quantifier that serves as their antecedent (for the cases of bound anaphora). This translation does capture the intuition that the truth conditions of a sentence containing a referential expression can only be evaluated after fixing the value of the referential expressions. It is also clear, however, that distinct propositions are obtained depending on the value assigned to these expressions, much as distinct propositions are obtained depending on the choice of an interpretation for lexical items, or of a scope for operators: in other words, a sentence which includes a referential expression is semantically ambiguous much in the way a sentence containing a lexically ambiguous item is.[36]

A complete discussion of reference interpretation would require introducing a formalization of context, so I will only consider here the issue of providing an underspecified treatment of intra-clausal and deictic anaphora. I propose that referential expressions are cases of semantic ambiguity, and translate into a special kind of underspecified object that I will call PARAMETERS. Semantically, a parameter is a type e expression that, in a discourse situation $\underline{\mathrm{d}}$, denotes a set of functions from situations to elements of D_e in $\underline{\mathrm{d}}$. For example, the pronoun *he* would translate into a parameter $\dot{x}$ which, in a discourse situation $\underline{\mathrm{d}}$ with constituents $\underline{\mathrm{a}_1} \ldots \underline{\mathrm{a}_n}$, and given the set $\mathcal{S}$ of situations, will denote a set of functions $\{\mathrm{f}_1, \ldots, \mathrm{f}_m, \ldots\}$ from situations in $\mathcal{S}$ to $\underline{\mathrm{a}_1} \ldots \underline{\mathrm{a}_n}$, including at least the set of all constant functions that map each situation $\underline{\mathrm{s}}$ into $\underline{\mathrm{a}_j}$ if $\underline{\mathrm{a}_j}$ is a constituent of that situation (see below), and the set of all variable denotations. The reader will immediately realize that parameters are the equivalent for type e expressions of 'underspecified predicates' like $\mathbf{croak}_U$ introduced above.[37]

More formally, I propose to extend the set of terms of $\mathcal{LSUL}$ with a new class of parameters, whose interpretation is defined as follows. First of all, let us reformulate the semantics of variables given before, and make variables functions from assignments to values (rather than the other way around). This involves again using as interpretation function one that maps expressions into functions

[36] Pinkal takes pretty much the same position in (Pinkal , 1995). He also introduces a distinction there between a 'speaker-oriented' perspective on meaning versus a 'hearer-oriented' perspective. A speaker may well have a single interpretation in mind for a particular anaphoric expressions, but the hearer may have to recover this interpretation among the many that are possible in that particular context. This is of course true of all kinds of ambiguities, also those which have a pragmatic rather than a semantic nature, but in the case of referential ambiguity, the alternative interpretations correspond to distinct propositions in the semantic sense as well.

[37] The term 'parameter' comes from Situation Semantics (e.g., (Gawron and Peters , 1990)), where the lexical items whose interpretation depends on context are called PARAMETRIC, in the sense that their interpretation depends on the value assigned in context to one or more parameters. Parameters are also used in situation theory to translate pronouns and other anaphoric expression; but although the name and the 'dotted' notation is preserved here, the parameters I have just introduced are an entirely different type of objects than the parameters of situation theory, which are a special sort of objects in the universe, entirely distinct from individuals.

from assignments to meanings, as done for lambda-abstracts.

- For α a variable of type τ, $[\![\alpha]\!]^{\mathrm{M,d}}$ = the singleton set {f}, where f is a function f : Ass $\rightarrow$ ($\mathcal{S} \rightarrow \mathrm{D}_{\mathcal{T}}$) such that for every assignment g, f(g) is a constant function from $\mathcal{S} \rightarrow \mathrm{D}_{\mathcal{T}}$.

This definition of the meaning of a variable allows us to abstract away from assignments. We can now define the semantics of parameters as follows:

- For α a parameter of type τ, $[\![\alpha]\!]^{\mathrm{M,g,d}} = \bigcup\limits_{\beta_\tau} [\![\beta_\tau]\!]^{\mathrm{M,d}}$ (the union of the denotation of all variables of type τ) $\bigcup$ {f | f : $\mathcal{S} \rightarrow \mathrm{D}_\tau$ is a function such that for each $\underline{\mathrm{s}}$, f($\underline{\mathrm{s}}$) = $\underline{\mathrm{a}}$, where $\underline{\mathrm{a}}$ is an object of τ that is a constituent of a situation $\underline{\mathrm{s}'}$ which in turn is a constituent of the discourse situation d, if $\underline{\mathrm{a}}$ is a constituent of $\underline{\mathrm{s}}$; f($\underline{\mathrm{s}}$) = undefined otherwise.}.

For example, if the subset of D_e in d consists of the two atoms $\underline{\mathrm{j}}$ and $\underline{\mathrm{b}}$, then $[\![\dot{x}_e]\!]^{\mathrm{M,g,d}}$ = {f_1,f_2,...f_i,...}, where f_1, f_2 etc. are the functions that may serve as the denotation of constants and variables—f_1 is the function that maps each situation of which $\underline{\mathrm{j}}$ is a constituent into $\underline{\mathrm{j}}$, f_2 is the function that maps each situation of which $\underline{\mathrm{b}}$ is a constituent into $\underline{\mathrm{b}}$— and the other functions represent all the possible denotations of objects that the parameter may be resolved to. Note that the discourse situations plays here the role played by the variable assignment in 'free variable' theories of context dependence.

The grammar presented in the previous section can be straightforwardly extended as follows to generate sentences such as *It croaked*:

- PRO $\rightarrow$ it; $\dot{x}$
- NP $\rightarrow$ PRO; $[_{\mathrm{NP}}\ [_{\mathrm{PRO}}\ \mathrm{PRO}']]$

The definition of the interpretation of logical forms given above already gives the correct results for these cases.

5.2.2 Parameters and Discourse Interpretation

Referential ambiguity gets 'resolved' by *anchoring* a parameter. A parameter is ANCHORED if only one among the functions in its denotation results in a consistent interpretation of the set of sentences in which the parameter occurs; a parameter can be anchored by means of equality statements of the form [$\dot{x}$ = *a*], where *a* is not parametric, or is already anchored: such equality statements make all but one of the interpretations of the parameter inadmissible. Once a parameter is anchored, it can be 'replaced' by a term that denotes the one function among those in the interpretation of the parameter that does not result in an inconsistent interpretation, much as in the previous discussion of lexical disambiguation, an H-type ambiguous predicate could be replaced by a disambiguated version. So, the discourse interpretation principles formalizing pronoun disambiguation involve a rewriting operation, just as the discourse interpretation principles formalizing lexical disambiguation.

An apparent disadvantage of the present theory with respect to the 'free variable' theory of context dependence is that we can derive from the latter that the value of referential expressions has to be fixed in order to get the meaning of the sentence in which they occur. A conversation is infelicitous unless the referents of all pronouns and definite descriptions have been identified, the domain of quantification of all quantifiers has been appropriately restricted, and so forth: so much so that listeners appear to be ready to ACCOMODATE new information (e.g., to introduce into the discourse some otherwise unspecified antecedent for a pronoun) rather than leave the interpretation unspecified (Lewis , 1979). But this fact about referential expressions also follows if we treat context dependence as a case of (H-type) semantic ambiguity; it is just a corollary of Pinkal's precisification imperative, from which I derived the Condition on Discourse Interpretation in section §4. Accomodation procedures can then be seen as a way of 'precisifying' in lack of sufficient information.

5.3 Syntactic Ambiguity

The one case of ambiguity that requires extending the framework introduced here considerably is syntactic ambiguity, as in *They saw her duck*. Furthermore, I haven't considered the problem of structural disambiguation in any detail. I refer the interested readers to (Poesio , 1994; Poesio , 1995) for a sketchy discussion of how to encode encoding syntactic ambiguity in an underspecified representation.[38]

6 Discussion

I have suggested that to develop a theory of discourse interpretation that is consistent with what we know about the problem of ambiguity, we need to look both at the grammar and at discourse interpretation. I proposed a theory of grammar consistent with what I have called the *Underspecification Hypothesis* and which is not based on the assumption that all natural language expressions can be disambiguated; and a theory of discourse interpretation according to which a perceived ambiguity occurs when defeasible interpretation principles result in conflicting hypothesis. The interpretation process is subject to two

[38] Most systems making use of underspecified representations perform structural disambiguation independently from the other forms of disambiguation (Schubert and Pelletier , 1982; Allen , 1987; Hobbs and Shieber , 1987; Alshawi , 1992). There is evidence, however, that structural disambiguation interacts at least with reference interpretation (Crain and Steedman , 1985; Altmann and Steedman , 1988) and a lot of the recent work on statistical parsing relies on the hypothesis that lexical interpretation affects parsing as well. Nothing in the proposal relies on structural disambiguation occurring prior to the other stages of disambiguation.

constraints: the *Anti-Random Hypothesis* (interpretations are not generated at random) and the *Condition on Discourse Interpretation*, derived from the *Precisification Imperative* (H-type ambiguity has to be resolved). Although treatments of disambiguation based on defeasible reasoning have been proposed elsewhere in the literature (e.g., in (Alshawi , 1992)), I am not aware of any discussion of the characteristics of this inferential process, the consequences of reasoning with an underspecified representation, or the need for constraints on the inference rules.

In the theory, semantic ambiguity is characterized model-theoretically in terms of multiplicity of sense, whereas perceived ambiguity is characterized in terms of inference. One may wonder if the distinction is really necessary; i.e., if it is really the case that the meaning of natural language expressions can be specified *a priori*. Two arguments in favor of a distinction are that it provides for a clean distinction between the role of grammar and the role of discourse interpretation; and that perceived ambiguity may also reflect non-semantic distinctions, e.g., distinctions in speech act interpretation; this question is not however totally resolved in the paper.

There are two obvious directions in which the present model needs to extended: to provide a model of syntactic ambiguity, and to account for the effect of incrementality in sentence processing. Preliminary work in this direction is discussed in (Poesio , 1995).

An issue that deserves further inspection is whether the formal similarity between the system used here to assign a denotation to indefinite sentences, and the systems developed by Hamblin for dealing with questions (Hamblin , 1973) and by Rooth for its alternative semantics (Rooth , 1985) has some significance. In particular, it would be interesting to explore the consequences of using parameters as the translation of focused elements.

Acknowledgments

I owe the realization of the importance of the phenomenon of deliberate ambiguity to my advisor Len Schubert and to Graeme Hirst. Special thanks to Robin Cooper, Richard Crouch, Kees van Deemter, Howard Kurtzman, David Milward, Manfred Pinkal for many discussions on the topic; and to two anonymous reviewers whose suggestions went well beyond the call of duty. I am also grateful to James Allen, Ariel Cohen, Tim Fernando, Janet Hitzeman, Peter Lasersohn, Barbara Partee, Enric Vallduvi, Sandro Zucchi, and the audiences at the University of Rochester, the University of Edinburgh, ICSI Berkeley, Carnegie-Mellon University, University of Stuttgart, SRI Cambridge, Tilburg University, and University of Saarbruecken. All errors are of course mine. This work was in part supported by the LRE Project 62-051 FraCaS.

References

Allen, J. F. 1987. *Natural Language Understanding.* Menlo Park, CA: Benjamin Cummings.

Allen, J. F., L. K. Schubert, G. Ferguson, P. Heeman, C. H. Hwang, T. Kato, M. Light, N. Martin, B. Miller, M. Poesio, and D. R. Traum. 1995. The TRAINS project: a case study in building a conversational planning agent. *Journal of Experimental and Theoretical AI* 7:7–48.

Alshawi, H. (ed.). 1992. *The Core Language Engine.* The MIT Press.

Altmann, G. T. M. (ed.). 1989. *Parsing and Interpretation.* Hove, East Sussex, UK: Lawrence Erlbaum.

Altmann, G. T. M., and M. Steedman. 1988. Interaction with Context during Human Sentence Processing. *Cognition* 30:191–238.

Asher, N., and M. Morreau. 1991. Common Sense Entailment: A Modal Theory of Commonsense Reasoning. In *Proc. 12th IJCAI.*

Barwise, J., and R. Cooper. 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4(2):159–219.

Barwise, J., and R. Cooper. 1993. Extended Kamp Notation. In *Situation Theory and its Applications, v.3*, ed. P. Aczel, D. Israel, Y. Katagiri, and S. Peters. Chap. 2, 29–54. CSLI.

Barwise, J., and J. Perry. 1983. *Situations and Attitudes.* Cambridge, MA: MIT Press, Cambridge Mass.

Charniak, E., and R. P. Goldman. 1988. A Logic for Semantic Interpretation. In *Proc. ACL-88*, 87–94. Buffalo, NY.

Chierchia, G., and S. McConnell-Ginet. 1990. *Meaning and Grammar: An Introduction to Semantics.* Cambridge, MA: The MIT Press.

Chomsky, N. 1981. *Lectures on Government and Binding.* Dordrecht: Foris.

Cooper, R. 1983. *Quantification and Syntactic Theory.* Dordrecht, Holland: D. Reidel Publishing Company.

Crain, S., and M. Steedman. 1985. On not being led up the garden path: the use of context by the psychological syntax processor. In *Natural Language Parsing: Psychological, Computational and Theoretical perspectives*, ed. D. R. Dowty, L. Karttunen, and A. M. Zwicky. 320–358. New York: Cambridge University Press.

Crouch, R. 1995. Ellipsis and Quantification: A Substitutional Approach. In *Proceedings 7th Conference of the European Chapter of the Association for Computational Linguistics*, 229–236. Dublin. Dublin City University.

Dalrymple, M., S. M. Shieber, and F. C. N. Pereira. 1991. Ellipsis and Higher-Order Unification. *Linguistics and Philosophy* 14(4):399–452.

Dowty, D. R., R. E. Wall, and S. Peters. 1981. *Introduction to Montague Semantics*. Dordrecht, Holland: D. Reidel.

Fenstad, J.E., P.K. Halvorsen, T. Langholm, and J. van Benthem. 1987. *Situations, Language and Logic*. Dordrecht: D.Reidel.

Fine, K. 1975. Vagueness, Truth, and Logic. *Synthese* 30:265–300.

Frazier, L., and J. D. Fodor. 1978. The sausage machine: A new two-stage parsing model. *Cognition* 6:291–295.

Gawron, J. M., and S. Peters. 1990. *Anaphora and Quantification in Situation Semantics*. Lecture Notes, Vol. 19. CSLI.

Gibson, E. 1991. *A Computational Theory of human linguistic processing: memory limitations and processing breakdown*. Doctoral dissertation, Carnegie Mellon University, Pittsburgh.

Gillon, B. 1990. Ambiguity, generality, and indeterminacy: Tests and definitions. *Synthese* 85:391–416.

Grosz, B.J., A.K. Joshi, and S. Weinstein. 1983. Providing a Unified Account of Definite Noun Phrases in Discourse. In *Proc. ACL-83*, 44–50.

Haegeman, L. 1991. *An Introduction to Government and Binding Theory*. Basil Blackwell. First edition.

Hamblin, C. 1973. Questions in Montague English. *Foundations of Language* 10:41–53.

Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Doctoral dissertation, University of Massachusetts at Amherst.

Hirst, G. 1987. *Semantic Interpretation and the Resolution of Ambiguity*. Studies in Natural Language Processing. Cambridge, UK: Cambridge University Press.

Hobbs, J. R. 1983. An Improper Treatment of Quantification in Ordinary English. In *Proc. ACL-83*, 57–63. Cambridge, MA, June.

Hobbs, J. R., and S. M. Shieber. 1987. An Algorithm for Generating Quantifier Scopings. *Computational Linguistics* 13(1-2):47–63.

Hobbs, J. R., M. Stickel, P. Martin, and D. Edwards. 1990. Interpretation as Abduction. Technical Note 499. Menlo Park, CA: SRI International, December.

Hwang, C. H., and L. K. Schubert. 1993. Episodic Logic: A Situational Logic for Natural Language Processing. In *Situation Theory and its Applications, v.3*, ed. P. Aczel, D. Israel, Y. Katagiri, and S. Peters. 303–338. CSLI.

Ioup, G. 1975. Some Universals for Quantifier Scope. In *Syntax and Semantics 4*, ed. J. Kimball. 37–58. New York: Academic Press.

Kamp, H., and U. Reyle. 1993. *From Discourse to Logic*. Dordrecht: D. Reidel.

Kaplan, D. 1977. Demonstratives. An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other indexicals. Unpublished manuscript, University of California, Los Angeles.

Keller, W. R. 1988. Nested Cooper Storage: The Proper Treatment of Quantification in Ordinary Noun Phrases. In *Natural Language Parsing and Linguistic Theories*, ed. U. Reyle and C. Rohrer. 432–447. Dordrecht: D. Reidel.

Kempson, R., and A. Cormack. 1981. Ambiguity and Quantification. *Linguistics and Philosophy* 4(2):259–310.

Kurtzman, H. 1985. *Studies in Syntactic Ambiguity Resolution*. Doctoral dissertation, MIT, Cambridge, MA.

Kurtzman, H. S., and M. C. MacDonald. 1993. Resolution of Quantifier Scope Ambiguities. *Cognition* 48:243–279.

Lakoff, G. P. 1970. A note on vagueness and ambiguity. *Linguistic Inquiry* 1(3):357–359.

Lascarides, A., N. Asher, and J. Oberlander. 1992. Inferring Discourse Relations in Context. In *Proc. ACL-92*, 1–8. University of Delaware.

Lewis, D. K. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8:339–359.

May, R. 1985. *Logical Form in Natural Language*. The MIT Press.

Montague, R. 1970. Universal Grammar. *Theoria* 36:373–398. Reprinted in (Thomason , 1974).

Partee, B. H., and M. Rooth. 1983. Generalized Conjunction and Type Ambiguity. In *Meaning, Use and Interpretation of Language*, ed. R. Bauerle, C. Schwarze, and A. von Stechow. Berlin, West Germany: Walter de Gruyter.

Pereira, F. C. N. 1990. Categorial Semantics and Scoping. *Computational Linguistics* 16(1):1–10.

Pereira, F. C. N., and M. E. Pollack. 1991. Incremental Interpretation. *Artificial Intelligence* 50:37–82.

Pinkal, M. 1985. *Logik und Lexikon: Die Semantik des Unbestimmten.* Berlin: de Gruyter.

Pinkal, M. 1995. *Logic and Lexicon.* London: Oxford.

Poesio, M. 1991. Relational Semantics and Scope Ambiguity. In *Situation Semantics and its Applications, vol.2*, ed. J. Barwise, J. M. Gawron, G. Plotkin, and S. Tutiya. Chap. 20, 469–497. Stanford, CA: CSLI.

Poesio, M. 1994. *Discourse Interpretation and the Scope of Operators.* Doctoral dissertation, University of Rochester, Department of Computer Science, Rochester, NY.

Poesio, M. 1995. A Model of Conversation Processing Based on Micro Conversational Events. In *Proceedings of the Annual Meeting of the Cognitive Science Society.* Pittsburgh.

Raskin, V. 1985. *Semantic Mechanisms of Humor.* Dordrecht and Boston: D. Reidel.

Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13(1–2):81–132.

Reyle, U. 1993. Dealing with ambiguities by underspecification: Construction, Representation and Deduction. *Journal of Semantics* 3.

Rodman, R. 1976. Scope Phenomena, "Movement Transformations," and Relative Clauses. In *Montague Grammar*, ed. Barbara Partee. 165–176. Academic Press.

Rooth, M. 1985. *Association with Focus.* Doctoral dissertation, University of Massachusetts, Amherst.

Schubert, L. K. 1986. Are There Preference Trade-Offs in Attachment Decisions. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 601–605. Philadelphia, Pennsylvania, August. American Association for Artificial Intelligence.

Schubert, L. K., and F. J. Pelletier. 1982. From English to Logic: Context-Free Computation of 'Conventional' Logical Translations. *American Journal of Computational Linguistics* 10:165–176.

Spivey-Knowlton, M., J. Sedivy, K. Eberhard, and M. Tanenhaus. 1994. Psycholinguistic Study of the Interaction between Language and Vision. In *Proceedings of 12th National Conference on Artificial Intelligence (AAAI-94)*. Seattle.

Stallard, D. 1987. The Logical Analysis of Lexical Ambiguity. In *Proceedings of the 25th Meeting of the ACL*, 179–185.

Su, S. P. 1994. *Lexical Ambiguity in Poetry*. London: Longman.

Swinney, D. A. 1979. Lexical Access During Sentence Comprehension: (Re)consideration of Context Effects. *Journal of Verbal Learning and Verbal Behavior* 18:545–567.

Thomason, R. H. (ed.). 1974. *Formal Philosophy: Selected Papers of Richard Montague*. New York: Yale University Press.

Turner, R. 1992. Properties, propositions and semantic theory. In *Computational Linguistics and Formal Semantics*, ed. M. Rosner and R. Johnson. CUP, Cambridge.

van Deemter, K. 1991. *On the Composition of Meaning*. Doctoral dissertation, University of Amsterdam.

Verkuyl, H. J. 1992. Some Issues in the Analysis of Multiple Quantification with Plural NPs. OTS Working Papers OTS-WP-TL-92-005. The Netherlands: University of Utrecht, Research Institute for Language and Speech. To appear in F. Hamm and E. Hinrichs, editors, *Plural Quantification*, Kluwer.

Zwicky, A., and J. Sadock. 1975. Ambiguity Tests and How to Fail Them. In *Syntax and Semantics 4*, ed. J. Kimball. 1–36. New York: Academic Press.